

---

# MPC8313E PowerQUICC™ II Pro Integrated Processor Family Reference Manual

**Supports**  
MPC8313E  
MPC8313

MPC8313ERM  
Rev. 2  
12/2008



## How to Reach Us:

### Home Page:

[www.freescale.com](http://www.freescale.com)

### Web Support:

<http://www.freescale.com/support>

### USA/Europe or Locations Not Listed:

Freescale Semiconductor, Inc.  
Technical Information Center, EL516  
2100 East Elliot Road  
Tempe, Arizona 85284  
1-800-521-6274 or  
+1-480-768-2130  
[www.freescale.com/support](http://www.freescale.com/support)

### Europe, Middle East, and Africa:

Freescale Halbleiter Deutschland GmbH  
Technical Information Center  
Schatzbogen 7  
81829 Muenchen, Germany  
+44 1296 380 456 (English)  
+46 8 52200080 (English)  
+49 89 92103 559 (German)  
+33 1 69 35 48 48 (French)  
[www.freescale.com/support](http://www.freescale.com/support)

### Japan:

Freescale Semiconductor Japan Ltd.  
Headquarters  
ARCO Tower 15F  
1-8-1, Shimo-Meguro, Meguro-ku  
Tokyo 153-0064  
Japan  
0120 191014 or  
+81 3 5437 9125  
[support.japan@freescale.com](mailto:support.japan@freescale.com)

### Asia/Pacific:

Freescale Semiconductor China Ltd.  
Exchange Building 23F  
No. 118 Jianguo Road  
Chaoyang District  
Beijing 100022  
China  
+86 10 5879 8000  
[support.asia@freescale.com](mailto:support.asia@freescale.com)

### For Literature Requests Only:

Freescale Semiconductor  
Literature Distribution Center  
P.O. Box 5405  
Denver, Colorado 80217  
1-800 441-2447 or  
+1-303-675-2140  
Fax: +1-303-675-2150  
[LDCForFreescaleSemiconductor@hibbertgroup.com](mailto:LDCForFreescaleSemiconductor@hibbertgroup.com)

Information in this document is provided solely to enable system and software implementers to use Freescale Semiconductor products. There are no express or implied copyright licenses granted hereunder to design or fabricate any integrated circuits or integrated circuits based on the information in this document.

Freescale Semiconductor reserves the right to make changes without further notice to any products herein. Freescale Semiconductor makes no warranty, representation or guarantee regarding the suitability of its products for any particular purpose, nor does Freescale Semiconductor assume any liability arising out of the application or use of any product or circuit, and specifically disclaims any and all liability, including without limitation consequential or incidental damages. "Typical" parameters which may be provided in Freescale Semiconductor data sheets and/or specifications can and do vary in different applications and actual performance may vary over time. All operating parameters, including "Typicals" must be validated for each customer application by customer's technical experts. Freescale Semiconductor does not convey any license under its patent rights nor the rights of others. Freescale Semiconductor products are not designed, intended, or authorized for use as components in systems intended for surgical implant into the body, or other applications intended to support or sustain life, or for any other application in which the failure of the Freescale Semiconductor product could create a situation where personal injury or death may occur. Should Buyer purchase or use Freescale Semiconductor products for any such unintended or unauthorized application, Buyer shall indemnify and hold Freescale Semiconductor and its officers, employees, subsidiaries, affiliates, and distributors harmless against all claims, costs, damages, and expenses, and reasonable attorney fees arising out of, directly or indirectly, any claim of personal injury or death associated with such unintended or unauthorized use, even if such claim alleges that Freescale Semiconductor was negligent regarding the design or manufacture of the part.

Portions of Chapter 15, "Universal Serial Bus Interface," relating to the EHCI specification are Copyright © Intel Corporation 1999-2001. The EHCI specification is provided "As Is" with no warranties whatsoever, including any warranty of merchantability, non-infringement, fitness for any particular purpose, or any warranty otherwise arising out of any proposal, specification or sample. Intel disclaims all liability, including liability for infringement of any proprietary rights, relating to use of information in the EHCI specification. Intel may make changes to the EHCI specifications at any time, without notice.

Freescale and the Freescale logo are trademarks or registered trademarks of Freescale Semiconductor, Inc. in the U.S. and other countries. All other product or service names are the property of their respective owners. The Power Architecture and Power.org word marks and the Power and Power.org logos and related marks are trademarks and service marks licensed by Power.org. The described product contains a PowerPC processor core. The PowerPC name is a trademark of IBM Corp. and used under license. IEEE 754, 802.1, 802.2, 802.1p, 802.1Q, 802.3, 802.3u, 802.3x, 802.3z, 802.3ab, 802.3ac, 802.11a/g, 802.11b, 802.11i, 802.11n, 1149.1, and 1588 are trademarks or registered trademarks of the Institute of Electrical and Electronics Engineers, Inc. (IEEE). This product is not endorsed or approved by the IEEE.

© Freescale Semiconductor, Inc., 2006–2008. All rights reserved.

Document Number: MPC8313ERM  
Rev. 2, 12/2008



# Contents

Paragraph Number	Title	Page Number
<b>About This Book</b>		
	Audience .....	lxxiii
	Organization.....	lxxiii
	Suggested Reading.....	lxxv
	General Information.....	lxxv
	Related Documentation.....	lxxvi
	Conventions .....	lxxvi
	Signal Conventions .....	lxxvii
	Acronyms and Abbreviations .....	lxxvii
<b>Chapter 1</b>		
<b>Overview</b>		
1.1	MPC8313E PowerQUICC II Pro Processor Overview .....	1-1
1.2	MPC8313E Architecture Overview .....	1-7
1.2.1	Power Architecture Core .....	1-7
1.2.2	Security Engine.....	1-10
1.2.3	DDR Memory Controller.....	1-10
1.2.4	Dual Enhanced Three-Speed Ethernet Controllers.....	1-11
1.2.5	PCI Controller.....	1-12
1.2.5.1	PCI Bus Arbitration Unit.....	1-13
1.2.6	Universal Serial Bus (USB) 2.0.....	1-13
1.2.6.1	USB Dual-Role Controller .....	1-14
1.2.7	Enhanced Local Bus Controller (eLBC).....	1-14
1.2.8	Integrated Programmable Interrupt Controller (IPIC).....	1-16
1.2.9	Dual I <sup>2</sup> C Interfaces .....	1-16
1.2.10	DMA Controller.....	1-17
1.2.11	Dual Universal Asynchronous Receiver/Transmitter (DUART).....	1-17
1.2.12	Serial Peripheral Interface (SPI).....	1-18
1.2.13	System Timers .....	1-18
1.3	Application Examples.....	1-18
1.3.1	Low-End Printer CPU and Interface ASIC.....	1-19
1.3.2	High-End Printer I/O Processor.....	1-20
1.3.3	IEEE Std. 1588 in Test and Measurement and Industrial Automation .....	1-21
1.3.4	IEEE Std. 802.11n WLAN Access Point .....	1-23
1.3.5	Media Server.....	1-24

# Contents

Paragraph Number	Title	Page Number
<b>Chapter 2</b>		
<b>Memory Map</b>		
2.1	Internal Memory Mapped Registers .....	2-1
2.2	Accessing IMMR Memory From the Local Processor .....	2-1
2.3	Complete IMMR Map .....	2-1
<b>Chapter 3</b>		
<b>Signal Descriptions</b>		
3.1	Signals Overview .....	3-1
3.2	Configuration Signals Sampled at Reset .....	3-12
3.3	Output Signal States During Reset .....	3-13
3.4	External Signal Description .....	3-14
<b>Chapter 4</b>		
<b>Reset, Clocking, and Initialization</b>		
4.1	External Signals .....	4-1
4.1.1	Reset Signals .....	4-1
4.1.2	Clock Signals .....	4-3
4.2	Functional Description .....	4-4
4.2.1	Reset Operations .....	4-4
4.2.1.1	Reset Causes .....	4-5
4.2.1.2	Reset Actions .....	4-5
4.2.2	Power-On Reset Flow .....	4-6
4.2.3	Hard Reset Flow .....	4-8
4.2.4	Soft Reset Flow .....	4-9
4.3	Reset Configuration .....	4-9
4.3.1	Reset Configuration Signals .....	4-9
4.3.1.1	Reset Configuration Word Source .....	4-10
4.3.1.2	SYS_CLK_IN Division .....	4-11
4.3.1.3	Selecting Reset Configuration Input Signals .....	4-11
4.3.2	Reset Configuration Words .....	4-12
4.3.2.1	Reset Configuration Word Low Register (RCWLR) .....	4-13
4.3.2.1.1	System PLL Configuration .....	4-13
4.3.2.2	Reset Configuration Word High Register (RCWHR) .....	4-14
4.3.2.2.1	PCI Host/Agent Configuration .....	4-16
4.3.2.2.2	Boot Memory Space (BMS) .....	4-17
4.3.2.2.3	Boot Sequencer Configuration .....	4-17
4.3.2.2.4	Boot ROM Location .....	4-18

# Contents

Paragraph Number	Title	Page Number
4.3.2.2.5	eTSEC1 Mode .....	4-19
4.3.2.2.6	eTSEC2 Mode .....	4-20
4.3.2.2.7	e300 Core True Little-Endian .....	4-20
4.3.2.2.8	LALE Configuration .....	4-21
4.3.3	Loading the Reset Configuration Words .....	4-21
4.3.3.1	Loading from Local Bus .....	4-21
4.3.3.1.1	Local Bus Controller Setting .....	4-22
4.3.3.2	Loading from I2C EEPROM .....	4-23
4.3.3.2.1	Using the Boot Sequencer Reset Configuration .....	4-23
4.3.3.2.2	EEPROM Calling Address .....	4-23
4.3.3.2.3	EEPROM Data Format in Reset Configuration Mode .....	4-23
4.3.3.2.4	Reset Configuration Load Fail .....	4-25
4.3.3.3	Default Reset Configuration Words .....	4-26
4.3.3.3.1	Examples for Hard-Coded Reset Configuration Words Usage .....	4-27
4.4	Clocking .....	4-28
4.4.1	Clocking in PCI Host Mode .....	4-30
4.4.1.1	PCI Clock Outputs (PCI_CLK_OUT[0:2]) .....	4-30
4.4.2	Clocking In PCI Agent Mode .....	4-30
4.4.3	System Clock Domains .....	4-30
4.4.4	USB Clocking .....	4-31
4.4.5	Ethernet Clocking .....	4-32
4.4.6	Real-Time Clock (RTC) .....	4-32
4.5	Memory Map/Register Definitions .....	4-32
4.5.1	Reset Configuration Register Descriptions .....	4-32
4.5.1.1	Reset Configuration Word Low Register (RCWLR) .....	4-33
4.5.1.2	Reset Configuration Word High Register (RCWHR) .....	4-33
4.5.1.3	Reset Status Register (RSR) .....	4-33
4.5.1.4	Reset Mode Register (RMR) .....	4-35
4.5.1.5	Reset Protection Register (RPR) .....	4-35
4.5.1.6	Reset Control Register (RCR) .....	4-36
4.5.1.7	Reset Control Enable Register (RCER) .....	4-37
4.5.2	Clock Configuration Registers .....	4-37
4.5.2.1	System PLL Mode Register (SPMR) .....	4-37
4.5.2.2	Output Clock Control Register (OCCR) .....	4-39
4.5.2.3	System Clock Control Register (SCCR) .....	4-40

## Chapter 5 System Configuration

5.1	Introduction .....	5-1
5.2	Local Memory Map Overview and Example .....	5-1

# Contents

Paragraph Number	Title	Page Number
5.2.1	Address Translation and Mapping .....	5-3
5.2.2	Window into Configuration Space.....	5-4
5.2.3	Local Access Windows.....	5-4
5.2.3.1	Local Access Register Memory Map .....	5-4
5.2.4	Local Access Register Descriptions .....	5-6
5.2.4.1	Internal Memory Map Registers Base Address Register (IMMRBAR).....	5-6
5.2.4.1.1	Updating IMMRBAR .....	5-6
5.2.4.2	Alternate Configuration Base Address Register (ALTCBAR).....	5-7
5.2.4.3	LBC Local Access Window <i>n</i> Base Address Registers (LBLAWBAR0–LBLAWBAR3) .....	5-8
5.2.4.3.1	LBLAWBAR0[BASE_ADDR] Reset Value .....	5-8
5.2.4.4	LBC Local Access Window <i>n</i> Attributes Registers (LBLAWAR0– LBLAWAR3).....	5-9
5.2.4.4.1	LBLAWAR0[EN] and LBLAWAR0[SIZE] Reset Value .....	5-9
5.2.4.5	PCI Local Access Window <i>n</i> Base Address Register (PCILAWBAR0–PCILAWBAR1) .....	5-10
5.2.4.5.1	PCILAWBAR0[BASE_ADDR] Reset Value .....	5-10
5.2.4.6	PCI Local Access Window <i>n</i> Attributes Registers (PCILAWAR0–PCILAWAR1) .....	5-11
5.2.4.6.1	PCILAWAR0[EN] and PCILAWAR0[SIZE] Reset Value .....	5-11
5.2.4.7	DDR Local Access Window <i>n</i> Base Address Registers (DDRLAWBAR0–DDRLAWBAR1).....	5-12
5.2.4.7.1	DDRLAWBAR0[BASE_ADDR] Reset Value.....	5-12
5.2.4.8	DDR Local Access Window <i>n</i> Attributes Registers (DDRLAWAR0– DDRLAWAR1).....	5-13
5.2.4.8.1	DDRLAWAR0[EN] and DDRLAWAR0[SIZE] Reset Value.....	5-13
5.2.5	Precedence of Local Access Windows .....	5-14
5.2.6	Configuring Local Access Windows .....	5-14
5.2.7	Distinguishing Local Access Windows from Other Mapping Functions .....	5-14
5.2.8	Outbound Address Translation and Mapping Windows.....	5-15
5.2.9	Inbound Address Translation and Mapping Windows .....	5-15
5.2.9.1	PCI Inbound Windows.....	5-15
5.2.10	Internal Memory Map.....	5-15
5.2.11	Accessing Internal Memory from External Masters.....	5-16
5.3	System Configuration .....	5-16
5.3.1	System Configuration Register Memory Map.....	5-16
5.3.2	System Configuration Registers .....	5-17
5.3.2.1	System General Purpose Register Low (SGPRL) .....	5-17
5.3.2.2	System General Purpose Register High (SGPRH).....	5-17
5.3.2.3	System Part and Revision ID Register (SPRIDR).....	5-18
5.3.2.3.1	SPRIDR[PARTID] Coding.....	5-18

# Contents

Paragraph Number	Title	Page Number
5.3.2.4	System Priority and Configuration Register (SPCR) .....	5-18
5.3.2.5	System I/O Configuration Register Low (SICRL) .....	5-20
5.3.2.6	System I/O Configuration Register High (SICRH) .....	5-23
5.3.2.7	Debug Configuration .....	5-26
5.3.2.7.1	DDR Debug Configuration.....	5-26
5.3.2.7.2	Local Bus Debug Configuration.....	5-27
5.3.2.8	DDR Control Driver Register (DDRCDR).....	5-27
5.3.2.9	DDR Debug Status Register (DDRDSR) .....	5-28
5.4	Software Watchdog Timer (WDT).....	5-29
5.4.1	WDT Overview.....	5-29
5.4.2	WDT Features.....	5-29
5.4.3	WDT Modes of Operation .....	5-30
5.4.4	WDT Memory Map/Register Definition .....	5-30
5.4.4.1	System Watchdog Control Register (SWCRR) .....	5-31
5.4.4.2	System Watchdog Count Register (SWCNR) .....	5-32
5.4.4.3	System Watchdog Service Register (SWSRR).....	5-32
5.4.5	Functional Description.....	5-33
5.4.5.1	Software Watchdog Timer Unit .....	5-33
5.4.5.2	Modes of Operation .....	5-34
5.4.6	Initialization/Application Information .....	5-35
5.4.6.1	WDT Programming Guidelines .....	5-35
5.5	Real Time Clock Module (RTC).....	5-35
5.5.1	RTC Overview .....	5-35
5.5.2	RTC Features .....	5-36
5.5.3	RTC Modes of Operation.....	5-36
5.5.4	RTC External Signal Description .....	5-36
5.5.5	RTC Memory Map/Register Definition.....	5-37
5.5.5.1	Real Time Counter Control Register (RTCNR) .....	5-37
5.5.5.2	Real Time Counter Load Register (RTLDR).....	5-38
5.5.5.3	Real Time Counter Prescale Register (RTPSR) .....	5-39
5.5.5.4	Real Time Counter Register (RTCTR) .....	5-39
5.5.5.5	Real Time Counter Event Register (RTEVR).....	5-40
5.5.5.6	Real Time Counter Alarm Register (RTALR) .....	5-40
5.5.6	Functional Description.....	5-41
5.5.6.1	Real Time Counter Unit.....	5-41
5.5.6.2	RTC Operational Modes .....	5-41
5.5.7	RTC Programming Guidelines.....	5-42
5.6	Periodic Interval Timer (PIT) .....	5-42
5.6.1	PIT Overview.....	5-42
5.6.2	PIT Features .....	5-43
5.6.3	PIT Modes of Operation .....	5-43

# Contents

Paragraph Number	Title	Page Number
5.6.4	PIT External Signal Description .....	5-43
5.6.5	PIT Memory Map/Register Definition .....	5-44
5.6.5.1	Periodic Interval Timer Control Register (PTCNR) .....	5-44
5.6.5.2	Periodic Interval Timer Load Register (PTLDR) .....	5-45
5.6.5.3	Periodic Interval Timer Prescale Register (PTPSR) .....	5-46
5.6.5.4	Periodic Interval Timer Counter Register (PTCTR) .....	5-46
5.6.5.5	Periodic Interval Timer Event Register (PTEVR) .....	5-46
5.6.6	Functional Description .....	5-47
5.6.6.1	Periodic Interval Timer Unit .....	5-47
5.6.6.2	PIT Operational Modes .....	5-48
5.6.7	PIT Programming Guidelines .....	5-48
5.7	General-Purpose Timers (GTM) .....	5-48
5.7.1	GTM Overview .....	5-48
5.7.2	GTM Features .....	5-49
5.7.3	GTM Modes of Operation .....	5-50
5.7.3.1	Cascaded Modes .....	5-50
5.7.3.2	Clock Source Modes .....	5-50
5.7.3.3	Reference Modes .....	5-50
5.7.3.4	Capture Modes .....	5-51
5.7.4	GTM External Signal Description .....	5-51
5.7.5	GTM Memory Map/Register Definition .....	5-52
5.7.5.1	Global Timers Configuration Registers (GTCFR <sub>n</sub> ) .....	5-54
5.7.5.2	Global Timers Mode Registers (GTMDR1–GTMDR4) .....	5-57
5.7.5.3	Global Timers Reference Registers (GTRFR1–GTRFR4) .....	5-58
5.7.5.4	Global Timers Capture Registers (GTCPR1–GTCPR4) .....	5-58
5.7.5.5	Global Timers Counter Registers (GTCNR1–GTCNR4) .....	5-59
5.7.5.6	Global Timers Event Registers (GTEVR1–GTEVR4) .....	5-59
5.7.5.7	Global Timers Prescale Registers (GTPSR1–GTPSR4) .....	5-60
5.7.6	Functional Description .....	5-61
5.7.6.1	General-Purpose Timer Units .....	5-61
5.7.6.2	Reference Modes .....	5-61
5.7.6.3	Capture Modes .....	5-61
5.7.6.4	Cascaded Modes .....	5-62
5.7.7	Initialization/Application Information .....	5-64
5.7.7.1	Programming Guidelines .....	5-64
5.7.7.1.1	GTM Registers .....	5-64
5.8	Power Management Control (PMC) .....	5-64
5.8.1	External Signal Description .....	5-65
5.8.2	PMC Memory Map/Register Definition .....	5-65
5.8.2.1	Power Management Controller Configuration Register (PMCCR) .....	5-66
5.8.2.2	Power Management Controller Event Register (PM CER) .....	5-67



# Contents

Paragraph Number	Title	Page Number
5.8.2.3	Power Management Controller Mask Register (PMCMR) .....	5-69
5.8.2.4	Power Management Controller Configuration Register 1 (PMCCR1).....	5-69
5.8.2.5	Power Management Controller Configuration Register 2 (PMCCR2).....	5-71
5.8.3	Functional Description.....	5-72
5.8.3.1	Dynamic Power Management.....	5-72
5.8.3.2	Shutting Down Unused Blocks.....	5-73
5.8.3.3	Software-Controlled Power-Down States.....	5-73
5.8.3.4	Software-Controlled Power Supply Switching.....	5-73
5.8.3.5	Support of PCI Power Management Interface Specification.....	5-74
5.8.3.5.1	Entering Low Power States—Core-Only Mode .....	5-77
5.8.3.5.2	Entering Low Power States—Core and System Mode .....	5-77
5.8.3.6	Exiting Core and System Low Power States .....	5-78
5.8.3.6.1	Exiting Low Power States—Core-Only Mode .....	5-78
5.8.3.6.2	Exiting Low Power States—Core and System Mode .....	5-78
5.8.3.7	MPC8313E-Specific PMC Low Power States .....	5-79
5.8.3.7.1	Power State Transitions from an ACPI Perspective .....	5-79
5.8.3.7.2	MPC8313E Low Power Sequencing .....	5-83
5.8.3.7.3	PMC External Power Supply Control .....	5-88
5.8.3.7.4	Low-Power Considerations .....	5-90
5.8.4	Initialization/Application Information .....	5-91
5.8.4.1	Core Disable in Low Power Mode .....	5-91

## Chapter 6 Arbiter and Bus Monitor

6.1	Overview.....	6-1
6.1.1	Coherent System Bus Overview .....	6-1
6.2	Arbiter Memory Map/Register Definition .....	6-2
6.2.1	Arbiter Configuration Register (ACR) .....	6-2
6.2.2	Arbiter Timers Register (ATR) .....	6-4
6.2.3	Arbiter Event Register (AER).....	6-5
6.2.4	Arbiter Interrupt Definition Register (AIDR).....	6-6
6.2.5	Arbiter Mask Register (AMR).....	6-7
6.2.6	Arbiter Event Attributes Register (AEATR).....	6-7
6.2.7	Arbiter Event Address Register (AEADR).....	6-9
6.2.8	Arbiter Event Response Register (AERR).....	6-10
6.3	Functional Description.....	6-10
6.3.1	Arbitration Policy .....	6-10
6.3.1.1	Address Bus Arbitration with $\overline{\text{PRIORITY}}[0:1]$ .....	6-11
6.3.1.2	Address Bus Arbitration with $\overline{\text{REPEAT}}$ .....	6-12
6.3.1.3	Address Bus Arbitration after $\overline{\text{ARTRY}}$ .....	6-13

# Contents

Paragraph Number	Title	Page Number
6.3.1.4	Address Bus Parking.....	6-13
6.3.1.5	Data Bus Arbitration.....	6-13
6.3.2	Bus Error Detection .....	6-13
6.3.2.1	Address Time Out.....	6-13
6.3.2.2	Data Time Out .....	6-14
6.3.2.3	Transfer Error .....	6-14
6.3.2.4	Address Only Transaction Type.....	6-14
6.3.2.5	Reserved Transaction Type.....	6-15
6.3.2.6	Illegal (eciwx/ecowx) Transaction Type.....	6-15
6.4	Initialization/Applications Information .....	6-16
6.4.1	Initialization Sequence.....	6-16
6.4.2	Error Handling Sequence.....	6-16

## Chapter 7 e300 Processor Core Overview

7.1	Overview.....	7-1
7.1.1	Features.....	7-3
7.1.2	Instruction Unit.....	7-6
7.1.2.1	Instruction Queue and Dispatch Unit .....	7-6
7.1.2.2	Branch Processing Unit (BPU).....	7-7
7.1.3	Independent Execution Units.....	7-7
7.1.3.1	Integer Unit (IU).....	7-7
7.1.3.2	Floating-Point Unit (FPU) .....	7-7
7.1.3.3	Load/Store Unit (LSU) .....	7-8
7.1.3.4	System Register Unit (SRU).....	7-8
7.1.4	Completion Unit .....	7-8
7.1.5	Memory Subsystem Support.....	7-8
7.1.5.1	Memory Management Units (MMUs).....	7-9
7.1.5.2	Cache Units.....	7-10
7.1.6	Bus Interface Unit (BIU) .....	7-10
7.1.7	System Support Functions .....	7-11
7.1.7.1	Power Management .....	7-11
7.1.7.2	Time Base/Decrementer .....	7-11
7.1.7.3	JTAG Test and Debug Interface.....	7-12
7.1.7.4	Clock Multiplier.....	7-12
7.1.7.5	Core Performance Monitor .....	7-12
7.2	PowerPC Architecture Implementation .....	7-13
7.3	Implementation-Specific Information.....	7-13
7.3.1	Register Model.....	7-14
7.3.1.1	UISA Registers .....	7-16

# Contents

Paragraph Number	Title	Page Number
7.3.1.1.1	General-Purpose Registers (GPRs) .....	7-16
7.3.1.1.2	Floating-Point Registers (FPRs).....	7-16
7.3.1.1.3	Condition Register (CR).....	7-16
7.3.1.1.4	Floating-Point Status and Control Register (FPSCR) .....	7-16
7.3.1.1.5	User-Level SPRs.....	7-16
7.3.1.2	VEA Registers .....	7-17
7.3.1.3	OEA Registers .....	7-17
7.3.1.3.1	Machine State Register (MSR).....	7-17
7.3.1.3.2	Segment Registers (SRs) .....	7-19
7.3.1.3.3	Supervisor-Level SPRs.....	7-19
7.3.2	Instruction Set and Addressing Modes .....	7-26
7.3.2.1	PowerPC Instruction Set and Addressing Modes.....	7-26
7.3.2.2	Implementation-Specific Instruction Set .....	7-27
7.3.3	Cache Implementation .....	7-28
7.3.3.1	PowerPC Cache Characteristics .....	7-28
7.3.3.2	Implementation-Specific Cache Organization.....	7-28
7.3.3.3	Instruction and Data Cache Way-Locking.....	7-30
7.3.4	Interrupt Model.....	7-30
7.3.4.1	PowerPC Interrupt Model.....	7-30
7.3.4.2	Implementation-Specific Interrupt Model .....	7-31
7.3.5	Memory Management.....	7-34
7.3.5.1	PowerPC Memory Management.....	7-34
7.3.5.2	Implementation-Specific Memory Management .....	7-34
7.3.6	Instruction Timing .....	7-35
7.3.7	Core Interface .....	7-36
7.3.7.1	Memory Accesses.....	7-37
7.3.7.2	Signals.....	7-37
7.3.8	Debug Features .....	7-38
7.3.8.1	Breakpoint Signaling .....	7-38
7.4	Differences Between Cores.....	7-39

## Chapter 8 Integrated Programmable Interrupt Controller (IPIC)

8.1	Introduction.....	8-1
8.2	Features .....	8-4
8.3	Modes of Operation .....	8-4
8.3.1	Core Enable Mode .....	8-4
8.3.2	Core Disable Mode .....	8-5
8.4	External Signal Description .....	8-5
8.4.1	Overview.....	8-5

# Contents

Paragraph Number	Title	Page Number
8.4.2	Detailed Signal Descriptions .....	8-5
8.5	Memory Map/Register Definition .....	8-6
8.5.1	System Global Interrupt Configuration Register (SICFR) .....	8-7
8.5.2	System Global Interrupt Vector Register (SIVCR).....	8-9
8.5.3	System Internal Interrupt Pending Registers (SIPNR_H and SIPNR_L).....	8-11
8.5.4	System Internal Interrupt Group A Priority Register (SIPRR_A).....	8-13
8.5.5	System Internal Interrupt Group D Priority Register (SIPRR_D).....	8-14
8.5.6	System Internal Interrupt Mask Register (SIMSR_H and SIMSR_L) .....	8-15
8.5.7	System Internal Interrupt Control Register (SICNR) .....	8-16
8.5.8	System External Interrupt Pending Register (SEPNR).....	8-18
8.5.9	System Mixed Interrupt Group A Priority Register (SMPRR_A).....	8-18
8.5.10	System Mixed Interrupt Group B Priority Register (SMPRR_B) .....	8-19
8.5.11	System External Interrupt Mask Register (SEMSR) .....	8-20
8.5.12	System External Interrupt Control Register (SECNR).....	8-21
8.5.13	System Error Status Register (SERSR) .....	8-22
8.5.14	System Error Mask Register (SERMR).....	8-23
8.5.15	System Error Control Register (SERCR) .....	8-24
8.5.16	System Internal Interrupt Force Registers (SIFCR_H and SIFCR_L) .....	8-25
8.5.17	System External Interrupt Force Register (SEFCR).....	8-26
8.5.18	System Error Force Register (SERFR).....	8-26
8.5.19	System Critical Interrupt Vector Register (SCVCR) .....	8-27
8.5.20	System Management Interrupt Vector Register (SMVCR) .....	8-27
8.6	Functional Description.....	8-28
8.6.1	Interrupt Types .....	8-28
8.6.2	Interrupt Configuration .....	8-29
8.6.3	Internal Interrupts Group Relative Priority .....	8-30
8.6.4	Mixed Interrupts Group Relative Priority.....	8-30
8.6.5	Highest Priority Interrupt.....	8-31
8.6.6	Interrupt Source Priorities.....	8-31
8.6.7	Masking Interrupt Sources.....	8-34
8.6.8	Interrupt Vector Generation and Calculation .....	8-35
8.6.9	Machine Check Interrupts.....	8-35

## Chapter 9 DDR Memory Controller

9.1	Introduction.....	9-1
9.2	Features .....	9-2
9.2.1	Modes of Operation .....	9-3
9.3	External Signal Descriptions .....	9-3
9.3.1	Signals Overview .....	9-3

# Contents

Paragraph Number	Title	Page Number
9.3.2	Detailed Signal Descriptions .....	9-5
9.3.2.1	Memory Interface Signals.....	9-5
9.3.2.2	Clock Interface Signals.....	9-7
9.3.2.3	Debug Signals.....	9-8
9.4	Memory Map/Register Definition .....	9-8
9.4.1	Register Descriptions.....	9-9
9.4.1.1	Chip Select Memory Bounds (CS <sub>n</sub> _BNDS).....	9-9
9.4.1.2	Chip Select Configuration (CS <sub>n</sub> _CONFIG).....	9-10
9.4.1.3	DDR SDRAM Timing Configuration 3 (TIMING_CFG_3).....	9-11
9.4.1.4	DDR SDRAM Timing Configuration 0 (TIMING_CFG_0).....	9-12
9.4.1.5	DDR SDRAM Timing Configuration 1 (TIMING_CFG_1).....	9-14
9.4.1.6	DDR SDRAM Timing Configuration 2 (TIMING_CFG_2).....	9-16
9.4.1.7	DDR SDRAM Control Configuration (DDR_SDRAM_CFG).....	9-18
9.4.1.8	DDR SDRAM Control Configuration 2 (DDR_SDRAM_CFG_2).....	9-21
9.4.1.9	DDR SDRAM Mode Configuration (DDR_SDRAM_MODE).....	9-22
9.4.1.10	DDR SDRAM Mode 2 Configuration (DDR_SDRAM_MODE_2).....	9-23
9.4.1.11	DDR SDRAM Mode Control Register (DDR_SDRAM_MD_CNTL).....	9-24
9.4.1.12	DDR SDRAM Interval Configuration (DDR_SDRAM_INTERVAL) .....	9-26
9.4.1.13	DDR SDRAM Data Initialization (DDR_DATA_INIT) .....	9-27
9.4.1.14	DDR SDRAM Clock Control (DDR_SDRAM_CLK_CNTL) .....	9-27
9.4.1.15	DDR Initialization Address (DDR_INIT_ADDR).....	9-28
9.4.1.16	DDR IP Block Revision 1 (DDR_IP_REV1).....	9-28
9.4.1.17	DDR IP Block Revision 2 (DDR_IP_REV2).....	9-29
9.5	Functional Description.....	9-29
9.5.1	DDR SDRAM Interface Operation.....	9-32
9.5.1.1	Supported DDR SDRAM Organizations .....	9-33
9.5.2	DDR SDRAM Address Multiplexing.....	9-34
9.5.3	JEDEC Standard DDR SDRAM Interface Commands .....	9-37
9.5.4	DDR SDRAM Interface Timing .....	9-39
9.5.4.1	Clock Distribution .....	9-42
9.5.5	DDR SDRAM Mode-Set Command Timing.....	9-42
9.5.6	DDR SDRAM Registered DIMM Mode .....	9-43
9.5.7	DDR SDRAM Write Timing Adjustments.....	9-43
9.5.8	DDR SDRAM Refresh .....	9-44
9.5.8.1	DDR SDRAM Refresh Timing.....	9-45
9.5.8.2	DDR SDRAM Refresh and Power-Saving Modes.....	9-45
9.5.8.2.1	Self-Refresh in Sleep Mode.....	9-47
9.5.9	DDR Data Beat Ordering.....	9-48
9.5.10	Page Mode and Logical Bank Retention .....	9-48

# Contents

Paragraph Number	Title	Page Number
9.6	Initialization/Application Information .....	9-49
9.6.1	Programming Differences Between Memory Types.....	9-50
9.6.2	DDR SDRAM Initialization Sequence .....	9-53

## Chapter 10 Enhanced Local Bus Controller

10.1	Introduction.....	10-1
10.1.1	Overview.....	10-2
10.1.2	Features.....	10-2
10.1.3	Modes of Operation .....	10-3
10.1.3.1	eLBC Bus Clock and Clock Ratios .....	10-3
10.1.3.2	Source ID Debug Mode .....	10-4
10.2	External Signal Descriptions .....	10-4
10.3	Memory Map/Register Definition .....	10-7
10.3.1	Register Descriptions.....	10-9
10.3.1.1	Base Registers (BR0–BR3) .....	10-9
10.3.1.2	Option Registers (OR0–OR3).....	10-11
10.3.1.2.1	Address Mask .....	10-11
10.3.1.2.2	Option Registers (OR $n$ )—GPCM Mode .....	10-13
10.3.1.2.3	Option Registers (OR $n$ )—FCM Mode .....	10-15
10.3.1.2.4	Option Registers (OR $n$ )—UPM Mode .....	10-18
10.3.1.3	UPM Memory Address Register (MAR).....	10-19
10.3.1.4	UPM Mode Registers (M $x$ MR) .....	10-20
10.3.1.5	Memory Refresh Timer Prescaler Register (MRTPR) .....	10-22
10.3.1.6	UPM/FCM Data Register (MDR) .....	10-22
10.3.1.7	Special Operation Initiation Register (LSOR).....	10-23
10.3.1.8	UPM Refresh Timer (LURT).....	10-24
10.3.1.9	Transfer Error Status Register (LTESR).....	10-25
10.3.1.10	Transfer Error Check Disable Register (LTEDR).....	10-27
10.3.1.11	Transfer Error Interrupt Enable Register (LTEIR) .....	10-28
10.3.1.12	Transfer Error Attributes Register (LTEATR) .....	10-29
10.3.1.13	Transfer Error Address Register (LTEAR).....	10-30
10.3.1.14	Transfer Error ECC Register (LTECCR).....	10-30
10.3.1.15	Local Bus Configuration Register (LBCR) .....	10-31
10.3.1.16	Clock Ratio Register (LCRR).....	10-33
10.3.1.17	Flash Mode Register (FMR).....	10-34
10.3.1.18	Flash Instruction Register (FIR) .....	10-35
10.3.1.19	Flash Command Register (FCR) .....	10-36
10.3.1.20	Flash Block Address Register (FBAR).....	10-37
10.3.1.21	Flash Page Address Register (FPAR) .....	10-37

# Contents

Paragraph Number	Title	Page Number
10.3.1.22	Flash Byte Count Register (FBCR) .....	10-39
10.3.1.23	Flash ECC Block $n$ Register (FECC0–FECC3) .....	10-39
10.4	Functional Description .....	10-40
10.4.1	Basic Architecture .....	10-41
10.4.1.1	Address and Address Space Checking .....	10-41
10.4.1.2	External Address Latch Enable Signal (LALE) .....	10-41
10.4.1.3	Data Transfer Acknowledge (TA) .....	10-43
10.4.1.4	Data Buffer Control (LBCTL) .....	10-44
10.4.1.5	Atomic Operation .....	10-44
10.4.1.6	Bus Monitor .....	10-44
10.4.2	General-Purpose Chip-Select Machine (GPCM) .....	10-45
10.4.2.1	GPCM Read Signal Timing .....	10-46
10.4.2.2	GPCM Write Signal Timing .....	10-47
10.4.2.3	Chip-Select Assertion Timing .....	10-49
10.4.2.3.1	Programmable Wait State Configuration .....	10-50
10.4.2.3.2	Chip-Select and Write Enable Negation Timing .....	10-50
10.4.2.3.3	Relaxed Timing .....	10-51
10.4.2.3.4	Output Enable (LOE) Timing .....	10-54
10.4.2.3.5	Extended Hold Time on Read Accesses .....	10-54
10.4.2.4	External Access Termination (LGTA) .....	10-55
10.4.2.5	GPCM Boot Chip-Select Operation .....	10-56
10.4.3	Flash Control Machine (FCM) .....	10-57
10.4.3.1	FCM Buffer RAM .....	10-59
10.4.3.1.1	Buffer Layout and Page Mapping for Small-Page NAND Flash Devices .....	10-59
10.4.3.1.2	Buffer Layout and Page Mapping for Large-Page NAND Flash Devices .....	10-60
10.4.3.1.3	Error Correcting Codes and the Spare Region .....	10-61
10.4.3.2	Programming FCM .....	10-62
10.4.3.2.1	FCM Command Instructions .....	10-63
10.4.3.2.2	FCM No-Operation Instruction .....	10-64
10.4.3.2.3	FCM Address Instructions .....	10-64
10.4.3.2.4	FCM Data Read Instructions .....	10-64
10.4.3.2.5	FCM Data Write Instructions .....	10-65
10.4.3.3	FCM Signal Timing .....	10-65
10.4.3.3.1	FCM Chip-Select Timing .....	10-65
10.4.3.3.2	FCM Command, Address, and Write Data Timing .....	10-66
10.4.3.3.3	FCM Ready/Busy Timing .....	10-67
10.4.3.3.4	FCM Read Data Timing .....	10-68
10.4.3.3.5	FCM Extended Read Hold Timing .....	10-69
10.4.3.4	FCM Boot Chip-Select Operation .....	10-69
10.4.3.4.1	FCM Bank 0 Reset Initialization .....	10-70
10.4.3.4.2	Boot Block Loading into the FCM Buffer RAM .....	10-70

# Contents

Paragraph Number	Title	Page Number
10.4.4	User-Programmable Machines (UPMs).....	10-72
10.4.4.1	UPM Requests .....	10-73
10.4.4.1.1	Memory Access Requests.....	10-74
10.4.4.1.2	UPM Refresh Timer Requests .....	10-74
10.4.4.1.3	Software Requests—RUN Command .....	10-75
10.4.4.1.4	Exception Requests.....	10-75
10.4.4.2	Programming the UPMs .....	10-75
10.4.4.2.1	UPM Programming Example (Two Sequential Writes to the RAM Array)....	10-76
10.4.4.2.2	UPM Programming Example (Two Sequential Reads from the RAM Array)	10-77
10.4.4.3	UPM Signal Timing .....	10-77
10.4.4.4	RAM Array .....	10-78
10.4.4.4.1	RAM Words.....	10-79
10.4.4.4.2	Chip-Select Signal Timing (CST <sub>n</sub> ) .....	10-82
10.4.4.4.3	Byte Select Signal Timing (BST <sub>n</sub> ).....	10-83
10.4.4.4.4	General-Purpose Signals (GnT <sub>n</sub> , GOn).....	10-83
10.4.4.4.5	Loop Control (LOOP) .....	10-83
10.4.4.4.6	Repeat Execution of Current RAM Word (REDO).....	10-84
10.4.4.4.7	Address Multiplexing (AMX) .....	10-84
10.4.4.4.8	Data Valid and Data Sample Control (UTA) .....	10-85
10.4.4.4.9	LGPL[0:5] Signal Negation (LAST).....	10-86
10.4.4.4.10	Wait Mechanism (WAEN).....	10-86
10.4.4.5	Synchronous Sampling of LUPWAIT for Early Transfer Acknowledge .....	10-87
10.4.4.6	Extended Hold Time on Read Accesses .....	10-87
10.5	Initialization/Application Information.....	10-88
10.5.1	Interfacing to Peripherals in Different Address Modes .....	10-88
10.5.1.1	Multiplexed Address/Data Bus for 26-Bit Addressing.....	10-88
10.5.1.2	Non-Multiplexed Address and Data Buses.....	10-88
10.5.1.3	Multiplexed Address and Data to Save Maximum Pins in 8- to 16-Bit Addressing .....	10-89
10.5.1.4	Peripheral Hierarchy on the Local Bus for High Bus Speeds .....	10-89
10.5.1.5	GPCM Timings.....	10-90
10.5.2	Bus Turnaround .....	10-91
10.5.2.1	Address Phase after Previous Read .....	10-91
10.5.2.2	Read Data Phase after Address Phase .....	10-91
10.5.2.3	Read-Modify-Write Cycle for Parity Protected Memory Banks .....	10-92
10.5.2.4	UPM Cycles with Additional Address Phases.....	10-92
10.5.3	Interface to Different Port-Size Devices.....	10-92
10.5.4	Command Sequence Examples for NAND Flash EEPROM.....	10-93
10.5.4.1	NAND Flash Soft Reset Command Sequence Example .....	10-94
10.5.4.2	NAND Flash Read Status Command Sequence Example .....	10-94
10.5.4.3	NAND Flash Read Identification Command Sequence Example .....	10-94



# Contents

Paragraph Number	Title	Page Number
10.5.4.4	NAND Flash Page Read Command Sequence Example .....	10-95
10.5.4.5	NAND Flash Block Erase Command Sequence Example .....	10-96
10.5.4.6	NAND Flash Program Command Sequence Example .....	10-96
10.5.5	Interfacing to Fast-Page Mode DRAM Using UPM .....	10-97
10.5.6	Interfacing to ZBT SRAM Using UPM.....	10-102

## Chapter 11 Sequencer

11.1	Overview.....	11-1
11.1.1	Features.....	11-2
11.2	External Signal Description .....	11-2
11.3	Memory Map/Register Definition .....	11-2
11.4	Register Descriptions .....	11-3
11.4.1	PCI Outbound Translation Address Registers (POTAR <sub>n</sub> ).....	11-3
11.4.2	PCI Outbound Base Address Registers (POBAR <sub>n</sub> ) .....	11-3
11.4.3	PCI Outbound Comparison Mask Registers (POCMR <sub>n</sub> ) .....	11-4
11.4.4	Power Management Control Register (PMCR) .....	11-5
11.4.5	Discard Timer Control Register (DTCR) .....	11-6
11.5	Functional Description.....	11-6
11.5.1	Transaction Forwarding .....	11-6
11.5.1.1	Transactions from the Coherency System Bus (CSB) Port .....	11-7
11.5.1.2	Transactions from the PCI Port .....	11-7
11.5.1.3	Transactions from the DMA Port .....	11-7
11.5.2	PCI Outbound Address Translation .....	11-7
11.5.3	Transaction Ordering .....	11-8

## Chapter 12 DMA/Messaging Unit

12.1	Features.....	12-1
12.2	Memory Map/Register Definition .....	12-2
12.3	Register Descriptions .....	12-3
12.3.1	Outbound Message Interrupt Status Register (OMISR).....	12-3
12.3.2	Outbound Message Interrupt Mask Register (OMIMR).....	12-4
12.3.3	Inbound Message Registers (IMR0–IMR1) .....	12-5
12.3.4	Outbound Message Registers (OMR0–OMR1).....	12-5
12.3.5	Doorbell Registers .....	12-6
12.3.5.1	Outbound Doorbell Register (ODR).....	12-6
12.3.5.2	Inbound Doorbell Register (IDR).....	12-7
12.3.6	Inbound Message Interrupt Status Register (IMISR) .....	12-7

# Contents

Paragraph Number	Title	Page Number
12.3.7	Inbound Message Interrupt Mask Register (IMIMR).....	12-8
12.3.8	DMA Registers .....	12-9
12.3.8.1	DMA Mode Register (DMAMR <sub>n</sub> ) .....	12-9
12.3.8.2	DMA Status Register (DMASR <sub>n</sub> ) .....	12-11
12.3.8.3	DMA Current Descriptor Address Register (DMACDAR <sub>n</sub> ) .....	12-12
12.3.8.4	DMA Source Address Register (DMASAR <sub>n</sub> ).....	12-13
12.3.8.5	DMA Destination Address Register (DMADAR <sub>n</sub> ).....	12-13
12.3.8.6	DMA Byte Count Register (DMABCR <sub>n</sub> ) .....	12-14
12.3.8.7	DMA Next Descriptor Address Register (DMANDAR <sub>n</sub> ).....	12-14
12.3.8.8	DMA General Status Register (DMAGSR).....	12-15
12.4	Functional Description.....	12-15
12.4.1	Message Unit .....	12-15
12.4.1.1	Messaging Registers (IMR0–IMR1, OMR0–OMR1) .....	12-15
12.4.1.2	Doorbell Registers (IDR and ODR) .....	12-16
12.4.2	DMA Controller.....	12-16
12.4.3	DMA Operation .....	12-16
12.4.3.1	DMA Coherency.....	12-17
12.4.3.2	Halt and Error Conditions.....	12-17
12.4.4	DMA Segment Descriptors.....	12-18
12.4.4.1	Descriptor in Big-Endian Mode.....	12-19
12.4.4.2	Descriptor in Little-Endian Mode.....	12-20
12.5	Initialization/Application Information .....	12-20
12.5.1	Initialization Steps in Direct Mode .....	12-20
12.5.2	Initialization Steps in Chaining Mode .....	12-20

## Chapter 13 PCI Bus Interface

13.1	Introduction.....	13-1
13.1.1	Features .....	13-3
13.1.2	Modes of Operation .....	13-3
13.1.2.1	Host/Agent Mode Configuration .....	13-3
13.1.2.2	PCI Arbiter Configuration .....	13-4
13.2	External Signal Description .....	13-4
13.3	Memory Map/Register Definitions .....	13-11
13.3.1	PCI Configuration Access Registers.....	13-12
13.3.1.1	PCI_CONFIG_ADDRESS .....	13-13
13.3.1.2	PCI_CONFIG_DATA.....	13-14
13.3.1.3	PCI Interrupt Acknowledge Register (PCI_INT_ACK).....	13-15
13.3.2	PCI Memory-Mapped Control and Status Registers .....	13-15
13.3.2.1	PCI Error Status Register (PCI_ESR) .....	13-15

# Contents

Paragraph Number	Title	Page Number
13.3.2.2	PCI Error Capture Disable Register (PCI_ECDDR).....	13-16
13.3.2.3	PCI Error Enable Register (PCI_EER).....	13-17
13.3.2.4	PCI Error Attributes Capture Register (PCI_EATCR).....	13-18
13.3.2.5	PCI Error Address Capture Register (PCI_EACR).....	13-19
13.3.2.6	PCI Error Extended Address Capture Register (PCI_EEACR).....	13-20
13.3.2.7	PCI Error Data Low Capture Register (PCI_EDLCR).....	13-20
13.3.2.8	PCI General Control Register (PCI_GCR).....	13-20
13.3.2.9	PCI Error Control Register (PCI_ECR).....	13-21
13.3.2.10	PCI General Status Register (PCI_GSR).....	13-22
13.3.2.11	PCI Inbound Translation Address Registers (PITAR <sub>n</sub> ).....	13-22
13.3.2.12	PCI Inbound Base Address Registers (PIBAR <sub>n</sub> ).....	13-23
13.3.2.13	PCI Inbound Extended Base Address Registers (PIEBAR <sub>n</sub> ).....	13-24
13.3.2.14	PCI Inbound Window Attribute Registers (PIWAR <sub>n</sub> ).....	13-24
13.3.3	PCI Configuration Space Registers.....	13-25
13.3.3.1	Vendor ID Configuration Register.....	13-27
13.3.3.2	Device ID Configuration Register.....	13-27
13.3.3.3	PCI Command Configuration Register.....	13-28
13.3.3.4	PCI Status Configuration Register.....	13-29
13.3.3.5	Revision ID Configuration Register.....	13-30
13.3.3.6	Standard Programming Interface Configuration Register.....	13-30
13.3.3.7	Subclass Code Configuration Register.....	13-31
13.3.3.8	Base Class Code Configuration Register.....	13-31
13.3.3.9	Cache Line Size Configuration Register.....	13-32
13.3.3.10	Latency Timer Configuration Register.....	13-32
13.3.3.11	Header Type Configuration Register.....	13-33
13.3.3.12	BIST Control Configuration Register.....	13-33
13.3.3.13	PIMMR Base Address Configuration Register.....	13-33
13.3.3.14	GPL Base Address Register 0.....	13-34
13.3.3.15	GPL Base Address Registers 1–2.....	13-34
13.3.3.16	GPL Extended Base Address Registers 1–2.....	13-35
13.3.3.17	Subsystem Vendor ID Configuration Register.....	13-36
13.3.3.18	Subsystem Device ID Configuration Register.....	13-36
13.3.3.19	Capabilities Pointer Configuration Register.....	13-36
13.3.3.20	Interrupt Line Configuration Register.....	13-37
13.3.3.21	Interrupt Pin Configuration Register.....	13-37
13.3.3.22	Minimum Grant Configuration Register.....	13-37
13.3.3.23	Maximum Latency Configuration Register.....	13-38
13.3.3.24	PCI Function Configuration Register.....	13-38
13.3.3.25	PCI Arbiter Control Register (PCIACR).....	13-39
13.3.3.26	Hot Swap Register Block.....	13-40
13.3.3.27	PCI Power Management Register 0 (PCIPMR0).....	13-41

# Contents

Paragraph Number	Title	Page Number
13.3.3.28	PCI Power Management Register 1 (PCIPMR1) .....	13-42
13.4	Functional Description .....	13-43
13.4.1	PCI Bus Arbitration .....	13-43
13.4.1.1	Bus Parking .....	13-44
13.4.1.2	Arbitration Algorithm .....	13-44
13.4.1.3	Broken Master Lock-Out .....	13-45
13.4.1.4	Master Latency Timer .....	13-45
13.4.2	Bus Commands .....	13-46
13.4.3	PCI Protocol Fundamentals .....	13-47
13.4.3.1	Basic Transfer Control .....	13-47
13.4.3.2	Addressing .....	13-47
13.4.3.3	Device Selection .....	13-48
13.4.3.4	Byte Enable Signals .....	13-48
13.4.3.5	Bus Driving and Turnaround .....	13-48
13.4.3.6	Bus Transactions .....	13-49
13.4.3.7	Read and Write Transactions .....	13-49
13.4.3.8	Transaction Termination .....	13-51
13.4.4	Other Bus Operations .....	13-53
13.4.4.1	Fast Back-to-Back Transactions .....	13-53
13.4.4.2	Dual Address Cycles .....	13-54
13.4.4.3	Data Streaming .....	13-54
13.4.4.4	Host Mode Configuration Access .....	13-54
13.4.4.5	Agent Mode Configuration Access .....	13-55
13.4.4.6	Special Cycle Command .....	13-55
13.4.4.7	Interrupt Acknowledge .....	13-56
13.4.5	Error Functions .....	13-57
13.4.5.1	Parity .....	13-57
13.4.5.2	Error Reporting .....	13-57
13.4.6	PCI Inbound Address Translation .....	13-59
13.4.7	CompactPCI Hot Swap Specification Support .....	13-60
13.5	Initialization/Application Information .....	13-60
13.5.1	Initialization Sequence for Host Mode .....	13-60
13.5.2	Initialization Sequence for Agent Mode .....	13-60

## Chapter 14 Security Engine (SEC) 2.2

14.1	SEC 2.2 Architecture Overview .....	14-2
14.1.1	Descriptors .....	14-3
14.1.2	Execution Units (EUs) .....	14-5
14.1.2.1	Data Encryption Standard Execution Unit (DEU) .....	14-5

# Contents

Paragraph Number	Title	Page Number
14.1.2.2	Advanced Encryption Standard Execution Unit (AESU).....	14-5
14.1.2.3	Message Digest Execution Unit (MDEU) .....	14-5
14.1.3	Channel .....	14-6
14.1.4	SEC Controller.....	14-7
14.1.4.1	Channel-Controlled Access .....	14-7
14.1.4.2	Host-Controlled Access .....	14-8
14.2	Configuration of Internal Memory Space .....	14-8
14.3	Descriptor Overview .....	14-10
14.3.1	Descriptor Structure .....	14-11
14.3.2	Descriptor Format: Header Dword .....	14-11
14.3.2.1	Selecting Execution Units—EU_SEL0 and EU_SEL1 .....	14-13
14.3.2.2	Selecting Descriptor Type—DESC_TYPE .....	14-14
14.3.3	Descriptor Format: Pointer Dwords.....	14-15
14.3.4	Link Table Format .....	14-16
14.3.5	Descriptor Types .....	14-18
14.4	Execution Units.....	14-19
14.4.1	Data Encryption Standard Execution Unit (DEU).....	14-19
14.4.1.1	DEU Mode Register (DEUMR) .....	14-19
14.4.1.2	DEU Key Size Register (DEUKSR).....	14-20
14.4.1.3	DEU Data Size Register (DEUDSR).....	14-21
14.4.1.4	DEU Reset Control Register (DEURCR).....	14-22
14.4.1.5	DEU Status Register (DEUSR) .....	14-23
14.4.1.6	DEU Interrupt Status Register (DEUISR) .....	14-24
14.4.1.7	DEU Interrupt Control Register (DEUICR).....	14-25
14.4.1.8	DEU End-of-Message Register (DEUEMR).....	14-27
14.4.1.9	DEU IV Register (DEUIV) .....	14-27
14.4.1.10	DEU Key Registers (DEUK <sub>n</sub> ).....	14-28
14.4.1.11	DEU FIFOs .....	14-28
14.4.2	Message Digest Execution Unit (MDEU) .....	14-28
14.4.2.1	MDEU Mode Register (MDEUMR) .....	14-28
14.4.2.2	Recommended Settings for MDEUMR.....	14-31
14.4.2.3	MDEU Key Size Register (MDEUKSR) .....	14-32
14.4.2.4	MDEU Data Size Register (MDEUDSR).....	14-32
14.4.2.5	MDEU Reset Control Register (MDEURCR).....	14-33
14.4.2.6	MDEU Status Register (MDEUSR) .....	14-34
14.4.2.7	MDEU Interrupt Status Register (MDEUISR).....	14-35
14.4.2.8	MDEU Interrupt Control Register (MDEUICR).....	14-36
14.4.2.9	MDEU ICV Size Register .....	14-37
14.4.2.10	MDEU End-of-Message Register (MDEUEMR).....	14-38
14.4.2.11	MDEU Context Registers .....	14-38
14.4.2.12	MDEU Key Registers .....	14-39

# Contents

Paragraph Number	Title	Page Number
14.4.2.13	MDEU FIFOs .....	14-40
14.4.3	Advanced Encryption Standard Execution Unit (AESU).....	14-40
14.4.3.1	AESU Mode Register (AESUMR).....	14-40
14.4.3.2	AESU Key Size Register (AESUKSR) .....	14-42
14.4.3.3	AESU Data Size Register (AESUDSR) .....	14-43
14.4.3.4	AESU Reset Control Register (AESURCR) .....	14-43
14.4.3.5	AESU Status Register (AESUSR).....	14-44
14.4.3.6	AESU Interrupt Status Register (AESUISR).....	14-45
14.4.3.7	AESU Interrupt Control Register (AESUICR).....	14-47
14.4.3.8	AESU End-of-Message Register (AESUEMR) .....	14-48
14.4.3.9	AESU Context Registers .....	14-49
14.4.3.9.1	Context for CBC Mode.....	14-50
14.4.3.9.2	Context for Counter Mode.....	14-50
14.4.3.9.3	Context for SRT Mode .....	14-50
14.4.3.9.4	Context for CCM Mode.....	14-50
14.4.3.9.5	AESU Key Registers .....	14-53
14.4.3.9.6	AESU FIFOs.....	14-53
14.5	Channel .....	14-53
14.5.1	Channel Registers .....	14-55
14.5.1.1	Crypto-Channel Configuration Register (CCCR) .....	14-55
14.5.1.2	Crypto-Channel Pointer Status Register (CCPSR).....	14-57
14.5.1.3	Crypto-Channel Current Descriptor Pointer Register (CDPR) .....	14-62
14.5.1.4	Fetch FIFO (FF).....	14-62
14.5.1.5	Descriptor Buffer (DB).....	14-63
14.5.2	Channel Interrupts.....	14-64
14.5.2.1	Channel Done Interrupt .....	14-64
14.5.2.2	Channel Error Interrupt.....	14-64
14.5.2.3	Channel Reset .....	14-64
14.6	Controller .....	14-64
14.6.1	Assignment of EUs to Channel.....	14-65
14.6.2	Bus Interface .....	14-65
14.6.2.1	Arbitration for Use of the Controller and Buses.....	14-65
14.6.2.2	Master Read .....	14-66
14.6.2.3	Master Write .....	14-66
14.6.3	Controller Interrupts .....	14-66
14.6.4	Controller Registers .....	14-67
14.6.4.1	EU Assignment Status Register (EUASR).....	14-67
14.6.4.2	Interrupt Mask Register (IMR).....	14-68
14.6.4.3	Interrupt Status Register (ISR) .....	14-70
14.6.4.4	Interrupt Clear Register (ICR).....	14-71
14.6.4.5	Identification Register (ID).....	14-73

# Contents

Paragraph Number	Title	Page Number
14.6.4.6	IP Block Revision Register.....	14-73
14.6.4.7	Master Control Register (MCR).....	14-74
14.6.5	Snooping by Caches.....	14-74
14.6.6	Interrupts.....	14-75
14.7	Power Saving Mode.....	14-75

## Chapter 15 Enhanced Three-Speed Ethernet Controllers

15.1	Overview.....	15-1
15.2	Features.....	15-2
15.3	Modes of Operation.....	15-4
15.4	External Signals Description.....	15-6
15.4.1	Detailed Signal Descriptions.....	15-8
15.5	Memory Map/Register Definition.....	15-11
15.5.1	Top-Level Module Memory Map.....	15-11
15.5.2	Detailed Memory Map.....	15-12
15.5.3	Memory-Mapped Register Descriptions.....	15-22
15.5.3.1	eTSEC General Control and Status Registers.....	15-22
15.5.3.1.1	Controller ID Register (TSEC_ID).....	15-22
15.5.3.1.2	Controller ID Register (TSEC_ID2).....	15-23
15.5.3.1.3	Interrupt Event Register (IEVENT).....	15-24
15.5.3.1.4	Interrupt Mask Register (IMASK).....	15-27
15.5.3.1.5	Error Disabled Register (EDIS).....	15-29
15.5.3.1.6	Ethernet Control Register (ECNTRL).....	15-31
15.5.3.1.7	Pause Time Value Register (PTV).....	15-33
15.5.3.1.8	DMA Control Register (DMACTRL).....	15-34
15.5.3.2	eTSEC Transmit Control and Status Registers.....	15-35
15.5.3.2.1	Transmit Control Register (TCTRL).....	15-35
15.5.3.2.2	Transmit Status Register (TSTAT).....	15-37
15.5.3.2.3	Default VLAN Control Word Register (DFVLAN).....	15-41
15.5.3.2.4	Transmit Interrupt Coalescing Register (TXIC).....	15-42
15.5.3.2.5	Transmit Queue Control Register (TQUEUE).....	15-43
15.5.3.2.6	TxBD Ring 0–3 Weighting Register (TR03WT).....	15-43
15.5.3.2.7	TxBD Ring 4–7 Weighting Register (TR47WT).....	15-44
15.5.3.2.8	Transmit Data Buffer Pointer High Register (TBDBPH).....	15-45
15.5.3.2.9	Transmit Buffer Descriptor Pointers 0–7 (TBPTR0–TBPTR7).....	15-45
15.5.3.2.10	Transmit Descriptor Base Address Registers (TBASE0–TBASE7).....	15-46
15.5.3.2.11	Transmit Timestamp Identification Register (TMR_TXTS1–2_ID).....	15-47
15.5.3.2.12	Transmit Timestamp Register (TMR_TXTS1–2_H/L).....	15-47

# Contents

Paragraph Number	Title	Page Number
15.5.3.3	eTSEC Receive Control and Status Registers .....	15-48
15.5.3.3.1	Receive Control Register (RCTRL) .....	15-48
15.5.3.3.2	Receive Status Register (RSTAT).....	15-50
15.5.3.3.3	Receive Interrupt Coalescing Register (RXIC) .....	15-52
15.5.3.3.4	Receive Queue Control Register (RQUEUE) .....	15-53
15.5.3.3.5	Receive Bit Field Extract Control Register (RBIFX).....	15-54
15.5.3.3.6	Receive Queue Filer Table Address Register (RQFAR) .....	15-55
15.5.3.3.7	Receive Queue Filer Table Control Register (RQFCR) .....	15-56
15.5.3.3.8	Receive Queue Filer Table Property Register (RQFPR) .....	15-57
15.5.3.3.9	Maximum Receive Buffer Length Register (MRBLR).....	15-61
15.5.3.3.10	Receive Data Buffer Pointer High Register (RBDBPH).....	15-62
15.5.3.3.11	Receive Buffer Descriptor Pointers 0–7 (RBPTR0–RBPTR7) .....	15-62
15.5.3.3.12	Receive Descriptor Base Address Registers (RBASE0–RBASE7) .....	15-63
15.5.3.3.13	Receive Stamp Register (TMR_RXTS_H/L).....	15-63
15.5.3.4	MAC Functionality.....	15-64
15.5.3.4.1	Configuring the MAC .....	15-64
15.5.3.4.2	Controlling CSMA/CD.....	15-64
15.5.3.4.3	Handling Packet Collisions .....	15-65
15.5.3.4.4	Controlling Packet Flow.....	15-65
15.5.3.4.5	Controlling PHY Links.....	15-66
15.5.3.5	MAC Registers .....	15-67
15.5.3.5.1	MAC Configuration 1 Register (MACCFG1).....	15-67
15.5.3.5.2	MAC Configuration 2 Register (MACCFG2).....	15-68
15.5.3.5.3	Inter-Packet Gap/Inter-Frame Gap Register (IPGIFG) .....	15-70
15.5.3.5.4	Half-Duplex Register (HAFDUP) .....	15-71
15.5.3.5.5	Maximum Frame Length Register (MAXFRM) .....	15-72
15.5.3.5.6	MII Management Configuration Register (MIIMCFG) .....	15-72
15.5.3.5.7	MII Management Command Register (MIIMCOM).....	15-73
15.5.3.5.8	MII Management Address Register (MIIMADD).....	15-74
15.5.3.5.9	MII Management Control Register (MIIMCON).....	15-75
15.5.3.5.10	MII Management Status Register (MIIMSTAT) .....	15-75
15.5.3.5.11	MII Management Indicator Register (MIIMIND).....	15-76
15.5.3.5.12	Interface Status Register (IFSTAT).....	15-76
15.5.3.5.13	MAC Station Address Part 1 Register (MACSTNADDR1) .....	15-77
15.5.3.5.14	MAC Station Address Part 2 Register (MACSTNADDR2) .....	15-78
15.5.3.5.15	MAC Exact Match Address 1–15 Part 1 Registers (MAC01ADDR1–MAC15ADDR1).....	15-78
15.5.3.5.16	MAC Exact Match Address 1–15 Part 2 Registers (MAC01ADDR2–MAC15ADDR2).....	15-79
15.5.3.6	MIB Registers .....	15-79
15.5.3.6.1	Transmit and Receive 64-Byte Frame Counter (TR64) .....	15-80



# Contents

Paragraph Number	Title	Page Number
15.5.3.6.2	Transmit and Receive 65- to 127-Byte Frame Counter (TR127) .....	15-80
15.5.3.6.3	Transmit and Receive 128- to 255-Byte Frame Counter (TR255) .....	15-81
15.5.3.6.4	Transmit and Receive 256- to 511-Byte Frame Counter (TR511) .....	15-81
15.5.3.6.5	Transmit and Receive 512- to 1023-Byte Frame Counter (TR1K) .....	15-82
15.5.3.6.6	Transmit and Receive 1024- to 1518-Byte Frame Counter (TRMAX).....	15-82
15.5.3.6.7	Transmit and Receive 1519- to 1522-Byte VLAN Frame Counter (TRMGV) .....	15-83
15.5.3.6.8	Receive Byte Counter (RBYT).....	15-83
15.5.3.6.9	Receive Packet Counter (RPKT).....	15-83
15.5.3.6.10	Receive FCS Error Counter (RFCS) .....	15-84
15.5.3.6.11	Receive Multicast Packet Counter (RMCA) .....	15-84
15.5.3.6.12	Receive Broadcast Packet Counter (RBCA) .....	15-85
15.5.3.6.13	Receive Control Frame Packet Counter (RXCF) .....	15-85
15.5.3.6.14	Receive Pause Frame Packet Counter (RXPF).....	15-86
15.5.3.6.15	Receive Unknown Opcode Packet Counter (RXUO).....	15-86
15.5.3.6.16	Receive Alignment Error Counter (RALN) .....	15-87
15.5.3.6.17	Receive Frame Length Error Counter (RFLR).....	15-87
15.5.3.6.18	Receive Code Error Counter (RCDE) .....	15-88
15.5.3.6.19	Receive Carrier Sense Error Counter (RCSE).....	15-88
15.5.3.6.20	Receive Undersize Packet Counter (RUND).....	15-89
15.5.3.6.21	Receive Oversize Packet Counter (ROVR).....	15-89
15.5.3.6.22	Receive Fragments Counter (RFRG) .....	15-90
15.5.3.6.23	Receive Jabber Counter (RJBR).....	15-90
15.5.3.6.24	Receive Dropped Packet Counter (RDRP).....	15-91
15.5.3.6.25	Transmit Byte Counter (TBYT) .....	15-91
15.5.3.6.26	Transmit Packet Counter (TPKT).....	15-92
15.5.3.6.27	Transmit Multicast Packet Counter (TMCA).....	15-92
15.5.3.6.28	Transmit Broadcast Packet Counter (TBCA).....	15-93
15.5.3.6.29	Transmit Pause Control Frame Counter (TXPF) .....	15-93
15.5.3.6.30	Transmit Deferral Packet Counter (TDFR) .....	15-94
15.5.3.6.31	Transmit Excessive Deferral Packet Counter (TEDF) .....	15-94
15.5.3.6.32	Transmit Single Collision Packet Counter (TSCL) .....	15-95
15.5.3.6.33	Transmit Multiple Collision Packet Counter (TMCL) .....	15-95
15.5.3.6.34	Transmit Late Collision Packet Counter (TLCL).....	15-96
15.5.3.6.35	Transmit Excessive Collision Packet Counter (TXCL).....	15-96
15.5.3.6.36	Transmit Total Collision Counter (TNCL) .....	15-97
15.5.3.6.37	Transmit Drop Frame Counter (TDRP).....	15-97
15.5.3.6.38	Transmit Jabber Frame Counter (TJBR) .....	15-98
15.5.3.6.39	Transmit FCS Error Counter (TFCS) .....	15-98
15.5.3.6.40	Transmit Control Frame Counter (TXCF).....	15-99
15.5.3.6.41	Transmit Oversize Frame Counter (TOVR) .....	15-99

# Contents

Paragraph Number	Title	Page Number
15.5.3.6.42	Transmit Undersize Frame Counter (TUND).....	15-100
15.5.3.6.43	Transmit Fragment Counter (TFRG).....	15-100
15.5.3.6.44	Carry Register 1 (CAR1).....	15-101
15.5.3.6.45	Carry Register 2 (CAR2).....	15-102
15.5.3.6.46	Carry Mask Register 1 (CAM1).....	15-103
15.5.3.6.47	Carry Mask Register 2 (CAM2).....	15-105
15.5.3.6.48	Receive Filer Rejected Packet Counter (RREJ).....	15-106
15.5.3.7	Hash Function Registers.....	15-106
15.5.3.7.1	Individual/Group Address Registers 0–7 (IGADDR $n$ ).....	15-107
15.5.3.7.2	Group Address Registers 0–7 (GADDR $n$ ).....	15-107
15.5.3.8	DMA Attribute Registers.....	15-108
15.5.3.8.1	Attribute Register (ATTR).....	15-108
15.5.3.9	Lossless Flow Control Configuration Registers.....	15-109
15.5.3.9.1	Receive Queue Parameters 0–7 (RQPRM0–PQPRM7).....	15-109
15.5.3.9.2	Receive Free Buffer Descriptor Pointer Registers 0–7 (RFBPTR0–RFBPTR7).....	15-109
15.5.3.10	Hardware Assist for IEEE1588 Compliant Timestamping.....	15-110
15.5.3.10.1	Timer Control Register (TMR_CTRL).....	15-110
15.5.3.10.2	Timer Event Register (TMR_TEVENT).....	15-112
15.5.3.10.3	Timer Event Mask Register (TMR_TEMASK).....	15-114
15.5.3.10.4	Timer PTP Packet Event Register (TMR_PEVENT).....	15-115
15.5.3.10.5	Timer Event Mask Register (TMR_PEMASK).....	15-115
15.5.3.10.6	Timer Status Register (TMR_STAT).....	15-116
15.5.3.10.7	Timer Counter Register (TMR_CNT_H/L).....	15-117
15.5.3.10.8	Timer Drift Compensation Addend Register (TMR_ADD).....	15-117
15.5.3.10.9	Timer Accumulator Register (TMR_ACC).....	15-118
15.5.3.10.10	Timer Prescale Register (TMR_PRSC).....	15-118
15.5.3.10.11	Timer Offset Register (TMROFF_H/L).....	15-119
15.5.3.10.12	Alarm Time Comparator Register (TMR_ALARM1–2_H/L).....	15-119
15.5.3.10.13	Timer Fixed Interval Period Register (TMR_FIPER1–3).....	15-120
15.5.3.10.14	External Trigger Stamp Register (TMR_ETTS1–2_H/L).....	15-121
15.5.4	Ten-Bit Interface (TBI).....	15-122
15.5.4.1	TBI Transmit Process.....	15-122
15.5.4.1.1	Packet Encapsulation.....	15-122
15.5.4.1.2	8B10B Encoding.....	15-122
15.5.4.1.3	Preamble Shortening.....	15-122
15.5.4.2	TBI Receive Process.....	15-122
15.5.4.2.1	Synchronization.....	15-123
15.5.4.2.2	Auto-Negotiation for 1000BASE-X.....	15-123
15.5.4.3	TBI MII Set Register Descriptions.....	15-123
15.5.4.3.1	Control Register (CR).....	15-124

# Contents

Paragraph Number	Title	Page Number
15.5.4.3.2	Status Register (SR).....	15-125
15.5.4.3.3	AN Advertisement Register (ANA) .....	15-126
15.5.4.3.4	AN Link Partner Base Page Ability Register (ANLPBPA).....	15-128
15.5.4.3.5	AN Expansion Register (ANEX) .....	15-129
15.5.4.3.6	AN Next Page Transmit Register (ANNPT).....	15-130
15.5.4.3.7	AN Link Partner Ability Next Page Register (ANLPANP) .....	15-130
15.5.4.3.8	Extended Status Register (EXST) .....	15-131
15.5.4.3.9	Jitter Diagnostics Register (JD).....	15-132
15.5.4.3.10	TBI Control Register (TBICON).....	15-133
15.6	Functional Description.....	15-134
15.6.1	Connecting to Physical Interfaces on Ethernet .....	15-134
15.6.1.1	Media-Independent Interface (MII).....	15-135
15.6.1.2	Reduced Media-Independent Interface (RMII) .....	15-135
15.6.1.3	Reduced Gigabit Media-Independent Interface (RGMII) .....	15-136
15.6.1.4	Reduced Ten-Bit Interface (RTBI) .....	15-137
15.6.1.5	Ethernet Physical Interfaces Signal Summary.....	15-139
15.6.1.6	SGMII Interface.....	15-143
15.6.2	Gigabit Ethernet Controller Channel Operation .....	15-143
15.6.2.1	Initialization Sequence.....	15-143
15.6.2.1.1	Hardware Controlled Initialization .....	15-144
15.6.2.1.2	User Initialization .....	15-144
15.6.2.2	Soft Reset and Reconfiguring Procedure.....	15-145
15.6.2.3	Gigabit Ethernet Frame Transmission .....	15-145
15.6.2.4	Gigabit Ethernet Frame Reception .....	15-147
15.6.2.5	Ethernet Preamble Customization .....	15-148
15.6.2.5.1	User-Defined Preamble Transmission .....	15-148
15.6.2.5.2	User-Visible Preamble Reception.....	15-149
15.6.2.6	RMON Support.....	15-150
15.6.2.7	Frame Recognition.....	15-150
15.6.2.7.1	Destination Address Recognition and Frame Filtering .....	15-151
15.6.2.7.2	Hash Table Algorithm.....	15-152
15.6.2.8	Magic Packet Mode .....	15-154
15.6.2.9	Flow Control.....	15-154
15.6.2.10	Interrupt Handling .....	15-155
15.6.2.10.1	Interrupt Coalescing .....	15-156
15.6.2.10.2	Interrupt Coalescing By Frame Count Threshold.....	15-156
15.6.2.10.3	Interrupt Coalescing By Timer Threshold .....	15-157
15.6.2.11	Inter-Frame Gap Time .....	15-158
15.6.2.12	Internal and External Loop Back.....	15-158
15.6.2.13	Error-Handling Procedure.....	15-158
15.6.3	TCP/IP Off-Load .....	15-160

# Contents

Paragraph Number	Title	Page Number
15.6.3.1	Frame Control Blocks.....	15-161
15.6.3.2	Transmit Path Off-Load and Tx PTP Packet Parsing .....	15-161
15.6.3.3	Receive Path Off-Load .....	15-163
15.6.4	Quality of Service (QoS) Provision .....	15-165
15.6.4.1	Receive Parser .....	15-165
15.6.4.2	Receive Queue Filer .....	15-167
15.6.4.2.1	Filing Rules .....	15-168
15.6.4.2.2	Comparing Properties with Bit Masks.....	15-169
15.6.4.2.3	Special-Case Rules .....	15-170
15.6.4.2.4	Filer Interrupt Events.....	15-170
15.6.4.2.5	Setting Up the Receive Queue Filer Table .....	15-171
15.6.4.2.6	Filer Example—802.1p Priority Filing.....	15-171
15.6.4.2.7	Filer Example—IP Diff-Serv Code Points Filing.....	15-172
15.6.4.2.8	Filer Example—TCP and UDP Port Filing .....	15-172
15.6.4.3	Transmission Scheduling.....	15-173
15.6.4.3.1	Priority-Based Queuing (PBQ).....	15-174
15.6.4.3.2	Modified Weighted Round-Robin Queuing (MWRR) .....	15-174
15.6.5	Lossless Flow Control .....	15-175
15.6.5.1	Back Pressure Determination through Free Buffers .....	15-175
15.6.5.2	Software Use of Hardware-Initiated Back Pressure .....	15-177
15.6.5.2.1	Initialization.....	15-177
15.6.5.2.2	Operation .....	15-177
15.6.6	Hardware Assist for IEEE Std. 1588-CompatibleTimestamping .....	15-178
15.6.6.1	Features.....	15-178
15.6.6.2	Timer Logic Overview.....	15-179
15.6.6.3	Timestamp Insertion on the Received Packets .....	15-179
15.6.6.3.1	Timestamp Point.....	15-179
15.6.6.4	PTP Packet Parsing .....	15-180
15.6.6.4.1	General Purpose Filer Rule.....	15-181
15.6.6.5	Timestamp Insertion on Transmit Packets.....	15-181
15.6.6.5.1	Interrupts.....	15-182
15.6.6.5.2	Error Condition.....	15-183
15.6.6.6	Tx PTP Packet Parsing.....	15-183
15.6.7	Buffer Descriptors.....	15-185
15.6.7.1	Data Buffer Descriptors .....	15-185
15.6.7.2	Transmit Data Buffer Descriptors (TxBD).....	15-186
15.6.7.3	Receive Buffer Descriptors (RxBD).....	15-190
15.7	Initialization/Application Information .....	15-192
15.7.1	Interface Mode Configuration .....	15-192
15.7.1.1	MII Interface Mode.....	15-193
15.7.1.2	RGMII Interface Mode .....	15-196

# Contents

Paragraph Number	Title	Page Number
15.7.1.3	RMII Interface Mode .....	15-200
15.7.1.4	RTBI Interface Mode .....	15-204
15.7.1.5	SGMII Interface Support .....	15-207

## Chapter 16 Universal Serial Bus Interface

16.1	Introduction.....	16-2
16.1.1	Overview.....	16-2
16.1.2	Features.....	16-2
16.1.3	Modes of Operation .....	16-3
16.2	External Signals .....	16-3
16.2.1	UTMI Interface.....	16-4
16.2.2	ULPI Interface .....	16-4
16.2.3	PHY Clocks .....	16-5
16.3	Memory Map/Register Definitions.....	16-5
16.3.1	Capability Registers.....	16-7
16.3.1.1	Capability Registers Length (CAPLENGTH) .....	16-8
16.3.1.2	Host Controller Interface Version (HCIVERSION).....	16-8
16.3.1.3	Host Controller Structural Parameters (HCSPARAMS) .....	16-9
16.3.1.4	Host Controller Capability Parameters (HCCPARAMS).....	16-10
16.3.1.5	Device Controller Interface Version (DCIVERSION)—Non-EHCI.....	16-10
16.3.1.6	Device Controller Capability Parameters (DCCPARAMS)—Non-EHCI.....	16-11
16.3.2	Operational Registers.....	16-11
16.3.2.1	USB Command Register (USBCMD).....	16-12
16.3.2.2	USB Status Register (USBSTS) .....	16-14
16.3.2.3	USB Interrupt Enable Register (USBINTR) .....	16-17
16.3.2.4	Frame Index Register (FRINDEX).....	16-18
16.3.2.5	Control Data Structure Segment Register (CTRLDSSEGMENT).....	16-19
16.3.2.6	Periodic Frame List Base Address Register (PERIODICLISTBASE).....	16-19
16.3.2.7	Device Address Register (DEVICEADDR)—Non-EHCI .....	16-20
16.3.2.8	Current Asynchronous List Address Register (ASYNCLISTADDR).....	16-20
16.3.2.9	Endpoint List Address Register (ENDPOINTLISTADDR)—Non-EHCI .....	16-21
16.3.2.10	Master Interface Data Burst Size Register (BURSTSIZE)—Non-EHCI.....	16-22
16.3.2.11	Transmit FIFO Tuning Controls Register (TXFILLTUNING)—Non-EHCI.....	16-22
16.3.2.12	ULPI Register Access (ULPI VIEWPORT).....	16-24
16.3.2.13	Configure Flag Register (CONFIGFLAG).....	16-26
16.3.2.14	Port Status and Control Register (PORTSC) .....	16-26
16.3.2.15	On-The-Go Status and Control (OTGSC)—Non-EHCI.....	16-31
16.3.2.16	USB Mode Register (USBMODE)—Non-EHCI .....	16-34
16.3.2.17	Endpoint Setup Status Register (ENDPTSETUPSTAT)—Non-EHCI .....	16-35

# Contents

Paragraph Number	Title	Page Number
16.3.2.18	Endpoint Initialization Register (ENDPTPRIME)—Non-EHCI.....	16-35
16.3.2.19	Endpoint Flush Register (ENDPTFLUSH)—Non-EHCI.....	16-36
16.3.2.20	Endpoint Status Register (ENDPTSTATUS)—Non-EHCI .....	16-36
16.3.2.21	Endpoint Complete Register (ENDPTCOMPLETE)—Non-EHCI.....	16-37
16.3.2.22	Endpoint Control Register 0 (ENDPTCTRL0)—Non-EHCI .....	16-38
16.3.2.23	Endpoint Control Register <i>n</i> (ENDPTCTRL <i>n</i> )—Non-EHCI .....	16-39
16.3.2.24	SNOOP1 and SNOOP2—Non-EHCI.....	16-40
16.3.2.25	Age Count Threshold Register (AGE_CNT_THRESH)—Non-EHCI .....	16-41
16.3.2.26	Priority Control Register (PRI_CTRL)—Non-EHCI.....	16-42
16.3.2.27	System Interface Control Register (SI_CTRL)—Non-EHCI.....	16-43
16.3.2.28	USB General Purpose Register (CONTROL)—Non-EHCI.....	16-43
16.4	Functional Description.....	16-45
16.4.1	System Interface .....	16-45
16.4.2	DMA Engine.....	16-46
16.4.3	FIFO RAM Controller .....	16-46
16.4.4	PHY Interface .....	16-46
16.5	Host Data Structures .....	16-47
16.5.1	Periodic Frame List.....	16-47
16.5.2	Asynchronous List Queue Head Pointer.....	16-49
16.5.3	Isochronous (High-Speed) Transfer Descriptor (iT <sub>D</sub> ).....	16-49
16.5.3.1	Next Link Pointer .....	16-50
16.5.3.2	iT <sub>D</sub> Transaction Status and Control List .....	16-51
16.5.3.3	iT <sub>D</sub> Buffer Page Pointer List (Plus) .....	16-51
16.5.4	Split Transaction Isochronous Transfer Descriptor (siT <sub>D</sub> ).....	16-53
16.5.4.1	Next Link Pointer .....	16-53
16.5.4.2	siT <sub>D</sub> Endpoint Capabilities/Characteristics.....	16-53
16.5.4.3	siT <sub>D</sub> Transfer State .....	16-54
16.5.4.4	siT <sub>D</sub> Buffer Pointer List (Plus).....	16-55
16.5.4.5	siT <sub>D</sub> Back Link Pointer .....	16-56
16.5.5	Queue Element Transfer Descriptor (qT <sub>D</sub> ) .....	16-56
16.5.5.1	Next qT <sub>D</sub> Pointer.....	16-57
16.5.5.2	Alternate Next qT <sub>D</sub> Pointer.....	16-57
16.5.5.3	qT <sub>D</sub> Token .....	16-58
16.5.5.4	qT <sub>D</sub> Buffer Page Pointer List .....	16-61
16.5.6	Queue Head.....	16-62
16.5.6.1	Queue Head Horizontal Link Pointer .....	16-62
16.5.6.2	Endpoint Capabilities/Characteristics.....	16-63
16.5.6.3	Transfer Overlay .....	16-65
16.5.7	Periodic Frame Span Traversal Node (FSTN).....	16-66
16.5.7.1	FSTN Normal Path Pointer.....	16-67
16.5.7.2	FSTN Back Path Link Pointer .....	16-67

# Contents

Paragraph Number	Title	Page Number
16.6	Host Operations .....	16-67
16.6.1	Host Controller Initialization .....	16-68
16.6.2	Power Port.....	16-69
16.6.3	Reporting Over-Current .....	16-69
16.6.4	Suspend/Resume .....	16-69
16.6.4.1	Port Suspend/Resume .....	16-70
16.6.5	Schedule Traversal Rules.....	16-71
16.6.6	Periodic Schedule Frame Boundaries vs. Bus Frame Boundaries.....	16-72
16.6.7	Periodic Schedule .....	16-75
16.6.8	Managing Isochronous Transfers Using iTDs .....	16-76
16.6.8.1	Host Controller Operational Model for iTDs .....	16-76
16.6.8.2	Software Operational Model for iTDs .....	16-78
16.6.8.2.1	Periodic Scheduling Threshold.....	16-79
16.6.9	Asynchronous Schedule.....	16-80
16.6.9.1	Adding Queue Heads to Asynchronous Schedule .....	16-81
16.6.9.2	Removing Queue Heads from Asynchronous Schedule.....	16-82
16.6.9.3	Empty Asynchronous Schedule Detection .....	16-84
16.6.9.4	Asynchronous Schedule Traversal: Start Event.....	16-85
16.6.9.5	Reclamation Status Bit (USBSTS Register).....	16-85
16.6.10	Managing Control/Bulk/Interrupt Transfers via Queue Heads.....	16-85
16.6.10.1	Buffer Pointer List Use for Data Streaming with qTDs .....	16-86
16.6.10.2	Adding Interrupt Queue Heads to the Periodic Schedule.....	16-88
16.6.10.3	Managing Transfer Complete Interrupts from Queue Heads .....	16-88
16.6.11	Ping Control.....	16-89
16.6.12	Split Transactions.....	16-90
16.6.12.1	Split Transactions for Asynchronous Transfers.....	16-90
16.6.12.1.1	Asynchronous—Do-Start-Split.....	16-91
16.6.12.1.2	Asynchronous—Do-Complete-Split .....	16-91
16.6.12.2	Split Transaction Interrupt .....	16-92
16.6.12.2.1	Split Transaction Scheduling Mechanisms for Interrupt .....	16-92
16.6.12.2.2	Host Controller Operational Model for FSTNs.....	16-95
16.6.12.2.3	Software Operational Model for FSTNs .....	16-97
16.6.12.2.4	Tracking Split Transaction Progress for Interrupt Transfers .....	16-98
16.6.12.2.5	Split Transaction Execution State Machine for Interrupt .....	16-98
16.6.12.2.6	Periodic Interrupt—Do-Start-Split .....	16-99
16.6.12.2.7	Periodic Interrupt—Do-Complete-Split .....	16-100
16.6.12.2.8	Managing the QH[FrameTag] Field .....	16-103
16.6.12.2.9	Rebalancing the Periodic Schedule .....	16-104
16.6.12.3	Split Transaction Isochronous .....	16-104
16.6.12.3.1	Split Transaction Scheduling Mechanisms for Isochronous .....	16-105
16.6.12.3.2	Tracking Split Transaction Progress for Isochronous Transfers.....	16-108

# Contents

Paragraph Number	Title	Page Number
16.6.12.3.3	Split Transaction Execution State Machine for Isochronous.....	16-110
16.6.12.3.4	Periodic Isochronous—Do-Start-Split.....	16-110
16.6.12.3.5	Periodic Isochronous—Do Complete Split .....	16-112
16.6.12.3.6	Complete-Split for Scheduling Boundary Cases 2a, 2b .....	16-115
16.6.12.3.7	Split Transaction for Isochronous—Processing Example .....	16-116
16.6.13	Port Test Modes .....	16-117
16.6.14	Interrupts.....	16-118
16.6.14.1	Transfer/Transaction Based Interrupts.....	16-119
16.6.14.1.1	Transaction Error .....	16-119
16.6.14.1.2	Serial Bus Babble .....	16-119
16.6.14.1.3	Data Buffer Error .....	16-120
16.6.14.1.4	USB Interrupt (Interrupt on Completion (IOC)) .....	16-121
16.6.14.1.5	Short Packet.....	16-121
16.6.14.2	Host Controller Event Interrupts .....	16-121
16.6.14.2.1	Port Change Events .....	16-121
16.6.14.2.2	Frame List Rollover.....	16-121
16.6.14.2.3	Interrupt on Async Advance.....	16-121
16.6.14.2.4	Host System Error .....	16-122
16.7	Device Data Structures .....	16-122
16.7.1	Endpoint Queue Head.....	16-123
16.7.1.1	Endpoint Capabilities/Characteristics.....	16-124
16.7.1.2	Transfer Overlay .....	16-125
16.7.1.3	Current dTD Pointer .....	16-125
16.7.1.4	Setup Buffer.....	16-125
16.7.2	Endpoint Transfer Descriptor (dTD) .....	16-126
16.8	Device Operational Model.....	16-128
16.8.1	Device Controller Initialization .....	16-128
16.8.2	Port State and Control.....	16-129
16.8.2.1	Bus Reset .....	16-131
16.8.2.2	Suspend/Resume.....	16-132
16.8.2.2.1	Suspend Description .....	16-132
16.8.2.2.2	Suspend Operational Model .....	16-132
16.8.2.2.3	Resume .....	16-132
16.8.3	Managing Endpoints .....	16-132
16.8.3.1	Endpoint Initialization .....	16-133
16.8.3.1.1	Stalling.....	16-133
16.8.3.2	Data Toggle.....	16-134
16.8.3.2.1	Data Toggle Reset.....	16-134
16.8.3.2.2	Data Toggle Inhibit.....	16-134
16.8.3.3	Device Operational Model For Packet Transfers .....	16-135
16.8.3.3.1	Priming Transmit Endpoints.....	16-135



# Contents

Paragraph Number	Title	Page Number
16.8.3.3.2	Priming Receive Endpoints .....	16-135
16.8.3.4	Interrupt/Bulk Endpoint Operational Model .....	16-136
16.8.3.4.1	Interrupt/Bulk Endpoint Bus Response Matrix .....	16-137
16.8.3.5	Control Endpoint Operation Model .....	16-138
16.8.3.5.1	Setup Phase .....	16-138
16.8.3.5.2	Data Phase .....	16-138
16.8.3.5.3	Status Phase .....	16-139
16.8.3.5.4	Control Endpoint Bus Response Matrix .....	16-139
16.8.3.6	Isochronous Endpoint Operational Model .....	16-140
16.8.3.6.1	Isochronous Pipe Synchronization .....	16-141
16.8.3.6.2	Isochronous Endpoint Bus Response Matrix .....	16-142
16.8.4	Managing Queue Heads .....	16-142
16.8.4.1	Queue Head Initialization .....	16-143
16.8.4.2	Operational Model for Setup Transfers .....	16-143
16.8.5	Managing Transfers with Transfer Descriptors .....	16-144
16.8.5.1	Software Link Pointers .....	16-144
16.8.5.2	Building a Transfer Descriptor .....	16-144
16.8.5.3	Executing a Transfer Descriptor .....	16-145
16.8.5.4	Transfer Completion .....	16-145
16.8.5.5	Flushing/De-Priming an Endpoint .....	16-146
16.8.5.6	Device Error Matrix .....	16-146
16.8.6	Servicing Interrupts .....	16-147
16.8.6.1	High-Frequency Interrupts .....	16-147
16.8.6.2	Low-Frequency Interrupts .....	16-147
16.8.6.3	Error Interrupts .....	16-148
16.9	Deviations from the EHCI Specifications .....	16-148
16.9.1	Embedded Transaction Translator Function .....	16-148
16.9.1.1	Capability Registers .....	16-149
16.9.1.2	Operational Registers .....	16-149
16.9.1.3	Discovery .....	16-149
16.9.1.4	Data Structures .....	16-150
16.9.1.5	Operational Model .....	16-150
16.9.1.5.1	Microframe Pipeline .....	16-150
16.9.1.5.2	Split State Machines .....	16-151
16.9.1.5.3	Asynchronous Transaction Scheduling and Buffer Management .....	16-151
16.9.1.5.4	Periodic Transaction Scheduling and Buffer Management .....	16-151
16.9.1.5.5	Multiple Transaction Translators .....	16-152
16.9.2	Device Operation .....	16-152
16.9.3	Non-Zero Fields the Register File .....	16-152
16.9.4	SOF Interrupt .....	16-152
16.9.5	Embedded Design .....	16-153

# Contents

Paragraph Number	Title	Page Number
16.9.5.1	Frame Adjust Register .....	16-153
16.9.6	Miscellaneous Variations from EHCI.....	16-153
16.9.6.1	Programmable Physical Interface Behavior .....	16-153
16.9.6.2	Discovery .....	16-153
16.9.6.2.1	Port Reset.....	16-153
16.9.6.2.2	Port Speed Detection .....	16-154
16.10	Timing Diagrams .....	16-154

## Chapter 17 I<sup>2</sup>C Interfaces

17.1	Introduction.....	17-1
17.1.1	Features.....	17-2
17.1.2	Modes of Operation .....	17-2
17.2	External Signal Descriptions .....	17-3
17.2.1	Signal Overview .....	17-3
17.2.2	Detailed Signal Descriptions .....	17-3
17.3	Memory Map/Register Definition .....	17-4
17.3.1	Register Descriptions.....	17-5
17.3.1.1	I <sup>2</sup> Cn Address Register (I2CnADR) .....	17-5
17.3.1.2	I <sup>2</sup> Cn Frequency Divider Register (I2CnFDR).....	17-6
17.3.1.3	I <sup>2</sup> Cn Control Register (I2CnCR) .....	17-7
17.3.1.4	I <sup>2</sup> Cn Status Register (I2CnSR) .....	17-8
17.3.1.5	I <sup>2</sup> Cn Data Register (I2CnDR).....	17-9
17.3.1.6	Digital Filter Sampling Rate Register (I2CnDFSRR) .....	17-9
17.4	Functional Description.....	17-10
17.4.1	Transaction Protocol .....	17-10
17.4.1.1	START Condition .....	17-11
17.4.1.2	Slave Address Transmission.....	17-11
17.4.1.3	Repeated START Condition .....	17-12
17.4.1.4	STOP Condition.....	17-12
17.4.1.5	Protocol Implementation Details .....	17-12
17.4.1.5.1	Transaction Monitoring—Implementation Details.....	17-12
17.4.1.5.2	Control Transfer—Implementation Details .....	17-12
17.4.1.6	Address Compare—Implementation Details .....	17-13
17.4.2	Arbitration Procedure .....	17-13
17.4.2.1	Arbitration Control .....	17-14
17.4.3	Handshaking .....	17-14
17.4.4	Clock Control.....	17-14
17.4.4.1	Clock Synchronization.....	17-15
17.4.4.2	Input Synchronization and Digital Filter .....	17-15

# Contents

Paragraph Number	Title	Page Number
17.4.4.2.1	Input Signal Synchronization .....	17-15
17.4.4.2.2	Filtering of SCL <sub>n</sub> and SDA <sub>n</sub> Lines .....	17-15
17.4.4.3	Clock Stretching .....	17-15
17.4.5	Boot Sequencer Mode.....	17-15
17.4.5.1	Using the Boot Sequencer for Reset Configuration .....	17-16
17.4.5.2	EEPROM Calling Address .....	17-16
17.4.5.3	EEPROM Data Format .....	17-16
17.4.5.4	Boot Sequencer Done Indication .....	17-19
17.5	Initialization/Application Information .....	17-19
17.5.1	Interrupt Service Routine Flowchart.....	17-19
17.5.2	Initialization Sequence.....	17-21
17.5.3	Generation of START .....	17-21
17.5.4	Post-Transfer Software Response .....	17-21
17.5.5	Generation of STOP.....	17-22
17.5.6	Generation of Repeated START .....	17-22
17.5.7	Generation of SCL <sub>n</sub> When SDA <sub>n</sub> is Negated .....	17-22
17.5.8	Slave Mode Interrupt Service Routine.....	17-22
17.5.8.1	Slave Transmitter and Received Acknowledge .....	17-23
17.5.8.2	Loss of Arbitration and Forcing of Slave Mode.....	17-23

## Chapter 18 DUART

18.1	Overview.....	18-1
18.1.1	Features.....	18-2
18.1.2	Modes of Operation .....	18-3
18.2	External Signal Descriptions .....	18-3
18.2.1	Signal Overview .....	18-3
18.2.2	Detailed Signal Descriptions .....	18-3
18.3	Memory Map/Register Definition .....	18-4
18.3.1	Register Descriptions.....	18-5
18.3.1.1	Receiver Buffer Registers (URBR1 and URBR2).....	18-5
18.3.1.2	Transmitter Holding Registers (UTHR1 and UTHR2).....	18-6
18.3.1.3	Divisor Most and Least Significant Byte Registers (UDMB and UDLB) .....	18-6
18.3.1.4	Interrupt Enable Registers (UIER1 and UIER2) .....	18-8
18.3.1.5	Interrupt ID Registers (UIIR1 and UIIR2) .....	18-9
18.3.1.6	FIFO Control Registers (UFCR1 and UFCR2) .....	18-10
18.3.1.7	Line Control Registers (ULCR1 and ULCR2) .....	18-11
18.3.1.8	MODEM Control Registers (UMCR1 and UMCR2).....	18-13
18.3.1.9	Line Status Registers (ULSR1 and ULSR2) .....	18-14
18.3.1.10	MODEM Status Registers (UMSR1 and UMSR2) .....	18-15

# Contents

Paragraph Number	Title	Page Number
18.3.1.11	Scratch Registers (USCR1 and USCR2) .....	18-16
18.3.1.12	Alternate Function Registers (UAFR1 and UAFR2).....	18-16
18.3.1.13	DMA Status Registers (UDSR1 and UDSR2).....	18-17
18.4	Functional Description.....	18-18
18.4.1	Serial Interface .....	18-19
18.4.1.1	START Bit .....	18-19
18.4.1.2	Data Transfer .....	18-19
18.4.1.3	Parity Bit.....	18-19
18.4.1.4	STOP Bit.....	18-20
18.4.2	Baud-Rate Generator Logic .....	18-20
18.4.3	Local Loopback Mode .....	18-20
18.4.4	Errors .....	18-21
18.4.4.1	Framing Error .....	18-21
18.4.4.2	Parity Error .....	18-21
18.4.4.3	Overrun Error.....	18-21
18.4.5	FIFO Mode .....	18-21
18.4.5.1	FIFO Interrupts .....	18-21
18.4.5.2	DMA Mode Select.....	18-22
18.4.5.3	Interrupt Control Logic.....	18-22
18.5	DUART Initialization/Application Information .....	18-22

## Chapter 19 Serial Peripheral Interface

19.1	Overview.....	19-1
19.2	Introduction.....	19-2
19.2.1	Features .....	19-2
19.2.2	SPI Transmission and Reception Process .....	19-3
19.2.3	Modes of Operation .....	19-3
19.2.3.1	SPI as a Master Device .....	19-3
19.2.3.2	SPI as a Slave Device .....	19-4
19.2.3.3	SPI in Multiple-Master Operation .....	19-5
19.3	External Signal Descriptions .....	19-6
19.3.1	Overview.....	19-7
19.3.2	Detailed Signal Descriptions .....	19-7
19.4	Memory Map/Register Definition .....	19-8
19.4.1	Register Descriptions.....	19-9
19.4.1.1	SPI Mode Register (SPMODE).....	19-9
19.4.1.2	SPI Event Register (SPIE) .....	19-11
19.4.1.3	SPI Mask Register (SPIM) .....	19-12
19.4.1.4	SPI Command Register (SPCOM) .....	19-14

# Contents

Paragraph Number	Title	Page Number
19.4.1.5	SPI Transmit Data Hold Register (SPITD).....	19-14
19.4.1.6	SPI Receive Data Hold Register (SPIRD).....	19-15
19.4.1.6.1	Reverse Mode SPMODE[REV] Examples .....	19-15
19.5	Initialization/Application Information .....	19-16
19.5.1	SPI Master Programming Example .....	19-16
19.5.2	SPI Slave Programming Example.....	19-16

## Chapter 20 JTAG/Testing Support

20.1	Overview.....	20-1
20.2	JTAG Signals .....	20-1
20.2.1	External Signal Descriptions .....	20-2
20.3	JTAG Registers and Scan Chains .....	20-3

## Chapter 21 General Purpose I/O (GPIO)

21.1	Introduction.....	21-1
21.1.1	Overview.....	21-1
21.1.2	Features .....	21-1
21.2	External Signal Description .....	21-2
21.2.1	Signals Overview .....	21-2
21.3	Memory Map/Register Definition .....	21-2
21.3.1	GPIO Direction Register (GPDIR).....	21-3
21.3.2	GPIO Open Drain Register (GPODR).....	21-3
21.3.3	GPIO Data Register (GPDAT).....	21-4
21.3.4	GPIO Interrupt Event Register (GPIER) .....	21-4
21.3.5	GPIO Interrupt Mask Register (GPIMR).....	21-4
21.3.6	GPIO Interrupt Control Register (GPICR) .....	21-5

## Appendix A Revision History

A.1	Changes From Revision 1 to Revision 2 .....	A-1
A.2	Changes From Revision 0 to Revision 1 .....	A-23

## Glossary

## Index



# Contents

**Paragraph  
Number**

**Title**

**Page  
Number**

# Figures

Figure Number	Title	Page Number
1-1	MPC8313E Block Diagram .....	1-1
1-2	MPC8313E Integrated e300c3 Core Block Diagram.....	1-9
1-3	Integrated Security Engine Functional Blocks.....	1-10
1-4	USB Controllers Port Configuration.....	1-13
1-5	MPC8313E Serving as the Main CPU in a Low-End Printer Application .....	1-19
1-6	MPC8313E I/O Processor Implementation in a High-End Printer Application .....	1-20
1-7	IEEE Std. 1588 in Test and Measurement .....	1-21
1-8	IEEE Std. 1588 in Industrial Control .....	1-22
1-9	MPC8313E as a WLAN Access Point.....	1-23
1-10	MPC8313E as a Media Server .....	1-24
3-1	MPC8313E Signal Groupings (1 of 2).....	3-2
3-2	MPC8313E Signal Groupings (2 of 2).....	3-3
4-1	Power-On Reset Flow .....	4-8
4-2	Hard Reset Flow.....	4-9
4-3	Reset Configuration Word Low Register (RCWLR).....	4-13
4-4	Reset Configuration Word High Register (RCWHR).....	4-14
4-5	EEPROM Data Format for Reset Configuration Words Preload Command .....	4-24
4-6	EEPROM Contents .....	4-25
4-7	Clock Subsystem Block Diagram .....	4-29
4-8	Reset Status Register (RSR).....	4-33
4-9	Reset Mode Register (RMR).....	4-35
4-10	Reset Protection Register (RPR).....	4-35
4-11	Reset Control Register (RCR).....	4-36
4-12	Reset Control Enable Register (RCER).....	4-37
4-13	System PLL Mode Register .....	4-38
4-14	Output Clock Control Register (OCCR).....	4-39
4-15	System Clock Control Register (SCCR).....	4-40
5-1	Local Memory Map Example .....	5-2
5-2	Internal Memory Map Registers' Base Address Register (IMMRBAR).....	5-6
5-3	Alternate Configuration Base Address Register (ALTCBAR) .....	5-7
5-4	LBC Local Access Window <i>n</i> Base Address Registers (LBLAWBAR0–LBLAWBAR3) ....	5-8
5-5	LBC Local Access Window <i>n</i> Attributes Registers (LBLAWAR0–LBLAWAR3) .....	5-9
5-6	PCI Local Access Window <i>n</i> Base Address Registers (PCILAWBAR0–PCILAWBAR1) .	5-10
5-7	PCI Local Access Window <i>n</i> Attributes Registers (PCILAWAR0–PCILAWAR1).....	5-11
5-8	DDR Local Access Window <i>n</i> Base Address Registers (DDRLAWBAR0– DDRLAWBAR1) .....	5-12
5-9	DDR Local Access Window <i>n</i> Attributes Registers (DDRLAWAR0–DDRLAWAR1).....	5-13
5-10	System General Purpose Register Low (SGPRL).....	5-17

# Figures

Figure Number	Title	Page Number
5-11	System General Purpose Register High (SGPRH) .....	5-17
5-12	System Part and Revision ID Register (SPRIDR) .....	5-18
5-13	System Priority Configuration Register (SPCR) .....	5-19
5-14	System I/O Configuration Register Low (SICRL) .....	5-21
5-15	System I/O Configuration Register High (SICRH) .....	5-24
5-16	DDR Control Driver Register (DDRCDR).....	5-27
5-17	DDR Debug Status Register (DDRDSR).....	5-28
5-18	Software Watchdog Timer High-Level Block Diagram .....	5-29
5-19	System Watchdog Control Register (SWCRR).....	5-31
5-20	System Watchdog Count Register (SWCNR).....	5-32
5-21	System Watchdog Service Register (SWSRR) .....	5-32
5-22	Software Watchdog Timer Service State Diagram.....	5-33
5-23	Software Watchdog Timer Functional Block Diagram.....	5-34
5-24	Real Time Clock Module High Level Block Diagram .....	5-36
5-25	Real Time Counter Control Register (RTCNR).....	5-37
5-26	Real Time Counter Load Register (RTLDR) .....	5-38
5-27	Real Time Counter Prescale Register (RTPSR).....	5-39
5-28	Real Time Counter Register (RTCTR).....	5-39
5-29	Real Time Counter Event Register (RTEVR) .....	5-40
5-30	Real Time Counter Alarm Register (RTALR) .....	5-40
5-31	Real Time Clock Module Functional Block Diagram .....	5-41
5-32	Periodic Interval Timer High Level Block Diagram.....	5-43
5-33	Periodic Interval Timer Control Register (PTCNR) .....	5-44
5-34	Periodic Interval Timer Load Register (PTLDR).....	5-45
5-35	Periodic Interval Timer Prescale Register (PTPSR) .....	5-46
5-36	Periodic Interval Timer Counter Register (PTCTR) .....	5-46
5-37	Periodic Interval Timer Event Register (PTEVR).....	5-47
5-38	Periodic Interval Timer Functional Block Diagram.....	5-47
5-39	Global Timers Block Diagram .....	5-49
5-40	Global Timers Configuration Register 1 (GTCFR1).....	5-54
5-41	Global Timers Configuration Register 2 (GTCFR2).....	5-55
5-42	Global Timers Mode Registers (GTMDR1–GTMDR4).....	5-57
5-43	Global Timers Reference Registers (GTRFR1–GTRFR4).....	5-58
5-44	Global Timers Capture Registers (GTCPR1–GTCPR4) .....	5-58
5-45	Global Timers Counter Registers (GTCNR1—GTCNR4).....	5-59
5-46	Global Timers Event Registers (GTEVR1—GTEVR4).....	5-59
5-47	Global Timers Prescale Registers (GTPSR1–GTPSR4).....	5-60
5-48	Timers Non-Cascaded Mode Block Diagram .....	5-62
5-49	Timer Pair-Cascaded Mode Block Diagram .....	5-63
5-50	Timers Super-Cascaded Mode Block Diagram.....	5-63
5-51	Power Management Controller Configuration Register .....	5-66



# Figures

Figure Number	Title	Page Number
5-52	Power Management Controller Event Register .....	5-67
5-53	Power Management Controller Mask Register .....	5-69
5-54	Power Management Controller Configuration Register 1 .....	5-70
5-55	Power Management Controller Configuration Register 2 .....	5-71
5-56	Power Segmentation in Deep Sleep Mode.....	5-74
5-57	Power State Transitions Supported .....	5-83
5-58	Example VDD Control of Device as Agent Using the Optional PMC_PWR_OK Signal ...	5-89
5-59	Example VDD Control of Device as Host.....	5-89
6-1	Arbiter Configuration Register (ACR) .....	6-2
6-2	Arbiter Timers Register (ATR) .....	6-4
6-3	Arbiter Event Register (AER).....	6-5
6-4	Arbiter Interrupt Definition Register (AIDR).....	6-6
6-5	Arbiter Mask Register (AMR).....	6-7
6-6	Arbiter Event Attributes Register (AEATR).....	6-8
6-7	Arbiter Event Address Register (AEADR).....	6-9
6-8	Arbiter Event Response Register (AERR).....	6-10
6-9	Address Bus Arbitration.....	6-11
6-10	An Example of Priority-Based Arbitration Algorithm .....	6-12
7-1	e300c3 Core Block Diagram.....	7-2
7-2	e300 Programming Model—Registers.....	7-15
7-3	e300c3 Data Cache Organization.....	7-29
7-4	Core Interface.....	7-37
8-1	Interrupt Sources Block Diagram .....	8-3
8-2	System Global Interrupt Configuration Register (SICFR) .....	8-7
8-3	System Global Interrupt Vector Register (SIVCR).....	8-9
8-4	System Internal Interrupt Pending Register (SIPNR_H) .....	8-11
8-5	System Internal Interrupt Pending Register (SIPNR_L).....	8-12
8-6	System Internal Interrupt Group A Priority Register (SIPRR_A) .....	8-14
8-7	System Internal Interrupt Group D Priority Register (SIPRR_D) .....	8-14
8-8	System Internal Interrupt Mask Register (SIMSR_H).....	8-15
8-9	System Internal Interrupt Mask Register (SIMSR_L) .....	8-16
8-10	System Internal Interrupt Control Register (SICNR) .....	8-17
8-11	System External Interrupt Pending Register (SEPNR).....	8-18
8-12	System Mixed Interrupt Group A Priority Register (SMPRR_A).....	8-18
8-13	System Mixed Interrupt Group B Priority Register (SMPRR_B) .....	8-19
8-14	System External Interrupt Mask Register (SEMSR) .....	8-20
8-15	System External Interrupt Control Register (SECNR) .....	8-21
8-16	System Error Status Register (SERSR).....	8-22
8-17	System Error Mask Register (SERMR).....	8-24
8-18	System Error Control Register (SERCR).....	8-24
8-19	System Internal Interrupt Force Register (SIFCR_H) .....	8-25

# Figures

Figure Number	Title	Page Number
8-20	System Internal Interrupt Force Register (SIFCR_L).....	8-25
8-21	System External Interrupt Force Register (SEFCR).....	8-26
8-22	System Error Status Register (SERFR).....	8-26
8-23	System Critical Interrupt Vector Register (SCVCR).....	8-27
8-24	System Management Interrupt Vector Register (SMVCR).....	8-28
8-25	Interrupt Structure.....	8-29
8-26	DDR Interrupt Request Masking.....	8-35
9-1	DDR Memory Controller Simplified Block Diagram.....	9-2
9-2	Chip Select Bounds Registers (CS <sub>n</sub> _BNDS).....	9-9
9-3	Chip Select Configuration Register (CS <sub>n</sub> _CONFIG).....	9-10
9-4	DDR SDRAM Timing Configuration 3 (TIMING_CFG_3).....	9-12
9-5	DDR SDRAM Timing Configuration 0 (TIMING_CFG_0).....	9-12
9-6	DDR SDRAM Timing Configuration 1 (TIMING_CFG_1).....	9-14
9-7	DDR SDRAM Timing Configuration 2 Register (TIMING_CFG_2).....	9-16
9-8	DDR SDRAM Control Configuration Register (DDR_SDRAM_CFG).....	9-18
9-9	DDR SDRAM Control Configuration Register 2 (DDR_SDRAM_CFG_2).....	9-21
9-10	DDR SDRAM Mode Configuration Register (DDR_SDRAM_MODE).....	9-22
9-11	DDR SDRAM Mode 2 Configuration Register (DDR_SDRAM_MODE_2).....	9-23
9-12	DDR SDRAM Mode Control Register (DDR_SDRAM_MD_CNTL).....	9-24
9-13	DDR SDRAM Interval Configuration Register (DDR_SDRAM_INTERVAL).....	9-26
9-14	DDR SDRAM Data Initialization Configuration Register (DDR_DATA_INIT).....	9-27
9-15	DDR SDRAM Clock Control Configuration Register (DDR_SDRAM_CLK_CNTL).....	9-27
9-16	DDR Initialization Address Configuration Register (DDR_INIT_ADDR).....	9-28
9-17	DDR IP Block Revision 1 (DDR_IP_REV1).....	9-28
9-18	DDR IP Block Revision 2 (DDR_IP_REV2).....	9-29
9-19	DDR Memory Controller Block Diagram.....	9-30
9-20	Typical Dual Data Rate SDRAM Internal Organization.....	9-31
9-21	Typical DDR SDRAM Interface Signals.....	9-31
9-22	Example 64-Mbyte DDR SDRAM Configuration.....	9-32
9-23	DDR SDRAM Burst Read Timing—ACTTORW = 3, MCAS Latency = 2.....	9-40
9-24	DDR SDRAM Single-Beat (Double Word) Write Timing—ACTTOR.....	9-41
9-25	DDR SDRAM Single-Beat (Double Word) Write Timing—ACTTORW = 3.....	9-41
9-26	DDR SDRAM Clock Distribution Example for x8 DDR SDRAMs.....	9-42
9-27	DDR SDRAM Mode-Set Command Timing.....	9-42
9-28	Registered DDR SDRAM DIMM Burst Write Timing.....	9-43
9-29	Write Timing Adjustments Example for Write Latency = 1.....	9-44
9-30	DDR SDRAM Bank Staggered Auto Refresh Timing.....	9-45
9-31	DDR SDRAM Power-Down Mode.....	9-46
9-32	DDR SDRAM Self-Refresh Entry Timing.....	9-47
9-33	DDR SDRAM Self-Refresh Exit Timing.....	9-47
10-1	Enhanced Local Bus Controller Block Diagram.....	10-1

# Figures

Figure Number	Title	Page Number
10-2	Base Registers (BR <sub>n</sub> ) .....	10-10
10-3	Option Registers (OR <sub>n</sub> ) in GPCM Mode.....	10-13
10-4	Option Registers (OR <sub>n</sub> ) in FCM Mode.....	10-15
10-5	Option Registers (OR <sub>n</sub> ) in UPM Mode .....	10-18
10-6	UPM Memory Address Register (MAR) .....	10-19
10-7	UPM Mode Registers (MxMR).....	10-20
10-8	Memory Refresh Timer Prescaler Register (MRTPR).....	10-22
10-9	UPM Data Register in UPM Mode (MDR) .....	10-23
10-10	FCM Data Register in FCM Mode (MDR).....	10-23
10-11	Special Operation Initiation Register (LSOR) .....	10-24
10-12	UPM Refresh Timer (LURT) .....	10-24
10-13	Transfer Error Status Register (LTESR) .....	10-25
10-14	Transfer Error Check Disable Register (LTEDR).....	10-27
10-15	Transfer Error Interrupt Enable Register (LTEIR).....	10-28
10-16	Transfer Error Attributes Register (LTEATR).....	10-29
10-17	Transfer Error Address Register (LTEAR) .....	10-30
10-18	Transfer Error ECC Register (LTECCR) .....	10-31
10-19	Local Bus Configuration Register.....	10-31
10-20	Clock Ratio Register (LCRR) .....	10-33
10-21	Flash Mode Register .....	10-34
10-22	Flash Instruction Register .....	10-36
10-23	Flash Command Register .....	10-36
10-24	Flash Block Address Register .....	10-37
10-25	Flash Page Address Register, Small Page Device (ORx[PGS] = 0) .....	10-37
10-26	Flash Page Address Register, Large Page Device (ORx[PGS] = 1) .....	10-38
10-27	Flash Byte Count Register .....	10-39
10-28	Flash ECC Blockn Register (FECC0–FECC3).....	10-40
10-29	Basic Operation of Memory Controllers in the eLBC .....	10-41
10-30	Example of 8-Bit GPCM Writing 32 Bytes to Address 0x5420 (LCRR[PBYP] = 0).....	10-43
10-31	Basic eLBC Bus Cycle with LALE, TA, and LCS <sub>n</sub> .....	10-43
10-32	Enhanced Local Bus to GPCM Device Interface.....	10-45
10-33	GPCM Basic Read Timing (XACS = 0, ACS = 1x, TRLX = 0, CLKDIV = 4,8).....	10-46
10-34	GPCM General Read Timing Parameters .....	10-46
10-35	GPCM General Write Timing Parameters .....	10-48
10-36	GPCM Basic Write Timing (XACS = 0, ACS = 00, CSNT = 1, SCY = 1, TRLX = 0, CLKDIV = 4, 8) .....	10-50
10-37	GPCM Relaxed Timing Back-to-Back Reads (XACS = 0, ACS = 1x, SCY = 1, CSNT = 0, TRLX = 1, EHTR = 0, CLKDIV = 4, 8) .....	10-52
10-38	GPCM Relaxed Timing Back-to-Back Writes (XACS = 0, ACS = 1x, SCY = 0, CSNT = 0, TRLX = 1, CLKDIV = 4, 8).....	10-52

# Figures

Figure Number	Title	Page Number
10-39	GPCM Relaxed Timing Write (XACS = 0, ACS = 10, SCY = 0, CSNT = 1, TRLX = 1, CLKDIV = 4, 8) .....	10-53
10-40	GPCM Relaxed Timing Write (XACS = 0, ACS = 00, SCY = 1, CSNT = 1, TRLX = 1, CLKDIV = 4, 8) .....	10-53
10-41	GPCM Read Followed by Read (TRLX = 0, EHTR = 0, Fastest Timing) .....	10-54
10-42	GPCM Read Followed by Write (TRLX = 0, EHTR = 1, One-Cycle Extended Hold Time on Reads) .....	10-55
10-43	External Termination of GPCM Access .....	10-56
10-44	Local Bus to 8-bit FCM Device Interface .....	10-58
10-45	FCM Basic Page Read Timing (PGS = 1, CSCT = 0, CST = 0, CHT = 1, RST = 1, SCY = 0, TRLX = 0, EHTR = 1) .....	10-58
10-46	FCM Buffer RAM Memory Map for Small-Page (512-byte page) NAND Flash Devices .....	10-60
10-47	FCM Buffer RAM Memory Map for Large-Page (2-Kbyte page) NAND Flash Devices .....	10-61
10-48	FCM ECC Calculation .....	10-61
10-49	ECC Placement in NAND Flash Spare Regions in Relation to FMR[ECCM] .....	10-62
10-50	FCM Instruction Sequencer Mechanism .....	10-63
10-51	Timing of FCM Command/Address and Write Data Cycles (for TRLX = 0, CHT = 0, CST = 1, SCY = 1, CLKDIV = 4*N) .....	10-66
10-52	Example of FCM Command and Address Timing with Minimum Delay Parameters (for TRLX = 0, CHT = 0, CST = 0, SCY = 0, CLKDIV = 4*N) .....	10-67
10-53	Example of FCM Command and Address Timing with Relaxed Parameters (for TRLX = 1, CHT = 0, CST = 1, SCY = 2, CLKDIV = 4*N) .....	10-67
10-54	FCM Delay Prior to Sampling LFRB State .....	10-68
10-55	FCM Read Data Timing (for TRLX = 0, RST = 0, SCY = 1, CLKDIV = 4*N) .....	10-68
10-56	FCM Read Data Timing with Extended Hold Time (for TRLX = 0, EHTR = 1, RST = 1, SCY = 1, CLKDIV = 4*N) .....	10-69
10-57	FCM Buffer RAM Memory Map During Boot Loading .....	10-71
10-58	User-Programmable Machine Functional Block Diagram .....	10-72
10-59	RAM Array Indexing .....	10-73
10-60	Memory Refresh Timer Request Block Diagram .....	10-74
10-61	UPM Clock Scheme for LCRR[CLKDIV] = 2 .....	10-78
10-62	UPM Clock Scheme for LCRR[CLKDIV] = 4 or 8 .....	10-78
10-63	RAM Array and Signal Generation .....	10-78
10-64	RAM Word Fields .....	10-79
10-65	LCS <sub>n</sub> Signal Selection .....	10-82
10-66	LBS Signal Selection .....	10-83
10-67	UPM Read Access Data Sampling .....	10-86
10-68	Effect of LUPWAIT Signal .....	10-87
10-69	Multiplexed Address/Data Bus for 26-Bit Addressing .....	10-88

# Figures

Figure Number	Title	Page Number
10-70	Non-Multiplexed Address and Data Buses .....	10-89
10-71	Local Bus Peripheral Hierarchy for High Bus Speeds.....	10-90
10-72	GPCM Address Timings .....	10-90
10-73	GPCM Data Timings.....	10-91
10-74	Interface to Different Port-Size Devices .....	10-92
10-75	Single-Beat Read Access to FPM DRAM .....	10-98
10-76	Single-Beat Write Access to FPM DRAM .....	10-99
10-77	Burst Read Access to FPM DRAM Using LOOP (Two Beats Shown).....	10-100
10-78	Refresh Cycle (CBR) to FPM DRAM .....	10-101
10-79	Exception Cycle .....	10-102
10-80	Interface to ZBT SRAM .....	10-103
11-1	I/O Sequencer Block Diagram .....	11-1
11-2	PCI Outbound Translation Address Registers (POTAR <sub>n</sub> ).....	11-3
11-3	PCI Outbound Base Address Registers (POBAR <sub>n</sub> ).....	11-3
11-4	PCI Outbound Comparison Mask Registers (POCMR <sub>n</sub> ) .....	11-4
11-5	Power Management Control Register (PMCR) .....	11-5
11-6	Discard Timer Control Register (DTCR).....	11-6
11-7	Outbound PCI Memory Address Translation .....	11-8
12-1	DMA/Messaging Unit Block Diagram .....	12-1
12-2	Outbound Message Interrupt Status Register (OMISR) .....	12-3
12-3	Outbound Message Interrupt Mask Register (OMIMR).....	12-4
12-4	Inbound Message Registers (IMR0, IMR1).....	12-5
12-5	Outbound Message Registers (OMR0–OMR1) .....	12-5
12-6	Outbound Doorbell Register (ODR) .....	12-6
12-7	Inbound Doorbell Register (IDR) .....	12-7
12-8	Inbound Message Interrupt Status Register (IMISR).....	12-7
12-9	Inbound Message Interrupt Mask Register (IMIMR) .....	12-8
12-10	DMA Mode Register (DMAMR <sub>n</sub> ) .....	12-9
12-11	DMA Status Register (DMASR <sub>n</sub> ) .....	12-11
12-12	DMA Current Descriptor Address Register (DMACDAR <sub>n</sub> ).....	12-12
12-13	DMA Source Address Register (DMASAR <sub>n</sub> ) .....	12-13
12-14	DMA Destination Address Register (DMADAR <sub>n</sub> ) .....	12-13
12-15	DMA Byte Count Register (DMABCR <sub>n</sub> ).....	12-14
12-16	DMA Next Descriptor Address Register (DMANDAR <sub>n</sub> ).....	12-14
12-17	DMA General Status Register (DMAGSR).....	12-15
12-18	DMA Controller Block Diagram .....	12-16
12-19	DMA Chain of Segment Descriptors .....	12-19
13-1	PCI Controller Block Diagram .....	13-2
13-2	PCI Interface External Signals.....	13-5
13-3	PCI_CONFIG_ADDRESS Register .....	13-13
13-4	PCI_CONFIG_DATA .....	13-15

# Figures

Figure Number	Title	Page Number
13-5	PCI Error Status Register (PCI_ESR).....	13-15
13-6	PCI Error Capture Disable Register (PCI_ECDR) .....	13-16
13-7	PCI Error Enable Register (PCI_EER) .....	13-17
13-8	PCI Error Attributes Capture Register (PCI_EATCR) .....	13-18
13-9	PCI Error Address Capture Register (PCI_EACR) .....	13-19
13-10	PCI Error Extended Address Capture Register (PCI_EEACR).....	13-20
13-11	PCI Error Data Low Capture Register (PCI_EDLCR) .....	13-20
13-12	PCI General Control Register (PCI_GCR) .....	13-21
13-13	PCI Error Control Register (PCI_ECR).....	13-21
13-14	PCI General Status Register (PCI_GSR) .....	13-22
13-15	PCI Inbound Translation Address Registers (PITAR <sub>n</sub> ) .....	13-23
13-16	PCI Inbound Base Address Registers (PIBAR <sub>n</sub> ).....	13-23
13-17	PCI Inbound Extended Base Address Registers (PIEBAR <sub>n</sub> ) .....	13-24
13-18	PCI Inbound Window Attribute Registers (PIWAR <sub>n</sub> ) .....	13-24
13-19	Vendor ID Configuration Register .....	13-27
13-20	Device ID Configuration Register .....	13-27
13-21	PCI Command Configuration Register .....	13-28
13-22	PCI Status Configuration Register .....	13-29
13-23	Revision ID Configuration Register .....	13-30
13-24	Standard Programming Interface Configuration Register.....	13-30
13-25	Subclass Code Configuration Register .....	13-31
13-26	Base Class Code Configuration Register .....	13-31
13-27	Cache Line Size Configuration Register.....	13-32
13-28	Latency Timer Configuration Register .....	13-32
13-29	Header Type Configuration Register .....	13-33
13-30	BIST Control Configuration Register .....	13-33
13-31	PIMMR Base Address Configuration Register.....	13-33
13-32	GPL Base Address Register 0.....	13-34
13-33	GPL Base Address Registers 1–2 .....	13-35
13-34	GPL Extended Base Address Registers 1–2 .....	13-35
13-35	Subsystem Vendor ID Configuration Register.....	13-36
13-36	Subsystem Device ID Configuration Register .....	13-36
13-37	Capabilities Pointer Configuration Register .....	13-37
13-38	Interrupt Line Configuration Register.....	13-37
13-39	Interrupt Pin Register .....	13-37
13-40	Minimum Grant Configuration Register.....	13-38
13-41	Maximum Latency Configuration Register .....	13-38
13-42	PCI Function Configuration Register .....	13-38
13-43	PCI Arbiter Control Register (PCIACR) .....	13-39
13-44	Hot Swap Register Block.....	13-40
13-45	PCI Power Management Register 0 (PCIPMR0).....	13-41

# Figures

Figure Number	Title	Page Number
13-46	PCI Power Management Register 1 (PCIPMR1).....	13-42
13-47	PCI Arbitration Example .....	13-45
13-48	Single Beat Read Example.....	13-49
13-49	Burst Read Example.....	13-50
13-50	Single Beat Write Example .....	13-50
13-51	Burst Write Example.....	13-51
13-52	Target-Initiated Terminations.....	13-52
13-53	PCI Parity Operation.....	13-58
13-54	Inbound PCI Memory Address Translation .....	13-59
14-1	SEC Connected to MPC8313E System Bus .....	14-2
14-2	SEC Functional Modules .....	14-3
14-3	Descriptor Format .....	14-11
14-4	Header Dword .....	14-12
14-5	Pointer Dword .....	14-15
14-6	Link Table Entry .....	14-17
14-7	DEU Mode Register (DEUMR).....	14-20
14-8	DEU Key Size Register (DEUKSR).....	14-21
14-9	DEU Data Size Register (DEUDSR).....	14-21
14-10	DEU Reset Control Register (DEURCR) .....	14-22
14-11	DEU Status Register (DEUSR).....	14-23
14-12	DEU Interrupt Status Register (DEUISR) .....	14-24
14-13	DEU Interrupt Control Register (DEUICR) .....	14-26
14-14	DEU End-of-Message Register (DEUEMR) .....	14-27
14-15	MDEU Mode Register (MDEUMR) in ‘Old’ Configuration .....	14-29
14-16	MDEU Mode Register (MDEUMR) in ‘New’ Configuration.....	14-30
14-17	MDEU Key Size Register (MDEUKSR).....	14-32
14-18	MDEU Data Size Register (MDEUDSR).....	14-33
14-19	MDEU Reset Control Register (MDEURCR).....	14-33
14-20	MDEU Status Register (MDEUSR).....	14-34
14-21	MDEU Interrupt Status Register (MDEUISR) .....	14-35
14-22	MDEU Interrupt Control Register (MDEUICR) .....	14-36
14-23	MDEU ICV Size Register.....	14-38
14-24	MDEU End-of-Message Register (MDEUEMR) .....	14-38
14-25	MDEU Context Registers .....	14-39
14-26	AESU Mode Register (AESUMR) .....	14-40
14-27	AESU Key Size Register (AESUKSR) .....	14-43
14-28	AESU Data Size Register (AESUDSR).....	14-43
14-29	AESU Reset Control Register (AESURCR).....	14-43
14-30	AESU Status Register (AESUSR) .....	14-44
14-31	AESU Interrupt Status Register (AESUISR).....	14-45
14-32	AESU Interrupt Control Register (AESUICR).....	14-47

# Figures

Figure Number	Title	Page Number
14-33	AESU End-of-Message Register (AESUEMR).....	14-49
14-34	AESU Context Registers.....	14-49
14-35	AESU CCM Context Registers.....	14-51
14-36	Crypto-Channel Configuration Register (CCCR).....	14-55
14-37	Crypto-Channel Pointer Status Register (CCPSR).....	14-57
14-38	Crypto-Channel Current Descriptor Pointer Register (CDPR).....	14-62
14-39	Fetch FIFO Register (FF).....	14-63
14-40	Descriptor Buffer (DB).....	14-63
14-41	EU Assignment Status Register (EUASR).....	14-68
14-42	Interrupt Mask Register (IMR).....	14-69
14-43	Interrupt Status Register (ISR).....	14-71
14-44	Interrupt Clear Register (ICR).....	14-72
14-45	ID Register (ID).....	14-73
14-46	IP Block Revision Register.....	14-73
14-47	Master Control Register (MCR).....	14-74
15-1	eTSEC Block Diagram.....	15-2
15-2	TSEC_ID Register.....	15-22
15-3	TSEC_ID2 Register.....	15-23
15-4	IEVENT Register Definition.....	15-25
15-5	IMASK Register Definition.....	15-28
15-6	EDIS Register Definition.....	15-29
15-7	ECNTRL Register Definition.....	15-31
15-8	PTV Register Definition.....	15-33
15-9	DMACTRL Register.....	15-34
15-10	TCTRL Register Definition.....	15-35
15-11	TSTAT Register Definition.....	15-37
15-12	DFVLAN Register Definition.....	15-41
15-13	TXIC Register Definition.....	15-42
15-14	TQUEUE Register Definition.....	15-43
15-15	TR03WT Register Definition.....	15-44
15-16	TR47WT Register Definition.....	15-44
15-17	TBDBPH Register Definition.....	15-45
15-18	TBPTR0–TBPTR7 Register Definition.....	15-46
15-19	TBASE Register Definition.....	15-46
15-20	TMR_TXTSn_ID Register Definition.....	15-47
15-21	TMR_TXTSn_H/L Register Definition.....	15-47
15-22	RCTRL Register Definition.....	15-48
15-23	RSTAT Register Definition.....	15-50
15-24	RXIC Register Definition.....	15-52
15-25	RQUEUE Register Definition.....	15-53
15-26	RBIFX Register Definition.....	15-54



# Figures

Figure Number	Title	Page Number
15-27	Receive Queue Filer Table Address Register Definition .....	15-56
15-28	Receive Queue Filer Table Control Register Definition .....	15-56
15-29	Receive Queue Filer Table Property IDs 0, 2–15 Register Definition.....	15-57
15-30	Receive Queue Filer Table Property ID1 Register Definition .....	15-58
15-31	MRBLR Register Definition .....	15-61
15-32	RBDBPH Register Definition .....	15-62
15-33	RBPTR0–RBPTR7 Register Definition .....	15-62
15-34	RBASE Register Definition .....	15-63
15-35	TMR_RXTS_H/L Register Definition.....	15-64
15-36	MACCFG1 Register Definition .....	15-67
15-37	MACCFG2 Register Definition .....	15-68
15-38	IPGIFG Register Definition .....	15-70
15-39	Half-Duplex Register Definition.....	15-71
15-40	Maximum Frame Length Register Definition.....	15-72
15-41	MII Management Configuration Register Definition .....	15-73
15-42	MIIMCOM Register Definition .....	15-73
15-43	MIIMADD Register Definition .....	15-74
15-44	MII Mgmt Control Register Definition.....	15-75
15-45	MIIMSTAT Register Definition .....	15-75
15-46	MII Mgmt Indicator Register Definition .....	15-76
15-47	Interface Status Register Definition .....	15-76
15-48	MAC Station Address Part 1 Register Definition .....	15-77
15-49	MAC Station Address Part 2 Register Definition .....	15-78
15-50	MAC Exact Match Address <i>n</i> Part 1 Register Definition .....	15-78
15-51	MAC Exact Match Address <i>x</i> Part 2 Register Definition .....	15-79
15-52	Transmit and Receive 64-Byte Frame Register Definition .....	15-80
15-53	Transmit and Receive 65- to 127-Byte Frame Register Definition .....	15-80
15-54	Transmit and Received 128- to 255-Byte Frame Register Definition .....	15-81
15-55	Transmit and Received 256- to 511-Byte Frame Register Definition.....	15-81
15-56	Transmit and Received 512- to 1023-Byte Frame Register Definition .....	15-82
15-57	Transmit and Received 1024- to 1518-Byte Frame Register Definition .....	15-82
15-58	Transmit and Received 1519- to 1522-Byte VLAN Frame Register Definition .....	15-83
15-59	Receive Byte Counter Register Definition.....	15-83
15-60	Receive Packet Counter Register Definition .....	15-83
15-61	Receive FCS Error Counter Register Definition.....	15-84
15-62	Receive Multicast Packet Counter Register Definition .....	15-84
15-63	Receive Broadcast Packet Counter Register Definition .....	15-85
15-64	Receive Control Frame Packet Counter Register Definition .....	15-85
15-65	Receive Pause Frame Packet Counter Register Definition .....	15-86
15-66	Receive Unknown OPCode Packet Counter Register Definition .....	15-86
15-67	Receive Alignment Error Counter Register Definition.....	15-87

# Figures

Figure Number	Title	Page Number
15-68	Receive Frame Length Error Counter Register Definition .....	15-87
15-69	Receive Code Error Counter Register Definition .....	15-88
15-70	Receive Carrier Sense Error Counter Register Definition .....	15-88
15-71	Receive Undersize Packet Counter Register Definition .....	15-89
15-72	Receive Oversize Packet Counter Register Definition .....	15-89
15-73	Receive Fragments Counter Register Definition .....	15-90
15-74	Receive Jabber Counter Register Definition.....	15-90
15-75	Receive Dropped Packet Counter Register Definition .....	15-91
15-76	Transmit Byte Counter Register Definition .....	15-91
15-77	Transmit Packet Counter Register Definition .....	15-92
15-78	Transmit Multicast Packet Counter Register Definition .....	15-92
15-79	Transmit Broadcast Packet Counter Register Definition .....	15-93
15-80	Transmit Pause Control Frame Counter Register Definition .....	15-93
15-81	Transmit Deferral Packet Counter Register Definition.....	15-94
15-82	Transmit Excessive Deferral Packet Counter Register Definition.....	15-94
15-83	Transmit Single Collision Packet Counter Register Definition .....	15-95
15-84	Transmit Multiple Collision Packet Counter Register Definition.....	15-95
15-85	Transmit Late Collision Packet Counter Register Definition .....	15-96
15-86	Transmit Excessive Collision Packet Counter Register Definition .....	15-96
15-87	Transmit Total Collision Counter Register Definition .....	15-97
15-88	Transmit Drop Frame Counter Register Definition .....	15-97
15-89	Transmit Jabber Frame Counter Register Definition .....	15-98
15-90	Transmit FCS Error Counter Register Definition .....	15-98
15-91	Transmit Control Frame Counter Register Definition .....	15-99
15-92	Transmit Oversized Frame Counter Register Definition .....	15-99
15-93	Transmit Undersize Frame Counter Register Definition .....	15-100
15-94	Transmit Fragment Counter Register Definition .....	15-100
15-95	Carry Register 1 (CAR1) Register Definition.....	15-101
15-96	Carry Register 2 (CAR2) Register Definition.....	15-102
15-97	Carry Mask Register 1 (CAM1) Register Definition.....	15-103
15-98	Carry Mask Register 2 (CAM2) Register Definition.....	15-105
15-99	Receive Filer Rejected Packet Counter Register Definition .....	15-106
15-100	IGADDR <sub>n</sub> Register Definition .....	15-107
15-101	GADDR <sub>n</sub> Register Definition.....	15-107
15-102	ATTR Register Definition.....	15-108
15-103	RQPRM Register Definition.....	15-109
15-104	RFBPTR0–RFBPTR7 Register Definition.....	15-110
15-105	TMR_CTRL Register Definition .....	15-111
15-106	TMR_TEVENT Register Definition.....	15-113
15-107	TMR_PEVENT Register Definition.....	15-115
15-108	TMR_PEMASK Register Definition .....	15-116

# Figures

Figure Number	Title	Page Number
15-109	TMR_CNT_H Register Definition .....	15-117
15-110	TMR_ADD Register Definition.....	15-118
15-111	TMR_ACC Register Definition .....	15-118
15-112	TMR_PRSC Register Definition .....	15-119
15-113	TMROFF_H/L Register Definition .....	15-119
15-114	TMR_ALARM1-2_H/L Register Definition .....	15-120
15-115	TMR_FIPER <sub>n</sub> Register Definition .....	15-121
15-116	TMR_ETTS1-2_H/L Register Definition.....	15-121
15-117	Control Register Definition.....	15-124
15-118	Status Register Definition .....	15-125
15-119	AN Advertisement Register Definition.....	15-126
15-120	AN Link Partner Base Page Ability Register Definition .....	15-128
15-121	AN Expansion Register Definition .....	15-129
15-122	AN Next Page Transmit Register Definition .....	15-130
15-123	AN Link Partner Ability Next Page Register Definition .....	15-130
15-124	Extended Status Register Definition .....	15-131
15-125	Jitter Diagnostics Register Definition .....	15-132
15-126	TBI Control Register Definition .....	15-133
15-127	eTSEC-MII Connection .....	15-135
15-128	eTSEC-RMII Connection .....	15-136
15-129	eTSEC-RGMII Connection.....	15-137
15-130	eTSEC-RTBI Connection .....	15-138
15-131	Definition of Custom Preamble Sequence .....	15-149
15-132	Definition of Received Preamble Sequence.....	15-149
15-133	Ethernet Address Recognition Flowchart .....	15-151
15-134	Sample C Code for Computing eTSEC Hash Table Indices.....	15-153
15-135	Location of Frame Control Blocks for TOE Parameters .....	15-161
15-136	Transmit Frame Control Block .....	15-161
15-137	Receive Frame Control Block.....	15-163
15-138	Structure of the Receive Queue Filer Table .....	15-168
15-139	1588 Timer Design Partition .....	15-179
15-140	Ethernet Sampling Points for 1588 .....	15-180
15-141	PTP Packet Format.....	15-181
15-142	Buffer Format for Transmit timestamp Insertion .....	15-183
15-143	Transmit Frame Control Block .....	15-183
15-144	Example of eTSEC Memory Structure for BDs .....	15-186
15-145	Buffer Descriptor Ring.....	15-186
15-146	Transmit Buffer Descriptor .....	15-187
15-147	Mapping of TxBDs to a C Data Structure.....	15-187
15-148	Receive Buffer Descriptor.....	15-190
15-149	Mapping of RxBDs to a C Data Structure .....	15-191

# Figures

Figure Number	Title	Page Number
16-1	USB Interface Block Diagram .....	16-2
16-2	Capability Registers Length (CAPLENGTH).....	16-8
16-3	Host Controller Interface Version (HCIVERSION) .....	16-8
16-4	Host Controller Structural Parameters (HCCPARAMS).....	16-9
16-5	Host Control Capability Parameters (HCCPARAMS) .....	16-10
16-6	Device Interface Version (DCIVERSION) .....	16-11
16-7	Device Control Capability Parameters (DCCPARAMS).....	16-11
16-8	USB Command Register (USBCMD) .....	16-12
16-9	USB Status Register (USBSTS).....	16-14
16-10	USB Interrupt Enable (USBINTR).....	16-17
16-11	USB Frame Index (FRINDEX).....	16-19
16-12	Periodic Frame List Base Address (PERIODICLISTBASE) .....	16-20
16-13	Device Address (DEVICEADDR).....	16-20
16-14	Current Asynchronous List Address (ASYNCLISTADDR) .....	16-21
16-15	Endpoint List Address (ENDPOINTLISTADDR).....	16-21
16-16	Master Interface Data Burst Size (BURSTSIZE) .....	16-22
16-17	Transmit FIFO Tuning Controls (TXFILLTUNING) .....	16-23
16-18	ULPI Register Access (ULPI VIEWPORT) .....	16-24
16-19	Configure Flag Register (CONFIGFLAG) .....	16-26
16-20	Port Status and Control (PORTSC).....	16-26
16-21	OTG Status Control (OTGSC).....	16-32
16-22	USB Mode (USBMODE) .....	16-34
16-23	Endpoint Setup Status (ENDPTSETUPSTAT) .....	16-35
16-24	Endpoint Initialization (ENDPTPRIME).....	16-35
16-25	Endpoint Flush (ENDPTFLUSH) .....	16-36
16-26	Endpoint Status (ENDPTSTATUS).....	16-36
16-27	Endpoint Complete (ENDPTCOMPLETE).....	16-37
16-28	Endpoint Control 0 (ENDPTCTRL0) .....	16-38
16-29	Endpoint Control 1 to 5 (ENDPTCTRL <sub>n</sub> ) .....	16-39
16-30	Snoop 1 and Snoop 2 (SNOOP <sub>n</sub> ).....	16-40
16-31	Age Count Threshold (AGE_CNT_THRESH).....	16-42
16-32	Priority Control (PRI_CTRL) .....	16-42
16-33	System Interface Control Register (SI_CTRL).....	16-43
16-34	USB General-Purpose Register (CONTROL) .....	16-44
16-35	Periodic Schedule Organization.....	16-48
16-36	Frame List Link Pointer Format.....	16-48
16-37	Asynchronous Schedule Organization .....	16-49
16-38	Isochronous Transaction Descriptor (iT <sub>D</sub> ) .....	16-50
16-39	Split-Transaction Isochronous Transaction Descriptor (siT <sub>D</sub> ) .....	16-53
16-40	Queue Element Transfer Descriptor (qT <sub>D</sub> ).....	16-57
16-41	Queue Head Layout .....	16-62

# Figures

Figure Number	Title	Page Number
16-42	Frame Span Traversal Node Structure .....	16-66
16-43	Derivation of Pointer into Frame List Array .....	16-72
16-44	General Format of Asynchronous Schedule List .....	16-72
16-45	Frame Boundary Relationship Between HS Bus and FS/LS Bus .....	16-73
16-46	Relationship of Periodic Schedule Frame Boundaries to Bus Frame Boundaries .....	16-74
16-47	Example Periodic Schedule .....	16-76
16-48	Example Association of iTDs to Client Request Buffer .....	16-79
16-49	Generic Queue Head Unlink Scenario .....	16-84
16-50	Asynchronous Schedule List with Annotation to Mark Head of List .....	16-85
16-51	Example Mapping of qTD Buffer Pointers to Buffer Pages .....	16-87
16-52	Host Controller Asynchronous Schedule Split-Transaction State Machine .....	16-90
16-53	Split Transaction, Interrupt Scheduling Boundary Conditions .....	16-93
16-54	General Structure of EHCI Periodic Schedule Utilizing Interrupt Spreading .....	16-94
16-55	Example Host Controller Traversal of Recovery Path via FSTNs .....	16-96
16-56	Split Transaction State Machine for Interrupt .....	16-99
16-57	Split Transaction, Isochronous Scheduling Boundary Conditions .....	16-105
16-58	siTD Scheduling Boundary Examples .....	16-107
16-59	Split Transaction State Machine for Isochronous .....	16-110
16-60	Endpoint Queue Head Organization .....	16-123
16-61	Endpoint Queue Head Layout .....	16-124
16-62	Endpoint Transfer Descriptor (dTD) .....	16-126
16-63	USB 2.0 Device States .....	16-130
16-64	Endpoint Queue Head Diagram .....	16-142
16-65	Software Link Pointers .....	16-144
16-66	ULPI Timing .....	16-154
16-67	Sending of RX CMD .....	16-155
16-68	ULPI Data Transmit (NOPID) .....	16-155
16-69	ULPI Data Transmit (PID) .....	16-155
16-70	ULPI Data Receive .....	16-156
16-71	ULPI Register Write .....	16-156
16-72	ULPI Register Read .....	16-156
17-1	I <sup>2</sup> C Block Diagram .....	17-1
17-2	I <sup>2</sup> C <sub>n</sub> Address Register (I2C <sub>n</sub> ADR) .....	17-5
17-3	I <sup>2</sup> C <sub>n</sub> Frequency Divider Register (I2C <sub>n</sub> FDR) .....	17-6
17-4	I <sup>2</sup> C <sub>n</sub> Control Register (I2C <sub>n</sub> CR) .....	17-7
17-5	I <sup>2</sup> C <sub>n</sub> Status Register (I2C <sub>n</sub> SR) .....	17-8
17-6	I <sup>2</sup> C <sub>n</sub> Data Register (I2C <sub>n</sub> DR) .....	17-9
17-7	I <sup>2</sup> C <sub>n</sub> Digital Filter Sampling Rate Register (I2C <sub>n</sub> DFSRR) .....	17-9
17-8	I <sup>2</sup> C Interface Transaction Protocol .....	17-10
17-9	EEPROM Contents .....	17-17
17-10	EEPROM Data Format for One Register Preload Command .....	17-18

# Figures

Figure Number	Title	Page Number
17-11	Example I <sup>2</sup> C Interrupt Service Routine Flowchart .....	17-20
18-1	UART Block Diagram .....	18-2
18-2	Receiver Buffer Registers (URBR1 and URBR2) .....	18-6
18-3	Transmitter Holding Registers (UTHR1 and UTHR2) .....	18-6
18-4	Divisor Most Significant Byte Registers (UDMB1 and UDMB2) .....	18-7
18-5	Divisor Least Significant Byte Registers (UDLB1 and UDLB2) .....	18-7
18-6	Interrupt Enable Registers (UIER1 and UIER2) .....	18-8
18-7	Interrupt ID Registers (UIIR1 and UIIR2) .....	18-9
18-8	FIFO Control Registers (UFCR1 and UFCR2) .....	18-11
18-9	Line Control Register (ULCR1 and ULCR2) .....	18-12
18-10	Modem Control Register (UMCR1 and UMCR2) .....	18-13
18-11	Line Status Register (ULSR1 and ULSR2) .....	18-14
18-12	Modem Status Register (UMSR1 and UMSR2) .....	18-15
18-13	Scratch Register (USCR) .....	18-16
18-14	Alternate Function Register (UAFR) .....	18-16
18-15	DMA Status Register (UDSR) .....	18-17
18-16	UART Bus Interface Transaction Protocol Example .....	18-19
19-1	SPI Block Diagram .....	19-2
19-2	Single-Master/Multi-Slave Configuration .....	19-4
19-3	Multiple-Master Configuration .....	19-6
19-4	SPMODE-SPI Mode Register Definition .....	19-9
19-5	SPI Transfer Format with SPMODE[CP] = 0 .....	19-11
19-6	SPI Transfer Format with SPMODE[CP] = 1 .....	19-11
19-7	SPIE—SPI Event Register Definition .....	19-12
19-8	SPIM—SPI Mask Register Definition .....	19-13
19-9	SPI Command Register Definition .....	19-14
19-10	SPI Transmit Data Hold Register Definition .....	19-14
19-11	SPI Receive Data Hold Register Definition .....	19-15
19-12	Example SPMODE[REV] = 0 SPMODE[LEN] = 7 LSB Sent First .....	19-15
19-13	Example SPMODE[REV] = 1 SPMODE[LEN] = 7 MSB Sent First .....	19-15
19-14	Example SPMODE[REV] = 1 SPMODE[LEN] = 15 MSB Sent First .....	19-16
19-15	Example SPMODE[REV] = 0 SPMODE[LEN] = 15 LSB Sent First .....	19-16
20-1	JTAG Interface Block Diagram .....	20-1
21-1	GPIO Module Block Diagram .....	21-1
21-2	GPIO Direction Register (GPDIR) .....	21-3
21-3	GPIO Open Drain Register (GPODR) .....	21-3
21-4	GPIO Data Register (GPDAT) .....	21-4
21-5	GPIO Interrupt Event Register (GPIER) .....	21-4
21-6	GPIO Interrupt Mask Register (GPIMR) .....	21-5
21-7	GPIO Interrupt Control Register (GPICR) .....	21-5

# Tables

Table Number	Title	Page Number
i	Acronyms and Abbreviated Terms.....	lxxvii
1-1	Supported eTSEC1 and eTSEC2 Configurations .....	1-12
2-1	Memory Map.....	2-2
3-1	MPC8313E Signal Reference by Functional Block.....	3-3
3-2	Reset Configuration Signals.....	3-12
3-3	Output Signal States During System Reset.....	3-13
3-4	Signals for Multiplexing .....	3-14
3-5	External Signals—Detailed Signal Descriptions .....	3-14
4-1	System Control Signals.....	4-1
4-2	External Clock Signals.....	4-3
4-3	Reset Causes .....	4-5
4-4	Reset Actions .....	4-6
4-5	Reset Configuration Words Source.....	4-10
4-6	SYS_CLK_IN Division .....	4-11
4-7	Selecting Reset Configuration Input Signals .....	4-11
4-8	RCWLR Bit Settings.....	4-13
4-9	System PLL Ratio .....	4-14
4-10	SPMF Maximum Values .....	4-14
4-11	Reset Configuration Word High Bit Settings.....	4-15
4-12	PCI Host/Agent Configuration.....	4-16
4-13	Boot Memory Space.....	4-17
4-14	Boot Sequencer Configuration.....	4-17
4-15	Boot ROM Location.....	4-18
4-16	eTSEC1 Mode Configuration .....	4-19
4-17	eTSEC2 Mode Configuration .....	4-20
4-18	e300 Core True Little-Endian .....	4-20
4-19	LALE Configuration .....	4-21
4-20	Local Bus Configuration EEPROM Addresses .....	4-21
4-21	Local Bus Reset Configuration Words Data Structure.....	4-21
4-22	Local Bus Controller Setting When Loading RCW.....	4-22
4-23	Hard Coded Reset Configuration Word Low Fields Values .....	4-26
4-24	Hard-Coded Reset Configuration Word High Field Values.....	4-27
4-25	Examples For Hard-Coded Reset Configuration Words Usage.....	4-27
4-26	Configurable Clock Units .....	4-31
4-27	Reset Configuration and Status Registers Memory Map.....	4-32
4-28	Reset Status Register Field Descriptions .....	4-33
4-29	RMR Field Descriptions .....	4-35
4-30	RPR Bit Descriptions.....	4-36

# Tables

Table Number	Title	Page Number
4-31	RCR Bit Settings .....	4-36
4-32	RCER Bit Settings .....	4-37
4-33	Clock Configuration Registers Memory Map .....	4-37
4-34	System PLL Mode Register Bit Settings .....	4-38
4-35	OCCR Bit Settings .....	4-39
4-36	SCCR Bit Descriptions .....	4-40
5-1	Local Access Windows Target Interface .....	5-1
5-2	Local Access Windows Example .....	5-2
5-3	Format of Window Definitions .....	5-3
5-4	Local Access Register Memory Map .....	5-4
5-5	IMMRBAR Bit Settings .....	5-7
5-6	ALTCBAR Bit Settings .....	5-7
5-7	LBLAWBAR0–LBLAWBAR3 Bit Settings .....	5-8
5-8	LBLAWBAR0[BASE_ADDR] Reset Value .....	5-8
5-9	LBLAWAR0–LBLAWAR3 Bit Settings .....	5-9
5-10	LBLAWAR0[EN] Reset Value .....	5-9
5-11	PCILAWBAR0–PCILAWBAR1 Bit Settings .....	5-10
5-12	PCILAWBAR0[BASE_ADDR] Reset Value .....	5-10
5-13	PCILAWAR0–PCILAWAR1 Bit Settings .....	5-11
5-14	PCILAWAR0[EN] Reset Value .....	5-11
5-15	DDRLAWBAR0–DDRLAWBAR1 Bit Settings .....	5-12
5-16	DDRLAWBAR0[BASE_ADDR] Reset Value .....	5-12
5-17	DDRLAWAR0–DDRLAWAR1 Bit Settings .....	5-13
5-18	DDRLAWAR0[EN] Reset Value .....	5-14
5-19	Overlapping Local Access Windows .....	5-14
5-20	System Configuration Register Memory Map .....	5-16
5-21	SGPRL Bit Settings .....	5-17
5-22	SGPRH Bit Settings .....	5-17
5-23	SPRIDR Bit Settings .....	5-18
5-24	PARTID Coding .....	5-18
5-25	REVID Coding .....	5-18
5-26	SPCR Bit Settings .....	5-19
5-27	SICRL Bit Settings .....	5-21
5-28	SICRH Bit Settings .....	5-24
5-29	SICRH[30–31] Bit Settings .....	5-26
5-30	DDRCDR Field Descriptions .....	5-27
5-31	DDRDSR Field Descriptions .....	5-28
5-32	WDT Register Address Map .....	5-30
5-33	SWCRR Bit Settings .....	5-31
5-34	SWCNR Bit Settings .....	5-32
5-35	SWSRR Bit Settings .....	5-33



# Tables

Table Number	Title	Page Number
5-36	RTC Signal Properties.....	5-36
5-37	RTC External Signal—Detailed Signal Description .....	5-37
5-38	RTC Register Address Map .....	5-37
5-39	RTCNR Bit Settings.....	5-38
5-40	RTLDR Bit Settings .....	5-38
5-41	RTPSR Bit Settings.....	5-39
5-42	RTCTR Bit Settings .....	5-39
5-43	RTEVR Bit Settings .....	5-40
5-44	RTALR Bit Settings .....	5-40
5-45	PIT Signal Properties .....	5-43
5-46	PIT External Signal—Detailed Signal Descriptions .....	5-44
5-47	PIT Register Address Map.....	5-44
5-48	PTCNR Bit Settings .....	5-45
5-49	PTLDR Bit Settings .....	5-45
5-50	PTPSR Bit Settings .....	5-46
5-51	PTCTR Bit Settings.....	5-46
5-52	PTEVR Bit Settings .....	5-47
5-53	GTM Signal Properties.....	5-51
5-54	GTM External Signals—Detailed Signal Descriptions.....	5-52
5-55	GTM Register Address Map .....	5-53
5-56	GTCFR1 Bit Settings .....	5-54
5-57	GTCFR2 Bit Settings .....	5-56
5-58	GTMDR Bit Settings.....	5-57
5-59	GTRFR Bit Settings .....	5-58
5-60	GTCPR <sub>n</sub> Bit Settings .....	5-59
5-61	GTCNR Bit Settings.....	5-59
5-62	GTEVR <sub>n</sub> Bit Settings.....	5-60
5-63	GTPSR <sub>n</sub> Bit Settings.....	5-60
5-64	System Control Signals—Detailed Signal Descriptions .....	5-65
5-65	Power Management Controller Registers Memory Map.....	5-65
5-66	PMCCR Bit Settings .....	5-66
5-67	PMCER Bit Settings .....	5-67
5-68	PMCMR Bit Settings .....	5-69
5-69	PMCCR1 Bit Settings .....	5-70
5-70	PMCCR2 Bit Settings .....	5-72
5-71	PCI Defined Power Management State Support.....	5-74
5-72	PCI Bus Power Management State Support.....	5-75
5-73	Software-Controller Power-Down States—Basic Description .....	5-76
5-74	MPC8313E Agent Mode Wake-Up Support.....	5-80
5-75	MPC8313E Host Mode Wake-Up Support .....	5-82
6-1	Arbiter Register Map .....	6-2

# Tables

Table Number	Title	Page Number
6-2	ACR Field Descriptions .....	6-3
6-3	ATR Field Descriptions.....	6-4
6-4	AER Field Descriptions .....	6-5
6-5	AIDR Field Descriptions .....	6-6
6-6	AMR Field Descriptions .....	6-7
6-7	AEATR Field Descriptions .....	6-8
6-8	AEADR Field Descriptions .....	6-9
6-9	AERR Field Descriptions.....	6-10
6-10	Address Only Transaction Type Encoding.....	6-14
6-11	Reserved Transaction Type Encoding.....	6-15
6-12	Illegal Transaction Type Encoding .....	6-15
7-1	MSR Bit Descriptions .....	7-17
7-2	e300 HID0 Bit Descriptions.....	7-20
7-3	Using HID0[ECLK] and HID0[SBCLK] to Configure <i>clk_out</i> .....	7-23
7-4	HID1 Bit Descriptions .....	7-23
7-5	e300HID2 Bit Descriptions.....	7-24
7-6	Interrupt Classifications .....	7-31
7-7	Exceptions and Interrupts.....	7-32
7-8	Differences Between e300 and G2_LE Cores .....	7-39
8-1	IPIC Signal Properties.....	8-5
8-2	IPIC External Signals—Detailed Signal Descriptions.....	8-5
8-3	IPIC Register Address Map .....	8-6
8-4	SICFR Field Descriptions .....	8-8
8-5	SIVCR Field Descriptions .....	8-9
8-6	IVEC/CVEC/MVEC Field Definition .....	8-9
8-7	SIPNR_H/SIFCR_H/SIMSR_H Bit Assignments.....	8-11
8-8	SIPNR_H Field Descriptions .....	8-12
8-9	SIPNR_L/SIFCR_L/SIMSR_L Bit Assignments .....	8-12
8-10	SIPNR_L Field Descriptions .....	8-13
8-11	SIPRR_A Field Descriptions .....	8-14
8-12	SIPRR_D Field Descriptions .....	8-15
8-13	SIMSR_H Field Descriptions .....	8-16
8-14	SIMSR_L Field Descriptions.....	8-16
8-15	SICNR Field Descriptions .....	8-17
8-16	SEPNR Field Descriptions.....	8-18
8-17	SMPRR_A Field Descriptions.....	8-19
8-18	SMPRR_B Field Descriptions .....	8-20
8-19	SEMSR Field Descriptions .....	8-21
8-20	SECNR Field Descriptions .....	8-22
8-21	SERSR/SERM/R/SERFR Bit Assignments.....	8-23
8-22	SERSR Field Descriptions .....	8-23

# Tables

Table Number	Title	Page Number
8-23	SERMCR Field Descriptions.....	8-24
8-24	SERCR Field Descriptions.....	8-24
8-25	SIFCR_H Field Descriptions.....	8-25
8-26	SIFCR_L Field Descriptions.....	8-25
8-27	SEFCR Field Descriptions.....	8-26
8-28	SERFR Field Descriptions.....	8-27
8-29	SCVCR Field Descriptions.....	8-27
8-30	SMVCR Field Descriptions.....	8-28
8-31	7Interrupt Source Priority Levels.....	8-31
9-1	DDR Memory Interface Signal Summary.....	9-3
9-2	Memory Address Signal Mappings.....	9-4
9-3	Memory Interface Signals—Detailed Signal Descriptions.....	9-5
9-4	Clock Signals—Detailed Signal Descriptions.....	9-7
9-5	DDR Memory Controller Memory Map.....	9-8
9-6	CS <sub>n</sub> _BNDS Field Descriptions.....	9-10
9-7	CS <sub>n</sub> _CONFIG Field Descriptions.....	9-10
9-8	TIMING_CFG_3 Field Descriptions.....	9-12
9-9	TIMING_CFG_0 Field Descriptions.....	9-13
9-10	TIMING_CFG_1 Field Descriptions.....	9-14
9-11	TIMING_CFG_2 Field Descriptions.....	9-17
9-12	DDR_SDRAM_CFG Field Descriptions.....	9-19
9-13	DDR_SDRAM_CFG_2 Field Descriptions.....	9-21
9-14	DDR_SDRAM_MODE Field Descriptions.....	9-23
9-15	DDR_SDRAM_MODE_2 Field Descriptions.....	9-23
9-16	DDR_SDRAM_MD_CNTL Field Descriptions.....	9-24
9-17	Settings of DDR_SDRAM_MD_CNTL Fields.....	9-25
9-18	DDR_SDRAM_INTERVAL Field Descriptions.....	9-26
9-19	DDR_DATA_INIT Field Descriptions.....	9-27
9-20	DDR_SDRAM_CLK_CNTL Field Descriptions.....	9-27
9-21	DDR_INIT_ADDR Field Descriptions.....	9-28
9-22	DDR_IP_REV1 Field Descriptions.....	9-28
9-23	DDR_IP_REV2 Field Descriptions.....	9-29
9-24	Byte Lane to Data Relationship.....	9-33
9-25	Supported DDR1 SDRAM Device Configurations.....	9-33
9-26	Supported DDR2 SDRAM Device Configurations.....	9-34
9-27	DDR1 Address Multiplexing for 32-Bit Data Bus with Interleaving Disabled.....	9-35
9-28	DDR2 Address Multiplexing.....	9-36
9-29	Example of Address Multiplexing for 32-Bit Data Bus Interleaving Between Two Banks with Partial Array Self Refresh Disabled.....	9-36
9-30	DDR SDRAM Command Table.....	9-38
9-31	DDR SDRAM Interface Timing Intervals.....	9-39

# Tables

Table Number	Title	Page Number
9-32	DDR SDRAM Power-Saving Modes Refresh Configuration.....	9-46
9-33	Memory Controller–Data Beat Ordering .....	9-48
9-34	Memory Interface Configuration Register Initialization Parameters.....	9-49
9-35	Programming Differences Between Memory Types.....	9-50
10-1	Signal Properties—Summary.....	10-4
10-2	Enhanced Local Bus Controller Detailed Signal Descriptions .....	10-5
10-3	Enhanced Local Bus Controller Registers .....	10-7
10-4	BR <sub>n</sub> Field Descriptions.....	10-10
10-5	Reset value of OR0 Register.....	10-11
10-6	Memory Bank Sizes in Relation to Address Mask .....	10-12
10-7	OR <sub>n</sub> —GPCM Field Descriptions .....	10-13
10-8	OR <sub>n</sub> —FCM Field Descriptions .....	10-16
10-9	OR <sub>n</sub> —UPM Field Descriptions .....	10-18
10-10	MAR Field Descriptions .....	10-19
10-11	MxMR Field Descriptions.....	10-20
10-12	MRTPR Field Descriptions.....	10-22
10-13	MDR Field Description.....	10-23
10-14	LSOR Field Description.....	10-24
10-15	LURT Field Descriptions .....	10-25
10-16	LTESR Field Descriptions .....	10-26
10-17	LTEDR Field Descriptions.....	10-27
10-18	LTEIR Field Descriptions .....	10-28
10-19	LTEATR Field Descriptions.....	10-30
10-20	LTEAR Field Descriptions.....	10-30
10-21	LTECCR Field Descriptions .....	10-31
10-22	LBCR Field Descriptions.....	10-32
10-23	LCRR Field Descriptions.....	10-33
10-24	FMR Field Descriptions.....	10-34
10-25	FIR Field Descriptions .....	10-36
10-26	FCR Field Descriptions.....	10-37
10-27	FBAR Field Descriptions.....	10-37
10-28	FPAR Field Descriptions, Small Page Device (OR <sub>x</sub> [PGS] = 0).....	10-38
10-29	FPAR Field Descriptions, Large Page Device (OR <sub>x</sub> [PGS] = 1).....	10-38
10-30	FBCR Field Descriptions .....	10-39
10-31	FECC <sub>n</sub> Field Descriptions .....	10-40
10-32	GPCM Read Control Signal Timing .....	10-47
10-33	GPCM Write Control Signal Timing .....	10-48
10-34	Boot Bank Field Values after Reset for GPCM as Boot Controller.....	10-56
10-35	FCM Chip-Select to First Command Timing.....	10-65
10-36	FCM Command, Address, and Write Data Timing Parameters.....	10-66
10-37	FCM Read Data Timing Parameters .....	10-69

# Tables

Table Number	Title	Page Number
10-38	Boot Bank Field Values after Reset for FCM as Boot Controller .....	10-70
10-39	UPM Routines Start Addresses .....	10-73
10-40	RAM Word Field Descriptions .....	10-79
10-41	MxMR Loop Field Use .....	10-84
10-42	UPM Address Multiplexing .....	10-85
10-43	Data Bus Drive Requirements For Read Cycles .....	10-93
10-44	FCM Register Settings for Soft Reset (ORn[PGS] = 1) .....	10-94
10-45	FCM Register Settings for Status Read (ORn[PGS] = 1) .....	10-94
10-46	FCM Register Settings for ID Read (ORn[PGS] = 1) .....	10-95
10-47	FCM Register Settings for Page Read (ORn[PGS] = 1) .....	10-95
10-48	FCM Register Settings for Block Erase (ORn[PGS] = 1) .....	10-96
10-49	FCM Register Settings for Page Program (ORn[PGS] = 1) .....	10-97
11-1	Sequencer Memory Map .....	11-2
11-2	POTAR <sub>n</sub> Field Descriptions .....	11-3
11-3	POBAR <sub>n</sub> Field Descriptions .....	11-4
11-4	POCMR <sub>n</sub> Field Descriptions .....	11-4
11-5	PMCR Field Descriptions .....	11-5
11-6	DTCR Field Descriptions .....	11-6
12-1	Module Memory Map .....	12-2
12-2	OMISR Field Descriptions .....	12-4
12-3	OMIMR Field Descriptions .....	12-4
12-4	IMR0 and IMR1 Field Descriptions .....	12-5
12-5	OMR0 and OMR1 Field Descriptions .....	12-5
12-6	ODR Field Descriptions .....	12-6
12-7	IDR Field Descriptions .....	12-7
12-8	IMISR Field Descriptions .....	12-8
12-9	IMIMR Field Descriptions .....	12-8
12-10	DMAMR <sub>n</sub> Field Descriptions .....	12-9
12-11	DMASR <sub>n</sub> Field Descriptions .....	12-11
12-12	DMACDAR <sub>n</sub> Field Descriptions .....	12-12
12-13	DMASAR <sub>n</sub> Field Descriptions .....	12-13
12-14	DMASAR <sub>n</sub> Field Descriptions .....	12-13
12-15	DMABCR <sub>n</sub> Field Descriptions .....	12-14
12-16	DMANDAR <sub>n</sub> Field Descriptions .....	12-14
12-17	DMA Segment Descriptor Fields .....	12-18
13-1	PCI Controller Modes .....	13-3
13-2	Signal Properties .....	13-4
13-3	PCI Interface Signals—Detailed Signal Descriptions .....	13-5
13-4	PCI Configuration Access Registers .....	13-11
13-5	PCI Memory-Mapped Registers .....	13-12
13-6	PCI_CONFIG_ADDRESS Field Descriptions .....	13-13

# Tables

Table Number	Title	Page Number
13-7	PCI_CONFIG_DATA Field Descriptions.....	13-15
13-8	PCI_ESR Field Descriptions.....	13-16
13-9	PCI_ECDR Field Descriptions .....	13-17
13-10	PCI_EER Field Descriptions .....	13-17
13-11	PCI_EATCR Field Descriptions .....	13-18
13-12	PCI_EACR Field Description.....	13-19
13-13	PCI_EEACR Field Description .....	13-20
13-14	PCI_EDLCR Field Description .....	13-20
13-15	PCI_GCR Field Descriptions.....	13-21
13-16	PCI_ECR Field Descriptions .....	13-22
13-17	PCI_GSR Field Descriptions .....	13-22
13-18	PITAR <sub>n</sub> Field Descriptions .....	13-23
13-19	PIBAR <sub>n</sub> Field Descriptions .....	13-23
13-20	PIEBAR <sub>n</sub> Field Descriptions .....	13-24
13-21	PIWAR <sub>n</sub> Field Descriptions.....	13-24
13-22	PCI Configuration Space Registers.....	13-26
13-23	Vendor ID Configuration Register Field Descriptions.....	13-27
13-24	Device ID Configuration Register Field Descriptions.....	13-27
13-25	PCI Command Configuration Register Field Descriptions.....	13-28
13-26	PCI Status Configuration Register Field Descriptions.....	13-29
13-27	Revision ID Configuration Register Field Descriptions .....	13-30
13-28	Standard Programming Interface Configuration Register Field Descriptions .....	13-30
13-29	Subclass Code Configuration Register Field Descriptions .....	13-31
13-30	Class Code Configuration Register Field Descriptions .....	13-31
13-31	Cache Line Size Configuration Register Field Descriptions .....	13-32
13-32	Latency Timer Configuration Register Field Descriptions .....	13-32
13-33	PIMMR Base Address Configuration Register Field Descriptions .....	13-33
13-34	GPL Base Address Register 0 Field Descriptions .....	13-34
13-35	GPL Base Address Register 1,2 Field Descriptions .....	13-35
13-36	GPL Extended Base Address Registers 1–2 Field Descriptions.....	13-35
13-37	Subsystem Vendor ID Configuration Register Field Descriptions .....	13-36
13-38	Subsystem Device ID Configuration Register Field Descriptions.....	13-36
13-39	Interrupt Line Configuration Register Field Descriptions .....	13-37
13-40	PCI Function Configuration Register Field Descriptions .....	13-38
13-41	PCI Arbiter Control Register (PCIACR) Field Descriptions.....	13-39
13-42	Hot Swap Register Block Field Descriptions .....	13-40
13-43	PCIPMR0 Field Descriptions.....	13-41
13-44	PCIPMR1 Field Descriptions.....	13-42
13-45	PCI Command Definitions.....	13-46
13-46	Special Cycle Commands .....	13-56
14-1	Example Descriptor.....	14-4

# Tables

Table Number	Title	Page Number
14-2	SEC Address Map.....	14-8
14-3	SEC Address Map.....	14-9
14-4	Header Dword Bit Definitions.....	14-12
14-5	Header Dword Writeback Bit Definitions.....	14-13
14-6	EU_SEL0 and EU_SEL1 Values.....	14-13
14-7	Descriptor Types.....	14-14
14-8	Pointer Dword Field Definitions.....	14-15
14-9	Link Table Field Definitions.....	14-17
14-10	Descriptor Format by Type.....	14-18
14-11	DEUMR Field Descriptions.....	14-20
14-12	DEUKSR Field Descriptions.....	14-21
14-13	DEURCR Field Descriptions.....	14-22
14-14	DEUSR Field Descriptions.....	14-23
14-15	DEUISR Field Descriptions.....	14-24
14-16	DEUICR Field Descriptions.....	14-26
14-17	MDEUMR in ‘Old’ Configuration.....	14-29
14-18	MDEUMR in ‘New’ Configuration.....	14-30
14-19	Mode Register—HMAC or SSL-MAC Generated by Single Descriptor.....	14-31
14-20	Mode Register—HMAC Generated Across a Sequence of Descriptors.....	14-32
14-21	MDEURCR Field Descriptions.....	14-33
14-22	MDEUSR Field Descriptions.....	14-34
14-23	MDEUISR Field Descriptions.....	14-35
14-24	MDEUICR Field Descriptions.....	14-37
14-25	AESUMR Field Descriptions.....	14-41
14-26	AES Cipher Modes.....	14-41
14-27	AESURCR Field Descriptions.....	14-44
14-28	AESUSR Field Descriptions.....	14-44
14-29	AESUISR Field Descriptions.....	14-46
14-30	AESUICR Field Descriptions.....	14-47
14-31	CCCR Field Descriptions.....	14-55
14-32	CCPSR Field Descriptions.....	14-57
14-33	G_STATE and S_STATE Field Values.....	14-59
14-34	CHN_STATE Field Values.....	14-60
14-35	Crypto-Channel Pointer Status Register Error Field Definitions.....	14-61
14-36	Crypto-Channel Pointer Status Register PAIR_PTR Field Values.....	14-61
14-37	CDPR Field Descriptions.....	14-62
14-38	Fetch FIFO Field Descriptions.....	14-63
14-39	Field Names in Interrupt Mask, Interrupt Status, and Interrupt Clear Registers.....	14-70
14-40	MCR Field Descriptions.....	14-74
15-1	eTSEC <sub>n</sub> Network Interface Signal Properties.....	15-6
15-2	eTSEC Signals—Detailed Signal Descriptions.....	15-8

# Tables

Table Number	Title	Page Number
15-3	Module Memory Map Summary.....	15-11
15-4	Module Memory Map .....	15-12
15-5	TSEC_ID Field Descriptions .....	15-23
15-6	TSEC_ID2 Field Descriptions .....	15-23
15-7	TSEC_ID2[TSEC_INT] Field Settings .....	15-23
15-8	IEVENT Field Descriptions.....	15-25
15-9	IMASK Field Descriptions .....	15-28
15-10	EDIS Field Descriptions .....	15-30
15-11	ECNTRL Field Descriptions.....	15-31
15-12	eTSEC Interface Configurations.....	15-32
15-13	PTV Field Descriptions.....	15-33
15-14	DMACTRL Field Descriptions.....	15-34
15-15	TCTRL Field Descriptions.....	15-36
15-16	TSTAT Field Descriptions.....	15-38
15-17	DFVLAN Field Descriptions .....	15-41
15-18	TXIC Field Descriptions .....	15-42
15-19	TQUEUE Field Descriptions .....	15-43
15-20	TR03WT Field Descriptions.....	15-44
15-21	TR47WT Field Descriptions.....	15-45
15-22	TBDBPH Field Descriptions .....	15-45
15-23	TBPTR <sub>n</sub> Field Descriptions .....	15-46
15-24	TBASE0–TBASE7 Field Descriptions .....	15-46
15-25	TMR_TXTS <sub>n</sub> _ID Register Field Descriptions .....	15-47
15-26	TMR_TXTS <sub>n</sub> _H/L Register Field Descriptions.....	15-47
15-27	RCTRL Field Descriptions .....	15-48
15-28	RSTAT Field Descriptions .....	15-51
15-29	RXIC Field Descriptions.....	15-52
15-30	RQUEUE Field Descriptions .....	15-53
15-31	RBIFX Field Descriptions .....	15-54
15-32	RQFAR Field Descriptions .....	15-56
15-33	RQFCR Field Descriptions .....	15-56
15-34	RQFPR Field Descriptions.....	15-58
15-35	MRBLR Field Descriptions .....	15-61
15-36	RBDBPH Field Descriptions .....	15-62
15-37	RBPTR <sub>n</sub> Field Descriptions.....	15-63
15-38	RBASE0–RBASE7 Field Descriptions .....	15-63
15-39	TMR_RXTS_H/L Register Field Descriptions.....	15-64
15-40	MACCFG1 Field Descriptions .....	15-67
15-41	MACCFG2 Field Descriptions .....	15-69
15-42	IPGIFG Field Descriptions .....	15-70
15-43	HAFDUP Field Descriptions .....	15-71



# Tables

<b>Table Number</b>	<b>Title</b>	<b>Page Number</b>
15-44	MAXFRM Descriptions.....	15-72
15-45	MIIMCFG Field Descriptions.....	15-73
15-46	MIIMCOM Descriptions.....	15-74
15-47	MIIMADD Field Descriptions.....	15-74
15-48	MIIMCON Field Descriptions.....	15-75
15-49	MIIMSTAT Field Descriptions.....	15-75
15-50	MIIMIND Field Descriptions.....	15-76
15-51	IFSTAT Field Descriptions.....	15-76
15-52	MACSTNADDR1 Field Descriptions.....	15-77
15-53	MACSTNADDR2 Field Descriptions.....	15-78
15-54	MACnADDR1 Field Descriptions.....	15-79
15-55	MAC01ADDR2–MAC15ADDR2 Field Descriptions.....	15-79
15-56	TR64 Field Descriptions.....	15-80
15-57	TR127 Field Descriptions.....	15-81
15-58	TR255 Field Descriptions.....	15-81
15-59	TR511 Field Descriptions.....	15-81
15-60	TR1K Field Descriptions.....	15-82
15-61	TRMAX Field Descriptions.....	15-82
15-62	TRMGV Field Descriptions.....	15-83
15-63	RBYT Field Descriptions.....	15-83
15-64	RPKT Field Descriptions.....	15-84
15-65	RFCS Field Descriptions.....	15-84
15-66	RMCA Field Descriptions.....	15-84
15-67	RBCA Field Descriptions.....	15-85
15-68	RXCF Field Descriptions.....	15-85
15-69	RXPF Field Descriptions.....	15-86
15-70	RXUO Field Descriptions.....	15-86
15-71	RALN Field Descriptions.....	15-87
15-72	RFLR Field Descriptions.....	15-87
15-73	RCDE Field Descriptions.....	15-88
15-74	RCSE Field Descriptions.....	15-88
15-75	RUND Field Descriptions.....	15-89
15-76	ROVR Field Descriptions.....	15-89
15-77	RFRG Field Descriptions.....	15-90
15-78	RJBR Field Descriptions.....	15-90
15-79	RDRP Field Descriptions.....	15-91
15-80	TBYT Field Descriptions.....	15-91
15-81	TPKT Field Descriptions.....	15-92
15-82	TMCA Field Descriptions.....	15-92
15-83	TBCA Field Descriptions.....	15-93
15-84	TXPF Field Descriptions.....	15-93

# Tables

Table Number	Title	Page Number
15-85	TDFR Field Descriptions .....	15-94
15-86	TEDF Field Descriptions .....	15-94
15-87	TSCL Field Descriptions .....	15-95
15-88	TMCL Field Descriptions .....	15-95
15-89	TLCL Field Descriptions .....	15-96
15-90	TXCL Field Descriptions .....	15-96
15-91	TNCL Field Descriptions .....	15-97
15-92	TDRP Field Descriptions .....	15-97
15-93	TJBR Field Descriptions .....	15-98
15-94	TFCS Field Descriptions .....	15-98
15-95	TXCF Field Descriptions .....	15-99
15-96	TOVR Field Descriptions .....	15-99
15-97	TUND Field Descriptions .....	15-100
15-98	TFRG Field Descriptions .....	15-100
15-99	CAR1 Field Descriptions .....	15-101
15-100	CAR2 Field Descriptions .....	15-102
15-101	CAM1 Field Descriptions .....	15-103
15-102	CAM2 Field Descriptions .....	15-105
15-103	RREJ Field Descriptions .....	15-106
15-104	IGADDR <sub>n</sub> Field Descriptions .....	15-107
15-105	GADDR <sub>n</sub> Field Descriptions .....	15-108
15-106	ATTR Field Descriptions .....	15-108
15-107	RQPRM Field Descriptions .....	15-109
15-108	RFBPTR0–RFBPTR7 Field Descriptions .....	15-110
15-109	TMR_CTRL Register Field Descriptions .....	15-111
15-110	TMR_TEVENT Register Field Descriptions .....	15-113
15-111	TMR_TEMASK Register Definition .....	15-114
15-112	TMR_TEMASK Register Field Descriptions .....	15-114
15-113	TMR_PEVENT Register Field Descriptions .....	15-115
15-114	TMR_PEMASK Register Field Descriptions .....	15-116
15-115	TMR_STAT Register Definition .....	15-116
15-116	TMR_STAT Register Field Descriptions .....	15-117
15-117	TMR_CNT_H/L Register Field Descriptions .....	15-117
15-118	TMR_ADD Register Field Descriptions .....	15-118
15-119	TMR_ACC Register Field Descriptions .....	15-118
15-120	TMR_PRSC Register Field Descriptions .....	15-119
15-121	TMROFF_H/L Register Field Descriptions .....	15-119
15-122	TMR_ALARM <sub>n</sub> _H/L Register Field Descriptions .....	15-120
15-123	TMR_FIPER Register Field Descriptions .....	15-121
15-124	TMR_ETTS1-2_H Register Field Descriptions .....	15-121
15-125	TBI MII Register Set .....	15-123

# Tables

Table Number	Title	Page Number
15-126	CR Field Descriptions .....	15-124
15-127	SR Descriptions.....	15-125
15-128	ANA Field Descriptions.....	15-126
15-129	PAUSE Priority Resolution.....	15-127
15-130	ANLPBPA Field Descriptions .....	15-128
15-131	ANEX Field Descriptions .....	15-129
15-132	ANNPT Field Descriptions .....	15-130
15-133	ANLPANP Field Descriptions .....	15-131
15-134	EXST Field Descriptions .....	15-132
15-135	JD Field Descriptions.....	15-133
15-136	TBICON Field Descriptions .....	15-134
15-137	GMII, MII, and RMII Signals Multiplexing .....	15-139
15-138	MII and RMII Signals Multiplexing .....	15-139
15-139	RGMII, TBI, and RTBI Signals Multiplexing .....	15-140
15-140	RGMII and RTBI Signals Multiplexing.....	15-141
15-141	RGMII Signals Multiplexing .....	15-142
15-142	Shared Signals.....	15-143
15-143	Steps for Minimum Register Initialization.....	15-144
15-144	Custom Preamble Field Descriptions.....	15-149
15-145	Received Preamble Field Descriptions .....	15-150
15-146	Flow Control Frame Structure .....	15-154
15-147	Non-Error Transmit Interrupts .....	15-156
15-148	Non-Error Receive Interrupts.....	15-156
15-149	Interrupt Coalescing Timing Threshold Ranges .....	15-157
15-150	Transmission Errors .....	15-158
15-151	Reception Errors .....	15-159
15-152	Tx Frame Control Block Description.....	15-162
15-153	Rx Frame Control Block Descriptions.....	15-164
15-154	Supported Stack L2 Ethernet Headers .....	15-166
15-155	Special Filer Rules .....	15-170
15-156	Receive Queue Filer Interrupt Events .....	15-170
15-157	Filer Table Example—802.1p Priority Filing .....	15-171
15-158	Filer Table Example—IP Diff-Serv Code Points Filing .....	15-172
15-159	Filer Table Example—TCP and UDP Port Filing.....	15-173
15-160	PTP Payload Special Fields.....	15-180
15-161	Timestamp Insertion Programming Requirements .....	15-182
15-162	Tx Frame Control Block Description.....	15-184
15-163	Transmit Data Buffer Descriptor (TxBD) Field Descriptions .....	15-188
15-164	Receive Buffer Descriptor Field Descriptions .....	15-191
15-165	MII Interface Mode Signal Configuration .....	15-193
15-166	Shared MII Signals.....	15-194

# Tables

<b>Table Number</b>	<b>Title</b>	<b>Page Number</b>
15-167	MII Mode Register Initialization Steps.....	15-194
15-168	RGMII Interface Mode Signal Configuration.....	15-196
15-169	Shared RGMII Signals.....	15-197
15-170	RGMII Mode Register Initialization Steps.....	15-197
15-171	RMII Interface Mode Signal Configuration.....	15-200
15-172	Shared RMII Signals.....	15-200
15-173	RMII Mode Register Initialization Steps.....	15-201
15-174	RTBI Interface Mode Signal Configuration.....	15-204
15-175	Shared RTBI Signals.....	15-204
15-176	RTBI Mode Register Initialization Steps.....	15-205
15-177	SGMII Interface Signal Configuration (4-Wire).....	15-207
15-178	SGMII Mode Register Initialization Steps.....	15-208
16-1	USB External Signals.....	16-3
16-2	ULPI Signal Descriptions.....	16-4
16-3	USB Interface Memory Map.....	16-6
16-4	CAPLENGTH Register Field Descriptions.....	16-8
16-5	HCIVERSION Register Field Descriptions.....	16-8
16-6	HCSPARAMS Register Field Descriptions.....	16-9
16-7	HCCPARAMS Register Field Descriptions.....	16-10
16-8	DCIVERSION Register Field Descriptions.....	16-11
16-9	DCCPARAMS Register Field Descriptions.....	16-11
16-10	USBCMD Register Field Descriptions.....	16-12
16-11	USBSTS Register Field Descriptions.....	16-15
16-12	USBINTR Register Field Descriptions.....	16-17
16-13	FRINDEX Register Field Descriptions.....	16-19
16-14	FRINDEX N Values.....	16-19
16-15	PERIODICLISTBASE Register Field Descriptions.....	16-20
16-16	DEVICEADDR Register Field Descriptions.....	16-20
16-17	ASYNCLISTADDR Register Field Descriptions.....	16-21
16-18	ENDPOINTLISTADDR Register Field Descriptions.....	16-22
16-19	BURSTSIZE Register Field Descriptions.....	16-22
16-20	TXFILLTUNING Register Field Descriptions.....	16-23
16-21	ULPI VIEWPORT Field Descriptions.....	16-25
16-22	CONFIGFLAG Register Field Descriptions.....	16-26
16-23	PORTSC Register Field Descriptions.....	16-27
16-24	OTGSC Register Field Descriptions.....	16-32
16-25	USBMODE Register Field Descriptions.....	16-34
16-26	ENDPTSETUPSTAT Register Field Descriptions.....	16-35
16-27	ENDPTPRIME Register Field Descriptions.....	16-35
16-28	ENDPTFLUSH Register Field Descriptions.....	16-36
16-29	ENDPTSTATUS Register Field Descriptions.....	16-37

# Tables

Table Number	Title	Page Number
16-30	ENDPTCOMPLETE Register Field Descriptions .....	16-37
16-31	ENDPTCTRL0 Register Field Descriptions .....	16-38
16-32	ENDPTCTRL $n$ Register Field Descriptions .....	16-39
16-33	SNOOP $n$ Register Field Descriptions .....	16-41
16-34	AGE_CNT_THRESH Register Field Descriptions .....	16-42
16-35	PRI_CTRL Register Field Descriptions .....	16-43
16-36	SI_CTRL Register Field Descriptions .....	16-43
16-37	CONTROL Field Descriptions .....	16-44
16-38	Supported PHY Interfaces .....	16-46
16-39	Typ Field Encodings .....	16-49
16-40	Next Schedule Element Pointer .....	16-50
16-41	iTD Transaction Status and Control .....	16-51
16-42	Buffer Pointer Page 0 (Plus) .....	16-52
16-43	iTD Buffer Pointer Page 1 (Plus) .....	16-52
16-44	Buffer Pointer Page 2 (Plus) .....	16-52
16-45	Buffer Pointer Page 3–6 .....	16-52
16-46	Next Link Pointer .....	16-53
16-47	Endpoint and Transaction Translator Characteristics .....	16-54
16-48	Micro-Frame Schedule Control .....	16-54
16-49	siTD Transfer Status and Control .....	16-54
16-50	siTD Buffer Pointer Page 0 (Plus) .....	16-55
16-51	siTD Buffer Pointer Page 1 (Plus) .....	16-56
16-52	siTD Back Link Pointer .....	16-56
16-53	qTD Next Element Transfer Pointer (DWord 0) .....	16-57
16-54	qTD Alternate Next Element Transfer Pointer (DWord 1) .....	16-57
16-55	qTD Token (DWord 2) .....	16-58
16-56	qTD Buffer Pointer .....	16-61
16-57	Queue Head DWord 0 .....	16-62
16-58	Endpoint Characteristics: Queue Head DWord 1 .....	16-63
16-59	Endpoint Capabilities: Queue Head DWord 2 .....	16-64
16-60	Current qTD Link Pointer .....	16-65
16-61	Host-Controller Rules for Bits in Overlay (DWords 5, 6, 8, and 9) .....	16-66
16-62	FSTN Normal Path Pointer .....	16-67
16-63	FSTN Back Path Link Pointer .....	16-67
16-64	Behavior During Wake-Up Events .....	16-71
16-65	Operation of FRINDEX and SOFV (SOF Value Register) .....	16-75
16-66	Example Periodic Reference Patterns for Interrupt Transfers .....	16-88
16-67	Ping Control State Transition Table .....	16-89
16-68	Interrupt IN/OUT Do Complete Split State Execution Criteria .....	16-103
16-69	Initial Conditions for OUT siTD TP and T-Count Fields .....	16-111
16-70	Transaction Position (TP)/Transaction Count (T-Count) Transition Table .....	16-111

# Tables

Table Number	Title	Page Number
16-71	Summary siTD Split Transaction State .....	16-115
16-72	Example Case 2a—Software Scheduling siTDs for an IN Endpoint.....	16-116
16-73	Summary of Transaction Errors .....	16-119
16-74	Summary Behavior on Host System Errors .....	16-122
16-75	Endpoint Capabilities/Characteristics .....	16-124
16-76	Current dTD Pointer.....	16-125
16-77	Multiple Mode Control .....	16-126
16-78	Next dTD Pointer .....	16-126
16-79	dTD Token .....	16-127
16-80	Buffer Pointer Page 0 .....	16-127
16-81	Buffer Pointer Page 1 .....	16-128
16-82	Buffer Pointer Pages 2–4 .....	16-128
16-83	Device Controller State Information Bits .....	16-130
16-84	Device Controller Endpoint Initialization.....	16-133
16-85	Device Controller Stall Response Matrix .....	16-134
16-86	Variable Length Transfer Protocol Example (ZLT = 0).....	16-136
16-87	Variable Length Transfer Protocol Example (ZLT = 1).....	16-136
16-88	Interrupt/Bulk Endpoint Bus Response Matrix.....	16-137
16-89	Control Endpoint Bus Response Matrix .....	16-139
16-90	Isochronous Endpoint Bus Response Matrix .....	16-142
16-91	Device Error Matrix .....	16-146
16-92	Error Descriptions .....	16-147
16-93	Interrupt Handling Order .....	16-147
16-94	Low Frequency Interrupt Events.....	16-148
16-95	Error Interrupt Events .....	16-148
16-96	Functional Differences Between EHCI and EHCI with Embedded TT .....	16-149
16-97	Emulated Handshakes .....	16-151
16-98	ULPI Timing .....	16-154
17-1	I <sup>2</sup> C Interface Signal Descriptions .....	17-3
17-2	I <sup>2</sup> C Interface Signals—Detailed Signal Descriptions .....	17-4
17-3	I <sup>2</sup> C Memory Map.....	17-4
17-4	I2C <sub>n</sub> ADR Field Descriptions.....	17-5
17-5	I2C <sub>n</sub> FDR Field Descriptions .....	17-6
17-6	I2C <sub>n</sub> CR Field Descriptions.....	17-7
17-7	I2C <sub>n</sub> SR Field Descriptions .....	17-8
17-8	I2C <sub>n</sub> DR Field Descriptions.....	17-9
17-9	I2C <sub>n</sub> DFSRR Field Descriptions.....	17-10
18-1	DUART Signal Overview .....	18-3
18-2	DUART Signals—Detailed Signal Descriptions .....	18-3
18-3	DUART Register Summary .....	18-4
18-4	URBR Field Descriptions .....	18-6

# Tables

Table Number	Title	Page Number
18-5	UTHR Field Descriptions .....	18-6
18-6	UDMB Field Descriptions .....	18-7
18-7	UDLB Field Descriptions .....	18-7
18-8	Baud Rate Examples .....	18-7
18-9	UIER Field Descriptions .....	18-9
18-10	UIIR Field Descriptions .....	18-10
18-11	UIIR IID Bits Summary .....	18-10
18-12	UFCR Field Descriptions .....	18-11
18-13	ULCR Field Descriptions .....	18-12
18-14	Parity Selection Using ULCR[PEN], ULCR[SP], and ULCR[EPS] .....	18-13
18-15	UMCR Field Descriptions .....	18-13
18-16	ULSR Field Descriptions .....	18-14
18-17	UMSR Field Descriptions .....	18-15
18-18	USCR Field Descriptions .....	18-16
18-19	UAFR Field Descriptions .....	18-16
18-20	UDSR Field Descriptions .....	18-17
18-21	UDSR[TXRDY] Set Conditions .....	18-17
18-22	UDSR[TXRDY] Cleared Conditions .....	18-17
18-23	UDSR[RXRDY] Set Conditions .....	18-18
18-24	UDSR[RXRDY] Cleared .....	18-18
19-1	Signal Properties .....	19-7
19-2	Detailed Signal Descriptions .....	19-7
19-3	SPI Register Summary .....	19-9
19-4	SPMODE Field Descriptions .....	19-9
19-5	SPIE Field Descriptions .....	19-12
19-6	SPIM Field Descriptions .....	19-13
19-7	SPCOM Field Descriptions .....	19-14
19-8	SPI Transmit Data Hold Field Descriptions .....	19-14
19-9	SPI Receive Data Hold Field Descriptions .....	19-15
20-1	JTAG Test Signals Summary .....	20-2
20-2	JTAG Test—Detailed Signal Descriptions .....	20-2
21-1	IPIC External Signals—Detailed Signal Descriptions .....	21-2
21-2	GPIO Register Address Map .....	21-2
21-3	GPDIR Bit Settings .....	21-3
21-4	GPODR Bit Settings .....	21-3
21-5	GPnDAT Bit Settings .....	21-4
21-6	GPIER Bit Settings .....	21-4
21-7	GPIMR Bit Settings .....	21-5
21-8	GPICR Bit Settings .....	21-5

# Tables

**Table  
Number**

**Title**

**Page  
Number**



## About This Book

This reference manual defines the functionality of the MPC8313E. It is written from the perspective of the MPC8313E, which is the superset device, and unless otherwise noted, the information applies also to the MPC8313. Note that the MPC8313 does not support a security engine.

The MPC8313E is a low-power, highly integrated host processor that addresses the requirements of several printing and imaging, consumer, and industrial applications, including main CPUs and I/O processors in printing systems, networking switches and line cards, wireless LANs (WLANs), network access servers (NAS), VPN routers, intelligent NIC, and industrial controllers. The MPC8313E extends the PowerQUICC™ family, adding higher CPU performance, additional functionality, and faster interfaces while addressing the requirements related to time-to-market, price, power consumption, and package size. The MPC8313E contains an embedded PowerPC™ e300c3 core built on Power Architecture™ technology.

## Audience

It is assumed that the reader understands operating systems, microprocessor system design, and the basic principles of RISC processing.

## Organization

Following is a summary and a brief description of the major parts of this reference manual:

- [Chapter 1, “Overview,”](#) provides a high-level description of features and functionality of the MPC8313E integrated host processor. It describes the MPC8313E, its interfaces, and the programming model. The functional operation of the MPC8313E, with emphasis on peripheral functions, is also described.
- [Chapter 2, “Memory Map,”](#) describes the memory map of the MPC8313E. An overview of the local address map is provided. Next, a complete listing of all memory-mapped registers is provided, with cross references to the sections detailing descriptions of each.
- [Chapter 3, “Signal Descriptions,”](#) provides a listing of all the external signals, cross-references for signals that serve multiple functions, their functional blocks, and I/O states. Also, these signals are listed by alphabetical order.
- [Chapter 4, “Reset, Clocking, and Initialization,”](#) describes the hard and soft resets, the power-on reset (POR) sequence, power-on reset configuration, clocking, and initialization of the MPC8313E.
- [Chapter 5, “System Configuration,”](#) provides an overview of several functions that control the local access windows, system configuration, software watchdog, real time clock, periodic and general purpose timers, power management, protection, and general utilities.

- [Chapter 6, “Arbiter and Bus Monitor,”](#) provides an overview of the arbiter in the MPC8313E device. Also, it describes the configuration, control, and status registers of the arbiter.
- [Chapter 7, “e300 Processor Core Overview,”](#) provides an overview of the basic functionality of the processor core and briefly describes how the functional units interact.
- [Chapter 8, “Integrated Programmable Interrupt Controller \(IPIC\),”](#) describes the IPIC interrupt protocol, various types of interrupt sources controlled by the IPIC unit, and the IPIC registers with some programming guidelines. It also provides a definition of the external interrupt signals and their functions. In addition, the interrupt configuration, control, and status registers are described in this chapter.
- [Chapter 9, “DDR Memory Controller,”](#) describes the 32-bit DDR SDRAM memory controller of the MPC8313E. This fully programmable controller supports most DDR memories available today, including both buffered and unbuffered devices. Dynamic power management and auto-precharge modes simplify memory system design.
- [Chapter 10, “Enhanced Local Bus Controller,”](#) describes the enhanced local bus controller (eLBC) of the MPC8313E. It describes the external signals and the memory-mapped registers as well as a functional description of the general-purpose chip-select machine (GPCM), Flash control machine (FCM), and user-programmable machines (UPMs) of the eLBC. Also, it includes an initialization and applications information section with many specific examples of its use.
- [Chapter 11, “Sequencer,”](#) describes how the I/O sequencer (IOS) switches transactions among its ports, using a buffer pool to minimize blocking. It also provides address translation on outbound PCI transactions.
- [Chapter 12, “DMA/Messaging Unit,”](#) describes the four-channel high speed general-purpose DMA controller of the MPC8313E. The channels share buffer space in the IOS to facilitate the gathering and sending of data. The DMA/messaging unit supports communication between two processors on different buses, for example, a local processor and a processor on a PCI bus. This communication unit operates with generic messages and doorbell registers. This block also provides a DMA controller that transfers blocks of data independent of the local processor or PCI hosts.
- [Chapter 13, “PCI Bus Interface,”](#) describes the PCI interface, which complies with the *PCI Local Bus Specification*, Rev. 2.3. This chapter provides a basic description of PCI bus operations. The specific emphasis is directed at how this device implements the PCI specification.
- [Chapter 14, “Security Engine \(SEC\) 2.2,”](#) describes the SEC 2.2, which is designed to offload computationally intensive security functions, such as key generation and exchange, authentication, and bulk encryption from the e300 core of the MPC8313E. It is optimized to process all the algorithms associated with IPsec. The SEC 2.2 (implemented in the MPC8313E) is derived from integrated security cores found in other members of the PowerQUICC family, including SEC 1.0, the version implemented in the MPC8272/MPC8248. Note that the MPC8313 does not support a security engine.
- [Chapter 15, “Enhanced Three-Speed Ethernet Controllers,”](#) describes the two enhanced three-speed Ethernet controllers on the MPC8313E. These controllers provide 10/100/1Gb Ethernet support with a set of media-independent interface options including MII, RMII, GMII, SGMII and RTBI. They are backward compatible with PowerQUICC III TSEC controllers.

- [Chapter 16, “Universal Serial Bus Interface,”](#) describes the universal serial bus (USB) interface. The USB DR module is a USB 2.0-compliant serial interface engine for implementing a USB interface. The registers and data structures are based on the *Enhanced Host Controller Interface Specification for Universal Serial Bus* (EHCI) from Intel Corporation. The USB DR module can act as a host, as a device, or as an on-the-go (OTG) negotiable host/device on the USB bus. The USB DR module allows either an internal PHY to be connected to the UTMI interface or an external PHY to be connected to the ULPI interface.
- [Chapter 17, “I<sup>2</sup>C Interface,”](#) describes the inter-IC (IIC or I<sup>2</sup>C) bus controllers of the MPC8313E. These synchronous, serial, bidirectional, multiple-master buses allow two-wire connection of devices, such as microcontrollers, EEPROMs, real-time clock devices, A/D converters, and LCDs. The MPC8313E powers up in boot sequencer mode, which allows the I<sup>2</sup>C controllers to initialize configuration registers.
- [Chapter 18, “DUART,”](#) describes the (dual) universal asynchronous receiver/transmitters (UARTs) which feature a PC16552D-compatible programming model. These independent UARTs are provided specifically to support system debugging.
- [Chapter 19, “Serial Peripheral Interface,”](#) describes the MPC8313E serial peripheral interface (SPI) that allows the exchange of data between MPC83xx family devices. The SPI can also be used to communicate with peripheral devices such as EEPROMs, real-time clocks, A/D converters, and ISDN devices.
- [Chapter 20, “JTAG/Testing Support,”](#) describes the joint test action group (JTAG) interface of the MPC8349E to facilitate boundary-scan testing. The JTAG interface complies to the IEEE 1149.1™ boundary-scan specification.
- [Chapter 21, “General Purpose I/O \(GPIO\),”](#) describes the general purpose I/O (GPIO) module in the MPC8313E device, including a definition of the external signals and functions they serve. Additionally, interrupt capabilities, pin description, and register settings are described.
- This reference manual also includes a revision history, glossary, and an index.

## Suggested Reading

This section lists additional reading that provides background for the information in this manual as well as general information about the architecture.

## General Information

The following documentation, published by Morgan-Kaufmann Publishers, 340 Pine Street, Sixth Floor, San Francisco, CA, provides useful information about the PowerPC architecture and computer architecture in general:

- *The PowerPC Architecture: A Specification for a New Family of RISC Processors*, Second Edition, by International Business Machines, Inc.
- *Computer Architecture: A Quantitative Approach*, Third Edition, by John L. Hennessy and David A. Patterson.
- *Computer Organization and Design: The Hardware/Software Interface*, Second Edition, by David A. Patterson and John L. Hennessy.

## Related Documentation

Freescale documentation is available from the sources listed on the back cover of this manual; the document order numbers are included in parentheses for ease in ordering:

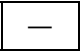
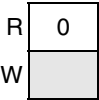
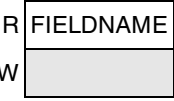
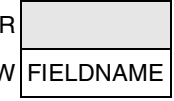
- *e300 Power Architecture™ Core Family Reference Manual*—This book provides a more detailed description of the e300 core.
- Reference manuals (formerly called user’s manuals)—These books provide details about individual implementations.
- Addenda/errata to reference or user’s manuals—Because some processors have follow-on devices, an addendum may be provided that describes the additional features and functionality changes. These addenda are intended for use with the corresponding reference or user’s manuals.
- Hardware specifications—Hardware specifications provide specific data regarding bus timing, signal behavior, and AC, DC, and thermal characteristics, as well as other design considerations.
- Product briefs—Each device has a product brief that provides an overview of its features.
- Application notes—These short documents address specific design issues useful to programmers and engineers working with Freescale processors.

Additional literature is published as new processors become available. For a current list of documentation, refer to <http://www.freescale.com>.

## Conventions

This document uses the following notational conventions:

cleared/set	When a bit takes the value zero, it is said to be cleared; when it takes a value of one, it is said to be set.
<b>mnemonics</b>	Instruction mnemonics are shown in lowercase bold
<i>italics</i>	Italics indicate variable command parameters, for example, <b>bcctrx</b>
	Book titles in text are set in italics
	Internal signals are set in lowercase italics, for example, <u>core_int</u>
0x0	Prefix to denote hexadecimal number
0b0	Prefix to denote binary number
rA, rB	Instruction syntax used to identify a source GPR
rD	Instruction syntax used to identify a destination GPR
REG[FIELD]	Abbreviations for registers are shown in uppercase text. Specific bits, fields, or ranges appear in brackets. For example, MSR[LE] refers to the little-endian mode enable bit in the machine state register.
x	In some contexts, such as signal encodings, an unitalicized x indicates a don’t care
<i>x</i>	An italicized <i>x</i> indicates an alphanumeric variable
<i>n</i>	An italicized <i>n</i> indicates a numeric variable
¬	NOT logical operator

&	AND logical operator
	OR logical operator
	Concatenation, for example TCR[WP]  TCR[WPEXT]
	Indicates a reserved bit field in a register. Although these bits can be written to as ones or zeros, they are always read as zeros.
	Indicates a reserved bit field in a memory-mapped register. Although these bits can be written to as ones or zeros, they are always read as zeros.
	Indicates a read-only bit field in a memory-mapped register.
	Indicates a write-only bit field in a memory-mapped register. Although these bits can be written to as ones or zeros, they are always read as zeros.

## Signal Conventions

<u>OVERBAR</u>	An overbar indicates that a signal is active-low.
<i>lowercase_italics</i>	Lowercase italics is used to indicate internal signals.
lowercase_plaintext	Lowercase plain text is used to indicate signals that are used for configuration. For more information, see <a href="#">Section 3.2, “Configuration Signals Sampled at Reset.”</a>

## Acronyms and Abbreviations

[Table i](#) contains acronyms and abbreviations used in this document.

**Table i. Acronyms and Abbreviated Terms**

Term	Meaning
AESU	Advanced encryption standard unit
AFEU	ARC four execution unit
BD	Buffer descriptor
BIST	Built-in self test
CD	Collision detect
COL	Collision
CPM	Communication processor module
CRC	Cyclic redundancy check
CRS	Carrier sense
CSB	Coherent system bus
CSMA	Carrier-sense multiple access

Table i. Acronyms and Abbreviated Terms (continued)

Term	Meaning
DDR	Double data rate
DEU	Data encryption standard execution unit
DMA	Direct memory access
DRAM	Dynamic random access memory
DTLB	Data translation lookaside buffers
DUART	Dual universal asynchronous receiver/transmitter
EA	Effective address
ECC	Error checking and correction
EHCI	Enhanced host controller interface
EHPI	Enhanced host port interface
EPROM	Erasable programmable read-only memory
FS	Full-speed
FCS	Frame-check sequence
GMII	Gigabit media independent interface
GPCM	General-purpose chip-select machine
GPIO	General-purpose I/O
GPR	General-purpose register
GTM	General purpose timers
IAD	Internet access device
I <sup>2</sup> C	Inter-integrated circuit
IEEE	Institute of Electrical and Electronics Engineers
IOS	I/O sequencer
IPG	Interpacket gap
ISDN	Integrated services digital network
ITLB	Instruction translation lookaside buffer
IU	Integer unit
JTAG	Joint Test Action Group
LALE	LBC external address latch enable
LBC	Local bus controller
LRU	Least recently used
LSB	Least-significant byte
lsb	Least-significant bit
LSU	Load/store unit

Table i. Acronyms and Abbreviated Terms (continued)

Term	Meaning
MAC	Multiply accumulate, media access control
MCP	Machine-check interrupt
MDI	Medium-dependent interface
MDEU	Message digest execution unit
MIB	Management information base
MII	Media independent interface
MMU	Memory management unit
MPH	Multi-port host
MSB	Most-significant byte
msb	Most-significant bit
OSI	Open systems interconnection
PCI	Peripheral component interconnect
PCS	Physical coding sublayer
PIC	Programmable interrupt controller
PIT	Periodic interval timer
PKEU	Public key execution unit
PMA	Physical medium attachment
PMD	Physical medium dependent
POR	Power-on reset
PRI	Primary rate interface
RGMII	Reduced gigabit media independent interface
RISC	Reduced instruction set computing
RMON	Remote monitoring
RMW	Read-modify-write
RNG	Random number generator
RTBI	Reduced ten-bit interface
RTC	Real time clock module
Rx	Receive
RxBD	Receive buffer descriptor
SCL	Serial clock
SDA	Serial data
SFD	Start frame delimiter
SGMII	Serial gigabit media independent interface

**Table i. Acronyms and Abbreviated Terms (continued)**

<b>Term</b>	<b>Meaning</b>
SI	Serial interface
SPI	Serial peripheral interface
SPR	Special-purpose register
SRAM	Static random access memory
TAP	Test access port
TBI	Ten-bit interface
TLB	Translation lookaside buffer
TSEC	Three-speed Ethernet controller
Tx	Transmit
TxBD	Transmit buffer descriptor
UART	Universal asynchronous receiver/transmitter
ULPI	USB low-pin count interface
UPM	User-programmable machine
USB	Universal serial bus
UTMI	USB transceiver macrocell interface
UTP	Unshielded twisted pair
WDT	Watchdog timer
ZBT	Zero bus turnaround



# Chapter 1

## Overview

This chapter provides an overview of the MPC8313E PowerQUICC™ II Pro processor features, including a block diagram showing the major functional components. The MPC8313E is a cost-effective, low-power, highly integrated host processor that addresses the requirements of several printing and imaging, consumer, and industrial applications, including main CPUs and I/O processors in printing systems, networking switches and line cards, wireless LANs (WLANs), network access servers (NAS), VPN routers, intelligent NIC, and industrial controllers. The MPC8313E extends the PowerQUICC family, adding higher CPU performance, additional functionality, and faster interfaces while addressing the requirements related to time-to-market, price, power consumption, and package size. This manual is written from the perspective of the MPC8313E, and unless otherwise noted, the information applies also to the MPC8313. Note that the MPC8313 does not support a security engine. The MPC8313E contains an embedded PowerPC™ e300c3 core built on Power Architecture™ technology.

### 1.1 MPC8313E PowerQUICC II Pro Processor Overview

Figure 1-1 shows the major functional units within the MPC8313E. The e300c3 core in the MPC8313E, with its 16 Kbytes of instruction and 16 Kbytes of data cache, implements the PowerPC user instruction set architecture and provides hardware and software debugging support. In addition, the MPC8313E offers dual three-speed 10, 100, and 1000 Mbps Ethernet controllers (eTSECs), a DDR1/DDR2 SDRAM memory controller, an enhanced local bus controller (eLBC), a 32-bit PCI-2.3 controller, a dedicated security engine, a programmable interrupt controller, dual I<sup>2</sup>C controllers, a 4-channel DMA controller, a general-purpose I/O port, and a USB 2.0 host and device controller with an on-chip high-speed USB 2.0 PHY. The high level of integration in the MPC8313E helps simplify board design and offers significant bandwidth and performance.




Figure 1-1. MPC8313E Block Diagram

The major features of this device are as follows:

- e300c3 Power Architecture™ processor core
  - High-performance, superscalar processor core with a four-stage pipeline and low interrupt latency times
  - Floating-point, dual integer units, load/store, system register, and branch processing units
  - 16-Kbyte instruction cache and 16-Kbyte data cache with lockable capabilities
  - Capable of completing two MACs every 3 cycles
  - Dynamic power management
  - Enhanced hardware program debug features
  - Software-compatible with Freescale processor families implementing Power Architecture technology
  - Separate PLL that is clocked by the system bus clock
  - Performance monitor
- Security engine optimized to handle all the algorithms associated with IPsec, SSL/TLS, SRTP, IEEE 802.11i® standard, iSCSI, and IKE processing. The security engine contains one crypto-channel, a controller, and a set of crypto execution units (EUs). The execution units are:
  - Data encryption standard execution unit (DEU)
    - DES and 3DES algorithms
    - Two key (K1, K2) or three key (K1, K2, K3) for 3DES
    - ECB and CBC modes for both DES and 3DES
  - Advanced encryption standard unit (AESU)
    - Implements the Rijndael symmetric-key cipher
    - Key lengths of 128-, 192-, and 256-bits
    - ECB, CBC, CCM, and counter (CTR) modes
  - Message digest execution unit (MDEU)
    - SHA with 160-, 224-, or 256-bit message digest
    - MD5 with 128-bit message digest
    - HMAC with either algorithm
  - One crypto-channel supporting multi-command descriptor chains
    - Dynamic assignment of crypto-execution units through an integrated controller
    - Buffer size of 256 bytes for each execution unit, with flow control for large data sizes
    - Support for multiple outstanding bus transactions
    - Scatter/gather capability
    - Fetch FIFOs in the crypto-channel
- DDR SDRAM memory controller
  - Programmable timing supporting both DDR1 and DDR2 SDRAM
  - 16-/32-bit data interface, up to 333-MHz data rate
  - 512-Mbyte addressable space

- Support for two x16 devices
- The following SDRAM configurations are supported:
  - Up to two physical banks (chip selects), each bank up to 1 Gbyte independently addressable
  - 64-Mbit to 1-Gbit devices with x8/x16/x32 data ports (no direct x4 support). Some 2-Gbit devices are supported depending on the internal device configuration.
  - One 16-bit device or two 8-bit devices on a 16-bit bus, or one 32-bit device or two 16-bit devices or four 8-bit devices on a 32-bit bus
- Support for up to 8 simultaneous open pages
- Sleep-mode support for SDRAM self refresh
- Supports auto refresh
- On-the-fly power management using CKE
- Registered DIMM support
- 2.5-V SSTL2 compatible I/O for DDR1, 1.8-V SSTL\_18 compatible I/O for DDR2
- Two enhanced three-speed Ethernet controllers (eTSECs)
  - Backward compatible with MPC8548 (PowerQUICC III) eTSEC
  - Three-speed support (10/100/1000 Mbps)
  - On-chip high-speed serial interface to external SGMII PHY interface
  - Two SGMII interfaces, two RGMII/RTBI/MII/RMII interfaces.
  - Two controllers designed to comply with IEEE Std 802.3®, 802.3u®, 802.3x®, 802.3z®, 802.3ac®, and 802.3ab®
  - Support for IEEE Std 1588™
  - Support for two full-duplex FIFO interface modes
  - Multiple PHY interface configurations
  - Support for Wake-on-Magic Packet™, a method to bring the device from standby to full operating mode
  - TCP/IP acceleration and QoS features available
    - IP v4 and IP v6 header recognition on receive
    - IP v4 header checksum verification and generation
    - TCP and UDP checksum verification and generation
    - Per-packet configurable acceleration
    - Recognition of VLAN, stacked (queue in queue) VLAN, 802.2, PPPoE session, MPLS stacks, and ESP/AH IP-security headers
    - Supported in all FIFO modes
    - Transmission from up to eight physical queues
    - Reception to up to eight physical queues
  - Full- and half-duplex Ethernet support (1000 Mbps supports only full-duplex):
    - IEEE Std. 802.3 full-duplex flow control (automatic PAUSE frame generation or software-programmed PAUSE frame generation and recognition)

- Programmable maximum frame length supports jumbo frames (up to 9.6 Kbytes) and IEEE Std 802.1Q virtual local area network (VLAN) tags and priority
- VLAN insertion and deletion
  - Per-frame VLAN control word or default VLAN for each eTSEC
  - Extracted VLAN control word passed to software separately
- Retransmission following a collision
- CRC generation and verification of inbound/outbound packets
- Programmable Ethernet preamble insertion and extraction of up to 7 bytes
- MAC address recognition:
  - Exact match on primary and virtual 48-bit unicast addresses
  - VRRP and HSRP support for seamless router fail-over
  - Up to 16 exact-match MAC addresses supported
  - Broadcast address (accept/reject)
  - Hash table match on up to 512 multicast addresses
  - Promiscuous mode
- Buffer descriptors backward compatible with MPC8260 and MPC860T 10/100 Ethernet programming models
- RMON statistics support
- 10-Kbyte internal transmit and 2-Kbyte receive FIFOs
- MII management interface for control and status
- Two SGMII PHY interfaces
  - Both interfaces share a single PLL
  - Each interface supports a 4-wire differential (Tx, Rx) interface
  - Support for auto-negotiation
- PCI interface
  - Designed to comply with *PCI Local Bus Specification, Revision 2.3*
  - 32-bit PCI interface operating at up to 66 MHz
  - PCI 3.3-V compatible
  - Support for host and agent modes
  - Memory prefetching of PCI read accesses and support for delayed read transactions
  - On-chip arbitration, supporting three external PCI masters
  - Selectable hardware-enforced coherency
  - Support for PCI Power Management 1.2
  - Support for PME generation (agent) and wake-on-PME
- Universal serial bus (USB) dual-role controller
  - Designed to comply with *Universal Serial Bus Revision 2.0 Specification*
  - Supports operation as a stand-alone USB host controller
    - Supports USB root hub with one downstream-facing port

- Enhanced host controller interface (EHCI) compatible
- Supports operation as a stand-alone USB device
  - Supports one upstream-facing port
  - Supports three programmable bidirectional USB endpoints
- Supports high-speed (480-Mbps), full-speed (12-Mbps), and low-speed (1.5-Mbps) operations. Low speed is only supported in host mode.
- Supports USB on-the-go mode when using an external ULPI (UTMI+ low-pin interface) PHY, which includes both device and host functionality
- On-chip USB-2.0 full-/high-speed PHY with ULPI (UTMI+ low-pin interface)
- Supports Wake-on-USB, a method for bringing the device from standby mode to full operating mode
- Host and device support
- Enhanced local bus controller (eLBC)
  - Multiplexed and non-multiplexed 26-bit address and 8-/16-bit data operating at up to 66 MHz
  - Four chip selects supporting four external slaves
  - Variable memory block sizes (32 Kbytes to 4 Gbytes in FCM mode, 32 Kbytes to 64 Mbytes in UPM mode, and 32 Kbytes to 64 Mbytes in GPCM mode)
  - Supports boot from parallel NOR Flash and parallel NAND Flash
  - Supports programmable clock ratio dividers
  - Up to eight-beat burst transfers
  - 16- and 8-bit ports
  - Three protocol engines available on a per chip select basis:
    - General-purpose chip select machine (GPCM)
    - Three user programmable machines (UPMs)
    - NAND Flash control machine (FCM)
  - Default boot ROM chip select with configurable bus width (8 or 16 bits)
- Integrated programmable interrupt controller (IPIC)
  - Functional and programming compatibility with the MPC8349 interrupt controller
  - Support for external and internal discrete interrupt sources
  - Programmable highest priority request
  - Four groups of interrupts with programmable priority
  - External and internal interrupts directed to host processor
  - Unique vector number for each interrupt source
- Dual I<sup>2</sup>C interfaces
  - Each controller operates up to 400 kHz
  - Two-wire interface
  - Multiple-master support
  - Master or slave I<sup>2</sup>C mode support

- On-chip digital filtering rejects spikes on the bus
- I/O sequencer
- DMA (Direct memory access) controller
  - Four independent virtual channels
  - Concurrent execution across multiple channels with programmable bandwidth control
  - All channels accessible by local core and remote PCI masters
  - Misaligned transfer capability for source/destination address
  - Data chaining and direct mode
  - Interrupt on completed segment and chain
- DUART
  - Two 4-wire interfaces (RxD, TxD,  $\overline{\text{RTS}}$ ,  $\overline{\text{CTS}}$ )
  - Programming model compatible with the original 16450 UART and the PC16550D
- Serial peripheral interface (SPI)
  - Master or slave support
- Power management controller (PMC)
  - Provides power management when the device is used in both host and agent modes
  - Supports PCI Power Management 1.2 D0, D1, D2, D3hot, and D3cold states
  - On-chip split power supply controlled through external power switch for minimum standby power
  - Supports PME generation in PCI agent mode and PME detection in PCI host mode
  - Supports wake-up from Ethernet Magic Packet, USB, GPIO, PCI (PME input as host), timer and external interrupts
  - Supports MPC8349E backward-compatibility mode
- Parallel I/O
  - General-purpose I/O (GPIO)
    - 32 parallel I/O pins multiplexed on various chip interfaces
    - One dedicated GPIOs
  - Open drain capability
  - Interrupt capability
- System timers
  - Periodic interrupt timer
  - Real-time clock
  - Software watchdog timer
  - Two general-purpose timers
- IEEE Std. 1149.1™ compliant JTAG boundary scan
- Integrated PCI bus and SDRAM clock generation

## 1.2 MPC8313E Architecture Overview

The following sections describe the major functional units of this device.

### 1.2.1 Power Architecture Core

The device contains the e300c3 Power Architecture processor core, which is an enhanced version of the MPC603e core (used in previous generations of PowerQUICC II processors). Enhancements include integrated parity checking, dual integer units, and other performance-enhancing features. The e300 core is upward software-compatible with existing MPC603e core-based products.

For detailed information regarding the processor core refer to the following:

- The *e300 Power Architecture™ Core Family Reference Manual* (chapters describing the programming model, cache model, memory management model, exception model, and instruction timing) (Document No. E300CORERM)
- The *Programming Environments Manual for 32-Bit Implementations of the PowerPC™ Architecture* (Document No. MPCFPE32B)

The e300 core is a low-power implementation of the family of microprocessors that implements Power Architecture technology. The core implements the 32-bit portion of the architecture, which provides 32-bit effective addresses, integer data types of 8, 16, and 32 bits, and floating-point data types of 32 and 64 bits.

The core is a superscalar processor that can issue three instructions (two plus a branch) and completes and retires as many as two instructions per clock cycle. Instructions can execute out of order for increased performance; however, the core makes completion appear sequential.

The e300c3 core integrates six execution units—two integer units (IU1 and IU2) with full multiply and divides, a floating-point unit (FPU), a branch processing unit (BPU) with static branch prediction, a load/store unit (LSU) for data transfers, a performance monitor, and a system register unit (SRU). The ability to execute five instructions in parallel and the use of simple instructions with rapid execution times yield high efficiency and throughput. Most integer instructions execute in one clock cycle; two integer instructions may be executed at the same time with the dual integer units. The FPU is pipelined so a single-precision multiply-add instruction can be issued and completed every clock cycle.

The e300c3 core provides independent on-chip, 16-Kbyte, eight-way set-associative, physically-addressed instruction and data caches with parity and integrated way lock capabilities. The processor also features independent on-chip instruction and data memory management units (MMUs). The MMUs contain 64-entry, two-way set-associative, data and instruction translation lookaside buffers (DTLB and ITLB) that provide support for demand-paged virtual memory address translation. The caches use a pseudo least recently used (PLRU) replacement algorithm; the TLBs use a least recently used (LRU) replacement algorithm. The processor also supports block address translation through the use of two independent instruction and data block address translation (IBAT and DBAT) arrays of eight entries each. Effective addresses are compared simultaneously with all eight entries in the BAT array during block translation. In accordance with the architecture, if an effective address hits in both the TLB and BAT array, the BAT translation takes priority.

As an added feature to the e300 core, the device can lock the contents of three of the four ways in the instruction and data cache (or an entire cache). For example, this allows embedded applications to lock

## Overview

interrupt routines or other important (time-sensitive) instruction sequences into the instruction cache. It allows data to be locked into the data cache, which may be important to code that must have deterministic execution.

The e300 core has high-performance 64-bit data bus and 32-bit address bus interfaces to the rest of the device. The e300 core supports single-beat and burst data transfers for memory accesses, and memory-mapped I/O operations.

[Figure 1-2](#) provides a block diagram of the e300 core that shows how the execution units (IU1, IU2, FPU, BPU, LSU, and SRU) operate independently and in parallel. Note that this is a conceptual diagram and does not attempt to show how these features are physically implemented on the chip.






Figure 1-2. MPC8313E Integrated e300c3 Core Block Diagram

## 1.2.2 Security Engine

A hardware encryption block is also integrated in the device. It supports many encryption algorithms allowing for high performance data encryption and authentication as required in today's SoHo/RoBo routers. The encryption block is compatible with the corresponding block in the MPC8280.

The security engine supports DES, 3DES, MD-5, SHA-1, AES, and RC-4 encryption algorithms in hardware.

A block diagram of the security engine's internal architecture is shown in Figure 1-3. The bus interface module is designed to transfer 64-bit words between the internal bus and any register inside the security engine.

An operation begins with a write of a pointer to a crypto-channel fetch register that points to a data packet descriptor. The channel requests the descriptor and decodes the operation to be performed. The channel then requests the controller to assign crypto execution units and fetch the keys, IVs, and data needed to perform the given operation. The controller satisfies the requests by assigning execution units to the channel and by making requests to the master interface. As data is processed, it is written to the individual execution unit's output buffer and then back to system memory through the bus interface module.

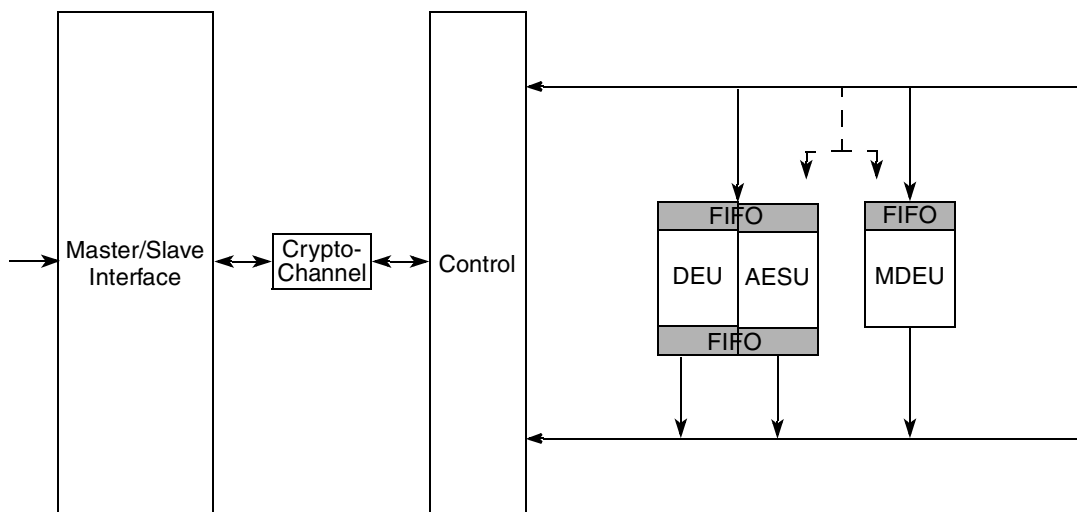


Figure 1-3. Integrated Security Engine Functional Blocks

## 1.2.3 DDR Memory Controller

This fully programmable DDR SDRAM controller supports most JEDEC standard x8 or x16 DDR1 or DDR2 memories available today, including buffered and unbuffered DIMMs. However, mixing nonregistered and registered DIMMs in the same system is not supported. Dynamic power management and auto-precharge modes simplify memory system design.

The DDR memory controller includes the following features:

- Support for DDR1 and DDR2 SDRAM
- 16- or 32-bit SDRAM data bus
- Programmable settings for meeting all SDRAM timing parameters

- Many different SDRAM configurations supported
  - Support for as many as two physical banks (chip selects), each bank independently addressable
  - Support for 64-Mbit to 1-Gbit devices with x8/x16/x32 data ports. Some 2-Gbit devices are supported depending on the internal device configuration.
  - Support for unbuffered and registered
- Support for data mask signals and read-modify-write operations for sub-double word writes
- Four-entry input request queue
- Open page management (dedicated entry for each sub-bank)
- Memory controller clock frequency of two times the SDRAM clock with support for sleep power management mode

## 1.2.4 Dual Enhanced Three-Speed Ethernet Controllers

The MPC8313E has two on-chip enhanced three-speed Ethernet controllers. The eTSECs incorporate a media access control (MAC) sublayer that supports 10- and 100-Mbps and 1-Gbps Ethernet/IEEE Std. 802.3 networks with MII, RMII, RGMII, RTBI, and SGMII physical interfaces. The eTSECs include 2-Kbyte receive and 10-Kbyte transmit FIFOs and DMA functions. They also support IEEE Std. 1588.

The buffer descriptors are based on the MPC8260 and MPC860T 10/100 Ethernet programming models. Each eTSEC can emulate a PowerQUICC III TSEC, allowing existing driver software to be re-used with minimal change.

The MPC8313E eTSECs support programmable CRC generation and checking, RMON statistics, and jumbo frames of up to 9.6 Kbytes. Frame headers and buffer descriptors can be forced into the L2 cache to speed classification or other frame processing.

Each eTSEC provides hardware support for accelerating TCP/IP packet transmission and reception. By default, TCP/IP acceleration is not enabled, and the eTSEC processes frames as pure Ethernet frames.

TCP/IP acceleration can be performed at a number of levels. The eTSEC can parse frames at layer 2 of the stack only (Ethernet headers and switching headers), layers 2 to 3 (including IP v4 or IP v6), or layers 2 to 4 (including TCP and UDP).

On receive, the eTSEC provides protocol header recognition, header verification (IP v4 header checksum verification), and TCP/UDP payload checksum verification including verification of associated pseudo-header checksums. On transmit, the eTSEC provides IP v4 and TCP/UDP header checksum generation. The eTSEC does not checksum transmitted packets with IP header options or IP fragments.

To provide for quality of service, transmission from up to eight queues is supported with priority-based queue selection. Arbitration is a modified weighted round-robin queue selection with fair bandwidth allocation.

On receive, packets may be distributed to any of the 64 virtual receive queues overlaid onto the 8 physical receive queues. A table-oriented queue filing strategy is provided based on 16 header fields or flags. Frame rejection is supported for filtering applications.

Filing can be based on Ethernet, IP, and TCP/UDP properties, including VLAN fields, Ether-type, IP protocol type, IP TOS or differentiated services, IP source and destination addresses, TCP/UDP port numbers, or user-defined bit fields.

Each eTSEC provides a full-duplex packet FIFO interface port that bypasses the Ethernet MAC but reuses the PHY interface pins. As a result, the FIFO interface normally does not impose the overheads of Ethernet framing. The FIFO interface operates synchronously, at up to 200MHz, providing up to 3.2-Gbps full-duplex transfer rates. Bare IP packets, with an optional 32-bit CRC check sequence, can be transferred to the eTSEC directly. The eTSEC Tx and Rx FIFOs, TCP/IP acceleration functions, and DMA continue to be used in packet FIFO mode.

### NOTE

While some of the Ethernet interfaces support either 2.5- or 3.3-V operation, the voltages of eTSEC1 and eTSEC2 must be the same.

Table 1-1 lists available configurations.

**Table 1-1. Supported eTSEC1 and eTSEC2 Configurations<sup>1</sup>**

Mode Option	eTSEC1	eTSEC2
Ethernet standard interfaces	TBI, GMII, or MII	TBI, GMII, or MII
Ethernet reduced interfaces	RTBI, RGMII, or RMII	RTBI, RGMII, or RMII
FIFO and mixed interfaces	8-bit FIFO	RTBI, RGMII, RMII, or 8-bit FIFO
	TBI, GMII, MII, RTBI, RGMII, RMII, or 8-bit FIFO	8-bit FIFO

<sup>1</sup> Both interfaces must use the same voltage.

## 1.2.5 PCI Controller

The 32-bit PCI controller is compatible with the *PCI Local Bus Specification, Rev. 2.3*. The PCI interface can function as a host bridge interface. The PCI interface can optionally function as an agent device. The PCI controller supports 32-bit addressing and 32-bit data buses.

As a host, the device supports read and write operations to the PCI memory space, the PCI I/O space, and the PCI configuration space. Also, the device can generate PCI special-cycle and interrupt acknowledge commands. As an agent, the device supports read and write operations to system memory, as well as PCI configuration space and the on-chip memory mapped configuration space.

The device PCI controller includes the following distinctive features:

- Address stepping on configuration transactions
- Fast back-to-back transactions
- Data streaming
- When in host mode, the PCI controller supports external signal isolation, thus enabling power shut off to external devices
- Supports PCI Power Management 1.2
- Supports PME generation (agent) and Wake on PME

### 1.2.5.1 PCI Bus Arbitration Unit

The PCI controller contains a PCI bus arbitration unit, which eliminates the need for an external unit, thus lowering system complexity and cost. It has the following features:

- Supports three  $\overline{\text{REQ}}/\overline{\text{GNT}}$  signal pairs, thus supporting three external masters. The device PCI controller is the fourth member of the arbitration pool.
- The bus arbitration unit allows fairness as well as a priority mechanism.
- A two-level round-robin scheme is used in which each device can be programmed within a pool of a high- or low-priority arbitration. One member of the low-priority pool is promoted to the high-priority pool. As soon as it is granted the bus, it returns to the low-priority pool.
- The unit can be disabled to allow a remote arbitration unit to be used.
- The unit can be isolated to allow power shut off of external devices.

### 1.2.6 Universal Serial Bus (USB) 2.0

The USB 2.0 controller offers operation as a host or device. The USB controller provides point-to-point connectivity, which complies with the *Universal Serial Bus Revision 2.0 Specification*. The USB controller can be configured to operate as a stand-alone host or stand-alone device. See [Figure 1-4](#) for more information.

The host and device functions are both configured to support the following four types of USB transfers:

- Bulk
- Control
- Interrupt
- Isochronous

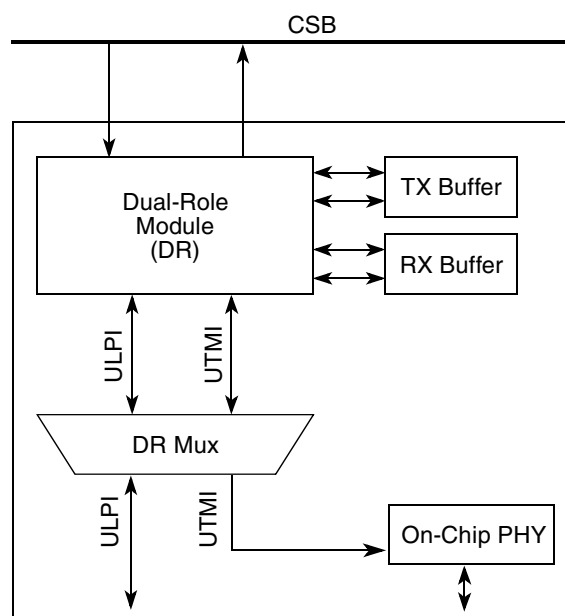


Figure 1-4. USB Controllers Port Configuration

### 1.2.6.1 USB Dual-Role Controller

- Designed to comply with *Universal Serial Bus Revision 2.0 Specification*
- Supports operation as a stand-alone USB host controller
  - Supports USB root hub with one downstream-facing port
  - Enhanced host controller interface (EHCI) compatible
- Supports operation as a stand-alone USB device
  - Supports one upstream-facing port
  - Supports three programmable bidirectional USB endpoints
- Supports high-speed (480-Mbps), full-speed (12-Mbps), and low-speed (1.5-Mbps) operations. Low speed is only supported in host mode.
- Supports USB on-the-go mode when using an external ULPI (UTMI+ low-pin interface) PHY, which includes both device and host functionality
- On-chip USB-2.0 full-/high-speed PHY with ULPI (UTMI+ low-pin interface)
- Supports Wake-on-USB, a method for bringing the device from standby mode to full operating mode
- Host and device support

### 1.2.7 Enhanced Local Bus Controller (eLBC)

The main component of the enhanced local bus controller (eLBC) is its memory controller, which provides a seamless interface to many types of memory devices and peripherals. The memory controller is responsible for controlling four memory banks shared by a NAND Flash control machine (FCM), a general-purpose chip-select machine (GPCM), and up to three user-programmable machines (UPMs). As such, it supports a minimal glue logic interface to SRAM, EPROM, NOR Flash EPROM, NAND Flash EPROM, Flash EPROM, burstable RAM, and other peripherals. The eLBC external address latch enable (LALE) signal allows multiplexing of addresses with data signals to reduce the device pin count.

The enhanced local bus controller also includes a number of data checking and protection features such as data parity generation and checking, write protection, and a bus monitor to ensure that each bus cycle is terminated within a user-specified period.

The main features of the enhanced local bus controller (eLBC) are as follows:

- Memory controller with four memory banks (chip selects)
  - 32-bit address decoding with mask
  - Variable memory block sizes (32 Kbytes to 4 Gbytes in FCM mode, 32 Kbytes to 64 Mbytes in UPM mode, and 32 Kbytes to 64 Mbytes in GPCM mode)
  - Selection of control signal generation on a per-bank basis
  - Data buffer controls activated on a per-bank basis
  - Up to 256-byte bursts, arbitrarily aligned
  - Automatic segmentation of large transactions into memory accesses optimized for bus width and addressing capability

- Write-protection capability
- Atomic operation
- General-purpose chip-select machine (GPCM)
  - Compatible with SRAM, EPROM, NOR Flash EEPROM, FEPRM, and peripherals
  - Global (boot) chip-select available at system reset
  - Boot chip-select support for 8- and 16-bit devices
  - Minimum three-clock access to external devices
  - Two byte-write-enable signals ( $\overline{\text{LWE}}[0:1]$ )
- NAND Flash control machine (FCM)
  - Compatible with small (512 + 16 bytes) and large (2048 + 64 bytes) page parallel NAND Flash EEPROM
  - Global (boot) chip-select available at system reset, with 4-Kbyte boot block buffer for execute-in-place boot loading
  - Boot chip-select support for 8-bit devices
  - Dual 2-Kbyte/eight 512-byte buffers allow simultaneous data transfer during Flash reads and programming
  - Interrupt-driven block transfer for reads and writes
  - Programmable command and data transfer sequences of up to eight steps supported
  - Generic command and address registers support proprietary Flash interfaces
  - Block write locking to ensure system security and integrity
- Three user-programmable machines (UPMs)
  - Programmable-array-based machine controls external signal timing with a granularity of up to one quarter of an external bus clock period
  - User-specified control-signal patterns run when an internal master requests a single-beat or burst read or write access
  - UPM refresh timer runs a user-specified control signal pattern to support refresh
  - User-specified control-signal patterns can be initiated by software
  - Each UPM can be defined to support DRAM devices with depths of 64, 128, 256, and 512 Kbytes, and 1, 2, 4, 8, 16, 32, 64, 128, and 256 Mbytes
  - Support for 8- and 16-bit devices
  - Page mode support for successive transfers within a burst
  - Internal address multiplexing supporting 64-, 128-, 256-, and 512-Kbyte, and 1-, 2-, 4-, 8-, 16-, 32-, 64-, 128-, and 256-Mbyte page banks
- Optional monitoring of transfers between local bus internal masters and local bus slaves (local bus error reporting)

## 1.2.8 Integrated Programmable Interrupt Controller (IPIC)

The IPIC implements the necessary functions to provide a flexible solution for general-purpose interrupt control. The IPIC includes the following features:

- Functional and programming models are compatible with the MPC8260 interrupt controller
- Support for external and internal discrete interrupt sources
- Support for one external (optional) and seven internal machine checkstop interrupt sources
- Programmable highest priority request
- Two programmable priority mixed groups of four on-chip and four external interrupt signals with two priority schemes for each group: grouped and spread
- Two programmable priority internal groups of eight on-chip interrupt signals with two priority schemes for each group: grouped and spread
- Priority interrupts can be programmed to support a critical ( $\overline{cint}$ ) or system management ( $\overline{smi}$ ) interrupt type
- External and internal interrupts directed to a host processor
- Unique vector number for each interrupt source
- Ability to redirect interrupts to external  $\overline{PCI\_INTA}$  pin when in core disable mode

## 1.2.9 Dual I<sup>2</sup>C Interfaces

The inter-IC (IIC or I<sup>2</sup>C) bus is a two-wire—serial data (SDA) and serial clock (SCL)—bidirectional serial bus that provides a simple, efficient method of data exchange between the system and other devices, such as microcontrollers, EEPROMs, real-time clock devices, A/D converters, and LCDs. The two-wire bus minimizes the interconnections between devices. The synchronous, multi-master bus of the I<sup>2</sup>C allows the connection of additional devices to the bus for expansion and system development.

The I<sup>2</sup>C controller is a true multi-master bus which includes collision detection and arbitration that prevents data corruption if two or more masters attempt to control the bus simultaneously. This feature allows for complex applications with multiprocessor control. The I<sup>2</sup>C controller consists of a transmitter/receiver unit, clocking unit, and control unit. The I<sup>2</sup>C unit supports general broadcast mode and on-chip filtering rejects spikes on the bus.

The I<sup>2</sup>C interfaces include the following features:

- Two-wire interface
- Multi-master operational
- Arbitration lost interrupt with automatic mode switching from master to slave
- Calling address identification interrupt
- Bus busy detection
- Software-programmable clock frequency
- Software-selectable acknowledge bit
- On-chip filtering for spikes on the bus
- Address broadcasting supported



- System initialization data is optionally loaded from I<sup>2</sup>C EPROM by boot sequencer embedded hardware

### 1.2.10 DMA Controller

The DMA engine is capable of transferring blocks of data from any legal address range to any other legal address range. Therefore, it can perform a DMA transfer between any of its I/O or memory ports, or even between two devices or locations on the same port.

The DMA controller offers the following features:

- Four high-speed/high-bandwidth channels accessible by local and remote masters
- Basic DMA operation modes (direct, simple chaining)
- Support for misaligned transfers
- Programmable bandwidth control between channels
- Interrupt on error and completed segment or chain

### 1.2.11 Dual Universal Asynchronous Receiver/Transmitter (DUART)

The device includes a DUART intended for use in maintenance, bring up, and debug systems. The device provides a standard four-wire handshake (TXD, RXD,  $\overline{\text{RTS}}$ ,  $\overline{\text{CTS}}$ ) for each port. The DUART is a slave interface. An interrupt is provided to the interrupt controller or optionally steered externally to allow device handshakes. Interrupts are generated for transmit, receive, and line status.

The DUART supports full-duplex operation. It is compatible with the PC16450 and PC16550 programming models. The transmitter and receiver both support 16-byte FIFOs.

Software programmable baud rate generators divide the system clock to generate a 16x clock. Serial interface data formats (data length, parity, 1/1.5/2 STOP bit, baud rate) are also software selectable.

The DUART includes the following features:

- Full-duplex operation
- Programming model compatible with the original PC16450 UART and the PC16550D (an improved version of the PC16450 that also operates in FIFO mode)
- PC16450 register reset values
- FIFO mode for both transmitter and receiver, providing 16-byte FIFOs
- Serial data encapsulation and decapsulation with standard asynchronous communication bits (START, STOP, and parity)
- Maskable transmit, receive, and line status interrupts
- Software-programmable baud rate generators that divide the system clock by 1 to  $(2^{16} - 1)$  and generate a 16x clock for the transmitter and receiver engines
- Clear to send ( $\overline{\text{CTS}}$ ) and ready to send ( $\overline{\text{RTS}}$ ) MODEM control functions
- Software-selectable serial-interface data format (data length, parity, 1/1.5/2 STOP bit, baud rate)
- Line status registers
- Line-break detection and generation

## Overview

- Internal diagnostic support, local loopback, and break functions
- Prioritized interrupt reporting
- Overrun, parity, and framing error detection

### 1.2.12 Serial Peripheral Interface (SPI)

The serial peripheral interface (SPI) allows the device to exchange data between other PowerQUICC family chips, Ethernet PHYs for configuration, and peripheral devices such as EEPROMs, real-time clocks, A/D converters, and ISDN devices.

The SPI is a full-duplex, synchronous, character-oriented channel that supports a four-wire interface (receive, transmit, clock, and slave select). The SPI block consists of transmitter and receiver sections, an independent baud-rate generator, and a control unit.

### 1.2.13 System Timers

The system includes the following timers:


- Periodic interrupt timer
- Real time clock
- Software watchdog timer
- Two general-purpose timer blocks, each supporting four 16-bit programmable timers, two cascaded 32-bit timers, or one cascaded 64-bit counter

## 1.3 Application Examples

The internal features of the MPC8313E make it suitable for a wide variety of printer and network communication applications as described in this section.

### 1.3.1 Low-End Printer CPU and Interface ASIC

Figure 1-5 illustrates how the MPC8313E can perform the function of the CPU + interface ASIC on a low-end printer application.



**Figure 1-5. MPC8313E Serving as the Main CPU in a Low-End Printer Application**

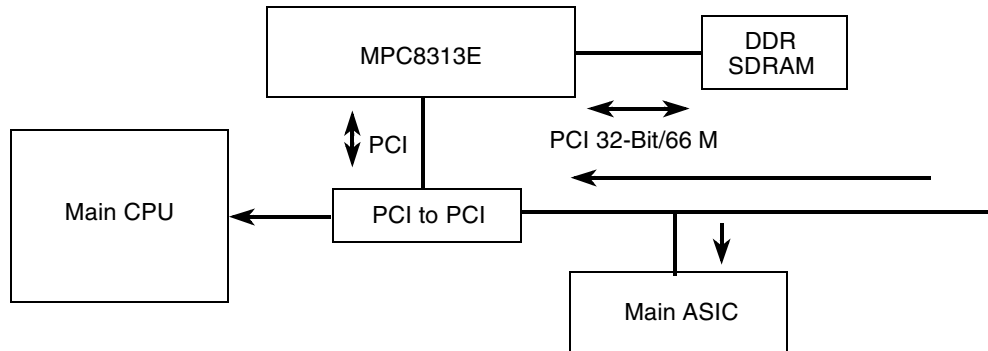
In this application, the CPU interfaces to the main ASIC through the high-bandwidth PCI bus. Low-end multi-function printers (MFPs) are able to share the same platform simply by adding a scanner/fax engine. The interface ASIC provides the various network interfaces that are used to access the printer.

Image data coming through the scanner or fax interface is sent to the main ASIC, which processes the image by implementing algorithms for image compression/decompression and rendering. The image data is then processed in the FPU in the CPU core at high speeds and sent to the printer engine. Likewise, image data or text data that are interfaced at the Ethernet interface on the interface ASIC are also manipulated at the main ASIC and CPU and sent to the printer engine.

Recent MFP systems require higher processor performance in order to manipulate large, high-quality images at high speeds. Required networking interfaces including USB, PCI, and Gigabit Ethernet are integrated on the MPC8313E, allowing for the CPU and interface ASIC to be consolidated in one device. As a result, an MFP application can be developed by combining the MPC8313E with the main ASIC (graphic processing ASIC) at a lower cost without the need to have separate a CPU and interface ASIC. At the same time, the system is required to consume low power. The MPC8313E provides several power management methods to reduce power consumption.

### 1.3.2 High-End Printer I/O Processor

The diagram in [Figure 1-6](#) illustrates how an MPC8313E can function as an I/O processor in a high-end printer application.



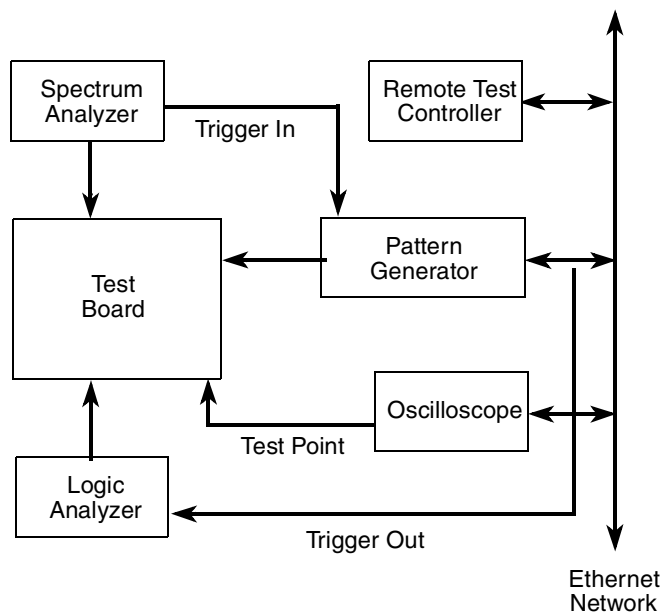
**Figure 1-6. MPC8313E I/O Processor Implementation in a High-End Printer Application**

In this example, the MPC8313E primarily functions as a controller for the main ASIC. The MPC8313E and the main ASIC connect to the main CPU through a PCI bridge.

The power consumption of high-end MFP systems is significant due to the required top-level processor performance and high-speed interfaces. In order to reduce the power consumption in stand-by mode, the MPC8313E as a secondary I/O processor shuts off the power to the main CPU and maintains the system at very low power. Once the I/O processor detects data transfer through the LAN, USB and PCI or an interrupt from the push of a button on the panel, it quickly boots up the main CPU.

### 1.3.3 IEEE Std. 1588 in Test and Measurement and Industrial Automation


Figure 1-7 shows how a test and measurement application can use the IEEE Std. 1588 precise time synchronization supported in the MPC8313E.



**Figure 1-7. IEEE Std. 1588 in Test and Measurement**

In this application, IEEE Std. 1588 allows coordination and control of test and measurement equipment over a distributed Ethernet network. Precise timing delivery allows test equipment to deliver patterns and measure responses at specific times, enabling accurate timestamping of measured data and allowing coordination of input stimuli and any associated measured data. The trigger inputs and outputs enable coordination of other devices.

Figure 1-8 illustrates how an industrial control application is able to take advantage of the IEEE Std. 1588 precise time synchronization.



**Figure 1-8. IEEE Std. 1588 in Industrial Control**

As shown in the example on the top of the figure, IEEE Std. 1588 allows precision control over a distributed Ethernet network. Precise timing delivery allows drive units to be placed where required. Traditional mechanical control mechanisms can limit the placement of systems.

In the bottom of the figure, timing synchronization at the drive enables flexibility in system configuration. Issues due to mismatched cable lengths are minimized. Servos can be added or deleted without having to rewire other servos. Industrial control applications typically augment IEEE Std. 1588 hardware to provide trigger inputs and outputs

In summary, IEEE Std. 1588 support in the MPC8313E enables accurate synchronization of clocks with varying precision, resolution, and oscillator stability in distributed systems. IEEE Std. 1588 synchronizes individual clocks to maintain accurate distribution-wide timing. It enhances applications that need local clocks at each control node and provides sub-microsecond synchronization over long distances using standard cabling. Target applications for systems with IEEE Std. 1588 are test and measurement appliances and industrial control and automation.

### 1.3.4 IEEE Std. 802.11n WLAN Access Point

Figure 1-9 illustrates the MPC8313E acting as an IEEE Std. 802.11n® WLAN access point.




Figure 1-9. MPC8313E as a WLAN Access Point


Current systems are being designed for IEEE Std. 802.11b® and 802.11a/g® as well as combo radios. The WiFi chipsets are on PCI in current systems. With the migration to IEEE 802.11n standard and the increased management features required by the IT community, the need for higher performance is growing.

These WLAN access points (WAPs) require low power and are usually powered exclusively by Power over Ethernet (POE). Each Ethernet line can power about 12 W. With the power the radios draw and the rest of the board this usually only leaves < 2.5 W for the embedded processor.

Integrated SGMII makes it possible to connect to low power Gigabit Ethernet PHYs. Power is a main consideration for WAPs deployed throughout the premise. Power is drawn from the AC outlet in the wall or power over Ethernet. Having an SGMII interface on the Gigabit Ethernet PHYs enables low overall power consumption. The MPC8313E also has superior PCI to memory performance.

### 1.3.5 Media Server

Figure 1-10 shows how the MPC8313E can be configured as a media server.



**Figure 1-10. MPC8313E as a Media Server**

Multimedia home networking emphasizes both audio and video streaming around the house. Since digital audio takes up relatively little bandwidth, just about any current home network can handle streaming, digital audio. Ultimately, it is video that is the test of a multimedia home network. Whether it's compressed MPEG 2 or MPEG 4 streams, or uncompressed, high-definition video, consumers are likely to demand some type of video streaming on a multimedia network. The trend in the industry will be more focused on providing ASSPs (MPEG processors, image processors, integrated digital TV processors, audio/video decoders).



## Chapter 2

# Memory Map

This chapter describes the MPC8313E memory map. The internal memory mapped registers are described, including a complete listing of all memory mapped registers with cross references to the sections detailing descriptions of each.

### 2.1 Internal Memory Mapped Registers

All of the memory mapped registers in the device are contained within a 1-Mbyte address region. To allow for flexibility, the base address of the memory mapped registers is relocatable in the local address space. The local address map location of this register block is controlled by the internal memory mapped registers base address register (IMMRBAR), see [Section 5.2.4.1, “Internal Memory Map Registers Base Address Register \(IMMRBAR\),”](#) for more information. The default value for IMMRBAR is 0xFF40\_0000.

### 2.2 Accessing IMMR Memory From the Local Processor

When the local e300 processor is used to configure IMMR space, the IMMR memory space should typically be marked as cache-inhibited and guarded.

In addition, many configuration registers affect accesses to other memory regions; therefore, writes to these registers must be guaranteed to have taken effect before accesses are made to the associated memory regions.

To guarantee that the results of any sequence of writes to configuration registers are in effect, the final configuration register write should be followed immediately by a read of the same register, and that should be followed by a sync instruction. Then accesses can safely be made to memory regions affected by the configuration register write.

### 2.3 Complete IMMR Map

[Table 2-1](#) lists the memory-mapped register regions (windows). Unless stated otherwise in a particular block, all accesses to and from the memory mapped registers must be made with 32-bit accesses. There is no support for accesses of sizes other than 32 bits.

Reading from address locations which appear as reserved in the memory map table is not guaranteed to return predictable data. Writing to address locations which appear as reserved in the memory map table is not allowed and could lead to unpredictable behavior of the device. Reserved bits in non-reserved registers will be read as zero unless the reset value of those bits is different due to internal logic considerations.

When writing to registers with reserved bits, those reserved bits should be cleared. By doing so, existing software would be able to run on a future modified device in which some reserved bits were allocated for enhanced modes. This would allow for maintaining the legacy functionality when set to zero.

## Memory Map

In certain specific cases, reserved bits should not be cleared but should keep their reset value. Thus, the software should perform a ‘read-modify-write’ and make sure that it does not change the reset value of those bits. The description of the specific bits will indicate when this is needed.

**Table 2-1. IMMR Memory Map**

Address	Use	Actual Size	Window
0x0_0000–0x0_01FF	System configuration	512 bytes	512 bytes
0x0_0200–0x0_02FF	Watchdog timer	16 bytes	256 bytes
0x0_0300–0x0_03FF	Real time clock	32 bytes	256 bytes
0x0_0400–0x0_04FF	Periodic interval timer	32 bytes	256 bytes
0x0_0500–0x0_05FF	Global timers module 1	64 bytes	256 bytes
0x0_0600–0x0_06FF	Global timers module 2	64 bytes	256 bytes
0x0_0700–0x0_07FF	Integrated programable interrupt controller (IPIC)	128 bytes	256 bytes
0x0_0800–0x0_08FF	System arbiter	30 bytes	256 bytes
0x0_0900–0x0_09FF	Reset module	44 bytes	256 bytes
0x0_0A00–0x0_0AFF	Clock module	44 bytes	256 bytes
0x0_0B00–0x0_0BFF	Power management control module	20 bytes	256 bytes
0x0_0C00–0x0_0CFF	GPIO	24 bytes	256 bytes
0x0_0D00–0x0_0DFF	Reserved	—	256 bytes
0x0_0E00–0x0_0EFF	Reserved	—	256 bytes
0x0_0F00–0x0_0FFF	Reserved	—	256 bytes
0x0_1000–0x0_10FF	Reserved	—	256 bytes
0x0_1100–0x0_11FF	Reserved	—	256 bytes
0x0_1200–0x0_12FF	Reserved	—	256 bytes
0x0_1300–0x0_13FF	Reserved	—	256 bytes
0x0_1400–0x0_17FF	Reserved	—	1 Kbyte
0x0_1800–0x0_1BFF	Reserved	—	1 Kbyte
0x0_1C00–0x0_1FFF	Reserved	—	1 Kbyte
0x0_2000–0x0_2FFF	DDR MEMC	3.8 Kbytes	4 Kbytes
0x0_3000–0x0_30FF	I <sup>2</sup> C1 controller	24 bytes	256 bytes
0x0_3100–0x0_31FF	I <sup>2</sup> C2 controller	24 bytes	256 bytes
0x0_3200–0x0_32FF	Reserved	—	256 bytes
0x0_3300–0x0_33FF	Reserved	—	256 bytes
0x0_3400–0x0_37FF	Reserved	—	1 Kbyte
0x0_3800–0x0_3BFF	Reserved	—	1 Kbyte

Table 2-1. IMMR Memory Map (continued)

Address	Use	Actual Size	Window
0x0_3C00–0x0_3FFF	Reserved	—	1 Kbyte
0x0_4000–0x0_44FF	Reserved	—	1.25 Kbytes
0x0_4500–0x0_46FF	DUART	18 bytes x 27	512 bytes
0x0_4700–0x0_4FFF	Reserved	—	2.3 Kbytes
0x0_5000–0x0_5FFF	eLBC	224 bytes	4 Kbytes
0x0_6000–0x0_6FFF	Reserved	—	4 Kbytes
0x0_7000–0x0_7FFF	SPI	24 bytes	4 Kbytes
0x0_8000–0x0_82FF	DMA	768 bytes	768 bytes
0x0_8300–0x0_837F	PCI configuration	16 bytes	128 bytes
0x0_8380–0x0_83FF	Reserved	—	128 bytes
0x0_8400–0x0_84FF	IOS	256 bytes	256 bytes
0x0_8500–0x0_85FF	PCI controller	128 bytes	256 bytes
0x0_8600–0x0_87FF	Reserved	—	512 bytes
0x0_8800–0x0_8FFF	Reserved	—	2 Kbytes
0x0_B000–0x0_BFFF	Reserved	—	4 Kbytes
0x0_C000–0x0_FFFF	Reserved	—	16 Kbytes
0x1_0000–0x1_3FFF	Reserved	—	16 Kbytes
0x1_4000–0x1_5FFF	Reserved	—	8 Kbytes
0x1_6000–0x1_60FF	Reserved	—	256 bytes
0x1_6100–0x1_617F	Reserved	—	128 bytes
0x1_6180–0x1_61FF	Reserved	—	128 bytes
0x1_6200–0x1_63FF	Reserved	—	512 bytes
0x1_6400–0x1_67FF	Reserved	—	1 Kbytes
0x1_6800–0x1_6FFF	Reserved	—	2 Kbytes
0x1_7000–0x1_7FFF	Reserved	—	4 Kbytes
0x1_8000–0x1_8FFF	Reserved	4 Kbytes	4 Kbytes
0x1_9000–0x1_9FFF	Reserved	4 Kbytes	4 Kbytes
0x1_A000–0x1_BFFF	Reserved	—	8 Kbytes
0x1_C000–0x1_FFFF	Reserved	—	16 Kbytes
0x2_0000–0x2_1FFF	Reserved	—	8 Kbytes
0x2_2000–0x2_2FFF	Reserved	—	4 Kbytes
0x2_3000–0x2_3FFF	USB_DR module	1 Kbyte	4 Kbytes

Table 2-1. IMMR Memory Map (continued)

Address	Use	Actual Size	Window
0x2_4000–0x2_4FFF	eTSEC 1	3 Kbytes + 1K reserved	4 Kbytes
0x2_5000–0x2_5FFF	eTSEC 2	3 Kbytes + 1K reserved	4 Kbytes
0x2_6000–0x2_7FFF	Reserved	—	8 Kbytes
0x2_8000–0x2_BFFF	Reserved	—	16 Kbytes
0x2_C000–0x2_DFFF	Reserved	8 Kbytes	8 Kbytes
0x2_E000–0x2_FFFF	Reserved	—	8 Kbytes
0x3_0000–0x3_FFFF	Security engine	52 Kbytes	64 Kbytes
0x4_0000–0x7_FFFF	Reserved	—	256 Kbytes
0x8_0000–0xB_FFFF	Reserved	—	256 Kbytes
0xC_0000–0xD_FFFF	Reserved	—	128 Kbytes
0xE_0000–0xE_1FFF	Reserved	—	8 Kbytes
0xE_2000–0xE_2FFF	Reserved	—	4 Kbytes
0xE_3000–0xE_30FF	Reserved	—	256 bytes
0xE_3100–0xE_31FF	Reserved	—	256 bytes
0xE_3200–0xE_33FF	Reserved	—	512 bytes
0xE_3400–0xE_37FF	Reserved	—	1 Kbyte
0xE_3800–0xE_3FFF	Reserved	—	2 Kbytes
0xE_4000–0xE_7FFF	Reserved	—	16 Kbytes
0xE_8000–0xE_FFFF	Reserved	—	32 Kbytes
0xF_0000–0xF_FFFF	Reserved	—	64 Kbytes

Table 2-2 lists the memory-mapped registers.

Table 2-2. Memory Map

Offset	Register	Access	Reset	Section/Page
<b>System Configuration Registers</b>				
0x0_0000	Internal memory map base address register (IMMRBAR)	R/W	0xFF40_0000	<a href="#">5.2.4.1/5-6</a>
0x0_0004	Reserved	—	—	—
0x0_0008	Alternate configuration base address register (ALTCBAR)	R/W	0x0000_0000	<a href="#">5.2.4.2/5-7</a>
0x0_000C– 0x0_001C	Reserved	—	—	—
0x0_0020	eLBC local access window 0 base address register (LBLAWBAR0)	R/W	0x0000_0000 <sup>1</sup>	<a href="#">5.2.4.3/5-8</a>

Table 2-2. Memory Map (continued)

Offset	Register	Access	Reset	Section/Page
0x0_0024	eLBC local access window 0 attribute register (LBLAWAR0)	R/W	0x0000_0000 <sup>2</sup>	5.2.4.4/5-9
0x0_0028	eLBC local access window 1 base address register (LBLAWBAR1)	R/W	0x0000_0000	5.2.4.3/5-8
0x0_002C	eLBC local access window 1 attribute register (LBLAWAR1)	R/W	0x0000_0000	5.2.4.4/5-9
0x0_0030	eLBC local access window 2 base address register (LBLAWBAR2)	R/W	0x0000_0000	5.2.4.3/5-8
0x0_0034	eLBC local access window 2 attribute register (LBLAWAR2)	R/W	0x0000_0000	5.2.4.4/5-9
0x0_0038	eLBC local access window 3 base address register (LBLAWBAR3)	R/W	0x0000_0000	5.2.4.3/5-8
0x0_003C	eLBC local access window 3 attribute register (LBLAWAR3)	R/W	0x0000_0000	5.2.4.4/5-9
0x0_0040– 0x0_005C	Reserved	—	—	—
0x0_0060	PCI local access window 0 base address register (PCILAWBAR0)	R/W	0x0000_0000 <sup>3</sup>	5.2.4.5/5-10
0x0_0064	PCI local access window 0 attribute register (PCILAWAR0)	R/W	0x0000_0000 <sup>4</sup>	5.2.4.6/5-11
0x0_0068	PCI local access window 1 base address register (PCILAWBAR1)	R/W	0x0000_0000 <sup>5</sup>	5.2.4.5/5-10
0x0_006C	PCI local access window 1 attribute register (PCILAWAR1)	R/W	0x0000_0000	5.2.4.6/5-11
0x0_0070– 0x0_009C	Reserved	—	—	—
0x0_00A0	DDR local access window 0 base address register (DDRLAWBAR0)	R/W	0x0000_0000 <sup>6</sup>	5.2.4.7/5-12
0x0_00A4	DDR local access window 0 attribute register (DDRLAWAR0)	R/W	0x0000_0000 <sup>7</sup>	5.2.4.8/5-13
0x0_00A8	DDR local access window 1 base address register (DDRLAWBAR1)	R/W	0x0000_0000	5.2.4.7/5-12
0x0_00AC	DDR local access window 1 attribute register (DDRLAWAR1)	R/W	0x0000_0000	5.2.4.8/5-13
0x0_00B0– 0x0_00FC	Reserved	—	—	—
0x00100	System general purpose register low (SGPRL)	R/W	0x0000_0000	5.3.2.1/5-17
0x00104	System general purpose register high (SGPRH)	R/W	0x0000_0000	5.3.2.2/5-17
0x00108	System part and revision ID register (SPRIDR)	R	0x80B0_0021	5.3.2.3/5-18
0x0010C	Reserved	—	—	—
0x00110	System priority configuration register (SPCR)	R/W	0x0000_0000	5.3.2.4/5-18
0x00114	System I/O configuration register low (SICRL)	R/W	0x0000_0000 <sup>8</sup>	5.3.2.5/5-21

Table 2-2. Memory Map (continued)

Offset	Register	Access	Reset	Section/Page
0x00118	System I/O configuration register high (SICRH)	R/W	0x0000_0000	<a href="#">5.3.2.6/5-23</a>
0x0011C– 0x00124	Reserved	—	—	—
0x00128	DDR control driver register (DDRCDR)	R/W	0x0004_0000	<a href="#">5.3.2.8/5-27</a>
0x0012C	DDR debug status register (DDRDSR)	R	0x3300_0000	<a href="#">5.3.2.9/5-29</a>
0x00130– 0x0014C	Reserved	—	—	—
0x00150– 0x001FC	Reserved	—	—	—
<b>Watchdog Timer (WDT) Registers</b>				
0x0–0x3	Reserved	—	—	—
0x4	System watchdog control register (SWCRR)	R/W	0xFFFF_0003 or 0xFFFF_0007 <sup>9</sup>	<a href="#">5.4.4.1/5-31</a>
0x8	System watchdog count register (SWCNR)	R	0x0000_FFFF	<a href="#">5.4.4.2/5-32</a>
0xC–0xD	Reserved	—	—	—
0xE	System watchdog service register (SWSRR)	R/W	0x0000	<a href="#">5.4.4.3/5-33</a>
<b>Real Time Clock Module Registers (RTC)</b>				
0x00	Real time counter control register (RTCNR)	R/W	0x0000_0000	<a href="#">5.5.5.1/5-38</a>
0x04	Real time counter load register (RTLDR)	R/W	0x0000_0000	<a href="#">5.5.5.2/5-39</a>
0x08	Real time counter prescale register (RTPSR)	R/W	0x0000_0000	<a href="#">5.5.5.3/5-40</a>
0x0C	Real time counter register (RTCTR)	R	0x0000_0000	<a href="#">5.5.5.4/5-40</a>
0x10	Real time counter event register (RTEVR)	w1c	0x0000_0000	<a href="#">5.5.5.5/5-41</a>
0x14	Real time counter alarm register (RTALR)	R/W	0xFFFF_FFFF	<a href="#">5.5.5.6/5-41</a>
0x18–0x1F	Reserved	—	—	—
<b>Periodic Interval Timer (PIT) Registers</b>				
0x00	Periodic interval timer control register (PTCNR)	R/W	0x0000_0000	<a href="#">5.6.5.1/5-45</a>
0x04	Periodic interval timer load register (PTLDR)	R/W	0x0000_0000	<a href="#">5.6.5.2/5-46</a>
0x08	Periodic interval timer prescale register (PTPSR)	R/W	0x0000_0000	<a href="#">5.6.5.3/5-47</a>
0x0C	Periodic interval timer counter register (PTCTR)	R	0x0000_0000	<a href="#">5.6.5.4/5-47</a>
0x10	Periodic interval timer event register (PTEVR)	w1c	0x0000_0000	<a href="#">5.6.5.5/5-47</a>
0x14–0x1F	Reserved	—	—	—
<b>Global Timers Module 1</b>				
0x00	Timer 1 and 2 global timers configuration register (GTCFR1)	R/W	0x00	<a href="#">5.7.5.1/5-55</a>
0x01–0x03	Reserved	—	—	—

Table 2-2. Memory Map (continued)

Offset	Register	Access	Reset	Section/Page
0x04	Timer 3 and 4 global timers configuration register (GTCFR2)	R/W	0x00	<a href="#">5.7.5.1/5-55</a>
0x05–0x0F	Reserved	—	—	—
0x10	Timer 1 global timers mode register (GTMDR1)	R/W	0x0000	<a href="#">5.7.5.2/5-58</a>
0x12	Timer 2 global timers mode register (GTMDR2)			
0x14	Timer 1 global timers reference register (GTRFR1)	R/W	0xFFFF	<a href="#">5.7.5.3/5-59</a>
0x16	Timer 2 global timers reference register (GTRFR2)			
0x18	Timer 1 global timers capture register (GTCPR1)	R/W	0x0000	<a href="#">5.7.5.4/5-59</a>
0x1A	Timer 2 global timers capture register (GTCPR2)			
0x1C	Timer 1 global timers counter register (GTCNR1)	R/W	0x0000	<a href="#">5.7.5.5/5-60</a>
0x1E	Timer 2 global timers counter register (GTCNR2)			
0x20	Timer 3 global timers mode register (GTMDR3)	R/W	0x0000	<a href="#">5.7.5.2/5-58</a>
0x22	Timer 4 global timers mode register (GTMDR4)			
0x24	Timer 3 global timers reference register (GTRFR3)	R/W	0xFFFF	<a href="#">5.7.5.3/5-59</a>
0x26	Timer 4 global timers reference register (GTRFR4)			
0x28	Timer 3 global timers capture register (GTCPR3)	R	0x0000	<a href="#">5.7.5.4/5-59</a>
0x2A	Timer 4 global timers capture register (GTCPR4)			
0x2C	Timer 3 global timers counter register (GTCNR3)	R/W	0x0000	<a href="#">5.7.5.5/5-60</a>
0x2E	Timer 4 global timers counter register (GTCNR4)			
0x30	Timer 1 global timers event register (GTEVR1)	w1c	0x0000	<a href="#">5.7.5.6/5-60</a>
0x32	Timer 2 global timers event register (GTEVR2)			
0x34	Timer 3 global timers event register (GTEVR3)			
0x36	Timer 4 global timers event register (GTEVR4)			
0x38	Timer 1 global timers prescale register (GTPSR1)	R/W	0x0003	<a href="#">5.7.5.7/5-61</a>
0x3A	Timer 2 global timers prescale register (GTPSR2)			
0x3C	Timer 3 global timers prescale register (GTPSR3)			
0x3E	Timer 4 global timers prescale register (GTPSR4)			
General Purpose (Global) Timer Module 2: All registers defined for GTM1 are also defined for GTM2; the base address of GTM2 registers is 0x0_06nn.				
<b>Integrated Programmable Interrupt Controller (IPIC)</b>				
0x00	System global interrupt configuration register (SICFR)	R/W	0x0000_0000	<a href="#">8.5.1/8-7</a>
0x04	System regular interrupt vector register (SIVCR)	R	0x0000_0000	<a href="#">8.5.2/8-9</a>
0x08	System internal interrupt pending register (SIPNR_H)	R	0x0000_0000	<a href="#">8.5.3/8-11</a>
0x0C	System internal interrupt pending register (SIPNR_L)	R	0x0000_0000	<a href="#">8.5.3/8-11</a>
0x10	System internal interrupt group A priority register (SIPRR_A)	R/W	0x0530_9770	<a href="#">8.5.4/8-13</a>

Table 2-2. Memory Map (continued)

Offset	Register	Access	Reset	Section/Page
0x1C	System internal interrupt group D priority register (SIPRR_D)	R/W	0x0530_9770	8.5.5/8-14
0x20	System internal interrupt mask register (SIMSR_H)	R/W	0x0000_0000	8.5.6/8-15
0x24	System internal interrupt mask register (SIMSR_L)	R/W	0x0000_0000	8.5.6/8-15
0x28	System internal interrupt control register (SICNR)	R/W	0x0000_0000	8.5.7/8-16
0x2C	System external interrupt pending register (SEPNR)	R/W	Special	8.5.8/8-18
0x30	System mixed interrupt group A priority register (SMPRR_A)	R/W	0x0530_9770	8.5.9/8-18
0x34	System mixed interrupt group B priority register (SMPRR_B)	R/W	0x0530_9770	8.5.10/8-19
0x38	System external interrupt mask register (SEMSR)	R/W	0x0000_0000	8.5.11/8-20
0x3C	System external interrupt control register (SECNR)	R/W	0x0000_0000	8.5.12/8-21
0x40	System error status register (SERSR)	R/W	0x0000_0000	8.5.13/8-22
0x44	System error mask register (SERMR)	R/W		8.5.14/8-23
0x48	System error control register (SERCR)	R/W	0x0000_0000	8.5.15/8-24
0x4C–0x4F	Reserved	—	—	—
0x50	System internal interrupt force register (SIFCR_H)	R/W	0x0000_0000	8.5.16/8-25
0x54	System internal interrupt force register (SIFCR_L)	R/W	0x0000_0000	8.5.16/8-25
0x58	System external interrupt force register (SEFCR)	R/W	0x0000_0000	8.5.17/8-26
0x5C	System error force register (SERFR)	R/W	0x0000_0000	8.5.18/8-26
0x60	System critical interrupt vector register (SCVCR)	R	0x0000_0000	8.5.19/8-27
0x64	System management interrupt vector register (SMVCR)	R	0x0000_0000	8.5.20/8-27
0x68–0xFF	Reserved	—	—	—
<b>System Arbiter Registers</b>				
0x00	Arbiter configuration register (ACR)	R/W	0x0000_0000/ 0x0010_0000 <sup>10</sup>	6.2.1/6-2
0x04	Arbiter timers register (ATR)	R/W	0xFFFF_FFFF	6.2.2/6-4
0x0C	Arbiter event register (AER)	w1c	0x0000_0000	6.2.3/6-5
0x10	Arbiter interrupt definition register (AIDR)	R/W	0x0000_0000	6.2.4/6-6
0x14	Arbiter mask register (AMR)	R/W	0x0000_0000	6.2.5/6-7
0x18	Arbiter event attributes register (AEATR)	R	0x0000_0000 <sup>11</sup>	6.2.6/6-7
0x1C	Arbiter event address register (AEADR)	R	0x0000_0000 <sup>11</sup>	6.2.7/6-9
0x20	Arbiter event response register (AERR)	R/W	0x0000_0000	6.2.8/6-10
<b>Reset Module</b>				
0x0_0900	Reset configuration word low register (RCWLR)	R	0x0000_0000	4.5.1.1/4-33
0x0_0904	Reset configuration word high register (RCWHR)	R	0x0000_0000	4.5.1.2/4-33



Table 2-2. Memory Map (continued)

Offset	Register	Access	Reset	Section/Page
0x0_0908	Reserved, should be cleared	—	—	—
0x0_090C	Reserved, should be cleared	—	—	—
0x0_0910	Reset status register (RSR)	R/W	0x0000_0000	<a href="#">4.5.1.3/4-33</a>
0x0_0914	Reset mode register (RMR)	R/W	0x0000_0000	<a href="#">4.5.1.4/4-35</a>
0x0_0918	Reset protection register (RPR)	R/W	0x0000_0000	<a href="#">4.5.1.5/4-35</a>
0x0_091C	Reset control register (RCR)	R/W	0x0000_0000	<a href="#">4.5.1.6/4-36</a>
0x0_0920	Reset control enable register (RCER)	R/W	0x0000_0000	<a href="#">4.5.1.7/4-37</a>
0x0_0924– 0x0_09FC	Reserved, should be cleared.	—	—	—
<b>Clock Module</b>				
0x0_0A00	System PLL mode register (SPMR)	R	0xn <sub>nnn</sub> _n <sub>nnn</sub>	<a href="#">4.5.2.1/4-37</a>
0x0_0A04	Output clock control register (OCCR)	R/W	0x0000_80C0	<a href="#">4.5.2.2/4-39</a>
0x0_0A08	System clock control register (SCCR)	R/W	0x7DDF_FFFF	<a href="#">4.5.2.3/4-40</a>
0x0_0A0C– 0x0_0AFC	Reserved, should be cleared	—	—	—
<b>Power Management Control Module</b>				
0x00B00	Power management controller configuration register (PMCCR)	R/W	0x0000_0000	<a href="#">5.8.2.1/5-67</a>
0x00B04	Power management controller event register (PM CER)	R/W	0x0000_0000	<a href="#">5.8.2.2/5-68</a>
0x00B08	Power management controller mask register (PMCMR)	R/W	0x0000_0000	<a href="#">5.8.2.3/5-70</a>
0x00B0C	Power management controller configuration register 1 (PMCCR1)	R/W	0x0000_0000	<a href="#">5.8.2.4/5-70</a>
0x00B10	Power management controller configuration register 2 (PMCCR2)	R/W	0x0002_0002	<a href="#">5.8.2.5/5-72</a>
0x00B14– 0x00BFC	Reserved	—	—	—
<b>GPIO Registers</b>				
0xC00	GPIO direction register (GPDIR)	R/W	0x0000_0000	<a href="#">21.3.1/21-3</a>
0xC04	GPIO open drain register (GPODR)	R/W	0x0000_0000	<a href="#">8.6/8-28</a>
0xC08	GPIO data register (GPDAT)	R/W	0x0000_0000	<a href="#">21.3.3/21-4</a>
0xC0C	GPIO interrupt event register (GPIER)	w1c	Undefined	<a href="#">21.3.4/21-4</a>
0xC10	GPIO interrupt mask register (GPIMR)	R/W	0x0000_0000	<a href="#">21.3.5/21-4</a>
0xC14	GPIO external interrupt control register (GPICR)	R/W	0x0000_0000	<a href="#">21.3.6/21-5</a>
<b>DDR Memory Controller Memory Map</b>				
0x000	CS0_BNDS—Chip select 0 memory bounds	R/W	0x0000_0000	<a href="#">9.4.1.1/9-9</a>
0x008	CS1_BNDS—Chip select 1 memory bounds	R/W	0x0000_0000	<a href="#">9.4.1.1/9-9</a>
0x080	CS0_CONFIG—Chip select 0 configuration	R/W	0x0000_0000	<a href="#">9.4.1.2/9-10</a>

Table 2-2. Memory Map (continued)

Offset	Register	Access	Reset	Section/Page
0x084	CS1_CONFIG—Chip select 1 configuration	R/W	0x0000_0000	9.4.1.2/9-10
0x100	TIMING_CFG_3—DDR SDRAM timing configuration 3	R/W	0x0000_0000	9.4.1.3/9-11
0x104	TIMING_CFG_0—DDR SDRAM timing configuration 0	R/W	0x0011_0105	9.4.1.4/9-12
0x108	TIMING_CFG_1—DDR SDRAM timing configuration 1	R/W	0x0000_0000	9.4.1.5/9-14
0x10C	TIMING_CFG_2—DDR SDRAM timing configuration 2	R/W	0x0000_0000	9.4.1.6/9-16
0x110	DDR_SDRAM_CFG—DDR SDRAM control configuration	R/W	0x0200_0000	9.4.1.7/9-18
0x114	DDR_SDRAM_CFG_2—DDR SDRAM control configuration 2	R/W	0x0000_0000	9.4.1.8/9-21
0x118	DDR_SDRAM_MODE—DDR SDRAM mode configuration	R/W	0x0000_0000	9.4.1.9/9-22
0x11C	DDR_SDRAM_MODE_2—DDR SDRAM mode configuration 2	R/W	0x0000_0000	9.4.1.10/9-23
0x120	DDR_SDRAM_MD_CNTL—DDR SDRAM mode control	R/W	0x0000_0000	9.4.1.11/9-24
0x124	DDR_SDRAM_INTERVAL—DDR SDRAM interval configuration	R/W	0x0000_0000	9.4.1.12/9-26
0x128	DDR_DATA_INIT—DDR SDRAM data initialization	R/W	0x0000_0000	9.4.1.13/9-27
0x130	DDR_SDRAM_CLK_CNTL—DDR SDRAM clock control	R/W	0x0200_0000	9.4.1.14/9-27
0x140– 0x144	Reserved	—	—	—
0x148	DDR_INIT_ADDR—DDR training initialization address	R/W	0x0000_0000	9.4.1.15/9-28
0x150– 0xBF4	Reserved	—	—	—
0xBF8	DDR_IP_REV1—DDR IP block revision 1	R	0xn <sup>nnn</sup> _n <sup>nnn</sup> <sup>12</sup>	9.4.1.16/9-28
0xBFC	DDR_IP_REV2—DDR IP block revision 2	R	0x00n <sup>n</sup> _00n <sup>n</sup> <sup>12</sup>	9.4.1.17/9-29
<b>I<sup>2</sup>C1 Controller</b>				
0x0_3000	I2C1ADR—I <sup>2</sup> C1 address register	R/W	0x00	17.3.1.1/17-5
0x0_3004	I2C1FDR—I <sup>2</sup> C1 frequency divider register	R/W	0x00	17.3.1.2/17-6
0x0_3008	I2C1CR—I <sup>2</sup> C1 control register	R/W	0x00	17.3.1.3/17-7
0x0_300C	I2C1SR—I <sup>2</sup> C1 status register	R/W	0x81	17.3.1.4/17-8
0x0_3010	I2C1DR—I <sup>2</sup> C1 data register	R/W	0x00	17.3.1.5/17-9
0x0_3014	I2C1DFSRR—I <sup>2</sup> C1 digital filter sampling rate register	R/W	0x10	17.3.1.6/17-9
0x0_301C– 0x0_30FF	Reserved, should be cleared	—	—	—
<b>I<sup>2</sup>C2 Controller</b>				
0x0_3100	I2C2ADR—I <sup>2</sup> C2 address register	R/W	0x00	17.3.1.1/17-5
0x0_3104	I2C2FDR—I <sup>2</sup> C2 frequency divider register	R/W	0x00	17.3.1.2/17-6
0x0_3108	I2C2CR—I <sup>2</sup> C2 control register	R/W	0x00	17.3.1.3/17-7
0x0_310C	I2C2SR—I <sup>2</sup> C2 status register	R/W	0x81	17.3.1.4/17-8

Table 2-2. Memory Map (continued)

Offset	Register	Access	Reset	Section/Page
0x0_3110	I2C2DR—I <sup>2</sup> C2 data register	R/W	0x00	<a href="#">17.3.1.5/17-9</a>
0x0_3114	I2C2DFSRR—I <sup>2</sup> C2 digital filter sampling rate register	R/W	0x10	<a href="#">17.3.1.6/17-9</a>
0x0_311C– 0x0_31FF	Reserved, should be cleared	—	—	—
<b>DUART</b>				
0x0_4500	URBR—ULCR[DLAB] = 0 UART1 receiver buffer register	R	0x00	<a href="#">18.3.1.1/18-5</a>
	UTHR—ULCR[DLAB] = 0 UART1 transmitter holding register	W	0x00	<a href="#">18.3.1.2/18-6</a>
	UDLB—ULCR[DLAB] = 1 UART1 divisor least significant byte register	R/W	0x00	<a href="#">18.3.1.3/18-6</a>
0x0_4501	UIER—ULCR[DLAB] = 0 UART1 interrupt enable register	R/W	0x00	<a href="#">18.3.1.4/18-8</a>
	UDMB—ULCR[DLAB] = 1 UART1 divisor most significant byte register	R/W	0x00	<a href="#">18.3.1.3/18-6</a>
0x0_4502	UIIR—ULCR[DLAB] = 0 UART1 interrupt ID register	R	0x01	<a href="#">18.3.1.5/18-9</a>
	UFCR—ULCR[DLAB] = 0 UART1 FIFO control register	W	0x00	<a href="#">18.3.1.6/18-10</a>
	UAFR—ULCR[DLAB] = 1 UART1 alternate function register	R/W	0x00	<a href="#">18.3.1.12/18-16</a>
0x0_4503	ULCR—ULCR[DLAB] = x UART1 line control register	R/W	0x00	<a href="#">18.3.1.7/18-11</a>
0x0_4504	UMCR—ULCR[DLAB] = x UART1 MODEM control register	R/W	0x00	<a href="#">18.3.1.8/18-13</a>
0x0_4505	ULSR—ULCR[DLAB] = x UART1 line status register	R	0x60	<a href="#">18.3.1.9/18-14</a>
0x0_4506	UMSR—ULCR[DLAB] = x UART1 MODEM status register	R	0x00	<a href="#">18.3.1.10/18-15</a>
0x0_4507	USCR—ULCR[DLAB] = x UART1 scratch register	R/W	0x00	<a href="#">18.3.1.11/18-16</a>
0x0_4510	UDSR—ULCR[DLAB] = x UART1 DMA status register	R	0x01	<a href="#">18.3.1.13/18-17</a>
0x0_4600	URBR—ULCR[DLAB] = 0 UART2 receiver buffer register	R	0x00	<a href="#">18.3.1.1/18-5</a>
	UTHR—ULCR[DLAB] = 0 UART2 transmitter holding register	W	0x00	<a href="#">18.3.1.2/18-6</a>
	UDLB—ULCR[DLAB] = 1 UART2 divisor least significant byte register	R/W	0x00	<a href="#">18.3.1.3/18-6</a>
0x0_4601	UIER—ULCR[DLAB] = 0 UART2 interrupt enable register	R/W	0x00	<a href="#">18.3.1.4/18-8</a>
	UDMB—ULCR[DLAB] = 1 UART2 divisor most significant byte register	R/W	0x00	<a href="#">18.3.1.3/18-6</a>
0x0_4602	UIIR—ULCR[DLAB] = 0 UART2 interrupt ID register	R	0x01	<a href="#">18.3.1.5/18-9</a>
	UFCR—ULCR[DLAB] = 0 UART2 FIFO control register	W	0x00	<a href="#">18.3.1.6/18-10</a>
	UAFR—ULCR[DLAB] = 1 UART2 alternate function register	R/W	0x00	<a href="#">18.3.1.12/18-16</a>
0x0_4603	ULCR—ULCR[DLAB] = x UART2 line control register	R/W	0x00	<a href="#">18.3.1.7/18-11</a>
0x0_4604	UMCR—ULCR[DLAB] = x UART2 MODEM control register	R/W	0x00	<a href="#">18.3.1.8/18-13</a>
0x0_4605	ULSR—ULCR[DLAB] = x UART2 line status register	R	0x60	<a href="#">18.3.1.9/18-14</a>

Table 2-2. Memory Map (continued)

Offset	Register	Access	Reset	Section/Page
0x0_4606	UMSR—ULCR[DLAB] = x UART2 MODEM status register	R	0x00	18.3.1.10/18-15
0x0_4607	USCR—ULCR[DLAB] = x UART2 scratch register	R/W	0x00	18.3.1.11/18-16
0x0_4610	UDSR—ULCR[DLAB] = x UART2 DMA status register	R	0x01	18.3.1.13/18-17
<b>Enhanced Local Bus Controller (eLBC) Registers</b>				
0x000	BR0—Base register 0	R/W	0x0000_nnnn	10.3.1.1/10-10
0x008	BR1—Base register 1	R/W	0x0000_0000	10.3.1.1/10-10
0x010	BR2—Base register 2	R/W	0x0000_0000	10.3.1.1/10-10
0x018	BR3—Base register 3	R/W	0x0000_0000	10.3.1.1/10-10
0x020–0x038	Reserved	R/W	0x0000_0000	10.3.1.1/10-10
0x004	OR0—Options register 0	R/W	0x0000_0FF7	10.3.1.2/10-11
0x00C	OR1—Options register 1	R/W	0x0000_0000	10.3.1.2/10-11
0x014	OR2—Options register 2	R/W	0x0000_0000	10.3.1.2/10-11
0x01C	OR3—Options register 3	R/W	0x0000_0000	10.3.1.2/10-11
0x024–0x064	Reserved	—	—	—
0x068	MAR—UPM address register	R/W	0x0000_0000	10.3.1.3/10-19
0x06C	Reserved	—	—	—
0x070	MAMR—UPMA mode register	R/W	0x0000_0000	10.3.1.4/10-20
0x074	MBMR—UPMB mode register	R/W	0x0000_0000	10.3.1.4/10-20
0x078	MCMR—UPMC mode register	R/W	0x0000_0000	10.3.1.4/10-20
0x07C–0x080	Reserved	—	—	—
0x084	MRTPR—Memory refresh timer prescaler register	R/W	0x0000_0000	10.3.1.5/10-22
0x088	MDR—UPM/FCM data register	R/W	0x0000_0000	10.3.1.6/10-22
0x08C	Reserved	—	—	—
0x090	LSOR—Special operation initiation register	R/W	0x0000_0000	10.3.1.7/10-23
0x094–0x09C	Reserved	—	—	—
0x0A0	LURT—UPM refresh timer	R/W	0x0000_0000	10.3.1.4/10-20
0x0A4–0x0AC	Reserved	—	—	—
0x0B0	LTESR—Transfer error status register	w1c	0x0000_0000	10.3.1.9/10-25
0x0B4	LTEDR—Transfer error disable register	R/W	0x0000_0000	10.3.1.10/10-27
0x0B8	LTEIR—Transfer error interrupt register	R/W	0x0000_0000	10.3.1.11/10-28

Table 2-2. Memory Map (continued)

Offset	Register	Access	Reset	Section/Page
0x0BC	LTEATR—Transfer error attributes register	R/W	0x0000_0000	<a href="#">10.3.1.12/10-29</a>
0x0C0	LTEAR—Transfer error address register	R/W	0x0000_0000	<a href="#">10.3.1.13/10-30</a>
0x0C4	LTECCR—Transfer error ECC register	w1c	0x0000_0000	<a href="#">10.3.1.14/10-30</a>
0x0C8– 0x0CC	Reserved	—	—	—
0x0D0	LBCR—Configuration register	R/W	0x0004_0000	<a href="#">10.3.1.15/10-31</a>
0x0D4	LCRR—Clock ratio register	R/W	0x8000_0008	<a href="#">10.3.1.16/10-33</a>
0x0D8– 0x0DC	Reserved	—	—	—
0x0E0	FMR—Flash mode register	R/W	0x0000_0n00	<a href="#">10.3.1.17/10-34</a>
0x0E4	FIR—Flash instruction register	R/W	0x0000_0000	<a href="#">10.3.1.18/10-35</a>
0x0E8	FCR—Flash command register	R/W	0x0000_0000	<a href="#">10.3.1.19/10-36</a>
0x0EC	FBAR—Flash block address register	R/W	0x0000_0000	<a href="#">10.3.1.20/10-37</a>
0x0F0	FPAR—Flash page address register	R/W	0x0000_0000	<a href="#">10.3.1.21/10-37</a>
0x0F4	FBCR—Flash byte count register	R/W	0x0000_0000	<a href="#">10.3.1.22/10-39</a>
0x0F8– 0x0FC	Reserved	—	—	—
0x100	FECC0—Flash ECC block 0 register	R	0x0000_0000	<a href="#">10.3.1.23/10-39</a>
0x104	FECC1—Flash ECC block 1 register	R	0x0000_0000	<a href="#">10.3.1.23/10-39</a>
0x108	FECC2—Flash ECC block 2 register	R	0x0000_0000	<a href="#">10.3.1.23/10-39</a>
0x10C	FECC3—Flash ECC block 3 register	R	0x0000_0000	<a href="#">10.3.1.23/10-39</a>
<b>Serial Peripheral Interface (SPI)</b>				
0x000– 0x01F	Reserved	—	—	—
0x020	SPI mode register (SPMODE)	R/W	0x0000_0000	<a href="#">Figure 19-4./19-9</a>
0x024	SPI event register (SPIE)	Mixed	0x0000_0000	<a href="#">Figure 19-7./19-12</a>
0x028	SPI mask register (SPIM)	R/W	0x0000_0000	<a href="#">Figure 19-8./19-13</a>
0x02C	SPI command register (SPCOM)	W	0x0000_0000	<a href="#">Figure 19-9./19-14</a>
0x030	SPI transmit register (SPITD)	W	0x0000_0000	<a href="#">19.4.1.5/19-14</a>
0x034	SPI receive register (SPIRD)	R	0xFFFF_FFFF	<a href="#">Figure 19-11./19-15</a>
0x038– 0xFFFF	Reserved	—	—	—
<b>DMA Registers</b>				
0x0_8030	OMISR—Outbound message interrupt status register	Mixed	0x0000_0000	<a href="#">12.3.1/12-3</a>
0x0_8034	OMIMR—Outbound message interrupt mask register	R/W	0x0000_0000	<a href="#">12.3.2/12-4</a>

Table 2-2. Memory Map (continued)

Offset	Register	Access	Reset	Section/Page
0x0_8050	IMR0—Inbound message register 0	R/W	0x0000_0000	<a href="#">12.3.3/12-5</a>
0x0_8054	IMR1—Inbound message register 1	R/W	0x0000_0000	<a href="#">12.3.3/12-5</a>
0x0_8058	OMR0—Outbound message register 0	R/W	0x0000_0000	<a href="#">12.3.4/12-5</a>
0x0_805C	OMR1—Outbound message register 1	R/W	0x0000_0000	<a href="#">12.3.4/12-5</a>
0x0_8060	ODR—Outbound doorbell register	R/W	0x0000_0000	<a href="#">12.3.5/12-6</a>
0x0_8068	IDR—Inbound doorbell register	R/W	0x0000_0000	<a href="#">12.3.5/12-6</a>
0x0_8080	IMISR—Inbound message interrupt status register	Mixed	0x0000_0000	<a href="#">12.3.6/12-7</a>
0x0_8084	IMIMR—Inbound message interrupt mask register	R/W	0x0000_0000	<a href="#">12.3.7/12-8</a>
0x0_8100	DMAMR0—DMA 0 mode register	R/W	0x0000_0000	<a href="#">12.3.8.1/12-9</a>
0x0_8104	DMASR0—DMA 0 status register	R/W	0x0000_0000	<a href="#">12.3.8.2/12-11</a>
0x0_8108	DMACDAR0—DMA 0 current descriptor address register	R/W	0x0000_0000	<a href="#">12.3.8.3/12-12</a>
0x0_8110	DMASAR0—DMA 0 source address register	R/W	0x0000_0000	<a href="#">12.3.8.4/12-13</a>
0x0_8118	DMADAR0—DMA 0 destination address register	R/W	0x0000_0000	<a href="#">12.3.8.5/12-13</a>
0x0_8120	DMABCR0—DMA 0 byte count register	R/W	0x0000_0000	<a href="#">12.3.8.6/12-14</a>
0x0_8124	DMANDAR0—DMA 0 next descriptor address register	R/W	0x0000_0000	<a href="#">12.3.8.7/12-14</a>
0x0_8180	DMAMR1—DMA 1 mode register	R/W	0x0000_0000	<a href="#">12.3.8.1/12-9</a>
0x0_8184	DMASR1—DMA 1 status register	R/W	0x0000_0000	<a href="#">12.3.8.2/12-11</a>
0x0_8188	DMACDAR1—DMA 1 current descriptor address register	R/W	0x0000_0000	<a href="#">12.3.8.3/12-12</a>
0x0_8190	DMASAR1—DMA 1 source address register	R/W	0x0000_0000	<a href="#">12.3.8.4/12-13</a>
0x0_8198	DMADAR1—DMA 1 destination address register	R/W	0x0000_0000	<a href="#">12.3.8.5/12-13</a>
0x0_81A0	DMABCR1—DMA 1 byte count register	R/W	0x0000_0000	<a href="#">12.3.8.6/12-14</a>
0x0_81A4	DMANDAR1—DMA 1 next descriptor address register	R/W	0x0000_0000	<a href="#">12.3.8.7/12-14</a>
0x0_8200	DMAMR2—DMA 2 mode register	R/W	0x0000_0000	<a href="#">12.3.8.1/12-9</a>
0x0_8204	DMASR2—DMA 2 status register	R/W	0x0000_0000	<a href="#">12.3.8.2/12-11</a>
0x0_8208	DMACDAR2—DMA 2 current descriptor address register	R/W	0x0000_0000	<a href="#">12.3.8.3/12-12</a>
0x0_8210	DMASAR2—DMA 2 source address register	R/W	0x0000_0000	<a href="#">12.3.8.4/12-13</a>
0x0_8218	DMADAR2—DMA 2 destination address register	R/W	0x0000_0000	<a href="#">12.3.8.5/12-13</a>
0x0_8220	DMABCR2—DMA 2 byte count register	R/W	0x0000_0000	<a href="#">12.3.8.6/12-14</a>
0x0_8224	DMANDAR2—DMA 2 next descriptor address register	R/W	0x0000_0000	<a href="#">12.3.8.7/12-14</a>
0x0_8280	DMAMR3—DMA 3 mode register	R/W	0x0000_0000	<a href="#">12.3.8.1/12-9</a>
0x0_8284	DMASR3—DMA 3 status register	R/W	0x0000_0000	<a href="#">12.3.8.2/12-11</a>
0x0_8288	DMACDAR3—DMA 3 current descriptor address register	R/W	0x0000_0000	<a href="#">12.3.8.3/12-12</a>
0x0_8290	DMASAR3—DMA 3 source address register	R/W	0x0000_0000	<a href="#">12.3.8.4/12-13</a>

Table 2-2. Memory Map (continued)

Offset	Register	Access	Reset	Section/Page
0x0_8298	DMADAR3—DMA 3 destination address register	R/W	0x0000_0000	12.3.8.5/12-13
0x0_82A0	DMABCR3—DMA 3 byte count register	R/W	0x0000_0000	12.3.8.6/12-14
0x0_82A4	DMANDAR3—DMA 3 next descriptor address register	R/W	0x0000_0000	12.3.8.7/12-14
0x0_82A8	DMAGSR—DMA general status register	R	0x0000_0000	12.3.8.8/12-15
0x0_82B0– 0x0_82FF	Reserved	—	—	—
<b>PCI Software Configuration Registers</b>				
0x0	PCI_CONFIG_ADDRESS	W	0x0000_0000	13.3.1.1/13-13
0x4	PCI_CONFIG_DATA	R/W	0x0000_0000	13.3.1.2/13-14
0x8	PCI_INT_ACK	R	0x0000_0000	13.3.1.3/13-15
0x0_830C– 0x0_837F	Reserved, should be cleared	—	—	—
0x0_8380	PCIPMR0—PCI power management register 0	R	0x7E4B_0001	13.3.3.27/13-41
0x0_8384	PCIPMR1—PCI power management register 1	R/W	0x00n0_0000	13.3.3.28/13-42
0x0_8388– 0x0_83FF	Reserved	—	—	—
<b>Sequencer (IOS)</b>				
0x00	POTAR0—PCI outbound translation address register 0	R/W	0x0000_0000	11.4.1/11-3
0x08	POBAR0—PCI outbound base address register 0	R/W	0x0000_0000	11.4.2/11-3
0x10	POCMR0—PCI outbound comparison mask register 0	R/W	0x0000_0000	11.4.3/11-4
0x18	POTAR1—PCI outbound translation address register 1	R/W	0x0000_0000	11.4.1/11-3
0x20	POBAR1—PCI outbound base address register 1	R/W	0x0000_0000	11.4.2/11-3
0x28	POCMR1—PCI outbound comparison mask register 1	R/W	0x0000_0000	11.4.3/11-4
0x30	POTAR2—PCI outbound translation address register 2	R/W	0x0000_0000	11.4.1/11-3
0x38	POBAR2—PCI outbound base address register 2	R/W	0x0000_0000	11.4.2/11-3
0x40	POCMR2—PCI outbound comparison mask register 2	R/W	0x0000_0000	11.4.3/11-4
0x48	POTAR3—PCI outbound translation address register 3	R/W	0x0000_0000	11.4.1/11-3
0x50	POBAR3—PCI outbound base address register 3	R/W	0x0000_0000	11.4.2/11-3
0x58	POCMR3—PCI outbound comparison mask register 3	R/W	0x0000_0000	11.4.3/11-4
0x60	POTAR4—PCI outbound translation address register 4	R/W	0x0000_0000	11.4.1/11-3
0x68	POBAR4—PCI outbound base address register 4	R/W	0x0000_0000	11.4.2/11-3
0x70	POCMR4—PCI outbound comparison mask register 4	R/W	0x0000_0000	11.4.3/11-4
0x78	POTAR5—PCI outbound translation address register 5	R/W	0x0000_0000	11.4.1/11-3
0x80	POBAR5—PCI outbound base address register 5	R/W	0x0000_0000	11.4.2/11-3

Table 2-2. Memory Map (continued)

Offset	Register	Access	Reset	Section/Page
0x88	POCMR5—PCI outbound comparison mask register 5	R/W	0x0000_0000	11.4.3/11-4
0xF0	PMCR—Power management control register	R/W	0x0000_0000	11.4.4/11-5
0xF8	DTCR—Discard timer control register	R/W	0x0000_0000	11.4.5/11-6
<b>PCI Error Management Registers</b>				
0x00	PCI error status register (PCI_ESR)	w1c	0x0000_0000	13.3.2.1/13-15
0x04	PCI error capture disable register (PCI_ECDR)	R/W	0x0000_0000	13.3.2.2/13-16
0x08	PCI error enable register (PCI_EER)	R/W	0x0000_0000	13.3.2.3/13-17
0x0C	PCI error attributes capture register (PCI_EATCR)	R/W	0x0000_0000	13.3.2.4/13-18
0x10	PCI error address capture register (PCI_EACR)	R	0x0000_0000	13.3.2.5/13-19
0x14	PCI error extended address capture register (PCI_EEACR)	R	0x0000_0000	13.3.2.6/13-20
0x18	PCI error data capture register (PCI_EDCR)	R/W	0x0000_0000	13.3.2.7/13-20
<b>PCI Control and Status Registers</b>				
0x20	PCI general control register (PCI_GCR)	R/W	0x0000_0000	13.3.2.8/13-20
0x24	PCI error control register (PCI_ECR)	R/W	0x0000_0000	13.3.2.9/13-21
0x28	PCI general status register (PCI_GSR)	R	0x0000_0000	13.3.2.10/13-22
<b>PCI Inbound ATU Registers</b>				
0x38	PCI inbound translation address register 2 (PITAR2)	R/W	0x0000_0000	13.3.2.11/13-22
0x3C	Reserved	—	—	—
0x40	PCI inbound base address register 2 (PIBAR2)	R/W	0x0000_0000	13.3.2.12/13-23
0x44	PCI inbound extended base address register 2 (PIEBAR2)	R/W	0x0000_0000	13.3.2.13/13-24
0x48	PCI inbound window attributes register 2 (PIWAR2)	R/W	0x0000_0000	13.3.2.14/13-24
0x50	PCI inbound translation address register 1 (PITAR1)	R/W	0x0000_0000	13.3.2.11/13-22
0x54	Reserved	—	—	—
0x58	PCI inbound base address register 1 (PIBAR1)	R/W	0x0000_0000	13.3.2.12/13-23
0x5C	PCI inbound extended base address register 1 (PIEBAR1)	R/W	0x0000_0000	13.3.2.13/13-24
0x60	PCI inbound window attributes register 1 (PIWAR1)	R/W	0x0000_0000	13.3.2.14/13-24
0x68	PCI inbound translation address register 0 (PITAR0)	R/W	0x0000_0000	13.3.2.11/13-22
0x6C	Reserved	—	—	—
0x70	PCI inbound base address register 0 (PIBAR0)	R/W	0x0000_0000	13.3.2.12/13-23
0x78	PCI inbound window attributes register 0 (PIWAR0)	R/W	0x0000_0000	13.3.2.13/13-24
0x7C– 0xFF	Reserved	—	—	—
<b>USB DR Controller Registers</b>				
0x2_3000– 0x2_30FF	Reserved, should be cleared	—	—	—



Table 2-2. Memory Map (continued)

Offset	Register	Access	Reset	Section/Page
0x2_3100	CAPLENGTH—Capability register length	R	0x40	16.3.1.1/16-8
0x2_3102	HCIVERSION—Host interface version number <sup>1</sup>	R	0x0100	16.3.1.2/16-8
0x2_3104	HCSPARAMS—Host ctrl. structural parameters <sup>1</sup>	R	0x0001_0011	16.3.1.3/16-9
0x2_3108	HCCPARAMS—Host ctrl. capability parameters <sup>1</sup>	R	0x0000_0006	16.3.1.4/16-10
0x2_3120	DCIVERSION—Device interface version number	R	0x0001	16.3.1.5/16-10
0x2_3124	DCCPARAMS—Device controller parameters	R	0x0000_0183	16.3.1.6/16-11
0x2_3140	USBCMD—USB command	Mixed	0x0008_0000	16.3.2.1/16-12
0x2_3144	USBSTS—USB status	Mixed	0x0000_0000	16.3.2.2/16-14
0x2_3148	USBINTR—USB interrupt enable	R/W	0x0000_0000	16.3.2.3/16-17
0x2_314C	FRINDEX—USB frame index	R/W	0x0000_ <i>nnnn</i>	16.3.2.4/16-18
0x2_3154	PERIODICLISTBASE—Frame list base address <sup>13</sup>	R/W	0x <i>nnnn</i> _0000	16.3.2.6/16-19
	DEVICEADDR—USB device address	R/W	0x0000_0000	16.3.2.7/16-20
0x2_3158	ASYNCLISTADDR—Next asynchronous list addr (host mode) <sup>13</sup>	R/W	0x0000_0000	16.3.2.8/16-20
	ENDPOINTLISTADDR—Address at endpoint list (device mode)	R/W	0x0000_0000	16.3.2.9/16-21
0x2_3160	BURSTSIZE—Programmable burst size	R/W	0x0000_1010	16.3.2.10/16-22
0x2_3164	TXFILLTUNING—Host TT transmit pre-buffer packet tuning	R/W	0x0000_0000	16.3.2.11/16-22
0x2_3170	ULPI VIEWPORT—ULPI Register Access	Mixed	0x0000_0000	16.3.2.12/16-24
0x2_3180	CONFIGFLAG—Configured flag register	R	0x0000'_0001	16.3.2.13/16-26
0x2_3184	PORTSC—Port status/control	Mixed	0x1000_0000	16.3.2.14/16-26
0x2_31A4	OTGSC—On-The-Go status and control <sup>1</sup>	Mixed	0x0000_0C20	16.3.2.15/16-31
0x2_31A8	USBMODE—USB device mode	R/W	0x0000_0000	16.3.2.16/16-34
0x2_31AC	ENDPTSETUPSTAT—Endpoint setup status	R/W	0x0000_0000	16.3.2.17/16-35
0x2_31B0	ENDPOINTPRIME—Endpoint initialization	R/W	0x0000_0000	16.3.2.18/16-35
0x2_31B4	ENDPTFLUSH—Endpoint de-initialize	R/W	0x0000_0000	16.3.2.19/16-36
0x2_31B8	ENDPTSTATUS—Endpoint status	R	0x0000_0000	16.3.2.20/16-36
0x2_31BC	ENDPTCOMPLETE—Endpoint complete	w1c	0x0000_0000	16.3.2.21/16-37
0x2_31C0	ENDPTCTRL0—Endpoint control 0	Mixed	0x0080_0080	16.3.2.22/16-38
0x2_31C4	ENDPTCTRL1—Endpoint control 1	R/W	0x0000_0000	16.3.2.23/16-39
0x2_31C8	ENDPTCTRL2—Endpoint control 2	R/W	0x0000_0000	16.3.2.23/16-39
0x2_31CA– 0x2_31D4	Reserved	—	—	—
0x2_3400	SNOOP1—Snoop 1	R/W	0x0000_0000	16.3.2.24/16-40
0x2_3404	SNOOP2—Snoop 2	R/W	0x0000_0000	16.3.2.24/16-40

Table 2-2. Memory Map (continued)

Offset	Register	Access	Reset	Section/Page
0x2_3408	AGE_CNT_THRESH—Age count threshold	R/W	0x0000_0000	16.3.2.25/16-41
0x2_340C	PRI_CTRL—Priority control	R/W	0x0000_0000	16.3.2.26/16-42
0x2_3410	SI_CTRL—System interface control	R/W	0x0000_0000	16.3.2.27/16-43
0x2_3500	CONTROL—Control	Mixed	0x0000_0000	16.3.2.28/16-43
0x2_3504– 0x2_3FFF	Reserved, should be cleared	—	—	—
<b>eTSEC General Control and Status Registers</b>				
0x2_4000	TSEC_ID*—Controller ID register	R	0x0124_0106	15.5.3.1.1/15-22
0x2_4004	TSEC_ID2*—Controller ID register	R	0x0030_00F0	15.5.3.1.2/15-23
0x2_4008– 0x2_400C	Reserved	—	—	—
0x2_4010	IEVENT—Interrupt event register	w1c	0x0000_0000	15.5.3.1.3/15-24
0x2_4014	IMASK—Interrupt mask register	R/W	0x0000_0000	15.5.3.1.4/15-27
0x2_4018	EDIS—Error disabled register	R/W	0x0000_0000	15.5.3.1.5/15-29
0x2_401C	Reserved	—	—	—
0x2_4020	ECNTRL—Ethernet control register	R/W	0x0000_0000	15.5.3.1.6/15-31
0x2_4024	Reserved	—	—	—
0x2_4028	PTV—Pause time value register	R/W	0x0000_0000	15.5.3.1.7/15-33
0x2_402C	DMACTRL—DMA control register	R/W	0x0000_0000	15.5.3.1.8/15-34
0x2_4030	Reserved	—	—	—
0x2_4034– 0x2_40FC	Reserved	—	—	—
<b>eTSEC Transmit Control and Status Registers</b>				
0x2_4100	TCTRL—Transmit control register	R/W	0x0000_0000	15.5.3.2.1/15-35
0x2_4104	TSTAT—Transmit status register	w1c	0x0000_0000	15.5.3.2.2/15-37
0x2_4108	DFVLAN*—Default VLAN control word	R/W	0x8100_0000	15.5.3.2.3/15-41
0x2_410C	Reserved	—	—	—
0x2_4110	TXIC—Transmit interrupt coalescing register	R/W	0x0000_0000	15.5.3.2.4/15-42
0x2_4114	TQUEUE*—Transmit queue control register	R/W	0x0000_8000	15.5.3.2.5/15-43
0x2_4118– 0x2_413C	Reserved	—	—	—
0x2_4140	TR03WT*—TxBD Rings 0–3 round-robin weightings	R/W	0x0000_0000	15.5.3.2.6/15-43
0x2_4144	TR47WT*—TxBD Rings 4–7 round-robin weightings	R/W	0x0000_0000	15.5.3.2.7/15-44
0x2_4148– 0x2_4180	Reserved	—	—	—
0x2_4180	TBDBPH*—Tx data buffer pointer high bits	R/W	0x0000_0000	15.5.3.2.8/15-45

Table 2-2. Memory Map (continued)

Offset	Register	Access	Reset	Section/Page
0x2_4184	TBPTR0—TxBD pointer for ring 0	R/W	0x0000_0000	<a href="#">15.5.3.2.9/15-45</a>
0x2_4188	Reserved	—	—	—
0x2_418C	TBPTR1*—TxBD pointer for ring 1	R/W	0x0000_0000	<a href="#">15.5.3.2.9/15-45</a>
0x2_4190	Reserved	—	—	—
0x2_4194	TBPTR2*—TxBD pointer for ring 2	R/W	0x0000_0000	<a href="#">15.5.3.2.9/15-45</a>
0x2_4198	Reserved	—	—	—
0x2_419C	TBPTR3*—TxBD pointer for ring 3	R/W	0x0000_0000	<a href="#">15.5.3.2.9/15-45</a>
0x2_41A0	Reserved	—	—	—
0x2_41A4	TBPTR4*—TxBD pointer for ring 4	R/W	0x0000_0000	<a href="#">15.5.3.2.9/15-45</a>
0x2_41A8	Reserved	—	—	—
0x2_41AC	TBPTR5*—TxBD pointer for ring 5	R/W	0x0000_0000	<a href="#">15.5.3.2.9/15-45</a>
0x2_41B0	Reserved	—	—	—
0x2_41B4	TBPTR6*—TxBD pointer for ring 6	R/W	0x0000_0000	<a href="#">15.5.3.2.9/15-45</a>
0x2_41B8	Reserved	—	—	—
0x2_41BC	TBPTR7*—TxBD pointer for ring 7	R/W	0x0000_0000	<a href="#">15.5.3.2.9/15-45</a>
0x2_41C0– 0x2_4200	Reserved	—	—	—
0x2_4204	TBASE0—TxBD base address of ring 0	R/W	0x0000_0000	<a href="#">15.5.3.2.10/15-46</a>
0x2_4208	Reserved	—	—	—
0x2_420C	TBASE1*—TxBD base address of ring 1	R/W	0x0000_0000	<a href="#">15.5.3.2.10/15-46</a>
0x2_4210	Reserved	—	—	—
0x2_4214	TBASE2*—TxBD base address of ring 2	R/W	0x0000_0000	<a href="#">15.5.3.2.10/15-46</a>
0x2_4218	Reserved	—	—	—
0x2_421C	TBASE3*—TxBD base address of ring 3	R/W	0x0000_0000	<a href="#">15.5.3.2.10/15-46</a>
0x2_4220	Reserved	—	—	—
0x2_4224	TBASE4*—TxBD base address of ring 4	R/W	0x0000_0000	<a href="#">15.5.3.2.10/15-46</a>
0x2_4228	Reserved	—	—	—
0x2_422C	TBASE5*—TxBD base address of ring 5	R/W	0x0000_0000	<a href="#">15.5.3.2.10/15-46</a>
0x2_4230	Reserved	—	—	—
0x2_4234	TBASE6*—TxBD base address of ring 6	R/W	0x0000_0000	<a href="#">15.5.3.2.10/15-46</a>
0x2_4238	Reserved	—	—	—
0x2_423C	TBASE7*—TxBD base address of ring 7	R/W	0x0000_0000	<a href="#">15.5.3.2.10/15-46</a>
0x2_4240– 0x2_427C	Reserved	—	—	—
0x2_4280	TMR_TXTS1_ID* - Tx timestamp identification tag (set 1)	R/W	0x0000_0000	<a href="#">15.5.3.2.11/15-47</a>

Table 2-2. Memory Map (continued)

Offset	Register	Access	Reset	Section/Page
0x2_4284	TMR_TXTS2_ID* - Tx timestamp identification tag (set 2)	R/W	0x0000_0000	15.5.3.2.11/15-47
0x2_4288– 0x2_42BC	Reserved	—	—	—
0x2_42C0	TMR_TXTS1_H* - Tx timestamp high (set 1)	R/W	0x0000_0000	15.5.3.2.12/15-47
0x2_42C4	TMR_TXTS1_L* - Tx timestamp high (set 1)	R/W	0x0000_0000	15.5.3.2.12/15-47
0x2_42C8	TMR_TXTS2_H* - Tx timestamp high (set 2)	R/W	0x0000_0000	15.5.3.2.12/15-47
0x2_42CC	TMR_TXTS2_L* - Tx timestamp high (set 2)	R/W	0x0000_0000	15.5.3.2.12/15-47
0x2_42D0– 0x2_42FC	Reserved	—	—	—
<b>eTSEC Receive Control and Status Registers</b>				
0x2_4300	RCTRL—Receive control register	R/W	0x0000_0000	15.5.3.3.1/15-48
0x2_4304	RSTAT—Receive status register	w1c	0x0000_0000	15.5.3.3.2/15-50
0x2_4308– 0x2_430C	Reserved	—	—	—
0x2_4310	RXIC—Receive interrupt coalescing register	R/W	0x0000_0000	15.5.3.3.3/15-52
0x2_4314	RQUEUE*—Receive queue control register.	R/W	0x0080_0080	15.5.3.3.4/15-53
0x2_4318– 0x2_432C	Reserved	—	—	—
0x2_4330	RBIFX*—Receive bit field extract control register	R/W	0x0000_0000	15.5.3.3.5/15-54
0x2_4334	RQFAR*—Receive queue filing table address register	R/W	0x0000_0000	15.5.3.3.6/15-55
0x2_4338	RQFCR*—Receive queue filing table control register	R/W	0xnnnn_nnnn	15.5.3.3.7/15-56
0x2_433C	RQFPR*—Receive queue filing table property register	R/W	0xnnnn_nnnn	15.5.3.3.8/15-57
0x2_4340	MRBLR—Maximum receive buffer length register	R/W	0x0000_0000	15.5.3.3.9/15-61
0x2_4344– 0x2_4380	Reserved	—	—	—
0x2_4380	RBDBPH*—Rx data buffer pointer high bits	R/W	0x0000_0000	15.5.3.3.10/15-62
0x2_4384	RBPTR0—RxBd pointer for ring 0	R/W	0x0000_0000	15.5.3.3.11/15-62
0x2_4388	Reserved	—	—	—
0x2_438C	RBPTR1*—RxBd pointer for ring 1	R/W	0x0000_0000	15.5.3.3.11/15-62
0x2_4390	Reserved	—	—	—
0x2_4394	RBPTR2*—RxBd pointer for ring 2	R/W	0x0000_0000	15.5.3.3.11/15-62
0x2_4398	Reserved	—	—	—
0x2_439C	RBPTR3*—RxBd pointer for ring 3	R/W	0x0000_0000	15.5.3.3.11/15-62
0x2_43A0	Reserved	—	—	—
0x2_43A4	RBPTR4*—RxBd pointer for ring 4	R/W	0x0000_0000	15.5.3.3.11/15-62
0x2_43A8	Reserved	—	—	—

Table 2-2. Memory Map (continued)

Offset	Register	Access	Reset	Section/Page
0x2_43AC	RBPTR5*—RxBd pointer for ring 5	R/W	0x0000_0000	15.5.3.3.11/15-62
0x2_43B0	Reserved	—	—	—
0x2_43B4	RBPTR6*—RxBd pointer for ring 6	R/W	0x0000_0000	15.5.3.3.11/15-62
0x2_43B8	Reserved	—	—	—
0x2_43BC	RBPTR7*—RxBd pointer for ring 7	R/W	0x0000_0000	15.5.3.3.11/15-62
0x2_43C0– 0x2_44400	Reserved	—	—	—
0x2_4404	RBASE0—RxBd base address of ring 0	R/W	0x0000_0000	15.5.3.3.12/15-63
0x2_4408	Reserved	—	—	—
0x2_440C	RBASE1*—RxBd base address of ring 1	R/W	0x0000_0000	15.5.3.3.12/15-63
0x2_4410	Reserved	—	—	—
0x2_4414	RBASE2*—RxBd base address of ring 2	R/W	0x0000_0000	15.5.3.3.12/15-63
0x2_4418	Reserved	—	—	—
0x2_441C	RBASE3*—RxBd base address of ring 3	R/W	0x0000_0000	15.5.3.3.12/15-63
0x2_4420	Reserved	—	—	—
0x2_4424	RBASE4*—RxBd base address of ring 4	R/W	0x0000_0000	15.5.3.3.12/15-63
0x2_4428	Reserved	—	—	—
0x2_442C	RBASE5*—RxBd base address of ring 5	R/W	0x0000_0000	15.5.3.3.12/15-63
0x2_4430	Reserved	—	—	—
0x2_4434	RBASE6*—RxBd base address of ring 6	R/W	0x0000_0000	15.5.3.3.12/15-63
0x2_4438	Reserved	—	—	—
0x2_443C	RBASE7*—RxBd base address of ring 7	R/W	0x0000_0000	15.5.3.3.12/15-63
0x2_4440– 0x2_44BC	Reserved	—	—	—
0x2_44C0	TMR_RXTS_H* - Rx timer timestamp register high	R/W	0x0000_0000	15.5.3.3.13/15-63
0x2_44C4	TMR_RXTS_L* - Rx timer timestamp register low	R/W	0x0000_0000	15.5.3.3.13/15-63
0x2_44C8– 0x2_44FC	Reserved	—	—	—
<b>eTSEC MAC Registers</b>				
0x2_4500	MACCFG1—MAC configuration register 1	R/W	0x0000_0000	15.5.3.5.1/15-67
0x2_4504	MACCFG2—MAC configuration register 2	R/W	0x0000_7000	15.5.3.5.2/15-68
0x2_4508	IPGIFG—Inter-packet/inter-frame gap register	R/W	0x4060_5060	15.5.3.5.3/15-70
0x2_450C	HAFDUP—Half-duplex control	R/W	0x00A1_F037	15.5.3.5.4/15-71
0x2_4510	MAXFRM—Maximum frame length	R/W	0x0000_0600	15.5.3.5.5/15-72
0x2_4514– 0x2_451C	Reserved	—	—	—

Table 2-2. Memory Map (continued)

Offset	Register	Access	Reset	Section/Page
0x2_4520	MIIMCFG—MII management configuration	R/W	0x0000_0007	<a href="#">15.5.3.5.6/15-72</a>
0x2_4524	MIIMCOM—MII management command	R/W	0x0000_0000	<a href="#">15.5.3.5.7/15-73</a>
0x2_4528	MIIMADD—MII management address	R/W	0x0000_0000	<a href="#">15.5.3.5.8/15-74</a>
0x2_452C	MIIMCON—MII management control	WO	0x0000_0000	<a href="#">15.5.3.5.9/15-75</a>
0x2_4530	MIIMSTAT—MII management status	R	0x0000_0000	<a href="#">15.5.3.5.10/15-75</a>
0x2_4534	MIIMIND—MII management indicator	R	0x0000_0000	<a href="#">15.5.3.5.11/15-76</a>
0x2_4538	Reserved	—	—	—
0x2_453C	IFSTAT—Interface status	R	0x0000_0000	<a href="#">15.5.3.5.12/15-76</a>
0x2_4540	MACSTNADDR1—MAC station address register 1	R/W	0x0000_0000	<a href="#">15.5.3.5.13/15-77</a>
0x2_4544	MACSTNADDR2—MAC station address register 2	R/W	0x0000_0000	<a href="#">15.5.3.5.14/15-78</a>
0x2_4548	MAC01ADDR1*—MAC exact match address 1, part 1	R/W	0x0000_0000	<a href="#">15.5.3.5.15/15-78</a> <a href="#">15.5.3.5.16/15-79</a>
0x2_454C	MAC01ADDR2*—MAC exact match address 1, part 2	R/W	0x0000_0000	
0x2_4550	MAC02ADDR1*—MAC exact match address 2, part 1	R/W	0x0000_0000	
0x2_4554	MAC02ADDR2*—MAC exact match address 2, part 2	R/W	0x0000_0000	
0x2_4558	MAC03ADDR1*—MAC exact match address 3, part 1	R/W	0x0000_0000	
0x2_455C	MAC03ADDR2*—MAC exact match address 3, part 2	R/W	0x0000_0000	
0x2_4560	MAC04ADDR1*—MAC exact match address 4, part 1	R/W	0x0000_0000	
0x2_4564	MAC04ADDR2*—MAC exact match address 4, part 2	R/W	0x0000_0000	
0x2_4568	MAC05ADDR1*—MAC exact match address 5, part 1	R/W	0x0000_0000	
0x2_456C	MAC05ADDR2*—MAC exact match address 5, part 2	R/W	0x0000_0000	

Table 2-2. Memory Map (continued)

Offset	Register	Access	Reset	Section/Page
0x2_4570	MAC06ADDR1*—MAC exact match address 6, part 1	R/W	0x0000_0000	15.5.3.5.15/15-78 15.5.3.5.16/15-79
0x2_4574	MAC06ADDR2*—MAC exact match address 6, part 2	R/W	0x0000_0000	
0x2_4578	MAC07ADDR1*—MAC exact match address 7, part 1	R/W	0x0000_0000	
0x2_457C	MAC07ADDR2*—MAC exact match address 7, part 2	R/W	0x0000_0000	
0x2_4580	MAC08ADDR1*—MAC exact match address 8, part 1	R/W	0x0000_0000	
0x2_4584	MAC08ADDR2*—MAC exact match address 8, part 2	R/W	0x0000_0000	
0x2_4588	MAC09ADDR1*—MAC exact match address 9, part 1	R/W	0x0000_0000	
0x2_458C	MAC09ADDR2*—MAC exact match address 9, part 2	R/W	0x0000_0000	
0x2_4590	MAC10ADDR1*—MAC exact match address 10, part 1	R/W	0x0000_0000	
0x2_4594	MAC10ADDR2*—MAC exact match address 10, part 2	R/W	0x0000_0000	
0x2_4598	MAC11ADDR1*—MAC exact match address 11, part 1	R/W	0x0000_0000	
0x2_459C	MAC11ADDR2*—MAC exact match address 11, part 2	R/W	0x0000_0000	
0x2_45A0	MAC12ADDR1*—MAC exact match address 12, part 1	R/W	0x0000_0000	
0x2_45A4	MAC12ADDR2*—MAC exact match address 12, part 2	R/W	0x0000_0000	
0x2_45A8	MAC13ADDR1*—MAC exact match address 13, part 1	R/W	0x0000_0000	
0x2_45AC	MAC13ADDR2*—MAC exact match address 13, part 2	R/W	0x0000_0000	
0x2_45B0	MAC14ADDR1*—MAC exact match address 14, part 1	R/W	0x0000_0000	
0x2_45B4	MAC14ADDR2*—MAC exact match address 14, part 2	R/W	0x0000_0000	
0x2_45B8	MAC15ADDR1*—MAC exact match address 15, part 1	R/W	0x0000_0000	
0x2_45BC	MAC15ADDR2*—MAC exact match address 15, part 2	R/W	0x0000_0000	
0x2_45C0– 0x2_467C	Reserved	—	—	—
<b>eTSEC Transmit and Receive Counters</b>				
0x2_4680	TR64—Transmit and receive 64-byte frame counter	R/W	0x0000_0000	15.5.3.6.1/15-80
0x2_4684	TR127—Transmit and receive 65- to 127-byte frame counter	R/W	0x0000_0000	15.5.3.6.2/15-80
0x2_4688	TR255—Transmit and receive 128- to 255-byte frame counter	R/W	0x0000_0000	15.5.3.6.3/15-81
0x2_468C	TR511—Transmit and receive 256- to 511-byte frame counter	R/W	0x0000_0000	15.5.3.6.4/15-81
0x2_4690	TR1K—Transmit and receive 512- to 1023-byte frame counter	R/W	0x0000_0000	15.5.3.6.5/15-82
0x2_4694	TRMAX—Transmit and receive 1024- to 1518-byte frame counter	R/W	0x0000_0000	15.5.3.6.6/15-82
0x2_4698	TRMGV—Transmit and receive 1519- to 1522-byte good VLAN frame count	R/W	0x0000_0000	15.5.3.6.7/15-83
<b>eTSEC Receive Counters</b>				
0x2_469C	RBYT—Receive byte counter	R/W	0x0000_0000	15.5.3.6.8/15-83

Table 2-2. Memory Map (continued)

Offset	Register	Access	Reset	Section/Page
0x2_46A0	RPKT—Receive packet counter	R/W	0x0000_0000	15.5.3.6.9/15-83
0x2_46A4	RFCS—Receive FCS error counter	R/W	0x0000_0000	15.5.3.6.10/15-84
0x2_46A8	RMCA—Receive multicast packet counter	R/W	0x0000_0000	15.5.3.6.11/15-84
0x2_46AC	RBCA—Receive broadcast packet counter	R/W	0x0000_0000	15.5.3.6.12/15-85
0x2_46B0	RXCF—Receive control frame packet counter	R/W	0x0000_0000	15.5.3.6.13/15-85
0x2_46B4	RXPF—Receive PAUSE frame packet counter	R/W	0x0000_0000	15.5.3.6.14/15-86
0x2_46B8	RXUO—Receive unknown OP code counter	R/W	0x0000_0000	15.5.3.6.15/15-86
0x2_46BC	RALN—Receive alignment error counter	R/W	0x0000_0000	15.5.3.6.16/15-87
0x2_46C0	RFLR—Receive frame length error counter	R/W	0x0000_0000	15.5.3.6.17/15-87
0x2_46C4	RCDE—Receive code error counter	R/W	0x0000_0000	15.5.3.6.18/15-88
0x2_46C8	RCSE—Receive carrier sense error counter	R/W	0x0000_0000	15.5.3.6.19/15-88
0x2_46CC	RUND—Receive undersize packet counter	R/W	0x0000_0000	15.5.3.6.20/15-89
0x2_46D0	ROVR—Receive oversize packet counter	R/W	0x0000_0000	15.5.3.6.21/15-89
0x2_46D4	RFRG—Receive fragments counter	R/W	0x0000_0000	15.5.3.6.22/15-90
0x2_46D8	RJBR—Receive jabber counter	R/W	0x0000_0000	15.5.3.6.23/15-90
0x2_46DC	RDRP—Receive drop counter	R/W	0x0000_0000	15.5.3.6.24/15-91
<b>eTSEC Transmit Counters</b>				
0x2_46E0	TBYT—Transmit byte counter	R/W	0x0000_0000	15.5.3.6.25/15-91
0x2_46E4	TPKT—Transmit packet counter	R/W	0x0000_0000	15.5.3.6.26/15-92
0x2_46E8	TMCA—Transmit multicast packet counter	R/W	0x0000_0000	15.5.3.6.27/15-92
0x2_46EC	TBCA—Transmit broadcast packet counter	R/W	0x0000_0000	15.5.3.6.28/15-93
0x2_46F0	TXPF—Transmit PAUSE control frame counter	R/W	0x0000_0000	15.5.3.6.29/15-93
0x2_46F4	TDFR—Transmit deferral packet counter	R/W	0x0000_0000	15.5.3.6.30/15-94
0x2_46F8	TEDF—Transmit excessive deferral packet counter	R/W	0x0000_0000	15.5.3.6.31/15-94
0x2_46FC	TSCL—Transmit single collision packet counter	R/W	0x0000_0000	15.5.3.6.32/15-95
0x2_4700	TMCL—Transmit multiple collision packet counter	R/W	0x0000_0000	15.5.3.6.33/15-95
0x2_4704	TLCL—Transmit late collision packet counter	R/W	0x0000_0000	15.5.3.6.34/15-96
0x2_4708	TXCL—Transmit excessive collision packet counter	R/W	0x0000_0000	15.5.3.6.35/15-96
0x2_470C	TNCL—Transmit total collision counter	R/W	0x0000_0000	15.5.3.6.36/15-97
0x2_4710	Reserved	—	—	—
0x2_4714	TDRP—Transmit drop frame counter	R/W	0x0000_0000	15.5.3.6.37/15-97
0x2_4718	TJBR—Transmit jabber frame counter	R/W	0x0000_0000	15.5.3.6.38/15-98
0x2_471C	TFCS—Transmit FCS error counter	R/W	0x0000_0000	15.5.3.6.39/15-98
0x2_4720	TXCF—Transmit control frame counter	R/W	0x0000_0000	15.5.3.6.40/15-99
0x2_4724	TOVR—Transmit oversize frame counter	R/W	0x0000_0000	15.5.3.6.41/15-99



Table 2-2. Memory Map (continued)

Offset	Register	Access	Reset	Section/Page
0x2_4728	TUND—Transmit undersize frame counter	R/W	0x0000_0000	15.5.3.6.42/15-100
0x2_472C	TFRG—Transmit fragments frame counter	R/W	0x0000_0000	15.5.3.6.43/15-100
<b>eTSEC Counter Control and TOE Statistics Registers</b>				
0x2_4730	CAR1—Carry register one register <sup>14</sup>	R	0x0000_0000	15.5.3.6.44/15-101
0x2_4734	CAR2—Carry register two register <sup>14</sup>	R	0x0000_0000	15.5.3.6.45/15-102
0x2_4738	CAM1—Carry register one mask register	R/W	0xFE03_FFFF	15.5.3.6.46/15-103
0x2_473C	CAM2—Carry register two mask register	R/W	0x000F_FFFD	15.5.3.6.47/15-105
0x2_4740	RREJ*—Receive filer rejected packet counter	R/W	0x0000_0000	15.5.3.6.48/15-106
0x2_4744– 0x2_47FC	Reserved	—	—	—
<b>Hash Function Registers</b>				
0x2_4800	IGADDR0—Individual/group address register 0	R/W	0x0000_0000	15.5.3.7.1/15-107
0x2_4804	IGADDR1—Individual/group address register 1	R/W	0x0000_0000	
0x2_4808	IGADDR2—Individual/group address register 2	R/W	0x0000_0000	
0x2_480C	IGADDR3—Individual/group address register 3	R/W	0x0000_0000	
0x2_4810	IGADDR4—Individual/group address register 4	R/W	0x0000_0000	
0x2_4814	IGADDR5—Individual/group address register 5	R/W	0x0000_0000	
0x2_4818	IGADDR6—Individual/group address register 6	R/W	0x0000_0000	
0x2_481C	IGADDR7—Individual/group address register 7	R/W	0x0000_0000	
0x2_4820– 0x2_487C	Reserved	—	—	—
0x2_4880	GADDR0—Group address register 0	R/W	0x0000_0000	15.5.3.7.2/15-107
0x2_4884	GADDR1—Group address register 1	R/W	0x0000_0000	
0x2_4888	GADDR2—Group address register 2	R/W	0x0000_0000	
0x2_488C	GADDR3—Group address register 3	R/W	0x0000_0000	
0x2_4890	GADDR4—Group address register 4	R/W	0x0000_0000	
0x2_4894	GADDR5—Group address register 5	R/W	0x0000_0000	
0x2_4898	GADDR6—Group address register 6	R/W	0x0000_0000	
0x2_489C	GADDR7—Group address register 7	R/W	0x0000_0000	
0x2_48A0– 0x2_4AFC	Reserved	—	—	—
<b>eTSEC DMA Attribute Registers</b>				
0x2_4B00– 0x2_4BF4	Reserved	—	—	—
0x2_4BF8	ATTR—Attribute register	R/W	0x0000_0000	15.5.3.8.1/15-108

Table 2-2. Memory Map (continued)

Offset	Register	Access	Reset	Section/Page
<b>eTSEC Lossless Flow Control Registers</b>				
0x2_4C00	RQPRM0*—Receive Queue Parameters register 0	R/W	0x0000_0000	15.5.3.9.1/15-109
0x2_4C04	RQPRM1*—Receive Queue Parameters register 1	R/W	0x0000_0000	
0x2_4C08	RQPRM2*—Receive Queue Parameters register 2	R/W	0x0000_0000	
0x2_4C0C	RQPRM3*—Receive Queue Parameters register 3	R/W	0x0000_0000	
0x2_4C10	RQPRM4*—Receive Queue Parameters register 4	R/W	0x0000_0000	
0x2_4C14	RQPRM5*—Receive Queue Parameters register 5	R/W	0x0000_0000	
0x2_4C18	RQPRM6*—Receive Queue Parameters register 6	R/W	0x0000_0000	
0x2_4C1C	RQPRM7*—Receive Queue Parameters register 7	R/W	0x0000_0000	
0x2_4C20– 0x2_4C40	Reserved	—	—	—
0x2_4C44	RFBPTR0*—Last Free RxBD pointer for ring 0	R/W	0x0000_0000	15.5.3.9.2/15-109
0x2_4C48	Reserved	—	—	—
0x2_4C4C	RFBPTR1*—Last Free RxBD pointer for ring 1	R/W	0x0000_0000	15.5.3.9.2/15-109
0x2_4C50	Reserved	—	—	—
0x2_4C54	RFBPTR2*—Last Free RxBD pointer for ring 2	R/W	0x0000_0000	15.5.3.9.2/15-109
0x2_4C58	Reserved	—	—	—
0x2_4C5C	RFBPTR3*—Last Free RxBD pointer for ring 3	R/W	0x0000_0000	15.5.3.9.2/15-109
0x2_4C60	Reserved	—	—	—
0x2_4C64	RFBPTR4*—Last Free RxBD pointer for ring 4	R/W	0x0000_0000	15.5.3.9.2/15-109
0x2_4C68	Reserved	—	—	—
0x2_4C6C	RFBPTR5*—Last Free RxBD pointer for ring 5	R/W	0x0000_0000	15.5.3.9.2/15-109
0x2_4C70	Reserved	—	—	—
0x2_4C74	RFBPTR6*—Last Free RxBD pointer for ring 6	R/W	0x0000_0000	15.5.3.9.2/15-109
0x2_4C78	Reserved	—	—	—
0x2_4C7C	RFBPTR7*—Last Free RxBD pointer for ring 7	R/W	0x0000_0000	15.5.3.9.2/15-109
<b>eTSEC Future Expansion Space</b>				
0x2_4CC0– 0x2_4D94	Reserved	—	—	—
<b>eTSEC IEEE 1588 Registers</b>				
0x2_4E00	TMR_CTRL*—Timer control register	R/W	0x0001_0001	15.5.3.10.1/15-110
0x2_4E04	TMR_TEVENT*—Timestamp event register	W1C	0x0000_0000	15.5.3.10.2/15-112
0x2_4E08	TMR_TEMASK*—Timer event mask register	R/W	0x0000_0000	15.5.3.10.3/15-114
0x2_4E0C	TMR_PEVENT*—Timestamp event register	R/W	0x0000_0000	15.5.3.10.4/15-115
0x2_4E10	TMR_PEMASK*—Timer event mask register	R/W	0x0000_0000	15.5.3.10.5/15-115

Table 2-2. Memory Map (continued)

Offset	Register	Access	Reset	Section/Page
0x2_4E14	TMR_STAT*—Timestamp status register	R/W	0x0000_0000	15.5.3.10.6/15-116
0x2_4E18	TMR_CNT_H*—Timer counter high register	R/W	0x0000_0000	15.5.3.10.7/15-117
0x2_4E1C	TMR_CNT_L*—Timer counter low register	R/W	0x0000_0000	15.5.3.10.7/15-117
0x2_4E20	TMR_ADD*—Timer drift compensation addend register	R/W	0x0000_0000	15.5.3.10.8/15-117
0x2_4E24	TMR_ACC*—Timer accumulator register	R/W	0x0000_0000	15.5.3.10.9/15-118
0x2_4E28	TMR_PRSC*—Timer prescale	R/W	0x0000_0002	15.5.3.10.10/15-118
0x2_4E2C	Reserved	—	—	—
0x2_4E30	TMROFF_H*—Timer offset high	R/W	0x0000_0000	15.5.3.10.11/15-119
0x2_4E34	TMROFF_L*—Timer offset low	R/W	0x0000_0000	15.5.3.10.11/15-119
0x2_4E40	TMR_ALARM1_H*—Timer alarm 1 high register	R/W	0xFFFF_FFFF	15.5.3.10.12/15-119
0x2_4E44	TMR_ALARM1_L*—Timer alarm 1 high register	R/W	0xFFFF_FFFF	
0x2_4E48	TMR_ALARM2_H*—Timer alarm 2 high register	R/W	0xFFFF_FFFF	
0x2_4E4C	TMR_ALARM2_L*—Timer alarm 2 high register	R/W	0xFFFF_FFFF	
0x2_4E50– 0x2_4E7C	Reserved	—	—	—
0x2_4E80	TMR_FIPER1*—Timer fixed period interval	R/W	0xFFFF_FFFF	15.5.3.10.13/15-120
0x2_4E84	TMR_FIPER2*—Timer fixed period interval	R/W	0xFFFF_FFFF	
0x2_4E88	TMR_FIPER*3—Timer fixed period interval	R/W	0xFFFF_FFFF	
0x2_4EA0	TMR_ETTS1_H*—Timestamp of general purpose external trigger	R/W	0x0000_0000	15.5.3.10.14/15-121
0x2_4EA4	TMR_ETTS1_L*—Timestamp of general purpose external trigger	R/W	0x0000_0000	
0x2_4EA8	TMR_ETTS2_H*—Timestamp of general purpose external trigger	R/W	0x0000_0000	
0x2_4EAC	TMR_ETTS2_L*—Timestamp of general purpose external trigger	R/W	0x0000_0000	
0x2_4EB0– 0x2_4FFF	Reserved	—	—	—
<b>Other eTSECs</b>				
0x2_5000– 0x2_5FFF	eTSEC2 REGISTERS <sup>15</sup>			
<b>Security Engine Address Map Registers</b>				
<b>Controller Registers</b>				
0x3_0000– 0x3_0FFF	Reserved, should be cleared	—	—	—
0x3_1008	IMR—Interrupt mask register	R/W	0x0000_0000 _0000_0000	14.6.4.2/14-68

Table 2-2. Memory Map (continued)

Offset	Register	Access	Reset	Section/Page
0x3_1010	ISR—Interrupt status register	R	0x0000_0000 _0000_0000	14.6.4.3/14-70
0x3_1018	ICR—Interrupt clear register	W	0x0000_0000 _0000_0000	14.6.4.4/14-71
0x3_1020	ID—Identification register	R	0000_0000 _0000_00A0	14.6.4.5/14-73
0x3_1028	EUASR—EU assignment status register	R	0xF0F0_F0F0 _00FF_F0F0	14.6.4.1/14-67
0x3_1030	MCR—Master control register	R/W	0000_0000 _0000_0000	14.6.4.7/14-74
<b>Channel</b>				
0x3_1108	CCCR—Crypto-channel configuration register	R/W	0x0000_0000 _0000_0000	14.5.1.1/14-55
0x3_1110	CCPSR—Crypto-channel pointer status register	R	0x0000_0000 _0000_0007	14.5.1.2/14-57
0x3_1140	CDPR—Crypto-channel current descriptor pointer register	R	0x0000_0000 _0000_0000	14.5.1.3/14-62
0x3_1148	FF—Crypto-channel fetch FIFO address register	W	0x0000_0000 _0000_0000	14.5.1.4/14-62
0x3_1180– 0x3_11BF	DB $n$ —Crypto-channel descriptor buffers [0–7]	R	0x0000_0000 _0000_0000	14.5.1.5/14-63
0x3_1BF8	IP block revision register	R	0x0000_0000 _0002_00A0	14.6.4.6/14-73
<b>Data Encryption Standard Execution Unit (DEU)</b>				
0x3_2000	DEUMR—DEU mode register	R/W	0x0000_0000 _0000_0000	14.4.1.1/14-19
0x3_2008	DEUKSR—DEU key size register	R/W	0x0000_0000 _0000_0000	14.4.1.2/14-20
0x3_2010	DEUDSR—DEU data size register	R/W	0x0000_0000 _0000_0000	14.4.1.3/14-21
0x3_2018	DEURCR—DEU reset control register	R/W	0x0000_0000 _0000_0000	14.4.1.4/14-22
0x3_2028	DEUSR—DEU status register	R	0x0000_0000 _0000_0000	14.4.1.5/14-23
0x3_2030	DEUIR—DEU interrupt status register	R	0x0000_0000 _0000_0000	14.4.1.6/14-24
0x3_2038	DEUICR—DEU interrupt control register	R/W	0x0000_0000 _0000_3000	14.4.1.7/14-25
0x3_2050	DEUEMR—DEU end-of-message register	W	0x0000_0000 _0000_0000	14.4.1.8/14-27
0x3_2100	DEUIV—DEU initialization vector register	R/W	0x0000_0000 _0000_0000	14.4.1.9/14-27

Table 2-2. Memory Map (continued)

Offset	Register	Access	Reset	Section/Page
0x3_2400	DEUK1—DEU key 1 register	W	—	14.4.1.10/14-28
0x3_2408	DEUK2—DEU key 2 register	W	—	14.4.1.10/14-28
0x3_2410	DEUK3—DEU key 3 register	W	—	14.4.1.10/14-28
0x3_2800– 0x3_2FFF	DEU FIFO	R/W	0x0000_0000 _0000_0000	14.4.1.11/14-28
<b>Advanced Encryption Standard Execution Unit (AESU)</b>				
0x3_4000	AESUMR—AESU mode register	R/W	0x0000_0000 _0000_0000	14.4.3.1/14-40
0x3_4008	AESUKSR—AESU key size register	R/W	0x0000_0000 _0000_0000	14.4.3.2/14-42
0x3_4010	AESUDSR—AESU data size register	R/W	0x0000_0000 _0000_0000	14.4.3.3/14-43
0x3_4018	AESURCR—AESU reset control register	R/W	0x0000_0000 _0000_0000	14.4.3.4/14-43
0x3_4028	AESUSR—AESU status register	R	0x0000_0000 _0000_0000	14.4.3.5/14-44
0x3_4030	AESUISR—AESU interrupt status register	R	0x0000_0000 _0000_0000	14.4.3.6/14-45
0x3_4038	AESUICR—AESU interrupt control register	R/W	0x0000_0000 _0000_1000	14.4.3.7/14-47
0x3_4050	AESUEMR—AESU end-of-message register	W	0x0000_0000 _0000_0000	14.4.3.8/14-48
0x3_4100	AESU context memory registers	R/W	0x0000_0000 _0000_0000	14.4.3.9/14-49
0x3_4400– 0x3_4408	AESU key memory	R/W	0x0000_0000 _0000_0000	
0x3_4800– 0x3_4FFF	AESU FIFO	R/W	0x0000_0000 _0000_0000	
<b>Message Digest Execution Unit (MDEU)</b>				
0x3_6000	MDEUMR—MDEU mode register	R/W	0x0000_0000 _0000_0000	14.4.2.1/14-28
0x3_6008	MDEUKSR—MDEU key size register	R/W	0x0000_0000 _0000_0000	14.4.2.3/14-32
0x3_6010	MDEUDSR—MDEU data size register	R/W	0x0000_0000 _0000_0000	14.4.2.4/14-32
0x3_6018	MDEURCR—MDEU reset control register	R/W	0x0000_0000 _0000_0000	14.4.2.5/14-33
0x3_6028	MDEUSR—MDEU status register	R	0x0000_0000 _0000_0000	14.4.2.6/14-34
0x3_6030	MDEUISR—MDEU interrupt status register	R	0x0000_0000 _0000_0000	14.4.2.7/14-35

Table 2-2. Memory Map (continued)

Offset	Register	Access	Reset	Section/Page
0x3_6038	MDEUICR—MDEU interrupt control register	R/W	0x0000_0000 _0000_1000	14.4.2.8/14-36
0x3_6040	MDEU ICV size register	R/W	0x0000_0000_0 000_0000	14.4.2.9/14-37
0x3_6050	MDEUEMR—MDEU end-of-message register	W	0x0000_0000 _0000_0000	14.4.2.10/14-38
0x3_6100– 0x3_6120	MDEU context memory registers	R/W	0x0000_0000 _0000_0000	14.4.2.11/14-38
0x3_6400– 0x3_647F	MDEU key memory	W	0x0000_0000 _0000_0000	14.4.2.12/14-39
0x3_6800– 0x3_6FFF	MDEU FIFO	W	0x0000_0000 _0000_0000	14.4.2.13/14-40

- <sup>1</sup> Depends on reset configuration word high values. See [Section 5.2.4.3.1, “LBLAWBAR0\[BASE\\_ADDR\] Reset Value,”](#) for details.
- <sup>2</sup> Depends on reset configuration word high values. See [Section 5.2.4.4.1, “LBLAWAR0\[EN\] and LBLAWAR0\[SIZE\] Reset Value,”](#) for details.
- <sup>3</sup> Depends on reset configuration word high values. See [Section 5.2.4.5.1, “PCILAWBAR0\[BASE\\_ADDR\] Reset Value,”](#) for details.
- <sup>4</sup> Depends on reset configuration word high values. See [Section 5.2.4.7.1, “DDRLAWBAR0\[BASE\\_ADDR\] Reset Value,”](#) for details.
- <sup>5</sup> Depends on reset configuration word high values. See [Section 5.2.4.6.1, “PCILAWAR0\[EN\] and PCILAWAR0\[SIZE\] Reset Value,”](#) for details.
- <sup>6</sup> Depends on reset configuration word high values. See [Section 5.2.4.7.1, “DDRLAWBAR0\[BASE\\_ADDR\] Reset Value,”](#) for details.
- <sup>7</sup> Depends on reset configuration word high values. See [Section 5.2.4.8.1, “DDRLAWAR0\[EN\] and DDRLAWAR0\[SIZE\] Reset Value,”](#) for details.
- <sup>8</sup> Depends on the reset configuration word high configuration values.
- <sup>9</sup> SWCRR[SWEN] reset value directly depends on RCWHR[SWEN] (reset configuration word high).
- <sup>10</sup> Reset value is determined from the core PLL configuration of the reset configuration word. See [Chapter 4, “Reset, Clocking, and Initialization,”](#) for details.
- <sup>11</sup> The registers AEATR and AEADR are affected only by the assertion of  $\overline{\text{PORESET}}$
- <sup>12</sup> Implementation-dependent reset values are listed in specified section/page.
- <sup>13</sup> This register has separate functions for the host and device operation; the host function is listed first in the table.
- <sup>14</sup> Cleared on read.
- <sup>15</sup> eTSEC2 has the same memory-mapped registers that are described for eTSEC1 from 0x2\_4000 to 0x2\_4FFF, except the offsets are from 0x2\_5000 to 0x2\_5FFF.

## Chapter 3

# Signal Descriptions

This chapter describes the external signals of the device. It is organized into the following sections:

- Overview of signals and cross references for signals that serve multiple functions, including two lists: one ordered by functional block and one alphabetical.
- List of reset configuration signals
- List of output signal states at reset

### NOTE

A bar over a signal name indicates that the signal is active low, such as  $\overline{\text{IRQ\_OUT}}$  (interrupt out). Active-low signals are referred to as asserted (active) when they are low and negated when they are high. Signals that are not active low, such as IRQ (interrupt input), are referred to as asserted when they are high and negated when they are low.

Internal signals throughout this document are shown as lower case and in italics. For example, *sys\_logic\_clk* is an internal signal. These are referenced only as necessary for understanding of the external functionality of the device.

## 3.1 Signals Overview

The signals are grouped as follows:

- DDR memory interface signals
- PCI interface signals
- DUART interface signals
- I<sup>2</sup>C interface signals
- Serial peripheral interface signals
- Ethernet management interface signals
- eTSEC1 and USB interface signals
- eTSEC1 and 1588 interface signals
- eTSEC2 interface signals
- SerDes interface signals
- Enhanced local bus interface signals
- USB PHY signals
- Global timers/USB interface signals
- PIC interface signals
- SPI, JTAG, PMC, configuration, system control signals
- SGMII PHY interface signals
- Clock signals

Figure 3-1 and Figure 3-2 show the external signals of the device and how the signals are grouped. Refer to the *MPC8313E Integrated Processor Hardware Specifications* for a pinout diagram showing pin numbers and a listing of all the electrical and mechanical specifications.

## Signal Descriptions

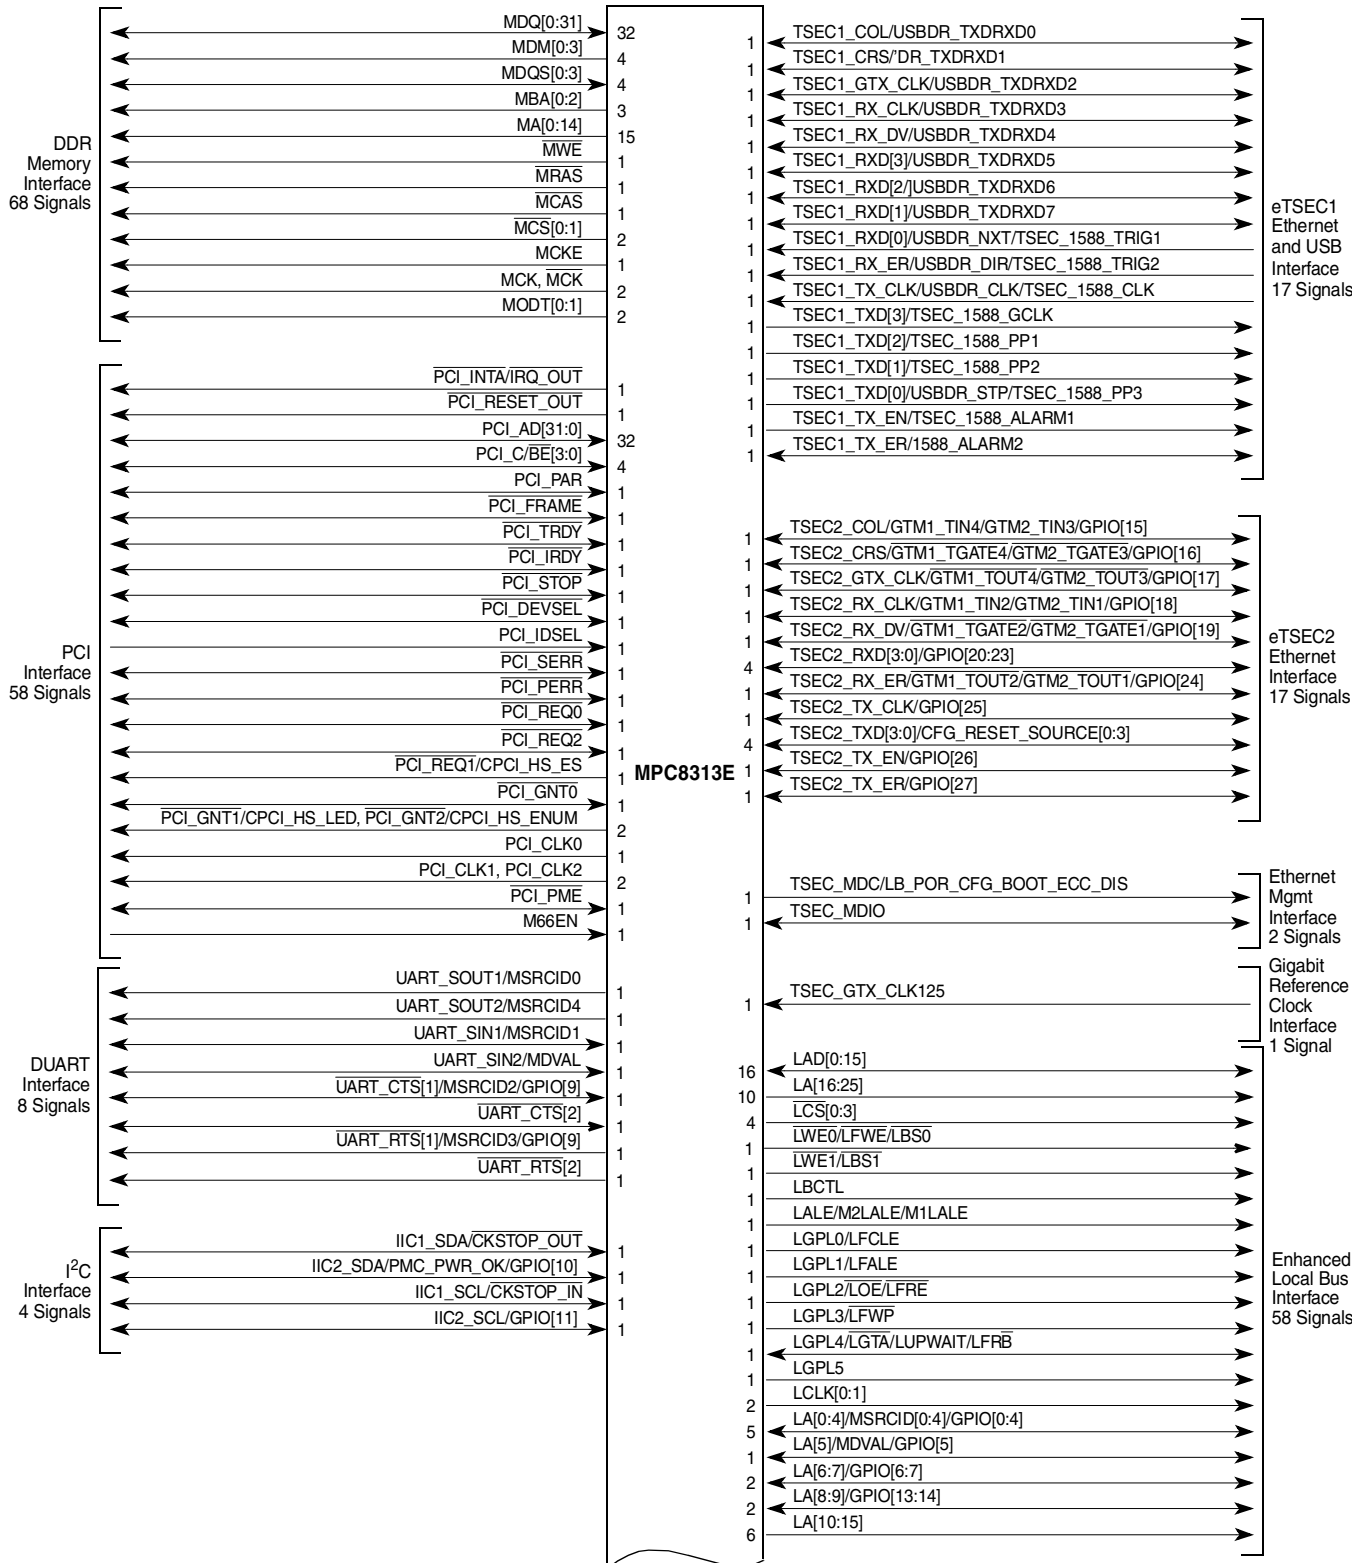


Figure 3-1. MPC8313E Signal Groupings (1 of 2)



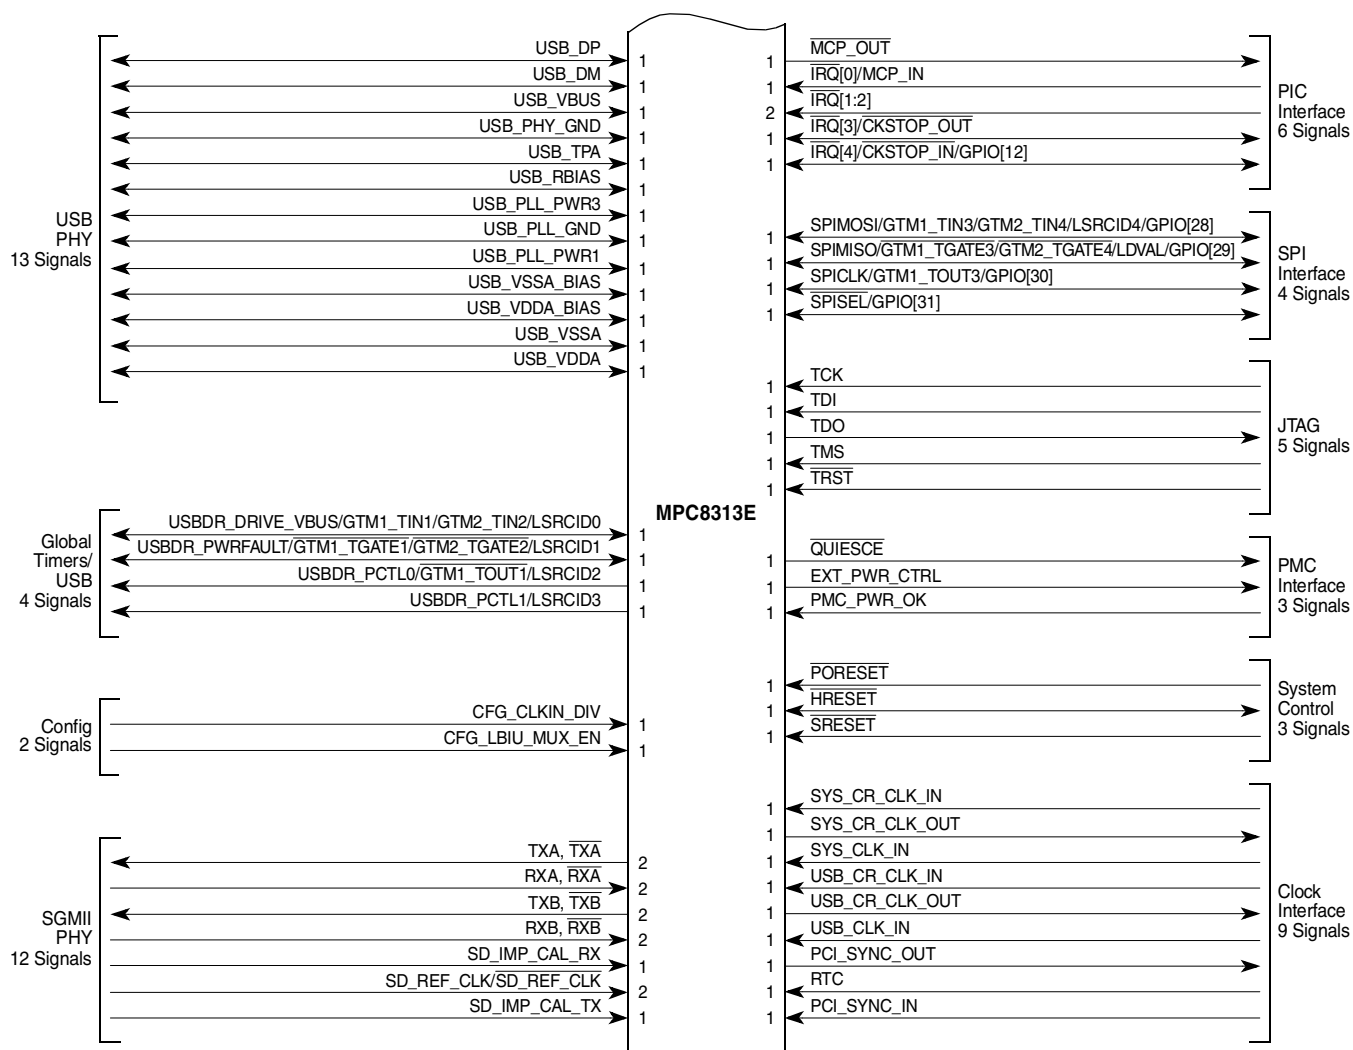


Figure 3-2. MPC8313E Signal Groupings (2 of 2)

Note that individual chapters of this document provide details for each signal, describing each signal's behavior when asserted and negated and when the signal is an input or an output.

The following tables provide summaries of signal functions. [Table 3-1](#) provides a summary of the signals grouped by function, and [Table 3-2](#) provides a summary of the signals grouped alphabetically. These tables detail the signal name, interface, alternate functions, number of signals, and whether the signal is an input, output, or bidirectional. Finally, the table provides a pointer to the table where the signal function is described.

Table 3-1. MPC8313E Signal Reference by Functional Block

Name	Description	Functional Block	No. of Signals	I/O	Table/ Page	Alternate Function(s)
MDQ[0:31]	DDR data	DDR	32	I/O	<a href="#">9-3/9-5</a>	—
MDM[0:3]	DDR data mask	DDR	4	O	<a href="#">9-3/9-5</a>	—

Table 3-1. MPC8313E Signal Reference by Functional Block (continued)

Name	Description	Functional Block	No. of Signals	I/O	Table/ Page	Alternate Function(s)
MDQS[0:3]	DDR data strobe	DDR	4	I/O	9-3/9-5	—
MBA[0:2]	DDR bank select	DDR	3	O	9-3/9-5	—
MA[0:14]	DDR address	DDR	15	O	9-3/9-5	—
$\overline{\text{MWE}}$	DDR write enable	DDR	1	O	9-3/9-5	—
$\overline{\text{MRAS}}$	DDR row address strobe	DDR	1	O	9-3/9-5	—
$\overline{\text{MCAS}}$	DDR column address strobe	DDR	1	O	9-3/9-5	—
$\overline{\text{MCS}}[0:1]$	DDR chip select (2/DIMM)	DDR	2	O	9-3/9-5	—
MCKE	DDR clock enable	DDR	1	O	9-4/9-7	—
MCK, $\overline{\text{MCK}}$	DDR differential clocks	DDR	2	O	9-4/9-7	—
MODT[0:1]	DRAM on-die termination	DDR	2	O	9-3/9-5	—
PCI_INTA	PCI interrupt output	PCI	1	O	13-3/13-5	—
$\overline{\text{PCI\_RESET\_OUT}}$	PCI reset	PCI	1	O	13-3/13-5	—
PCI_AD[0:31]	PCI address/data	PCI	32	I/O	13-3/13-5	—
PCI_C_BE[0:3]	PCI command/byte enable	PCI	4	I/O	13-3/13-5	—
PCI_PAR	PCI parity	PCI	1	I/O	13-3/13-5	—
$\overline{\text{PCI\_FRAME}}$	PCI frame	PCI	1	I/O	13-3/13-5	—
$\overline{\text{PCI\_TRDY}}$	PCI target ready	PCI	1	I/O	13-3/13-5	—
$\overline{\text{PCI\_IRDY}}$	PCI initiator ready	PCI	1	I/O	13-3/13-5	—
$\overline{\text{PCI\_STOP}}$	PCI stop	PCI	1	I/O	13-3/13-5	—
$\overline{\text{PCI\_DEVSEL}}$	PCI device select	PCI	1	I/O	13-3/13-5	—
PCI_IDSEL	PCI initial device select	PCI	1	I	13-3/13-5	—
$\overline{\text{PCI\_SERR}}$	PCI system error	PCI	1	I/O	13-3/13-5	—
$\overline{\text{PCI\_PERR}}$	PCI parity error	PCI	1	I/O	13-3/13-5	—
$\overline{\text{PCI\_REQ0}}$	PCI request 0	PCI	1	I/O	13-3/13-5	—
$\overline{\text{PCI\_REQ1}}$	PCI request 1	PCI	1	I	13-3/13-5	CPCI_HS_ES
$\overline{\text{PCI\_REQ2}}$	PCI request 2	PCI	1	I	13-3/13-5	—
$\overline{\text{PCI\_GNT0}}$	PCI grant 0	PCI	1	I/O	13-3/13-5	—
$\overline{\text{PCI\_GNT1}}$	PCI grant 1	PCI	1	O	13-3/13-5	CPCI_HS_LED

Table 3-1. MPC8313E Signal Reference by Functional Block (continued)

Name	Description	Functional Block	No. of Signals	I/O	Table/ Page	Alternate Function(s)
PCI_GNT2	PCI grant 2	PCI	1	O	13-3/13-5	CPCI_HS_ENUM
M66EN	66-MHz system configuration	PCI	1	I	13-3/13-5	—
PCI_CLK0	PCI clock 0	PCI	1	O	4-2/4-3	—
PCI_CLK1	PCI clock 1	PCI	1	O	4-2/4-3	—
PCI_CLK2	PCI clock 2	PCI	1	O	4-2/4-3	—
PCI_PME	PCI PME assertion request	PCI	1	I/O	13-3/13-5	—
CPCI_HS_ENUM	CompactPCI hot swap enumerator	PCI	1	O	13-3/13-5	PCI_GNT2
CPCI_HS_ES	CompactPCI hot swap ejector switch	PCI	1	I	13-3/13-5	PCI_REQ1
CPCI_HS_LED	CompactPCI hot swap LED	PCI	1	O	13-3/13-5	PCI_GNT1
UART_SOUT1	DUART serial data out	DUART	1	I/O	18-2/18-3	MSRCID0
UART_SOUT2	DUART serial data out	DUART	1	I/O	18-2/18-3	TSEC_1588_CLK/ MSRCID4
UART_SIN1	DUART serial data in	DUART	1	I/O	18-2/18-3	MSRCID1
UART_SIN2	DUART serial data in	DUART	1	I/O	18-2/18-3	TSEC_1588_GCLK/ MDVAL
UART_CTS1	DUART clear to send	DUART	1	I/O	18-2/18-3	MSRCID2/GPIO8
UART_CTS2	DUART clear to send	DUART	1	I/O	18-2/18-3	TSEC_1588_PP1
UART_RTS1	DUART ready to send	DUART	1	I/O	18-2/18-3	MSRCID3/GPIO9
UART_RTS2	DUART ready to send	DUART	1	I/O	18-2/18-3	TSEC_1588_PP2
IIC1_SDA	I <sup>2</sup> C serial data	I <sup>2</sup> C1	1	I/O	17-1/17-3	CKSTOP_OUT/ TSEC_1588_TRIG1
IIC1_SCL	I <sup>2</sup> C serial clock	I <sup>2</sup> C1	1	I/O	17-1/17-3	CKSTOP_IN/ TSEC_1588_ALARM2
IIC2_SDA	I <sup>2</sup> C serial data	I <sup>2</sup> C2	1	I/O	17-1/17-3	PMC_PWR_OK GPIO10
IIC2_SCL	I <sup>2</sup> C serial clock	I <sup>2</sup> C2	1	I/O	17-1/17-3	GPIO11

Table 3-1. MPC8313E Signal Reference by Functional Block (continued)

Name	Description	Functional Block	No. of Signals	I/O	Table/ Page	Alternate Function(s)
SPIMOSI	SPI master-out slave-in	SPI	1	I/O	19-2/19-7	GTM1_TIN3/ GTM2_TIN4/LSRCID4 /GPIO28
SPIMISO	SPI master-in slave-out	SPI	1	I/O	19-2/19-7	GTM1_TGATE3/ GTM2_TGATE4/ LDVAL/GPIO29
SPICLK	SPI clock	SPI	1	I/O	19-2/19-7	$\overline{\text{GTM1\_TOUT3}}$ / GPIO30
SPISEL	SPI slave select	SPI	1	I	19-2/19-7	GPIO31
TSEC1_COL	eTSEC1 collision detect	eTSEC1	1	I/O	15-2/15-8	USBDR_TXDRXD0
TSEC1_CRS	eTSEC1 carrier sense	eTSEC1	1	I/O	15-2/15-8	USBDR_TXDRXD1
TSEC1_GTX_CLK	eTSEC1 transmit clock out	eTSEC1	1	I/O	15-2/15-8	USBDR_TXDRXD2
TSEC1_RX_CLK	eTSEC1 receive clock	eTSEC1	1	I/O	15-2/15-8	USBDR_TXDRXD3
TSEC1_RX_DV	eTSEC1 receive data valid	eTSEC1	1	I/O	15-2/15-8	USBDR_TXDRXD4
TSEC1_RXD[3:1]	eTSEC1 receive data 3–1	eTSEC1	3	I/O	15-2/15-8	USBDR_TXDRXD [5:7]
TSEC1_RXD0	eTSEC1 receive data 0	eTSEC1	1	I	15-2/15-8	USBDR_NXT
TSEC1_RX_ER	eTSEC1 receiver error	eTSEC1	1	I	15-2/15-8	USBDR_DIR/ TSEC_1588_TRIG2
TSEC1_TX_ER	eTSEC1 transmit error	eTSEC1	1	O	15-2/15-8	TSEC_1588_ALARM2
TSEC1_TX_CLK	eTSEC1 transmit clock in	eTSEC1	1	I	15-2/15-8	USBDR_CLK/ TSEC_1588_CLK
TSEC1_TXD3	eTSEC1 transmit data 3	eTSEC1	1	O	15-2/15-8	TSEC_1588_GCLK
TSEC1_TXD2	eTSEC1 transmit data 2	eTSEC1	1	O	15-2/15-8	TSEC_1588_PP1
TSEC1_TXD1	eTSEC1 transmit data 2	eTSEC1	1	O	15-2/15-8	TSEC_1588_PP2
TSEC1_TXD0	eTSEC1 transmit data 2	eTSEC1	1	O	15-2/15-8	USBDR_STP/ TSEC_1588_PP3
TSEC1_TX_EN	eTSEC1 transmit enable	eTSEC1	1	O	15-2/15-8	—

Table 3-1. MPC8313E Signal Reference by Functional Block (continued)

Name	Description	Functional Block	No. of Signals	I/O	Table/ Page	Alternate Function(s)
TSEC2_COL	eTSEC2 collision detect	eTSEC2	1	I/O	15-2/15-8	GTM1_TIN4/ GTM2_TIN3/ GPIO15
TSEC2_CRS	eTSEC2 carrier sense	eTSEC2	1	I/O	15-2/15-8	GTM1_TGATE4/ GTM2_TGATE3/ GPIO16
TSEC2_GTX_CLK	eTSEC2 transmit clock out	eTSEC2	1	I/O	15-2/15-8	$\overline{\text{GTM1\_TOUT4}}$ / $\overline{\text{GTM2\_TOUT3}}$ / GPIO17
TSEC2_RX_CLK	eTSEC2 receive clock	eTSEC2	1	I/O	15-2/15-8	GTM1_TIN2/ GTM2_TIN1/GPIO18
TSEC2_RX_DV	eTSEC2 receive data valid	eTSEC2	1	I/O	15-2/15-8	$\overline{\text{GTM1\_TGATE2}}$ / $\overline{\text{GTM2\_TGATE1}}$ / GPIO19
TSEC2_RXD[3:0]	eTSEC2 receive data 3–0	eTSEC2	4	I/O	15-2/15-8	GPIO[20:23]
TSEC2_RX_ER	eTSEC2 receiver error	eTSEC2	1	I/O	15-2/15-8	$\overline{\text{GTM1\_TOUT2}}$ / $\overline{\text{GTM2\_TOUT1}}$ / GPIO24
TSEC2_TXD[3:0]	eTSEC2 transmit data 3–0	eTSEC2	4	I/O	15-2/15-8	CFG_RESET_ SOURCE[0:3]
TSEC2_TX_ER	eTSEC2 transmit error	eTSEC2	1	I/O	15-2/15-8	GPIO27
TSEC2_TX_EN	eTSEC2 transmit enable	eTSEC2	1	I/O	15-2/15-8	GPIO26
TSEC2_TX_CLK	eTSEC2 transmit clock in	eTSEC2	1	I/O	15-2/15-8	GPIO25
TSEC_GTX_CLK125	Gigabit reference clock	eTSEC1/ eTSEC2	1	I	15-2/15-8	—
TSEC_MDC	Ethernet management data clock	Ethernet management	1	O	15-2/15-8	LB_POR_CFG_ BOOT_ECC_DIS <sup>1</sup>
TSEC_MDIO	Ethernet management data in/out	Ethernet management	1	I/O	15-2/15-8	—
TSEC_1588_CLK	1588 external timer reference clock input	eTSEC	1	I	15-2/15-8	UART_SOUT2/LA10
TSEC_1588_GCLK	1588 timers	eTSEC	1	O	15-2/15-8	UART_SIN2/LA11
TSEC_1588_PP1	1588 timer pulse per period 1	eTSEC	1	O	15-2/15-8	$\overline{\text{UART\_CTS2}}$ /LA12

Table 3-1. MPC8313E Signal Reference by Functional Block (continued)

Name	Description	Functional Block	No. of Signals	I/O	Table/ Page	Alternate Function(s)
TSEC_1588_PP2	1588 timer pulse per period 2	eTSEC	1	O	15-2/15-8	UART_RTS2/LA13
TSEC_1588_PP3	1588 timer pulse per period 3	eTSEC	1	O	15-2/15-8	GPIO14/LA9
TSEC_1588_TRIG1	1588 external timer trigger input 1	eTSEC	1	I	15-2/15-8	IIC1_SDA/LA14
TSEC_1588_TRIG2	1588 external timer trigger input 2	eTSEC	1	I	15-2/15-8	LA7
TSEC_1588_ALARM1	1588 timer; current time is equal to or greater than alarm time comparator register	eTSEC	1	O	15-2/15-8	LA8
TSEC_1588_ALARM2	1588 timer; current time is equal to or greater than alarm time comparator register	eTSEC	1	O	15-2/15-8	IIC1_SCL/LA15
RXA	Serial receiver, lane A, positive data	SGMII PHY	1	I	15-2/15-8	—
$\overline{\text{RXA}}$	Serial receiver, lane A, negative data (complement)	SGMII PHY	1	I	15-2/15-8	—
RXB	Serial receiver, lane B, positive data	SGMII PHY	1	I	15-2/15-8	—
$\overline{\text{RXB}}$	Serial receiver, lane B, negative data (complement)	SGMII PHY	1	I	15-2/15-8	—
SDAVDD	Analog supply for SerDes PLL	SGMII PHY <sup>2</sup>	1	I/O	15-2/15-8	—
SDAVSS	Analog ground for SerDes PLL	SGMII PHY	1	I/O	15-2/15-8	—
SD_IMP_CAL_RX	Receiver impedance control signal	SGMII PHY <sup>3</sup>	1	I	15-2/15-8	—
SD_IMP_CAL_TX	Transmitter impedance control signal	SGMII PHY <sup>3</sup>	1	I	15-2/15-8	—
SD_PLL_TPA_ANA	Analog test point for SerDes PLL testing	SGMII PHY	1	O	15-2/15-8	—
SD_PLL_TPD	Digital test point for SerDes PLL testing	SGMII PHY	1	O	15-2/15-8	—

Table 3-1. MPC8313E Signal Reference by Functional Block (continued)

Name	Description	Functional Block	No. of Signals	I/O	Table/ Page	Alternate Function(s)
SD_REF_CLK	SerDes PLL reference clock	SGMII PHY	1	I	15-2/15-8	—
$\overline{\text{SD\_REF\_CLK}}$	SerDes PLL reference clock (complement)	SGMII PHY	1	I	15-2/15-8	—
TXA	Serial transmitter, lane A, positive data	SGMII PHY	1	O	15-2/15-8	—
$\overline{\text{TXA}}$	Serial transmitter, lane A, negative data (complement)	SGMII PHY	1	O	15-2/15-8	—
TXB	Serial transmitter, lane B, positive data	SGMII PHY	1	O	15-2/15-8	—
$\overline{\text{TXB}}$	Serial transmitter, lane B, negative data (complement)	SGMII PHY	1	O	15-2/15-8	—
XCOREVDD[0:2]	SerDes transceiver core supply	SGMII PHY <sup>2</sup>	3	PWR	15-2/15-8	—
XCOREVSS[0:2]	SerDes transceiver core ground	SGMII PHY	3	GND	15-2/15-8	—
XPADVDD[0:1]	SerDes transceiver pad supply	SGMII PHY <sup>2</sup>	2	PWR	15-2/15-8	—
XPADVSS[0:1]	SerDes transceiver pad ground	SGMII PHY	2	GND	15-2/15-8	—
LAD[0:15]	LBC address/data	eLBC	16	I/O	10-2/10-5	—
LA[16:25]	LBC port address	eLBC	10	O	10-2/10-5	—
LA[0:4]	LBC port address	eLBC	5	I/O	10-2/10-5	MSRCID[0:4]/ GPIO[0:4]
LA[5]	LBC port address	eLBC	1	I/O	10-2/10-5	MDVAL/GPIO5
LA6	LBC port address	eLBC	1	I/O	10-2/10-5	GPIO6
LA7	LBC port address	eLBC	1	I/O	10-2/10-5	GPIO7/ TSEC_1588_TRIG2
LA8	LBC port address	eLBC	1	I/O	10-2/10-5	GPIO13/ TSEC_1588_ALARM1
LA9	LBC port address	eLBC	1	I/O	10-2/10-5	GPIO14/ TSEC_1588_PP3
LA10	LBC port address	eLBC	1	O	10-2/10-5	TSEC_1588_CLK
LA11	LBC port address	eLBC	1	O	10-2/10-5	TSEC_1588_GCLK
LA12	LBC port address	eLBC	1	O	10-2/10-5	TSEC_1588_PP1
LA13	eLBC port address	eLBC	1	O	10-2/10-5	TSEC_1588_PP2

Table 3-1. MPC8313E Signal Reference by Functional Block (continued)

Name	Description	Functional Block	No. of Signals	I/O	Table/ Page	Alternate Function(s)
LA14	eLBC port address	eLBC	1	O	10-2/10-5	TSEC_1588_TRIG1
LA15	eLBC port address	eLBC	1	O	10-2/10-5	TSEC_1588_ALARM2
$\overline{\text{LCS}}[0:3]$	eLBC chip select 0–3	eLBC	4	O	10-2/10-5	—
$\overline{\text{LWE}}0/\overline{\text{LFWE}}/\overline{\text{LBS}}0$	eLBC write enable 0/FCM write enable/UPM byte (lane) select 0	eLBC	1	O	10-2/10-5	—
$\overline{\text{LWE}}1/\overline{\text{LBS}}1$	eLBC write enable 1/UPM byte (lane) select 1	eLBC	1	O	10-2/10-5	—
LBCTL	eLBC data buffer control	eLBC	1	O	10-2/10-5	—
LALE	eLBC address latch enable	eLBC	1	O	10-2/10-5	M1LALE
LGPL0/LFCLE	eLBC UPM general purpose line 0/Flash command latch enable	eLBC	1	O	10-2/10-5	—
LGPL1/LFALE	eLBC GP line 1/Flash address latch enable	eLBC	1	O	10-2/10-5	—
LGPL2/ $\overline{\text{LOE}}/\overline{\text{LFRE}}$	eLBC output enable/GP line 2/FCM read enable	eLBC	1	O	10-2/10-5	—
LGPL3/ $\overline{\text{LFWP}}$	eLBC GP line 3/Flash write project	eLBC	1	O	10-2/10-5	—
LGPL4/ $\overline{\text{LGTA}}/\overline{\text{LUPWAIT}}/\overline{\text{LFRB}}$	eLBC GP line 4/GPCM terminate access/UPM wait/Flash read/ $\overline{\text{busy}}$ , open-drain shared pin	eLBC	1	I/O	10-2/10-5	—
LGPL5	eLBC GP line 5	eLBC	1	O	10-2/10-5	—
LCLK[0:1]	eLBC clocks 0–1	eLBC	2	O	10-2/10-5	—
LB_POR_CFG_BOOT_ECC_DIS	Boot time ECC checking	eLBC	1	1	—	TSEC_MDC
USBDR_DRIVE_VBUS	USB VBus power enable	USB	1	I/O	16-1/16-3	GTM1_TIN1/ GTM2_TIN2/ LSRCID0



Table 3-1. MPC8313E Signal Reference by Functional Block (continued)

Name	Description	Functional Block	No. of Signals	I/O	Table/ Page	Alternate Function(s)
USBDR_PWRFAULT	USB VBus fault	USB	1	I/O	16-1/16-3	GTM1_TGATE1/ GTM2_TGATE2/ LSRCID1
USBDR_PCTL0	USB status LED 0 control	USB	1	O	16-1/16-3	GTM1_TOUT1/ LSRCID2
USBDR_PCTL1	USB status LED 0 control	USB	1	O	16-1/16-3	LSRCID3/ LBC_PM_REF_10
USBDR_TXDRXD0	USB data bus 0	USB	1	O	16-1/16-3	TSEC1_COL
USBDR_TXDRXD1	USB data bus 1	USB	1	O	16-1/16-3	TSEC1_CRS
USBDR_TXDRXD2	USB data bus 2	USB	1	O	16-1/16-3	TSEC1_GTX_CLK
USBDR_TXDRXD3	USB data bus 3	USB	1	O	16-1/16-3	TSEC1_RX_CLK
USBDR_TXDRXD4	USB data bus 4	USB	1	O	16-1/16-3	TSEC1_RX_DV
USBDR_TXDRXD[5:7]	USB data bus 5-7	USB	1	O	16-1/16-3	TSEC1_RXD[3:1]
USBDR_NXT	USB next data	USB	1	O	16-1/16-3	TSEC1_RXD0
USBDR_DIR	USB direction	USB	1	O	16-1/16-3	TSEC1_RX_ER
USBDR_CLK	USB clock	USB	1	O	16-1/16-3	TSEC1_TX_CLK
USBDR_STP	USB stop	USB	1	O	16-1/16-3	TSEC1_TXD0/ TSEC_1588_PP3
USB_DP	USB 2.0 D+ line	USB PHY	1	I/O	16-1/16-3	—
USB_DM	USB 2.0 D– line	USB PHY	1	I/O	16-1/16-3	—
USB_RBIAS	USB bias input <sup>4</sup>	USB PHY	1	I/O	16-1/16-3	—
USB_VBUS	USB 2.0 VBUS line	USB PHY	1	I/O	16-1/16-3	—
USB_TPA	Dedicated analog test signal	USB PHY	1	I/O	16-1/16-3	—
USB_VDDA	Dedicated power for USB transceiver	USB PHY	1	I/O	16-1/16-3	—
USB_VSSA	Dedicated ground for USB transceiver	USB PHY	1	I/O	16-1/16-3	—
USB_PLL_PWR1	Dedicated 1.0 V analog power for USB PLL	USB PHY	1	I/O	16-1/16-3	—
USB_PLL_PWR3	Dedicated 3.3 V analog power for USB PLL	USB PHY	1	I/O	16-1/16-3	—
USB_PLL_GND	Dedicated analog ground for USB PLL	USB PHY	1	I/O	16-1/16-3	—
USB_VSSA_BIAS	Dedicated power for USB bias circuit	USB PHY	1	I/O	16-1/16-3	—

Table 3-1. MPC8313E Signal Reference by Functional Block (continued)

Name	Description	Functional Block	No. of Signals	I/O	Table/ Page	Alternate Function(s)
USB_VDDA_BIAS	Dedicated ground for USB bias circuit	USB PHY	1	I/O	16-1/16-3	—
GTM1_TIN3/ GTM2_TIN4	Timer in 3/4	Global Timers	1	I/O	5-54/5-52	GPIO28/LSRCID4 SPIMOSI
GTM1_TGATE3/ GTM2_TGATE4	Timer gate 3/4	Global Timers	1	I/O	5-54/5-52	LDVAL/GPIO29/ SPIMISO
GTM1_TIN4/ GTM2_TIN3	Timer in 4/3	Global Timers	1	I/O	5-54/5-52	GPIO15/TSEC2_COL
GTM1_TGATE4/ GTM2_TGATE3	Timer gate 4/3	Global Timers	1	I/O	5-54/5-52	GPIO16/TSEC2_CRS
GTM1_TOUT3	Timer out 3	Global Timers	1	I/O	5-54/5-52	GPIO30/SPICLK
GTM1_TOUT4/ GTM2_TOUT3	Timer out 4/3	Global Timers	1	I/O	5-54/5-52	GPIO17/ TSEC2_GTX_CLK
GTM1_TIN2/ GTM2_TIN1	Timer in 2/1	Global Timers	1	I/O	5-54/5-52	GPIO18/ TSEC2_Rx_CLK
GTM1_TGATE2/ GTM2_TGATE1	Timer gate 2/1	Global Timers	1	I/O	5-54/5-52	GPIO19/ TSEC2_RX_DV
GTM1_TOUT2/ GTM2_TOUT1	Timer out 2/1	Global Timers	1	I/O	5-54/5-52	GPIO24/ TSEC2_RX_ER
GTM1_TIN1/ GTM2_TIN2	Timer in 1/2	Global Timers	1	I/O	5-54/5-52	LSRCID0/USBDR_ DRIVE_VBUS
GTM1_TGATE1/ GTM2_TGATE2	Timer gate 1/2	Global Timers	1	I/O	5-54/5-52	LSRCID1/USBDR_ PWRFAULT
GTM1_TOUT1	Timer out 1	Global Timers	1	I/O	5-54/5-52	LSRCID2/ USBDR_PCTLO
MCP_OUT	Machine check interrupt output	IPIC	1	O	8-2/8-5	—
IRQ0/MCP_IN	External interrupt 0	IPIC	1	I	8-2/8-5	—
IRQ1	External interrupt 1	IPIC	1	I	8-2/8-5	—
IRQ2	External interrupt 2	IPIC	1	I	8-2/8-5	—
IRQ3	External interrupt 3	IPIC	1	I/O	8-2/8-5	CKSTOP_OUT
IRQ4	External interrupt 4	IPIC	1	I/O	8-2/8-5	CKSTOP_IN/ GPIO12
TCK	Test clock	JTAG	1	I	20-2/20-2	—
TDI	Test data in	JTAG	1	I	20-2/20-2	—
TDO	Test data out	JTAG	1	O	20-2/20-2	—
TMS	Test mode select	JTAG	1	I	20-2/20-2	—

Table 3-1. MPC8313E Signal Reference by Functional Block (continued)

Name	Description	Functional Block	No. of Signals	I/O	Table/ Page	Alternate Function(s)
TRST	Test reset	JTAG	1	I	20-2/20-2	—
GPIO[0:4]	General-purpose I/O signals 0-4	GPIO	5	I/O	—	LA[0:4]/MSRCID[0:4]
GPIO5	General-purpose I/O signal 5	GPIO	1	I/O	—	LA5/MDVAL
GPIO6	General-purpose I/O signal 6	GPIO	1	I/O	—	LA6
GPIO7	General-purpose I/O signal 7	GPIO	1	I/O	—	LA7/ TSEC_1588_TRIG2
GPIO8	General-purpose I/O signal 8	GPIO	1	I/O	—	UART_CTS1/ MSRCID2
GPIO9	General-purpose I/O signal 9	GPIO	1	I/O	—	UART_RTS1/ MSRCID3
GPIO10	General-purpose I/O signal 10	GPIO	1	I/O	—	IIC2_SDA/ PMC_PWR_OK
GPIO11	General-purpose I/O signal 11	GPIO	1	I/O	—	IIC2_SCL
GPIO12	General-purpose I/O signal 12	GPIO	1	I/O	—	IRQ4/CKSTOP_IN
GPIO13	General-purpose I/O signal 13	GPIO	1	I/O	—	LA8/ TSEC_1588_ALARM1
GPIO14	General-purpose I/O signal 14	GPIO	1	I/O	—	LA9/ TSEC_1588_PP3
GPIO15	General-purpose I/O signal 15	GPIO	1	I/O	—	GTM1_TIN4/ GTM2_TIN3/ TSEC2_COL
GPIO16	General-purpose I/O signal 16	GPIO	1	I/O	—	GTM1_TGATE4/ GTM2_TGATE3/ TSEC2_CRS
GPIO17	General-purpose I/O signal 17	GPIO	1	I/O	—	GTM1_TOUT4/ GTM2_TOUT3/ TSEC2_GTX_CLK
GPIO18	General-purpose I/O signal 18	GPIO	1	I/O	—	GTM1_TIN2/ GTM2_TIN1/ TSEC2_RX_CLK
GPIO19	General-purpose I/O signal 19	GPIO	1	I/O	—	GTM1_TGATE2/ GTM2_TGATE1/ TSEC2_RX_DV
GPIO[20:23]	General-purpose I/O signal 20-23	GPIO	1	I/O	—	TSEC2_RXD[3:0]

Table 3-1. MPC8313E Signal Reference by Functional Block (continued)

Name	Description	Functional Block	No. of Signals	I/O	Table/ Page	Alternate Function(s)
GPIO24	General-purpose I/O signal 24	GPIO	1	I/O	—	$\overline{\text{GTM1\_TOUT2}}$ / $\overline{\text{GTM2\_TOUT1}}$ / $\overline{\text{TSEC2\_RX\_ER}}$
GPIO25	General-purpose I/O signal 25	GPIO	1	I/O	—	TSEC2_TX_CLK
GPIO26	General-purpose I/O signal 26	GPIO	1	I/O	—	TSEC2_TX_EN
GPIO27	General-purpose I/O signal 27	GPIO	1	I/O	—	TSEC2_TX_ER
GPIO28	General-purpose I/O signal 28	GPIO	1	I/O	—	GTM1_TIN3/ GTM2_TIN4/ LSRCID4/SPI MOSI
GPIO29	General-purpose I/O signal 29	GPIO	1	I/O	—	$\overline{\text{GTM1\_TGATE3}}$ / $\overline{\text{GTM2\_TGATE4}}$ / LDVAL/SPI MISO
GPIO30	General-purpose I/O signal 30	GPIO	1	I/O	—	$\overline{\text{GTM1\_TOUT3}}$ / SPICLK
GPIO31	General-purpose I/O signal 31	GPIO	1	I/O	—	SPIS $\overline{\text{E}}$ L
THERM[0:1]	Thermal resistor access	TEST	2	I	—	—
TEST_MODE	Test mode	TEST	1	I	<a href="#">20-2/20-2</a>	—
EXT_PWR_CTRL	External power control	PMC	1	O	<a href="#">5-65/5-66</a>	—
PMC_PWR_OK	Stable power	PMC	1	I	<a href="#">5-65/5-66</a>	IIC2_SDA/GPIO10
$\overline{\text{QUIESCE}}$	Quiesce state	PMC	1	O	<a href="#">5-65/5-66</a>	—
$\overline{\text{PORESET}}$	Power on reset	System control	1	I	<a href="#">4-1/4-1</a>	—
$\overline{\text{HRESET}}$	Hard reset	System control	1	I/O	<a href="#">4-1/4-1</a>	—
$\overline{\text{SRESET}}$	Soft reset to the e300 core	System control	1	I	<a href="#">4-1/4-1</a>	—
$\overline{\text{CKSTOP\_OUT}}$	Clock stop out	Reset and clock	1	I/O	<a href="#">4-3/4-5</a>	IIC1_SDA/ TSEC_1588_TRIG1
$\overline{\text{CKSTOP\_IN}}$	Clock stop in	Reset and clock	1	I/O	<a href="#">4-3/4-5</a>	IIC1_SCL/ TSEC_1588_ALARM2
$\overline{\text{CFG\_CLKIN\_DIV}}$	Configuration clock in division selection	Reset and clock	1	I	<a href="#">3-3/3-29</a>	—

Table 3-1. MPC8313E Signal Reference by Functional Block (continued)

Name	Description	Functional Block	No. of Signals	I/O	Table/ Page	Alternate Function(s)
CFG_LBIU_MUX_EN	Selects multiplexing for address/data bus of the eLBC block	Reset and clock	1	I	3-6/3-31	—
SYS_CR_CLK_IN	Crystal clock input	Clocks	1	I	4-2/4-3	—
SYS_CR_CLK_OUT	Crystal clock output	Clocks	1	O	4-2/4-3	—
SYS_CLK_IN	Clock input	Clocks	1	I	4-2/4-3	—
USB_CR_CLK_IN	USB crystal clock input	Clocks	1	I	4-2/4-3	—
USB_CR_CLK_OUT	USB crystal clock output	Clocks	1	O	4-2/4-3	—
USB_CLK_IN	USB clock input	Clocks	1	I	4-2/4-3	—
PCI_SYNC_IN	PCI clock/PCI clock sync input	Clocks	1	I	4-2/4-3	—
PCI_SYNC_OUT	PCI clock sync output	Clocks	1	O	4-2/4-3	—
RTC_PIT_CLOCK	Timer clock/Real time clock	PIT/RTC	1	I	4-2/4-3	—
MSRCID[0:4]	Memory debug source ID 0–4	Debug	5	I/O	10-2/10-5	LA[0:4]/GPIO[0:4]
MSRCID0	Memory debug source ID 0	Debug	1	I/O	10-2/10-5	UART_SOUT1
MSRCID1	Memory debug source ID 1	Debug	1	I/O	10-2/10-5	UARTSIN1
MSRCID2	Memory debug source ID 2	Debug	1	I/O	10-2/10-5	UART_CTS1/GPIO8
MSRCID3	Memory debug source ID 3	Debug	1	I/O	10-2/10-5	UART_RTS1/GPIO9
MSRCID4	Memory debug source ID 4	Debug	1	I/O	10-2/10-5	UART_SOUT2/ TSEC_1588_CLK
MDVAL	Memory debug data valid	Debug	1	I/O	10-2/10-5	LA5/GPIO5
MDVAL	Memory debug data valid	Debug	1	I/O	10-2/10-5	UART_SIN2/ TSEC_1588_GCLK
LSRCID0	Memory debug source ID 0	Debug	1	I/O	10-2/10-5	GTM1_TIN1/ GTM2_TIN2/USBDR _DRIVE_VBUS
LSRCID1	Memory debug source ID 1	Debug	1	I/O	10-2/10-5	GTM1_TGATE1/ GTM2_TGATE2/ USBDR_DRIVE_ VBUS

Table 3-1. MPC8313E Signal Reference by Functional Block (continued)

Name	Description	Functional Block	No. of Signals	I/O	Table/ Page	Alternate Function(s)
LSRCID2	Memory debug source ID 2	Debug	1	I/O	10-2/10-5	GTM1_TOUT1/ USBDR_PCTL0
LSRCID3	Memory debug source ID 3	Debug	1	I/O	10-2/10-5	LBC_PM_REF_10/ USBDR_PCTL1
LSRCID4	Memory debug source ID 4	Debug	1	I/O	10-2/10-5	GTM1_TIN3/ GTM2_TIN4/ GPIO28/SPIMOSI
LDVAL	Memory debug data valid	Debug	1	I/O	10-2/10-5	GTM1_TGATE3/ GTM2_TGATE4/ GPIO29/SPIMISO

<sup>1</sup> The LB\_POR\_CFG\_BOOT\_ECC\_DIS function will be selected on the TSEC\_MDC pin whenever HRESET is asserted; the pin will act as TSEC\_MDC at all other times. The reset block will sample this signal on PORESET negation only; the sampled value is then passed to the eLBC controller to enable/disable ECC checking during boot time. An internal pull-down resistor has been added to this pad to enable the boot-time ECC checking by default. A pull-up resistor, via a three-state buffer, is needed during HRESET assertion period to disable ECC checking during boot time.

<sup>2</sup> See Table 3-3 for proper connection to power.

<sup>3</sup> See the MPC8313E Hardware Specification for resistor values.

<sup>4</sup> Must be connected to a 10K ±1% precision resistor if using the integrated USB PHY through the UTMI.

Table 3-2 lists the signals in alphabetical order.

Table 3-2. MPC8313E Signal Reference by Signal Name

Name	Description	Functional Block	No. of Signals	I/O	Table/ Page	Alternate Function(s)
CFG_CLKIN_DIV	Configuration clock in division selection	Reset and clock	1	I	3-3/3-29	—
CFG_LBIU_MUX_EN	Selects multiplexing for address/data bus of the eLBC block	Reset and clock	1	I	3-6/3-31	—
CKSTOP_IN	Clock stop in	Reset and clock	1	I/O	4-3/4-5	IIC1_SCL/ TSEC_1588_ALARM2
CKSTOP_OUT	Clock stop out	Reset and clock	1	I/O	4-3/4-5	IIC1_SDA/ TSEC_1588_TRIG1
CPCI_HS_ENUM	CompactPCI hot swap enumerator	PCI	1	O	13-3/13-5	PCI_GNT2
CPCI_HS_ES	CompactPCI hot swap ejector switch	PCI	1	I	13-3/13-5	PCI_REQ1
CPCI_HS_LED	CompactPCI hot swap LED	PCI	1	O	13-3/13-5	PCI_GNT1
EXT_PWR_CTRL	External power control	PMC	1	O	5-65/5-66	—

Table 3-2. MPC8313E Signal Reference by Signal Name (continued)

Name	Description	Functional Block	No. of Signals	I/O	Table/ Page	Alternate Function(s)
GPIO[0:4]	General-purpose I/O signals 0-4	GPIO	5	I/O	—	LA[0:4]/MSRCID[0:4]
GPIO[20:23]	General-purpose I/O signal 20-23	GPIO	1	I/O	—	TSEC2_RXD[3:0]
GPIO10	General-purpose I/O signal 10	GPIO	1	I/O	—	IIC2_SDA/ PMC_PWR_OK
GPIO11	General-purpose I/O signal 11	GPIO	1	I/O	—	IIC2_SCL
GPIO12	General-purpose I/O signal 12	GPIO	1	I/O	—	IRQ4/CKSTOP_IN
GPIO13	General-purpose I/O signal 13	GPIO	1	I/O	—	LA8/ TSEC_1588_ALARM1
GPIO14	General-purpose I/O signal 14	GPIO	1	I/O	—	LA9/ TSEC_1588_PP3
GPIO15	General-purpose I/O signal 15	GPIO	1	I/O	—	GTM1_TIN4/ GTM2_TIN3/ TSEC2_COL
GPIO16	General-purpose I/O signal 16	GPIO	1	I/O	—	GTM1_TGATE4/ GTM2_TGATE3/ TSEC2_CRS
GPIO17	General-purpose I/O signal 17	GPIO	1	I/O	—	GTM1_TOUT4/ GTM2_TOUT3/ TSEC2_GTX_CLK
GPIO18	General-purpose I/O signal 18	GPIO	1	I/O	—	GTM1_TIN2/ GTM2_TIN1/ TSEC2_RX_CLK
GPIO19	General-purpose I/O signal 19	GPIO	1	I/O	—	GTM1_TGATE2/ GTM2_TGATE1/ TSEC2_RX_DV
GPIO24	General-purpose I/O signal 24	GPIO	1	I/O	—	GTM1_TOUT2/ GTM2_TOUT1/ TSEC2_RX_ER
GPIO25	General-purpose I/O signal 25	GPIO	1	I/O	—	TSEC2_TX_CLK
GPIO26	General-purpose I/O signal 26	GPIO	1	I/O	—	TSEC2_TX_EN
GPIO27	General-purpose I/O signal 27	GPIO	1	I/O	—	TSEC2_TX_ER
GPIO28	General-purpose I/O signal 28	GPIO	1	I/O	—	GTM1_TIN3/ GTM2_TIN4/ LSRCID4/SPI MOSI

Table 3-2. MPC8313E Signal Reference by Signal Name (continued)

Name	Description	Functional Block	No. of Signals	I/O	Table/ Page	Alternate Function(s)
GPIO29	General-purpose I/O signal 29	GPIO	1	I/O	—	$\overline{\text{GTM1\_TGATE3}}$ / $\overline{\text{GTM2\_TGATE4}}$ / LDVAL/SPIMSO
GPIO30	General-purpose I/O signal 30	GPIO	1	I/O	—	$\overline{\text{GTM1\_TOUT3}}$ / SPICLK
GPIO31	General-purpose I/O signal 31	GPIO	1	I/O	—	$\overline{\text{SPISEL}}$
GPIO5	General-purpose I/O signal 5	GPIO	1	I/O	—	LA5/MDVAL
GPIO6	General-purpose I/O signal 6	GPIO	1	I/O	—	LA6
GPIO7	General-purpose I/O signal 7	GPIO	1	I/O	—	LA7/ TSEC_1588_TRIG2
GPIO8	General-purpose I/O signal 8	GPIO	1	I/O	—	$\overline{\text{UART\_CTS1}}$ / MSRCID2
GPIO9	General-purpose I/O signal 9	GPIO	1	I/O	—	$\overline{\text{UART\_RTS1}}$ / MSRCID3
$\overline{\text{GTM1\_TGATE1}}$ / $\overline{\text{GTM2\_TGATE2}}$	Timer gate 1/2	Global Timers	1	I/O	5-54/5-52	LSRCID1/USBDR_ PWRFAULT
$\overline{\text{GTM1\_TGATE2}}$ / $\overline{\text{GTM2\_TGATE1}}$	Timer gate 2/1	Global Timers	1	I/O	5-54/5-52	GPIO19/ TSEC2_RX_DV
$\overline{\text{GTM1\_TGATE3}}$ / $\overline{\text{GTM2\_TGATE4}}$	Timer gate 3/4	Global Timers	1	I/O	5-54/5-52	LDVAL/GPIO29/ SPIMISO
$\overline{\text{GTM1\_TGATE4}}$ / $\overline{\text{GTM2\_TGATE3}}$	Timer gate 4/3	Global Timers	1	I/O	5-54/5-52	GPIO16/TSEC2_CRS
$\overline{\text{GTM1\_TIN1}}$ / $\overline{\text{GTM2\_TIN2}}$	Timer in 1/2	Global Timers	1	I/O	5-54/5-52	LSRCID0/USBDR_ DRIVE_VBUS
$\overline{\text{GTM1\_TIN2}}$ / $\overline{\text{GTM2\_TIN1}}$	Timer in 2/1	Global Timers	1	I/O	5-54/5-52	GPIO18/ $\overline{\text{TSEC2\_Rx\_CLK}}$
$\overline{\text{GTM1\_TIN3}}$ / $\overline{\text{GTM2\_TIN4}}$	Timer in 3/4	Global Timers	1	I/O	5-54/5-52	GPIO28/LSRCID4 SPIMOSI
$\overline{\text{GTM1\_TIN4}}$ / $\overline{\text{GTM2\_TIN3}}$	Timer in 4/3	Global Timers	1	I/O	5-54/5-52	GPIO15/TSEC2_COL
$\overline{\text{GTM1\_TOUT1}}$	Timer out 1	Global Timers	1	I/O	5-54/5-52	LSRCID2/ USBDR_PCTL0
$\overline{\text{GTM1\_TOUT2}}$ / $\overline{\text{GTM2\_TOUT1}}$	Timer out 2/1	Global Timers	1	I/O	5-54/5-52	GPIO24/ TSEC2_RX_ER
$\overline{\text{GTM1\_TOUT3}}$	Timer out 3	Global Timers	1	I/O	5-54/5-52	GPIO30/SPICLK



Table 3-2. MPC8313E Signal Reference by Signal Name (continued)

Name	Description	Functional Block	No. of Signals	I/O	Table/ Page	Alternate Function(s)
GTM1_TOUT4/ GTM2_TOUT3	Timer out 4/3	Global Timers	1	I/O	5-54/5-52	GPIO17/ TSEC2_GTX_CLK
HRESET	Hard reset	System control	1	I/O	4-1/4-1	—
IIC1_SCL	I <sup>2</sup> C serial clock	I <sup>2</sup> C1	1	I/O	17-1/17-3	CKSTOP_IN/ TSEC_1588_ALARM2
IIC1_SDA	I <sup>2</sup> C serial data	I <sup>2</sup> C1	1	I/O	17-1/17-3	CKSTOP_OUT/ TSEC_1588_TRIG1
IIC2_SCL	I <sup>2</sup> C serial clock	I <sup>2</sup> C2	1	I/O	17-1/17-3	GPIO11
IIC2_SDA	I <sup>2</sup> C serial data	I <sup>2</sup> C2	1	I/O	17-1/17-3	PMC_PWR_OK GPIO10
IRQ0/MCP_IN	External interrupt 0	IPIC	1	I	8-2/8-5	—
IRQ1	External interrupt 1	IPIC	1	I	8-2/8-5	—
IRQ2	External interrupt 2	IPIC	1	I	8-2/8-5	—
IRQ3	External interrupt 3	IPIC	1	I/O	8-2/8-5	CKSTOP_OUT
IRQ4	External interrupt 4	IPIC	1	I/O	8-2/8-5	CKSTOP_IN/ GPIO12
LA[0:4]	LBC port address	eLBC	5	I/O	10-2/10-5	MSRCID[0:4]/ GPIO[0:4]
LA[16:25]	LBC port address	eLBC	10	O	10-2/10-5	—
LA[5]	LBC port address	eLBC	1	I/O	10-2/10-5	MDVAL/GPIO5
LA10	LBC port address	eLBC	1	O	10-2/10-5	TSEC_1588_CLK
LA11	LBC port address	eLBC	1	O	10-2/10-5	TSEC_1588_GCLK
LA12	LBC port address	eLBC	1	O	10-2/10-5	TSEC_1588_PP1
LA13	eLBC port address	eLBC	1	O	10-2/10-5	TSEC_1588_PP2
LA14	eLBC port address	eLBC	1	O	10-2/10-5	TSEC_1588_TRIG1
LA15	eLBC port address	eLBC	1	O	10-2/10-5	TSEC_1588_ALARM2
LA6	LBC port address	eLBC	1	I/O	10-2/10-5	GPIO6
LA7	LBC port address	eLBC	1	I/O	10-2/10-5	GPIO7/ TSEC_1588_TRIG2
LA8	LBC port address	eLBC	1	I/O	10-2/10-5	GPIO13/ TSEC_1588_ALARM1
LA9	LBC port address	eLBC	1	I/O	10-2/10-5	GPIO14/ TSEC_1588_PP3
LAD[0:15]	LBC address/data	eLBC	16	I/O	10-2/10-5	—

Table 3-2. MPC8313E Signal Reference by Signal Name (continued)

Name	Description	Functional Block	No. of Signals	I/O	Table/ Page	Alternate Function(s)
LALE	eLBC address latch enable	eLBC	1	O	10-2/10-5	M1LALE
LB_POR_CFG_BOOT_ECC_DIS	Boot time ECC checking	eLBC	1	1	—	TSEC_MDC
LBCTL	eLBC data buffer control	eLBC	1	O	10-2/10-5	—
LCLK[0:1]	eLBC clocks 0–1	eLBC	2	O	10-2/10-5	—
LCS[0:3]	eLBC chip select 0–3	eLBC	4	O	10-2/10-5	—
LDVAL	Memory debug data valid	Debug	1	I/O	10-2/10-5	GTM1_TGATE3/ GTM2_TGATE4/ GPIO29/SPIMISO
LGPL0/LFCLE	eLBC UPM general purpose line 0/Flash command latch enable	eLBC	1	O	10-2/10-5	—
LGPL1/LFALE	eLBC GP line 1/Flash address latch enable	eLBC	1	O	10-2/10-5	—
LGPL2/LOE/LFRE	eLBC output enable/GP line 2/FCM read enable	eLBC	1	O	10-2/10-5	—
LGPL3/LFWP	eLBC GP line 3/Flash write project	eLBC	1	O	10-2/10-5	—
LGPL4/LGTA/ LUPWAIT/LFRB	eLBC GP line 4/GPCM terminate access/UPM wait/Flash read/busy, open-drain shared pin	eLBC	1	I/O	10-2/10-5	—
LGPL5	eLBC GP line 5	eLBC	1	O	10-2/10-5	—
LSRCID0	Memory debug source ID 0	Debug	1	I/O	10-2/10-5	GTM1_TIN1/ GTM2_TIN2/USBDR_DRIVE_VBUS
LSRCID1	Memory debug source ID 1	Debug	1	I/O	10-2/10-5	GTM1_TGATE1/ GTM2_TGATE2/ USBDR_DRIVE_VBUS
LSRCID2	Memory debug source ID 2	Debug	1	I/O	10-2/10-5	GTM1_TOUT1/ USBDR_PCTL0
LSRCID3	Memory debug source ID 3	Debug	1	I/O	10-2/10-5	LBC_PM_REF_10/ USBDR_PCTL1

Table 3-2. MPC8313E Signal Reference by Signal Name (continued)

Name	Description	Functional Block	No. of Signals	I/O	Table/ Page	Alternate Function(s)
LSRCID4	Memory debug source ID 4	Debug	1	I/O	10-2/10-5	GTM1_TIN3/ GTM2_TIN4/ GPIO28/SPIMOSI
$\overline{\text{LWE0/LFWE/LBS0}}$	eLBC write enable 0/FCM write enable/UPM byte (lane) select 0	eLBC	1	O	10-2/10-5	—
$\overline{\text{LWE1/LBS1}}$	eLBC write enable 1/UPM byte (lane) select 1	eLBC	1	O	10-2/10-5	—
M66EN	66-MHz system configuration	PCI	1	I	13-3/13-5	—
MA[0:14]	DDR address	DDR	15	O	9-3/9-5	—
MBA[0:2]	DDR bank select	DDR	3	O	9-3/9-5	—
$\overline{\text{MCAS}}$	DDR column address strobe	DDR	1	O	9-3/9-5	—
MCK, $\overline{\text{MCK}}$	DDR differential clocks	DDR	2	O	9-4/9-7	—
MCKE	DDR clock enable	DDR	1	O	9-4/9-7	—
$\overline{\text{MCP\_OUT}}$	Machine check interrupt output	IPIC	1	O	8-2/8-5	—
$\overline{\text{MCS}}[0:1]$	DDR chip select (2/DIMM)	DDR	2	O	9-3/9-5	—
MDM[0:3]	DDR data mask	DDR	4	O	9-3/9-5	—
MDQ[0:31]	DDR data	DDR	32	I/O	9-3/9-5	—
MDQS[0:3]	DDR data strobe	DDR	4	I/O	9-3/9-5	—
MDVAL	Memory debug data valid	Debug	1	I/O	10-2/10-5	LA5/GPIO5
MDVAL	Memory debug data valid	Debug	1	I/O	10-2/10-5	UART_SIN2/ TSEC_1588_GCLK
MODT[0:1]	DRAM on-die termination	DDR	2	O	9-3/9-5	—
$\overline{\text{MRAS}}$	DDR row address strobe	DDR	1	O	9-3/9-5	—
MSRCID[0:4]	Memory debug source ID 0–4	Debug	5	I/O	10-2/10-5	LA[0:4]/GPIO[0:4]
MSRCID0	Memory debug source ID 0	Debug	1	I/O	10-2/10-5	UART_SOUT1

Table 3-2. MPC8313E Signal Reference by Signal Name (continued)

Name	Description	Functional Block	No. of Signals	I/O	Table/ Page	Alternate Function(s)
MSRCID1	Memory debug source ID 1	Debug	1	I/O	10-2/10-5	UARTSIN1
MSRCID2	Memory debug source ID 2	Debug	1	I/O	10-2/10-5	UART_CTS1/GPIO8
MSRCID3	Memory debug source ID 3	Debug	1	I/O	10-2/10-5	UART_RTS1/GPIO9
MSRCID4	Memory debug source ID 4	Debug	1	I/O	10-2/10-5	UART_SOUT2/ TSEC_1588_CLK
MWE	DDR write enable	DDR	1	O	9-3/9-5	—
PCI_AD[0:31]	PCI address/data	PCI	32	I/O	13-3/13-5	—
PCI_C_BE[0:3]	PCI command/byte enable	PCI	4	I/O	13-3/13-5	—
PCI_CLK0	PCI clock 0	PCI	1	O	4-2/4-3	—
PCI_CLK1	PCI clock 1	PCI	1	O	4-2/4-3	—
PCI_CLK2	PCI clock 2	PCI	1	O	4-2/4-3	—
PCI_DEVSEL	PCI device select	PCI	1	I/O	13-3/13-5	—
PCI_FRAME	PCI frame	PCI	1	I/O	13-3/13-5	—
PCI_GNT0	PCI grant 0	PCI	1	I/O	13-3/13-5	—
PCI_GNT1	PCI grant 1	PCI	1	O	13-3/13-5	CPCI_HS_LED
PCI_GNT2	PCI grant 2	PCI	1	O	13-3/13-5	CPCI_HS_ENUM
PCI_IDSEL	PCI initial device select	PCI	1	I	13-3/13-5	—
PCI_INTA	PCI interrupt output	PCI	1	O	13-3/13-5	—
PCI_IRDY	PCI initiator ready	PCI	1	I/O	13-3/13-5	—
PCI_PAR	PCI parity	PCI	1	I/O	13-3/13-5	—
PCI_PERR	PCI parity error	PCI	1	I/O	13-3/13-5	—
PCI_PME	PCI PME assertion request	PCI	1	I/O	13-3/13-5	—
PCI_REQ0	PCI request 0	PCI	1	I/O	13-3/13-5	—
PCI_REQ1	PCI request 1	PCI	1	I	13-3/13-5	CPCI_HS_ES
PCI_REQ2	PCI request 2	PCI	1	I	13-3/13-5	—
PCI_RESET_OUT	PCI reset	PCI	1	O	13-3/13-5	—
PCI_SERR	PCI system error	PCI	1	I/O	13-3/13-5	—
PCI_STOP	PCI stop	PCI	1	I/O	13-3/13-5	—

Table 3-2. MPC8313E Signal Reference by Signal Name (continued)

Name	Description	Functional Block	No. of Signals	I/O	Table/ Page	Alternate Function(s)
PCI_SYNC_IN	PCI clock/PCI clock sync input	Clocks	1	I	4-2/4-3	—
PCI_SYNC_OUT	PCI clock sync output	Clocks	1	O	4-2/4-3	—
PCI_TRDY	PCI target ready	PCI	1	I/O	13-3/13-5	—
PMC_PWR_OK	Stable power	PMC	1	I	5-65/5-66	IIC2_SDA/GPIO10
PORESET	Power on reset	System control	1	I	4-1/4-1	—
QUIESCE	Quiesce state	PMC	1	O	5-65/5-66	—
RTC_PIT_CLOCK	Timer clock/Real time clock	PIT/RTC	1	I	4-2/4-3	—
RXA	Serial receiver, lane A, positive data	SGMII PHY	1	I	15-2/15-8	—
$\overline{\text{RXA}}$	Serial receiver, lane A, negative data (complement)	SGMII PHY	1	I	15-2/15-8	—
RXB	Serial receiver, lane B, positive data	SGMII PHY	1	I	15-2/15-8	—
$\overline{\text{RXB}}$	Serial receiver, lane B, negative data (complement)	SGMII PHY	1	I	15-2/15-8	—
SD_IMP_CAL_RX	Receiver impedance control signal	SGMII PHY <sup>1</sup>	1	I	15-2/15-8	—
SD_IMP_CAL_TX	Transmitter impedance control signal	SGMII PHY <sup>3</sup>	1	I	15-2/15-8	—
SD_PLL_TPA_ANA	Analog test point for SerDes PLL testing	SGMII PHY	1	O	15-2/15-8	—
SD_PLL_TPD	Digital test point for SerDes PLL testing	SGMII PHY	1	O	15-2/15-8	—
SD_REF_CLK	SerDes PLL reference clock	SGMII PHY	1	I	15-2/15-8	—
$\overline{\text{SD\_REF\_CLK}}$	SerDes PLL reference clock (complement)	SGMII PHY	1	I	15-2/15-8	—
SDAVDD	Analog supply for SerDes PLL	SGMII PHY <sup>2</sup>	1	I/O	15-2/15-8	—
SDAVSS	Analog ground for SerDes PLL	SGMII PHY	1	I/O	15-2/15-8	—

Table 3-2. MPC8313E Signal Reference by Signal Name (continued)

Name	Description	Functional Block	No. of Signals	I/O	Table/ Page	Alternate Function(s)
SPICLK	SPI clock	SPI	1	I/O	19-2/19-7	GTM1_TOUT3/ GPIO30
SPIMISO	SPI master-in slave-out	SPI	1	I/O	19-2/19-7	GTM1_TGATE3/ GTM2_TGATE4/ LDVAL/GPIO29
SPIMOSI	SPI master-out slave-in	SPI	1	I/O	19-2/19-7	GTM1_TIN3/ GTM2_TIN4/LSRCID4 /GPIO28
SPISEL	SPI slave select	SPI	1	I	19-2/19-7	GPIO31
SRESET	Soft reset to the e300 core	System control	1	I	4-1/4-1	—
SYS_CLK_IN	Clock input	Clocks	1	I	4-2/4-3	—
SYS_CR_CLK_IN	Crystal clock input	Clocks	1	I	4-2/4-3	—
SYS_CR_CLK_OUT	Crystal clock output	Clocks	1	O	4-2/4-3	—
TCK	Test clock	JTAG	1	I	20-2/20-2	—
TDI	Test data in	JTAG	1	I	20-2/20-2	—
TDO	Test data out	JTAG	1	O	20-2/20-2	—
TEST_MODE	Test mode	TEST	1	I	20-2/20-2	—
THERM[0:1]	Thermal resistor access	TEST	2	I	—	—
TMS	Test mode select	JTAG	1	I	20-2/20-2	—
TRST	Test reset	JTAG	1	I	20-2/20-2	—
TSEC_1588_ALARM1	1588 timer; current time is equal to or greater than alarm time comparator register	eTSEC	1	O	15-2/15-8	LA8
TSEC_1588_ALARM2	1588 timer; current time is equal to or greater than alarm time comparator register	eTSEC	1	O	15-2/15-8	IIC1_SCL/LA15
TSEC_1588_CLK	1588 external timer reference clock input	eTSEC	1	I	15-2/15-8	UART_SOUT2/LA10
TSEC_1588_GCLK	1588 timers	eTSEC	1	O	15-2/15-8	UART_SIN2/LA11
TSEC_1588_PP1	1588 timer pulse per period 1	eTSEC	1	O	15-2/15-8	UART_CTS2/LA12
TSEC_1588_PP2	1588 timer pulse per period 2	eTSEC	1	O	15-2/15-8	UART_RTS2/LA13

Table 3-2. MPC8313E Signal Reference by Signal Name (continued)

Name	Description	Functional Block	No. of Signals	I/O	Table/ Page	Alternate Function(s)
TSEC_1588_PP3	1588 timer pulse per period 3	eTSEC	1	O	15-2/15-8	GPIO14/LA9
TSEC_1588_TRIG1	1588 external timer trigger input 1	eTSEC	1	I	15-2/15-8	IIC1_SDA/LA14
TSEC_1588_TRIG2	1588 external timer trigger input 2	eTSEC	1	I	15-2/15-8	LA7
TSEC_GTX_CLK125	Gigabit reference clock	eTSEC1/ eTSEC2	1	I	15-2/15-8	—
TSEC_MDC	Ethernet management data clock	Ethernet management	1	O	15-2/15-8	LB_POR_CFG_BOOT_ECC_DIS <sup>3</sup>
TSEC_MDIO	Ethernet management data in/out	Ethernet management	1	I/O	15-2/15-8	—
TSEC1_COL	eTSEC1 collision detect	eTSEC1	1	I/O	15-2/15-8	USBDR_TXDRXD0
TSEC1_CRS	eTSEC1 carrier sense	eTSEC1	1	I/O	15-2/15-8	USBDR_TXDRXD1
TSEC1_GTX_CLK	eTSEC1 transmit clock out	eTSEC1	1	I/O	15-2/15-8	USBDR_TXDRXD2
TSEC1_RX_CLK	eTSEC1 receive clock	eTSEC1	1	I/O	15-2/15-8	USBDR_TXDRXD3
TSEC1_RX_DV	eTSEC1 receive data valid	eTSEC1	1	I/O	15-2/15-8	USBDR_TXDRXD4
TSEC1_RX_ER	eTSEC1 receiver error	eTSEC1	1	I	15-2/15-8	USBDR_DIR/ TSEC_1588_TRIG2
TSEC1_RXD[3:1]	eTSEC1 receive data 3–1	eTSEC1	3	I/O	15-2/15-8	USBDR_TXDRXD [5:7]
TSEC1_RXD0	eTSEC1 receive data 0	eTSEC1	1	I	15-2/15-8	USBDR_NXT
TSEC1_TX_CLK	eTSEC1 transmit clock in	eTSEC1	1	I	15-2/15-8	USBDR_CLK/ TSEC_1588_CLK
TSEC1_TX_EN	eTSEC1 transmit enable	eTSEC1	1	O	15-2/15-8	—
TSEC1_TX_ER	eTSEC1 transmit error	eTSEC1	1	O	15-2/15-8	TSEC_1588_ALARM2
TSEC1_TXD0	eTSEC1 transmit data 2	eTSEC1	1	O	15-2/15-8	USBDR_STP/ TSEC_1588_PP3
TSEC1_TXD1	eTSEC1 transmit data 2	eTSEC1	1	O	15-2/15-8	TSEC_1588_PP2

Table 3-2. MPC8313E Signal Reference by Signal Name (continued)

Name	Description	Functional Block	No. of Signals	I/O	Table/ Page	Alternate Function(s)
TSEC1_TXD2	eTSEC1 transmit data 2	eTSEC1	1	O	15-2/15-8	TSEC_1588_PP1
TSEC1_TXD3	eTSEC1 transmit data 3	eTSEC1	1	O	15-2/15-8	TSEC_1588_GCLK
TSEC2_COL	eTSEC2 collision detect	eTSEC2	1	I/O	15-2/15-8	GTM1_TIN4/ GTM2_TIN3/ GPIO15
TSEC2_CRS	eTSEC2 carrier sense	eTSEC2	1	I/O	15-2/15-8	GTM1_TGATE4/ GTM2_TGATE3/ GPIO16
TSEC2_GTX_CLK	eTSEC2 transmit clock out	eTSEC2	1	I/O	15-2/15-8	GTM1_TOUT4/ GTM2_TOUT3/ GPIO17
TSEC2_RX_CLK	eTSEC2 receive clock	eTSEC2	1	I/O	15-2/15-8	GTM1_TIN2/ GTM2_TIN1/GPIO18
TSEC2_RX_DV	eTSEC2 receive data valid	eTSEC2	1	I/O	15-2/15-8	GTM1_TGATE2/ GTM2_TGATE1/ GPIO19
TSEC2_RX_ER	eTSEC2 receiver error	eTSEC2	1	I/O	15-2/15-8	GTM1_TOUT2/ GTM2_TOUT1/ GPIO24
TSEC2_RXD[3:0]	eTSEC2 receive data 3–0	eTSEC2	4	I/O	15-2/15-8	GPIO[20:23]
TSEC2_TX_CLK	eTSEC2 transmit clock in	eTSEC2	1	I/O	15-2/15-8	GPIO25
TSEC2_TX_EN	eTSEC2 transmit enable	eTSEC2	1	I/O	15-2/15-8	GPIO26
TSEC2_TX_ER	eTSEC2 transmit error	eTSEC2	1	I/O	15-2/15-8	GPIO27
TSEC2_TXD[3:0]	eTSEC2 transmit data 3–0	eTSEC2	4	I/O	15-2/15-8	CFG_RESET_ SOURCE[0:3]
TXA	Serial transmitter, lane A, positive data	SGMII PHY	1	O	15-2/15-8	—
$\overline{\text{TXA}}$	Serial transmitter, lane A, negative data (complement)	SGMII PHY	1	O	15-2/15-8	—
TXB	Serial transmitter, lane B, positive data	SGMII PHY	1	O	15-2/15-8	—
$\overline{\text{TXB}}$	Serial transmitter, lane B, negative data (complement)	SGMII PHY	1	O	15-2/15-8	—



Table 3-2. MPC8313E Signal Reference by Signal Name (continued)

Name	Description	Functional Block	No. of Signals	I/O	Table/ Page	Alternate Function(s)
UART_CTS1	DUART clear to send	DUART	1	I/O	18-2/18-3	MSRCID2/GPIO8
UART_CTS2	DUART clear to send	DUART	1	I/O	18-2/18-3	TSEC_1588_PP1
UART_RTS1	DUART ready to send	DUART	1	I/O	18-2/18-3	MSRCID3/GPIO9
UART_RTS2	DUART ready to send	DUART	1	I/O	18-2/18-3	TSEC_1588_PP2
UART_SIN1	DUART serial data in	DUART	1	I/O	18-2/18-3	MSRCID1
UART_SIN2	DUART serial data in	DUART	1	I/O	18-2/18-3	TSEC_1588_GCLK/ MDVAL
UART_SOUT1	DUART serial data out	DUART	1	I/O	18-2/18-3	MSRCID0
UART_SOUT2	DUART serial data out	DUART	1	I/O	18-2/18-3	TSEC_1588_CLK/ MSRCID4
USB_CLK_IN	USB clock input	Clocks	1	I	4-2/4-3	—
USB_CR_CLK_IN	USB crystal clock input	Clocks	1	I	4-2/4-3	—
USB_CR_CLK_OUT	USB crystal clock output	Clocks	1	O	4-2/4-3	—
USB_DM	USB 2.0 D– line	USB PHY	1	I/O	16-1/16-3	—
USB_DP	USB 2.0 D+ line	USB PHY	1	I/O	16-1/16-3	—
USB_PLL_GND	Dedicated analog ground for USB PLL	USB PHY	1	I/O		—
USB_PLL_PWR1	Dedicated 1.0 V analog power for USB PLL	USB PHY	1	I/O		—
USB_PLL_PWR3	Dedicated 3.3 V analog power for USB PLL	USB PHY	1	I/O		—
USB_RBIAS	USB bias input <sup>4</sup>	USB_PHY	1	I/O	16-1/16-3	—
USB_TPA	Dedicated analog test signal	USB PHY	1	I/O		—
USB_VBUS	USB 2.0 VBUS line	USB PHY	1	I/O		—
USB_VDDA	Dedicated power for USB transceiver	USB PHY	1	I/O		—
USB_VDDA_BIAS	Dedicated ground for USB bias circuit	USB PHY	1	I/O		—

Table 3-2. MPC8313E Signal Reference by Signal Name (continued)

Name	Description	Functional Block	No. of Signals	I/O	Table/ Page	Alternate Function(s)
USB_VSSA	Dedicated ground for USB transceiver	USB PHY	1	I/O		—
USB_VSSA_BIAS	Dedicated power for USB bias circuit	USB PHY	1	I/O		—
USBDR_CLK	USB clock	USB	1	O	16-1/16-3	TSEC1_TX_CLK
USBDR_DIR	USB direction	USB	1	O	16-1/16-3	TSEC1_RX_ER
USBDR_DRIVE_VBUS	USB VBus power enable	USB	1	I/O	16-1/16-3	GTM1_TIN1/ GTM2_TIN2/ LSRCID0
USBDR_NXT	USB next data	USB	1	O	16-1/16-3	TSEC1_RXD0
USBDR_PCTL0	USB status LED 0 control	USB	1	O	16-1/16-3	$\overline{\text{GTM1\_TOUT1}}$ / LSRCID2
USBDR_PCTL1	USB status LED 0 control	USB	1	O	16-1/16-3	LSRCID3/ LBC_PM_REF_10
USBDR_PWRFAULT	USB VBus fault	USB	1	I/O	16-1/16-3	$\overline{\text{GTM1\_TGATE1}}$ / $\overline{\text{GTM2\_TGATE2}}$ / LSRCID1
USBDR_STP	USB stop	USB	1	O	16-1/16-3	TSEC1_TXD0/ TSEC_1588_PP3
USBDR_TXDRXD[5:7]	USB data bus 5-7	USB	1	O	16-1/16-3	TSEC1_RXD[3:1]
USBDR_TXDRXD0	USB data bus 0	USB	1	O	16-1/16-3	TSEC1_COL
USBDR_TXDRXD1	USB data bus 1	USB	1	O	16-1/16-3	TSEC1_CRS
USBDR_TXDRXD2	USB data bus 2	USB	1	O	16-1/16-3	TSEC1_GTX_CLK
USBDR_TXDRXD3	USB data bus 3	USB	1	O	16-1/16-3	TSEC1_RX_CLK
USBDR_TXDRXD4	USB data bus 4	USB	1	O	16-1/16-3	TSEC1_RX_DV
XCOREVDD[0:2]	SerDes transceiver core supply	SGMII PHY <sup>2</sup>	3	PWR	15-2/15-8	—
XCOREVSS[0:2]	SerDes transceiver core ground	SGMII PHY	3	GND	15-2/15-8	—
XPADVDD[0:1]	SerDes transceiver pad supply	SGMII PHY <sup>2</sup>	2	PWR	15-2/15-8	—
XPADVSS[0:1]	SerDes transceiver pad ground	SGMII PHY	2	GND	15-2/15-8	—

<sup>1</sup> See the *MPC8313E Hardware Specification* for resistor values.

<sup>2</sup> See [Table 3-3](#) for proper connection to power.

- <sup>3</sup> The LB\_POR\_CFG\_BOOT\_ECC\_DIS function will be selected on the TSEC\_MDC pin whenever  $\overline{\text{HRESET}}$  is asserted; the pin will act as TSEC\_MDC at all other times. The reset block will sample this signal on  $\overline{\text{PORESET}}$  negation only; the sampled value is then passed to the eLBC controller to enable/disable ECC checking during boot time. An internal pull-down resistor has been added to this pad to enable the boot-time ECC checking by default. A pull-up resistor, via a three-state buffer, is needed during HRESET assertion period to disable ECC checking during boot time.
- <sup>4</sup> Must be connected to a 10K  $\pm$ 1% precision resistor if using the integrated USB PHY through the UTMI.

## 3.2 Configuration Signals Sampled at Reset

The signals that serve alternate functions as configuration input signals during system reset are summarized in [Table 3-3](#). The detailed interpretation of their voltage levels during reset is described in [Chapter 4, “Reset, Clocking, and Initialization.”](#)

**Table 3-3. Reset Configuration Signals**

Functional Interface	Functional Signal Name	Reset Configuration Name
eTSEC2	TSEC2_TXD[3:0]	CFG_RESET_SOURCE[0:3]
Dedicated pin	None (dedicated pin)	$\overline{\text{CFG\_CLKIN\_DIV}}$
Dedicated pin	None (dedicated pin)	CFG_LBIU_MUX_EN

## 3.3 Output Signal States During Reset

When a system reset is recognized ( $\overline{\text{PORESET}}$  or  $\overline{\text{HRESET}}$  are asserted), the device aborts all current internal and external transactions and releases all bidirectional I/O signals to a high-impedance state. See [Chapter 4, “Reset, Clocking, and Initialization,”](#) for a complete description of the reset functionality.

During reset, the device ignores most input signals (except for the reset configuration signals) and drives most of the output-only signals to an inactive state. [Table 3-4](#) shows the states of the output-only signals.

**Table 3-4. Output Signal States During System Reset**

Interface	Signal	State During Reset
MDM[0:3]	DDR data mask	All 'Z'
MBA[0:2]	DDR bank select	All 'Z'
MA[0:14]	DDR address	All 'Z'
$\overline{\text{MWE}}$	DDR write enable	'Z'
$\overline{\text{MRAS}}$	DDR row address strobe	'Z'
$\overline{\text{MCAS}}$	DDR column address strobe	'Z'
$\overline{\text{MCS}}$ [0:1]	DDR chip select (2/DIMM)	Both 'Z'
MCKE	DDR clock enable	'0'
MCK	DDR differential clocks	'0'
$\overline{\text{MCK}}$	DDR differential clocks	'1'
MODT[0:1]	DRAM on-die termination	Both '0'

Table 3-4. Output Signal States During System Reset (continued)

Interface	Signal	State During Reset
$\overline{\text{PCI\_INTA}}$	PCI interrupt output	'Z'
$\overline{\text{PCI\_RESET\_OUT}}$	PCI reset output	'0'
$\overline{\text{PCI\_GNT}}[0:2]$	PCI grant 0–2	All 'Z'
UART_SOUT[0:1]	DUART serial data out	All 'Z'
$\overline{\text{UART\_RTS}}[0:1]$	DUART ready to send	All 'Z'
TSEC1_GTX_CLK	eTSEC1 transmit clock out	'0'
TSEC1_TXD[3:0] <sup>1</sup>	eTSEC1 transmit data 2–0	All 'Z'
TSEC1_TX_EN	eTSEC1 transmit enable	'0'
TSEC2_GTX_CLK	eTSEC2 transmit clock out	'Z'
TSEC2_TXD[3:0]	eTSEC2 transmit data 6-4	All 'Z'
TSEC_MDC	Ethernet management data clock	'0'
LA[16:25]	eLBC port address	Active—used to load reset configuration word
$\overline{\text{LCS0}}$	eLBC chip select 0	Active—used to load reset configuration word
$\overline{\text{LCS1}}$	eLBC chip select 1	'1'
$\overline{\text{LWE0/LFWE/LBS0}}$	eLBC write enable 0/FCM write enable 0/UPM byte (lane) select 0	'1'
$\overline{\text{LWE1/LBS1}}$	eLBC write enable 1/UPM byte (lane) select 1	'1'
LBCTL	eLBC data buffer control	Active—used to load reset configuration word
LALE	eLBC address latch enable	Active—used to load reset configuration word
$\overline{\text{LOE/LGPL2/LFRE}}$	eLBC output enable/GP line 2/FCM read enable	Active—used to load reset configuration word
LCLK[0:1]	eLBC clock	All '0'
USBDR_PCTL0	USB port control 0	'1'
USBDR_PCTL1	USB port control 1	'0'
$\overline{\text{MCP\_OUT}}$	Machine check interrupt output	'Z'
TDO	Test data out	'Z'
$\overline{\text{QUIESCE}}$	Quiesce state	'1'
PCI_SYNC_OUT	PCI sync output	Active clock
USBDR_STP_SUSPEND <sup>2</sup>	USB host 0 data stop/suspend	Z

<sup>1</sup> State of TSEC1\_TXD0 depends on RTBI mode in TSEC1M bits in RCWH. Its state changes from 1 to 0 after RCWH is loaded during reset sequence.

<sup>2</sup> This pin should be pulled high during a hard reset for proper functionality of the device since it has a weak internal pull up. No external pull-down resistors are allowed to be mounted on this net.

Table 3-5 provides the list of registers that control the device's signal multiplexing.

**Table 3-5. Signals for Multiplexing**

Signal Group	Multiplexing is Controlled By	Table/Page
eLBC (LA[0:15]/GPIO)	CFG_LBIU_MUX_EN	3-6/3-31
PCI/CPCI	RCWH[PCIARB]	4.3.2.2/4-14
All others	SICRL/SICRH	5.3.2.5/5-21/5.3.2.6/5-23

### 3.4 External Signal Description

Table 3-6 provides a detailed description of the external CFG\_LBIU\_MUX\_EN (LA[0:15]) signals.

**Table 3-6. External Signals—Detailed Signal Descriptions**

Signal	I/O	Description	
CFG_LBIU_MUX_EN (LA[0:15]/GPIO)	I	<b>State Meaning</b>	Asserted—GPIO ports function according to the values in their configuration registers (this may be eLBC or non-eLBC functionality). Negated—GPIO ports function as LA[0:15] irrespective of the value in their configuration registers.
		<b>Timing</b>	This signal should be connected to pull-up or pull-down on the board. The signal needs to be valid at all times and static.



# Chapter 4

## Reset, Clocking, and Initialization

The reset, clocking, and control signals offer many options for operating the device. Various modes and features can be configured during hard reset or power-on reset. Most configurable features are loaded to the device through a reset configuration word, and a few device signals are used as reset configuration inputs during the reset sequence.

### 4.1 External Signals

The following sections describe the reset and clock signals in detail.

#### 4.1.1 Reset Signals

Table 4-1 describes the reset signals of the device. Section 4.3.2, “Reset Configuration Words,” describes the signals that also function as reset configuration signals.

**Table 4-1. System Control Signals**

Signal	I/O	Description
PORESET	I	Power-on reset. Initiates the power-on reset flow that resets the device and configures various attributes of the device, including its clock modes.
		<b>State Meaning</b> Asserted—An external agent has triggered a power-on reset sequence. Negated—No power-on reset.
		<b>Timing</b> See the hardware specifications for timing information.
		<b>Reset State</b> Always input.
HRESET	I/O	Hard reset. Causes the device to abort all current internal and external transactions and set most registers to their default values. HRESET can be asserted completely asynchronously with respect to all other signals. The device can detect an external assertion of HRESET while the device is not asserting hard reset. HRESET is an open-drain signal.
		<b>State Meaning</b> Asserted—An external agent or internal hardware has triggered a hard reset. The internal hardware drives HRESET until the sequence completes. Negated—No hard reset.
		<b>Timing</b> Assertion—Occur at any time, asynchronously to any clock. Negation—Must be asserted for at least 32 SYS_CLK_IN (PCI host mode) or PCI_CLK (PCI agent mode) cycles.
		<b>Requirements</b> An open-drain signal. An external pull-up is required.
		<b>Reset State</b> Output, driven low during power-on and hard reset flows. High impedance after reset flow completes.

Table 4-1. System Control Signals (continued)

Signal	I/O	Description	
$\overline{\text{SRESET}}$	I	The $\overline{\text{SRESET}}$ input is connected directly to the soft reset input of the e300c3 core. The assertion of the e300 soft reset input, $\overline{\text{SRESET}}$ , causes a high priority interrupt to the e300 core as described in <a href="#">Section 4.2.1, "Reset Operations."</a> It does not reset the e300 core or any other portion of the MPC8313E. The $\overline{\text{SRESET}}$ input is not registered in the reset status register (RSR).	
		<b>State Meaning</b>	Asserted—Indicates that the processor must initiate a system reset interrupt. Negated—Indicates that the interrupt is not being requested.
		<b>Timing</b>	Assertion—Occurs at any time, asynchronously to any clock. Negation—Occurs after being serviced. (PCI host mode) or PCI_CLK (PCI agent mode)
		<b>Requirements</b>	An open-drain signal. An external pull-up is required.
		<b>Reset State</b>	Always input
CFG_RESET_SOURCE[0:3]	I	Reset configuration word source selection. These signals are on device pins that have other functions when the device is not in reset. They are sampled during the assertion of $\overline{\text{PORESET}}$ to determine the interface from which the device loads the reset configuration words.	
		<b>State Meaning</b>	See <a href="#">Section 4.3.1.1, "Reset Configuration Word Source."</a>
		<b>Timing</b>	These signals are sampled during the assertion of $\overline{\text{PORESET}}$ after a stable clock is supplied ( $\overline{\text{PORESET}}$ flow) and must be pulled high or low by external resistors as long as $\overline{\text{HRESET}}$ is asserted.
		<b>Requirements</b>	During $\overline{\text{PORESET}}$ and $\overline{\text{HRESET}}$ flows, all other signal drivers connected to these signals must be in the high-impedance state. Refer to the hardware specifications for proper resistor values to pull reset configuration signals high or low.
		<b>Reset State</b>	Input during power-on and hard reset flows. Functional signal after reset flow completes.
$\overline{\text{CFG\_CLKIN\_DIV}}$	I	Clock in division selection. This signal is located on a device pin that has another function when the device is not in reset state. This signal is sampled during the assertion of $\overline{\text{PORESET}}$ to determine whether SYS_CLK_IN is divided by two.	
		<b>State Meaning</b>	See <a href="#">Section 4.3.1.2, "SYS_CLK_IN Division."</a>
		<b>Timing</b>	This signal is sampled during the assertion of $\overline{\text{PORESET}}$ after a stable clock is supplied ( $\overline{\text{PORESET}}$ flow), and it must be pulled high or low by external resistors as long as $\overline{\text{HRESET}}$ is asserted.
		<b>Requirements</b>	During $\overline{\text{PORESET}}$ and $\overline{\text{HRESET}}$ flows, all other signal drivers connected to this signal must be in the high-impedance state. Refer to the hardware specifications for proper resistors values to pull reset configuration signals high or low.
		<b>Reset State</b>	Always input



## 4.1.2 Clock Signals

In Table 4-2, some clock signals are specific to blocks within the device. Although some of their functionality is described in Section 4.4, “Clocking,” they are defined in detail in their respective chapters. See Figure 4-7 for the internal distribution of clocks in the device.

**Table 4-2. External Clock Signals**

Signal	I/O	Description	
SYS_CLK_IN	I	System clock. In PCI host mode, SYS_CLK_IN is the primary input clock. SYS_CLK_IN directly feeds the PCI output clock dividers and is driven out on the PCI_SYNC_OUT signal for de-skewing external PCI clocks routing. If the device is used as PCI agent, PCI_CLK can be provided from a PCI source. In this case, SYS_CLK_IN may no longer be needed and should be tied low. If SYS_CLK_IN is the clock source, SYS_CR_CLK_IN should be tied low and SYS_CR_CLK_OUT should be left unconnected.	
		<b>Timing</b>	Assertion/Negation—See the hardware specifications for timing information.
		<b>Requirements</b>	Should be tied low if unused, for example in PCI agent mode or when the clock is provided through SYS_CR_CLK_IN.
		<b>Reset State</b>	Always input.
SYS_CR_CLK_IN	I	Crystal input. SYS_CR_CLK_IN/SYS_CR_CLK_OUT allows the system clock to be provided using an external crystal oscillator. If a crystal source is used, SYS_CLK_IN should be tied low.	
		<b>Timing</b>	Assertion/Negation—See the hardware specifications for timing information.
		<b>Requirements</b>	Should be tied low if unused, for example in PCI agent mode or when the clock is provided through SYS_CLK_IN.
		<b>Reset State</b>	Always input.
SYS_CR_CLK_OUT	O	Crystal output. SYS_CR_CLK_IN/SYS_CR_CLK_OUT allows the system clock to be provided through an external crystal oscillator. If a crystal source is used, SYS_CLK_IN should be tied low.	
		<b>Timing</b>	Assertion/Negation—See the hardware specifications for timing information.
		<b>Requirements</b>	Should be left unconnected if unused, for example in PCI agent mode or when the clock is provided through SYS_CLK_IN.
		<b>Reset State</b>	Always output.
USB_CLK_IN	I	USB PHY clock. USB_CLK_IN feeds into the USB PHY PLL.	
		<b>Timing</b>	Assertion/Negation—See the hardware specifications for timing information.
		<b>Requirements</b>	Should be tied low if unused, for example when the clock is provided through USB_CR_CLK_IN or when derived from the system clock.
		<b>Reset State</b>	Always input.
USB_CR_CLK_IN	I	USB crystal input. USB_CR_CLK_IN/USB_CR_CLK_OUT allows the USB clock to be provided using an external crystal oscillator. If a crystal source is used, USB_CLK_IN should be tied low.	
		<b>Timing</b>	Assertion/Negation—See the hardware specifications for timing information.
		<b>Requirements</b>	Should be tied low if unused, for example when the clock is provided through USB_CR_CLK_IN or when derived from the system clock.
		<b>Reset State</b>	Always input.

Table 4-2. External Clock Signals (continued)

Signal	I/O	Description	
USB_CR_CLK_OUT	O	USB crystal output. USB_CR_CLK_IN/USB_CR_CLK_OUT allows the USB clock to be provided through an external crystal oscillator. If a crystal source is used, USB_CLK_IN should be tied low.	
		<b>Timing</b>	Assertion/Negation—See the hardware specifications for timing information
		<b>Requirements</b>	Should be left unconnected if unused, for example when the clock is provided through USB_CR_CLK_IN or when derived from the system clock.
		<b>Reset State</b>	Always output.
PCI_CLK/ PCI_SYNC_IN	I	PCI clock/ PCI synchronization clock (PCI_CLK/PCI_SYNC_IN). PCI_CLK is the primary clock input to the device, the reference clock for the system APLL. When the device is in PCI host mode SYS_CLK_IN or SYS_CR_CLK_IN is used as the clock source. In this case the PCI_CLK_OUT[0:2] signals are driven and PCI_SYNC_IN should be tied to PCI_SYNC_OUT. When the device is in PCI agent mode PCI_CLK will be tied directly to a PCI system clock source.	
		<b>Timing</b>	Assertion/Negation—See the hardware specifications for timing information
		<b>Reset State</b>	Always input.
PCI_SYNC_OUT	O	Reference PCI output synchronization clock (PCI_SYNC_OUT). In PCI host mode with the PCI_CLK_OUT[0:2] signals driven, PCI_SYNC_OUT is connected externally to PCI_SYNC_IN signal for de-skewing external PCI clocks routing. PCI_SYNC_OUT has the same frequency as CLKIN or CLKIN/2 depending on the state of CFG_CLKIN_DIV at reset. See <a href="#">Section 4.3.1.2, “SYS_CLK_IN Division.”</a> In PCI agent mode, this signal is typically not used.	
		<b>Timing</b>	Assertion/Negation—See the hardware specifications for timing information.
		<b>Reset State</b>	Always output, toggling in PCI host mode.
PCI_CLK_OUT[0:2]	O	PCI output clocks bank. In PCI host mode, the device provides three separate clock output signals for feeding PCI agent devices.	
		<b>Timing</b>	Assertion/Negation—See the hardware specifications for timing information.
		<b>Reset State</b>	Always output. Drive ‘0’ and after power-on reset flow. Enabled by a memory-mapped register.

## 4.2 Functional Description

This section describes the various ways to reset the device, the power-on reset configurations, and clocking.

### 4.2.1 Reset Operations

The device has several inputs to the reset logic:

- Power-on reset ( $\overline{\text{PORESET}}$ )
- External hard reset ( $\overline{\text{HRESET}}$ )
- External soft reset ( $\overline{\text{SRESET}}$ )
- Software watchdog reset
- System bus monitor reset
- Checkstop reset

- JTAG reset
- Software hard reset

All of these reset sources are fed into the reset controller and, depending on the source of the reset, different actions are taken. The reset status register, described in [Section 4.5.1.3, “Reset Status Register \(RSR\),”](#) indicates the last sources to cause a reset.

#### 4.2.1.1 Reset Causes

[Table 4-3](#) describes reset causes.

**Table 4-3. Reset Causes**

Name	Description
Power-on reset (PORESET)	Input signal. Asserting this signal initiates the power-on reset flow that resets the entire device and configures various attributes of the device including its clock modes.
Hard reset ( $\overline{\text{HRESET}}$ )	A bidirectional I/O signal. The device can detect an external assertion of $\overline{\text{HRESET}}$ only while it is not asserting hard reset. $\overline{\text{HRESET}}$ is an open-drain signal.
Soft reset ( $\overline{\text{SRESET}}$ )	Input signal. Connected to the $\overline{sreset}$ input of the e300 core, and when asserted, causes a high priority interrupt to the e300 core.
Software watchdog reset	After the device watchdog counts to zero, a software watchdog reset is signaled. The enabled software watchdog event then generates an internal hard reset sequence.
System bus monitor reset	After the device CSB bus monitor reaches a timeout condition, a bus monitor reset is asserted. The enabled bus monitor event then generates an internal hard reset sequence.
Checkstop reset	If the core enters checkstop state and the checkstop reset is enabled ( $\text{RMR}[\text{CSRE}] = 1$ ), the checkstop reset is asserted. The enabled checkstop event then generates an internal hard reset sequence.
JTAG reset	When JTAG logic asserts the JTAG soft reset signal, an internal soft reset sequence is generated.
Software hard reset	A hard reset sequence can be initialized by writing to a memory-mapped register (RCR).

#### 4.2.1.2 Reset Actions

The reset control logic determines the cause of reset, synchronizes it if necessary, and resets the appropriate internal hardware. The  $\overline{\text{SRESET}}$  input is not routed to the reset control logic, but directly to the  $\overline{sreset}$  input on the e300 core. Each reset flow has a different impact on the device logic:

- Power-on reset has the greatest impact, resetting the entire device, including clock logic and error capture registers.
- Hard reset resets the entire device excluding clock logic and error capture registers.
- Soft reset initializes the internal logic while maintaining the system configuration.

The memory controller, system protection logic, interrupt controller, and I/O signals are initialized only on hard reset. A soft reset causes a reset exception to the e300 core but does not reset other device logic. [Table 4-4](#) identifies the reset actions for each reset source.

Table 4-4. Reset Actions

Action	Reset Source			
	Power-On Reset	External Hard Reset Software Watchdog Bus Monitor Checkstop Software Hard Reset	JTAG Reset	External Soft Reset
Resets: PLLs, clocks, RTC unit, and error capture registers	Yes	No	No	No
Resets: DDR, LBC, I/O multiplexors, GTM, PIT, GPIO, system configuration, and local access windows	Yes	Yes	No	No
Resets other internal logic	Yes	Yes	Yes	No
Reset configuration words loaded	Yes	Yes	No	No
$\overline{\text{HRESET}}$ driven	Yes	Yes	No	No
Hard reset to e300 core	Yes	Yes	Yes	Yes
High priority interrupt to the e300 core	No	No	No	Yes

## 4.2.2 Power-On Reset Flow

Assertion of the  $\overline{\text{PORESET}}$  external signal initiates the power-on reset flow.  $\overline{\text{PORESET}}$  should be asserted externally for at least 32 input clock cycles after stable external power to the device is applied.

Directly after the negation of  $\overline{\text{PORESET}}$ , the device starts the configuration process. The device asserts  $\overline{\text{HRESET}}$  throughout the power-on reset process, including configuration. Configuration time varies according to the configuration source and  $\text{SYS\_CLK\_IN}$  (PCI host mode) or  $\text{PCI\_CLK}$  (PCI agent mode) frequency. Initially, the reset configuration inputs are sampled to determine the configuration source and the input clock division mode. Next, the device starts loading the reset configuration words. The system PLL begins to lock according to the clock mode values in the reset configuration word low. When the system PLL is locked, the clock unit starts distributing clock signals in the device. At this stage, the core PLL begins to lock. When it is locked and the reset configuration words are loaded,  $\overline{\text{HRESET}}$  is released.

The detailed power-on reset (POR) flow for the device is as follows:

1. Power is applied to meet the specifications in the device data sheet.
2. The system asserts  $\overline{\text{PORESET}}$  and  $\overline{\text{TRST}}$ , causing all registers to be initialized to their default states and most I/O drivers to be released to high-impedance.  
Some clock, clock enabled, and system control signals remain active.
3. The system applies a stable  $\text{SYS\_CLK\_IN}$  (PCI host mode) or  $\text{PCI\_CLK}$  (PCI agent mode) signal and stable reset configuration inputs ( $\text{CFG\_RESET\_SOURCE}$ ,  $\text{CFG\_CLKIN\_DIV}$ ).
4. The system negates  $\overline{\text{PORESET}}$  after at least 32 stable  $\text{SYS\_CLK\_IN}$  (PCI host mode) or  $\text{PCI\_CLK}$  (PCI agent mode) clock cycles.

5. The device samples the reset configuration input signals to determine the clock division and the reset configuration words source.
6. The device starts loading the reset configuration words.  
Loading time depends on the reset configuration word source.
7. When the reset configuration word low is loaded, the system PLL begins to lock.  
When the system PLL is locked, *csb\_clk* is supplied to the core PLL.
8. The core PLL begins to lock.
9. The device drives  $\overline{\text{HRESET}}$  asserted until the e300 PLL is locked and the reset configuration words are loaded.
10. The user optionally negates  $\overline{\text{HRESET}}$  if it was not negated earlier.  
JTAG logic must always be initialized by asserting  $\overline{\text{TRST}}$ . If the JTAG signals are not used,  $\overline{\text{TRST}}$  should be connected directly to  $\overline{\text{PORESET}}$ .  $\overline{\text{TRST}}$  must not remain asserted after the negation of  $\overline{\text{PORESET}}$ . There is no need to assert the  $\overline{\text{SRESET}}$  signal when  $\overline{\text{HRESET}}$  is asserted.
11. The internal reset to the core and the rest of the logic is negated. I/O drivers are enabled. The PCI interface can assert  $\overline{\text{DEVSEL}}$  in response to configuration cycles.
12. The device stops driving  $\overline{\text{HRESET}}$ . The reset to the e300 core is negated and the core is enabled. The boot sequencer, if enabled, is released, causing it to load configuration data from serial ROMs, as described in [Section 17.4.5, “Boot Sequencer Mode.”](#)
13. Before the boot sequencer finishes, it can enable the PCI interface to accept external requests, if required, by clearing the `CFG_LOCK` bit in the PCI function configuration register as described in [Table 13-40](#). If the e300 core is required to proceed, the boot sequencer should enable boot vector fetch by clearing `ACR[COREDIS]` as described in [Section 6.2.1, “Arbiter Configuration Register \(ACR\).”](#)
14. The PCI interface can now accept external requests, if enabled, and the boot vector fetch by the core can proceed, if enabled.

The device is now in its ready state.

Figure 4-1 shows a timing diagram of the power-on reset flow.

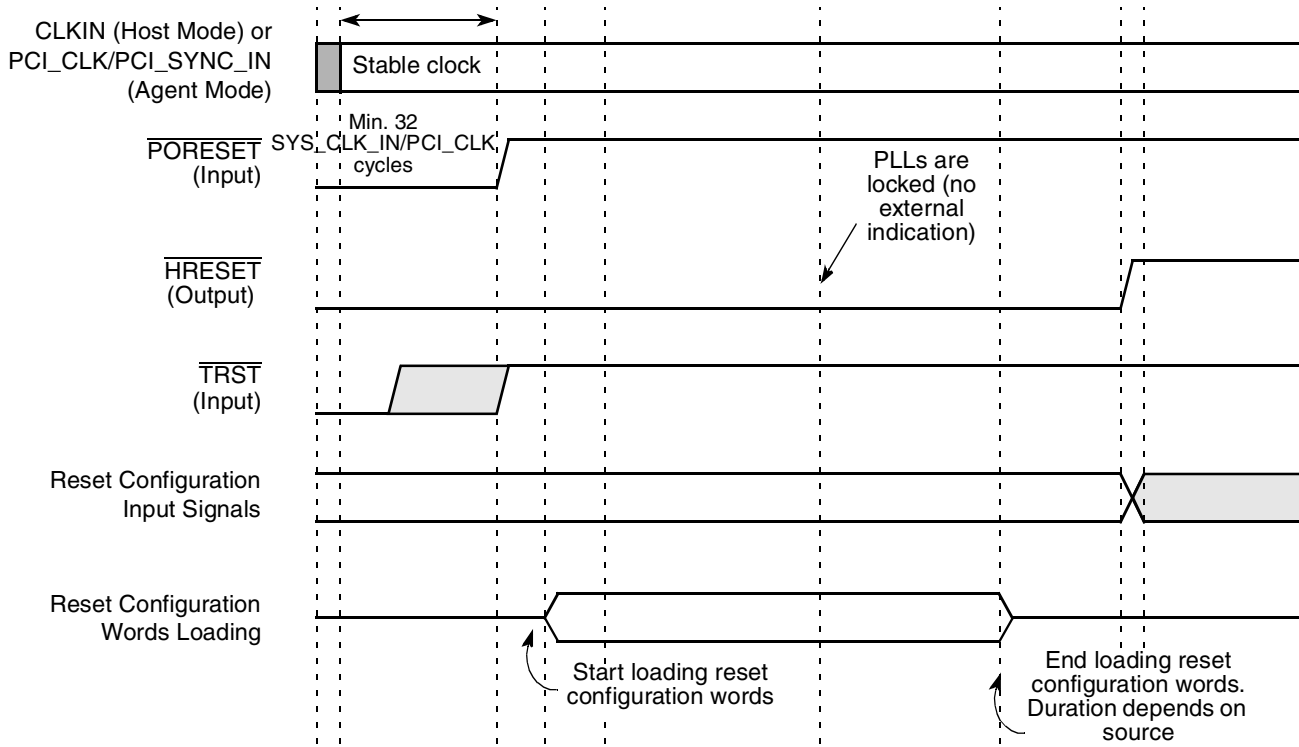


Figure 4-1. Power-On Reset Flow

### 4.2.3 Hard Reset Flow

The  $\overline{\text{HRESET}}$  signal is initiated externally by asserting  $\overline{\text{HRESET}}$  or internally when the device detects a reason to generate an internal hard reset sequence. In both cases, the device continues asserting  $\overline{\text{HRESET}}$  throughout the  $\overline{\text{HRESET}}$  state. The hard reset sequence time varies according to the configuration source and  $\text{SYS\_CLK\_IN}$  (PCI host mode) or  $\text{PCI\_CLK}$  (PCI agent mode) frequency. The reset configuration input signals ( $\text{CFG\_RESET\_SOURCE}$  and  $\overline{\text{CFG\_CLKIN\_DIV}}$ ) are not sampled by hard reset (only by power-on reset), so the device immediately starts loading the reset configuration words and configures the device as explained in Section 4.3.3, "Loading the Reset Configuration Words." After the configuration sequence completes, the device releases the  $\overline{\text{HRESET}}$  signal and exits the  $\overline{\text{HRESET}}$  state. An external pull-up resistor should negate the signals. After negation is detected, a 16-cycle period is taken before testing for the presence of an external (hard) reset.

#### NOTE

Because the device does not sample the reset configuration input signals ( $\text{CFG\_RESET\_SOURCE}$ ,  $\overline{\text{CFG\_CLKIN\_DIV}}$ ) during a hard reset flow, setting a new value on those signals (other than that set during power-on reset) has no effect.

Figure 4-2 shows a timing diagram of the hard reset flow.




Figure 4-2. Hard Reset Flow

## 4.2.4 Soft Reset Flow

The  $\overline{\text{SRESET}}$  signal can be initiated externally by asserting  $\overline{\text{SRESET}}$ .  $\overline{\text{SRESET}}$  assertion asserts the *sreset* input to the e300 core. No other hardware is reset.

## 4.3 Reset Configuration

The device is initialized using two complementary methods, latching `CFG_RESET_SOURCE` and loading the reset configuration words. Initially, a few input signals are sampled during the assertion of the  $\overline{\text{PORESET}}$  signal. These signals determine whether a reset configuration word is required and the device source interface from which it is loaded. According to the value on these signals, the device continues loading the reset configuration word.

### 4.3.1 Reset Configuration Signals

Reset configuration input signals are on device pins that have other functions when the device is not in reset state. These input signals are sampled into registers during the assertion of  $\overline{\text{PORESET}}$ , after a stable clock is supplied (`PCI_CLK`), and must be pulled high or low by external resistors as long as  $\overline{\text{HRESET}}$  is asserted. While the  $\overline{\text{PORESET}}$  and  $\overline{\text{HRESET}}$  signals are asserted, all other signal drivers connected to these signals must be in the high-impedance state. Refer to the hardware specifications for proper resistor values for pulling reset configuration signals high or low.

This section describes the modes configured by the reset configuration signals. Note that the reset configuration inputs sampled values are accessible to software through memory-mapped registers described in Section 4.5.1.3, “Reset Status Register (RSR),” and Section 4.5.2.1, “System PLL Mode Register (SPMR).”

### NOTE

Implement one of the following methods to control the selection between the reset and non-reset function of these pins.

- Resistors. Use pullup or pulldown resistors to set the desired value on the reset configuration input signals. During the power-on and hard reset sequences, these signals are inputs to the device.
- Active driving device. Use  $\overline{\text{HRESET}}$  to control the driving device. When  $\overline{\text{HRESET}}$  is asserted, drive reset configuration values on the pins; when  $\overline{\text{HRESET}}$  is negated, stop driving the reset configuration input signals.

#### 4.3.1.1 Reset Configuration Word Source

The reset configuration word source options, shown in Table 4-5, select whether the device loads a reset configuration word from NOR Flash, NAND Flash, or an I<sup>2</sup>C EEPROM (I<sup>2</sup>C #1) or uses hard-coded default options. The value of these signals also affects the duration of power-on and hard reset sequences. In any case, the reset sequence does not exceed 1 ms.

**Table 4-5. Reset Configuration Words Source**

CFG_RESET_SOURCE[0:3]	Meaning
0000	Reset configuration word is loaded from NOR Flash
0001	Reset configuration word is loaded from NAND Flash memory (8-bit small page).
0010	Reserved
0011	Reserved
0100	Reset configuration word is loaded from an I <sup>2</sup> C EEPROM. PCI_CLK/PCI_SYNC_IN is valid for any PCI frequency up to 66.666 MHz (range of 24–66.666 MHz).
0101	Reset configuration word is loaded from NAND Flash memory (8-bit large page).
0110	Reserved
0111	Reserved
1000	Hard-coded option 0. Reset configuration word is not loaded.
1001	Hard-coded option 1. Reset configuration word is not loaded.
1010	Hard-coded option 2. Reset configuration word is not loaded.
1011	Hard-coded option 3. Reset configuration word is not loaded.
1100	Hard-coded option 4. Reset configuration word is not loaded.
1101	Reserved
1110	Reserved
1111	Reserved



### 4.3.1.2 SYS\_CLK\_IN Division

When the device is configured as a PCI host, the  $\overline{\text{CFG\_CLKIN\_DIV}}$  configuration input selects the relationship between SYS\_CLK\_IN and PCI\_SYNC\_OUT as shown in Table 4-7. As a PCI host, the device supports three output signals. The frequency of the output clocks will be equal to the PCI\_SYNC\_OUT frequency.

When the device is configured as a PCI agent, the  $\overline{\text{CFG\_CLKIN\_DIV}}$  configuration input can be used to double the internal clock frequencies, if sampled as '0' during power-on reset assertion. This feature is useful if a fixed internal frequency is desired regardless of whether the PCI clock is running at 33 or 66 MHz. PCI specifications require the PCI clock frequency information to be provided by the M66EN signal.

When the device is configured as PCI host, there are two scenarios for connecting the  $\overline{\text{CFG\_CLKIN\_DIV}}$  configuration input. If the frequency of SYS\_CLK\_IN is 33 MHz (that is, the PCI system is running on a 33-MHz clock),  $\overline{\text{CFG\_CLKIN\_DIV}}$  should be connected high. If the frequency of SYS\_CLK\_IN is 66 MHz (that is, the PCI system can run at 33- or 66-MHz clock signaled by M66EN),  $\overline{\text{CFG\_CLKIN\_DIV}}$  should be connected to the M66EN signal.

Table 4-6. SYS\_CLK\_IN Division

CFG_CLKIN_DIV	Description
1	In PCI host mode, SYS_CLK_IN: PCI_SYNC_OUT = 1:1 and all PCI_CLK_OUT[0:2] clocks are running at a frequency which is equal to the SYS_CLK_IN frequency
0	In PCI agent mode, SYS_CLK_IN: PCI_SYNC_OUT = 2:1 and all PCI_CLK_OUT[0:2] clocks are running at a frequency which is equal to the PCI_SYNC_OUT frequency.

### 4.3.1.3 Selecting Reset Configuration Input Signals

The example described in Table 4-7 shows how the user should pull down or pull up the reset configuration input signals (CFG\_RESET\_SOURCE,  $\overline{\text{CFG\_CLKIN\_DIV}}$ ). The reset sequence duration is measured from the negation of  $\overline{\text{PORESET}}$  to the negation of  $\overline{\text{HRESET}}$ . Note that the duration mentioned in this table is typical and does not represent cases in which the process of loading the reset configuration word had to be retried due to errors.

Table 4-7. Selecting Reset Configuration Input Signals

I <sup>2</sup> C EEPROM Configuration Words	SYS_CLK_IN Frequency (Host Mode)	$\overline{\text{CFG\_CLKIN\_DIV}}$ (Host Mode)	PCI_CLK Frequency (Agent Mode)	CFG_RESET_SOURCE[0:3]	Reset Sequence Duration in SYS_CLK_IN/ PCI_CLK Cycles	Duration
No	33 MHz	1	33 MHz	0000 (RCW loaded from NOR Flash)	15210	456 $\mu\text{s}$
No	66 MHz	1	66 MHz	1000–1100 (use hard coded RCW)	15380	231 $\mu\text{s}$

Table 4-7. Selecting Reset Configuration Input Signals (continued)

I <sup>2</sup> C EEPROM Configuration Words	SYS_CLK_IN Frequency (Host Mode)	CFG_CLKIN_DIV (Host Mode)	PCI_CLK Frequency (Agent Mode)	CFG_RESET_SOURCE[0:3]	Reset Sequence Duration in SYS_CLK_IN/PCI_CLK Cycles	Duration
No	66 MHz	0	33 MHz	0000 (RCW loaded from NOR Flash)	30420/15210	456 $\mu$ s
Yes	33 MHz	1	33 MHz	0100 (I <sup>2</sup> C EEPROM)	106534	3196 $\mu$ s
Yes	66 MHz	1	66 MHz	0100 (I <sup>2</sup> C EEPROM)	106534	1598 $\mu$ s
Yes	66 MHz	0	33 MHz	0100 (I <sup>2</sup> C EEPROM)	213068/106534	3196 $\mu$ s
No	66 MHz	1	66 MHz	0001 (RCW loaded from 8-bit small page NAND Flash)	23024	345 $\mu$ s
No	66 MHz	1	66 MHz	0101 (RCW loaded from 8-bit large page NAND Flash)	45284	679 $\mu$ s

### 4.3.2 Reset Configuration Words

The reset configuration words control the clock ratios and other basic device functions such as PCI host or agent mode, boot location, and endian mode. The reset configuration words are loaded from NOR Flash, NAND Flash, or the I<sup>2</sup>C interfaces or from hard-coded values during the power-on or hard reset flows. See [Section 4.3.1, “Reset Configuration Signals,”](#) for information on the reset configuration word source.

Although the configuration reset words are loaded during hard reset flows, the clocks and PLL modes are reset only when PORESET is asserted during a power-on reset flow. See [Section 4.2.1.2, “Reset Actions.”](#) The values of fields in the reset configuration words registers (RCWLR and RCWHR) reflect only their state during the reset flow. Some of these parameters and modes can be modified by changing their values in the memory-mapped registers of other units, which does not affect RCWLR and RCWHR.


The reset configuration settings are accessible to software through the following read-only memory-mapped registers:

- Reset configuration word low register (RCWLR)
- Reset configuration word high register (RCWHR)
- Reset status register (RSR)
- System PLL mode register (SPMR)

See [Section 4.5, “Memory Map/Register Definitions.”](#)

### 4.3.2.1 Reset Configuration Word Low Register (RCWLR)

RCWLR is shown in [Figure 4-3](#). This read-only register gets its values according to the reset configuration word low loaded during the reset flow.



**Figure 4-3. Reset Configuration Word Low Register (RCWLR)**

[Table 4-8](#) defines the RCWLR bit fields.

**Table 4-8. RCWLR Bit Settings**

Bits	Name	Description
0	LBCM	Local bus memory controller clock mode. Selects the local bus controller clock ratio. The local bus memory controller operates with a frequency equal to the frequency of <i>csb_clk</i> . This bit should be cleared.
1	DDRCM	DDR SDRAM memory controller clock mode. Selects the DDR SDRAM memory controller clock ratio. The DDR SDRAM memory controller operates at twice the frequency of the <i>csb_clk</i> . 0 <i>csb_clk</i> ratio is 1:1 1 <i>csb_clk</i> ratio is 2:1
2–3	—	Reserved. Must be configured as 2'b10.
4–7	SPMF	System PLL multiplication factor. See <a href="#">Section , “,”</a> for more information.
8	—	Reserved, should be cleared
9–15	COREPLL	Core PLL configuration. COREPLL sets the ratio between the e300 core clock and the internal <i>csb_clk</i> of the device. The encodings for COREPLL are given in the hardware specifications for this device.
16–31	—	Reserved, should be cleared.

#### 4.3.2.1.1 System PLL Configuration

The system PLL ratio reset, shown in [Table 4-9](#), establishes the clock ratio between the SYS\_CLK\_IN signal and the internal *csb\_clk* of the device. *csb\_clk* drives internal units and feeds the e300 core PLL.

**Table 4-9. System PLL Ratio**

RCWLR Bits	Field Name	Value (Binary)	<i>csb_clk</i> : CLKIN (PCI Host Mode) <i>csb_clk</i> : (PCI_CLK x (1+~sampled_cfg_clkdiv)) (PCI Agent Mode)
4–7	SPMF	0000	Reserved
		0001	Reserved
		0010	2 : 1
		0011	3 : 1
		0100	4 : 1
		0101	5 : 1
		0110	6 : 1
		0111–1111	Reserved, should not be set

**NOTE**

In PCI host mode, the SPMF field described in [Table 4-9](#) always selects the *csb\_clk*:SYS\_CLK\_IN ratio regardless of the CFG\_CLKIN\_DIV reset configuration input value during reset flow.

The SPMF field maximum allowed value is dependent on the value sampled on CFG\_CLKIN\_DIV during power-on reset. [Table 4-10](#) defines the upper limit of SPMF with respect to these values. Values for SPMF are as follows:

**Table 4-10. SPMF Maximum Values**

CFG_CLKIN_DIV	LBCM	DDRCM	Maximum SPMF Value (Decimal)
1	0	1	6
0	0	1	3

**4.3.2.2 Reset Configuration Word High Register (RCWHR)**

RCWHR is shown in [Figure 4-4](#). This read-only register gets its values according to the reset configuration word high loaded during the reset flow.

Offset 0x0\_0904

Access: Read/Write

	0	1	2	3	4	5	6	7	8	9	11	12	13	14	15
Field	PCIHOST	—	PCIARB	—	COREDIS	BMS	BOOTSEQ	SWEN	ROMLOC		RLEXT		—	—	
	16	18	19		21	22		27	28	29	30	31			
Field	TSEC1M		TSEC2M			—			TLE	LALE	—				

**Figure 4-4. Reset Configuration Word High Register (RCWHR)**

Table 4-11 defines the reset configuration word high bit fields.

**Table 4-11. Reset Configuration Word High Bit Settings**

Bits	Name	Description								
0	PCIHOST	PCI host mode. See <a href="#">Section 4.3.2.2.1, “PCI Host/Agent Configuration,”</a> for more information.								
1	—	Reserved, should be cleared.								
2	PCIARB	<p>PCI internal arbiter mode. Enables the on-chip PCI arbiter.</p> <p>0 On-chip PCI arbiter is disabled. External arbitration is required. 1 On-chip PCI arbiter is enabled.</p> <p>The value of PCIARB also defines the function of the PCI arbitration signals that are multiplexed with CompactPCI signals, as follows:</p> <table border="1" data-bbox="699 632 1174 863"> <thead> <tr> <th>Pin Function When PCIARB = 0</th> <th>Pin Function When PCIARB = 1</th> </tr> </thead> <tbody> <tr> <td>CPCI_HS_ES</td> <td><math>\overline{\text{PCI\_REQ}}[1]</math></td> </tr> <tr> <td>CPCI_HS_LED</td> <td><math>\overline{\text{PCI\_GNT}}[1]</math></td> </tr> <tr> <td>CPCI_HS_ENUM</td> <td><math>\overline{\text{PCI\_GNT}}[2]</math></td> </tr> </tbody> </table>	Pin Function When PCIARB = 0	Pin Function When PCIARB = 1	CPCI_HS_ES	$\overline{\text{PCI\_REQ}}[1]$	CPCI_HS_LED	$\overline{\text{PCI\_GNT}}[1]$	CPCI_HS_ENUM	$\overline{\text{PCI\_GNT}}[2]$
Pin Function When PCIARB = 0	Pin Function When PCIARB = 1									
CPCI_HS_ES	$\overline{\text{PCI\_REQ}}[1]$									
CPCI_HS_LED	$\overline{\text{PCI\_GNT}}[1]$									
CPCI_HS_ENUM	$\overline{\text{PCI\_GNT}}[2]$									
3	—	Reserved, should be cleared.								
4	COREDIS	<p>Core disable mode. Specifies the e300 core mode out of reset. If COREDIS is set, the core cannot fetch boot code until it is configured by an external master. The external master frees the core to boot by clearing the COREDIS bit in the arbiter configuration register as described in <a href="#">Section 6.2.1, “Arbiter Configuration Register (ACR).”</a></p> <p>This bit must be set when the boot sequencer is enabled to initiate the device (BOOTSEQ is not 0b00). Otherwise, unpredictable behavior occurs.</p> <p>0 The core can boot without waiting for configuration by an external master. 1 Core boot holdoff mode. The core is prevented from booting until it is configured by an external master.</p>								
5	BMS	Boot memory space. See <a href="#">Section 4.3.2.2.2, “Boot Memory Space (BMS),”</a> for more information.								
6–7	BOOTSEQ	Boot sequencer configuration. See <a href="#">Section 4.3.2.2.3, “Boot Sequencer Configuration,”</a> for more information.								
8	SWEN	Software watchdog enable. Selects whether the software watchdog is enabled to start counting down immediately when coming out of reset. The user can override this value by writing to the system watchdog control register (SWCRR[SWEN]) during system initialization. 0 Disabled 1 Enabled								
9–11	ROMLOC	Boot ROM interface location. This bit combined with bit RLEXT determines where the device boots from. See <a href="#">Section 4.3.2.2.4, “Boot ROM Location,”</a> for more information.								

Table 4-11. Reset Configuration Word High Bit Settings (continued)

Bits	Name	Description
12–13	RLEXT	Boot ROM location extension. This bit combined with bit ROMLOC determines where the device boots from. See <a href="#">Section 4.3.2.2.4, “Boot ROM Location,”</a> for more information. 00 Legacy mode—allows for booting from on-chip peripherals. Refer to <a href="#">Table 4-15</a> for more information. 01 NAND Flash mode—allows for booting from NAND flash devices. Refer to <a href="#">Table 4-15</a> for more information. 10 Reserved 11 Reserved
14–15	—	Reserved, should be cleared.
16–18	TSEC1M	TSEC1 mode. See <a href="#">Section 4.3.2.2.5, “eTSEC1 Mode,”</a> for more information.
19–21	TSEC2M	TSEC2 mode. See <a href="#">Section 4.3.2.2.6, “eTSEC2 Mode,”</a> for more information.
22–27	—	Reserved, should be cleared.
28	TLE	True little-endian. See <a href="#">Section 4.3.2.2.7, “e300 Core True Little-Endian,”</a> for more information.
29	LALE	Local bus LALE signal timing. See <a href="#">Section 4.3.2.2.8, “LALE Configuration,”</a> for more information.
30–31	—	Reserved, should be cleared.

#### 4.3.2.2.1 PCI Host/Agent Configuration

The PCIHOST configuration parameter, shown in [Table 4-12](#), configures the device to act as a PCI host or as a PCI agent device. In host mode, the device can immediately master transactions to the PCI interface. If the device is a PCI agent device, the device is disabled from mastering PCI transactions until the external host enables it to do so. The external host does this by setting the control registers of the device’s interfaces appropriately. See details in the PCI programming model described in [Section 13.3, “Memory Map/Register Definitions.”](#)

Table 4-12. PCI Host/Agent Configuration

RCWHR Bit	Field Name	Value (Binary)	Meaning
0	PCIHOST	0	The device acts as a PCI agent device.
		1	The device acts as the host processor (default).

#### NOTE

If the device is a PCI agent, and the e300 core is not in holdoff mode (as described in [Section 4.3.2.2, “Reset Configuration Word High Register \(RCWHR\)”](#)), the boot ROM should not be located on the PCI interface because the device is not enabled to master reads onto the PCI bus.

### 4.3.2.2.2 Boot Memory Space (BMS)

BMS defines the initial value of the e300 core MSR[IP] bit, which specifies the location of the interrupt vectors (including the hard reset exception vector). The device defines the default boot ROM memory space to be 8 Mbytes at addresses 0x0000\_0000 to 0x007F\_FFFF or 0xFF80\_0000 to 0xFFFF\_FFFF. When the core comes out of reset, if it is enabled to boot, it begins fetching boot code from one of two addresses: 0x0000\_0100 or 0xFFFF0\_0100, and exceptions are vectored to physical addresses 0x000n\_nnnn or 0xFFFFn\_nnnn appropriately. This bit specifies whether an interrupt vector offset is prepended with 0xFFFF or 0x000. In the description below, *n\_nnnn* is the offset of the exception vector.

The boot memory space reset configuration word field, shown in [Table 4-13](#), specifies both the device boot ROM address window and the initial e300 core boot address.

**Table 4-13. Boot Memory Space**

RCWHR Bit	Field Name	Value (Binary)	Meaning
5	BMS	0	Boot memory space is 8 Mbytes at 0x0000_0000 to 0x007F_FFFF. e300 core register MSR[IP] initial value is 0b0. The core, if enabled to boot, begins fetching boot code from address 0x0000_0100 and exceptions are vectored to the physical address of 0x000n_nnnn.
		1	Boot memory space is 8 Mbytes at 0xFF80_0000 to 0xFFFF_FFFF. e300 core register MSR[IP] initial value is 0b1. The core, if enabled to boot, begins fetching boot code from address 0xFFFF0_0100 and exceptions are vectored to the physical address of 0xFFFFn_nnnn.

### 4.3.2.2.3 Boot Sequencer Configuration

The boot sequencer configuration options, shown in [Table 4-14](#), allow the boot sequencer to load configuration data from the serial ROM located on the I<sup>2</sup>C port before the host tries to configure the device. These options also specify normal or extended I<sup>2</sup>C addressing modes. See [Section 17.4.5, “Boot Sequencer Mode.”](#)

**Table 4-14. Boot Sequencer Configuration**

RCWHR Bits	Field Name	Value (Binary)	Meaning
6–7	BOOTSEQ	00	Boot sequencer is disabled. No I <sup>2</sup> C ROM is accessed.
		01	Normal I <sup>2</sup> C addressing mode is used. Boot sequencer is enabled and loads configuration information from a ROM on the I <sup>2</sup> C interface. A valid ROM must be present.
		10	Extended I <sup>2</sup> C addressing mode is used. Boot sequencer is enabled and loads configuration information from a ROM on the I <sup>2</sup> C interface. A valid ROM must be present.
		11	Reserved, should be cleared.

**NOTE**

When the boot sequencer is enabled, the e300 core must be prevented from fetching boot code, by setting the core disable reset configuration word field (COREDIS) as described in [Section 4.3.2.2, “Reset Configuration Word High Register \(RCWHR\).”](#) If the e300 core is required to proceed, the boot sequencer should enable boot vector fetch by clearing ACR[COREDIS] as described in [Section 6.2.1, “Arbiter Configuration Register \(ACR\).”](#)

**4.3.2.2.4 Boot ROM Location**

The device defines the default boot ROM address range to be 8 Mbytes at addresses 0x0000\_0000 to 0x007F\_FFFF or 0xFF80\_0000 to 0xFFFF\_FFFF (selected by the BMS reset configuration word field). However, the on-chip peripheral that manages these boot ROM accesses can be selected at power up.

The boot ROM location reset configuration word field, shown in [Table 4-15](#), establishes the location of boot ROM. The exact boot ROM location table to be used is defined by the setting of RCWHR[RLEXT] bits, as shown in [Table 4-11](#). Accesses to the boot vector and the default boot ROM region of the local address map are directed to the interface specified by this field.

**Table 4-15. Boot ROM Location**

RCWHR Bits	Field Name	Value (Binary)	Meaning	
			Legacy Mode (RLEXT = 00)	NAND Flash Mode (RLEXT = 01)
9–11	ROMLOC	000	DDR SDRAM	Reserved
		001	PCI	Local bus NAND Flash—8-bit small page ROM
		010	Reserved, should be cleared.	Reserved
		011	Reserved, should be cleared.	Reserved
		100	Reserved	Reserved
		101	Local bus GPCM—8-bit ROM	Local bus NAND Flash—8-bit large page ROM
		110	Local bus GPCM—16-bit ROM	Reserved
		111	Reserved	Reserved

The local access window of the selected boot ROM interface is enabled and initialized with the proper base address and size, as described in [Section 5.2, “Local Memory Map Overview and Example.”](#)



#### 4.3.2.2.5 eTSEC1 Mode

The TSEC1 mode reset configuration word field, shown in [Table 4-16](#), selects the protocol used by the eTSEC1 controller (enhanced three-speed Ethernet controller interface).

**Table 4-16. eTSEC1 Mode Configuration**

Reset Configuration Word High Register (RCWHR) Bits	Field Name	Value (Binary)	Meaning
16–18	TSEC1M	000	The eTSEC1 controller operates in the MII protocol, using only four transmit data signals and four receive data signals.
		001	The eTSEC1 controller operates in the RMII protocol, using only two transmit data signals and two receive data signals.
		010	Reserved
		011	The eTSEC1 controller operates in the RGMII protocol, using four transmit data signals and four receive data signals.
		100	Reserved
		101	The eTSEC1 controller operates in the RTBI protocol, using only four transmit data signals and four receive data signals.
		110	The eTSEC1 controller operates in the SGMII protocol, using the on-chip PHY.
		111	Reserved

#### NOTE

The reset value of the system I/O configuration register high (SICRH) depends on the reset configuration word high TSEC1M field setting. This is used to avoid contention in systems not using the RTBI modes. In non-TBI modes, device signals which have additional functions are set to be in a non-TSEC function, thus not driven during and after reset. The function of these signals can be changed by writing to this register during system initialization. See [Section 5.3.2.6, “System I/O Configuration Register High \(SICRH\).”](#)

#### 4.3.2.2.6 eTSEC2 Mode

The eTSEC2 mode reset configuration word field, shown in [Table 4-17](#), selects the protocol used by the eTSEC2 controller (enhanced three-speed Ethernet controller interface).

**Table 4-17. eTSEC2 Mode Configuration**

Reset Configuration Word High Register (RCWHR) Bits	Field Name	Value (Binary)	Meaning
19–21	TSEC2M	000	The eTSEC2 controller operates in the MII protocol, using only four transmit data signals and four receive data signals.
		001	The eTSEC2 controller operates in the RMII protocol, using only two transmit data signals and two receive data signals.
		010	Reserved
		011	The eTSEC2 controller operates in the RGMII protocol, using four transmit data signals and four receive data signals.
		100	Reserved
		101	The eTSEC2 controller operates in the RTBI protocol, using only four transmit data signals and four receive data signals.
		110	The eTSEC2 controller operates in the SGMII protocol, using the on-chip PHY.
		111	Reserved

#### NOTE

The reset value of the system I/O configuration register high (SICRH) depends on the setting of TSEC2M. This is used to avoid contention in systems not using TBI or RTBI modes. In non-TBI modes, device signals which have additional functions are set to be in a non-TSEC function, thus not driven during and after reset. The function of these signals can be changed by writing to this register during system initialization. See [Section 5.3.2.6, “System I/O Configuration Register High \(SICRH\).”](#)

#### 4.3.2.2.7 e300 Core True Little-Endian

The true little endian reset configuration word field, shown in [Table 4-18](#), selects whether the e300 core operates in big-endian mode or true little-endian mode at reset.

**Table 4-18. e300 Core True Little-Endian**

Reset Configuration Word High Register (RCWHR) Bit	Field Name	Value (Binary)	Meaning
28	TLE	0	Big-endian mode
		1	True little-endian mode

### 4.3.2.2.8 LALE Configuration

The LALE reset configuration word field, shown in [Table 4-19](#), configures the timing of the local bus LALE signal. Refer to the hardware specifications for specific timing information.

**Table 4-19. LALE Configuration**

Reset Configuration Word High Register (RCWHR) Bit	Field Name	Value (Binary)	Meaning
29	LALE	0	Normal LALE timing
		1	LALE is negated 1/2 a LBC_controller_clk earlier

## 4.3.3 Loading the Reset Configuration Words

The device loads the reset configuration words from a local bus EEPROM, a local bus NAND Flash, or an I<sup>2</sup>C serial EEPROM, or uses hard-coded configuration, as selected by the reset configuration inputs described in [Section 4.3.1, “Reset Configuration Signals.”](#) The following sections describe each of these options.

### 4.3.3.1 Loading from Local Bus

The reset configuration words are assumed to reside in an EEPROM or NOR Flash or NAND Flash device connected to  $\overline{\text{LCS0}}$  of the device local bus. Because the port size of this EEPROM is unknown, the device reads all configuration words byte-by-byte only from locations that are independent of port size.

[Table 4-20](#) shows addresses that should be used to contain the reset configuration words. Byte addresses that do not appear in this table have no effect on the configuration of the device. The values of the bytes in [Table 4-20](#) are always read on byte lane LAD[0:7] regardless of the port size.

**Table 4-20. Local Bus Configuration EEPROM Addresses**

Reset Configuration Word	Bits [0:7] Address	Bits [8:15] Address	Bits [16:23] Address	Bits [24:31] Address
Low	0x00	0x08	0x10	0x18
High	0x20	0x28	0x30	0x38

[Table 4-21](#) shows the data structure of the local bus device containing the reset configuration words (RCWL and RCWH).

**Table 4-21. Local Bus Reset Configuration Words Data Structure**

EEPROM Address	EEPROM Data Bits			
	[0:7]	[8:15]	[16:23]	[24:31]
0x00	RCWL[0:7]			
0x04				
0x08	RCWL[8:15]			
0x0C				

**Table 4-21. Local Bus Reset Configuration Words Data Structure (continued)**

EEPROM Address	EEPROM Data Bits			
	[0:7]	[8:15]	[16:23]	[24:31]
0x10	RCWL[16:23]			
0x14				
0x18	RCWL[24:31]			
0x1C				
0x20	RCWH[0:7]			
0x24				
0x28	RCWH[8:15]			
0x2C				
0x30	RCWH[16:23]			
0x34				
0x38	RCWH[24:31]			
0x3C				

#### 4.3.3.1.1 Local Bus Controller Setting

The device will use GPCM to load the reset configuration from EEPROM or NOR Flash. The device will read 64 bytes in this case. The local bus controller's registers setting will be set according to [Table 4-22](#).

The device will use FCM to load the reset configuration from NAND Flash. The device will read 512 bytes if small size NAND Flash is used, or 2048 bytes if large page NAND Flash is used. The local bus controller's registers setting will be set according to [Table 4-22](#).

The device will use PCI\_SYNC\_IN clock to generate the internal LCLK, which will run at half the frequency of PCI\_SYNC\_IN.

**Table 4-22. Local Bus Controller Setting When Loading RCW**

CFG_RESET_SOURCE	Meaning	BR0[PS]	BR0[MSEL]	OR0[SCY]	OR0[PGS]
0000	NOR Flash	10	000	1111	NA
0001	NAND Flash, 8 bit, small page	01	001	0010	0
0101	NAND Flash, 8 bit, large page	01	001	0010	1

### 4.3.3.2 Loading from I<sup>2</sup>C EEPROM

The device is capable of loading the reset configuration word from the I<sup>2</sup>C interfaces (the device has two I<sup>2</sup>C interfaces, but only I<sup>2</sup>C #1 can be used for this purpose). If the device is configured to load the reset configuration word from the I<sup>2</sup>C interface, according to the reset configuration input signals, it uses the I<sup>2</sup>C unit boot sequencer in a special mode. In this mode, the I<sup>2</sup>C boot sequencer is activated while the rest of the device is still in reset state ( $\overline{\text{HRESET}}$  asserted) to load the reset configuration words from an I<sup>2</sup>C serial EEPROM.

Note that this does not prevent using the I<sup>2</sup>C boot sequencer to initiate the device in the normal functional mode after reset state has completed. The only restriction is that the first two EEPROM data structures contain dedicated reset information.

#### 4.3.3.2.1 Using the Boot Sequencer Reset Configuration

For a detailed description about the I<sup>2</sup>C interface and the boot sequencer refer to [Section 17.4.5, “Boot Sequencer Mode.”](#)

#### NOTE

When reset configuration words are loaded from an I<sup>2</sup>C EEPROM, an I<sup>2</sup>C serial EEPROM of extended addressing type must be used.

If the I<sup>2</sup>C interface is used for loading the reset configuration words, the I<sup>2</sup>C module addresses the EEPROM and reads the first two data structures (after reading the preamble). Upon being read, the reset configuration words are latched inside the device and the I<sup>2</sup>C module enters its reset state until  $\overline{\text{HRESET}}$  is negated. There should be no other I<sup>2</sup>C traffic when the boot sequencer is active.

After  $\overline{\text{HRESET}}$  is negated, the functional boot sequencer, in extended I<sup>2</sup>C addressing mode, may be activated if the BOOTSEQ field of the reset configuration word high is set to 0b10.

#### 4.3.3.2.2 EEPROM Calling Address

The device uses 0b101\_0000 for the EEPROM calling address. The EEPROM to be addressed must contain the reset configuration information and be programmed to respond to this address. No additional EEPROMs are accessed by the boot sequencer in reset configuration mode.

#### 4.3.3.2.3 EEPROM Data Format in Reset Configuration Mode

The I<sup>2</sup>C module expects that a particular data format be used for data in the EEPROM. A preamble should be the first 3 bytes programmed into the EEPROM. It should have a value of 0xAA55AA. The I<sup>2</sup>C module checks to ensure that this preamble is correctly detected before proceeding further. Following the preamble, there should be the two reset configuration words, programmed according to a particular format, as shown in [Figure 4-5](#).

The first 3 bytes hold the attributes and address offset. The addresses of the two reset configuration words must be programmed to the offset of the reset configuration word low register (RCWLR) and reset configuration word high register (RCWHR) respectively (see [Section 4.5.1.1, “Reset Configuration Word Low Register \(RCWLR\),”](#) and [Section 4.5.1.2, “Reset Configuration Word High Register \(RCWHR\)”](#)).

The attributes should be programmed as follows: alternate configuration space (ACS) should be cleared (0b0), byte enables should be all ones, and continue (CONT) should be set.

After the first 3 bytes, 4 bytes of data should hold the desired value of the reset configuration word. The boot sequencer assumes that a big-endian address is stored in the EEPROM.

IMMRBAR value is prepended to the EEPROM address to generate the complete memory-mapped register's address.

When the I<sup>2</sup>C operates in reset configuration mode, the cyclic redundancy check (CRC) is ignored, as well as any registers following the first two reset configuration words.

0	1	4	5	6	7
ACS (0)	BYTE_EN (1111)			CONT (1)	RCWLR ADDR[12–13]
RCWLR ADDR[14:21]					
RCWLR ADDR[22:29]					
Reset configuration word low [0–7]					
Reset configuration word low [8–15]					
Reset configuration word low [16–23]					
Reset configuration word low [24–31]					
ACS (0)	BYTE_EN (1111)			CONT (1)	RCWHR ADDR[12–13]
RCWHR ADDR[14–21]					
RCWHR ADDR[22–29]					
Reset configuration word high [0–7]					
Reset configuration word high [8–15]					
Reset configuration word high [16–23]					
Reset configuration word high [24–31]					

**Figure 4-5. EEPROM Data Format for Reset Configuration Words Preload Command**

Figure 4-6 shows an example of the EEPROM contents, including the preamble, reset configuration words and additional initialization data, and CRC. In this example, it is assumed that the EEPROM contains information additional to the reset configuration words, which should be loaded in the functional state after the device completes its reset flow.

0	1	2	3	4	5	6	7	
1	0	1	0	1	0	1	0	Preamble
0	1	0	1	0	1	0	1	
1	0	1	0	1	0	1	0	
0	1	1	1	1	1	RCWLR ADDR[12:13]		
RCWLR ADDR[14–21]								Reset configuration word low preload command
RCWLR ADDR[22–29]								
Reset configuration word low [0–7]								
Reset configuration word low [8–15]								
Reset configuration word low [16–23]								
Reset configuration word low [24–31]								
0	1	1	1	1	1	RCWHR ADDR[12:13]		Reset configuration word high preload command
RCWHR ADDR[14–21]								
RCWHR ADDR[22–29]								
Reset configuration word high [0–7]								
Reset configuration word high [8–15]								
Reset configuration word high [16–23]								
Reset configuration word high [24–31]								
*								
ACS	BYTE_EN				1	ADDR[12–13]		Last configuration preload command
ADDR[14–21]								
ADDR[22–29]								
DATA[0–7]								
DATA[8–15]								
DATA[16–23]								
DATA[24–31]								
0	0	0	0	0	0	0	0	End command
0	0	0	0	0	0	0	0	
0	0	0	0	0	0	0	0	
CRC[0–7]								Cyclic redundancy check
CRC[8–15]								
CRC[16–23]								
CRC[24–31]								

Figure 4-6. EEPROM Contents

#### 4.3.3.2.4 Reset Configuration Load Fail

Failure of reset configuration load by the I<sup>2</sup>C boot sequencer can be caused by an incorrect EEPROM data structure or I<sup>2</sup>C bus problem. If a reset configuration load failure occurs, due to preamble fail or any other

I<sup>2</sup>C bus error detection, the device will continuously attempt to reload the hard reset configuration words from the I<sup>2</sup>C bus. The device does not negate  $\overline{\text{HRESET}}$  and remains in hard reset state until the HRCWs are successfully loaded or the PORESET flow is restarted.

### 4.3.3.3 Default Reset Configuration Words

If the device is configured not to load the reset configuration words from NOR Flash, NAND Flash, or an I<sup>2</sup>C EEPROM, it can also be initialized with one of five hard-coded default options, selected by the reset configuration input signals, CFG\_RESET\_SOURCE[0:3]. In this mode, the device is assumed to be a PCI agent, and therefore only clock modes differ among the four options.

The reset configuration words are driven internally with the values shown in [Table 4-23](#) and [Table 4-24](#).

#### NOTE

In this mode the device is also configured to accept PCI configuration cycles when completing its reset sequence (In PCI function configuration register, the CFG\_LOCK bit is cleared). In addition, the inbound window size of the PCI inbound window attribute registers (PIWAR<sub>n</sub>[IWS]) is set to 0b010011, defining 1-Mbyte ( $2^{(19+1)}$ ) memory windows. See [Section 13.3.3.24, “PCI Function Configuration Register.”](#)

**Table 4-23. Hard Coded Reset Configuration Word Low Fields Values**

RCWL Bits:	0	1	2–3	4–7	8	9–15	16–31
Field:	LBCM	DDRCM	Res	SPMF	Res	COREPLL	Res
Meaning:	LBC controller clock:	DDR controller clock:	—	<i>csb_clk</i> : PCI_CLK ratio SPMF:1	—	Core clock: <i>csb_clk</i> ratio	—
CFG_RESET_SOURCE Value	0 1:1 1 2:1	<i>csb_clk</i> 0 1:1 1 2:1					
1000	0	1	10	0100	0	0000100	16'b0
1001	0	1	10	0010	0	0000101	16'b0
1010	0	1	10	0101	0	0000011	16'b0
1011	0	1	10	0010	0	0000100	16'b0
1100	0	1	10	0101	0	0000100	16'b0

[Table 4-24](#) defines the hard-coded reset configuration word high fields values. These values select hard-coded reset configuration words options, as described in [Section 4.3.1.1, “Reset Configuration Word Source.”](#)



Table 4-24. Hard-Coded Reset Configuration Word High Field Values

Bits	Name	Field Values when CFG_RESET_SOURCE[0:3] = 1000–1100				Meaning
		1000	1001	1010	1100	
0	PCIHOST	0				PCI agent mode
1	Reserved	1				—
2	PCIARB	0				External arbiter is used
3	Reserved	0				—
4	COREDIS	1				e300 core is disabled (boot holdoff)
5	BMS	1				Boot memory space is 0xFF80_0000–0xFFFF_FFFF. MSR[IP] initial value is 0b1.
6–7	BOOTSEQ	00				Boot sequencer is disabled.
8	SWEN	0				Software watchdog disabled.
9–11	ROMLOC	000				Boot ROM interface location.
12–13	RLEXT	00				Legacy mode.
14–15	Reserved	00				—
16–18	TSEC1M	101	011	000	001	000 = MII mode 001 = RMII 011 = RGMII mode 101 = RTBI mode
19–21	TSEC2M	000	101	001	101	
22–27	Reserved	000000				—
28	TLE	0				Big-endian mode
29	LALE	0				Normal timing
30–31	Reserved	00				—

#### 4.3.3.3.1 Examples for Hard-Coded Reset Configuration Words Usage

Examples for various clock modes are listed in [Table 4-25](#).

Table 4-25. Examples For Hard-Coded Reset Configuration Words Usage

CFG_RESET_SOURCE[0:3]	1000	1001	1010	1011	1100
PCI_CLK (MHz)	33	66	33	66	33
<i>csb_clk</i> (MHz)	133	133	167	133	167
DDR controller clock (MHz)	266	266	333	266	333
Core clock (MHz)	266	333	250	266	333

## 4.4 Clocking

The following external clock sources are utilized on the MPC8313E:

- System clock (SYS\_CLK\_IN or PCI\_CLK)
- Ethernet clock (GTX\_CLK125 for eTSEC $n$  or SD\_REF\_CLK/SD\_REF\_CLK for SGMII PHY interface)
- USB clock (USB\_CLK\_IN for USB PHY)
- Real-time clock (RTC\_CLK)

All clock inputs can be supplied using an external canned oscillator, a clock generation chip, or some other source that provides a standard CMOS square wave input.

The system and USB clock sources can alternatively be provided using an external crystal (on-chip crystal oscillator provided). Separate crystal (CR) inputs/outputs are provided for this case. The crystal oscillator dissipates 50 mW of current, so this may not be the best choice for very low power applications. When a crystal is not used, a square wave clock input must be supplied.

If SYS\_CLK\_IN is used as the system clock, the SYS\_CR\_CLK\_IN (crystal input) must be tied to quiet ground; otherwise, SYS\_CLK\_IN must be tied to quiet ground. Similarly, if USB\_CLK\_IN is used as the USB clock, USB\_CR\_CLK\_IN (crystal input) must be tied to quiet ground; otherwise, USB\_CLK\_IN must be tied to quiet ground.

Figure 4-7 shows the internal distribution of clocks within the device.




Figure 4-7. Clock Subsystem Block Diagram

The primary clock input to this device is PCI\_CLK. This clock is the reference to the system APLL.

### 4.4.1 Clocking in PCI Host Mode

When the device is configured as a PCI host device ( $RCWH[PCIHOST] = 1$ ),  $SYS\_CLK\_IN$  is the primary input clock.  $SYS\_CLK\_IN$  feeds the PCI clock divider ( $\div 2$ ) and the  $PCI\_SYNC\_OUT$  and  $PCI\_CLK\_OUT$  multiplexors. The  $\overline{CFG\_CLKIN\_DIV}$  configuration input selects whether  $SYS\_CLK\_IN$  or  $SYS\_CLK\_IN/2$  is driven out on the  $PCI\_SYNC\_OUT$  and  $PCI\_CLK\_OUT$  signals.

$PCI\_SYNC\_OUT$  is connected externally to  $PCI\_SYNC\_IN$  to allow the internal clock subsystem to synchronize to the system PCI clocks.  $PCI\_SYNC\_OUT$  must be connected properly to  $PCI\_SYNC\_IN$ , with equal delay to all PCI agent devices in the system.

#### 4.4.1.1 PCI Clock Outputs ( $PCI\_CLK\_OUT[0:2]$ )

When the device is configured as a PCI host, it provides three clock output signals,  $PCI\_CLK\_OUT[0:2]$ , for external PCI agents.

When the device comes out of reset, the PCI clock outputs are disabled and are actively driven to a steady low state. Each of the individual clock outputs can be enabled (enable toggling of the clock) by setting its corresponding  $OCCR[PCICOEn]$  bit. All output clocks are phase aligned to each other and to  $PCI\_SYNC\_OUT$ .

### 4.4.2 Clocking In PCI Agent Mode

When the device is configured as a PCI agent,  $PCI\_CLK$  is the primary input clock. In agent mode, the  $SYS\_CLK\_IN$  signal may not be used. If it is unused,  $SYS\_CLK\_IN$  should be tied to GND, and the clock output signals,  $PCI\_CLK\_OUTn$  and  $PCI\_SYNC\_OUT$ , are not used.

In agent mode, the  $\overline{CFG\_CLKIN\_DIV}$  configuration input can be used to double the internal clock frequencies, if sampled as 0 during PORESET assertion. This feature is useful if a fixed internal frequency is desired regardless of whether the PCI clock is running at 33 or 66 MHz. PCI specifications require that the signal M66EN provides the PCI clock frequency information.

### 4.4.3 System Clock Domains

As shown in Figure 4-7, the primary clock input ( $PCI\_CLK/PCI\_SYNC\_IN$ ) frequency is multiplied up by the system phase-locked loop (PLL) and the clock unit to create three major clock domains:

- The coherent system bus clock (*csb\_clk*)
- The internal clock for the DDR controller (*ddr\_clk*)
- The internal clock for the local bus interface unit (*lbc\_clk*)

The *csb\_clk* frequency is derived from a complex set of factors that can be simplified into the following equation:

$$csb\_clk = [PCI\_SYNC\_IN \times (1 + \overline{CFG\_CLKIN\_DIV})] \times SPMF$$

In PCI host mode,  $PCI\_SYNC\_IN \times (1 + \overline{CFG\_CLKIN\_DIV})$  is the  $SYS\_CLK\_IN$  frequency.

The *csb\_clk* serves as the clock input to the e300 core. A second PLL inside the core multiplies up the *csb\_clk* frequency to create the internal clock for the core (*core\_clk*). The system and core PLL multipliers

are selected by the SPMF and COREPLL fields in the reset configuration word low (RCWL), which is loaded at power-on reset or by one of the hard-coded reset options. See [Section 4.3, “Reset Configuration.”](#)

The DDR SDRAM memory controller will operate with a frequency equal to twice the frequency of *csb\_clk*. Note that *ddr\_clk* is not the external memory bus frequency; *ddr\_clk* passes through the DDR clock divider ( $\div 2$ ) to create the differential DDR memory bus clock outputs (MCK and  $\overline{\text{MCK}}$ ). However, the data rate is the same frequency as *ddr\_clk*.

The local bus memory controller will operate with a frequency equal to the frequency of *csb\_clk*. Note that *lbc\_clk* is not the external local bus frequency; *lbc\_clk* passes through the LBC clock divider to create the external local bus clock output (LCLK[0:1]). The LBC clock divider ratio is controlled by LCCR[CLKDIV]. See [Section 10.1.3.1, “eLBC Bus Clock and Clock Ratios,”](#) for more information.

In addition, some of the internal units may be required to be shut off or operate at lower frequency than the *csb\_clk* frequency. These units have a default clock ratio that can be configured by a memory mapped register after the device comes out of reset. [Table 4-26](#) specifies which units have a configurable clock frequency. Refer to [Section 4.5.2.3, “System Clock Control Register \(SCCR\).”](#)

**Table 4-26. Configurable Clock Units**

Unit	Default Frequency	Options
eTSEC1 and eTSEC2	<i>csb_clk</i>	Off, <i>csb_clk</i> , <i>csb_clk/2</i> , <i>csb_clk/3</i>
Security core, I <sup>2</sup> C1	<i>csb_clk</i>	Off, <i>csb_clk</i> , <i>csb_clk/2</i> , <i>csb_clk/3</i>
USB DR	<i>csb_clk</i>	Off, <i>csb_clk</i> , <i>csb_clk/2</i> , <i>csb_clk/3</i>
PCI and DMA complex	<i>csb_clk</i>	Off, <i>csb_clk</i>

#### NOTE

The clock ratios of these units must be set before they are accessed.

eTSEC1 and eTSEC2 share the same clock control, and therefore the same clock ratio. They can be independently switched off.

#### NOTE

A portion of the eTSEC<sub>n</sub> and USB DR controller run at the line rate on a clock provided by the SGMII and USB PHY blocks, respectively. Synchronization between the line rate clock and the *csb\_clk* is provided in these controllers.

### 4.4.4 USB Clocking

If the on-chip USB PHY is utilized, the reference input can be provided externally using a separate clock source, either a crystal or an external oscillator. The PHY supplies the clock to the USB DR controller in UTMI mode (when the on-chip PHY is used). Synchronization between the PHY clock domain and the CSB clock domain occurs in the USB controller.

An option is provided to supply the USB reference clock from the SYS\_CLK\_IN or SYS\_CR\_CLK\_IN inputs. This allows for a single crystal or clock input to supply both system and USB references. The USB reference clock can be provided with a divide by 1 or 2 from these inputs (see [Figure 4-7](#)). When using the

single crystal option, the frequency for SYS\_CLK\_IN must be chosen such that the USB reference will be 24 or 48 MHz when utilizing the divide by 1 or 2 option, that is, the SYS\_CLK\_IN must be 24 or 48 MHz.

If this option is used in PCI agent mode, the PCI clock supplied to device must still be chosen at the appropriate frequencies mentioned above. In this case the PCI source clock would be tied to both SYS\_CLK\_IN and PCI\_CLK inputs. The clock source will need to meet the input clock specifications for both SYS\_CLK\_IN and USB\_CLK\_IN.

For more details refer to [Chapter 16, “Universal Serial Bus Interface.”](#)

#### 4.4.5 Ethernet Clocking

The device contains an integrated high speed serial/deserializing (SerDes) PHY block that provides an SGMII interface to an external PHY. The SGMII PHY interface has its own PLL and requires a reference clock which is 1 V signal that can either be single ended or differential. The SerDes PHY generates its own transmit and receive sampling clock. The receive clock is re-created from the receive data.

When running in RGMII or RTBI modes (not using the SerDes) the Gigabit reference clock is the GTX\_CLK125 input on the eTSEC1 interface which is common to both eTSECs. This can either be a 2.5- or 3.3-V signal.

For more information refer to [Chapter 15, “Enhanced Three-Speed Ethernet Controllers.”](#)

#### 4.4.6 Real-Time Clock (RTC)

The source for the RTC can come from the internal system clock (csb\_clk) through a divider circuit or from an off-chip source. For more information refer to [Section 5.5, “Real Time Clock Module \(RTC\).”](#)

### 4.5 Memory Map/Register Definitions

This section presents the memory maps and register descriptions for both reset and clocking.

#### 4.5.1 Reset Configuration Register Descriptions

The reset configuration and status registers are shown in [Table 4-27](#).

**Table 4-27. Reset Configuration and Status Registers Memory Map**

Address	Register	Access	Reset	Section/Page
0x0_0900	Reset configuration word low register (RCWLR)	R	0x0000_0000	<a href="#">4.5.1.1/4-33</a>
0x0_0904	Reset configuration word high register (RCWHR)	R	0x0000_0000	<a href="#">4.5.1.2/4-33</a>
0x0_0908	Reserved, should be cleared	—	—	—
0x0_090C	Reserved, should be cleared	—	—	—
0x0_0910	Reset status register (RSR)	R/W	0x0000_0000	<a href="#">4.5.1.3/4-33</a>
0x0_0914	Reset mode register (RMR)	R/W	0x0000_0000	<a href="#">4.5.1.4/4-35</a>
0x0_0918	Reset protection register (RPR)	R/W	0x0000_0000	<a href="#">4.5.1.5/4-35</a>

**Table 4-27. Reset Configuration and Status Registers Memory Map (continued)**

Address	Register	Access	Reset	Section/Page
0x0_091C	Reset control register (RCR)	R/W	0x0000_0000	4.5.1.6/4-36
0x0_0920	Reset control enable register (RCER)	R/W	0x0000_0000	4.5.1.7/4-37
0x0_0924– 0x0_09FC	Reserved, should be cleared.	—	—	—

### 4.5.1.1 Reset Configuration Word Low Register (RCWLR)

The reset configuration word low register (RCWLR) is shown in [Figure 4-3](#) and described in [Section 4.3.2.1, “Reset Configuration Word Low Register \(RCWLR\).”](#)

### 4.5.1.2 Reset Configuration Word High Register (RCWHR)

The reset configuration word high register (RCWHR) is shown in [Figure 4-4](#) and described in [Section 4.3.2.2, “Reset Configuration Word High Register \(RCWHR\).”](#)

### 4.5.1.3 Reset Status Register (RSR)

RSR, shown in [Figure 4-8](#), captures various reset events in the device. The RSR accumulates reset events. For example, because software watchdog expiration results in a hard reset, SWRS and HRS are all set after a software watchdog reset. This register returns to its reset value only when power-on reset occurs.

Address 0x0\_0910

Access: User read/write

	0	3	4							14	15			
R	RSTSRC			—						BSF				
W	n1			0						0				
Reset	n1			0						0				
	16	17	18	19	20	22	23	24	26	27	28	29	30	31
R	—		SWHR	—		JSRS	—		CSHR	SWRS	BMRS	—		HRS
W	—		SWHR	—		JSRS	—		CSHR	SWRS	BMRS	—		HRS
Reset	All zeros													

<sup>1</sup> The reset value of this field is determined according to the reset configuration input signals CFG\_RESET\_SOURCE[0:3] sampled during the reset flow.

**Figure 4-8. Reset Status Register (RSR)**

[Table 4-28](#) defines the reset status register bit fields.

**Table 4-28. Reset Status Register Field Descriptions**

Bits	Name	Description
0–3	RSTSRC	Reset configuration word source. Reflects the value of CFG_RESET_SOURCE input signal during the reset flow. See <a href="#">Section 4.3.1.1, “Reset Configuration Word Source,”</a> on page 4-10. Changing this field has no effect.
4–14	—	Reserved, should be cleared.

Table 4-28. Reset Status Register Field Descriptions (continued)

Bits	Name	Description
15	BSF	Boot sequencer fail. If set, indicates that the I <sup>2</sup> C boot sequencer has failed while loading the reset configuration words. Cleared by writing a 1 to it (writing zero has no effect).
16–18	—	Reserved, should be cleared.
19	SWHR	Software hard reset. If set, indicates a software hard reset. SWHR is cleared by writing a 1 to it (writing zero has no effect).
20–22	—	Reserved, should be cleared.
23	JSRS	JTAG soft reset status. Set when the JTAG reset request is set and remains set until software clears it. JSRS is cleared by writing a 1 to it (writing zero has no effect). 0 No JTAG reset event. 1 JTAG reset event.
24–26	—	Reserved, should be cleared.
27	CSHR	Check stop reset status. When the core enters a checkstop state and the checkstop reset is enabled by the RMR[CSRE], CSRS is set and it remains set until software clears it. CSRS is cleared by writing a 1 to it (writing zero has no effect). 0 No enabled check stop reset event. 1 Enabled check stop reset event.
28	SWRS	Software watchdog reset status. When a software watchdog expire event (which causes a reset) is detected, SWRS is set and remains that way until the software clears it. SWRS is cleared by writing a 1 to it (writing zero has no effect). 0 No software watchdog reset event. 1 Software watchdog reset event.
29	BMRS	Bus monitor reset status. When a bus monitor expire event (which causes a reset) is detected, BMRS is set and remains set until the software clears it. BMRS can be cleared by writing a 1 to it (writing zero has no effect). 0 No bus monitor reset event. 1 Bus monitor reset event.
30	—	Reserved
31	HRS	Hard reset status. When an external or internal hard reset event is detected, HRS is set and remains set until software clears it. HRS is cleared by writing a 1 (writing zero has no effect). 0 No hard reset event. 1 Hard reset event.



#### 4.5.1.4 Reset Mode Register (RMR)

RMR, shown in Figure 4-9, enables a hard reset sequence on the device when the e300 core enters checkstop state.




Figure 4-9. Reset Mode Register (RMR)

Table 4-29 describes the RMR fields.

Table 4-29. RMR Field Descriptions

Bits	Name	Function
0–30	—	Reserved, should be cleared.
31	CSRE	Checkstop reset enable. The core can enter checkstop mode as the result of several exception conditions. Setting CSRE configures the device to perform a hard reset sequence when the core enters checkstop state. 0 Reset not generated when core enters checkstop state. 1 Reset generated when core enters checkstop state.

#### 4.5.1.5 Reset Protection Register (RPR)

RPR, shown in Figure 4-10, prevents unintended software reset requests caused by writes to the reset control register (RCR). To disable a write to the reset control register (RCR), the user should write a 1 to RCER[CRE].




Figure 4-10. Reset Protection Register (RPR)

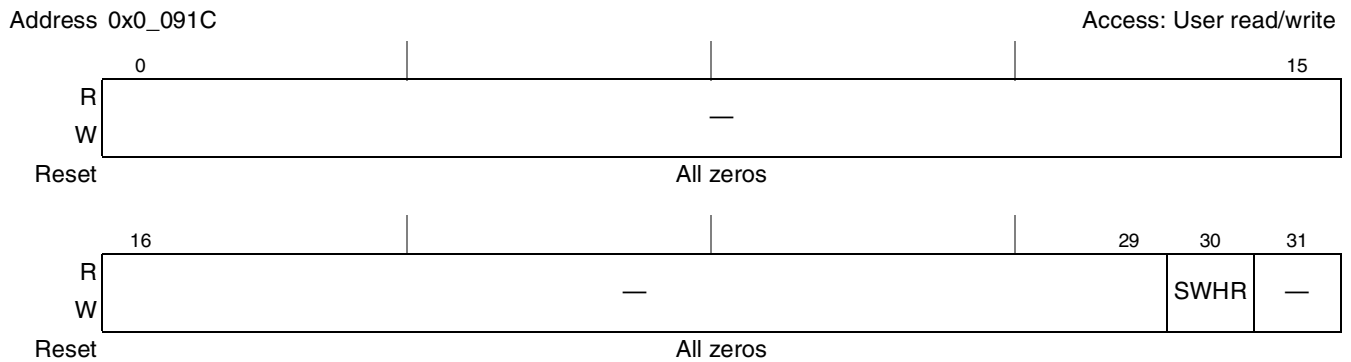
Table 4-30 defines the bit fields of RPR.

**Table 4-30. RPR Bit Descriptions**

Bits	Name	Description
0–31	RCPW	Reset control protection word. Prevents unintended software reset requests because of a write to the RCR. The user should write the value 0x5253_5445 (RSTE in ASCII) to enable. Enable indication appears in the reset control enable register (RCER[CRE]). Reading this register always returns all zeros.

### 4.5.1.6 Reset Control Register (RCR)

RCR, shown in Figure 4-11, can be used by software to initiate a soft or hard reset sequence. To allow writing to this register, the user must enable it by writing the value 0x5253\_5445 to the RPR.



**Figure 4-11. Reset Control Register (RCR)**

Table 4-31 defines the bit fields of RCR.

**Table 4-31. RCR Bit Settings**

Bits	Name	Description
0–29	—	Reserved, should be cleared.
30	SWHR	Software hard reset. Setting this bit causes the device to begin a hard reset flow. This bit returns to its reset state during the reset sequence, so reading it always returns all zeros.
31	—	Reserved. This bit should never be set.

### 4.5.1.7 Reset Control Enable Register (RCER)

RCER, shown in Figure 4-12, indicates by the CRE field that the RPR is accessed with a value that enables RCR.




Figure 4-12. Reset Control Enable Register (RCER)

Table 4-32 defines the bit fields of RCER.

Table 4-32. RCER Bit Settings

Bits	Name	Description
0–30	—	Reserved, should be cleared.
31	CRE	Control register enabled. When set, indicates that the RPR was accessed with a value that enables the RCR. Writing 1 to this bit disables the RCR and clears this bit. Writing zero has no effect.

## 4.5.2 Clock Configuration Registers

The clock configuration and status registers are shown in Table 4-33.

Table 4-33. Clock Configuration Registers Memory Map

Address	Register	Access	Reset	Section/Page
0x0_0A00	System PLL mode register (SPMR)	R	0xn <sub>nnn</sub> _n <sub>nnn</sub>	4.5.2.1/4-37
0x0_0A04	Output clock control register (OCCR)	R/W	0x0000_80C0	4.5.2.2/4-39
0x0_0A08	System clock control register (SCCR)	R/W	0x7DDF_FFFF	4.5.2.3/4-40
0x0_0A0C– 0x0_0AFC	Reserved, should be cleared	—	—	—

### 4.5.2.1 System PLL Mode Register (SPMR)

SPMR is shown in Figure 4-13, gets its values according to the  $\overline{\text{CFG\_CLKIN\_DIV}}$  reset configuration input signal and the reset configuration word low loaded during the reset flow. Note that this register is updated only during a power-on reset sequence and not by a hard reset sequence. It may hold values different than those in the RCWLR after a a hard reset sequence.

Address 0x0\_0A00

Access: Read only




Figure 4-13. System PLL Mode Register

<sup>1</sup> See Table 4-34 for reset values.

Table 4-34 defines the system PLL mode register bit fields.

Table 4-34. System PLL Mode Register Bit Settings

Bits	Name	Meaning	Description
0	LBCM	Local bus memory controller clock mode.	Section 4.3.2.1, “Reset Configuration Word Low Register (RCWLR)”
1	DDRCM	DDR SDRAM memory controller clock mode.	Section 4.3.2.1, “Reset Configuration Word Low Register (RCWLR)”
2–3	—	Reserved, should be cleared.	—
4–7	SPMF	System PLL multiplication factor	Section 4.3.2.1.1, “System PLL Configuration”
8	CKID	SYS_CLK_IN division factor. Reflects the value of CFG_CLKIN_DIV input signal during the reset flow.	Section 4.3.1.2, “SYS_CLK_IN Division”
9–15	COREPLL	Core PLL configuration.	See the hardware specifications for this device
16–31	—	Reserved, should be cleared.	—

### 4.5.2.2 Output Clock Control Register (OCCR)

The OCCR shown in Figure 4-14, controls the device output clocks. It is possible to control some output clock modes by writing to this memory mapped register as described below.

Address 0x0\_0A04

Access: Read/Write




Figure 4-14. Output Clock Control Register (OCCR)

Table 4-35 defines the bit fields of OCCR.

Table 4-35. OCCR Bit Settings

Bits	Name	Description
0	PCICOE0	PCI_CLK_OUT0 enable. 0 PCI_CLK_OUT0 signal is disabled (drive constant zero). 1 PCI_CLK_OUT0 signal is enabled to toggle.
1	PCICOE1	PCI_CLK_OUT1 enable. 0 PCI_CLK_OUT1 signal is disabled (drive constant zero). 1 PCI_CLK_OUT1 signal is enabled to toggle.
2	PCICOE2	PCI_CLK_OUT2 enable. 0 PCI_CLK_OUT2 signal is disabled (drive constant zero). 1 PCI_CLK_OUT2 signal is enabled to toggle.
3–15	—	Reserved, should be cleared
16	MCKOE, MCK_BOE	Enable/Disable MCK pin clock out 0 Disable MCK and $\overline{\text{MCK}}$ 1 Enable MCK and $\overline{\text{MCK}}$
17–23	—	Reserved
24	LCLK0OE	Enable/Disable LCLK[0] pin clock out 0 Disable LCLK[0] 1 Enable LCLK[0]
25	LCLK1OE	Enable/Disable LCLK[1] pin clock out 0 Disable LCLK[1] 1 Enable LCLK[1]
26–31	—	Reserved

### 4.5.2.3 System Clock Control Register (SCCR)

SCCR, shown in Figure 4-15, controls device units that have a configurable clock ratio.

Address 0x0\_0A08

Access: Read/Write




Figure 4-15. System Clock Control Register (SCCR)

Table 4-36 defines the bit fields of SCCR.

Table 4-36. SCCR Bit Descriptions

Bits	Name	Description
0–1	TSECCM	TSEC1 and TSEC2 clock mode. 00 Reserved. Write not allowed. If written, treated as ratio 1:1. 01 TSEC1&2 clock/ <i>csb_clk</i> ratio is 1:1. 10 TSEC1&2 clock/ <i>csb_clk</i> ratio is 1:2 ( <i>csb_clk</i> has higher frequency than TSEC1&2). 11 TSEC1&2 clock/ <i>csb_clk</i> ratio is 1:3 ( <i>csb_clk</i> has higher frequency than TSEC1&2).
2	TSEC1ON	TSEC1 clock switch off 0 TSEC1 clock is disabled. 1 TSEC1 clock is enabled and in the ratio as specified by TSECCM.
3	TSEC2ON	TSEC2 clock switch off 0 TSEC2 clock is disabled. 1 TSEC2 clock is enabled and in the ratio as specified by TSECCM.
4–5	—	Reserved
6–7	ENCCM	Encryption core and I <sup>2</sup> C1 clock mode. 00 Encryption core clock is disabled. 01 Encryption core clock/ <i>csb_clk</i> ratio is 1:1. 10 Encryption core clock/ <i>csb_clk</i> ratio is 1:2 ( <i>csb_clk</i> has higher frequency than the encryption core). 11 Encryption core clock/ <i>csb_clk</i> ratio is 1:3 ( <i>csb_clk</i> has higher frequency than the encryption core).
8–9	—	Reserved
10–11	USB DRCM	USB DR clock mode. 00 USB DR clock is disabled. 01 USB DR clock/ <i>csb_clk</i> ratio is 1:1. 10 USB DR clock/ <i>csb_clk</i> ratio is 1:2 ( <i>csb_clk</i> has higher frequency than the USB DR). <b>Note:</b> 11USB DR clock/ <i>csb_clk</i> ratio is 1:3 ( <i>csb_clk</i> has higher frequency than the USB DR).
12–14	—	Reserved

**Table 4-36. SCCR Bit Descriptions (continued)**

Bits	Name	Description
15	PCICM	PCI clock mode. Define the clock mode for all of the PCI complex - PCI and DMA. 0 PCI complex clocks are disabled. 1 PCI complex clocks are enabled.
16–31	—	Reserved





# Chapter 5

## System Configuration

### 5.1 Introduction

This chapter describes several functions that control the local access windows, system configuration, protection, and general utilities. These functions are discussed in the following sections:

- [Section 5.2, “Local Memory Map Overview and Example”](#)
- [Section 5.3, “System Configuration”](#)
- [Section 5.4, “Software Watchdog Timer \(WDT\)”](#)
- [Section 5.5, “Real Time Clock Module \(RTC\)”](#)
- [Section 5.6, “Periodic Interval Timer \(PIT\)”](#)
- [Section 5.7, “General-Purpose Timers \(GTM\)”](#)
- [Section 5.8, “Power Management Control \(PMC\)”](#)

### 5.2 Local Memory Map Overview and Example

The device provides a flexible local memory map. The local memory map refers to the 32-bit address space seen by the processor as it accesses memory and I/O space. Internal DMA engines also see this same local memory map. All memory accessed by the DDR SDRAM and local bus memory controllers exists in this memory map, as do all memory-mapped configuration, control, and status registers.

The local memory map is defined by a set of nine local access windows. Each of these windows maps a region of memory to a particular target interface, such as the DDR SDRAM controller or the PCI controller. The DSP subsystem is not operational in the MSC7104. Note that the local access windows do not perform any address translation. The size of each window can be configured from 4 Kbytes to 2 Gbytes. Each local access window is assigned to a specific target interface as specified in [Table 5-1](#).

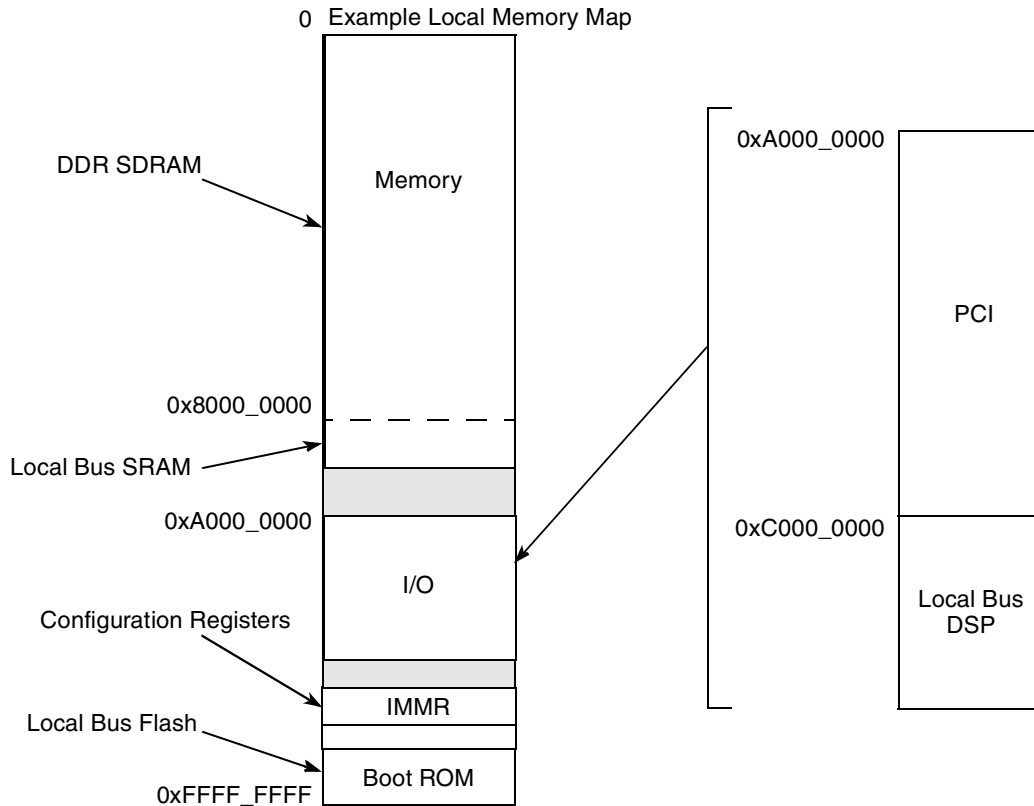
**Table 5-1. Local Access Windows Target Interface**

Window Number	Target Interface	Comments
0	Configuration registers (IMMR)	Fixed 1-Mbyte window size
1	Local bus	—
2	Local bus	—
3	Local bus	—
4	Local bus	—
5	PCI	—
6	PCI	—

**Table 5-1. Local Access Windows Target Interface (continued)**

Window Number	Target Interface	Comments
7	DDR SDRAM	—
8	DDR SDRAM	—

Figure 5-1 shows an example memory map.



**Figure 5-1. Local Memory Map Example**

Table 5-2 shows one example of local access window settings.

**Table 5-2. Local Access Windows Example**

Window	Base Address	Size	Target Interface
7	0x0000_0000	2 Gbytes	DDR SDRAM
2	0x8000_0000	1 Mbyte	Local bus
5	0xA000_0000	512 Mbytes	PCI
3	0xC000_0000	256 Mbytes	Local bus
0	0xFF40_0000	1 Mbyte	Configuration registers (IMMR)
1	0xFF80_0000	8 Mbytes	Local bus boot ROM Flash
4, 6, 8	Unused		

In this example, the local access window of the boot ROM is defined as window number 1, on a local bus device, in the highest 8 Mbytes of memory as set by the reset configuration word high during the reset sequence (see [Section 4.3.2.2.4, “Boot ROM Location”](#)) and [Section 5.2.4.3.1, “LAWBAR0\[BASE\\_ADDR\] Reset Value.”](#) The local access window, which describes the range of memory used for memory-mapped registers (IMMR), is a fixed 1-Mbyte space pointed to by the IMMRBAR register, using its default value (0xFF40\_0000). See [Section 5.2.4.1, “Internal Memory Map Registers Base Address Register \(IMMRBAR\).”](#)

## 5.2.1 Address Translation and Mapping

In addition to any address translation performed by the e300c3 core MMU, three distinct types of translation and mapping operations are performed on transactions at the integrated device level. These are as follows:

- Mapping a local address to a target interface
- Translating the local 32-bit address to an external address space
- Translating external addresses to the local 32-bit address space

The local access windows perform target mapping for transactions within the local address space. The local access windows do not perform any address translation.

Outbound windows perform the mapping from the local 32-bit address space to the address space of PCI, which may be much larger than the local space.

Inbound windows perform address translation from the external address spaces of PCI to the local address space.

The target mappings created by an inbound window must be consistent with those of the local access windows. That is, if an inbound window maps a transaction to a given local address, a valid local access window for that address must be set independently.

All of the configuration registers that define mapping of local access windows follow the same register format. [Table 5-3](#) summarizes the general format of these window definitions.

**Table 5-3. Format of Window Definitions**

Register	Function
Base address	High-order address bits defining location of the window in the initial address space
Window size/attributes	Window enable, window size <sup>1</sup>

<sup>1</sup> An exception is the IMMR window, which is always enabled and has a fixed 1-Mbyte size.

Windows must be a power-of-two size. To perform a mapping function, the address of the transaction is compared with the base address register of each window. The number of bits used in the comparison is dictated by each window’s size attribute. When an address hits within a window, the transaction is directed to the appropriate target.

## 5.2.2 Window into Configuration Space

The internal memory map registers' base address register (IMMRBAR) defines a window that is used to access all memory-mapped configuration, control, and status registers, referred to as internal memory map registers or IMMR. This window is always enabled with a fixed size of 1 Mbyte, and no other attributes are attached so there is no associated size/attribute register. This window always takes precedence over all local access windows. The IMMRBAR always come out of reset with a default base address value of 0xFF40\_0000, and this base address value can be modified by writing to this register. For more information, see [Section 5.2.4.1, “Internal Memory Map Registers Base Address Register \(IMMRBAR\).”](#)

### NOTE

Although it is legal to use the 3-Mbyte space consecutive to the 1 Mbyte of the IMMR (for example, if IMMRBAR is 0xFF40\_0000, the 3-Mbyte address space consecutive to it is 0xFF50\_0000–0xFF7F\_FFFF), it is not recommended. This space may be used in future derivatives of the device that require a larger internal memory space.

## 5.2.3 Local Access Windows

As demonstrated in the address map overview in [Section 5.2, “Local Memory Map Overview and Example,”](#) local access windows associate a range of the local 32-bit address space with a particular target interface. This allows the internal interconnections of the device to route a transaction from its source to the proper target. No address translation is performed. The base address defines the high order address bits that give the location of the window in the local address space. The window attributes enable the window and define its size, while the window number specifies the target interface.

With the exception of configuration space (mapped by IMMRBAR), all addresses used by the system must be mapped by a local access window. This includes addresses that are mapped by PCI inbound windows.

The local access window registers exist as part of the local access block in the system configuration registers. See [Section 5.3.2, “System Configuration Registers.”](#) A detailed description of the local access window registers is given in the following sections. Note that the minimum size of a window is 4 Kbytes, so the low order 12 bits of the base address cannot be specified.

### 5.2.3.1 Local Access Register Memory Map

[Table 5-4](#) shows the memory map for the local access registers.

**Table 5-4. Local Access Register Memory Map**

Local Memory Offset (Hex)	Register	Access	Reset	Section/Page
0x0_0000	Internal memory map base address register (IMMRBAR)	R/W	0xFF40_0000	<a href="#">5.2.4.1/5-6</a>
0x0_0004	Reserved	—	—	—
0x0_0008	Alternate configuration base address register (ALTCBAR)	R/W	0x0000_0000	<a href="#">5.2.4.2/5-7</a>
0x0_000C– 0x0_001C	Reserved	—	—	—
0x0_0020	eLBC local access window 0 base address register (LBLAWBAR0)	R/W	0x0000_0000 <sup>1</sup>	<a href="#">5.2.4.3/5-8</a>

Table 5-4. Local Access Register Memory Map (continued)

Local Memory Offset (Hex)	Register	Access	Reset	Section/Page
0x0_0024	eLBC local access window 0 attribute register (LBLAWAR0)	R/W	0x0000_0000 <sup>2</sup>	5.2.4.4/5-9
0x0_0028	eLBC local access window 1 base address register (LBLAWBAR1)	R/W	0x0000_0000	5.2.4.3/5-8
0x0_002C	eLBC local access window 1 attribute register (LBLAWAR1)	R/W	0x0000_0000	5.2.4.4/5-9
0x0_0030	eLBC local access window 2 base address register (LBLAWBAR2)	R/W	0x0000_0000	5.2.4.3/5-8
0x0_0034	eLBC local access window 2 attribute register (LBLAWAR2)	R/W	0x0000_0000	5.2.4.4/5-9
0x0_0038	eLBC local access window 3 base address register (LBLAWBAR3)	R/W	0x0000_0000	5.2.4.3/5-8
0x0_003C	eLBC local access window 3 attribute register (LBLAWAR3)	R/W	0x0000_0000	5.2.4.4/5-9
0x0_0040– 0x0_005C	Reserved	—	—	—
0x0_0060	PCI local access window 0 base address register (PCILAWBAR0)	R/W	0x0000_0000 <sup>3</sup>	5.2.4.5/5-10
0x0_0064	PCI local access window 0 attribute register (PCILAWAR0)	R/W	0x0000_0000 <sup>4</sup>	5.2.4.6/5-11
0x0_0068	PCI local access window 1 base address register (PCILAWBAR1)	R/W	0x0000_0000 <sup>5</sup>	5.2.4.5/5-10
0x0_006C	PCI local access window 1 attribute register (PCILAWAR1)	R/W	0x0000_0000	5.2.4.6/5-11
0x0_0070– 0x0_009C	Reserved	—	—	—
0x0_00A0	DDR local access window 0 base address register (DDRLAWBAR0)	R/W	0x0000_0000 <sup>6</sup>	5.2.4.7/5-12
0x0_00A4	DDR local access window 0 attribute register (DDRLAWAR0)	R/W	0x0000_0000 <sup>7</sup>	5.2.4.8/5-13
0x0_00A8	DDR local access window 1 base address register (DDRLAWBAR1)	R/W	0x0000_0000	5.2.4.7/5-12
0x0_00AC	DDR local access window 1 attribute register (DDRLAWAR1)	R/W	0x0000_0000	5.2.4.8/5-13
0x0_00B0– 0x0_00FC	Reserved	—	—	—

<sup>1</sup> Depends on reset configuration word high values. See [Section 5.2.4.3.1](#), “[LBLAWBAR0\[BASE\\_ADDR\] Reset Value](#),” for details.

<sup>2</sup> Depends on reset configuration word high values. See [Section 5.2.4.4.1](#), “[LBLAWAR0\[EN\] and LBLAWAR0\[SIZE\] Reset Value](#),” for details.

<sup>3</sup> Depends on reset configuration word high values. See [Section 5.2.4.5.1](#), “[PCILAWBAR0\[BASE\\_ADDR\] Reset Value](#),” for details.

<sup>4</sup> Depends on reset configuration word high values. See [Section 5.2.4.7.1](#), “[DDRLAWBAR0\[BASE\\_ADDR\] Reset Value](#),” for details.

<sup>5</sup> Depends on reset configuration word high values. See [Section 5.2.4.6.1](#), “[PCILAWAR0\[EN\] and PCILAWAR0\[SIZE\] Reset Value](#),” for details.

<sup>6</sup> Depends on reset configuration word high values. See [Section 5.2.4.7.1](#), “[DDRLAWBAR0\[BASE\\_ADDR\] Reset Value](#),” for details.

<sup>7</sup> Depends on reset configuration word high values. See [Section 5.2.4.8.1](#), “[DDRLAWAR0\[EN\] and DDRLAWAR0\[SIZE\] Reset Value](#),” for details.

## 5.2.4 Local Access Register Descriptions

### 5.2.4.1 Internal Memory Map Registers Base Address Register (IMMRBAR)


The IMMR window contains configuration, control, and status registers, as well as internal device memory arrays. The internal memory map occupies a 1-Mbyte region of memory space. Its location is programmable using the internal memory map register (IMMR). The default base address for the internal memory map register is 0xFF40\_0000. Because IMMRBAR is at offset 0x0 from the beginning of the local access registers, IMMRBAR always points to itself.

#### 5.2.4.1.1 Updating IMMRBAR

Updates to IMMRBAR that relocate the entire 1-Mbyte region of the internal memory block require special treatment. The effect of the update must be guaranteed to be visible by the mapping logic before an access to the new location is seen. To make sure this happens, the following guidelines should be followed:

- IMMRBAR should be updated during initial configuration of the device when only one host or controller has access to the device as follows:
  - If an external host on PCI is configuring the device, it should set IMMRBAR to the desired final location before the e300c3 core is released to boot.
  - If the core is initializing the device, it should set IMMRBAR to the desired final location before enabling other I/O devices to access the device.
- When the e300 core is writing to IMMRBAR, it should use the following sequence:
  - Read the current value of IMMRBAR using a load word instruction followed by an **isync**. This forces all accesses to configuration space to complete.
  - Write the new value to IMMRBAR.
  - Perform a load of an address that does not access configuration space or the on-chip SRAM, but has an address mapping already in effect (for example, boot ROM). Follow this load with an **isync**.
  - Read the contents of IMMRBAR from its new location, followed by another **isync**.

The IMMRBAR is shown in [Figure 5-2](#).



**Figure 5-2. Internal Memory Map Registers' Base Address Register (IMMRBAR)**

Table 5-5 defines the bit fields of IMMRBAR.

**Table 5-5. IMMRBAR Bit Settings**

Bits	Name	Description
0–11	BASE_ADDR	Identifies the 12 most-significant address bits of the base of the 1-Mbyte internal memory window.
12–31	—	Reserved. Software must write all zeros.


### 5.2.4.2 Alternate Configuration Base Address Register (ALTCBAR)

The alternate configuration base address register (ALTCBAR) is used to define the base address for an alternate 1-Mbyte region of configuration space to be used by the boot sequencer. By loading the proper boot sequencer command in the serial ROM, the base address in the ALTCBAR can be combined with the 20 bits of address offset supplied from the serial ROM to generate a 32-bit address. Thus, by configuring this register, the boot sequencer has access to the entire memory map, one 1-Mbyte block at a time. See Section 17.4.5, “Boot Sequencer Mode,” for more information.

#### NOTE

ALTCBAR is not considered a local access window on its own, so the boot sequencer must configure one of the other eight local access windows properly to reach the desired target peripherals.

The alternate configuration base address register is shown in Figure 5-3.



**Figure 5-3. Alternate Configuration Base Address Register (ALTCBAR)**


Table 5-6 defines the bit fields of ALTCBAR.

**Table 5-6. ALTCBAR Bit Settings**

Bits	Name	Description
0–11	BASE_ADDR	Identifies the 12 most-significant address bits of an alternate base address used for boot sequencer configuration accesses.
12–31	—	Reserved. Write has no effect, read returns 0.

### 5.2.4.3 LBC Local Access Window *n* Base Address Registers (LBLAWBAR0–LBLAWBAR3)

The LBC local access window *n* base address registers (LBLAWBAR0–LBLAWBAR3) are shown in Figure 5-4.



- The LBLAWBAR0[BASE\_ADDR] reset value depends on the reset configuration word high values. See Section 5.2.4.3.1, “LBLAWBAR0[BASE\_ADDR] Reset Value,” for a detailed description.

**Figure 5-4. LBC Local Access Window *n* Base Address Registers (LBLAWBAR0–LBLAWBAR3)**

Table 5-7 defines the bit fields of LBLAWBAR0–LBLAWBAR3.

**Table 5-7. LBLAWBAR0–LBLAWBAR3 Bit Settings**

Bits	Name	Description
0–19	BASE_ADDR	Identifies the 20 most-significant address bits of the base of local access window <i>n</i> . The specified base address should be aligned to the window size, as defined by LBLAWARn[SIZE].
20–31	—	Reserved. Write has no effect, read returns 0.

#### 5.2.4.3.1 LBLAWBAR0[BASE\_ADDR] Reset Value

The core may also use a local bus peripheral device to fetch its boot vector. For this purpose, the LBLAWBAR0[BASE\_ADDR] reset value is set according to the value set in the reset configuration word high BMS field.

Table 5-8 defines the reset value of LBLAWBAR0[BASE\_ADDR].

**Table 5-8. LBLAWBAR0[BASE\_ADDR] Reset Value**


RCWHR[BMS]	BASE_ADDR Reset Value
0	0x00000
1	0xFF800





### 5.2.4.5 PCI Local Access Window $n$ Base Address Register (PCILAWBAR0–PCILAWBAR1)

The PCI local access window  $n$  base address registers (PCILAWBAR0–PCILAWBAR1) are shown in Figure 5-6.



<sup>1</sup> The reset value of PCILAWBAR0[BASE\_ADDR] depends on the reset configuration word high values. See Section 5.2.4.5.1, “PCILAWBAR0[BASE\_ADDR] Reset Value,” for a detailed description.

**Figure 5-6. PCI Local Access Window  $n$  Base Address Registers (PCILAWBAR0–PCILAWBAR1)**

Table 5-11 defines the bit fields of PCILAWBAR0–PCILAWBAR1.

**Table 5-11. PCILAWBAR0–PCILAWBAR1 Bit Settings**

Bits	Name	Description
0–19	BASE_ADDR	Identifies the 20 most-significant address bits of the base of local access window $n$ . The specified base address should be aligned to the window size, as defined by PCILAWAR $n$ [SIZE].
20–31	—	Reserved. Write has no effect, read returns 0.

#### 5.2.4.5.1 PCILAWBAR0[BASE\_ADDR] Reset Value

The core may use a PCI peripheral device to fetch its boot vector. For this purpose, the PCILAWBAR0[BASE\_ADDR] reset value is set according to the value set in the reset configuration word high BMS field.


Table 5-12 defines the reset value of PCILAWBAR0[BASE\_ADDR].

**Table 5-12. PCILAWBAR0[BASE\_ADDR] Reset Value**

RCWHR[BMS]	PCILAWBAR0[BASE_ADDR] Reset Value
0	0x00000
1	0xFF800

## 5.2.4.6 PCI Local Access Window $n$ Attributes Registers (PCILAWAR0–PCILAWAR1)

The PCI local access window  $n$  attributes registers (PCILAWAR0–PCILAWAR1) are shown in [Figure 5-7](#).



<sup>1</sup> The reset value of PCILAWAR0[EN] depends on RCWH[RLEXT] and RCWH[ROMLOC]. See [Section 5.2.4.6.1, “PCILAWAR0\[EN\] and PCILAWAR0\[SIZE\] Reset Value,”](#) for a detailed description.

<sup>2</sup> The reset value of PCILAWAR0[SIZE] is always 0b010110, meaning an 8-Mbyte local access window. See [Section 5.2.4.6.1, “PCILAWAR0\[EN\] and PCILAWAR0\[SIZE\] Reset Value,”](#) for a detailed description

**Figure 5-7. PCI Local Access Window  $n$  Attributes Registers (PCILAWAR0–PCILAWAR1)**

[Table 5-13](#) defines the bit fields of PCILAWAR0–PCILAWAR1.

**Table 5-13. PCILAWAR0–PCILAWAR1 Bit Settings**

Bits	Name	Description
0	EN	0 The PCI local access window $n$ is disabled. 1 The PCI local access window $n$ is enabled and other PCILAWAR $n$ and PCILAWBAR $n$ fields combine to identify an address range for this window.
1–25	—	Reserved. Write has no effect, read returns 0.
26–31	SIZE	Identifies the size of the window from the starting address. Window size is $2^{(\text{SIZE}+1)}$ bytes. 000000–001010 Reserved. Window is undefined. 001011 4 Kbytes 001100 8 Kbytes 001101 16 Kbytes ..... $2^{(\text{SIZE}+1)}$ bytes 011110 2 Gbytes 011111–111111 Reserved. Window is undefined.

### 5.2.4.6.1 PCILAWAR0[EN] and PCILAWAR0[SIZE] Reset Value

The core may use a PCI peripheral device to fetch its boot vector. For this purpose an 8-Mbyte ( $2^{(22+1)}$ ) local access window is defined by the PCILAWBAR0[SIZE] reset value, and PCILAWAR0 is enabled according to the value set in the reset configuration word high ROMLOC and RLEXT fields.


[Table 5-14](#) defines the reset value of PCILAWAR0[EN].

**Table 5-14. PCILAWAR0[EN] Reset Value**

RCWHR[RLEXT]/RCWHR [ROMLOC]	PCILAWR0[EN] Reset Value	Description
000, 011–111	0	e300c3 core boot not performed from a PCI device.
001	1	e300c3 core boot performed from a PCI device. PCI 8-Mbyte ( $2^{(22+1)}$ ) local access window is enabled.

### 5.2.4.7 DDR Local Access Window $n$ Base Address Registers (DDRLAWBAR0–DDRLAWBAR1)

The DDR local access window  $n$  base address registers (DDRLAWBAR0–DDRLAWBAR1) are shown in Figure 5-8.



<sup>1</sup> The reset value of DDRLAWBAR0[BASE\_ADDR] depends on the reset configuration word high values. See Section 5.2.4.7.1, “DDRLAWBAR0[BASE\_ADDR] Reset Value,” for a detailed description

**Figure 5-8. DDR Local Access Window  $n$  Base Address Registers (DDRLAWBAR0–DDRLAWBAR1)**

Table 5-15 defines the bit fields of DDRLAWBAR0–DDRLAWBAR1.

**Table 5-15. DDRLAWBAR0–DDRLAWBAR1 Bit Settings**

Bits	Name	Description
0–19	BASE_ADDR	Identifies the 20 most-significant address bits of the base of local access window $n$ . The specified base address should be aligned to the window size, as defined by DDRLAWAR $n$ [SIZE].
20–31	—	Reserved. Write has no effect, read returns 0.

#### 5.2.4.7.1 DDRLAWBAR0[BASE\_ADDR] Reset Value

The core may use a DDR SDRAM device to fetch its boot vector. For this purpose, the DDRLAWBAR0[BASE\_ADDR] reset value is set according to the value set in the reset configuration word high BMS field.


Table 5-16 defines the reset value DDRLAWBAR0.

**Table 5-16. DDRLAWBAR0[BASE\_ADDR] Reset Value**

RCWHR[BMS]	DDRLAWBAR0[BASE_ADDR] Reset Value
0	0x00000
1	0xFF800

## 5.2.4.8 DDR Local Access Window $n$ Attributes Registers (DDRLAWAR0–DDRLAWAR1)

The DDR local access window  $n$  attributes registers (DDRLAWAR0–DDRLAWAR1) are shown in Figure 5-9.



<sup>1</sup> The reset value of DDRLAWAR0[EN] depends on the reset configuration word high values. See Section 5.2.4.8.1, “DDRLAWAR0[EN] and DDRLAWAR0[SIZE] Reset Value,” for a detailed description.

<sup>2</sup> The reset value of DDRLAWAR0[SIZE] is always 0b010110, meaning an 8-Mbyte local access window. See Section 5.2.4.8.1, “DDRLAWAR0[EN] and DDRLAWAR0[SIZE] Reset Value,” for a detailed description.

**Figure 5-9. DDR Local Access Window  $n$  Attributes Registers (DDRLAWAR0–DDRLAWAR1)**

Table 5-17 defines the bit fields of DDRLAWAR0–DDRLAWAR1.

**Table 5-17. DDRLAWAR0–DDRLAWAR1 Bit Settings**

Bits	Name	Description
0	EN	0 The DDR local access window $n$ is disabled. 1 The DDR local access window $n$ is enabled and other DDRLAWAR $n$ and DDRLAWBAR $n$ fields combine to identify an address range for this window.
1–25	—	Reserved. Write has no effect, read returns 0.
26–31	SIZE	Identifies the size of the window from the starting address. Window size is $2^{(SIZE+1)}$ bytes. 000000–001010 Reserved. Window is undefined. 001011 4 Kbytes 001100 8 Kbytes 001101 16 Kbytes ..... $2^{(SIZE+1)}$ bytes 011110 2 Gbytes 011111–111111 Reserved. Window is undefined.

### 5.2.4.8.1 DDRLAWAR0[EN] and DDRLAWAR0[SIZE] Reset Value

The core may use a DDR SDRAM device to fetch its boot vector. For this purpose an 8-Mbyte ( $2^{(22+1)}$ ) local access window is defined by DDRLAWBAR0[SIZE] reset value, and DDRLAWAR0 is enabled according to the value set in the reset configuration word high ROMLOC and RLEXT field.

Table 5-18 defines the reset value DDRLAWAR0[EN] and DDRLAWAR0[SIZE].

**Table 5-18. DDRLAWAR0[EN] Reset Value**

RCWHR[RLEXT]/ RCWHR[ROMLOC]	DDRLAWAR0[EN] Reset Value	Description
000	1	e300c3 core boot performed from a DDR SDRAM device. DDR 8-Mbyte ( $2^{(22+1)}$ ) local access window is enabled.
Else	0	e300c3 core boot not performed from a DDR SDRAM device.

## 5.2.5 Precedence of Local Access Windows

If two local access windows overlap, the lower numbered window takes precedence (see Table 5-1 for window numbers). For instance, if two windows are set up as shown in Table 5-19, local access window 1 governs the mapping of the 1-Mbyte region from 0x7FF0\_0000 to 0x7FFF\_FFF, even though the window described in local access window 7 also encompasses that memory region.

**Table 5-19. Overlapping Local Access Windows**

Window	Base Address	Size	Target Interface
1	0x7FF0_0000	1 Mbyte	Local bus
7	0x0000_0000	2 Gbytes	DDR SDRAM

## 5.2.6 Configuring Local Access Windows

After a local access window is enabled, it should not be modified while any device in the system may be using the window. Accordingly, a new window should not be used until the effect of the write to the window is visible to all blocks that use the window. This can be guaranteed by completing a read of the last local access window configuration register before enabling any other devices to use the window. For instance, if local bus local access windows 1–3 are being configured in order during the initialization process, the last write (to LBLAWAR3) should be followed by a read of LBLAWAR3 before any devices try to use any of these windows. If the configuration is being done by the local e300c3 core, the read of LBLAWAR3 should be followed by an **isync** instruction.

## 5.2.7 Distinguishing Local Access Windows from Other Mapping Functions

It is important to distinguish between the mapping function performed by the local access windows and the additional mapping functions that happen at the target interface. The local access windows define how a transaction is routed through the device internal interconnects from the transaction's source to its target. Once the transaction has arrived at its target interface, that interface controller may perform additional mapping. For instance, the DDR SDRAM controller has chip select registers that map a memory request to a particular external device. The local bus controller has base registers that perform a similar function. The PCI interface has outbound address translation units that map the local address into an external address space.

These other mapping functions are configured by programming the configuration, control, and status registers of the individual interfaces. Note that there is no need to have a one-to-one correspondence between local access windows and chip select regions or outbound windows. A single local access window can be further decoded to any number of chip selects or to any number or outbound windows at the target interface.

## 5.2.8 Outbound Address Translation and Mapping Windows

Outbound address translation and mapping refers to the translation of addresses from the local 32-bit address space to the external address space and attributes of a particular I/O interface. On this device, the PCI block has an outbound address translation unit.

The PCI controller has six outbound windows plus a default window. See [Section 4.5, “Memory Map/Register Definitions,”](#) for a detailed description of the PCI outbound windows.

## 5.2.9 Inbound Address Translation and Mapping Windows

Inbound address translation and mapping refers to the translation of an address from the external address space of an I/O interface (such as PCI address space) to the local address space understood by the internal interfaces of this processor. It also refers to the mapping of transactions to a particular target interface and the assignment of transaction attributes. The PCI controller has inbound address translation unit.

### 5.2.9.1 PCI Inbound Windows

The PCI controller has three general inbound windows plus a dedicated window for memory mapped configuration accesses (PIMMR). These windows have a one-to-one correspondence with the base address registers in the PCI programming model. Updating one automatically updates the other. There is no default inbound window; if a PCI address does not match one of the inbound windows, this processor does not respond with an assertion of PCI\_DEVSEL. See [Section 13.4.6, “PCI Inbound Address Translation,”](#) for a detailed description of the PCI inbound windows.

## 5.2.10 Internal Memory Map

All of the memory mapped configuration, control, and status registers in the device are contained within a 1-Mbyte address region, referred as the IMMR. To allow for flexibility, the internal memory map block can be relocated in the local address space. The local address map location of this register block is controlled by the internal memory map registers’ base address register (IMMRBAR); see [Section 5.2.4.1, “Internal Memory Map Registers Base Address Register \(IMMRBAR\).”](#) The default value for the IMMRBAR is 0xFF40\_0000.

### NOTE

The internal memory map window is always the highest priority local access window.

## 5.2.11 Accessing Internal Memory from External Masters

In addition to being accessible by the e300 processor, the IMMR memory window is accessible from external interfaces. This allows external masters on the I/O ports to configure the device.

External masters do not need to know the location of the IMMR memory in the local address map. Rather, they access this region of the local memory map through a window defined by a register in the interface's programming model that is accessible to the external master from its external memory map.

The PCI base address for accessing the local IMMR memory is selectable through the PCI internal memory map register (PIMMR), at offset 0x10, described in [Section 13.3.2.12, "PCI Inbound Base Address Registers \(PIBARn\)."](#) When the device is a PCI agent, an external PCI master sets this register by running a PCI configuration cycle. Subsequent memory accesses by a PCI master to the PCI address range indicated by PIMMR are translated to the local address indicated by the current setting of IMMRBAR.

## 5.3 System Configuration

The following sections describe some general information and configuration options that affect system behavior and performance.

### 5.3.1 System Configuration Register Memory Map

[Table 5-20](#) shows the memory map for the system configuration registers.

**Table 5-20. System Configuration Register Memory Map**

Local Memory Offset (Hex)	Register	Access	Reset	Section/Page
0x00100	System general purpose register low (SGPRL)	R/W	0x0000_0000	<a href="#">5.3.2.1/5-17</a>
0x00104	System general purpose register high (SGPRH)	R/W	0x0000_0000	<a href="#">5.3.2.2/5-17</a>
0x00108	System part and revision ID register (SPRIDR)	R	0x80B0_0021	<a href="#">5.3.2.3/5-18</a>
0x0010C	Reserved	—	—	—
0x00110	System priority configuration register (SPCR)	R/W	0x0000_0000	<a href="#">5.3.2.4/5-18</a>
0x00114	System I/O configuration register low (SICRL)	R/W	0x0000_0000 <sup>1</sup>	<a href="#">5.3.2.5/5-21</a>
0x00118	System I/O configuration register high (SICRH)	R/W	0x0000_0000	<a href="#">5.3.2.6/5-23</a>
0x0011C–0x00124	Reserved	—	—	—
0x00128	DDR control driver register (DDRCDR)	R/W	0x0004_0000	<a href="#">5.3.2.8/5-27</a>
0x0012C	DDR debug status register (DDRDSR)	R	0x3300_0000	<a href="#">5.3.2.9/5-29</a>
0x00130–0x0014C	Reserved	—	—	—
0x00150–0x001FC	Reserved	—	—	—


<sup>1</sup> Depends on the reset configuration word high configuration values.



## 5.3.2 System Configuration Registers

### 5.3.2.1 System General Purpose Register Low (SGPRL)

The system general purpose register low (SGPRL), shown in [Figure 5-10](#), can be used by software for any purpose. The values set in this register have no effect on the internal hardware.



**Figure 5-10. System General Purpose Register Low (SGPRL)**


[Table 5-21](#) defines the bit fields of SGPRL.

**Table 5-21. SGPRL Bit Settings**

Bits	Name	Description
0–31	GP	General purpose

### 5.3.2.2 System General Purpose Register High (SGPRH)

The system general purpose register high (SGPRH), shown in [Figure 5-11](#), can be used by software for any purpose. The values set in this register have no effect on the internal hardware.



**Figure 5-11. System General Purpose Register High (SGPRH)**


[Table 5-22](#) defines the bit fields of SGPRH.

**Table 5-22. SGPRH Bit Settings**

Bits	Name	Description
0–31	GP	General purpose

### 5.3.2.3 System Part and Revision ID Register (SPRIDR)

SPRIDR, shown in [Figure 5-12](#), provides information about the device and revision numbers.



**Figure 5-12. System Part and Revision ID Register (SPRIDR)**

[Table 5-23](#) defines the bit fields of SPRIDR.

**Table 5-23. SPRIDR Bit Settings**

Bits	Name	Description
0–15	PARTID	Part identification. This read-only field is mask-programmed with a code corresponding to the device number. It is intended to help factory test and user code that is sensitive to device changes. The device number changes according to manufacturing considerations. See <a href="#">Table 5-24</a> for values of this field.
16–31	REVID	Revision identification. This read-only field is mask-programmed with a code corresponding to the revision number of the part defined in PARTID field. It is intended to help factory test and user code that is sensitive to device changes. The mask number is programmed in a commonly changed layer, and changes with each mask set change. See <a href="#">Table 5-25</a> for values of this field.

#### 5.3.2.3.1 SPRIDR[PARTID] Coding

[Table 5-24](#) defines the reset values of SPRIDR[PARTID].

**Table 5-24. PARTID Coding**

PARTID	Device Name	Package Type
0x80B1	MPC8313	PBGA
0x80B0	MPC8313E	PBGA

[Table 5-25](#) defines the reset values of SPRIDR[REVID].

**Table 5-25. REVID Coding**

REVID	Device Revision
0x0021	2.1

### 5.3.2.4 System Priority and Configuration Register (SPCR)

The system priority and configuration register (SPCR), shown in [Figure 5-13](#), controls the priority of requests for transactions on the internal system bus. This priority is considered by the system arbiter

whenever an internal unit requests mastership of the coherent system bus (CSB). The SPCR also includes some other control functions.

Offset 0x00110		Access: Read/Write															
		0	2	3	4	5	6	7	8	9	10	11	12	15			
R	W	BUFGTX 125	BUFMDIO	—	PCIHPE	—	PCIPR	OPT	TBEN	COREPR	—						
Reset		$n^1$	$n^1$	0	0	0	0	0	0	0	0	0	0	0	0	0	0
		16	17	18	19	20	21	22	23	24	25	26	27	28	29	30	31
R	W	TSEC1588		TSECDP	TSECBDP	TSECEP	—										
Reset		All zeros															

**Figure 5-13. System Priority Configuration Register (SPCR)**

<sup>1</sup> Default values depend on CFG\_RESET\_SRC value. See [Table 5-27](#).

[Table 5-26](#) defines the bit fields of SPCR.

**Table 5-26. SPCR Bit Settings**

Bits	Name	Description
0	BUFGTX125	0 A signal is supplied from an external PHY or oscillator to TSEC1_GTX_CLK125 has the same voltage level as the LVddb power supply. 1 A 2.5 V signal is supplied to TSEC1_GTX_CLK125 from an external PHY or oscillator, and the LVddb power supply has a 3.3 V voltage level.
1	BUFMDIO	0 A 3.3 V signal is supplied from external phy to TSEC1_MDIO. 1 A 2.5 V signal is supplied from external phy to TSEC1_MDIO.
2	—	Reserved. Should be cleared.
3	PCIHPE	PCI highest priority enable. If this bit is set, the PCI bridge is permitted to request the coherent system bus (CSB) with highest priority, regardless of SPCR[PCIPR] value, when it needs to complete a posted write transaction from an external PCI master. To follow PCI ordering rules specifications, the PCI bridge must flush any outstanding write transactions before it can start a new read transaction. Setting this bit allows faster flushing of the outstanding write transactions coming from the PCI bus onto the CSB and to the device targets, such as DDR SDRAM and local bus memories.
4–5	—	Reserved. Should be cleared.
6–7	PCIPR	PCI bridge CSB request priority. The level of priority can be chosen from 4 possible levels. 00 Level 0 (lowest priority) 01 Level 1 10 Level 2 11 Level 3 (highest priority) <b>Note:</b> DMA has the same priority as PCI.
8	OPT	Optimize. Setting this bit may enhance the performance of transactions issued to the internal coherent system bus (CSB) by the security engine (SEC) and the USB controller. Performance is enhanced by reading more bytes on the bus than actually needed by the master in the case that this is more efficient. The user may set this bit only if it is known that USB transactions sent to the internal CSB are not accessing devices in which speculative reads may change the state of the device (for example, FIFOs in which reading a byte may advance some internal counter). 0 No performance enhancement. 1 Performance enhancement by speculative reading is enabled.

Table 5-26. SPCR Bit Settings (continued)

Bits	Name	Description
9	TBEN	e300c3 core time base unit enable 0 Time base unit is disabled. 1 Time base unit is enabled.
10–11	COREPR	e300c3 core CSB request priority. The priority level for the core in accessing the CSB can be chosen from 4 possible levels. 00 Level 0 (lowest priority) 01 Level 1 10 Level 2 11 Level 3 (highest priority)
12–15	—	Reserved. Should be cleared.
16–17	TSEC1588	00 Selects 1588 pins muxed with eTSEC1 (default) 01 Selects 1588 pins muxed with LA[7:15] pads. 10 Selects 1588 pins muxed with UART2 + I2C1 pads. 11 Reserved
18–19	TSECDP	eTSEC data priority. Selects the CSB request priority driven by eTSEC1 and eTSEC2 when they require to transfer data on this bus. The level of priority can be chosen from 4 possible levels. 00 Level 0 (lowest priority) 01 Level 1 10 Level 2 11 Level 3 (highest priority)
20–21	TSECBDP	eTSEC buffer descriptor priority. Selects the CSB request priority driven by eTSEC1 and eTSEC2 when they require to transfer a buffer descriptor (BD) on this bus. The level of priority can be chosen from 4 possible levels. 00 Level 0 (lowest priority) 01 Level 1 10 Level 2 11 Level 3 (highest priority)
22–23	TSECEP	eTSEC emergency priority. Selects the CSB request priority driven by eTSEC1 and eTSEC2 when an emergency condition occurs. The level of priority can be chosen from 4 possible levels. 00 Level 0 (lowest priority) 01 Level 1 10 Level 2 11 Level 3 (highest priority)
24-31	—	Reserved, should be cleared.

Table 5-27. CFG\_RESET\_SRC Values

CFG_RESET_SRC	SPCR[0]	SPCR[1]
0000	0	0
1000	0	0
1001	0	1
1010	0	0
1011	1	0
1100	1	0

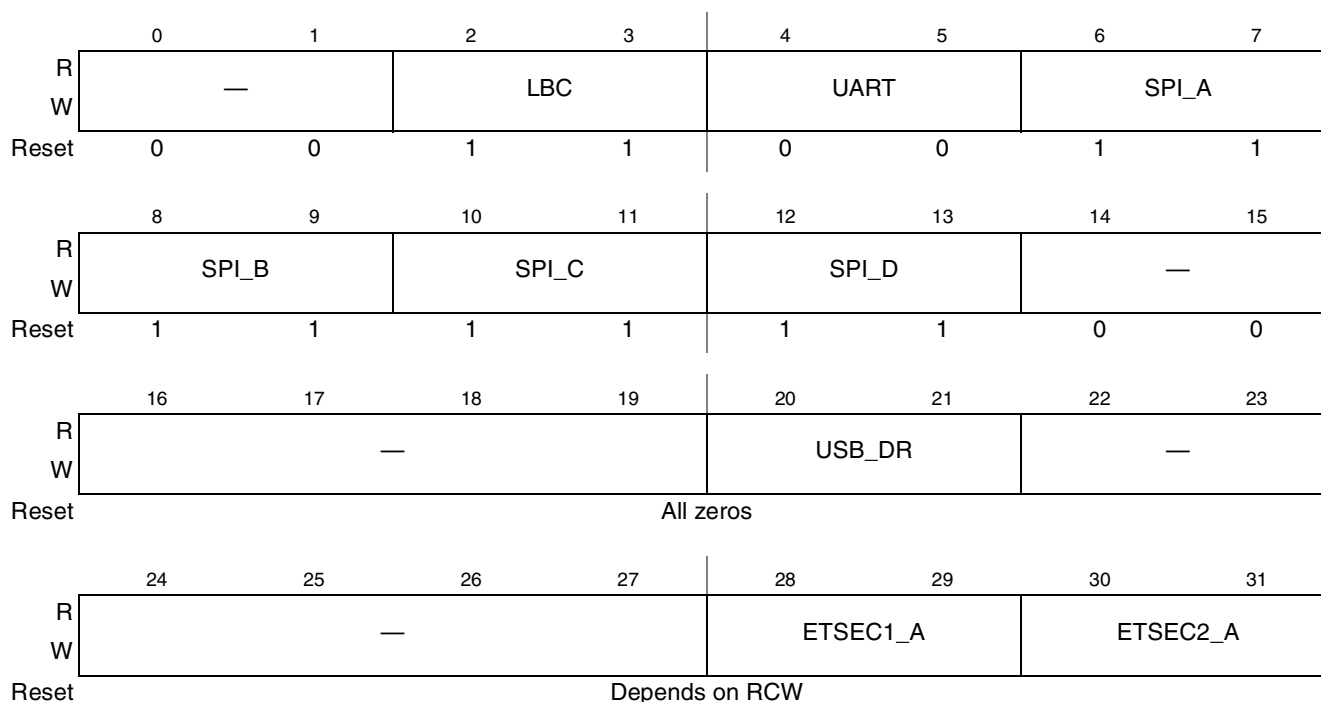
### 5.3.2.5 System I/O Configuration Register Low (SICRL)

The system I/O configuration register low (SICRL) controls the multiplexing of some of the device I/O pins. Each bit or set of bits in the SICRL selects which function is used by a certain group of the device pins.

Figure 5-14 shows SICRL.

Offset 0x00114

Access: Read/Write



**Figure 5-14. System I/O Configuration Register Low (SICRL)**

Table 5-28 defines the bit fields of SICRL. Each Pin Function column lists the name of the multi-function pin used in this option. Some groups have only two options (shown as Pin Function 0 and Pin Function 1) and therefore, only one control bit. In this case they can only have a value of 0b0 or 0b1. Other groups may have four options (shown as Pin Function 0, Pin Function 1, Pin Function 2, and Pin Function 3) and therefore, two control bits. In this case they can have a value of 0b00, 0b01, 0b10, or 0b11. Use the notations '0bN' or '0bNN' according to whether a group has one or two control bits, respectively.

**Table 5-28. SICRL Bit Settings**

SICRL[Bits] Value		0b0/0b00	0b1/0b01	0b10	0b11	Reset Value
Bits	Group	Pin Function 0	Pin Function 1	Pin Function 2	Pin Function 3	
0-1	Reserved	—	—	—	—	00

Table 5-28. SICRL Bit Settings (continued)

SICRL[Bits] Value		0b0/0b00	0b1/0b01	0b10	0b11	Reset Value
Bits	Group	Pin Function 0	Pin Function 1	Pin Function 2	Pin Function 3	
2–3	LBC	LA0	MSRCID0 (DDRID)	—	GPIO[0]	11
		LA1	MSRCID1 (DDRID)	—	GPIO[1]	
		LA2	MSRCID2 (DDRID)	—	GPIO[2]	
		LA3	MSRCID3 (DDRID)	—	GPIO[3]	
		LA4	MSRCID4 (DDRID)	—	GPIO[4]	
		LA5	MDVAL (DDRID)	—	GPIO[5]	
		LA6	—	—	GPIO[6]	
		LA7	—	TSEC_TMR_TRIG2	GPIO[7]	
		LA8	—	TSEC_TMR_ALARM1	GPIO[13]	
		LA9	—	TSEC_TMR_PP3	GPIO[14]	
<b>NOTE: CFG_LBIU_MUX_EN = 0 bypasses the register and selects function 0</b>						
4–5	UART	UART_SOUT1	MSRCID0 (DDRID)	—	—	00
		UART_SIN1	MSRCID1 (DDRID)	—	—	
		UART_CTS_B1	MSRCID2 (DDRID)	—	GPIO[8]	
		UART_RTS_B1	MSRCID3 (DDRID)	—	GPIO[9]	
		UART_SOUT2	MSRCID4 (DDRID)	—	TSEC_TMR_CLK	
		UART_SIN2	MDVAL (DDRID)	—	TSEC_TMR_GCLK	
		UART_CTS_B2	—	—	TSEC_TMR_PP1	
		UART_RTS_B2	—	—	TSEC_TMR_PP2	
6–7	SPI_A	SPIMOSI	GTM1_TIN3/GTM2_TIN4	LSRCID4	GPIO[28]	11
8–9	SPI_B	SPIMOSO	GTM1_TGATE3/GTM2_TGATE4	LDVAL	GPIO[29]	11
10–11	SPI_C	SPICLK	GTM1_TOUT3	—	GPIO[30]	11
12–13	SPI_D	SPISEL	—	—	GPIO[31]	11
14–19	Reserved	—	—	—	—	0
20–21	USBDR	GTM1_TIN1/GTM2_TIN2	LSRCID0	USBDR_DRIVE_VBUS	—	00
		GTM1_TGATE1/GTM2_TGATE2	LSRCID1	USBDR_PWRFAULT	—	
		GTM1_TOUT1	LSRCID2	USBDR_PCTL0	—	
		—	LSRCID3	USBDR_PCTL1	—	

Table 5-28. SICRL Bit Settings (continued)

SICRL[Bits] Value		0b0/0b00	0b1/0b01	0b10	0b11	Reset Value
Bits	Group	Pin Function 0	Pin Function 1	Pin Function 2	Pin Function 3	
22–27	Reserved	—	—	—	—	0
28–29	ETSEC1 _A	TSEC1_TX_EN	—	TSEC_1588_ALARM1	—	00
		TSEC1_TX_ER	—	TSEC_1588_ALARM2	—	
30–31	ETSEC2 _A	TSEC2_RXD3	—	—	GPIO[20]	00 RTBI 11 Else
		TSEC2_RXD2	—	—	GPIO[21]	
		TSEC2_RXD1	—	—	GPIO[22]	
		TSEC2_RXD0	—	—	GPIO[23]	
		TSEC2_TX_EN	—	—	GPIO[26]	
		TSEC2_TX_ER	—	—	GPIO[27]	

### 5.3.2.6 System I/O Configuration Register High (SICRH)

The system I/O configuration register high, shown in [Figure 5-15](#), controls the multiplexing of the rest of the device I/O pins. Each bit or set of bits in this register select which function is used by a certain group of the device pins. The reset value of this register depends on TSEC1M and TSEC2M fields setting in the reset configuration word high. This is used to avoid contention in systems not using TBI or RTBI modes. In these systems, the device pins are not driven during and after reset. The function of these pins can be changed by writing to this register during system initialization. The reset value of TSOBI1 and TSOBI2 also depends on the TSEC1M and TSEC2M fields settings in the reset configuration word high in order to select the correct output buffer impedance for full or reduced TSEC pin mode.

Offset 0x00118

Access: Read/Write




Figure 5-15. System I/O Configuration Register High (SICRH)

Table 5-29 defines the bit fields of SICRH. Each Pin Function column lists the name of the multi-function pin used in this option. Some groups have only two options (shown as Pin Function 0 and Pin Function 1) and therefore, only one control bit. In this case they can have the value of 0b0 or 0b1 only. Other groups may have three options (shown as Pin Function 0, Pin Function 1, and Pin Function 2) and therefore, two control bits. In this case they can have the value of 0b00, 0b01, or 0b10. A value of 0b11 is illegal for all groups. Use the 0bN or 0bNN notations according to a group having one or two control bits, respectively.

Note that bits 30 and 31 (TSOBI1 and TSOBI2), which control TSEC output buffer impedance, are described in Table 5-30

Table 5-29. SICRH Bit Settings

SICRH[Bits] Value:		0b0/0b00	0b1/0b01	0b10	0b11	Reset Value
Bits	Group	Pin Function 0	Pin Function 1	Pin Function 2	Pin Function 3	
0-5	Reserved	—	—	—	—	0
6	INTR_A	IRQ_B3	$\overline{\text{CKSTOP\_OUT}}$	—	—	0



Table 5-29. SICRH Bit Settings (continued)

SICRH[Bits] Value:		0b0/0b00	0b1/0b01	0b10	0b11	Reset Value
Bits	Group	Pin Function 0	Pin Function 1	Pin Function 2	Pin Function 3	
7	eLBC	LA7	—	TSEC_1588_TRIG2 <sup>1</sup>	GPIO7	1
		LA8	—	TSEC_1588_ALARM1 <sup>1</sup>	GPIO13	
		LA9	—	TSEC_1588_PP3 <sup>1</sup>	GPIO14	
		LA10	—	—	—	
		LA11	—	TSEC_1588_CLK <sup>1</sup>	—	
		LA12	—	TSEC_TMR_GCLK <sup>1</sup>	—	
		LA13	—	TSEC_TMR_PP1 <sup>1</sup>	—	
		LA14	—	TSEC_TMR_PP2 <sup>1</sup>	—	
		LA15	—	TSEC_TMR_ALARM2 <sup>1</sup>	—	
8–9	INTR_B	IRQ_B4	$\overline{\text{CKSTOP\_IN}}$	—	GPIO[12]	00
10–11	I <sup>2</sup> C	IIC1_SDA	$\overline{\text{CKSTOP\_OUT}}$	—	TSEC_TMR_TRIG1	00
		IIC1_SCL	$\overline{\text{CKSTOP\_IN}}$	—	TSEC_TMR_ALARM2	
		IIC2_SDA	PMC_PWR_OK	—	GPIO[10]	
		IIC2_SCL	—	—	GPIO[11]	
12–13	ETSEC2_B	TSEC2_COL	GTM1_TIN4/GTM2_TIN3	—	GPIO[15]	00 RTBI 11 Else
14–15	ETSEC2_C	TSEC2_CRS	$\overline{\text{GTM1\_GATE4/GTM2\_TGATE3}}$	—	GPIO[16]	00 RTBI 11 Else
16–17	ETSEC2_D	TSEC2_GTX_CLK	$\overline{\text{GTM1\_TOUT4}}$	$\overline{\text{GTM2\_TOUT3}}$	GPIO[17]	00 RTBI 11 Else
18–19	ETSEC2_E	TSEC2_RX_CLK	GTM1_TIN2/GTM2_TIN1	—	GPIO[18]	00 RTBI 11 Else
20–21	ETSEC2_F	TSEC2_RX_DV	$\overline{\text{GTM1\_TGATE2/GTM2\_TGATE1}}$	—	GPIO[19]	00 RTBI 11 Else
22–23	ETSEC2_G	TSEC2_RX_ER	$\overline{\text{GTM1\_TOUT2}}$	$\overline{\text{GTM2\_TOUT1}}$	GPIO[24]	00 RTBI 11 Else
		TSEC2_TX_CLK	—	—	GPIO[25]	—

Table 5-29. SICRH Bit Settings (continued)

SICRH[Bits] Value:		0b0/0b00	0b1/0b01	0b10	0b11	Reset Value
Bits	Group	Pin Function 0	Pin Function 1	Pin Function 2	Pin Function 3	
24	ETSEC1_B	TSEC1_COL	USBDR_TXDRXD0	—	—	0 RTBI 1 Else
		TSEC1_CRS	USBDR_TXDRXD1	—	—	
		TSEC1_GTX_CLK	USBDR_TXDRXD2	—	—	
		TSEC1_RX_CLK	USBDR_TXDRXD3	—	—	
		TSEC1_RX_DV	USBDR_TXDRXD4	—	—	
		TSEC1_RXD3	USBDR_TXDRXD5	—	—	
		TSEC1_RXD2	USBDR_TXDRXD6	—	—	
		TSEC1_RXD1	USBDR_TXDRXD7	—	—	
25-26	ETSEC1_C	TSEC1_RXD0	USBDR_NXT	TSEC_1588_TRIG1	—	00 RTBI 01 Else
		TSEC1_RX_ER	USBDR_DIR	TSEC_1588_TRIG2	—	
		TSEC1_TX_CLK	USBDR_CLK	TSEC_1588_CLK	—	
		TSEC1_TXD3	—	TSEC_1588_GCLK	—	
		TSEC1_TXD2	—	TSEC_1588_PP1	—	
		TSEC1_TXD1	—	TSEC_1588_PP2	—	
		TSEC1_TXD0	USBDR_STP	TSEC_1588_PP3	—	
27	Reserved	—	—	—	—	0
28-29	Reserved	Writes to SICRH[28-29] should preserve the reset value				0
30	TSOBI1	See <a href="#">Table 5-30</a> for description and reset value.				
31	TSOBI2	See <a href="#">Table 5-30</a> for description and reset value.				

<sup>1</sup> CFG\_LBIU\_MUX\_EN must be asserted during power-on reset to select timer functionality. CFG\_LBIU\_MUX\_EN=0 bypasses SICRH/SICRL register control and selects Pin Function 0 for these pads.

Table 5-30. SICRH[30–31] Bit Settings

Bits	Name	Description	Reset Value
30	TSOBI1	TSEC1 output buffer impedance. This bit controls the output buffer impedance of the TSEC1 output signals, used for reduced pin mode interfaces (RTBI/RGMII). The output buffer impedance should be correlated to the voltage supplied to the TSEC1 I/O pins (lvddb). For non-eTSEC mode of operation, this bit must be cleared. 0 Output buffer is set for 40 Ω, 3.3 V. 1 Output buffer is set for 40 Ω, 2.5 V.	0 Else 1 RGMII, RTBI
31	TSOBI2	TSEC2 output buffer impedance. This bit controls the output buffer impedance of TSEC2 output signals, used for reduced pin mode interfaces (RTBI/RGMII). The output buffer impedance should be correlated to the voltage supplied to the TSEC2 I/O pins (lvdda). 0 Output buffer is set for 40 Ω, 3.3 V. 1 Output buffer is set for 40 Ω, 2.5 V.	0 Else 1 RGMII, RTBI

### 5.3.2.7 Debug Configuration

Debug information may be driven on the device pins. This information can identify the internal source of a transaction that reached the DDR SDRAM or local bus interfaces. The device can be configured to drive the MSRCID[0:4] and MDVAL, or LSRCID[0:4] and LDVAL signals, respectively on other device pins. The coding of the source ID debug information is the same as the coding of the MSTR\_ID field in the AEATR register of the arbiter (See [Section 6.2.6, “Arbiter Event Attributes Register \(AEATR\)”](#)).

#### 5.3.2.7.1 DDR Debug Configuration

The DDR debug configuration enables a DDR memory controller to enter debug mode in which the DDR SDRAM source ID field and data valid strobe are driven onto one of two optional sets of pins:

- UART pins. UART operation is disabled, and any signals driven by UART devices must be electrically disconnected from the UART I/O pins. Set SICRL[4–5] to 0b01 to select this mode.
- LBC pins. LBC operation is disabled, and any signals driven by LBC must be electrically disconnected from the LBC I/O pins. Set SICRL[2–3] to 0b01 to select this mode.

#### 5.3.2.7.2 Local Bus Debug Configuration

The local bus debug configuration enables a LBC debug mode in which the SDRAM source ID field and data valid strobe for LBC memory accesses are driven onto USB and SPI pins. USB and SPI operation must be disabled, and any signals driven by USB and SPI devices must be electrically disconnected from the I/O pins in this case. Set SICRL[6–7] and SICRL[8–9] to 0b10 and SICRL[20–21] to 0b01 to select this mode.

### 5.3.2.8 DDR Control Driver Register (DDRCDR)

The DDR control driver register (DDRCDR) contains bits that allow control over the driver of the DDR SDRAM controller.

DDRCDR is shown in Figure 5-16.




Figure 5-16. DDR Control Driver Register (DDRCDR)

Table 5-31 shows the bit definition of the DDRCDR.

Table 5-31. DDRCDR Field Descriptions

Bits	Name	Description
0	—	Reserved
1	DSO_EN	0 DDR driver software override disable 1 DDR driver software override enable
2–5	DSO_PZ	DDR driver software p-impedance override 0000 Half strength—Highest Z 1000 Much higher Z than nominal 1100 Higher Z than nominal 1110 Nominal impedance setting 1111 Lower Z than nominal
6–9	DSO_NZ	DDR driver software n-impedance override 0000 Half strength—Highest Z 1000 Much higher Z than nominal 1100 Higher Z than nominal 1110 Nominal impedance setting 1111 Lower Z than nominal
10–11	—	Reserved. Should be cleared.
12	ODT	ODT termination value for I/Os 0 75 Ω 1 150 Ω
13	DDR_cfg	Selects voltage level for DDR pads 0 DDR2 (1.8V mode) nominal impedance—18 Ω 1 DDR1 (2.5V mode) nominal impedance—18 Ω <b>Note:</b> DDR_cfg must be set according to the logical type of the DDR memory devices, as it effects logic behavior of the DDR controller as well as the physical parameters of the DDR I/O pads.
14–29	—	Reserved
30	M_odr	Disable memory transaction reordering 0 Memory transaction reordering enabled 1 Memory transaction reordering disabled
31	Q_DRN	0 Drain queue before sleep disable 1 Drain queue before sleep enable

### 5.3.2.9 DDR Debug Status Register (DDRDSR)

Figure 5-17 contains the debug status bits from the DDR SDRAM controller.




Figure 5-17. DDR Debug Status Register (DDRDSR)

<sup>1</sup> Reset value of bits 0-9 depends on the actual state of the signals being monitored.

Table 5-32 shows the bit settings of the DDRDSR.

Table 5-32. DDRDSR Field Descriptions

Bits	Name	Description
0–1	—	Reserved
2–5	PZ	Current setting of PFET driver impedance 0000 Half strength—highest Z 1000 Higher Z than nominal 1100 Nominal impedance setting 1110 Lower Z than nominal 1111 Much lower Z than nominal
6–9	NZ	Current setting of NFET driver impedance 0000 Half strength—highest Z 1000 Higher Z than nominal 1100 Nominal impedance setting 1110 Lower Z than nominal 1111 Much lower Z than nominal
10–31	—	Reserved

## 5.4 Software Watchdog Timer (WDT)

The following sections describe the theory of operation of the software watchdog timer (WDT) in the device, including a definition of the external signals and the functions they serve. Additionally, the configuration, control, and status registers are also described. Note that individual chapters in this book describe specific initialization aspects for each individual block.

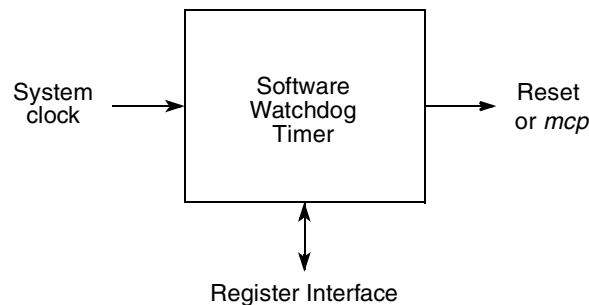
### 5.4.1 WDT Overview

The device provides a software watchdog timer (WDT) feature to prevent system lock in case the software becomes trapped in loops with no controlled exit. Watchdog timer operations are configured in the system watchdog control register (SWCRR).

The watchdog counter is a free-running down-counter that generates a reset or a non-maskable interrupt on underflow. To prevent a reset, software must periodically restart the countdown. The WDT is responsible for asserting a hardware reset or machine-check interrupt (*mcp*) if the software fails to service

the software watchdog timer for a certain period of time (for example, because software is lost or trapped in a loop with no controlled exit).

Figure 5-18 shows a high-level block diagram of the WDT.



**Figure 5-18. Software Watchdog Timer High-Level Block Diagram**

The software watchdog timer is enabled after reset to cause a hardware reset if it times out. The user has the option of disabling the software watchdog if it is not needed. If used, the software watchdog timer requires a special service sequence to be executed periodically. Without this periodic servicing, the software watchdog timer times out and issues a reset or a non-maskable interrupt.

## 5.4.2 WDT Features

The WDT includes the following key features:

- Based on 16-bit prescaler and 16-bit down-counter
- Provides a selectable range for the time-out period
- Provides ~12.8-sec maximum software time-out delay for 333-MHz input clock
- Functional and programming compatibility with MPC8260 watchdog timer

## 5.4.3 WDT Modes of Operation

The WDT unit can operate in the following modes:

- WDT enable/disable mode:  
If the software watchdog timer is not needed, the user can disable it with software after a system reset. When the watchdog timer is disabled, the watchdog counter and prescaler counter are held in a stopped state.
- WDT output reset/interrupt mode:  
Without software periodic servicing, the software watchdog timer times out and issues a reset or a nonmaskable interrupt (*mcp*)
- WDT prescaled/non-prescaled clock mode:  
The WDT counter clock can be prescaled by programming the SWCRR[SWPR] bit, which controls the divide-by-65,536 of the WDT counter.

## 5.4.4 WDT Memory Map/Register Definition

The WDT programmable register map occupies 16 bytes of memory-mapped space. Reading undefined portions of the memory map returns all zeros; writing has no effect.

All WDT registers are 16- or 32-bits wide, located on 16-bit address boundaries, and should be accessed as 16- or 32-bit quantities. All addresses used in this chapter are offsets from the WDT base, as defined in Chapter 2, “Memory Map.”

Table 5-33 shows the WDT memory map.


**Table 5-33. WDT Register Address Map**

Offset	Register	Access	Reset Value	Section/ Page
0x0–0x3	Reserved	—	—	—
0x4	System watchdog control register (SWCRR)	R/W	0xFFFF_0003 or 0xFFFF_0007 <sup>1</sup>	5.4.4.1/5-31
0x8	System watchdog count register (SWCNR)	R	0x0000_FFFF	5.4.4.2/5-32
0xC–0xD	Reserved	—	—	—
0xE	System watchdog service register (SWSRR)	R/W	0x0000	5.4.4.3/5-33

<sup>1</sup> SWCRR[SWEN] reset value directly depends on RCWHR[SWEN] (reset configuration word high).

### 5.4.4.1 System Watchdog Control Register (SWCRR)

The system watchdog control register (SWCRR), shown in Figure 5-19, controls the software watchdog period and configures watchdog timer operation. SWCRR can be read at any time but can be written only once after system reset.



<sup>1</sup> SWCRR[SWEN] reset value directly depends on RCWHR[SWEN] (reset configuration word high).

**Figure 5-19. System Watchdog Control Register (SWCRR)**


Table 5-34 defines the bit fields of SWCRR.

**Table 5-34. SWCRR Bit Settings**

Bits	Name	Description
0–15	SWTC	Software watchdog time count The SWTC field contains the modulus that is reloaded into the watchdog counter by a service sequence. When a new value is loaded into SWCRR[SWTC], the software watchdog timer is not updated until the servicing sequence is written to the SWSRR. If SWCRR[SWEN] is loaded with 0, the modulus counter does not count. The new value is also used at the next and all subsequent reloads. Reading the SWCRR register returns the value in the system watchdog control register. Reset initializes the SWCRR[SWTC] field to 0xFFFF. <b>Note:</b> The prescaler counter is reset any time a new value is loaded into the watchdog counter and also during reset.
16–28	—	Write reserved, read = 0
29	SWEN	Watchdog enable bit Enables the watchdog timer. The reset value directly depends on the value of the RCWHR[SWEN] bit. It should be cleared by software after a system reset to disable the software watchdog timer. When the watchdog timer is disabled, the watchdog counter and prescaler counter are held in a stopped state. 0 Watchdog timer disabled 1 Watchdog timer enabled <b>Note:</b> After software writes the SWRI bit, the state of SWEN cannot be changed.
30	SWRI	Software watchdog reset/interrupt select bit A WDT timer out causes either a hard reset or machine check interrupt to the core. 0 Software watchdog timer causes a machine check interrupt to the core 1 Software watchdog timer causes a hard reset
31	SWPR	Software watchdog counter prescale bit Controls the divide-by-65,536 WDT counter prescaler 0 The WDT counter is not prescaled. 1 The WDT counter clock is prescaled.

### 5.4.4.2 System Watchdog Count Register (SWCNR)

The system watchdog count register (SWCNR), shown in Figure 5-20, provides visibility to the watchdog counter value. SWCNR is a read-only register. Writes to SWCNR have no effect and terminate without transfer error exception.



**Figure 5-20. System Watchdog Count Register (SWCNR)**



Table 5-35 defines the bit fields of SWCNR.


**Table 5-35. SWCNR Bit Settings**

Bits	Name	Description
0–15	—	Write reserved, read = 0
16–31	SWCN	Software watchdog count field. The read-only SWCNR[SWCN] field reflects the current value in the watchdog counter. Writing to the SWCNR register has no effect, and write cycles are terminated normally. Reset initializes the SWCNR[SWCN] field to 0xFFFF. <b>Note:</b> Reading the 16 least-significant bits of 32-bit SWCNR register with two 8-bit reads is not guaranteed to return a coherent value.

#### 5.4.4.3 System Watchdog Service Register (SWSRR)

The system watchdog service register (SWSRR) is shown in Figure 5-21. When the watchdog timer is enabled, a write of 0x556C followed by a write 0xAA39 to the SWSRR register before the watchdog counter times out prevents a device reset. If the SWSRR register is not serviced before the timeout, a signal from the watchdog timer to the reset or interrupt controller module asserts a system reset or interrupt (depending on the setting of SWCRR[SWRI]).

Both writes must occur before the timeout in the order listed, but any number of instructions can be executed between the two writes. However, writing any value other than 0x556C or 0xAA39 to the SWSRR register resets the servicing sequence, requiring both values to be written to keep the watchdog timer from causing a reset. Reset initializes the SWSRR[WS] field to 0x0000. SWSRR can be written at any time, but returns all zeros when read.



**Figure 5-21. System Watchdog Service Register (SWSRR)**

Table 5-36 defines the bit fields of SWCNR.

**Table 5-36. SWSRR Bit Settings**

Bits	Name	Description
0–15	WS	Software watchdog service field. The user should periodically write 0x556C followed by 0xAA39 to this register to prevent a software watchdog timer timeout. SWSRR[WS] can be written at any time, but returns all zeros when read.

## 5.4.5 Functional Description

### 5.4.5.1 Software Watchdog Timer Unit

The device provides a software watchdog timer (WDT) feature to prevent system lock in case the software becomes trapped in loops with no controlled exit. Watchdog timer operations are configured in the system watchdog control register (SWCRR).

The software watchdog timer is enabled after reset to cause a soft reset or non-maskable interrupt (MCP) if it times out. If the software watchdog timer is not needed, the user must clear SWCRR[SWEN] to disable it. If used, the software watchdog timer requires a special service sequence to be executed periodically. Without this periodic servicing, the software watchdog timer times out and issues a reset or a nonmaskable interrupt, as programmed in SWCRR[SWRI]. Once software writes SWRI, the state of SWEN cannot be changed.

The software watchdog timer service sequence consists of the following two steps:

- Write 0x556C to the system watchdog service register (SWSRR)
- Write 0xAA39 to SWSRR

The service sequence reloads the watchdog timer and the timing process begins again. If a value other than 0x556C or 0xAA39 is written to the SWSRR, the entire sequence must start over. Although the writes must occur in the correct order before a time-out, any number of instructions can be executed between the writes. This allows interrupts and exceptions to occur between the two writes when necessary. Figure 5-22 shows a state diagram for the watchdog timer.

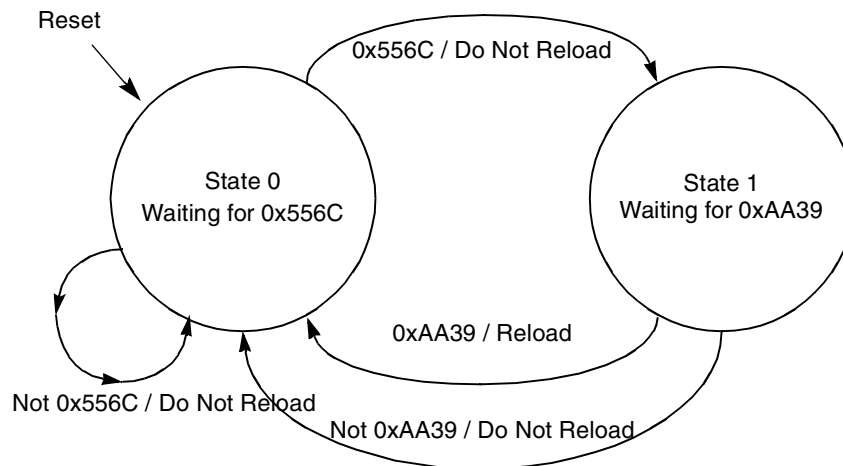



Figure 5-22. Software Watchdog Timer Service State Diagram

Although most software disciplines permit or even encourage the watchdog concept, some systems require a selection of time-out periods. For this reason, the software watchdog timer must provide a selectable range for the time-out period. Figure 5-23 shows how to handle this need.



**Figure 5-23. Software Watchdog Timer Functional Block Diagram**

Figure 5-23 shows that the range is determined by SWCRR[SWTC]. The value in SWTC is then loaded into a 16-bit decremter clocked by the system clock. An additional divide-by-65,536 prescaler value is used when needed.

The decremter begins counting when loaded with a value from SWTC. After the timer reaches 0x0, a software watchdog expiration request is issued to the reset or *mcp* (machine check) control logic. Upon reset, SWTC is set to the maximum value and is again loaded into the system watchdog service register (SWSRR), starting the process over. When a new value is loaded into SWTC, the software watchdog timer is not updated until the servicing sequence is written to the SWSRR. If SWCRR[SWEN] is loaded with 0, the modulus counter does not count.

### 5.4.5.2 Modes of Operation

The WDT unit can operate in the following modes:

- WDT enable/disable mode:

If the software watchdog timer is not needed, the user can disable it. The SWCRR[SWEN] bit enables the watchdog timer. It should be cleared by software after a system reset to disable the software watchdog timer. When the watchdog timer is disabled, the watchdog counter and prescaler counter are held in a stopped state.

— WDT enable mode (SWCRR[SWEN] = 1)

This is the default value after soft reset.

— WDT disable mode (SWCRR[SWEN] = 0)

- WDT reset/interrupt output mode

Without software periodic servicing, the software watchdog timer times out and issues a reset or a nonmaskable interrupt (*mcp*), programmed in SWCRR[SWRI].

According to the value of SWCRR[SWRI], the WDT timer causes a hard reset or machine check interrupt to the core.

- Reset mode (SWCRR[SWRI] = 1).  
Software watchdog timer causes a hard reset (this is the default value after hard reset).
- Interrupt mode (SWCRR[SWRI] = 0).  
Software watchdog timer causes a machine check interrupt to the core.
- WDT prescaled/non-prescaled clock mode  
The WDT counter clock can be prescaled by programming the SWCRR[SWPR] bit that controls the divide-by-65,536 of the WDT counter.
  - Prescale mode (SWCRR[SWPR] = 1)  
The WDT clock is prescaled.
  - Non-prescale mode (SWCRR[SWPR] = 0)  
The WDT clock is not prescaled.

## 5.4.6 Initialization/Application Information

### 5.4.6.1 WDT Programming Guidelines

The software watchdog timer is enabled (by the default value of SWCRR[SWEN]) after reset. The following initialization sequence of WDT is required:

- WDT disabling  
If the software watchdog timer is not needed, the user must clear SWCRR[SWEN] bit to disable the WDT not later than its timer times out (~12.8 sec. for a 333-MHz system clock).
- WDT initial servicing  
If the software watchdog timer is to be used, the special service sequence, described in [Section 5.4.5.1, “Software Watchdog Timer Unit,”](#) must be executed after system reset and not later than the first WDT time-out (~12.8 sec. for a 333-MHz system clock).  
Subsequently, periodical WDT servicing should be performed according to the programming guidelines given in [Section 5.4.5.1, “Software Watchdog Timer Unit.”](#)

## 5.5 Real Time Clock Module (RTC)

The following sections describe the theory of operation of the real time clock module (RTC) including a definition of the external signals and the functions it serves. Additionally, the configuration, control, and status registers are described. Note that individual chapters in this reference manual describe additional specific initialization aspects for each individual block.

### 5.5.1 RTC Overview

The device platform provides a real time clock (RTC) timer suitable for timestamping or time and calendar generation. It can maintain a one-second count which is unique over a period of approximately 136 years.

The RTC can be initialized by software with an initial count value using the real time counter load register (RTLDR). It can also be programmed to generate an interrupt every second. The real time counter control register (RTCTR) is used to enable or disable the various timer functions. The real time counter event

register (RTEVR) is used to report the interrupt source. The RTC counter is initialized by software and can be disabled if needed.

Figure 5-24 shows the high level RTC block diagram.




Figure 5-24. Real Time Clock Module High Level Block Diagram

## 5.5.2 RTC Features

The key features of the RTC include the following:

- Maintains a one-second count, unique over a period of thousand of years
- 32-bit RTC counter can be initialized by software to specific initial count value
- Provides an alarm function with programmable and maskable alarm interrupt
- Provides programmable and maskable every second interrupt
- Uses two possible clock sources: the CSB bus clock or an external RTC clock
- RTC function can be disabled if needed

## 5.5.3 RTC Modes of Operation

The RTC unit can operate in the following modes:

- RTC enable/disable mode
- RTC every-second interrupt enable/disable mode
- RTC alarm interrupt enable/disable mode
- RTC internal/external input clock mode

## 5.5.4 RTC External Signal Description

This section provides an overview and detailed descriptions of the RTC signals.

There is one distinct external RTC clock input signal, defined in [Table 5-37](#).

Table 5-37. RTC Signal Properties

Name	Port	Function	I/O	Reset	Pull Up
RTC_CLK	RTC_CLK	Real time clock input.	I	N/A	—

[Table 5-38](#) provides a detailed description of the external RTC signal.

**Table 5-38. RTC External Signal—Detailed Signal Description**

Signal	I/O	Description	
RTC_CLK	I	This signal is used as the timebase for the real time clock module.	
		State Meaning	—
		Timing	—

## 5.5.5 RTC Memory Map/Register Definition

The RTC programmable register map occupies 32 bytes of memory-mapped space. Reading undefined portions of the memory map returns all zeros; writing has no effect.

All RTC registers are 32 bits wide that are located on 32-bit address boundaries and should only be accessed as a 32-bit quantities.

All addresses used in this section are offsets from the RTC base, as defined in [Chapter 2, “Memory Map.”](#)

[Table 5-33](#) shows the memory map of the RTC.

**Table 5-39. RTC Register Address Map**

Offset	Register	Access	Reset Value	Section/ Page
0x00	Real time counter control register (RTCNR)	R/W	0x0000_0000	<a href="#">5.5.5.1/5-38</a>
0x04	Real time counter load register (RTLDR)	R/W	0x0000_0000	<a href="#">5.5.5.2/5-39</a>
0x08	Real time counter prescale register (RTPSR)	R/W	0x0000_0000	<a href="#">5.5.5.3/5-40</a>
0x0C	Real time counter register (RTCTR)	R	0x0000_0000	<a href="#">5.5.5.4/5-40</a>
0x10	Real time counter event register (RTEVR)	w1c	0x0000_0000	<a href="#">5.5.5.5/5-41</a>
0x14	Real time counter alarm register (RTALR)	R/W	0xFFFF_FFFF	<a href="#">5.5.5.6/5-41</a>
0x18–0x1F	Reserved	—	—	

### 5.5.5.1 Real Time Counter Control Register (RTCNR)

The real time counter control register (RTCNR), shown in [Figure 5-25](#), is used to enable RTC functions. The register can be read at any time.


**Figure 5-25. Real Time Counter Control Register (RTCNR)**


Table 5-40 defines the bit fields of RTCNR.

**Table 5-40. RTCNR Bit Settings**

Bits	Name	Description
0–23	—	Write reserved, read = 0
24	CLEN	Clock enable control bit. This bit controls the counting of the RTC. When the RTC's clock is disabled, the counter maintains its old value. When the counter's clock is enabled, it continues counting using the previous value. 0 Disable counter. 1 Enable counter.
25	CLIN	Input clock control bit. The input clock to the RTC may be either the CSB clock or an external RTC clock. 0 The input clock to the periodic interrupt timer is CSB input clock. 1 The input clock to the periodic interrupt timer is the external RTC clock.
26–29	—	Write reserved, read = 0
30	AIM	Alarm interrupt mask bit. Used to enable or disable (mask) the RTC alarm interrupt when the RTC's 32-bit counter reaches RTALR[ALR] value. 0 Alarm interrupt generation disabled. 1 Alarm interrupt generation enabled.
31	SIM	Second interrupt mask bit. Used to enable or disable (mask) the RTC periodic interrupt. 0 Periodic interrupt generation disabled. 1 Periodic interrupt generation enabled.

### 5.5.5.2 Real Time Counter Load Register (RTLDR)

The real time counter load register (RTLDR), shown in Figure 5-26, contains the 32-bit value to be loaded in the 32-bit RTC counter.



**Figure 5-26. Real Time Counter Load Register (RTLDR)**


Table 5-41 defines the bit fields of RTLDR.

**Table 5-41. RTLDR Bit Settings**

Bits	Name	Description
0–31	CLDV	Contains the 32-bit value to be loaded in the 32-bit RTC counter.

### 5.5.5.3 Real Time Counter Prescale Register (RTPSR)

The real time counter prescale register (RTPSR), shown in [Figure 5-27](#), is a read/write register used to configure the RTC prescaler's value.



**Figure 5-27. Real Time Counter Prescale Register (RTPSR)**

[Table 5-42](#) defines the bit fields of RTPSR.


**Table 5-42. RTPSR Bit Settings**

Bits	Name	Description
0–31	PRSC	RTC prescaler bits. Select the input clock divider for the RTC counter clock. The prescaler is programmed to divide the RTC clock input by values from 1 to 4,294,967,296. The value 0x0000 divides the clock by 1 and 0xFFFF_FFFF divides the clock by 4,294,967,296. To accurately predict the timing of the next count, change the RTPSR[PRSC] field only when the enable bit RTCNR[CLE] is clear. Changing the RTPSR[PRSC] bits resets the prescaler counter. System reset and the loading of a new value into the counter both reset the prescaler counter. Clearing RTCNR[CLE] stops the prescaler counter.

### 5.5.5.4 Real Time Counter Register (RTCTR)

The real time counter register (RTCTR), shown in [Figure 5-28](#), is a read-only register that shows the current value in the RTC counter.

The CNTV value is not affected by reads or writes to RTCTR.



**Figure 5-28. Real Time Counter Register (RTCTR)**

[Table 5-43](#) defines the bit fields of RTCTR.

**Table 5-43. RTCTR Bit Settings**

Bits	Name	Description
0–31	CNTV	RTC counter value field. RTCTR[CNTV] contains the current value of the time counter. This is a read-only field. Writes have no effect on RTCTR[CNTV].



### 5.5.5.5 Real Time Counter Event Register (RTEVR)

The real time counter event register (RTEVR), shown in Figure 5-29, is used to report the source of the interrupts. The register can be read at any time.



**Figure 5-29. Real Time Counter Event Register (RTEVR)**

RTEVR bits are cleared by writing ones. Writing zeros does not affect the value of the status bits.


Table 5-44 defines the bit fields of RTEVR.

**Table 5-44. RTEVR Bit Settings**

Bits	Name	Description
0–29	—	Write reserved, read = 0
30	AIF	Alarm interrupt flag bit. Used to indicate the alarm interrupt. It is set if the RTC issues an interrupt after the RTC counter counts to zero.
31	SIF	Second interrupt flag bit. Used to indicate the every-second interrupt. This status bit is set each time that the prescaler count reaches zero and should be cleared by software.

### 5.5.5.6 Real Time Counter Alarm Register (RTALR)

The real time counter alarm register (RTALR), shown in Figure 5-30, contains the 32-bit alarm (ALRM) value. When the value of the RTC counter equals the RTALR[ALRM] value, a maskable interrupt is generated.



**Figure 5-30. Real Time Counter Alarm Register (RTALR)**

Table 5-45 defines the bit fields of RTALR.

**Table 5-45. RTALR Bit Settings**

Bits	Name	Description
0–31	ALRM	RTC alarm value. The alarm interrupt is generated when the value of the RTC counter equals RTALR[ALRM].

## 5.5.6 Functional Description

### 5.5.6.1 Real Time Counter Unit

The real time clock (RTC) timer is suitable for time stamping or time and calendar generation. It can maintain a one-second count which is unique over a period of approximately 136 years. Software can convert this count into time-of-day or calendar information if required. An alarm function is also provided. The RTC can be clocked by the internal system bus clock or by an external clock source. The RTC consists of 32-bit up-counter which is incremented by an one-second count clock derived from the RTC input clock. The RTC can be programmed to generate a maskable interrupt when the time value matches the value in its associated alarm register.

The RTC can be initialized by software with an initial count value in the real time counter load register (RTLDR). It can also be programmed to generate an interrupt every second. The real time counter control register (RTCTR) is used to enable or disable the various timer functions. The real time counter event register (RTEVR) is used to report the interrupt source. The RTC counter is reset to zero on hard reset but is not affected by soft reset. It is initialized by the software. The RTC function can be disabled.

Figure 5-31 shows the functional RTC block diagram.




Figure 5-31. Real Time Clock Module Functional Block Diagram

### 5.5.6.2 RTC Operational Modes

The RTC unit can operate in the following modes:

- RTC enable/disable mode:

RTCNR[CLEN] enables the RTC timer. It should be set by software after a system reset to enable the RTC timer.

— RTC disable mode (RTCNR[CLEN] = 0)

When the RTC's clock is disabled, counter maintains its old value (default).

— RTC enable mode (RTCNR[CLEN] = 1)

When the counter's clock is enabled, it continues counting using the previous value.

- RTC every-second interrupt enable/disable mode:
  - RTC every-second interrupt enable mode (RTCNR[SIM] = 1)  
In this mode the RTC sets the RTEVR[SIF] flag and generate an interrupt after the RTC's 32-bit counter reaches zero.
  - RTC every-second interrupt disable mode (RTCNR[SIM] = 0)  
In this mode the RTC sets the RTEVR[SIF] flag but does not generate an interrupt after the RTC's 32-bit counter reaches zero.
- RTC alarm interrupt enable/disable mode:
  - RTC alarm interrupt enable mode (RTCNR[AIM] = 1)  
In this mode, the RTC sets the RTEVR[AIF] flag and generates an interrupt each time when the RTC's 32-bit counter reaches the RTALR[ALR] value.
  - RTC alarm interrupt disable mode (RTCNR[AIM] = 0)  
In this mode the RTC sets the RTEVR[AIF] flag but does not generate an interrupt when the RTC's 32-bit counter reaches the RTALR[ALR] value.
- RTC internal/external input clock mode:  
The input clock to the RTC may be the CSB clock or an external 32.768-kHz crystal.
  - RTC uses the internal input clock mode (RTCNR[CLIN] = 0)
  - RTC uses the external 32.768-kHz crystal clock (RTCNR[CLIN] = 1)

### 5.5.7 RTC Programming Guidelines

The following initialization sequence for the RTC is recommended:

1. Write to RTPSR to set the RTC prescaler to the desired value
2. Write to RTLDR to initialize the RTC initial value
3. Write to RTALR to program the RTC alarm value, if needed
4. Write to RTCNR to configure and start the RTC operation: RTC input clock source, second/alarm interrupt mask, RTC clock enable.

## 5.6 Periodic Interval Timer (PIT)

The following sections describe theory of operation of the periodic interval timer (PIT) including a definition of the external signals and the functions it serves. Additionally, the configuration, control, and status registers are described. Note that individual chapters in this reference manual describe additional specific initialization aspects for each individual block.

### 5.6.1 PIT Overview

The periodic interval timer (PIT) that generates periodic interrupts for a real-time operating system or an application software.

The PIT consists of a 32-bit down-counter which is decremented by a clock derived from a CSB clock or from an external 32.768-kHz crystal. The 32-bit counter decrements to zero when loaded with a initial value from the periodic interval timer load register (PTLDR). The periodic interval timer control register (PTCTR) is used to enable or disable the various timer functions. The periodic interval timer event register (PTEVR) is used to report the interrupt source. The PIT function can be disabled if needed.

Figure 5-32 shows the functional PIT block diagram.




Figure 5-32. Periodic Interval Timer High Level Block Diagram

## 5.6.2 PIT Features

The key features of the PIT include the following:

- Maintains a 32-bit down-counter, clocked by a 32-bit prescaled input clock
- 32-bit PIT counter can be initialized by software to specific initial count value
- Provides programmable and maskable periodic interrupt
- Uses two possible clock sources: the CSB clock or an external PIT clock
- PIT function can be disabled

## 5.6.3 PIT Modes of Operation

The PIT unit can operate in the following modes:

- PIT enable/disable mode
- PIT periodic interrupt enable/disable mode
- PIT internal/external input clock mode

## 5.6.4 PIT External Signal Description

This section provides an overview and detailed descriptions of the PIT signals.

There is one distinct external input signal (PIT clock), defined in [Table 5-46](#).

Table 5-46. PIT Signal Properties

Name	Port	Function	I/O	Reset	Pull Up
PIT_CLK	PIT_CLK	Periodic interval timer.	I	N/A	—

Table 5-47 describes of the external PIT signal.

**Table 5-47. PIT External Signal—Detailed Signal Descriptions**

Signal	I/O	Description	
PIT_CLK	I	This signal is used as the timebase for the periodic interval timer module.	
		<b>State Meaning</b>	—
		<b>Timing</b>	—

## 5.6.5 PIT Memory Map/Register Definition

The PIT programmable register map occupies 32 bytes of memory-mapped space. Reading undefined portions of the memory map returns all zeros; writing has no effect.

All PIT registers are 32 bits wide and reside on 32-bit address boundaries and should only be accessed as 32-bit quantities.

All addresses used in this chapter are offsets from PIT base, as defined in Chapter 2, “Memory Map.”


Table 5-48 shows the PIT memory map.

**Table 5-48. PIT Register Address Map**

Offset	Register	Access	Reset Value	Section/ Page
0x00	Periodic interval timer control register (PTCNR)	R/W	0x0000_0000	5.6.5.1/5-45
0x04	Periodic interval timer load register (PTLDR)	R/W	0x0000_0000	5.6.5.2/5-46
0x08	Periodic interval timer prescale register (PTPSR)	R/W	0x0000_0000	5.6.5.3/5-47
0x0C	Periodic interval timer counter register (PTCTR)	R	0x0000_0000	5.6.5.4/5-47
0x10	Periodic interval timer event register (PTEVR)	w1c	0x0000_0000	5.6.5.5/5-47
0x14–0x1F	Reserved	—	—	

### 5.6.5.1 Periodic Interval Timer Control Register (PTCNR)

The periodic interval timer control register (PTCNR), shown in Figure 5-33, is used to enable the different PIT functions. The register can be read at any time.



**Figure 5-33. Periodic Interval Timer Control Register (PTCNR)**


Table 5-49 defines the bit fields of PTCNR.

**Table 5-49. PTCNR Bit Settings**

Bits	Name	Description
0–23	—	Write reserved, read = 0
24	CLEN	Clock enable control bit. Controls the counting of the PIT. When the PIT's clock is disabled, the counter maintains its old value. When the counter's clock is enabled, it continues counting using the previous value. 0 Disable counter. 1 Enable counter.
25	CLIN	Input clock control bit. The input clock to the PIT can be either an internal system clock or an external PIT clock. 0 The input clock to the periodic interrupt timer is internal system clock. 1 The input clock to the periodic interrupt timer is external PIT clock.
26–30	—	Write reserved, read = 0
31	PIM	Periodic interrupt mask bit. Used to enable or disable (mask) the PIT periodic interrupt. 0 Periodic interrupt generation disabled. 1 Periodic interrupt generation enabled.

### 5.6.5.2 Periodic Interval Timer Load Register (PTLDR)

The periodic interval timer load register (PTLDR), shown in Figure 5-34, contains the 32-bit value to be loaded in a 32-bit PIT counter.



**Figure 5-34. Periodic Interval Timer Load Register (PTLDR)**

Table 5-50 defines the bit fields of PTLDR.

**Table 5-50. PTLDR Bit Settings**

Bits	Name	Description
0–31	CLDV	Contains the 32-bit value to be loaded in a 32-bit PIT counter.

### 5.6.5.3 Periodic Interval Timer Prescale Register (PTPSR)

The periodic interval timer prescale register (PTPSR), shown in Figure 5-35, is a read/write register that used to configure the PIT prescaler's value.




Figure 5-35. Periodic Interval Timer Prescale Register (PTPSR)

Table 5-51 defines the bit fields of PTPSR.

Table 5-51. PTPSR Bit Settings

Bits	Name	Description
0–31	PRSC	PIT prescaler bits. Selects the input clock divider to generate the PIT counter clock. The prescaler is programmed to divide the PIT clock input by values from 1 to 4,294,967,296. The value 0x0000 divides the clock by 1 and 0xFFFF_FFFF divides the clock by 4,294,967,296. To accurately predict the timing of the next count, change the PRSC bit only when the enable bit PTCNR[CLE] is clear. Changing PRSC resets the prescaler counter. System reset and the loading of a new value into the counter also reset the prescaler counter. Clearing the PTCNR[CLE] bit stops the prescaler counter.

### 5.6.5.4 Periodic Interval Timer Counter Register (PTCTR)

The periodic interval timer counter register (PTCTR), shown in Figure 5-36, is a read-only register that shows the current value in the PIT counter. The PTCTR counter is not affected by reads or writes.




Figure 5-36. Periodic Interval Timer Counter Register (PTCTR)

Table 5-52 defines the bit fields of PTCTR.

Table 5-52. PTCTR Bit Settings


Bits	Name	Description
0–31	CNTV	PIT counter value field. Contains the current value of the time counter. This is a read-only field. Writes have no effect on PTCTR[CNTV].

### 5.6.5.5 Periodic Interval Timer Event Register (PTEVR)

The periodic interval timer event register (PTEVR), shown in Figure 5-37, is used to report the source of the interrupts. The register can be read at any time.

PTEVR bits are cleared by writing ones. Writing zeros does not affect the value of the status bits.

## System Configuration



**Figure 5-37. Periodic Interval Timer Event Register (PTEVR)**

Table 5-53 defines the bit fields of PTEVR.

**Table 5-53. PTEVR Bit Settings**

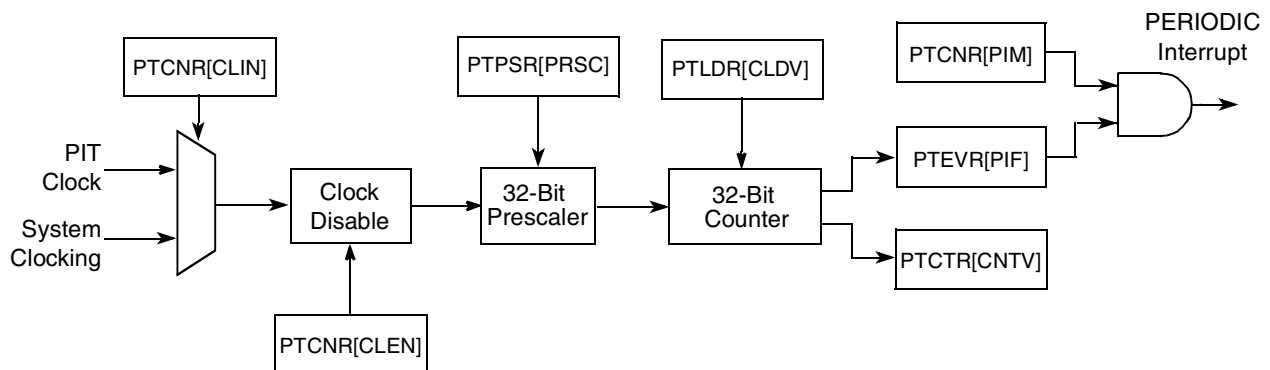
Bits	Name	Description
0–30	—	Write reserved, read = 0
31	PIF	Periodic interrupt flag bit. Used to indicate the periodic interrupt. Its asserted if the PIT issues an interrupt after the SPMPIT counter counts to zero. This status bit should be cleared by software.

## 5.6.6 Functional Description

### 5.6.6.1 Periodic Interval Timer Unit

The PIT generates periodic interrupts for use with a real-time operating system or the application software. It consists of a 32-bit down-counter which is decremented by a clock derived from the CSB clock or from the PIT clock. The 32-bit counter decrements to zero when loaded with a initial value from the periodic interval timer load register (PTLDR). After the timer reaches zero, PTEVR[PIF] is set and an interrupt is generated if PTCNR[PIM] = 1. At the next count cycle, the value in the PTLDR[CLDV] is loaded into the counter and the process repeats. When a new value is loaded into the PTLDR[CLDV], the PIT is updated, the prescaler counter is reset, and the counter begins counting. Setting of PTEVR[PIF] generates an interrupt, that remains pending until PTEVR[PIF] is cleared. If PTEVR[PIF] is set again before being cleared, the interrupt remains pending until PTEVR[PIF] is cleared. Any write to the PTLDR[CLDV] stops the current countdown and the count resumes with the new value in PTLDR[CLDV]. If PTCNR[CLEN] = 0, the PIT cannot count and retains the old count value. PTCTR contain the PIT current value. The PIT function can be disabled if needed.

Figure 5-38 shows the functional PIT block diagram.



**Figure 5-38. Periodic Interval Timer Functional Block Diagram**



### 5.6.6.2 PIT Operational Modes

The PIT unit can operate in the following modes:

- PIT enable/disable mode:
 

The PTCNR[CLEN] bit enables the PIT timer. It should be set by software after a system reset to enable the PIT timer.

  - PIT disable mode (PTCNR[CLEN] = 0). When the PIT's clock is disabled, counter maintains its old value.
  - PIT enable mode (PTCNR[CLEN] = 1). When the counter's clock is enabled, it continues counting using the previous value.
- PIT periodic interrupt enable/disable mode:
  - PIT periodic interrupt enable mode (PTCNR[PIM] = 1). After the PIT's 32-bit counter reaches zero, the PIT sets the PTEVR[PIF] flag and generates an interrupt.
  - PIT periodic interrupt disable mode (PTCNR[PIM] = 0). After the PIT's 32-bit counter reaches zero, the PIT sets the PTEVR[PIF] flag but does not generate an interrupt.
- PIT internal/external input clock mode:
 

The input clock to the PIT may be an internal system clock or the PIT clock.

  - PIT use the internal input clock mode (PTCNR[CLIN] = 0)
  - PIT use the PIT clock (PTCNR[CLIN] = 1)

### 5.6.7 PIT Programming Guidelines

The following initialization sequence of PIT is recommended:

1. Write to PTPSR to set the PIT prescaler to the desired value
2. Write to PTLDR to initialize the PIT initial value
3. Write to PTCNR to configure and start the PIT operation: PIT input clock source, periodic interrupt mask, PIT clock enable.

See [Section 5.5.7, “RTC Programming Guidelines,”](#) for real-time clock programming guidelines.

## 5.7 General-Purpose Timers (GTMs)

The following sections describe theory of operation of the general purpose (global) timer module, including a definition of the external signals and the functionality. Additionally, the configuration, control, and status registers are described. Note that individual chapters in this book describe additional specific initialization aspects for each individual block.

### 5.7.1 GTM Overview

Each global timer module (GTM) includes four identical 16-bit general-purpose timers, two 32-bit timers or one 64-bit timer. Each GTM timer consists of a timer prescale register (GTPSR), a timer mode register (GTMDR) a timer capture register (GTCPR), a timer counter register (GTCNR), a timer reference register

(GTRFR), a timer event register (GTEVR), and a timer global configuration register (GTCFR). The GTPSRs and the GTMDRs contain the primary and secondary prescalers, programmed by the user.

Figure 5-39 shows the functional GTM block diagram.




Figure 5-39. Global Timers Block Diagram

## 5.7.2 GTM Features

The key features of the timer include the following:

- The maximum input clock is the system bus clock
- Four 16-bit programmable timers
- Two timers cascaded internally or externally to form a 32-bit timer
- One timer cascaded internally or externally to form a 64-bit timer
- Maximum period of ~412.3 seconds (at 167-MHz bus clock and prescaler = 256) for 16-bit timer
- Maximum period of ~12.8825 seconds (at 167-MHz bus clock and prescaler = 256) for 32-bit timer
- Maximum period of thousands of years (at 167-MHz bus clock and prescaler = 256) for 64-bit timer
- 3-nanosecond timer resolution (at 167-MHz bus clock and no prescaler)
- Resolution and maximum period can be traded off by selecting prescaler divisor
- Three programmable input clock sources for the timer prescalers

- Input capture capability
- Output compare with programmable mode for the output pin
- Free run and restart modes
- Functional and programming compatibility with MPC8260 timers

### 5.7.3 GTM Modes of Operation

The GTM unit can operate in the following modes.

#### 5.7.3.1 Cascaded Modes

$GTCFR_n[PCAS]$  and  $GTCFR_2[SCAS]$  are used to put the timers into different cascaded modes:

- Non-cascaded mode: Each timer (timer 1, timer 2, timer 3, and timer 4), function as a independent 16-bit timer with a 16-bit  $GTRFR$ ,  $GTCPR$ ,  $GTMDR$ , and  $GTCNR$ . In this mode, the non-cascaded  $GTRFR$ ,  $GTCPR$ , and  $GTCNR$  should be referenced with corresponding 16-bit bus cycles.
- Pair-cascaded mode: In this mode, two 16-bit timers can be internally cascaded to form a 32-bit counter: timer 1 can be internally cascaded to timer 2 and timer 3 may be internally cascaded to timer 4. Because the decision to cascade timers is made independently, the user has the option of selecting two 16-bit timers and one 32-bit timer, or two 32-bit timers. When working in the pair-cascaded mode, the cascaded  $GTRFR$ ,  $GTCPR$ , and  $GTCNR$  should be referenced with 32-bit bus cycles.
- Super-cascaded mode: In this mode, all four 16-bit timers can be internally cascaded to form a 64-bit counter. When working in the super-cascaded mode, the cascaded  $GTRFR$ ,  $GTCPR$ , and  $GTCNR$  should be referenced with two 32-bit bus cycles.

#### 5.7.3.2 Clock Source Modes

The clock input to the timer's prescaler can be selected from three sources:

- The system clock
- The system slow go clock (system bus clock internally divided by 16)
- The corresponding  $TIN_x$  pin

#### 5.7.3.3 Reference Modes

Each timer can be configured to count until a reference is reached and then either begin a new time count immediately or continue to run. The  $FRR$  bit of the corresponding  $GTMRR$  selects each mode.

- Free run reference mode. The corresponding timer count continues to increment after the reference value is reached.
- Reset reference mode. The corresponding timer count is reset immediately after the reference value is reached.

### 5.7.3.4 Capture Modes

Each timer has a 16-bit field in GTCPR, used to latch the value of the counter when a defined transition of TINx is sensed by the corresponding input capture edge detector.

- Normal gate mode enables the count on a falling edge of the  $\overline{\text{TGATE}}$  pin and disables the count on the rising edge of  $\overline{\text{TGATE}}$ . This mode allows the timer to count conditionally, based on the state of  $\overline{\text{TGATE}}$ .
- The restart gate mode performs the same function as normal mode, except it also resets the counter on the falling edge of the  $\overline{\text{TGATE}}$  pin. This mode has applications in pulse interval measurement and bus monitoring.

### 5.7.4 GTM External Signal Description

This section provides an overview and detailed descriptions of the GTM signals.

There are four distinct external input timer capture signals (TIN1, TIN2, TIN3, and TIN4), four distinct external input timer gate signals ( $\overline{\text{TGATE1}}$ ,  $\overline{\text{TGATE2}}$ ,  $\overline{\text{TGATE3}}$ , and  $\overline{\text{TGATE4}}$ ), and four distinct external timer output signals ( $\overline{\text{TOUT1}}$ ,  $\overline{\text{TOUT2}}$ ,  $\overline{\text{TOUT3}}$ , and  $\overline{\text{TOUT4}}$ ). The GTM interface signals are defined in [Table 5-54](#).

**Table 5-54. GTM Signal Properties**

Name	Port	Function	I/O	Reset	Require Pull Up
TIN1	TIN1	Global timer 1 capture control signal	I	0	No
TIN2	TIN2	Global timer 2 capture control signal	I	0	No
TIN3	TIN3	Global timer 3 capture control signal	I	0	No
TIN4	TIN4	Global timer 4 capture control signal	I	0	No
$\overline{\text{TGATE1}}$	$\overline{\text{TGATE1}}$	Global timer 1 counter gate control signal	I	0	No
$\overline{\text{TGATE2}}$	$\overline{\text{TGATE2}}$	Global timer 2 counter gate control signal	I	0	No
$\overline{\text{TGATE3}}$	$\overline{\text{TGATE3}}$	Global timer 3 counter gate control signal	I	0	No
$\overline{\text{TGATE4}}$	$\overline{\text{TGATE4}}$	Global timer 4 counter gate control signal	I	0	No
$\overline{\text{TOUT1}}$	$\overline{\text{TOUT1}}$	Global timer 1 counter output signal	O	1	No
$\overline{\text{TOUT2}}$	$\overline{\text{TOUT2}}$	Global timer 2 counter output signal	O	1	No
$\overline{\text{TOUT3}}$	$\overline{\text{TOUT3}}$	Global timer 3 counter output signal	O	1	No
$\overline{\text{TOUT4}}$	$\overline{\text{TOUT4}}$	Global timer 4 counter output signal	O	1	No

Table 5-55 provides detailed descriptions of the external GTM signals.

**Table 5-55. GTM External Signals—Detailed Signal Descriptions**

Signal	I/O	Description
TIN <sub>n</sub>	I	Global timer capture control signal. Used to latch the value of the counter when a defined transition of TIN <sub>n</sub> is sensed by the corresponding input capture edge detector.
		<b>State Meaning</b> Asserted/Negated—According to the programmed polarity by the corresponding GTMDR <sub>n</sub> [CE]. Each timer has a 16-bit GTCPR used to latch the value of the counter when a defined transition of TIN <sub>n</sub> is sensed by the corresponding input capture edge detector. Upon a capture or reference event, the corresponding GTEVR bit is set and a maskable interrupt request is issued to the interrupt controller.
		<b>Timing</b> Assertion/Negation—Asynchronous to internal bus clock. TIN <sub>n</sub> is internally synchronized to the system bus clock. If TIN <sub>n</sub> meets the asynchronous input setup time, the value of counter is captured after one system bus clock when working with the internal clock.
$\overline{\text{TGATE}}_n$	I	Global timer counter gate control signal. Used to gate/restart the counter when a defined transition of $\overline{\text{TGATE}}_n$ is sensed by the corresponding input capture edge detector.
		<b>State Meaning</b> Asserted/Negated—According to the programmed polarity by the corresponding GTCFR[GMx] bits. In a reset gate mode (GTCFR[GM <sub>n</sub> ] = 0), the $\overline{\text{TGATE}}_n$ pin is used to enable/disable count. A falling $\overline{\text{TGATE}}_n$ pin enables and restarts the count and a rising edge of $\overline{\text{TGATE}}_n$ disables the count. In a normal gate mode (GTCFR[GM <sub>n</sub> ] = 1), the $\overline{\text{TGATE}}_n$ have similar functionality, except the falling edge of $\overline{\text{TGATE}}_n$ does not restart the appropriate count value in GTCNR <sub>n</sub> [CNV <sub>n</sub> ].
		<b>Timing</b> Assertion/Negation—Asynchronous to internal bus clock. $\overline{\text{TGATE}}_n$ is internally synchronized to the system bus clock. If $\overline{\text{TGATE}}_n$ meets the asynchronous input setup time, the counter begins counting after one system bus clock when working with the internal clock.
$\overline{\text{TOUT}}_n$	O	Global timer counter output signal. The GTM output a signal on the timer output pin $\overline{\text{TOUT}}_n$ when the reference value is reached.
		<b>State Meaning</b> Asserted/Negated—According to the programmed polarity by the corresponding GTMDR <sub>n</sub> [OM <sub>n</sub> ]. <ol style="list-style-type: none"> <li>Active-low pulse on <math>\overline{\text{TOUT}}_n</math> for one timer input clock cycle as defined by the GTMDR<sub>n</sub>[ICLK<sub>n</sub>] bits (GTMDR<sub>n</sub>[OM<sub>n</sub>] = 1). Thus, <math>\overline{\text{TOUT}}_n</math> may be low for one general system clock period, one general system slow go clock period, or one TIN<sub>n</sub> pin clock cycle period.</li> <li>Toggle the <math>\overline{\text{TOUT}}_n</math> pin (GTMDR<sub>n</sub>[OM<sub>n</sub>] = 0). <math>\overline{\text{TOUT}}_n</math> changes occur on the rising edge of the system clock.</li> </ol>
		<b>Timing</b> Assertion/Negation— $\overline{\text{TOUT}}_n$ changes occur on the rising edge of the system clock.

## 5.7.5 GTM Memory Map/Register Definition

The GTM programmable register map occupies 64 bytes of memory-mapped space. Reading undefined portions of the memory map returns all zeros; writing has no effect.

All GTM registers are 8 or 16 bits wide, located on 8-bit or 16-bit address boundaries, and should only be accessed as 8-bit or 16-bit quantities. All addresses used in this chapter are offsets from GPT Base, as defined in Chapter 2, “Memory Map.”

Table 5-56 shows memory map of the GTM.

**Table 5-56. GTM Register Address Map**

Offset	Register	Access	Reset Value	Section/ Page
<b>General Purpose (Global) Timer Module 1—Block Base Address 0x0_0500</b>				
0x00	Timer 1 and 2 global timers configuration register (GTCFR1)	R/W	0x00	<a href="#">5.7.5.1/5-55</a>
0x01–0x03	Reserved	—	—	—
0x04	Timer 3 and 4 global timers configuration register (GTCFR2)	R/W	0x00	<a href="#">5.7.5.1/5-55</a>
0x05–0x0F	Reserved	—	—	—
0x10	Timer 1 global timers mode register (GTMDR1)	R/W	0x0000	<a href="#">5.7.5.2/5-58</a>
0x12	Timer 2 global timers mode register (GTMDR2)			
0x14	Timer 1 global timers reference register (GTRFR1)	R/W	0xFFFF	<a href="#">5.7.5.3/5-59</a>
0x16	Timer 2 global timers reference register (GTRFR2)			
0x18	Timer 1 global timers capture register (GTCPR1)	R/W	0x0000	<a href="#">5.7.5.4/5-59</a>
0x1A	Timer 2 global timers capture register (GTCPR2)			
0x1C	Timer 1 global timers counter register (GTCNR1)	R/W	0x0000	<a href="#">5.7.5.5/5-60</a>
0x1E	Timer 2 global timers counter register (GTCNR2)			
0x20	Timer 3 global timers mode register (GTMDR3)	R/W	0x0000	<a href="#">5.7.5.2/5-58</a>
0x22	Timer 4 global timers mode register (GTMDR4)			
0x24	Timer 3 global timers reference register (GTRFR3)	R/W	0xFFFF	<a href="#">5.7.5.3/5-59</a>
0x26	Timer 4 global timers reference register (GTRFR4)			
0x28	Timer 3 global timers capture register (GTCPR3)	R	0x0000	<a href="#">5.7.5.4/5-59</a>
0x2A	Timer 4 global timers capture register (GTCPR4)			
0x2C	Timer 3 global timers counter register (GTCNR3)	R/W	0x0000	<a href="#">5.7.5.5/5-60</a>
0x2E	Timer 4 global timers counter register (GTCNR4)			
0x30	Timer 1 global timers event register (GTEVR1)	w1c	0x0000	<a href="#">5.7.5.6/5-60</a>
0x32	Timer 2 global timers event register (GTEVR2)			
0x34	Timer 3 global timers event register (GTEVR3)			
0x36	Timer 4 global timers event register (GTEVR4)			
0x38	Timer 1 global timers prescale register (GTPSR1)	R/W	0x0003	<a href="#">5.7.5.7/5-61</a>
0x3A	Timer 2 global timers prescale register (GTPSR2)			
0x3C	Timer 3 global timers prescale register (GTPSR3)			
0x3E	Timer 4 global timers prescale register (GTPSR4)			
General Purpose (Global) Timer Module 2: All registers defined for GTM1 are also defined for GTM2; the base address of GTM2 registers is 0x0_06nn.				

### 5.7.5.1 Global Timers Configuration Registers (GTCFR $n$ )

The global timers configuration registers (GTCFR1 and GTCFR2), shown in Figure 5-40 and Figure 5-41, contain configuration parameters used by the timers. These registers allow simultaneous starting, stopping and resetting of a pair of timers (1 and 2 or 3 and 4) or of a groups of timers (1, 2, 3, and 4) if one bus cycle is used. GTCFR is cleared by reset.

#### NOTE

For proper operation of the timers, do not change the modes of operation and enable the timer in the same register write operation. The modes can be changed when GTCFR $n$ [RST $n$ ] is cleared. However, when GTCFR $n$ [RST $n$ ] are set, they are the only bits that can be changed.

Offset 0x00

Access: Read/Write




Figure 5-40. Global Timers Configuration Register 1 (GTCFR1)

Table 5-57 defines the bit fields of GTCFR1.


Table 5-57. GTCFR1 Bit Settings

Bits	Name	Description
0	PCAS	Pair-cascade mode 0 Normal operation 1 Timers 1 and 2 cascade to form a 32-bit timer. <b>Note:</b> This bit is ignored in super-cascade mode (GTCFR2[SCAS] = 1). <b>Note:</b> It is allowed to change the value of this bit only when the corresponding timers are in reset mode. Thus, the user should first clear the RST1 and RST2 bits (without changing PCAS) and then, in a separate write to the register, change the value of PCAS.
1	BCM	Backward compatible mode 0 Provide backward compatibility to PowerQUICC II family timers. In this mode GTCFR1[GM2] bit will control the gate mode for timers 1 and 2 and GTCFR2[GM4] bit will control the gate mode for timers 3 and 4. GTCFR1[GM1] and GTCFR2[GM3] bits are ignored. 1 Normal operational mode
2	STP2	Stop timer 2 0 Normal operation 1 Reduce power consumption of the corresponding timer. This bit stops all clocks to the timer 2, except the Register Interface clock, which allows to read and write timer registers. The clocks to the timer remain stopped until the user clears this bit or a hardware reset occurs.
3	RST2	Reset timer 2 0 Reset the timer 2, including GTMDR2, GTRFR2, GTCNR2, GTCPR2, and GTEVR2 (a software reset is identical to an external reset). 1 Enable the corresponding timer if the STP2 bit is cleared.

**Table 5-57. GTCFR1 Bit Settings (continued)**

Bits	Name	Description
4	GM2	<p>Gate mode for <math>\overline{\text{TGATE2}}</math></p> <p>0 Restart gate mode. The <math>\overline{\text{TGATE2}}</math> pin is used to enable/disable count. A low level of <math>\overline{\text{TGATE2}}</math> enables and a falling edge of <math>\overline{\text{TGATE2}}</math> restarts the count (reset the dynamic counter's count value to 0) and a high level of <math>\overline{\text{TGATE2}}</math> disables the count.</p> <p>1 Normal gate mode. This mode is the same as 0, except the falling edge of <math>\overline{\text{TGATE2}}</math> does not restart the appropriate count value in GTCNR2[CNV2].</p>
5	GM1	<p>Gate mode for <math>\overline{\text{TGATE1}}</math></p> <p>0 Restart gate mode. The <math>\overline{\text{TGATE1}}</math> is used to enable/disable count. A low level of <math>\overline{\text{TGATE1}}</math> enables and a falling edge of <math>\overline{\text{TGATE1}}</math> restarts the count (reset the dynamic counter's count value to 0) and a high level of <math>\overline{\text{TGATE1}}</math> disables the count.</p> <p>1 Normal gate mode. This mode is the same as 0, except the falling edge of <math>\overline{\text{TGATE1}}</math> does not restart the appropriate count value in GTCNR1[CNV1].</p> <p><b>Note:</b> In backward compatible mode (GTCFR1[BCM] = 0) this bit is ignored. GTCFR1[GM2] bit will control the gate mode for timers 1 and 2.</p>
6	STP1	<p>Stop timer 1</p> <p>0 Normal operation</p> <p>1 Reduce power consumption of the corresponding timer. This bit stops all clocks to the timer 1, except the register interface clock, which allows to read and write timer registers. The clocks to the timer remain stopped until the user clears this bit or a hardware reset occurs.</p>
7	RST1	<p>Reset timer 1</p> <p>0 Reset the timer 1, including GTMDR1, GTRFR1, GTCNR1, GTCPR1, and GTEVR1 (a software reset is identical to an external reset).</p> <p>1 Enable the corresponding timer if the STP1 bit is cleared.</p>

The GTCFR2 register is shown in [Figure 5-41](#).



**Figure 5-41. Global Timers Configuration Register 2 (GTCFR2)**



Table 5-58 defines the bit fields of GTCFR2.

**Table 5-58. GTCFR2 Bit Settings**

Bits	Name	Description
0	PCAS	<p>Pair-cascade mode</p> <p>0 Normal operation.</p> <p>1 Timers 3 and 4 cascade to form a 32-bit timer.</p> <p><b>Note:</b> This bit is ignored in super-cascade mode (GTCFR2[SCAS] = 1).</p> <p><b>Note:</b> It is allowed to change the value of this bit only when the corresponding timers are in reset mode. Thus, the user should first clear the RST3 and RST4 bits (without changing PCAS) and then, in a separate write to the register, change the value of PCAS.</p>
1	SCAS	<p>Super cascade mode</p> <p>0 Normal operation</p> <p>1 Timers 1, 2, 3 and 4 cascade to form a 64-bit timer.</p> <p><b>Note:</b> In super-cascade mode (GTCFR2[SCAS] = 1) the pair-cascade mode bits are ignored, (GTCFR1/2[PCAS] = Don't Care).</p> <p><b>Note:</b> It is allowed to change the value of this bit only when the corresponding timers are in reset mode. Thus, the user should first clear the RST1, RST2, RST3, and RST4 bits (without changing SCAS) and then, in a separate write to the register, change the value of SCAS.</p>
2	STP4	<p>Stop timer 4</p> <p>0 Normal operation</p> <p>1 Reduce power consumption of the corresponding timer. This bit stops all clocks to the timer 4, except the register interface clock, which allows to read and write timer registers. The clocks to the timer remain stopped until the user clears this bit or a hardware reset occurs.</p>
3	RST4	<p>Reset timer 4</p> <p>0 Reset the timer 4, including GTMDR4, GTRFR4, GTCNR4, GTCPR4, and GTEVR4 (a software reset is identical to an external reset).</p> <p>1 Enable the corresponding timer if the STP4 bit is cleared.</p>
4	GM4	<p>Gate mode for <math>\overline{\text{TGATE4}}</math></p> <p>0 Restart gate mode. The <math>\overline{\text{TGATE4}}</math> is used to enable/disable count. A low level of <math>\overline{\text{TGATE4}}</math> enables and a falling edge of <math>\overline{\text{TGATE4}}</math> restarts the count (reset the dynamic counter's count value to 0) and a high level of <math>\overline{\text{TGATE4}}</math> disables the count.</p> <p>1 Normal gate mode. This mode is the same as 0, except the falling edge of <math>\overline{\text{TGATE4}}</math> does not restart the appropriate count value in GTCNR4[CNV4].</p>
5	GM3	<p>Gate mode for <math>\overline{\text{TGATE3}}</math></p> <p>0 Restart gate mode. The <math>\overline{\text{TGATE3}}</math> is used to enable/disable count. A low level of <math>\overline{\text{TGATE3}}</math> enables and a falling edge of <math>\overline{\text{TGATE3}}</math> restarts the count (reset the dynamic counter's count value to 0) and a high level of <math>\overline{\text{TGATE3}}</math> disables the count.</p> <p>1 Normal gate mode. This mode is the same as 0, except the falling edge of <math>\overline{\text{TGATE3}}</math> does not restart the appropriate count value in GTCNR3[CNV3].</p> <p><b>Note:</b> In backward compatible mode (GTCFR1[BCM] = 0) this bit is ignored. The GTCFR2[GM4] bit controls the gate mode for timers 3 and 4.</p>
6	STP3	<p>Stop timer 3</p> <p>0 Normal operation</p> <p>1 Reduce power consumption of the corresponding timer. This bit stops all clocks to the timer 3, except the register interface clock, which allows to read and write timer registers. The clocks to the timer remain stopped until the user clears this bit or a hardware reset occurs.</p>
7	RST3	<p>Reset timer 3</p> <p>0 Reset the timer 3, including GTMDR3, GTRFR3, GTCNR3, GTCPR3, and GTEVR3 (a software reset is identical to an external reset).</p> <p>1 Enable the corresponding timer if the STP3 bit is cleared.</p>

### 5.7.5.2 Global Timers Mode Registers (GTMDR1–GTMDR4)

The global timers mode registers (GTMDR1, GTMDR2, GTMDR3, and GTMDR4) are shown in Figure 5-42.

Erratic behavior may occur if GTCFR1 and GTCFR2 are not initialized before the GTMDR $n$ . Only GTCFR $n$ [RST $n$ ] and GTCFR $n$ [STP $n$ ] can be modified at any time.




Figure 5-42. Global Timers Mode Registers (GTMDR1–GTMDR4)

Table 5-59 defines the bit fields of GTMDR.

Table 5-59. GTMDR Bit Settings

Bits	Name	Description
0–7	SPS	Secondary prescaler value The secondary prescaler is programmed to divide the clock input to corresponding timer by values from 1 to 256. The value 0x00 divides the clock by 1 and 0xFF divides the clock by 256.
8–9	CE	Capture edge and enable interrupt 00 Disable interrupt on capture event; capture function is disabled 01 Capture on rising TIN $n$ edge only and enable interrupt on capture event. 10 Capture on falling TIN $n$ edge only and enable interrupt on capture event. 11 Capture on any TIN $n$ edge and enable interrupt on capture event. <b>Note:</b> The frequency of TIN $n$ should be slower than system clock (TIN $n$ is sampled internally by system clock to detect TIN $n$ 's rising/falling edge before updating the counter)
10	OM	Output mode 0 Toggle $\overline{\text{TOUT}}_n$ every time when the corresponding timer matches its reference value. 1 Active-low pulse on $\overline{\text{TOUT}}_n$ for one timer input clock cycle (4 input clock cycles for the system clock) as defined by the ICLK $n$ bits. Thus, $\overline{\text{TOUT}}_n$ may be low for four general system clocks, one general system slow go clock period, or one TIN $n$ pin clock cycle period. <b>Note:</b> $\overline{\text{TOUT}}_n$ changes are internally synchronized to the rising edge of the system clock
11	ORI	Output reference interrupt enable 0 Disable interrupt for reference reached (does not affect interrupt on capture function). 1 Enable interrupt on reaching the reference value.
12	FRR	Free run/restart mode 0 Free run. The timer count continues to increment after the reference value is reached. 1 Restart. The timer count is reset immediately after the reference value is reached.

Table 5-59. GTMDR Bit Settings (continued)

Bits	Name	Description
13–14	ICLK	Input clock source for the timer. 00 Internally cascaded input. This selection means: For ICLK1, the timer 1 input is the output of timer 2. For ICLK2, the timer 1 input is the output of timer 2, the timer 2 input is the output of timer 3, the timer 3 input is the output of timer 4. For ICLK3, the timer 3 input is the output of timer 4. For ICLK4 this selection means no input clock is provided to the timer. 01 Internal general system bus clock. 10 Internal slow go clock (divided by 16 system bus clock). 11 TIN $n$ : corresponding TIN1, TIN2, TIN3, or TIN4 pin (falling edge).
15	GE	Gate enable 0 The $\overline{\text{TGATE}}_n$ signal is ignored. 1 The $\overline{\text{TGATE}}_n$ signal is used to control the timer.

### 5.7.5.3 Global Timers Reference Registers (GTRFR1–GTRFR4)

Global timers reference registers, shown in Figure 5-43, are 16-bit memory-mapped, read/write registers containing the 16-bit reference values for each timer's timeout. The reference value is not reached until  $\text{GTCNR}_n[\text{CNV}]$  increments to the value in  $\text{GTRFR}_n[\text{TRV}]$ .




Figure 5-43. Global Timers Reference Registers (GTRFR1–GTRFR4)

Table 5-60 defines the bit fields of GTRFR.

Table 5-60. GTRFR Bit Settings

Bits	Name	Description
0–15	TRV	Timeout reference value. 16-bit timeout reference value for the corresponding timer. Set to all ones by reset.

### 5.7.5.4 Global Timers Capture Registers (GTCPR1–GTCPR4)

Global timers capture registers ( $\text{GTCPR}_1$ ,  $\text{GTCPR}_2$ ,  $\text{GTCPR}_3$ , and  $\text{GTCPR}_4$ ), shown in Figure 5-44, are used to latch the value of the counters according to  $\text{GTMDR}_n[\text{CE}]$ .




Figure 5-44. Global Timers Capture Registers (GTCPR1–GTCPR4)

Table 5-61 defines the bit fields of  $GTCPR_n$ .

Table 5-61.  $GTCPR_n$  Bit Settings

Bits	Name	Description
0–15	LCV	Latched counter value. Corresponding timer's 16-bit latched value.

### 5.7.5.5 Global Timers Counter Registers (GTCNR1–GTCNR4)

Global timers counter registers (GTCNR1, GTCNR2, GTCNR3, and GTCNR4), shown in Figure 5-45, are four 16-bit, memory-mapped, read/write up-counters. A read cycle to a  $GTCNR_n[CNV]$  fields yields the current value of the appropriate timer but does not affect the counting operation. A write cycle to a  $GTCNR_n[CNV]$  field sets the register to the written value, causing its corresponding primary and secondary prescaler counters to be reset.




Figure 5-45. Global Timers Counter Registers (GTCNR1—GTCNR4)

Table 5-62 defines the bit fields of GTCNR.

Table 5-62. GTCNR Bit Settings

Bits	Name	Description
0–15	CNV	Counter value. Corresponding timer's 16-bit read/write up-counter value.

### 5.7.5.6 Global Timers Event Registers (GTEVR1–GTEVR4)

Global timers event registers (GTEVR1, GTEVR2, GTEVR3, and GTEVR4), shown in Figure 5-46, are used to report events recognized by any of the timers. On recognition of an output reference event, the appropriate timer sets  $GTEVR_n[REF]$ , regardless of the corresponding  $GTMDR_n[ORI]$ . The capture event is only set if it is enabled by  $GTMDR_n[CE]$ . GTEVRs appear as memory-mapped registers to users, which can be read at any time.

$GTEVR_n$  bits are cleared by writing ones to them (writing zeros does not affect bit values). Both bits must be reset before the timer negates the interrupt to the interrupt controller.




Figure 5-46. Global Timers Event Registers (GTEVR1—GTEVR4)

Table 5-63 defines the bit fields of GTEVR $n$ .

Table 5-63. GTEVR $n$  Bit Settings

Bits	Name	Description
0–13	—	Reserved, should be cleared.
14	REF	Output reference event 0 No event 1 The counter reached the GTRFR $n$ [TRV] value. GTMDR $n$ [ORI] is used to enable the interrupt request caused by this event.
15	CAP	Counter capture event Corresponding timer's 16-bit read/write up-counter value. 0 No event 1 The counter value has been latched into the GTCPR $n$ [LCV]. GTMDR $n$ [CE] is used to enable generation of this event.

### 5.7.5.7 Global Timers Prescale Registers (GTPSR1–GTPSR4)

The global timers prescale registers (GTPSR1, GTPSR2, GTPSR3, and GTPSR4) are shown in Figure 5-47.

Erratic behavior may occur if GTPSR $n$  is not initialized before the corresponding GTMDR $n$ .



Figure 5-47. Global Timers Prescale Registers (GTPSR1–GTPSR4)

Table 5-64 defines the bit fields of GTPSR $n$ .

Table 5-64. GTPSR $n$  Bit Settings

Bits	Name	Description
0–7	—	Reserved, should be cleared.
8–15	PPS	Primary prescaler bits The primary prescaler is programmed to divide the clock input to corresponding timer by values from 1 to 256. The value 0x00 divides the clock by 1 and 0xFF divides the clock by 256.

#### NOTE

The total timer prescale value is calculated as follows:

$$\text{GTM}_{n_{\text{prescaler}}} = (\text{GTPSR}_{n[\text{PPS}]} + 1) \cdot (\text{GTMDR}_{n[\text{SPS}]} + 1)$$

This gives a total prescale range from 1 (GTPSR $n$ [PPS] = 0x00, GTMDR $n$ [SPS] = 0x00) to 65,536 (GTPSR $n$ [PPS] = 0xFF, GTMDR[SPS] = 0xFF).

## 5.7.6 Functional Description

### 5.7.6.1 General-Purpose Timer Units

The clock input to the timer's prescaler can be selected from the following sources:

- The system clock
- The system slow go clock (internally divided by 16)

The general system clock is generated in the clock synthesizer and defaults to the system frequency. However, the general system clock has the option to be divided before it leaves the clock synthesizer. This mode, called slow go, is used to save power. Whatever the resulting frequency of the general system clock, the user can either choose that frequency or the frequency divided by 16 as the input to the prescaler of each timer. Alternatively, the user may prefer  $TIN_n$  to be the clock source.  $TIN_n$  is internally synchronized to the internal clock. If the user has chosen to internally cascade two 16-bit timers to a 32-bit timer, then a timer can use the clock generated by the output of another timer.

The clock input source is selected by the corresponding  $GTMDR_n[ICLK]$  bits. The prescalers ( $GTMDR_n[SPS]$  and  $GTPSR_n[PPS]$ ) can be programmed to divide the clock input by values from 1 to 65,537 and the output of the prescaler is used as an input to the 16-bit counters. The best resolution of the timer is one clock cycle (3 ns at a 333-MHz system clock, for example). The maximum period (when the reference value is all ones and the prescaler divides by 256) for one 16-bit timer is ~50 ms at 333 MHz.

### 5.7.6.2 Reference Modes

Each timer can be configured to count until a reference is reached and then either begin a new time count immediately or continue to run. The FRR bit of the corresponding  $GTMRR$  selects each mode.

- Free run reference mode ( $GTMDR_n[FRR] = 0$ )  
The corresponding timer count continues to increment after the reference value is reached.
- Reset reference mode ( $GTMDR_n[FRR] = 1$ )  
The corresponding timer count is reset immediately after the reference value is reached.

Upon reaching the reference value, the corresponding  $GTEVR_n[REF]$  bit is set and an interrupt is issued if  $GTMDR_n[ORI] = 1$ . The timers can output a signal on the timer output pin  $\overline{TOUT}_n$  if the reference value is reached (selected by the corresponding  $GTMDR_n[OM]$ ). This signal can be an active-low pulse or a toggle of the current output. The output can also be connected internally to the input of another timer, resulting in a 32- or 64-bit timer.

### 5.7.6.3 Capture Modes

In addition, each timer has a 16-bit field in  $GTCPR$ , used to latch the value of the counter when a defined transition of  $TIN_n$  is sensed by the corresponding input capture edge detector. The timers may be gated/restarted by an external gate signals ( $\overline{TGATE}_n$ ) that controls the timers. The type of transition triggering the capture is selected by the corresponding  $GTMDR_n[CE]$  bits. Upon a capture or reference

event, corresponding  $GTEVR_n[REF]$  or  $GTEVR_n[CAP]$  is set and a maskable interrupt request is issued to the interrupt controller.

- Normal gate mode enables the count on a falling edge of  $\overline{TGATE}$  and disables the count on the rising edge of  $\overline{TGATE}$ . This mode allows the timer to count conditionally, based on the state of  $\overline{TGATE}$ .
- The restart gate mode performs the same function as normal mode, except it also resets the counter on the falling edge of  $\overline{TGATE}$ .

This mode has applications in pulse interval measurement and bus monitoring as follows:

- Pulse measurement—The restart gate mode can measure a low pulse on  $\overline{TGATE}$ . The rising edge of  $\overline{TGATE}$  completes the measurement and if  $\overline{TGATE}_n$  is connected externally to  $TIN_n$ , it causes the timer to capture the count value and generate a rising-edge interrupt.
- Bus monitoring—The restart gate mode can detect a signal that is stuck abnormally low. The bus signal should be connected to  $\overline{TGATE}$ . The timer count is reset on the falling edge of the bus signal and if the bus signal does not go high again within the number of user-defined clocks, an interrupt can be generated.

The gate function is enabled in the  $GTMDR$ ; the gate operating mode is selected in the  $GTCFR_n$ .

#### NOTE

$\overline{TGATE}$  is internally synchronized to the system clock. If  $\overline{TGATE}$  meets the asynchronous input setup time, the counter begins counting after one system clock when working with the internal clock.

#### 5.7.6.4 Cascaded Modes

$GTCFR_n[PCAS]$  and  $GTCFR2[SCAS]$  are used to put the timers into different cascaded modes:

- Non-cascaded mode ( $GTCFR_n[PCAS] = 0$  and  $GTGCF2[SCAS] = 0$ )  
If  $GTCFR_n[PCAS] = 0$  and  $GTCFR2[SCAS] = 0$ , the each timer (timer 1, timer 2, timer 3, and timer 4), function as a independent 16-bit timer with a 16-bit  $GTRFR$ ,  $GTCPR$ ,  $GTMDR$ , and  $GTCNR$  for each one (Figure 5-48). When working in the none-cascaded mode, the non-cascaded  $GTRFR$ ,  $GTCPR$ , and  $GTCNR$  should be referenced with appropriate 16-bit bus cycles.




Figure 5-48. Timers Non-Cascaded Mode Block Diagram

- Pair-cascaded mode ( $GTCFR1[PCAS] = 1$  and/or  $GTCFR2[PCAS] = 1$ ,  $GTCFR2[SCAS] = 0$ )

In this mode, two 16-bit timers can be internally cascaded to form a 32-bit counter: timer 1 may be internally cascaded to timer 2 and timer 3 may be internally cascaded to timer 4, as shown in Figure 5-49. Since the decision to cascade timers is made independently, the user has the option of selecting two 16-bit timers and one 32-bit timer ( $GTCFR1[PCAS] = 1$ ,  $GTCFR2[PCAS] = 0$  or  $GTCFR1[PCAS] = 0$ ,  $GTCFR2[PCAS] = 1$ ), or two 32-bit timers ( $GTCFR1[PCAS] = 1$  and  $GTCFR2[PCAS] = 1$ ).

If  $GTCFR1[PCAS] = 1$  and/or  $GTCFR2[SCAS] = 1$ , the two 16-bit timers (timer 1 and timer 2 or timer 3 and timer 4) function as a 32-bit timer with a 32-bit GTRFR, GTCPR, and GTCNR. In this case, GTMDR1/GTMDR3 is ignored, and the modes and functions are defined using GTMDR2/GTMDR4, and GTCFR1/GTCFR2. The capture are controlled from TIN2, and the interrupts are generated from GTEVR2. When working in the pair-cascaded mode, the cascaded GTRFR, GTCPR, and GTCNR should be referenced with 32-bit bus cycles.




Figure 5-49. Timer Pair-Cascaded Mode Block Diagram

- Super-cascaded mode ( $GTCFR2[SCAS] = 1$ )

In this mode, all four 16-bit timers can be internally cascaded to form a 64-bit counter, as shown in Figure 5-50.

If  $GTCFR2[SCAS] = 1$ , the all four 16-bit timers function as a 64-bit timer with a cascaded 32-bit GTRFR, GTCPR, and GTCNR. In this case, registers GTMDR1, GTMDR2, GTMDR3, and GTCFR1 are ignored, and the modes and functions are defined using GTMDR4 and GTCFR2 only. The capture are controlled from TIN4, and the interrupts are generated from GTEVR4. When working in the super-cascaded mode, the cascaded GTRFR, GTCPR, and GTCNR should be referenced with two 32-bit bus cycles.




Figure 5-50. Timers Super-Cascaded Mode Block Diagram



## 5.7.7 Initialization/Application Information

### 5.7.7.1 Programming Guidelines

#### 5.7.7.1.1 GTM Registers

The following initialization sequence of GTM is recommended:

- Write to  $GTCFR_n$  in order to reset, to stop or to configure the appropriate timer's operation: cascaded timers configuration, gate mode configuration.
- Write to  $GTPSR_n[PPS]$  fields in order to program the appropriate timer's clock primary prescaler.
- Write to  $GTMDR_n$  in order to choose an input clock, to program the secondary prescaler and to set a desirable appropriate timer's operational mode.

#### NOTE

Erratic behavior may occur if  $GTCFR_n$  and  $GTPSR$  are not initialized before the  $GTMDR$ . Only  $GTCFR_n[RST_n]$  can be modified at any time

- Clear  $GTEVR_n[REF]$  and  $GTEVR_n[CAP]$  by writing 1s in order to clear the previous events.
- Write to  $GTRFR$  and to  $GTCNR_n$  according to appropriate timer's  $GTMDR_n$  programming.

#### NOTE

A write cycle to a  $GTCNR_n[CNV]$  fields sets the register to the written value, causing its corresponding primary and secondary prescalers, ( $GTPSR_n[PPS]$  and  $GTMDR_n[SPS]$ ), to be reset.

- Write to  $GTCFR_n[STP_n]$  and to  $GTCFR_n[RST_n]$  in order to initialize the appropriate timer's operation.

## 5.8 Power Management Control (PMC)

The device provides a power management control (PMC) unit, which enables the device to smoothly enter and exit low power modes. Low power modes may be used when internal units in the device temporarily or permanently do not perform any action.

The device uses one or more of the following methods for power saving:

- Dynamic power management
- Shutting down unused blocks
- Software-controlled power-down states
- Low power state, which is reached by powering down the e300 and other non-essential blocks (DDR controller, eLBC). This mode utilizes an on-chip split power supply allowing power to be removed from a portion of the device.
- Support for the PCI Power Management Interface Specification in both host and agent modes. When the device is in either mode, the PMC is capable of placing the device into one of the supported low-power states and supporting the power management event (PME) signaling protocol.

## 5.8.1 External Signal Description

Table 5-65 describes the power management signals.

**Table 5-65. System Control Signals—Detailed Signal Descriptions**

Signal	I/O	Description
$\overline{\text{QUIESCE}}$	O	Quiesce state. Indicates that the processor system and PowerPC core are in low power state.
		<b>State Meaning</b> Asserted—The system and PowerPC core are in low power state. Negated—The system and PowerPC core are not in low power state.
		<b>Timing</b> The timing between a quiesce request from the PowerPC core and the assertion of the external indication or between negation of the core's quiesce request and negation of the external indication depends on the current state of the internal system units and may vary accordingly.
EXT_PWR_CTRL	O	External power control. Enables the external switch to provide power to the device.
		<b>State Meaning</b> Asserted—Power is supplied to the switchable power supply. Negated—Power is removed from the switchable power supply.
		<b>Timing</b> —
PMC_PWR_OK	I	Stable power. Indicates whether the power supply on the board is stable.
		<b>State Meaning</b> Asserted—The external power supply is stable to specifications. Negated—The external power supply is off or not stable to specifications.
		<b>Timing</b> —

## 5.8.2 PMC Memory Map/Register Definition


Table 5-66 shows the memory map for the power management controller registers.

**Table 5-66. Power Management Controller Registers Memory Map**

Offset	Register	Access	Reset	Section/Page
0x00B00	Power management controller configuration register (PMCCR)	R/W	0x0000_0000	5.8.2.1/5-67
0x00B04	Power management controller event register (PM CER)	R/W	0x0000_0000	5.8.2.2/5-68
0x00B08	Power management controller mask register (PM C MR)	R/W	0x0000_0000	5.8.2.3/5-70
0x00B0C	Power management controller configuration register 1 (PMCCR1)	R/W	0x0000_0000	5.8.2.4/5-70
0x00B10	Power management controller configuration register 2 (PMCCR2)	R/W	0x0002_0002	5.8.2.5/5-72
0x00B14–0x00BFC	Reserved	—	—	—

### 5.8.2.1 Power Management Controller Configuration Register (PMCCR)

The power management controller configuration register (PMCCR), shown in Figure 5-51, controls whether only the PowerPC core will enter low power state upon quiesce request or additional parts of the device will also enter low power state.



**Figure 5-51. Power Management Controller Configuration Register**


Table 5-5 defines the bit fields of PMCCR.

**Table 5-67. PMCCR Bit Settings**

Bits	Name	Description
0–29	—	Reserved. Write has no effect, read returns 0.
30	DLPEN	DDR SDRAM low power enable 0 The DDR SDRAM memory controller is prevented from entering low power state. 1 The DDR SDRAM memory controller will enter low power state when the rest of the system enters low power, according to SLPEN setting. DDR SDRAM will enter self-refresh mode (if enabled by DDR_SDRAM_CFG[SREN] memory controller register) and DDR clocks ( $MCK_n$ ) are shut off. This bit is cleared when the device exits from low power state. Note that setting this bit without setting SLPEN has no effect.
31	SLPEN	System low power enable 0 The system is prevented from entering low power state. 1 The system will enter low power state when a quiesce request from the PowerPC core arrives. This bit is cleared when the device exits from low power state.

## 5.8.2.2 Power Management Controller Event Register (PMCER)

The power management controller event register (PMCER), shown in [Figure 5-52](#), indicates with the PMCI bit that the power management controller has detected a wake-up event, that the system is not in idle state anymore, and that the device should exit low power state. If PMCMR[PMCIIE] is set, the PMC interrupt request to the PowerPC core is driven. When set, bits 23–30 indicate the sources of various wake-up events.



**Figure 5-52. Power Management Controller Event Register**

[Table 5-68](#) defines the bit fields of PMCER.

**Table 5-68. PMCER Bit Settings**

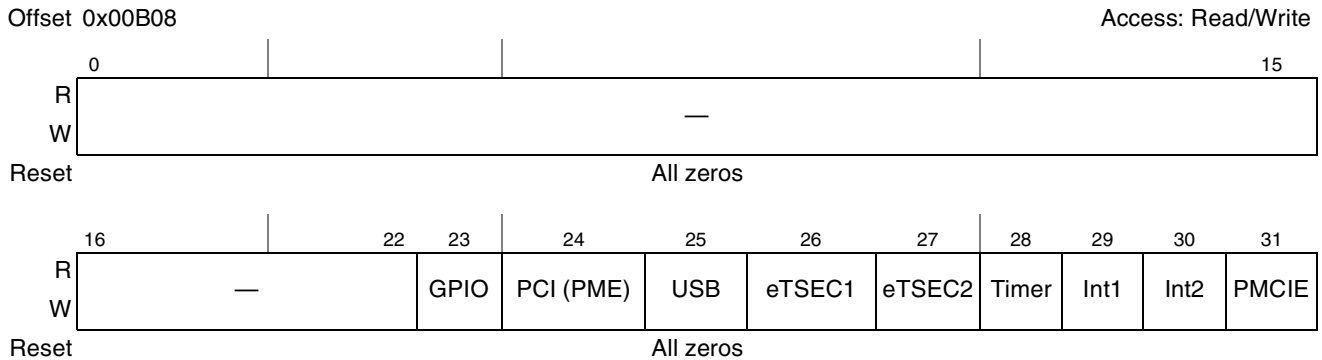
Bits	Name	Description
0–22	—	Reserved. Write has no effect, read returns 0.
23	GPIO	Wake-up event detected. 0 A wake-up event did not occur from this wake-up source. 1 A wake-up event occurred on GPIO. This wake-up event was caused by an unmasked event of GPIO module. See <a href="#">Chapter 21, “General Purpose I/O (GPIO),”</a> for more details. If the corresponding PMCMR bit is set, the PMC will assert interrupt request to the PowerPC core or external PME to the remote host, depending on the state of PMCCR1[PME_EN]. This bit can be cleared by writing a 1 to the bit location (writing zero has no effect). <b>Note:</b> This bit will not be affected by the wake-up event if the corresponding mask bit in PMCMR is cleared.
24	PCI (PME)	Wake-up event detected. 0 A wake-up event did not occur from this wake-up source. 1 A wake-up event occurred. This wake up event was caused by an active state of the $\overline{\text{PCI\_PME}}$ input signal. See <a href="#">Table 13-3, “PCI Interface Signals—Detailed Signal Descriptions,”</a> for more details. If the corresponding PMCMR bit is set, the PMC will assert interrupt request to the PowerPC core. This bit can be cleared by writing a 1 to the bit location (writing zero has no effect). <b>Note:</b> This bit will not be affected by the wake-up event if the corresponding mask bit in PMCMR is cleared.
25	USB	Wake-up event detected. 0 A wake-up event did not occur from this wake-up source. 1 A wake-up event occurred on USB. This wake-up event was caused by a detection of a non idle state on USB interface. See <a href="#">Chapter 16, “Universal Serial Bus Interface,”</a> for more details. If the corresponding PMCMR bit is set, the PMC will assert interrupt request to the PowerPC core or external PME to the remote host, depending on the state of PMCCR1[PME_EN]. This bit can be cleared by writing a 1 to the bit location (writing zero has no effect). <b>Note:</b> This bit will not be affected by the wake-up event if the corresponding mask bit in PMCMR is cleared.

Table 5-68. PMCER Bit Settings (continued)

Bits	Name	Description
26	eTSEC1	<p>Wake-up event detected.</p> <p>0 A wake-up event did not occur from this wake-up source.</p> <p>1 A wake-up event occurred on eTSEC1. This wake-up event was caused by a detection of a Magic Packet on the receive path of eTSEC1. See <a href="#">Chapter 15, “Enhanced Three-Speed Ethernet Controllers,”</a> for more details. If the corresponding PMCMR bit is set, the PMC will assert interrupt request to the PowerPC core or external PME to the remote host, depending on the state of PMCCR1[PME_EN]. This bit can be cleared by writing a 1 to the bit location (writing zero has no effect).</p> <p><b>Note:</b> This bit will not be affected by the wake-up event if the corresponding mask bit in PMCMR is cleared.</p>
27	eTSEC2	<p>Wake-up event detected.</p> <p>0 A wake-up event did not occur from this wake-up source.</p> <p>1 A wake-up event occurred on eTSEC2. This wake-up event was caused by a detection of a Magic Packet on the receive path of eTSEC2. See <a href="#">Chapter 15, “Enhanced Three-Speed Ethernet Controllers,”</a> for more details. If the corresponding PMCMR bit is set, the PMC will assert interrupt request to the PowerPC core or external PME to the remote host, depending on the state of PMCCR1[PME_EN]. This bit can be cleared by writing a 1 to the bit location (writing zero has no effect).</p> <p><b>Note:</b> This bit will not be affected by the wake-up event if the corresponding mask bit in PMCMR is cleared.</p>
28	TIMER	<p>Wake-up event detected.</p> <p>0 A wake-up event did not occur from this wake-up source.</p> <p>1 A wake-up event occurred on General Purpose Timer 1. This wake up event was caused by a detection of match between timer current value and timer reference value of the fourth 16 bit unit of GTM1. See <a href="#">Section 5.7, “General-Purpose Timers (GTMs),”</a> for more details. If the corresponding PMCMR bit is set, the PMC will assert interrupt request to the PowerPC core or external PME to the remote host, depending on the state of PMCCR1[PME_EN]. This bit can be cleared by writing a 1 to the bit location (writing zero has no effect).</p> <p><b>Note:</b> This bit will not be affected by the wake-up event if the corresponding mask bit in PMCMR is cleared.</p>
29	INT1	<p>Wake-up event detected.</p> <p>0 A wake-up event did not occur from this wake-up source.</p> <p>1 A wake-up event occurred on external interrupt request 1. This wake-up event was caused by a detection of an active state of the <math>\overline{IRQ1}</math> external pin. If the corresponding PMCMR bit is set, the PMC will assert interrupt request to the PowerPC core or external PME to the remote host, depending on the state of PMCCR1[PME_EN]. This bit can be cleared by writing a 1 to the bit location (writing zero has no effect).</p> <p><b>Note:</b> This bit will not be affected by the wake-up event if the corresponding mask bit in PMCMR is cleared.</p>
30	INT2	<p>Wake-up event detected.</p> <p>0 A wake-up event did not occur from this wake-up source.</p> <p>1 A wake-up event occurred on external interrupt request 2. This wake-up event was caused by a detection of an active state of the <math>\overline{IRQ2}</math> external pin. If the corresponding PMCMR bit is set, the PMC will assert interrupt request to the PowerPC core or external PME to the remote host, depending on the state of PMCCR1[PME_EN]. This bit can be cleared by writing a 1 to the bit location (writing zero has no effect).</p> <p><b>Note:</b> This bit will not be affected by the wake-up event if the corresponding mask bit in PMCMR is cleared.</p>
31	PMCI	<p>Power management controller interrupt.</p> <p>When set, indicates that one of the following events has occurred:</p> <ul style="list-style-type: none"> <li>• One of the unmasked wake-up events (bits 23–30) occurred and PMCCR1[PME_vPEN] is cleared, or</li> <li>• PM current state (as indicated in PMCCR1[CURR_STATE]) is different than PM next state (as written to PCIPMR1[Power_State] and indicated in PMCCR1[NEXT_STATE]) and PMCCR1[USE_STATE] is set, or</li> <li>• CSB platform is in low power mode and a new CSB bus request is detected</li> </ul> <p>If PMCMR[PMCI] is set, the PMC interrupt request to the PowerPC core is driven, causing the PowerPC core to exit its low power state. PMCI can be cleared by writing a 1 to it (writing zero has no effect).</p>

### 5.8.2.3 Power Management Controller Mask Register (PMCMR)

The power management controller mask register (PMCMR), shown in [Figure 5-53](#), controls through the PMCIE bit whether the PMC interrupt request to the PowerPC core is enabled. The PMC interrupt request causes the PowerPC core to exit its low power state before any transaction on the system bus occurs. Bits 23–30 are mask bits for the defined low power wake-up events.



**Figure 5-53. Power Management Controller Mask Register**

[Table 5-69](#) defines the bit fields of PMCMR.

**Table 5-69. PMCMR Bit Settings**

Bits	Name	Description
0–22	—	Reserved. Write has no effect, read returns 0.
23–30	GPIO, PCI(PME), USB, eTSEC1, eTSEC2, Timer, Int1, Int2	Wake-up event masking. 0 Mask wake-up events from Int2, Int1, Timer, eTSEC2, eTSEC1, USB, PCI, or GPIO, respectively. 1 Do not mask wake-up events
31	PMCIE	Power management controller interrupt enable. 0 PMC interrupt request (PMCI) is disabled. 1 PMC interrupt request (PMCI) is enabled.

#### NOTE

The user is also required to enable the PMC interrupt in the programmable interrupt controller by setting SIMSR\_L[PMC].

### 5.8.2.4 Power Management Controller Configuration Register 1 (PMCCR1)

The power management controller configuration register 1 (PMCCR1), shown in [Figure 5-54](#), controls the sequencing of the device into its low power state including PME (power management event) signaling, toggling of the external power switch, and indication of current and desired power states.

Offset 0x00B0C

Access: Read/Write


**Figure 5-54. Power Management Controller Configuration Register 1**

Table 5-70 defines the bit fields of PMCCR1.

**Table 5-70. PMCCR1 Bit Settings**


Bits	Name	Description
0	USE_STATE	Controls whether the next and current state values should be used. This typically depends on the device operating mode (PCI host or agent). 0 Ignore the next and current state values (host mode). 1 Use the next and current state values (agent mode). When next state does not equal current state, an interrupt will be sent to the e300 CPU.
1–22	—	Reserved, must be cleared.
23	LLPEN	SerDes low power enable bit. This bit dictates whether the SerDes is intended to be switched off or not when the chip goes into low power mode. 0 SerDes will be switched off when the chip goes into low power mode. 1 SerDes will not be switched off when the chip goes into low power mode. (This would be needed when wake-up through SGMII is intended.)
24	PME_EN	PME (power management event) signaling enable. Clearing this bit typically indicates the device is acting as PCI host, meaning that $\overline{\text{PCI\_PME}}$ is a wake-up input to device. Setting this bit typically indicates that the device is a PCI agent and will assert $\overline{\text{PCI\_PME}}$ on wake-up. PME_EN does not qualify the assertion of $\overline{\text{PCI\_PME}}$ when setting the PMCCR1[ASSERT_PME] bit to 1. 0 Mask PME signaling when wake-up events occur. 1 Allow PME signaling when wake-up events occur.
25	ASSERT_PME	Normally $\overline{\text{PCI\_PME}}$ output is asserted automatically by PMC when a defined wake-up event occurs assuming PMCCR1[PME_EN] = 1 and PCIPMCR1[PME_EN] = 1. A defined wake-up event refers to those events that are registered in bits 23–30 of PMCER. ASSERT_PME allows $\overline{\text{PCI\_PME}}$ to be asserted manually, by the e300. This would be done to inform the host of a power state change when PMC does not generate $\overline{\text{PCI\_PME}}$ automatically, for example waking from D1 due to CSB bus activity. ASSERT_PME is not qualified by PMCCR1[PME_EN]. ASSERT_PME is cleared when the $\overline{\text{PCI\_PME}}$ signal is asserted. 0 $\overline{\text{PCI\_PME}}$ will only be asserted automatically when a defined wake-up event occurs. 1 Assert the $\overline{\text{PCI\_PME}}$ signal under software control (manually).
26	POWER_OFF	Distinguishes between D3Hot and D3Warm. If this bit is set, EXT_PWR_CTRL will be toggled in D3 causing VDD to be switched off (if implemented externally). 0 Do not use D3warm state. Always assert the EXT_PWR_CTRL output. 1 On transitions to D3 state, negate the EXT_PWR_CTRL signal to switch external power low (VDD = 0). On wake-up assert EXT_PWR_CTRL to switch VDD power on.
27	—	Reserved

**Table 5-70. PMCCR1 Bit Settings (continued)**

Bits	Name	Description
28–29	NEXT_STATE	Indicate the power state as programmed by the PCI host in PCIPMR1[Power_State]. The host will write the Power_State bits to request that the device enter a certain low power state. The host may also write the Power_State to request that the device return to D0 from some low power state. When the NEXT_STATE field is different than the CURR_STATE field, an interrupt is asserted to the e300 processor through the IPIC. This field is read-only. 00 Host's desired power state is D0 01 Host's desired power state is D1 10 Host's desired power state is D2 11 Host's desired power state is D3 (either D3Hot or D3Warm)
30–31	CURR_STATE	Indicate the current power state of the device. These bits are written by the e300 just before entering a requested low power state, or when the device has returned to the full on state (D0). Writing these bits causes the PCIPMR1[Power_State] field to be updated informing the host that the device has entered the requested power state. 00 Current power state is D0 01 Current power state is D1 10 Current power state is D2 11 Current power state is D3Hot (also used for D3Warm)

### 5.8.2.5 Power Management Controller Configuration Register 2 (PMCCR2)

The power management controller configuration register 2 (PMCCR2), shown in [Figure 5-55](#), contains count values used for power-up and power-down timers.



**Figure 5-55. Power Management Controller Configuration Register 2**



Table 5-71 defines the bit fields of PMCCR2.

**Table 5-71. PMCCR2 Bit Settings**

Bits	Name	Description
0–3	—	Reserved, must be cleared.
4–15	RCNT	<p>Reset count value. When waking up from D3Warm, power (VDD) is reapplied to a portion of the die. This value determines the duration of the reset signal applied to this logic when power is re-applied. If PMCCR1[POWER_OFF] = 0, this field has no effect.</p> <p>Reset is applied to the powered-off region upon entering D3Warm. The RCNT value is copied into a decremter that counts down once every 32,000 CSB clock cycles. When a wake-up event occurs and the PMCCR1[NEXT_STATE] is set to 00 (D0), PMC will wait for the PMC_PWR_OK signal to be asserted then begin decrementing the reset counter. When the counter reaches 0 reset is removed. In some systems the PMC_PWR_OK signal will not be provided externally and will be tied active. In this case the counter needs to include time for the VDD power supply to become stable.</p> <p>Software needs to set the RCNT value based on the PMC clock frequency, the amount of time required for the for the VDD power supply to ramp (if PMC_PWR_OK is not used), and the amount of time required for the e300 PLL to lock.</p> <p><b>WARNING:</b> If the value placed in this register is too small, the reset may not assert long enough to allow the chip to function properly. The default value is larger than the time it takes for the e300 PLLs to re-lock.</p>
16–19	—	Reserved, must be cleared.
20–31	PDCNT	<p>Power-down count value. This counter establishes a minimum time for which power can be removed to the VDD supply during D3Warm. When the device enters D3Warm, the EXT_PWR_CTRL signal is negated. At this point this counter is loaded with the PDCNT value and begins to decrement, once every 32,000 CSB clock cycles. PMC will not respond to a wake-up request and assert EXT_PWR_CTRL until this counter has expired. The count value is reloaded each time the VDD power is removed. If PMCCR1[POWER_OFF] = 0 this field has no effect.</p> <p>Software needs to set this register based on the PMC clock frequency and the requirements of the power supply.</p> <p><b>WARNING:</b> If the value placed in this register is too small, the power supply may cycle too quickly and the chip may not function properly.</p>

### 5.8.3 Functional Description

The device has features to minimize power consumption at several levels. Dynamic power management locally minimizes power consumption when a block is idle. Software can also shut down clocks to individual blocks when they are not needed through a memory-mapped register in the clock unit (SCCR). Additionally, software running on the PowerPC core can access the core's SPRs to put the device into doze, nap, or sleep power down states. Finally, software can access the PMCCR register to enable the device to go to low power state whenever the PowerPC core enters nap or sleep states. The device supports a low power mode where power is removed from a portion of the die. The PMC supports features that work in concert with the PCI power management (PM) block (PME context) to provide support for PCI power management capabilities such as asserting or responding to power management events (PMEs). These power management features are described in further detail in this section.

#### 5.8.3.1 Dynamic Power Management

Many blocks in the device can dynamically turn off clocks within the block when sections of the block are idle. This feature is always enabled and occurs automatically.

### 5.8.3.2 Shutting Down Unused Blocks

As described in [Section 4.5.2.3, “System Clock Control Register \(SCCR\),”](#) SCCR provides a way to shut down certain functional blocks within the device when they are not needed in a particular system. SCCR can be written by the PowerPC core or by an external master. Powering down a block in this way turns off all clocks to that block. It does not remove power. It is required that the SCCR is written to shut down a certain functional block only when that block is idle.

#### NOTE

Functional blocks disabled using SCCR cannot respond to configuration accesses. Any access to configuration, control, and status registers of a disabled block is a programming error.

### 5.8.3.3 Software-Controlled Power-Down States

PowerPC software can place the core in doze, nap, or sleep power-down states by writing to HID0 in the core, as described in detail in the section “Hardware Implementation Register 0 (HID0),” of the *e300 PowerPC Core Reference Manual*. In addition, if PMCCR[SLPEN] is set when the PowerPC core request to enter nap or sleep modes, it will also cause the system internal logic units to enter low power mode.

### 5.8.3.4 Software-Controlled Power Supply Switching

The device has additional low power features that allow power to be removed to a portion of the die, allowing significant additional power savings. This mode is referred to as D3Warm (described below). [Figure 5-56](#) illustrates the power segmentation provided on the device. Sequencing in and out of D3Warm will be described in the following sections.

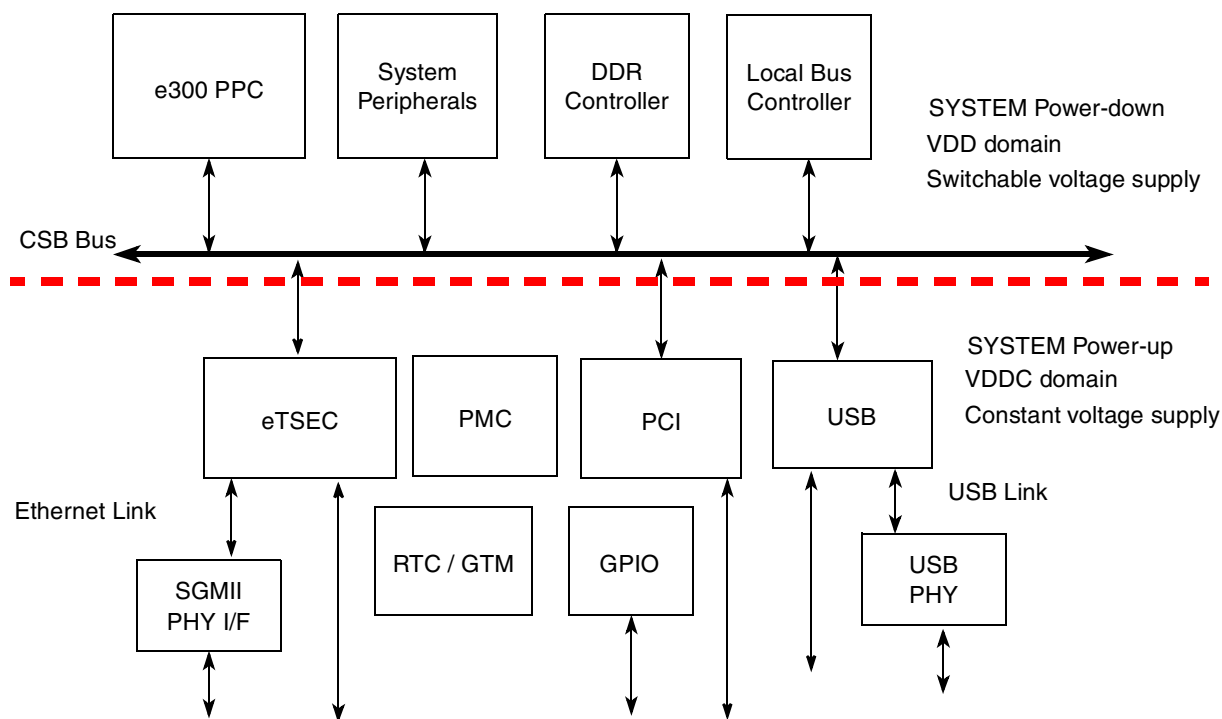


Figure 5-56. Power Segmentation in Deep Sleep Mode

### 5.8.3.5 Support of PCI Power Management Interface Specification

The device will support the PCI power states D0, D1, D2, D3Hot, and D3Cold as defined in the Rev. 1.2 of the PCI Power Management Interface Specification. Table 5-72 defines these D<sub>x</sub> states. The PCI power management specification defines a power management event, or PME, as the process by which a PCI agent signals a request to the host for a change in its power consumption state. When in agent mode, the device has the capability to generate PME signaling through the assertion of an external  $\overline{\text{PCI\_PME}}$  signal. As host, the device is able to respond to PME signaling as a wake-up event.

It is assumed that in D3Cold all power will be removed from the device, so  $\overline{\text{PCI\_PME}}$  signaling is not supported from D3Cold. Instead, an MPC8313E-specific D3Warm state is defined. The difference between D3Hot and D3Warm is that in D3Hot the device's entire core region is supplied with the nominal 1-V VDD supply. In D3Warm, a portion of the core region can be powered off. This partial power-down mode allows the device to achieve significant reduction in power dissipation while still maintaining the capability to respond to wake-up events.

Table 5-72 defines the PCI power states.

Table 5-72. PCI Defined Power Management State Support

PCID <sub>x</sub> Power State	Supported	How Supported
D0	Yes	Default state; full power; PME signaling supported through software control
D1	Yes	e300 in Doze mode; e300 PLL running; PME signaling supported
D2	Yes	e300 in Nap state; e300 PLL running; PME signaling supported

**Table 5-72. PCI Defined Power Management State Support (continued)**

PCI D <sub>x</sub> Power State	Supported	How Supported
D3Hot	Yes	e300 in Sleep mode. In D3Hot VDD is still applied to the entire device. In this mode you have an option to power-off a portion of the device VDD. This is referred to as D3Warm. In D3Warm, power is removed to the e300, DDR, LBC and IPIC. Power is not removed to Ethernet, USB, PCI, GPIO and timers. PME signaling is supported in either D3Hot or D3Warm.
D3Cold	Yes*	*PME signaling is not supported in D3Cold. For PME signaling, use the D3Warm state, which is similar to PCI D3Hot state but with a portion of the chip powered off. According to PCI PM Specification 1.2, in D3Cold the PCI clock is removed and all devices will be reset when power is restored. The device can be placed into D3Cold state, but wake-up events cannot be recognized or generated and when power is restored, the device must go through a normal power-on reset boot sequence as it needs to be re-initialized.

For completeness, [Table 5-73](#) shows the PCI bus power management states (B0–B3) that the device supports. If the device is used as a PCI agent, the supported bus states are B0 and B1, since the PCI clock is running in these two states. Otherwise, as a host, the device could be configured to support all the bus states.

**Table 5-73. PCI Bus Power Management State Support**

PCI B <sub>x</sub> Bus State	PCI Bus VDD	PCI Bus Clock	PCI Bus Activity	Support in PCI Agent Mode	Support in PCI Host Mode
B0 (full on)	On	Yes	Any PCI transaction, function interrupt, or PME event	Yes	Yes
B1	On	Yes	PME event; bus is idle	Yes	Yes <sup>1</sup>
B2	On	No	PME event	No <sup>2</sup> (no bus clock)	Yes <sup>3</sup>
B3	Off	No	PME event	No <sup>4</sup> (no bus clock or VDD)	Yes <sup>5</sup>

<sup>1</sup> Requires the device to hold the PCI bus in idle state

<sup>2</sup> The device cannot process wake-up events with the PCI bus in this state and the bus must be returned to B0 state through a **PORESET**.

<sup>3</sup> Requires the device to turn off PCI bus clock

<sup>4</sup> The device cannot process wake-up events with the PCI bus in this state and the bus must be returned to B0 state through a **PORESET**.

<sup>5</sup> Requires the device to turn off the PCI bus clock through the output clock control register (OCCR), and to turn off the PCI bus VDD through some customer-defined method (perhaps GPIO).

Below is a summary of the new *PCI Power Management Interface Specification* capabilities supported on the device:

- Generation of  $\overline{\text{PCI\_PME}}$  signal to an external PCI host when the device is operating in agent mode.
- Responding to  $\overline{\text{PCI\_PME}}$  as an input wake-up event when the device is operating as PCI host.
- Responding to other wake-up events from various sources: Ethernet Magic Packet, USB, GPIO, internal timer, external interrupt, PCI PME# ( $\overline{\text{PCI\_PME}}$ ).
- Properly sequencing the device into and out its lowest power mode where VDD is removed to a portion of the die.

- Generating or responding to appropriate control signals relative to low power modes: VDD switching control (EXT\_PWR\_CTRL) output signal and input signal indicating external power is stable (PMC\_PWR\_OK)

The relationship between the e300 core and the defined PMC power states is illustrated in [Table 5-74](#).

### NOTE

The relationships shown below between e300 core doze, nap, and sleep modes and PCI D-states are only suggested. There is no forced hardware relationship between these states. For example, if a PCI host sets PCIPMR1[Power\_State] = 01, it is still up to the software running on the e300 to put itself into doze mode.

**Table 5-74. Software-Controller Power-Down States—Basic Description**

System Mode	Core Mode	Suggested PCI D- state	Description	Core Responds to		DDR SDRAM strl state	Quiesce Signal State
				Snoop	Interrupt		
Full On (PMCCR[SLPEN]=0)	Full On	D0 PCIPMR1[PowerState]=00	All units operating normally	Yes	Yes	Active	Negated
	Doze	D1 PCIPMR1[PowerState]=01	Core stops dispatching new instructions (core is halted), and most of the core functional units are disabled. System operates normally.	Yes	Yes	Active	Negated
	Nap	D2 PCIPMR1[PowerState]=10	Core is stopped with its clocks off except to time base. System operates normally.	No	Yes	Active	Negated
	Sleep	D2 PCIPMR1[PowerState]=10	Core is stopped with its clocks off. Core clocks powered down to all blocks (including core time base) except to the interrupt unit. System operates normally.	No	Yes	Active	Negated

**Table 5-74. Software-Controller Power-Down States—Basic Description (continued)**

System Mode	Core Mode	Suggested PCI D- state	Description	Core Responds to		DDR SDRAM strl state	Quiesce Signal State
				Snoop	Interrupt		
Low Power (PMCCR[SLPEN]=1)	Nap	D3 PCIPMR1[PowerState]=11	Core operation as described above. System is in idle state, DDR SDRAM memory operates in self-refresh mode if enabled.	No	Yes	According to PMCCR[DLPEN]	Asserted
	Sleep	D3 PCIPMR1[PowerState]=11	Core operation as described above. System is in idle state, DDR SDRAM memory operates in self-refresh mode if enabled.	No	Yes	According to PMCCR[DLPEN]	Asserted
Lowest Power (PMCCR[SLPEN]=1), PMCCR1[POWER_OFF]=1	Power Off, Deep Sleep	D3Warm PCIPMR1[PowerState]=11	This state is an extension of the above low power modes where power can be removed to a portion of the device die using an external power switch. Software enters core "Sleep" mode with PMCCR1[POWER_OFF] = 1. Wake-up is in response to defined Wake-up events.	No	No	PMCCR[DLPEN] = 1	Asserted

### 5.8.3.5.1 Entering Low Power States—Core-Only Mode

Entering Doze mode is controlled only by the e300 PowerPC core itself, and does not involve the power management controller or other blocks. For a more detailed description, see [Table 7-1](#).

Entering Nap or Sleep modes occurs by writing to HID0 in the core, causing the core to make a quiesce request to the power management controller while PMCCR[SLPEN] is cleared. The core is immediately enabled to enter low power state, regardless of the system status. Note that since the core does not snoop the bus in this mode, it is the user's responsibility to keep the cache coherent. Other device peripheral and internal units continue to operate in full-on mode while the core is in low power state in this mode.

### 5.8.3.5.2 Entering Low Power States—Core and System Mode

Core and system mode is achieved when the core makes a quiesce request to the power management controller after PMCCR[SLPEN] is set. To preserve cache coherency and otherwise avoid loss of system state, the core's transition to low-power modes is coordinated with other functional blocks. The power management controller allows the core to enter power down mode only when the rest of the system is idle.

When the power management controller detects that the internal system bus is idle, and there are no outstanding transactions, it signals the internal logic units to enter low power state.

If PMCCR[DLPEN] is set, the DDR SDRAM is first set to self-refresh mode (if enabled by DDR\_SDRAM\_CFG[SREN] memory controller register) before the memory controller stops driving refresh commands. Self-refresh mode guarantees that the memory content will remain valid while the memory controller and its clocks are off. The DDR clocks are then disabled. Finally the DDR SDRAM memory controller enters low power state and acknowledges the power management controller.

The power management controller then signals the core and acknowledges its request to enter power down mode. Finally the  $\overline{\text{QUIESCE}}$  output signal is asserted.

### 5.8.3.6 Exiting Core and System Low Power States

The device can exit low power state and return to full-on mode for one of the following reasons:

- The core internal time base unit invokes a request to exit low power state.
- The core has received an interrupt request.
- The device is a PCI host, and the power management controller has detected that the system is not idle and there are outstanding requests for transactions on the internal bus.
- The device is a PCI agent, and the power management controller has detected that the PCI power next state does not equal the PCI power current state (meaning that the remote PCI host has requested a change from non D0 to D0 state).

The actions taken to exit low power state depend on the mode and whether the system or only the core are in this state. The following sections describe the various scenarios.

#### 5.8.3.6.1 Exiting Low Power States—Core-Only Mode

Exit from Doze mode is controlled only by the core itself and does not involve the power management controller or other blocks in the device. For a detailed description, see [Table 7-1](#).

Nap or Sleep modes are exited when the core has received an interrupt request, or according to the internal time base unit of the core (Nap mode only). The source of the interrupt can be an internal block or external signal. When the core returns to full-on state, it signals to the power management controller that it is ready and is immediately acknowledged to access the rest of the system.

#### 5.8.3.6.2 Exiting Low Power States—Core and System Mode

The power management controller decides to exit low power state when it detects that the system is not idle anymore. The device may exit idle state when one of the peripheral interfaces makes a request to access the internal bus or when the core returns to full-on state, as described above, and makes a request to access the internal bus. For example, the TSEC receives an Ethernet frame, and requires to store it on the DDR SDRAM memory.

If the DDR SDRAM memory controller is in low power state (PMCCR[DxLPEN] was set when entering low power state), the power management controller initially enables the DDR SDRAM memory controller. DDR SDRAM clocks (MCK $n$ ) are enabled and the memory controller exit self-refresh and returns to auto-refresh mode.

The power management controller then enables other internal units and interrupts the PowerPC core. When all internal units, including the core, are ready, the power management controller enables the device to return to full-on state, negate the  $\overline{\text{QUIESCE}}$  output, and clear PMCCR[SLPEN]. Outstanding requests for transactions are now granted to execute on the internal bus.

#### NOTE

Software is required to enable PMCI interrupt by setting PMCMR[PMCIIE], otherwise exiting from low power state is not possible.

#### NOTE

It is the software's responsibility to clear PMCER[PMCI].

### 5.8.3.7 MPC8313E-Specific PMC Low Power States

This section will describe the MPC8313E D0–D3 low power states, including D3Warm, from a PCI power management perspective. The sequence of entering and exiting D0–D3 states will also be described.

The power management implementation assumes the following:

- The core region of the device will have a segmented power supply: a constant supply segment VDDC, and switchable supply segment VDD.
- When the device is used as a PCI agent, the PCI clock (PCI\_CLK / PCI\_SYNC\_IN) must remain available, as this is the source of the device’s system clock. This applies to D3Warm as well. The PCI interface will reflect this requirement in the PCI’s PMC[PME\_Clock] register bit.
- When the device is used as a PCI host, the PCI\_CLK must also remain available.
- An external host is not required to reset the device through the PCI interface when transitioning from low-power states D1–D3Warm to the fully-powered state. Reset will need to be applied if the device is put into D3Cold state (all power removed to the chip).

#### 5.8.3.7.1 Power State Transitions from an ACPI Perspective

The ACPI 3.0 specification defines the power management interface between host and agent from a board-level perspective. The device is intended to be usable as an agent in an ACPI-compliant system through the PCI interface. For the agent mode, the ACPI specification defines the power-down/wake-up sequence to be, generally:

1. The host determines the capabilities of the agent. The agent may or may not have the capability to inform the host of a wake-up event from a low-power state. This capability is determined through operating system software that identifies all PCI function power capabilities by traversing the capabilities structure in the PCI configuration space and reading the power management capabilities register.
2. The host enables the agent to generate wake-up signaling to the host.
3. The host signals the agent to transition into a low power state that it deems appropriate.
4. When the agent receives a wake-up event, it signals the host by asserting  $\overline{\text{PCI\_PME}}$ . The agent then waits for the host to request that it transition into the active power state.
5. The host requests the agent to transition to full power state by writing the new power state into the PCI configuration space.

[Table 5-75](#) defines the support for wake-up when the device is used as an agent. A “defined wake-up event” refers to wake-up events that are recognized by the PMC controller as agent: USB, internal timer, external interrupt, GPIO, or eTSEC. The events must be unmasked at the PMC to cause a wake-up.

Note that the device will also begin the wake-up process if the host sets PCIPMCR1[Power\_State] = 00. In this case, the PMC will assert an interrupt to the e300 core, but this action will not set any of the wake-up events in the PMC event register (PMCER).



Table 5-75. MPC8313E Agent Mode Wake-Up Support

Case	Power State	Remote Wake-up (PME) Signaling Enabled?	Wake-up Event Source	Action
1	D0	Not available	Not applicable	Active state, normal system activity
2	D1	Yes	Defined wake-up event: USB, internal timer, external interrupt, TSEC, GPIO	<p>PME signaling to the host from the D1 state is not likely to be used in an application. When a defined wake-up event occurs PMC will assert an interrupt to the e300 (if not masked) which will cause it to wake-up directly.</p> <p>To support PME signaling the following sequence should be followed: When a defined wake-up event occurs in D1 and PME signaling is enabled (PMCCR1[PME_EN] = 1), PCI_PME will be asserted to the host. Since the e300 has not been powered off, it will return to D0 (full-on) mode directly. To support PME signaling the e300 should not continue normal processing until the host has instructed the device to return to D0. The host would do this by writing PCIPMCR1[Power_State] = 00b in the PCI config space. This will cause the PMCCR1[NEXT_STATE] field to become 00b (D0) causing an interrupt to the e300. The e300 can then respond to the host that it had returned to D0 by writing PMCCR1[CURRENT_STATE] = 00, which will then be reflected in PCIPMCR1[POWER_STATE], indicating to the host that it has returned to D0. Normal processing can then continue.</p>
3	D2	Yes	USB, internal timer, external interrupt, TSEC, GPIO	Same as D1 state (Case 2), except e300 transitions to Full On mode from Nap state.
4	D3Hot	Yes	USB, internal timer, external interrupt, TSEC, GPIO	Same as D1 state (Case 2), except e300 transitions to Full On mode from Sleep mode.
5	D3Warm	Yes	USB, internal timer, external interrupt, TSEC, GPIO	In D3Warm the e300 is powered off. When the wake-up event occurs e300 does not return to D0 directly. Rather, PCI_PME is asserted to the host. The host will then instruct the device to return to D0 (writing PCIPMCR1[POWER_STATE] = 00). This will set PMCCR1[NEXT_STATE] = 00 and the device will continue to wake-up. Once awake the e300 should signal to the host that it has returned to D0 by writing PMCCR1[CURRENT_STATE] = 00b which will set PCIPMCR1[POWER_STATE] = 00. When the wake-up event occurs the device wakes to a partially uninitialized state (DDRC, IPIC etc. will need to be initialized). NOTE: PCI_PME signaled to the host occurs from D3Warm without powering up the device through the external power switch.
6	D1	Yes	Other (e.g., e300 decremter timer, e300 snoop hit, I <sup>2</sup> C)	<p>PME signaling to the host from the D1 state is not likely to be used in an application.</p> <p>e300 is in Doze mode. If an e300 interrupt occurs, the e300 will transition to Full On mode directly. PME signalling to the host can be supported if desired (see Case 2).</p>
7	D2	Yes	Other (e.g., e300 decremter timer, I <sup>2</sup> C)	Same as D1 state (Case 6), except e300 transitions to Full On mode from Nap state.

Table 5-75. MPC8313E Agent Mode Wake-Up Support (continued)


Case	Power State	Remote Wake-up (PME) Signaling Enabled?	Wake-up Event Source	Action
8	D3Hot	Yes	Other (e.g., I <sup>2</sup> C)	Same as D1 (Case 6), with the following differences: - e300 is in Sleep mode; - the e300 may not be awakened by a decremter timer interrupt.
9	D3Warm	Yes	Other	Not applicable, because the e300, IPIC, and other non-essential blocks are powered off. Only defined wake-up events can wake the device from D3Warm (see Case 5).
10	D1	No	USB, internal timer, external interrupt, TSEC, GPIO	When a defined wake-up event occurs PMC will assert an interrupt to the e300 (if not masked) which will cause it to wake-up directly (return to D0). Since PME signaling is disabled (PMCCR1[PME_EN] = 0) <u>PCI_PME</u> will not be asserted to the host. In this case the host will not be aware of the device power state change. To support PME signaling, see case 2.
11	D2	No	USB, internal timer, external interrupt, TSEC, GPIO	Same as Case 10, except e300 transitions to D0 from nap mode. The state change is not reported to the host.
12	D3Hot	No	USB, internal timer, external interrupt, TSEC, GPIO	Same as Case 10, except e300 transitions to D0 from sleep mode. The state change is not reported to the host.
13	D3Warm	No	USB, internal timer, external interrupt, TSEC, GPIO	Same as Case 10, except e300 transitions to D0 from sleep mode. The state change is not reported to the host.
14	D1	No	Other (e.g., e300 decremter timer, e300 snoop hit, I <sup>2</sup> C)	Same as Case 10. When PME signaling is disabled wake-up from Other sources is the same as wake-up from defined wake-up events.
15	D2	No	Other (e.g., e300 decremter timer, I <sup>2</sup> C)	Same as Case 10. When PME signaling is disabled wake-up from Other sources is the same as wake-up from defined wake-up events.
16	D3Hot	No	Other (e.g., I <sup>2</sup> C)	Same as Case 10. When PME signaling is disabled wake-up from Other sources is the same as wake-up from defined wake-up events.
17	D3Warm	No	Other	Not applicable, because the e300, IPIC, and other non-essential blocks are powered off.
18	D3Cold	Not Available	Not applicable	It is assumed that in D3Cold VDDC and VDD power supplies are powered off by the host and PORESET is applied. Wake-up events or PME signaling is not supported in this case. Returning to D0 implies a POR reset and a complete initialization of the device.

Table 5-76 defines the support for wake-up events when the device is operating as a host. A “defined wake-up event” refers to an interrupt that is recognized by the PMC in host mode: USB, GPIO, eTSEC, internal timer, external interrupt or PCI (PME input). Since the device is the host, PME signaling refers to external agents asserting the PCI\_PME input to the device, which is one of the defined wake-up events.

Table 5-76. MPC8313E Host Mode Wake-Up Support

Case	Power State	Wake-up Event Source	Action
1	D0	Not applicable	Active state, normal system activity
2	D1	Defined wake-up event: USB, internal timer, external interrupt, TSEC, GPIO, PCI (PCI_PME)	The wake-up event causes PMC to interrupt the e300 through the IPIC. The e300 transitions from Doze mode to Full On mode. If the wake-up source was the PCI_PME input the e300 OS may direct a power state change to the remote agent.
3	D2	USB, internal timer, external interrupt, TSEC, GPIO, PCI (PCI_PME)	Same as Case 2, except e300 transitions to Full On mode from Nap mode.
4	D3Hot	USB, internal timer, external interrupt, TSEC, GPIO, PCI (PCI_PME)	Same as Case 2, except e300 transitions to Full On mode from Sleep mode
5	D3Warm	USB, internal timer, external interrupt, TSEC, GPIO, PCI (PCI_PME)	Same as Case 2, except the e300 is powered off. External power will be re-applied and the e300 will be required to initialize a portion of the device (DDR, IPIC etc.) before the transition to D0 is complete.
6	D1	Other (e.g., e300 decremter timer, e300 snoop hit, I <sup>2</sup> C)	e300 is in Doze mode; when an e300 interrupt occurs, the e300 transitions to Full On mode.
7	D2	Other (e.g., e300 decremter timer, I <sup>2</sup> C)	e300 is in Nap mode; when an e300 interrupt occurs, the e300 transitions to Full On mode.
8	D3Hot	Other (e.g., I <sup>2</sup> C)	e300 is in Sleep mode; when an e300 interrupt occurs, the e300 transitions to Full On mode.
9	D3Warm	Other	Not applicable; e300, IPIC, and other non-essential blocks are powered off.
10	D3Cold	Not available	It is assumed that in D3Cold VDDC and VDD power supplies are powered off. Wake-up events are not recognized by the device in this case. Returning to D0 implies a POR reset and a complete initialization of the device.

The Dx power state transitions are supported in the PMC as indicated by the following state machine, which includes both hardware and software implications. Note that when the device is agent and in one of its low power modes, the host may decide to wake it up independent of a wake-up event or interrupt by writing the D0 state into the PCI configuration registers (PCIPMR1[Power\_State] = 00. This action will cause an interrupt to the e300 (D1–D3Hot) or cause the PMC to begin the wake-up process (D3Warm).

**NOTES:**

- 1) Wake-up case #1: Wake-up will occur as a result of one of the following: e300 interrupt or bus activity,  $PMCR[X]$  &  $PMCMR[X] = 1$ , where X is one of {USB, TSEC, GPIO, PCI, interrupt, timer},  $PMCCR1[NEXT\_STATE] = 00$ . The device will wake-up directly as host or agent.  $\overline{PCI\_PME}$  can be generated manually to host if desired.
- 2) Wake-up case #2: As host ( $PMCCR[PME\_EN] = 0$ ), Wake-up will occur as a result of one of the following:  $PMCR[X]$  &  $PMCMR[X] = 1$ , where X is one of {USB, TSEC, GPIO, PCI, interrupt, timer}. As agent ( $PMCCR[PME\_EN]=1$ ) wake-up event will cause  $\overline{PCI\_PME}$  to be asserted; the device will wake up when the host sets  $PCIPMR1[Power\_State] = 00$ , which causes  $PMCCR1[NEXT\_STATE] = 00$ .
- 3) Before the transition to D3Warm, VDD power is switched off. VDDC remains constant.
- 4) Before the transition from D3Warm, VDD power is switched on.
- 5) D1, D2 & D3Hot modes are supported through use of the e300 Doze, Nap, and Sleep modes respectively. Transitions from D1–D3Hot are triggered as follows: the e300 internal time base unit invokes a request to exit low power state or e300 receives an interrupt request, including interrupt from a defined wake-up event.
- 6) Transitions to D3Cold implies power-down of VDDC and VDD supplies.
- 7) Transitions from D3Warm go to the D0-non-initialized state, meaning a portion of the chip (e300, DDR etc.) will need to be initialized. Portions of the chip which are not powered-off may not need to be initialized.

**Figure 5-57. Power State Transitions Supported**

### 5.8.3.7.2 MPC8313E Low Power Sequencing

The following sections describe the sequence of events that occur when entering or exiting the device's lowest power state (D3Warm). Host and agent cases are described separately.

#### PMC Power-Down When the MPC8313E is a PCI Agent

In this case the PCI device driver runs on an external PCI host. It interfaces to the device through PCI configuration interface.

The e300 device driver runs locally on the e300 processor and interfaces to PMC internally

At initial power-up (POR), external PCI host software is required to determine the power management capabilities of all agents and initialize their PME context, including the PME support bits. To support PME signaling the PCI device driver must set the PCI function's PCIPMR1[PME\_En] register bit and the PCIPMR1[PME\_Status] register bit (clearing the PME\_Status bit). The e300 device driver must also set the PMC's PMCCR1[PME\_EN].

Steps to power down the device to the lowest power state, D3Warm, are as follows.

1. Configure IPIC registers to mask interrupts to core.
2. Disable MSR[EE] to disable external interrupts. The software must complete servicing any pending interrupts before doing so.
3. An e300 driver routine programs PMC to allow wake-up on one of the PMC wake-up events (eTSEC magic packet, USB, GPIO internal timer, external interrupt) by writing a "1" in the appropriate PMCMR[] mask register bit.
4. The PCI device driver executes code to save any MPC8313E context that would not otherwise survive the transition to the new power state (any MPC8313E context beyond what e300 software would configure on reset).
5. PCI device driver enables the PCI function to generate  $\overline{\text{PCI\_PME}}$ , then programs the D3Hot state into the PCI function's PCIPMR1[Power\_State] field.

Note: In order for the device to assert  $\overline{\text{PCI\_PME}}$ , both PCIPMR1[PME\_EN] and PMCCR1[PME\_EN] must be set.

6. This PCI Configuration PowerState register setting is detected and also reflected into the PMCCR1[NEXT\_STATE] register. This change generates an interrupt to the e300 through the IPIC.

Note: PCI PM 1.2 specification requires PCI's PCIPMR1[Power\_State] field to reflect the current state until the PCI agent is in the next state. Therefore the PCIPMR1[Power\_State] field will still read as 00b until the e300 updates the power state (see Step 11)

Note: The PMC is designed such that any time the PMCCR1[NEXT\_STATE] field is different than the PMCCR1[CURRENT\_STATE] field and interrupt will be asserted to the e300 core via IPIC.

7. An e300 interrupt routine detects PMC's PMCCR1[NEXT\_STATE] notification and begins the process of power down. It stops all of the masters on the CSB bus, including:
  - Security engine (by not providing buffer descriptors to work on)
  - eTSECs (gracefully stop through DMACTRL[GRS] and DMACTRL[GTS])
  - USB (placing USB into suspend mode)
 Then the e300 enables any desired wake-up sources: eTSEC to wake on Magic Packet™ reception by setting eTSEC MACCFG2[MPEN], USB link activity, GPIO transition, internal timer expires, assertion of external interrupt.

Note: that the PMC does not automatically sequence the device into low power mode when the next\_state bits are set to D3. This next\_state change only generates an interrupt to the e300. the e300 will then sequence the device into the requested low power mode.

8. The e300 sets existing PMC registers that allow power-down when the e300 enters nap or Sleep mode (PMCCR[DLPEN] and PMCCR[SLPEN]). Settings are based on the PMCCR1[NEXT\_STATE] field.
9. The e300 stores context in SDRAM through MEMC, and sets MEMC up to allow SDRAM to self-refresh during Sleep mode (DDR\_SDRAM\_CFG[SREN] = 1b).
10. The e300 writes the PMCCR1[NEXT\_STATE] value into PMCCR1[Curr\_State], which is reported in the PCIPMR1[Power\_State]. e300 then clears the PMCIE event in the PMCER register. It is important to note that the clearing of the PMCIE event should be done after changing the current state to match the next state value. Only then PMCIE event can be cleared. e300 then sets the PMCCR1[POWER\_OFF] bit to 1 indicating that the external EXT\_PWR\_CTRL signal will be toggled at the appropriate time to switch off external power.
11. The e300 asserts core\_qreq\_b, which is detected by the PMC. The PMC asserts STOP to the CSB arbiter, and waits for STOP\_ACK. When the PMC receives STOP\_ACK from the CSB arbiter, the PMC asserts STOP to the DDR memory controller and DDR xlb2mg gasket and waits for their STOP\_ACKs. When the DDR controller's STOP\_ACK is received, the PMC asserts an existing pmc\_stop\_ddr\_clk signal to stop the DDR clock.
12. Before shutting off power to a portion of the die, PMC will turn on conditioning logic to isolate the powered-on region from the powered-off region.
13. The PMC asserts the core\_qack\_b signal to the e300 indicating it should enter Sleep mode. The external  $\overline{\text{QUIESCE}}$  signal asserts. If PMCCR1[POWER\_OFF] is set, EXT\_PWR\_CTRL signal is negated to disable an external power switch shutting off VDD to a portion of the die.

### PMC Power-Down When the MPC8313E is a PCI Host

In this case the PCI device driver runs on the device e300 core and directly controls the PCI interface. The e300 device driver also runs on e300 and controls the PMC module internally.

At initial power-up (POR) the host software running on the e300 is required to determine the power management capabilities of all agents and initialize their PME context. If an agent is able to generate PME its PME status must be cleared and its PME capability initialized.

Steps to power down an external PCI agent, and the device, are as follows:

1. The PCI device driver executes code to save any PCI function context that would not otherwise survive the transition to the new PM state. The OS is required to disable the I/O and memory space as well as bus mastering via the PCI Command register.
2. The PCI device driver enables external PCI agents to generate  $\overline{\text{PCI\_PME}}$ . The driver clears PME\_Status in the agent and programs the D3Hot state into the PCI agent's Power\_State field.
3. In response to this PowerState field change the external PCI agent transitions itself to the new power state and then updates its own PowerState field in its PCI configuration registers.
4. The device detects the external agent's PowerState change by polling the agent's configuration registers. This process is repeated until all external agents are in their low power states.

5. Once all agents have transitioned appropriately the device may enter the D3Hot state. If the device wants to transition to the D3Warm state (power off a portion of the die) it must set the PMCCR1[POWER\_OFF] bit so that the EXT\_PWR\_CTRL signal will toggle, turning of the VDD externally.
6. As host, the device transitions to D3Warm state by putting the e300 into deep sleep. This action causes the qreq signal to be asserted which causes PMC to sequence in to D3Warm. If the PMCCR1[POWER\_OFF] bit is set the EXT\_PWR\_CTRL signal will transition low causing power to be removed to a portion of the die.

Note the following:

- In host mode, the PMCCR1[PME\_EN] bit should be cleared. Also, the PMCCR1[USE\_STATE] bit should be cleared to indicate that the next\_state and curr\_state fields are ignored. The next\_state and curr\_state fields in that register are NOT used in host mode. The curr\_state field should be left as 00.
- When powering down in host mode, the e300 should ensure that the PMCER is cleared. Any pending interrupt will prevent PMC from entering the low power state and asserting EXT\_PWR\_CTRL.
- Note that the next\_state field in PMCCR1 is not writable by the e300. These bits are a reflection of the values in the corresponding PCIPMR1[Power\_State] field.

### PMC Wake-Up When the MPC8313E is a PCI Agent

This sequence assumes the device is a PCI agent in D3Warm state, with power supplied only by the VDDC power rail. The VDD power rail has been externally switched to 0. Assume the PMC is set to wake on one of the defined wake-up events and that the appropriate PMCMR[x] is set to 1. Assume that PME signaling is desired and that PMCCR1[PME\_EN] = 1.

Steps to wake up are as follows.

1. A wake-up event occurs setting one of the PMCER[x] bits. The wake-up event will remain set in the PMCER[x] register until cleared by the e300.  
Alternatively, the PCI host may request a change of state in the device through the PCI's PCIPMR1[Power\_State] register, for example setting it to 00b (D0 mode). This state change will be reflected into the PMC's PMCCR1[NEXT\_STATE] register field, and the wake-up flow would continue at step 4.
2. In response to the wake-up event, PME signaling is generated asserting the  $\overline{\text{PCI\_PME}}$  signal to the host.  
Note: In order for the device to assert  $\overline{\text{PCI\_PME}}$ , the PME\_EN bit in the PCIPMR1 register must be set to 1, AND the PMCCR1[PME\_EN] bit must be set to 1.
3. The PCI host recognizes the  $\overline{\text{PCI\_PME}}$  signal. The host may opt to leave the device in its current state, in which case the sequence stops at this point. If the host decides to wake up the device, it will change the power state field, PCIPMR1[Power\_State] in the device PCI PME context block to "00" (D0). This change will be reflected in the PMCCR1[NEXT\_STATE] register bits in the PMC module.

Note that the PCIPMR1[Power\_State] register field maintains its “D3Hot” coding until the device transitions to D0. This is according to PCI specification. Once the device is awake the e300 will write the D0 state to the PMCCR1[CURR\_STATE] field, which will then be reflected into the PCIPMR1[Power\_State] field. This indicates to the host that the device has completed the transition to D0.

4. The PMCCR1[NEXT\_STATE] transition triggers the PMC to power up the device. PMC toggles the external EXT\_PWR\_CTRL signal to apply voltage to the VDD supply rail. The PMC also asserts an interrupt to the IPIC. this interrupt as well as the source of the wake-up event will be stored in PMC registers until cleared by the e300.
5. The PMC then waits for the PMC\_PWR\_OK signal to be asserted. PMC\_PWR\_OK is an indication that the external VDD supply is stable and within spec. The PMC\_PWR\_OK signal is an input to the PMC that can be supplied from an external source, or can be tied active internally. The source is determined by GPIO mux logic. If the system is designed to supply PMC\_PWR\_OK externally the mux logic would be programmed to bring PMC\_PWR\_OK in on the IIC2\_SDA input. If the system does not supply this signal mux logic would be programmed to supply the IIC2\_SDA signal on this input and PMC\_PWR\_OK would be forced active internally.
6. When the PMC\_PWR\_OK signal asserts, the PMC’s reset timer begins to count down. If the PMC\_PWR\_OK signal is programmed to be asserted externally but never asserts, this is a failure condition. The host must detect this failure through time-out by polling the agent’s state, which will never change.
7. The PMC asserts a reset to the logic blocks in the powered-off region upon entering the low power mode (D3Warm). This reset signal is ORed with the hard reset asserted during chip POR. When a wake-up event occurs the reset signal continues to be asserted for the duration of the reset timer count. This is to allow the e300 PLL to lock. During wake-up the reset configuration word (RCW) is not reloaded. The e300 PLL locks with previous RCW settings. The length of the PMC reset is programmable via software and must be initialized by an e300 device driver before power down. Guidelines will be provided as to the required length of this PMC reset.  
Note that if the PMC\_PWR\_OK signal is not used, permanently asserted, the value of the reset time should be increased to account for VDD becoming stable.
8. When the PMC reset timer expires, the pmc\_reset is de-asserted and the conditioning logic is removed allowing the powered on and powered off regions of the die to operate normally.
9. The e300 begins fetching instructions just as it did when initially reset (POR), starting at the memory address specified by MSRP[IP]. This address is decoded by the local address window logic (eLBC) and directed to a particular boot device.
10. The e300's initialization code checks the PMCCR1[POWER\_OFF] bit to see that this is a reset from the D3Warm state, not a full POR. The e300 initializes the DDR controller so that it does not do an initialization of the DDR when coming out of self-refresh (DDR\_SDRAM\_CFG[BI] = 1).
11. The initialization code initializes the IPIC controller and the MSR[EE] such that the e300 can see pending interrupts including those from the PMC. After responding, the e300 clears the IPIC, PMC, and eTSEC interrupts.
12. The PMCCR1[CURR\_STATE] register field is updated to reflect the new active state. This update is reflected in the PCIPMR1[Power\_State] field indicating to the PCI host that the device has returned to its active state.



## PMC Wake-Up When the MPC8313E is a PCI Host

This sequence assumes the device is a PCI host in D3Warm state. Power is off (0v) on the VDD voltage supply. The PMC is programmed to wake on one of its wake-events: Ethernet magic packet, GPIO, USB, internal timer, external interrupt, or  $\overline{\text{PCI\_PME}}$  from an external PCI agent. The PMCCR1[PME\_EN] bit is reset, indicating the device is a host ( $\overline{\text{PCI\_PME}}$  is an input).


Steps to wake up the device are as follows.

1. A wake-up event occurs and sets the corresponding wake-up event bit in the PMCER[x] register. This event will trigger an interrupt from the PMC to the IPIC controller if not masked. This interrupt will not be serviced by the e300 until it wakes up from D3Warm. The PMCER[x] bit remains set until cleared by e300 software.
2. The wake-up event in the PMCER[x] register triggers the PMC to begin the wake-up process. It asserts EXT\_PWR\_CTRL to the external power supply switch. This applies power to the VDD power rail.
3. From this point the sequence is similar to the agent power-up sequence described in the previous section, Step 5 on.
4. Once the e300 boots up and the IPIC has been initialized, the source of the interrupt can be detected from the PMC registers. The e300 will then reset the source of the wake-up. The source of the interrupt may be an on-chip device (eTSEC, USB, or GPIO) or an off chip device (agent) that asserted  $\overline{\text{PCI\_PME}}$ . An external PCI agent may need to have its power state changed and be initialized.
5. The e300 clears the on-chip IPIC, PMC interrupts, and wake-up events.

### 5.8.3.7.3 PMC External Power Supply Control

The following diagrams show the assumed scenarios for controlling the power supply to the device. An external means is required to switch off the VDD supply in low power mode. A commercial power switch can be used but there is usually a switching delay associated with these devices, sometimes around 1 ms. A cheaper and faster solution is to use a FET, as shown in [Figure 5-58](#) and [Figure 5-59](#). The use of the PMCCR2[RCNT] and PMCCR2[PDCNT] timer fields will allow calibration of the device to the switching characteristics of the power switch.

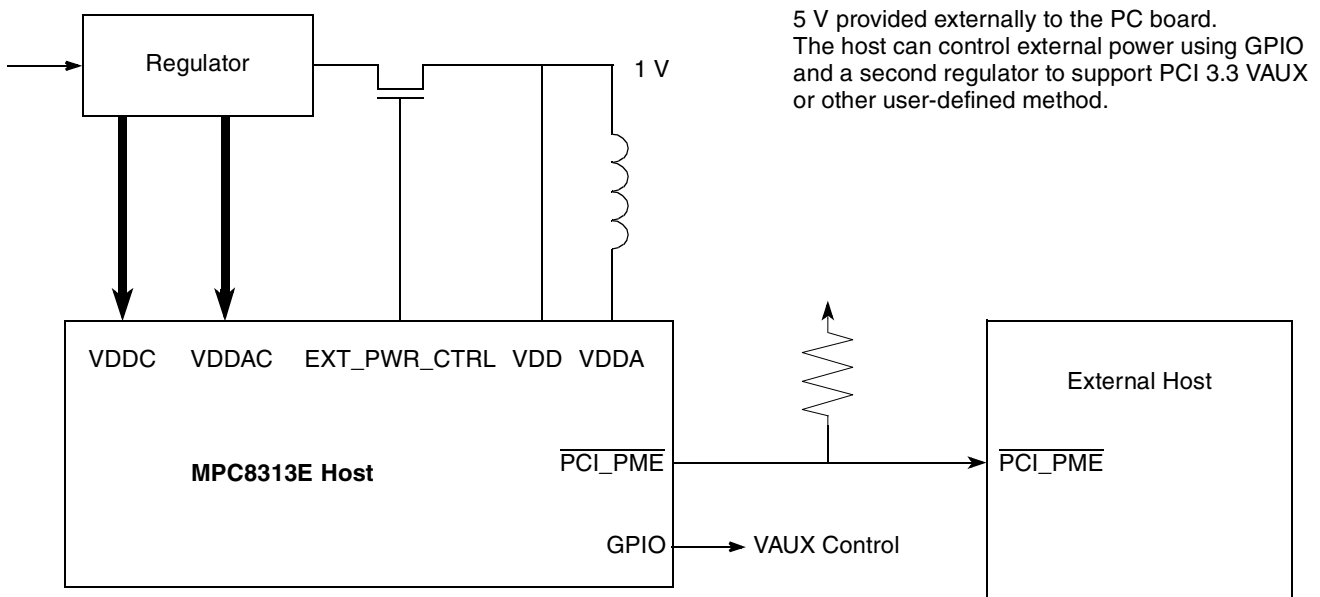
## System Configuration



### NOTES:

- N-FET: ON Semi NTR4501NT1G or similar.
- VDDC: Constant VDD for digital I/O, partial SoC digital core region.
- VDDAC: Constant VDD for analog I/O, partial SoC analog region.
- VDD: Switched VDD for internal digital logic, partial SoC digital core region.
- VDDA: Switched VDD for internal analog circuits, partial SoC analog region.

**Figure 5-58. Example VDD Control of Device as Agent Using the Optional PMC\_PWR\_OK Signal**



**Figure 5-59. Example VDD Control of Device as Host**

### 5.8.3.7.4 Low-Power Considerations

The following should be considered in a low-power system implementation.

#### PME Generation

If the host wants  $\overline{\text{PCI\_PME}}$  to be asserted it must make sure both the  $\text{PMCCR1}[\text{PME\_EN}]$  and the  $\text{PCIPMR1}[\text{PME\_EN}]$  bits are set to 1.

If the external host does not want  $\overline{\text{PCI\_PME}}$  to be asserted it can set  $\text{PCIPMR1}[\text{PME\_EN}]$  to 0. However if the  $\text{PMCCR1}[\text{PME\_EN}]$  is set to 1, the PMC will assume PME signaling. If a wake-up event occurs PMC will assume  $\text{PCI\_PME}$  will be asserted and that the host will respond requesting it to change its power state. If  $\text{PCIPMR1}[\text{PME\_EN}] = 0$ ,  $\overline{\text{PCI\_PME}}$  will not be asserted and the host will never know a wake-up event has occurred. In this case the host can still wake-up the device by writing a new power state in the  $\text{PCIPMR1}[\text{Power\_State}]$  field, but it will never know about the device wake-up event.

The  $\text{PCIPMR1}[\text{PME\_EN}]$  in the PCI configuration space will default to 0 (per spec). The  $\text{PMCCR1}[\text{PME\_EN}]$  in the PMC block also defaults to 0. If there is no PCI host in the system,  $\overline{\text{PCI\_PME}}$  generation will most likely not be required. However if  $\overline{\text{PCI\_PME}}$  assertion is desired on wake-up for some reason, the e300 can assert the two  $\text{PME\_EN}$ 's and still generate  $\overline{\text{PCI\_PME}}$ .

If  $\text{PMCCR1}[\text{PME\_EN}] = 1$  the implication is that the device is an agent and  $\overline{\text{PCI\_PME}}$  is not a wake-up event. If  $\text{PMCCR1}[\text{PME\_EN}] = 0$  the implication is that the device is the host and will wake up directly without  $\overline{\text{PCI\_PME}}$  signaling.

The device has the ability to assert  $\overline{\text{PCI\_PME}}$  under software control. The e300 can write  $\text{PMCCR1}[\text{ASSERT\_PME}] = 1$  and  $\overline{\text{PCI\_PME}}$  will be asserted (assuming  $\text{PCIPMR1}[\text{PME\_EN}] = 1$ ). This capability is provided to support PCI protocol when waking up from D1, D2 or D3Hot when the cause of the wake-up is not a defined wake-up event. In this case the e300 will wake-up or return to D0 when an interrupt is asserted or there is CSB bus activity independent of PMC. Software can manually assert  $\overline{\text{PCI\_PME}}$  in this way to inform the host that a power state change has occurred.

If when waking up from D1–D3Hot the wake-up event is a defined wake-up event, it will be registered in the  $\text{PMCER}$  register. When the event occurs, if the interrupt to the e300 is enabled, an interrupt will be asserted and e300 will wake-up directly. If  $\text{PMCCR1}[\text{PME\_EN}]$  and  $\text{PCIPMR1}[\text{PME\_EN}]$  are both set, the device will assert  $\overline{\text{PCI\_PME}}$ . In this case the software on the e300 may be written such that when it wakes up due to the PMC interrupt it waits until the host request that it go to D0 before continuing operation so that it follows PCI protocol.

If multiple wake-up events occur before the device is instructed to go to D0, multiple  $\overline{\text{PCI\_PME}}$  assertions will occur. All interrupt sources will be stored in the  $\text{PMCER}$  until cleared. As an agent ( $\text{PMCCR1}[\text{PME\_EN}] = 1$ ), a wake-up interrupt will not be asserted to the e300 until the host has instructed the device to wake-up by writing D0 into the  $\text{PCIPMR1}[\text{Power\_State}]$  field. As a host ( $\text{PMCCR1}[\text{PME\_EN}] = 0$ ) the interrupt will be asserted to the e300 as soon as the wake-up event occurs.

The  $\text{PMCCR1}[\text{USE\_STATE}]$  bit is there to tell the PMC not to respond to differences between the  $\text{next\_state}$  and  $\text{current\_state}$  values. This is mainly to prevent PMC from waking up in Host mode if there are inadvertent changes made to the  $\text{PCIPMCR1}[\text{Power\_State}]$  register.

In general, EXT\_PWR\_CTRL should only be toggled when going to D3Warm. This is according to general MPC8313E definitions for power levels. However, the implementation is flexible and allows for the assertion of EXT\_PWR\_CTRL even in D1 or D2 if desired.

### PMC\_PWR\_OK Signal

PMC\_PWR\_OK is an external input indication that the VDD that was switched off in the device D3Warm mode has returned to specified levels after a wake-up event occurs. If an external power switch device is used (not a FET), it will typically provide such a signal. When a wake-up event occurs and PMC asserts the EXT\_PWR\_CTRL signal to turn on power, it will wait until the PMC\_PWR\_OK signal is asserted before it will proceed to wait for the e300 pll to lock. If there is no external source of PMC\_PWR\_OK, that is, an external FET is used to switch power and there's no way to indicate that power is stable, then the customer will not select PMC\_PWR\_OK in the I/O multiplexing configuration. In this case PMC\_PWR\_OK will always be asserted to the PMC by the multiplexing logic and the user will need to add additional time to the PMC reset counter to insure there is enough time for power to become stable and for the e300 PLL to lock.

### PMCCR1[POWER\_OFF] Bit

The PMCCR1[POWER\_OFF] bit is set when the user wants to remove power to a portion of the die in low power mode (D3Warm). This bit also allows boot code to distinguish between boot from power-on reset and boot from D3Warm. It is referenced relative to the IMMRBAR register value. If the IMMRBAR register is modified from its default location (the device configuration registers are moved to a different location in memory), boot software must take care to ensure it can still find the PMCCR1[POWER\_OFF] bit. It may be necessary before entering D3Warm to change the IMMRBAR register back to its default location.

### Debugger in Low Power Mode

In D3Warm, the JTAG debug logic is powered down such that JTAG debug accesses will not be possible. The conditioning logic in D3Warm causes JTAG inputs (TDI) to be ignored and holds TDO constant. This will cause JTAG accesses to time out.

### SerDes

The SerDes PHY is kept in low power mode when SGMII mode is not selected through the reset configuration word. The SerDes PHY can also be put in low power mode under software control by clearing the PMCCR1[l1pen] bit.

## 5.8.4 Initialization/Application Information

### 5.8.4.1 Core Disable in Low Power Mode

If the device is required to operate with the core permanently disabled, the following steps must be taken:

1. Initialize the device with the core enable.

2. Clear PMCCR[SLPEN] and disable the core time base unit by clearing SPCR[TBEN]. See [Section 5.3.2.4, “System Priority and Configuration Register \(SPCR\),”](#) for more information.
3. The e300 core enters low power state by accessing the HID0 register.
4. Set ACR[COREDISE] in the system arbiter register with an external master (that is, PCI). This steers all device interrupts to the  $\overline{\text{PCI\_INTA}}$  so the core cannot receive any interrupt requests.

#### NOTE

Make sure that the core cannot receive any interrupt requests during the time interval between setting HID0 and setting ACR[COREDISE]. This can be achieved if the rest of the system is idle (during the initialization).

By following this flow, the e300 core remains in low power state while the rest of the system is operational, and does not get out of this state as a result of any interrupt or time-based event.



# Chapter 6

## Arbiter and Bus Monitor

This chapter describes operation theory of the arbiter in the device. In addition, it describes configuration, control, and status registers of the arbiter.

### 6.1 Overview

The arbiter is responsible for providing coherent system bus arbitration. It tracks all address and data tenures and provides all the arbitration signals to masters and slaves. In addition, it monitors the bus and reports on errors and protocol violations.

The arbiter includes the following features:

- Supports a programmable pipeline depth (from 1 to 4)
- Supports four levels of priority for bus arbitration
- Supports repeat-request mode: number of programmable consecutive transactions from the same master (up to eight transactions)
- Supports data streaming operations
- Supports programmable address bus parking mode: disable, park to last bus owner, park to software selected master
- Claims address only, reserved and illegal transaction types, report on it and can raise maskable interrupt
- Provides timers for address tenure time out and data tenure time out detection and can issue maskable interrupt, if any timer expired
- Reports on transfer error and can issue maskable interrupt
- Can issue regular or machine check interrupt for each type of error event (programmable)

#### 6.1.1 Coherent System Bus Overview

The coherent system bus is the central bus of the device. Any data transaction from master to slave in the device passes through the coherent system bus. The device coherent system bus supports pipelined transactions. It has independent address and data tenures. Pipeline depth determines the number of address tenures that can be started before the first data tenure is finished.

Basic burst size is equal to cache line length of the core, which is 32 bytes. Using repeat request mode enables up to eight consecutive bursts to be executed by the same master. Maximum number of consecutive transactions can be limited by programming arbiter configuration register. See [Section 6.2.1, “Arbiter Configuration Register \(ACR\)”](#) for more details.

**NOTE**

Write accesses to different interfaces are not guaranteed to finish in order.

**6.2 Arbiter Memory Map/Register Definition**

Table 6-1 shows the memory map for arbiter's configuration, control and status registers.

**Table 6-1. Arbiter Register Map**

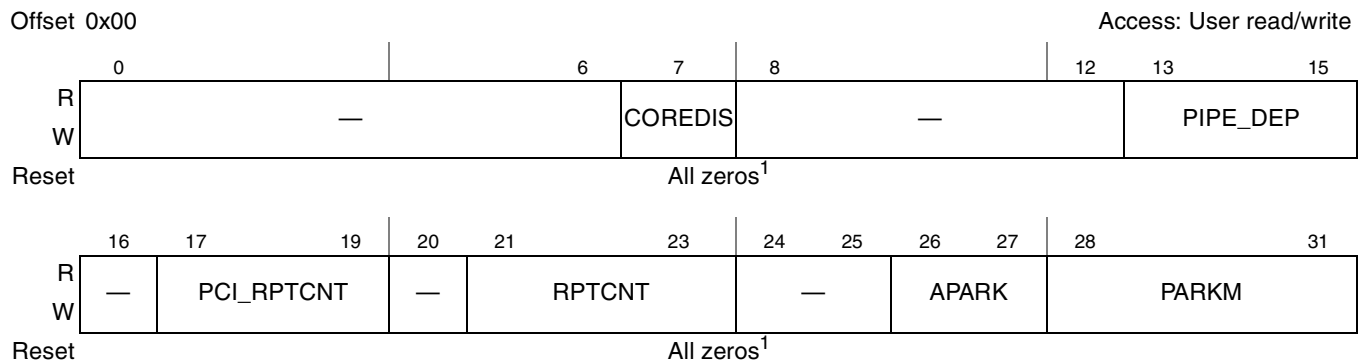
Memory Offset (Hex)	Register	Access	Reset	Section/Page
0x00	Arbiter configuration register (ACR)	R/W	0x0000_0000/ 0x0010_0000 <sup>1</sup>	6.2.1/6-2
0x04	Arbiter timers register (ATR)	R/W	0xFFFF_FFFF	6.2.2/6-4
0x0C	Arbiter event register (AER)	w1c	0x0000_0000	6.2.3/6-5
0x10	Arbiter interrupt definition register (AIDR)	R/W	0x0000_0000	6.2.4/6-6
0x14	Arbiter mask register (AMR)	R/W	0x0000_0000	6.2.5/6-7
0x18	Arbiter event attributes register (AEATR)	R	0x0000_0000 <sup>2</sup>	6.2.6/6-7
0x1C	Arbiter event address register (AEADR)	R	0x0000_0000 <sup>2</sup>	6.2.7/6-9
0x20	Arbiter event response register (AERR)	R/W	0x0000_0000	6.2.8/6-10

<sup>1</sup> Reset value is determined from the core PLL configuration of the reset configuration word. See Chapter 4, "Reset, Clocking, and Initialization," for details.

<sup>2</sup> The registers AEATR and AEADR are affected only by the assertion of  $\overline{\text{PORESET}}$

**6.2.1 Arbiter Configuration Register (ACR)**

Arbiter configuration register (ACR) defines the arbiter modes and parked master on the bus. Figure 6-1 shows the fields of ACR.



<sup>1</sup> Note that the reset value of COREDIS and bits 10–11 are determined from reset configuration word. (See Section 4.3.2, "Reset Configuration Words," for more details on reset configuration word.)

**Figure 6-1. Arbiter Configuration Register (ACR)**



Table 6-2 describes ACR fields.

**Table 6-2. ACR Field Descriptions**

Bits	Name	Description
0–6	—	Reserved, write should preserve reset value.
7	COREDIS	Core disable. Specifies whether CPU is disabled. When CPU is disabled, it cannot be granted on the bus by the arbiter. After reset, this bit receives its value from the reset configuration bit of COREDIS and can be configured by software. Also, if boot source is boot sequencer, COREDIS must be set to 1 at reset and the last transaction of the boot sequencer must set COREDIS to 0, if CPU enable is needed. 0 CPU enabled. 1 CPU disabled.
8–9	—	Write reserved, read = 0
10–11	—	Reserved. Write should preserve reset value. The reset value is a function of the core PLL configuration, which is part of the reset configuration word. When the core is set to operate at 1:1 or 3:2 bus clock, these bits are set to '01' during reset; otherwise, they are set to '00'.
12	—	Reserved, write should preserve reset value.
13–15	PIPE_DEP	Pipeline depth (number of outstanding transactions). 000 Pipeline depth 1 (1 outstanding transaction) 001 Pipeline depth 2 (2 outstanding transactions) 010 Pipeline depth 3 (3 outstanding transactions) 011 Pipeline depth 4 (4 outstanding transactions) 1xx Reserved
16	—	Reserved, write should preserve reset value.
17–19	PCI_RPTCNT	PCI repeat count. Specifies the maximum number of consecutive transactions, that PCI master can perform, using $\overline{\text{REPEAT}}$ request mode. 000 One consecutive transaction ( $\overline{\text{REPEAT}}$ request mode disable) 001 Two consecutive transactions 010 Three consecutive transactions 011 Four consecutive transactions 100 Five consecutive transactions 101 Six consecutive transactions 110 Seven consecutive transactions 111 Eight consecutive transactions
20	—	Reserved, write should preserve reset value.
21–23	RPTCNT	Repeat count. Specifies the maximum number of consecutive transactions, that any master (except PCI) can perform, using $\overline{\text{REPEAT}}$ request mode. 000 1 consecutive transactions ( $\overline{\text{REPEAT}}$ request mode disable) 001 2 consecutive transactions 010 3 consecutive transactions 011 4 consecutive transactions 100 5 consecutive transactions 101 6 consecutive transactions 110 7 consecutive transactions 111 8 consecutive transactions <b>Note:</b> It is recommended not to program this field for more than four consecutive transactions.
24–25	—	Reserved, write should preserve reset value.

**Table 6-2. ACR Field Descriptions (continued)**

Bits	Name	Description
26–27	APARK	Address parking. Specifies arbiter bus parking mode. 00 Park to master. Arbiter parks the address bus to the master, that is selected by numeric value of PARKM field. 01 Park to last owner. Arbiter parks the address bus to last bus owner. 10 Disable. Arbiter does not assert $\overline{BG}$ to any master, if no $\overline{BR}$ is present. 11 Reserved
28–31	PARKM	Parking master. 0000 e300 core 0001 PCI, DMA 0010 TSEC1, TSEC2 0011 Encryption core 0100 USB 0101–1111Reserved

## 6.2.2 Arbiter Timers Register (ATR)

The arbiter timers register (ATR) defines the arbiter address time out (ATO) and data time out (DTO) values. [Figure 6-2](#) shows the fields of ATR.

**Figure 6-2. Arbiter Timers Register (ATR)**

[Table 6-3](#) describes ATR fields.

**Table 6-3. ATR Field Descriptions**

Bits	Name	Description
0–15	DTO	Data time out. Specifies the time-out period for the data tenure. The granularity of this field is bus clocks. The maximum value is coherent system bus clocks. Data time_out occurs if the data tenure does not end before the specified time-out period. When DTO = n, the timeout cycle is n*128. 0000 Reserved 0001 128 clock cycles 0002 256 clock cycles 0003 384 clock cycles ... FFFF 8355840 clock cycles
16–31	ATO	Address time out. Specifies the time-out period for the address tenure. The granularity of this field is bus clocks. Maximum value is coherent system bus clocks. Address time-out occurs if the address tenure did not end before the specified time-out period. When ATO = n, the timeout cycle is n*128. 0000 Reserved 0001 128 clock cycles 0002 256 clock cycles 0003 384 clock cycles ... FFFF 8355840 clock cycles

## 6.2.3 Arbiter Event Register (AER)

The arbiter uses arbiter event register (AER) to report on erroneous transactions. This register is cleared by writing ones to the fields to be cleared. [Figure 6-3](#) shows the fields of AER.

Offset 0x0C

Access: User w1c



**Figure 6-3. Arbiter Event Register (AER)**

[Table 6-4](#) describes AER fields.

**Table 6-4. AER Field Descriptions**

Bits	Name	Description
0–25	—	Write reserved, read = 0
26	ETEA	Transfer error. Reports on detection of transfer error by one of the slaves. 0 No transfer error detected by one of the slaves. 1 Transfer error detected by one of the slaves.
27	RES	Reserved transfer type. Reports on transaction with reserved transfer type. See <a href="#">Section 6.3.2.5, “Reserved Transaction Type,”</a> for more information. 0 No transaction with reserved transfer type occurred. 1 Transaction with reserved transfer type occurred.
28	ECW	External control word transfer type. Reports on transaction with external control word transfer type. See <a href="#">Section 6.3.2.6, “Illegal (eciwx/ecowx) Transaction Type,”</a> for more information. 0 No transaction with external control word transfer type occurred. 1 Transaction with external control word transfer type occurred.
29	AO	Address Only transfer type. Reports on transaction with address only transfer type. See <a href="#">Section 6.3.2.4, “Address Only Transaction Type,”</a> for more information. 0 No transaction with address only transfer type occurred. 1 Transaction with address only transfer type occurred.
30	DTO	Data time out. Reports on data tenure time out. 0 Data time out timer is not expired. 1 Data time out timer is expired.
31	ATO	Address time out. Reports on address tenure time out. 0 Address time out timer is not expired. 1 Address time out timer is expired.

## 6.2.4 Arbiter Interrupt Definition Register (AIDR)

Arbiter interrupt definition register (AIDR) determines the interrupt that responds to different error conditions. Setting a bit defines the corresponding interrupt as MCP interrupt; clearing a bit defines the corresponding interrupt as regular interrupt. Figure 6-4 shows the fields of AIDR.

Offset 0x10

Access: User read/write




Figure 6-4. Arbiter Interrupt Definition Register (AIDR)

Table 6-5 describes AIDR fields.

Table 6-5. AIDR Field Descriptions

Bits	Name	Description
0–25	—	Write reserved, read = 0
26	ETEA	Transfer error. Detection of transfer error by one of the slaves interrupt definition. 0 Detection of transfer error by one of the slaves causes regular interrupt. 1 Detection of transfer error by one of the slaves causes MCP interrupt.
27	RES	Reserved transfer type. Transaction with reserved transfer type interrupt definition. 0 Transaction with reserved transfer type causes regular interrupt. 1 Transaction with reserved transfer type causes MCP interrupt.
28	ECW	External control word transfer type. Transaction with external control word transfer type interrupt definition. 0 Transaction with external control word transfer type causes regular interrupt. 1 Transaction with external control word transfer type causes MCP interrupt.
29	AO	Address only transfer type. Transaction with address only transfer type interrupt definition. 0 Transaction with address only transfer type causes regular interrupt. 1 Transaction with address only transfer type causes MCP interrupt.
30	DTO	Data time out. Data tenure time out interrupt definition. 0 Data tenure time out causes regular interrupt. 1 Data tenure time out causes MCP interrupt.
31	ATO	Address time out. Address tenure time out interrupt definition. 0 Address tenure time out causes regular interrupt. 1 Address tenure time out causes MCP interrupt.

## 6.2.5 Arbiter Mask Register (AMR)

Arbiter mask register (AMR) is used to mask interrupts or reset requests. Setting a mask bit enables the corresponding interrupt or reset request; clearing a bit masks it. Regular interrupts, MCP interrupts and reset requests can be masked by AMR register. Figure 6-5 shows the fields of AMR.

Offset 0x14

Access: User read/write




Figure 6-5. Arbiter Mask Register (AMR)

Table 6-6 describes AMR fields.

Table 6-6. AMR Field Descriptions

Bits	Name	Description
0–25	—	Write reserved, read = 0
26	ETEA	Transfer error. Detection of transfer error by one of the slaves interrupt mask bit. 0 Detection of transfer error by one of the slaves interrupt disabled. 1 Detection of transfer error by one of the slaves interrupt enabled.
27	RES	Reserved transfer type. Transaction with reserved transfer type interrupt mask bit. 0 Transaction with reserved transfer type interrupt disabled. 1 Transaction with reserved transfer type interrupt enabled.
28	ECW	External control word transfer type. Transaction with external control word transfer type interrupt mask bit. 0 Transaction with external control word transfer type interrupt disabled. 1 Transaction with external control word transfer type interrupt enabled.
29	AO	Address only transfer type. Transaction with address only transfer type interrupt mask bit. 0 Transaction with address only transfer type interrupt disabled. 1 Transaction with address only transfer type interrupt enabled.
30	DTO	Data time out. Data tenure time out interrupt mask bit. 0 Data tenure time out interrupt disabled. 1 Data tenure time out interrupt enabled.
31	ATO	Address time out. Address tenure time out interrupt mask bit. 0 Address tenure time out interrupt disabled. 1 Address tenure time out interrupt enabled.

## 6.2.6 Arbiter Event Attributes Register (AEATR)

Arbiter event attributes register (AEATR) reports the type of transaction that causes error, which is specified in the event register. See Section 6.2.3, “Arbiter Event Register (AER),” for more information. AEATR is cleared only by power-on reset. The attributes of the first error event are stored. Note that this means that AEATR does not change its value when AER is not clear. As AEATR is not affected by soft or hard reset, software can read this register and determine the cause of the bus failure, even if the failure caused a deadlock situation. Refer to Section 6.4.2, “Error Handling Sequence,” for more information.

Figure 6-6 shows the fields of AEATR.




Figure 6-6. Arbiter Event Attributes Register (AEATR)

Table 6-7 describes AEATR fields.

Table 6-7. AEATR Field Descriptions

Bits	Name	Description
0–4	—	Write reserved, read = 0
5–7	EVENT	Event type. 000 Address time out <span style="float: right;">100 Reserved transfer type</span> 001 Data time out <span style="float: right;">101 Transfer error</span> 010 Address only transfer type <span style="float: right;">11c Reserved</span> 011 External control word transfer type
8–10	—	Write reserved, read = 0
11–15	MSTR_ID	Master Id. 00000 e300 core data transaction <span style="float: right;">01001 I2C (boot sequencer)</span> 00001 Reserved <span style="float: right;">01010 JTAG</span> 00010 e300 core instruction fetch <span style="float: right;">01011 Reserved</span> 00011 Reserved <span style="float: right;">01100 Reserved</span> 00100 eTSEC1 <span style="float: right;">01101 PCI</span> 00101 eTSEC2 <span style="float: right;">01110 Reserved</span> 00110 Reserved <span style="float: right;">01111 DMA</span> 00111 USB DR <span style="float: right;">10000–11111 Reserved</span> 01000 Encryption core <b>Note:</b> Master Id reflects the source of transaction and is used for debug purpose.
16–19	—	Write reserved, read = 0
20	TBST	Transfer burst. 0 Burst transaction. Transfer size is greater than 8 bytes 1 Single-beat transaction. Transfer size is up to 8 bytes
21–23	TSIZE	Transfer size. Transfer size encoding depends on the value of the field TBST. TBST = 1: 001 1 byte <span style="float: right;">TBST = 0:</span> 010 2 bytes <span style="float: right;">000 16 bytes</span> 011 3 bytes <span style="float: right;">001 24 bytes</span> 100 4 bytes <span style="float: right;">010 32 bytes</span> 101 5 bytes <span style="float: right;">011–111 Reserved</span> 110 6 bytes 111 7 bytes 000 8 bytes
24–26	—	Write reserved, read = 0

Table 6-7. AEATR Field Descriptions (continued)

Bits	Name	Description
27–31	TTYPE	Transfer type.
	00000	Address-only
	00001	Address-only
	00010	Single-beat or burst write
	00011	Reserved
	00100	Address-only
	00101	Reserved
	00110	Burst write
	00111	Reserved
	0100x	Address-only
	0101x	Single-beat or burst read
	0110x	Address-only
	01110	Burst read
	01111	Reserved
	10000	Address-only
	1XX01	Reserved
	10010	Single-beat write
	1XX11	Reserved
	10100	<b>ecowx</b> —Illegal single-beat write
	10110	Reserved
	11000	Address-only
	11010	Single-beat or burst read
	11100	<b>eciwx</b> —Illegal single-beat read
	11110	Burst read

## 6.2.7 Arbiter Event Address Register (AEADR)

Arbiter event address register (AEADR) reports the address of transaction that causes the error, which is specified in the event register. See [Section 6.2.3, “Arbiter Event Register \(AER\),”](#) for more information. AEADR is cleared only by power-on reset. The address of the first error event is stored. Note that this means that AEADR does not change its value when AER is not clear. As AEADR is not effected by soft or hard reset, software can read this register and determine the cause of the bus failure, even if the bus failure had caused a deadlock situation. Refer to [Section 6.4.2, “Error Handling Sequence,”](#) for more information.

Figure 6-7 shows the fields of AEADR.



Figure 6-7. Arbiter Event Address Register (AEADR)

Table 6-8 describes AEADR fields.

Table 6-8. AEADR Field Descriptions

Bits	Name	Description
0–31	ADDR	Address of the event reported in AEATR register. See <a href="#">Section 6.2.6, “Arbiter Event Attributes Register (AEATR),”</a> for more information.

## 6.2.8 Arbiter Event Response Register (AERR)

Arbiter event response register (AERR) determines whether different error conditions cause interrupt or reset request. Setting a bit defines the corresponding error condition to cause reset request; clearing a bit defines the corresponding error condition to cause interrupt. Figure 6-8 shows the fields of AERR.

Offset 0x20

Access: User read/write




Figure 6-8. Arbiter Event Response Register (AERR)

Table 6-9 describes AERR field.

Table 6-9. AERR Field Descriptions

Bits	Name	Description
0–25	—	Write reserved, read = 0
26	ETEA	Transfer error. Detection of transfer error by one of the slaves event response. 0 Detection of transfer error by one of the slaves causes interrupt. 1 Detection of transfer error by one of the slaves causes reset request.
27	RES	Reserved transfer type. Transaction with reserved transfer type interrupt definition. 0 Transaction with reserved transfer type causes interrupt. 1 Transaction with reserved transfer type causes reset request.
28	ECW	External control word transfer type. Transaction with external control word transfer type interrupt definition. 0 Transaction with external control word transfer type causes interrupt. 1 Transaction with external control word transfer type causes reset request.
29	AO	Address only transfer type. Transaction with address only transfer type interrupt definition. 0 Transaction with address only transfer type causes interrupt. 1 Transaction with address only transfer type causes reset request.
30	DTO	Data time out. Data tenure time out interrupt definition. 0 Data tenure time out causes interrupt. 1 Data tenure time out causes reset request.
31	ATO	Address time out. Address tenure time out interrupt definition. 0 Address tenure time out causes interrupt. 1 Address tenure time out causes reset request.

## 6.3 Functional Description

The following sections describe arbiter functionality: arbitration policy and bus error detection.

### 6.3.1 Arbitration Policy

The arbitration process involves the masters and the arbiter. Masters arbitrate on the privilege to own an address tenure. For data tenures, the arbiter uses the same order of transactions as address tenures.



Figure 6-9 shows the interface signals between the arbiter and masters that are involved in address bus arbitration.

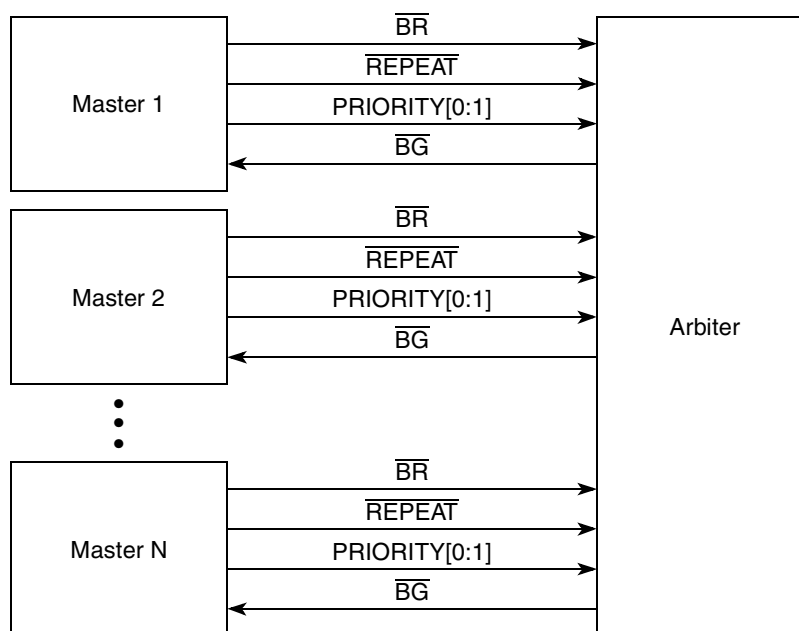


Figure 6-9. Address Bus Arbitration

A master has to acquire address bus ownership before it starts any transaction. The master asserts its own bus request signal along with the arbitration attribute signals  $\overline{REPEAT}$  &  $PRIORITY[0:1]$ . The arbiter later asserts the corresponding address bus grant signal to the requesting master depending on the system states and arbitration scheme. See Section 6.3.1.1, “Address Bus Arbitration with  $PRIORITY[0:1]$ ,” for details on arbitration scheme. When address bus grant is received the master can start the address tenure.

### 6.3.1.1 Address Bus Arbitration with $PRIORITY[0:1]$

Whenever a master asserts its bus request to acquire address bus ownership, it can drive its  $PRIORITY[0:1]$  signals to indicate request priority. The master would be served sooner because of its higher priority level. The arbiter takes this extra information into consideration in order to yield better service for a higher priority request than a lower priority request. Therefore, the arbiter operates according to the following priority based arbitration scheme:

1. For every priority level a fair arbitration scheme is used (a simple round robin scheme)
2. For every priority level other than 0, one place is reserved as a place holder for lower level arbitration rings.
3. Each master can change its priority level at any time.

Figure 6-10 shows an example of priority-based arbitration algorithm with four priority levels. In this example, if all masters request the bus continuously, the following order of bus grants occurs with the specific bandwidth:

- M6 gets 1/2 of the bus bandwidth
- M4 and M5 each gets 1/6 of the bus bandwidth

- M0 and M3 each gets 1/18 of the bus bandwidth
- M1 and M2 each gets 1/36 of the bus bandwidth




Figure 6-10. An Example of Priority-Based Arbitration Algorithm

**NOTE**

See each bus master’s chapter and [Section 5.3.2.4, “System Priority and Configuration Register \(SPCR\),”](#) for more details about priority programming.

**6.3.1.2 Address Bus Arbitration with REPEAT**

When a master owns the address bus and wants to perform another transaction, it can assert bus request along with  $\overline{\text{REPEAT}}$ , to make a repeat request to the arbiter. Consequently, the arbiter asserts bus grant to the same master if the current address tenure is not being  $\overline{\text{ARTRY}}$ ed. This happens regardless of the priority level of bus request from other masters. In another word, “repeat request” overrides the priority scheme.

Even though repeat request can improve the page hit ratio and the overall memory bandwidth efficiency, it can increase the worst case latency of individual master. Therefore, the arbiter has programmable counter to limit the maximum number of consecutive transactions that are performed by masters. Whenever the counter expires, arbiter ignores the  $\overline{\text{REPEAT}}$  signal and falls back to the regular arbitration scheme. PCI master has a dedicated repeat counter as it might need more repeated transactions before accepting read requests. PCI ordering rules require that the PCI bridge should empty all queued write operations before any new read operation can begin. See Section 3.2.5, “Transaction Ordering and Posting,” of the *PCI Local Bus Specifications Rev 2.2*, for more information.

See [Section 6.2.1, “Arbiter Configuration Register \(ACR\),”](#) for more details about programming ACR[RPTCNT]PCI.

### 6.3.1.3 Address Bus Arbitration after $\overline{\text{ARTRY}}$

The  $\overline{\text{ARTRY}}$  protocol is used primarily by the CPU to interrupt a transaction that hits to a modified line in its D-cache, so that it can maintain data coherency by performing the snoop copyback. When CPU asserts  $\overline{\text{ARTRY}}$ , the bus is immediately granted to the CPU to perform snoop copyback. After the completion of snoop copyback, the arbiter grants the bus back to the master that had its transaction  $\overline{\text{ARTRY}}$ ed.

### 6.3.1.4 Address Bus Parking

The arbiter supports address bus parking. This feature implies that when no master is requesting the bus (all bus requests are negated), the arbiter can choose to park the address bus (or assert the address bus grant) to a master. The parked master can skip the bus request and assume the bus mastership directly. This reduces the access latency for parked master.

See [Section 6.2.1, “Arbiter Configuration Register \(ACR\),”](#) for more details about ACR[APARK] and ACR[PARKM].

### 6.3.1.5 Data Bus Arbitration

For every committed address tenure a data tenure is required to complete the transaction.

In the device system, the arbiter controls the issuing of data bus grants to a master and a slave, which are involved in a data tenure of a previously performed address tenure.

## 6.3.2 Bus Error Detection

The arbiter is responsible for tracking the following cases on the bus:

- Address time out
- Data time out
- Transfer error
- Address only transaction type
- Reserved transaction type
- Illegal (**eciwx/ecowx**) transaction type

### 6.3.2.1 Address Time Out

Address time out occurs, if the address tenure was not ended before the specified time-out period (programmed by ATR[ATO]). In this case, the arbiter performs as follows:

1. Ends the address tenure.
2. Starts data tenure and ends it by asserting transfer error.
3. Reports on the event to AER[ATO].

4. Issues reset request, MCP or regular interrupt according to AERR[ATO] and AIDR[ATO], if enabled by AMR[ATO].
5. Updates transaction attributes and address of AEATR and AEADR for the first error event.

### 6.3.2.2 Data Time Out

Data time out occurs, if the data tenure was not ended before the specified time-out period (programmed by ATR[DTO]). In this case, the arbiter performs as follows:

1. Ends the data tenure by asserting transfer error.
2. Reports on this event in AER[DTO].
3. Issues reset request, MCP or regular interrupt according to AERR[DTO] and AIDR[DTO], if enabled by AMR[DTO].
4. Updates transaction attributes and address of AEATR and AEADR for the first error event.

### 6.3.2.3 Transfer Error

The arbiter tracks the transfer error asserted by one of the slaves. In this case, the arbiter performs as follows:

1. Reports on the event to AER[ETEA].
2. Issues reset request, MCP or regular interrupt according to AERR[ETEA] and AIDR[ETEA] if enabled by AMR[ETEA].
3. Updates transaction attributes and address of AEATR and AEADR for the first error event.

### 6.3.2.4 Address Only Transaction Type

Table 6-10 shows transaction types, which are defined as address only:

**Table 6-10. Address Only Transaction Type Encoding**

ttype[0:4]	Bus Commands
00000	Clean block
00100	Flush block
01000	Sync
01100	Kill block
10000	eieio
11000	TLB Invalidate
00001	lwarx reservation set
01001	tlbsync
01101	icbi

The arbiter allows address-only (AO) transactions on the bus and the G2 core has ability to issue address-only (AO) transactions (see HID0 [ABE] in the *G2 PowerPC Core Reference Manual*, Rev 1). As there is no advantage in using AO transaction in this system, the bus monitor allows the detection of AO transactions and treats them as an error.

The transaction with an address only transfer type, the arbiter performs as follows:

1. Ends the address tenure by asserting  $\overline{\text{AACK}}$ .
2. Reports on the event to AER[AO].
3. Issues reset request, MCP or regular interrupt according to AERR[AO] and AIDR[AO] if enabled by AMR[AO].
4. Updates transaction attributes and address of AEATR and AEADR for the first error event.

### 6.3.2.5 Reserved Transaction Type

Table 6-11 shows transaction types defined as reserved.

**Table 6-11. Reserved Transaction Type Encoding**

ttype[0:4]	Bus Commands
00101	Reserved
1xx01	Reserved for customer
10110	Reserved
00011	Reserved
00111	Reserved
01111	Reserved
1xx11	Reserved for customer

The transaction with a reserved transfer type, the arbiter performs as follows:

1. Ends the address tenure by asserting  $\overline{\text{AACK}}$ .
2. Reports on the event to AER[RES].
3. Issues reset request, MCP or regular interrupt according to AERR[RES] and AIDR[RES], if enabled by AMR[RES].
4. Updates transaction attributes and address of AEATR and AEADR for the first error event.

### 6.3.2.6 Illegal (eciwx/ecowx) Transaction Type

Table 6-12 shows transaction types defined as illegal.

**Table 6-12. Illegal Transaction Type Encoding**

ttype[0:4]	Bus command
10100	External control word write ( <b>ecowx</b> )
11100	External control word read ( <b>eciwx</b> )

The transaction with an illegal (eciwx, ecowx) transfer type, the arbiter performs as follows:

1. Ends the address tenure by asserting  $\overline{\text{AACK}}$ .
2. Starts data tenure and ends data tenure by asserting  $\overline{\text{TEA}}$ .
3. Reports on the event in AER[ECW].
4. Issues reset request, MCP or regular interrupt according to AERR[ECW] and AIDR[ECW], if enabled by AMR[ECW].
5. Updates transaction attributes and address of AEATR and AEADR for the first error event.

See Section 6.2.3, “Arbiter Event Register (AER),” Section 6.2.4, “Arbiter Interrupt Definition Register (AIDR),” Section 6.2.5, “Arbiter Mask Register (AMR),” Section 6.2.6, “Arbiter Event Attributes Register (AEATR),” Section 6.2.7, “Arbiter Event Address Register (AEADR),” and Section 6.2.8, “Arbiter Event Response Register (AERR),” for more information.

## 6.4 Initialization/Applications Information

The following sections describe the initialization and error handling sequences for the arbiter.

### 6.4.1 Initialization Sequence

The following initialization sequence is recommended:

1. Write to ACR to configure pipeline depth, address bus parking mode, global maximum repeat count PCI maximum.
2. Write to AERR defines whether different error events cause a reset request or an interrupt.
3. Write to AIDR defines the kind of interrupt (regular or MCP) caused by each error event. Note that this is necessary only if interrupts are enabled and AERR defines error events to cause interrupt.
4. Write to AMR to enable interrupts.
5. Write to ATR to set the ATO and DTO timers. Note that this is only necessary if the required timers are less than the maximum value (which is default).

### 6.4.2 Error Handling Sequence

The following error handling sequence is recommended:

1. Read to AER to find out about the error that occurred in the system. Also, read the values of AEATR and AEADR to check on the first error event in the system.
2. If those registers are not accessible because of a bus deadlock situation, reset the chip and read the values of the AEATR and AEADR registers to check on the event that causes this problem to the system. Use  $\overline{\text{HRESET}}$  to reset the chip to guarantee that the information stored in AEATR and AEADR is not lost.
3. Clear all the previous events by writing ones to the AER. This register is also cleared after reset.

# Chapter 7

## e300 Processor Core Overview

This chapter provides an overview of features for the embedded microprocessors in the e300 core family, which are PowerPC microprocessors built on Power Architecture technology. Throughout this chapter, the terms ‘e300 core’, ‘core’, and ‘processor’ are used interchangeably. The term ‘e300c3’ is used when describing an implementation-specific feature or when a difference exists between different configurations. The term ‘e300’ is used when describing a feature that pertains to the family of e300 processors. The MPC8313E uses an e300c3 core.

### 7.1 Overview

This section describes the details of the e300 core, provides a block diagram showing the major functional units, and briefly describes how these units interact. All differences between the e300 and previous PowerPC implementations derived from the MPC603e processor are noted. For additional information, please refer to the *e300 PowerPC Core Family Reference Manual*.

The e300 core is a low-power implementation of this microprocessor family of reduced instruction set computing (RISC) microprocessors. The core implements the 32-bit portion of the architecture, which defines 32-bit effective addresses, integer data types of 8, 16, and 32 bits, and floating-point data types of 32 and 64 bits.

The core is a superscalar processor that can issue and retire as many as three instructions per clock cycle. Instructions can execute out of program order for increased performance; however, the core makes completion appear sequential.

The e300 core integrates independent execution units including: an integer unit (IU) a floating-point unit (FPU), a branch processing unit (BPU), a load/store unit (LSU), and a system register unit (SRU). The e300c3 integrates an additional integer unit for a total of two IUs. The ability to execute instructions in parallel and the use of simple instructions with rapid execution times yield high efficiency and throughput for e300-core-based systems. Most integer instructions execute in one clock cycle. The additional IUs along with enhanced multipliers in the e300c3 improve multiply instructions to a maximum two-cycle latency, a significant improvement from previous processors. In the e300c3 core, the FPU is pipelined so a single-precision multiply-add instruction can be issued and completed every clock cycle. The e300c3 core provide hardware support for all single- and double-precision floating-point operations for most value representations and all rounding modes.

[Figure 7-1](#) shows a block diagram of the e300c3 core. Note that the e300c3 supports floating-point operations and includes two integer units.




Figure 7-1. e300c3 Core Block Diagram



The e300c3 includes 16-Kbyte, four way set-associative instruction and data caches. The MMUs contain 64-entry, two-way, set-associative, data and instruction translation lookaside buffers (DTLB and ITLB) that provide support for demand-paged, virtual-memory, address translation, and variable-sized block translation. The TLBs use a least recently used (LRU) replacement algorithm and the caches use a pseudo least recently used algorithm (PLRU).

The core also supports block address translation through the use of two independent instruction and data block address translation (IBAT and DBAT) arrays, each containing eight pairs of BATs, an increase from four pairs of each type of BATs in the G2 core. This increase provides more flexibility in protecting accesses and providing translation on a segment, block, or page basis for memory accesses and I/O accesses. Effective addresses are compared simultaneously with all eight entries in the BAT array during block translation. In accordance with the PowerPC architecture, if an effective address hits in both the TLB and BAT array, the BAT translation takes priority.

As part of the coherent system bus (CSB), the e300 core has a 64-bit data bus and a 32-bit address bus. During normal operation, the e300 core provides a three-state (modified, exclusive, and invalid) coherency protocol which is a compatible subset of a four-state (modified/exclusive/shared/invalid) MESI protocol. However, the e300 data cache contains a programmable MESI extension that supports the shared cache coherency state (similar to other PowerPC processors). Both protocols operate coherently in systems that contain four-state caches. Although MESI is supported by the e300 core, it is not implemented on the MPC8313E. The core also supports single-beat and burst data transfers for memory accesses and supports memory-mapped I/O operations.

The true little-endian mode is another enhanced capability of the e300 core. Unlike the PowerPC little-endian mode (which manipulates only the address bits), no longer supported on the e300, the true little-endian mode actually operates on true little-endian instructions and data from memory.

The critical interrupt is an additional interrupt in the e300 core and has higher priority order than the system management interrupt. Also, debug features are improved in the e300. Additional SPRG interrupt handling registers are provided for enhancing flexibility for the operating system.

The e300c3 include a performance monitor facility that provides the ability to monitor and count predefined events such as core clocks, misses in the instruction cache, data cache, or L2 cache, types of instructions dispatched, mispredicted branches, and other occurrences. The count of such events (which may be an approximation) can be used to trigger the performance monitor interrupt. [Section 7.1.7.5, “Core Performance Monitor,”](#) describes the operation of the performance monitor diagnostic tool.

## 7.1.1 Features

This section describes the major features of the e300 core:

- High-performance, superscalar microprocessor core
  - As many as three instructions issued and retired per clock (two instructions plus one branch instruction)
  - As many as five instructions in execution per clock
  - Single-cycle execution for most instructions
  - Pipelined floating-point unit (FPU) for all single- and double-precision operations

- Independent execution units and two register files
  - Branch processing unit (BPU) featuring static branch prediction
  - Two 32-bit integer units (IU) in the e300c3 .
  - FPU based on the IEEE Std 754™ for both single- and double-precision operations
  - Load/store unit (LSU) for data transfer between data-cache and general-purpose registers (GPRs) and floating-point registers (FPRs)
  - System register unit (SRU) that executes condition register (CR), special-purpose register (SPR), and integer add/compare instructions. Add/compare instructions are also executed in the IUs.
  - Thirty-two 32-bit GPRs for integer operands
  - Thirty-two 64-bit FPRs for single- or double-precision operands
- High instruction and data throughput
  - Zero-cycle branch capability (branch folding)
  - Programmable static branch prediction on unresolved conditional branches
  - Two integer units with enhanced multipliers in the e300c3 for increased integer instruction throughput and a maximum two-cycle latency for multiply instructions
  - Instruction fetch unit capable of fetching two instructions per clock from the instruction cache
  - A six-entry instruction queue (IQ) that provides lookahead capability
  - Independent pipelines with feed-forwarding that reduces data dependencies in hardware
  - 16-Kbyte, four-way set-associative instruction and data caches on the e300c3.
  - Cache write-back or write-through operation programmable on a per-page or per-block basis
  - Features for instruction and data cache locking and protection
  - BPU that performs CR lookahead operations
  - Address translation facilities for 4-Kbyte page size, variable block size, and 256-Mbyte segment size
  - A 64-entry, two-way, set-associative ITLB and DTLB
  - Eight-entry data and instruction BAT arrays providing 128-Kbyte to 256-Mbyte blocks
  - Software table search operations and updates supported through fast trap mechanism
  - 52-bit virtual address; 32-bit physical address
- Facilities for enhanced system performance
  - A 64-bit split-transaction internal data bus interface to the coherent system bus (CSB) with burst transfers
  - Support for one-level address pipelining on the CSB interface
  - True little-endian mode for compatibility with other true little-endian devices
  - Critical interrupt support
  - Hardware support for misaligned little-endian accesses
  - Configurable processor bus frequency multipliers as defined in the *MPC8313E Integrated Processor Hardware Specifications*

- Integrated power management
  - Internal processor/bus clock multiplier ratios
  - Three power-saving modes: doze, nap, and sleep
  - Automatic dynamic power reduction when internal functional units are idle
- In-system testability and debugging features through JTAG boundary-scan capability

Features specific to the e300 core not present on the G2 processors follow:

- Enhancements to the register set
  - The e300 core has one more HID0 bit than the G2:
    - The enable cache parity checking (ECPE) bit, HID0[1], gives the e300 core the ability to enable the taking of a machine check interrupt based on the detection of a cache parity error
- Enhancements to cache implementation
  - 16-Kbyte, four-way set-associative instruction and data caches on the e300c3.
  - Full parity checking is performed on both instruction and data cache memory arrays
  - Lockable L1 instruction and data caches—entire cache or on a per-way basis up to 3 of 4 ways on the e300c3
  - New **icbt** instruction supports initialization of instruction cache
  - Data cache supports four-state MESI coherency protocol (not implemented on MPC8313E)
  - The instruction cache is blocked only until the critical load completes (hit under reloads allowed)
  - Instruction cancel mechanism improves utilization of instruction cache by supporting hits-under-cancels and misses-under-cancels.
  - The critical double word is simultaneously written to the cache and forwarded to the requesting unit, thus minimizing stalls due to load delays.
  - Data cache queue sharing makes cast-outs and snoop pushes more efficient
  - Provides for an optional data cache operation broadcast feature (enabled by HID0[ABE]) that allows for coherent system management. All of the data cache control instructions, except **dcbz** (**dcbi**, **dcbf**, and **dcbst**) require that HID0[ABE] be enabled to broadcast.
  - Instruction fetch burst feature allows all instruction fetches from caching-inhibited space to be performed on the bus as burst transactions
- Interrupts
  - The e300 core offers hardware support for misaligned little-endian accesses. Little-endian load/store accesses that are not on a word boundary, except for strings and multiples, generate interrupts under the same circumstances as big-endian accesses.
  - The e300 core supports true little-endian mode to minimize the impact on software porting from true little-endian systems.
  - An input interrupt signal,  $\overline{cint}$ , is provided to trigger the critical interrupt exception on the e300 core. The pm\_event\_in input signal can be used by the performance monitor counters to trigger an interrupt upon overflow on the e300c3 .
- Bus clock—PLL configuration signals include seven signals for settings and control: *pll\_cfg[0:6]*.

- Debug features
  - Breakpoint status recorded in DBCR and IBCR control registers
  - Two signals for the debug interface: *stopped* and *ext\_halt*
  - Performance monitor registers for system analysis in the e300c3

[Figure 7-1](#) provides a block diagram of the e300 core that shows how the execution units—IU, FPU, BPU, LSU, and SRU—operate independently and in parallel. It should be noted that this is a conceptual diagram and does not attempt to show how these features are physically implemented on the device.

The e300 core provides address translation and protection facilities, including an ITLB, DTLB, and instruction and data BAT arrays. Instruction fetching and issuing are handled in the instruction unit. Translation of addresses for cache or external memory accesses are handled by the MMUs. Both units are discussed in more detail in [Section 7.1.2, “Instruction Unit,”](#) and [Section 7.1.5.1, “Memory Management Units \(MMUs\).”](#)

## 7.1.2 Instruction Unit

As shown in [Figure 7-1](#), the e300 core instruction unit, containing a fetch unit, instruction queue, dispatch unit, and BPU, provides centralized control of instruction flow to the execution units. The instruction unit determines the address of the next instruction to be fetched based on information from the sequential fetcher and from the BPU.

The instruction unit fetches the instructions from the instruction cache into the instruction queue. The BPU receives branch instructions from the fetcher and uses static branch prediction to allow fetching from a predicted instruction stream while a conditional branch is evaluated. The BPU folds out for unconditional branch instructions and conditional branch instructions unaffected by instructions in the execution pipeline.

Instructions issued beyond a predicted branch cannot complete execution until the branch is resolved, preserving the programming model of sequential execution. If any of these are branch instructions, they are decoded but not issued. Instructions to be executed by the FPU, IU, LSU, and SRU are issued and allowed to progress up to the register write-back stage. Write-back is allowed when a correctly predicted branch is resolved, and execution continues along the predicted path.

If branch prediction is incorrect, the instruction unit flushes all predicted path instructions, and instructions are issued from the correct path.

### 7.1.2.1 Instruction Queue and Dispatch Unit

The instruction queue (IQ), shown in [Figure 7-1](#), holds as many as six instructions and loads up to two instructions from the instruction unit during a single cycle. The instruction fetch unit continuously loads as many instructions as space in the IQ allows. Instructions are dispatched to their respective execution units from the dispatch unit at a maximum rate of two instructions per cycle. Dispatching is facilitated to the IUs, FPU, LSU, and SRU by the provision of a reservation station at each unit. The dispatch unit performs source and destination register dependency checking, determines dispatch serializations, and inhibits subsequent instruction dispatching as required.

For a more detailed overview of instruction dispatch, see [Section 7.3.6, “Instruction Timing.”](#)

### 7.1.2.2 Branch Processing Unit (BPU)

The BPU receives branch instructions from the fetch unit and performs CR lookahead operations on conditional branches to resolve them early, achieving the effect of a zero-cycle branch in many cases.

The BPU uses a bit in the instruction encoding to predict the direction of the conditional branch. Therefore, when an unresolved conditional branch instruction is encountered, the core fetches instructions from the predicted target stream until the conditional branch is resolved.

The BPU contains an adder to compute branch target addresses and three user-control registers: the link register (LR), the count register (CTR), and the conditional register (CR). The BPU calculates the return pointer for sub-routine calls and saves it into the LR for certain types of branch instructions. The LR also contains the branch target address for the Branch Conditional to Link Register (**bclrx**) instruction. The CTR contains the branch target address for the Branch Conditional to Count Register (**bctrx**) instruction. The contents of the LR and CTR can be copied to or from any GPR. Because the BPU uses dedicated registers rather than GPRs or FPRs, execution of branch instructions is largely independent from execution of integer and floating-point instructions.

### 7.1.3 Independent Execution Units

The PowerPC architecture's support for independent execution units allows implementation of processors with out-of-order instruction execution. For example, because branch instructions do not depend on GPRs or FPRs, branches can often be resolved early, eliminating stalls caused by taken branches.

The four other execution units and the completion unit are described in the following sections.

#### 7.1.3.1 Integer Unit (IU)

The IU executes all integer instructions. The IU executes one integer instruction at a time, performing computations with its arithmetic logic unit (ALU), multiplier, divider, and XER register. Most integer instructions are single-cycle instructions. The 32 GPRs hold integer operands. Stalls due to contention for GPRs are minimized by the automatic allocation of rename registers. The core writes the contents of the rename registers to the appropriate GPR when integer instructions are retired by the completion unit. The e300c3 provides two integer units for greater integer instruction throughput along with enhanced multipliers in each IU for faster multiply-instruction execution.

#### 7.1.3.2 Floating-Point Unit (FPU)

The FPU contains a single-precision multiply-add array and the floating-point status and control register (FPSCR). The multiply-add array allows the core to efficiently implement multiply and multiply-add operations. The FPU is pipelined so that single- and double-precision instructions can be issued back-to-back. The 32 FPRs are provided to support floating-point operations. Stalls due to contention for FPRs are minimized by the automatic allocation of rename registers. The core writes the contents of the rename registers to the appropriate FPR when floating-point instructions are retired by the completion unit.

The e300c3 core supports all floating-point data types based on the IEEE 754 standard (normalized, denormalized, NaN, zero, and infinity) in hardware, eliminating the latency incurred by software interrupt routines.

### 7.1.3.3 Load/Store Unit (LSU)

The LSU executes all load and store instructions and provides the data transfer interface between the GPRs, FPRs, and the cache/memory subsystem. The LSU calculates effective addresses, performs data alignment, and provides sequencing for load/store string and multiple instructions.

Load and store instructions are issued and executed in program order; however, the memory accesses can occur out of order. Synchronizing instructions are provided to enforce strict ordering.

Cacheable loads, when free of data bus dependencies, can execute out of order with a maximum throughput of one per cycle and with a two-cycle total latency. Data returned from the cache is held in a rename register until the completion logic commits the value to a GPR or FPR. Stores cannot be executed in a predicted manner and are held in the store queue until the completion logic signals that the store operation is to be completed to memory. The core executes store instructions with a maximum throughput of one per cycle and with a three-cycle total latency. The time required to perform the actual load or store depends on whether the operation involves the cache, system memory, or an I/O device.

### 7.1.3.4 System Register Unit (SRU)

The SRU executes various system-level instructions, including condition register logical operations and move to/from special-purpose register instructions. It also executes integer add/compare instructions. In order to maintain system state, most instructions executed by the SRU are completion-serialized; that is, the instruction is held for execution in the SRU until all prior instructions issued have completed. Results from completion-serialized instructions executed by the SRU are not available or forwarded for subsequent instructions until they complete.

## 7.1.4 Completion Unit

The completion unit tracks instructions in program order from dispatch through execution and then completes. Completing an instruction commits the core to any architectural register changes caused by that instruction. In-order completion ensures the correct architectural state when the core must recover from a mispredicted branch or an interrupt.

Instruction state and other information required for completion is kept in a five-entry FIFO completion queue. A single completion queue entry is allocated for each instruction once it enters the execution unit from the dispatch unit. An available completion queue entry is a required resource for dispatch; if no completion entry is available, dispatch stalls. A maximum of two instructions per cycle are completed in order from the queue.

## 7.1.5 Memory Subsystem Support

The core provides separate instruction and data caches and MMUs. The core also provides an efficient processor bus interface to facilitate access to main memory and other bus subsystems. The memory subsystem support functions are described in the following sections.

### 7.1.5.1 Memory Management Units (MMUs)

The core MMUs support up to 4 Petabytes ( $2^{52}$ ) of virtual memory and 4 Gigabytes ( $2^{32}$ ) of physical memory (referred to as real memory in the architecture specification) for instruction and data. The MMUs also control access privileges for these spaces on block and page granularities. Referenced and changed status is maintained by the processor for each page to assist implementation of a demand-paged virtual memory system. Note that software assistance is required for the device to maintain reference and changed status. A key bit is implemented to provide information about memory protection violations prior to page table search operations.

The LSU calculates effective addresses for data loads and stores, performs data alignment to and from cache memory, and provides the sequencing for load and store string and multiple word instructions. The instruction unit calculates effective addresses for instruction fetching.

After an EA is generated, its higher-order bits are translated by the appropriate MMU into physical address bits. The lower-order EA bits are the same on the physical address which are directed to the on-chip cache and formed the index into a four-way set-associative tag array. After translating the address, the MMU passes the higher-order physical address bits to the cache and the cache lookup completes. For caching-inhibited accesses or accesses that miss in the cache, the untranslated lower-order address bits are concatenated with the translated higher-order address bits; the resulting 32-bit physical address is then used by the memory unit and the core interface to access external memory.

The MMU also directs the address translation and enforces the protection hierarchy programmed by the operating system in relation to the supervisor/user privilege level of the access and in relation to whether the access is a load or store.

For instruction fetches, the IMMU looks for the address in the ITLB and in the IBAT array. If an address hits both, the IBAT array translation is used. Data accesses cause a lookup in the DTLB and DBAT array. In most cases, the translation is in a TLB and the physical address bits are available to the on-chip cache.

The e300 core implements four more IBAT and four more DBAT entries than the G2.

When the EA misses in the TLBs, the core provides hardware assistance for software to perform a search of the translation tables in memory. The hardware assist consists of the following features:

- Automatic storage of the missed effective address in IMISS and DMISS
- Automatic generation of the primary and secondary hashed real addresses of the page-table entry group (PTEG), which are readable from the HASH1 and HASH2 register locations.  
The HASH data is generated from the contents of the IMISS or DMISS register. The register that is selected depends on the miss (instruction or data) that was last acknowledged.
- Automatic generation of the first word of the page table entry (PTE) of the tables being searched
- A real page address (RPA) register that matches the format of the lower word of the PTE
- TLB access instructions (**tlbli** and **tlbld**) that are used to load an address translation into the instruction or data TLBs
- Shadow registers for GPR0–GPR3 that allow miss code to execute without corrupting the state of any of the existing GPRs. Shadow registers are used only for servicing a TLB miss.

See [Section 7.3.5.2, “Implementation-Specific Memory Management,”](#) for more information about memory management for the core.

### 7.1.5.2 Cache Units

The e300c3 provides 16-Kbyte, four-way set-associative instruction and data caches. The cache block is 32 bytes long. The caches adhere to a write-back policy, but the e300 core allows control of cacheability, write policy, and memory coherency at the page and block levels. The caches use a pseudo LRU replacement policy.

As shown in [Figure 7-1](#), the caches provide a 64-bit interface to the instruction fetch unit and LSU. The surrounding logic selects, organizes, and forwards the requested information to the requesting unit. Write operations to the cache can be performed on a byte basis, and a complete read-modify-write operation to the cache can occur in each cycle.

The load/store and instruction fetch units provide the caches with the address of the data or instruction to be fetched. In the case of a cache hit, the cache returns two words to the requesting unit.

Because the data cache tags are single-ported, simultaneous load/store and snoop accesses cause resource contention. Snoop accesses have the highest priority and are given first access to the tags, unless the snoop access coincides with a tag write; in this case the snoop is retried and must re-arbitrate for cache access. Loads or stores deferred due to snoop accesses are performed on the clock cycle following the snoop.

The e300 core includes a new instruction cancel extension. The instruction cancel extension improves utilization of the instruction cache during cancel operations. It allows a new instruction fetch to be issued to the cache or to the bus if a canceled instruction fetch is pending or active on the bus. This supports hit-under-cancel and miss-under-cancel instruction fetch operations.

### 7.1.6 Bus Interface Unit (BIU)

Because the caches are on-chip, write-back caches, the most common transactions are burst-read memory operations, burst-write memory operations, and single-beat (noncacheable or write-through) memory read and write operations. There can also be address-only operations, variants of the burst and single-beat operations, (for example, global memory operations that are snooped and atomic memory operations), and address retry activity (for example, when a snooped read access hits a modified cache block).

Memory accesses can occur in single-beat (1–8 bytes) and four-beat burst (32 bytes) data transfers on the 64-bit data bus. The address and data buses operate independently to support pipelining and split transactions during memory accesses.

The e300 bus interface unit (BIU) has been enhanced to allow a pipeline slot to become available once a previous transaction has been granted the data bus (that is, as early as when the data tenure starts rather than after the data tenure completes), thus allowing for greater bus utilization in systems that support it. This is sometimes referred to as 1 1/2-level pipelining.

Typically, memory accesses are weakly ordered, meaning that sequences of operations, including load/store string and multiple instructions, do not necessarily complete in the order they begin. This weak ordering maximizes the efficiency of the bus without sacrificing coherency of the data. The core allows read operations to precede store operations (except when a dependency exists, or in cases where a noncacheable access is performed), and provides support for a write operation to proceed a previously queued read data tenure (for example, allowing a snoop push to be enveloped by the address and data



tenures of a read operation). Because the processor can dynamically optimize run-time ordering of load/store traffic, overall performance is improved.

## 7.1.7 System Support Functions

The e300 core implements several support functions that include power management, time base/decrementer registers for system timing tasks, a JTAG (based on IEEE Std 1149.1™) interface, hardware debug, and a phase-locked loop (PLL) clock multiplier. These system support functions are described in the following sections.

### 7.1.7.1 Power Management

The e300 core provides four power modes, selectable by setting the appropriate control bits in the machine state register (MSR) and the hardware implementation register 0 (HID0). When entering into a power mode other than full-power, the core will request entry via a *qreq* signal and will only enter another power mode after an acknowledge (*qack*) is received. The four power modes are as follows:

- Full-power—This is the default power state of the e300 core. The e300 core is fully powered and the internal functional units are operating at the full processor clock speed. If the dynamic power management mode is enabled, functional units that are idle will automatically enter a low-power state without affecting performance, software execution, or external hardware.
- Doze—All the functional units of the e300 core are disabled except for the time base/decrementer registers and the bus snooping logic. When the processor is in doze mode, an external asynchronous interrupt, system management interrupt, decremter interrupt, hard or soft reset, or machine check brings the e300 core into the full-power state. The core in doze mode maintains the PLL in a fully-powered state and locked to the system external clock input (*sysclk*), so a transition to the full-power state takes only a few processor clock cycles.
- Nap—The nap mode further reduces power consumption by disabling bus snooping, leaving only the time base register and the PLL in a powered state. The core returns to the full-power state on receipt of an external asynchronous interrupt, system management interrupt, decremter interrupt, hard or soft reset, or machine check input (*mcp*) signal. A return to full-power state from a nap state takes only a few processor clock cycles.
- Sleep—Sleep mode reduces power consumption to a minimum by disabling all internal functional units; then external system logic may disable the PLL and *sysclk*. Returning the core to the full-power state requires the enabling of the PLL and *sysclk*, followed by the assertion of an external asynchronous interrupt, system management interrupt, hard or soft reset, or *mcp* signal after the time required to relock the PLL.

### 7.1.7.2 Time Base/Decrementer

The time base is a 64-bit register (accessed as two 32-bit registers) that is incremented once every four bus clock cycles; external control of the time base is provided through the time base/decrementer clock base enable (*tben*) signal. The decremter is a 32-bit register that generates a decremter interrupt after a programmable delay. The contents of the decremter register are decremented once every four bus clock cycles, and the decremter interrupt is generated as the count passes through zero.

### 7.1.7.3 JTAG Test and Debug Interface

The core provides JTAG and hardware debug functions for facilitating board testing and chip debugging. The JTAG test interface (based on IEEE 1149.1) provides a means for boundary-scan testing of the core and the attached system logic. The hardware debug function accesses the JTAG test port, providing a means for executing test routines and facilitating chip and software debugging.

All instruction and data address breakpoints are accessible in the IBCR and DBCR. See [Section 7.3.8, “Debug Features,”](#) for more information.

### 7.1.7.4 Clock Multiplier

The internal clocking of the e300 core is generated from and synchronized to the external clock signal, *sysclk*, by means of a voltage-controlled, oscillator-based PLL. The PLL provides programmable internal processor clock multiplier ratios which multiply the externally supplied clock frequency. The bus clock is the same frequency and is synchronous with *sysclk*. The configuration of the PLL can be read by software from the hardware implementation register 1 (HID1).

### 7.1.7.5 Core Performance Monitor

The performance monitor provides the ability to count predefined events and processor clocks associated with particular operations, such as cache misses, mispredicted branches, or the number of cycles an execution unit stalls. The count of such events can be used to trigger the performance monitor interrupt.

The performance monitor can be used to do the following:

- Improve system performance by monitoring software execution and then recoding algorithms for more efficiency. For example, memory hierarchy behavior can be monitored and analyzed to optimize task scheduling or data distribution algorithms.
- Characterize processors in environments not easily characterized by benchmarking.
- Help system developers bring up and debug their systems.

The performance monitor uses the following resources:

- The performance monitor mark bit in the MSR (MSR[PMM]). This bit controls which programs are monitored.
- The move to/from performance monitor registers (PMR) instructions, **mtpmr** and **mfpmr**.
- The external core input, *pm\_event\_in*.
- PMRs:
  - The performance monitor counter registers (PMC0–PMC3) are 32-bit counters used to count software-selectable events. Each counter counts up to 128 events. UPMC0–UPMC3 provide user-level read access to these registers. They are identified in [Table 7-2](#).
  - The performance monitor global control register (PMGC0) controls the counting of performance monitor events. It takes priority over all other performance monitor control registers. UPMGC0 provides user-level read access to PMGC0.

- The performance monitor local control registers (PMLCa0–PMLCa3) control each individual performance monitor counter. Each counter has a corresponding PMLCa register. UPMLCa0–UPMLCa3 provide user-level read access to PMLCa0–PMLCa3).
- The performance monitor interrupt is assigned to interrupt vector 0x0F00.

Software communication with the performance monitor is achieved through PMRs rather than SPRs. The PMRs are used for enabling conditions that can trigger the performance monitor interrupt.

## 7.2 PowerPC Architecture Implementation

The PowerPC architecture consists of the following layers, and adherence to the PowerPC architecture can be measured in terms of which of the following levels of the architecture is implemented:

- User instruction set architecture (UISA)—Defines the base user-level instruction set, user-level registers, data types, floating-point interrupt model, memory models for a uniprocessor environment, and programming model for a uniprocessor environment.
- Virtual environment architecture (VEA)—Describes the memory model for a multiprocessor environment, defines cache control instructions, and describes other aspects of virtual environments. Implementations that conform to the VEA also adhere to the UISA but may not necessarily adhere to the OEA.
- Operating environment architecture (OEA)—Defines the memory management model, supervisor-level registers, synchronization requirements, and interrupt model. Implementations that conform to the OEA also adhere to the UISA and VEA.

The PowerPC architecture allows a wide range of designs for such features as cache and core interface implementations.

## 7.3 Implementation-Specific Information

This section describes the PowerPC architecture in general and specific details about the implementation of the e300 core as a low-power, 32-bit member of this PowerPC core family. The main topics addressed are as follows:

- [Section 7.3.1, “Register Model,”](#) describes the registers for the operating environment architecture common among e300 cores that implement the PowerPC architecture and describes the programming model. It also describes the additional registers that are unique to the core.
- [Section 7.3.2, “Instruction Set and Addressing Modes,”](#) describes the PowerPC instruction set and addressing modes for the OEA, and defines and describes the instructions implemented in the core.
- [Section 7.3.3, “Cache Implementation,”](#) describes the cache model that is defined generally for cores that implement the PowerPC architecture by the VEA. It also provides specific details about the e300 core cache implementation.
- [Section 7.3.4, “Interrupt Model,”](#) describes the interrupt model of the OEA and the differences in the core interrupt model.
- [Section 7.3.5, “Memory Management,”](#) describes generally the conventions for memory management among these cores. This section also describes the core implementation of the 32-bit PowerPC memory management specification.

- [Section 7.3.6, “Instruction Timing,”](#) provides a general description of the instruction timing provided by the superscalar, parallel execution supported by the PowerPC architecture and the e300 core.
- [Section 7.1.6, “Bus Interface Unit \(BIU\),”](#) describes the signals implemented on the core.

The e300 core is a high-performance, superscalar processor core. The PowerPC architecture allows optimizing compilers to schedule instructions to maximize performance through efficient use of the PowerPC instruction set and register model. The multiple, independent execution units allow compilers to optimize instruction throughput. Compilers that take advantage of the flexibility of the PowerPC architecture can additionally optimize system performance.

The following sections summarize the features of the core, including both those that are defined by the architecture and those that are unique to the various core implementations.

Specific features of the core are listed in [Section 7.1.1, “Features.”](#)

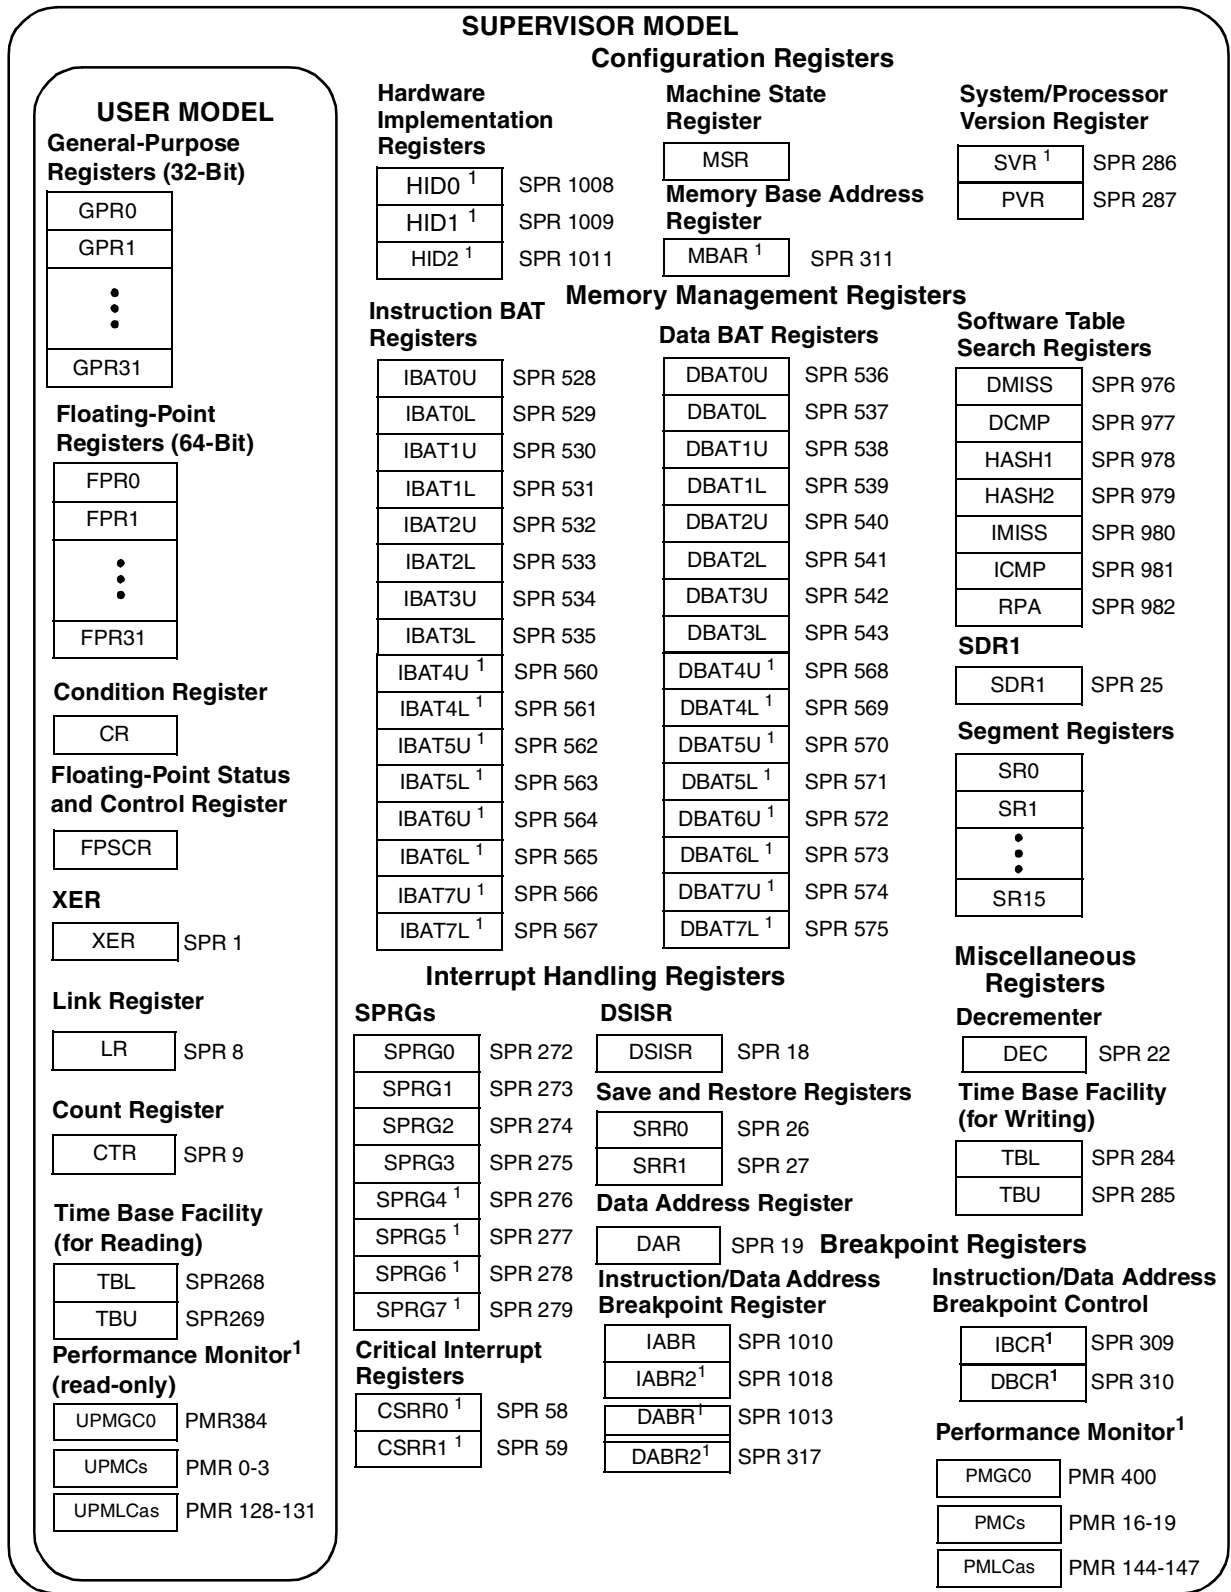
### 7.3.1 Register Model

The PowerPC architecture defines register-to-register operations for most computational instructions. Source operands for these instructions are accessed from the registers or are provided as immediate values embedded in the instruction opcode. The three-register instruction format allows specification of a target register distinct from the two-source operands. Load and store instructions transfer data between registers and memory.

The e300 core has two levels of privilege: supervisor mode of operation (typically used by the operating system) and user mode of operation (used by the application software). The programming models incorporate 32 GPRs, 32 FPRs, special-purpose registers (SPRs), and several miscellaneous registers. Each core also has its own unique set of hardware implementation (HID) registers.

Having access to privileged instructions, registers, and other resources allows the operating system to control the application environment (providing virtual memory and protecting operating system and critical machine resources). Instructions that control the state of the e300 core, the address translation mechanism, and supervisor registers can be executed only when the core is operating in supervisor mode.

[Figure 7-2](#) shows all the core registers available at the user and supervisor level. The numbers to the right of the SPRs indicate the number that is used in the syntax of the instruction operands for the move to/from SPR instructions.



<sup>1</sup> These registers are e300 core implementation-specific (not defined by the PowerPC architecture).

Figure 7-2. e300 Programming Model—Registers

The following sections describe the e300-core-implementation-specific features as they apply to registers.

### 7.3.1.1 UISA Registers

UISA registers are user-level registers that include the following.

#### 7.3.1.1.1 General-Purpose Registers (GPRs)

The PowerPC architecture defines 32 user-level GPRs that are 32 bits wide in 32-bit cores. The GPRs serve as the data source or destination for all integer instructions.

#### 7.3.1.1.2 Floating-Point Registers (FPRs)

The PowerPC architecture also defines 32 user-level, 64-bit FPRs. The FPRs serve as the data source or destination for floating-point instructions. These registers can contain data objects of either single- or double-precision floating-point formats.

#### 7.3.1.1.3 Condition Register (CR)

The CR is a 32-bit user-level register that provides a mechanism for testing and branching. It consists of eight 4-bit fields that reflect the results of certain operations, such as move, integer and floating-point comparisons, arithmetic, and logical operations.

#### 7.3.1.1.4 Floating-Point Status and Control Register (FPSCR)

The user-level FPSCR contains all floating-point exception signal bits, exception summary bits, exception enable bits, and rounding control bits needed for compliance with the IEEE 754 standard.

#### 7.3.1.1.5 User-Level SPRs

The PowerPC architecture defines numerous special purpose registers that serve a variety of functions, such as providing controls, indicating status, configuring the core, and performing special operations. During normal execution, a program can access the registers, as shown in [Figure 7-2](#), depending on the program's access privilege (supervisor or user, determined by the privilege-level bit, MSR[PR]). Note that GPRs and FPRs are accessed through operands that are part of the instructions. Access to registers can be explicit (that is, through the use of specific instructions for that purpose such as Move to Special-Purpose Register (**mtspr**) and Move from Special-Purpose Register (**mfspir**) instructions) or implicit, as the part of the execution of an instruction. Some registers are accessed both explicitly and implicitly. In the e300 core, all SPRs are 32 bits wide.

The following SPRs are accessible by user-level software:

- Link register (LR)—The LR can be used to provide the branch target address and to hold the return address after branch and link instructions. The LR is 32 bits wide in 32-bit implementations.
- Count register (CTR)—The CTR is decremented and tested automatically as a result of branch-and-count instructions. The CTR is 32 bits wide in 32-bit implementations.
- XER register—The 32-bit XER contains the summary overflow bit, integer carry bit, overflow bit, and a field specifying the number of bytes to be transferred by a Load String Word Indexed (**lswx**) or Store String Word Indexed (**stswx**) instruction.

### 7.3.1.2 VEA Registers

The VEA introduces the time base facility (TB) for reading. The TB is a 64-bit register pair whose contents are incremented once every four core input clock cycles. The TB consists of two 32-bit registers—time base upper (TBU) and time base lower (TBL). Note that the time base registers are read-only in user state.

### 7.3.1.3 OEA Registers

OEA registers are supervisor-level registers that include the following.

#### 7.3.1.3.1 Machine State Register (MSR)

The MSR is a supervisor-level register that defines the state of the core. The contents of this register are saved when an interrupt is taken, and restored when the interrupt handling completes. A critical interrupt is taken in the e300 core when the *cint* signal is asserted and MSR[CE] is set. The e300 core implements the MSR as a 32-bit register.

Table 7-1 shows the bit definitions for MSR.

**Table 7-1. MSR Bit Descriptions**

Bits	Name	Description
0 <sup>1</sup>	—	Reserved. Full function.
1–4 <sup>1</sup>	—	Reserved. Partial function.
5–9 <sup>1</sup>	—	Reserved. Full function.
10–12 <sup>1</sup>	—	Reserved. Partial function.
13	POW	Power management enable (implementation-specific) 0 Disables programmable power modes (normal operation mode) 1 Enables programmable power modes (nap, doze, or sleep mode). This bit controls the programmable power modes only; it has no effect on dynamic power management (DPM). MSR[POW] may be altered with an <b>mtmsr</b> instruction only. Also, when altering the POW bit, software may alter only this bit in the MSR and no others. The <b>mtmsr</b> instruction must be followed by a context-synchronizing instruction.
14	TGPR	Temporary GPR remapping (implementation-specific) 0 Normal operation 1 TGPR mode. GPR0–GPR3 are remapped to TGPR0–TGPR3 for use by TLB miss routines. The contents of GPR0–GPR3 remain unchanged while MSR[TGPR] = 1. Attempts to use GPR4–GPR31 with MSR[TGPR] = 1 yield undefined results. Temporarily replaces TGPR0–TGPR3 with GPR0–GPR3 for use by TLB miss routines. The TGPR bit is set when either an instruction TLB miss, data read miss, or data write miss interrupt is taken. The TGPR bit is cleared by an <b>rfi</b> instruction.
15	ILE	Interrupt little-endian mode. When an interrupt occurs, this bit is copied into MSR[LE] to select the endian mode for the context established by the interrupt.
16	EE	External interrupt enable 0 The processor ignores external interrupts, system management interrupts, and decremter interrupts. 1 The processor is enabled to take an external interrupt, system management interrupt, or decremter interrupt.
17	PR	Privilege level 0 The processor can execute both user- and supervisor-level instructions 1 The processor can only execute user-level instructions

Table 7-1. MSR Bit Descriptions (continued)

Bits	Name	Description
18	FP	Floating-point available 0 The processor prevents dispatch of floating-point instructions, including floating-point loads, stores, and moves. 1 The processor can execute floating-point instructions and can take floating-point enabled exception type program interrupts.
19	ME	Machine check enable 0 Machine check interrupts are disabled 1 Machine check interrupts are enabled
20	FE0	Floating-point exception mode 0
21	SE	Single-step trace enable 0 The processor executes instructions normally 1 The processor generates a trace interrupt upon the successful completion of the next instruction
22	BE	Branch trace enable 0 The processor executes branch instructions normally 1 The processor generates a trace interrupt upon the successful completion of a branch instruction
23	FE1	Floating-point exception mode 1
24	CE	Critical interrupt enable 0 Critical interrupts disabled 1 Critical interrupts enabled; critical interrupt and <b>rfci</b> instruction enabled The critical interrupt is an asynchronous implementation-specific interrupt. The critical interrupt vector offset is 0x00A00. The <b>rfci</b> instruction is implemented to return from these interrupt handlers. Also, CSRR0 and CSRR1 are used to save and restore the processor state for critical interrupts.
25	IP	Interrupt prefix. The setting of this bit specifies whether an interrupt vector offset is prepended with Fs or 0s. In the following description, <i>nnnn</i> is the offset of the interrupt. 0 Interrupts are vectored to the physical address 0x000 <i>n_nnnn</i> 1 Interrupts are vectored to the physical address 0xFFF <i>n_nnnn</i>
26	IR	Instruction address translation 0 Instruction address translation is disabled 1 Instruction address translation is enabled
27	DR	Data address translation 0 Data address translation is disabled 1 Data address translation is enabled
28–29 <sup>1</sup>	—	Reserved. Full function. Bit 29 not reserved on e300c3 .
29	PMM	Performance monitor mark bit (e300c3 ). System software can set PMM when a marked process is running to enable statistics to be gathered only during the execution of the marked process. MSR[PR] and MSR[PMM] together define a state that the processor (supervisor or user) and the process (marked or unmarked) may be in at any time. If this state matches an individual state specified in the PMLC <i>an</i> , the state for which monitoring is enabled, counting is enabled.
30	RI	Recoverable interrupt (for system reset and machine check interrupts) 0 Interrupt is not recoverable 1 Interrupt is recoverable
31	LE	Little-endian mode enable 0 The processor runs in big-endian mode 1 The processor runs in little-endian mode.

<sup>1</sup> All reserved bits should be set to zero for future compatibility.



### 7.3.1.3.2 Segment Registers (SRs)

For memory management, 32-bit processors implement sixteen 32-bit SRs. To speed access, the core implements the SRs as two arrays: a main array, for data memory accesses, and a shadow array, for instruction memory accesses. Loading a segment entry with the Move to Segment Register (**mtsr**) instruction loads both arrays.

### 7.3.1.3.3 Supervisor-Level SPRs

The e300 core, like the G2\_LE core, has additional supervisor-level SPRs, which are shown in [Figure 7-3](#). Two critical interrupt SPRs (CSRR0 and CSRR1), eight SPRGs (SPRG0–SPRG7), eight pairs of instruction BATs (IBAT0–IBAT7), eight pairs of data BATs (DBAT0–DBAT7), one system version register (SVR), one system memory base address (MBAR), one instruction address breakpoint control (IBCR), one data address breakpoint control (DBCR), a new instruction breakpoint register (IABR2), and two data address breakpoint registers (DABR and DABR2) are integrated into the core.

Supervisor-level SPRs include the following:

- The DSISR defines the cause of data access and alignment interrupts. The cause of a DSI interrupt for a data breakpoint (match with DABR and DABR2) can be determined by the value of the DSISR[DABR] bit (bit 9).
- The data address register (DAR) holds the address of an access after an alignment or DSI interrupt. For example, it contains the address of the breakpoint match condition.
- The decremter register (DEC) is a 32-bit decrementing counter that provides a mechanism for causing a decremter interrupt after a programmable delay.
- SDR1 specifies the page table format used in virtual-to-physical address translation for pages. (Note that physical address is referred to as ‘real address’ in the architecture specification.)
- The machine status save/restore register 0 (SRR0) is used for saving the address of the instruction that caused the interrupt, and the address to return to when a Return from Interrupt (**rfi**) instruction is executed.
- The machine status save/restore register 1 (SRR1) is used to save machine status on interrupts and to restore machine status when an **rfi** instruction is executed.
- The SPRG0–SPRG7 registers are provided for operating system use. They reduce the latency that may be incurred in the saving of registers to memory while in a handler. Note that the e300 implements four more SPRGs than the G2 (SPRG0–SPRG3).
- The time base register (TB) is a 64-bit register that maintains the time of day and operates interval timers. It consists of two 32-bit fields: time base upper (TBU) and time base lower (TBL).
- The processor version register (PVR) is a read-only register that identifies the version (model) and revision level of the processor. See [Table 7-8](#) for the version and revision level of the PVR for the e300 processor core.
- Block address translation (BAT) arrays—The PowerPC architecture defines 16 BAT registers. The e300 core includes a total of eight pairs of DBAT and eight pairs of IBAT registers. See [Figure 7-2](#) for a list of the SPR numbers for the BAT arrays.

The following supervisor-level SPRs are implementation-specific (not defined in the PowerPC architecture):

- DMISS and IMISS are read-only registers that are loaded automatically on an instruction or data TLB miss.
- HASH1 and HASH2 contain the physical addresses of the primary and secondary page table entry groups (PTEGs).
- ICMP and DCMP contain a duplicate of the first word in the page table entry (PTE) for which the table search is looking.
- The required physical address (RPA) register is loaded by the core with the second word of the correct PTE during a page table search.
- The system version register (SVR) is available on the e300 core, which identifies the specific version (model) and revision level of the system-on-a-chip (SOC) integration.
- System memory base address (MBAR) is an implementation-specific register available on the e300 core. It supports a temporary storage for the system-level memory map.
- The instruction and data address breakpoint registers (IABR, IABR2, DABR, DABR2) are loaded with an instruction or data address, respectively, that is compared to instruction addresses in the dispatch queue or to the data address in the LSU. When an address match occurs, a breakpoint interrupt is generated.
- One instruction breakpoint control register (IBCR) and one data breakpoint control register (DBCR) are implemented in the e300 core.
- To support critical interrupts, two registers (CSRR0 and CSRR1) are included in the e300 core.
- Eight SPRG registers (SPRG0–SPRG7) are in the e300 core.
- Block address translation (BAT) arrays—The e300 core has eight instruction and eight data BAT registers.
- The hardware implementation (HID0 and HID1) registers provide the means for enabling core checkstops and features and allow software to read the configuration of the PLL configuration signals. The HID2 register enables the true little-endian mode, cache way-locking, and the additional BAT registers.

Table 7-2 shows the bit definitions for HID0.

**Table 7-2. e300 HID0 Bit Descriptions**

Bits	Name	Function
0	EMCP	Enable $\overline{mcp}$ . The purpose of this bit is to mask out machine check interrupts caused by assertion of $\overline{mcp}$ , similar to how MSR[EE] can mask external interrupts. 0 Masks $\overline{mcp}$ . Asserting $\overline{mcp}$ does not generate a machine check interrupt or a checkstop. 1 Asserting $\overline{mcp}$ causes checkstop if MSR[ME] = 0 or a machine check interrupt if ME = 1
1	ECPE	Enable cache parity errors. 0 Disables instruction and data cache parity error reporting 1 Allows a detected cache parity error to cause a machine check interrupt if MSR[ME] = 1 or a checkstop if MSR[ME] = 0

Table 7-2. e300 HID0 Bit Descriptions (continued)

Bits	Name	Function
2	EBA	Enable $\overline{ap\_in}[0:3]$ and $\overline{ape}$ for address parity checking. 0 Disables address parity checking during a snoop operation 1 Allows an address parity error during snoop operations to cause a checkstop if MSR[ME] = 0 or a machine check interrupt if MSR[ME] = 1
3	EBD	Enable $\overline{dpe}$ for data parity checking. 0 Disables data parity checking 1 Allows a data parity error during reads to cause a checkstop if MSR[ME] = 0 or a machine check interrupt if MSR[ME] = 1
4	SBCLK	$clk\_out$ output enable. Used in conjunction with HID0[ECLK] and $\overline{hreset}$ to configure $clk\_out$ . See Table 7-3 for settings.
5	—	Reserved, should be cleared
6	ECLK	$clk\_out$ output enable. Used in conjunction with HID0[SBCLK] and the $\overline{hreset}$ signal to configure $clk\_out$ . See Table 7-3 for settings.
7	PAR	Disable precharge of $\overline{artry\_out}$ 0 Precharge of $\overline{artry\_out}$ enabled 1 Alters bus protocol slightly by preventing the processor from driving $\overline{artry\_out}$ to high (negated) state. If this is done, the integrated device must restore the signals to the high state.
8	DOZE	Doze mode enable. Operates in conjunction with MSR[POW]. 0 Doze mode disabled 1 Doze mode enabled. Doze mode is invoked by setting MSR[POW] while this bit is set. In doze mode, the PLL, time base, and snooping remain active.
9	NAP	Nap mode enable. Operates in conjunction with MSR[POW]. 0 Nap mode disabled 1 Nap mode enabled. Nap mode is invoked by setting MSR[POW] while this bit is set. In nap mode, the PLL and time base remain active.
10	SLEEP	Sleep mode enable. Operates in conjunction with MSR[POW]. 0 Sleep mode disabled 1 Sleep mode enabled. Sleep mode is invoked by setting MSR[POW] while this bit is set. $\overline{qreq}$ is asserted to indicate that the processor is ready to enter sleep mode. If the system logic determines that the processor may enter sleep mode, the quiesce acknowledge signal, $\overline{qack}$ , is asserted back to the processor. Once $\overline{qack}$ assertion is detected, the processor enters sleep mode after several processor clocks. At this point, the system logic may turn off the PLL by first configuring $pll\_cfg[0:6]$ to PLL bypass mode, then disabling $sysclk$ .
11	DPM	Dynamic power management enable 0 Dynamic power management is disabled 1 Functional units enter a low-power mode automatically if the unit is idle. This does not affect operational performance and is transparent to software or any external hardware.
12–15	—	Reserved, should be cleared.
16	ICE	Instruction cache enable 0 The instruction cache is neither accessed nor updated. All pages are accessed as if they were marked cache-inhibited (WIM = x1x). Potential cache accesses from the bus (snoop and cache operations) are ignored. In the disabled state for the L1 caches, the cache tag state bits are ignored and all instruction fetches are propagated to the coherent system bus (CSB) as single-beat transactions. For those transactions, however, $\overline{ci}$ reflects the state of the I bit in the MMU for that page regardless of cache disabled status. ICE is zero at power-up. 1 The instruction cache is enabled

Table 7-2. e300 HID0 Bit Descriptions (continued)

Bits	Name	Function
17	DCE	<p>Data cache enable</p> <p>0 The data cache is neither accessed nor updated. All pages are accessed as if they were marked cache-inhibited (WIM = x1x). Potential cache accesses from the bus (snoop and cache operations) are ignored. In the disabled state for the L1 caches, the cache tag state bits are ignored and all data read and write accesses are propagated to the CSB as single-beat transactions. For those transactions, however, <math>\bar{c}_i</math> reflects the state of the I bit in the MMU for that page regardless of cache disabled status. DCE is zero at power-up.</p> <p>1 The data cache is enabled</p>
18	ILOCK	<p>Instruction cache lock</p> <p>0 Normal operation</p> <p>1 The entire instruction cache is locked (that is, all eight ways of the cache are locked). A locked cache supplies data normally on a hit, but the access is treated as a cache-inhibited transaction on a miss. On a miss, the transaction to the bus is single-beat; however, <math>\bar{c}_i</math> still reflects the state of the I bit in the MMU for that page independent of cache locked or disabled status.</p> <p>To prevent locking during a cache access, an <b>isync</b> instruction must precede the setting of ILOCK.</p>
19	DLOCK	<p>Data cache lock</p> <p>0 Normal operation</p> <p>1 The entire data cache is locked (that is, all eight ways of the cache are locked). A locked cache supplies data normally on a hit, but is treated as a cache-inhibited transaction on a miss. On a miss, the transaction to the bus is single-beat; however, <math>\bar{c}_i</math> still reflects the state of the I bit in the MMU for that page independent of cache locked or disabled status. A snoop hit to a locked L1 data cache performs as if the cache were not locked. A cache block invalidated by a snoop remains invalid until the cache is unlocked.</p> <p>To prevent locking during a cache access, a <b>sync</b> instruction must precede the setting of DLOCK.</p>
20	ICFI	<p>Instruction cache Flash invalidate</p> <p>0 The instruction cache is not invalidated. The bit is cleared when the invalidation operation begins (usually the next cycle after the write operation to the register). The instruction cache must be enabled for the invalidation to occur.</p> <p>1 An invalidate operation is issued that marks the state of each instruction cache block as invalid. Cache access is blocked during this time. Setting ICFI clears all the valid bits of the blocks and the PLRU bits to point to way L0 of each set.</p> <p>For the e300 core, the proper use of the ICFI and DCFI bits is to set and clear them with two consecutive <b>mtspr</b> operations.</p>
21	DCFI	<p>Data cache Flash invalidate</p> <p>0 The data cache is not invalidated. The bit is cleared when the invalidation operation begins (usually the next cycle after the write operation to the register). The data cache must be enabled for the invalidation to occur.</p> <p>1 An invalidate operation is issued that marks the state of each data cache block as invalid without writing back modified cache blocks to memory. Cache access is blocked during this time. Bus accesses to the cache are signaled as a miss during invalidate-all operations. Setting DCFI clears all the valid bits of the blocks and the PLRU bits to point to way L0 of each set.</p> <p>For the e300 core, the proper use of the ICFI and DCFI bits is to set and clear them with two consecutive <b>mtspr</b> operations.</p>
22–23	—	Reserved, should be cleared.
24	IFEM	<p>Enable M bit on bus for instruction fetches</p> <p>0 M bit not reflected on bus for instruction fetches. Instruction fetches are treated as nonglobal on the bus.</p> <p>1 Instruction fetches reflect the M bit from the WIM settings</p>

Table 7-2. e300 HID0 Bit Descriptions (continued)

Bits	Name	Function
25	DECAREN	Decrementer auto reload 0 Normal operation. 1 Decrementer loads last mtdec value for precise periodic interrupt.
26	—	Reserved, should be cleared.
27	FBIOB	Force branch indirect on the bus 0 Register indirect branch targets are fetched normally 1 Forces register indirect branch targets to be fetched externally
28	ABE	Address broadcast enable. Controls whether certain address-only operations (such as cache operations) are broadcast on the bus. 0 Address-only operations affect only local caches and are not broadcast 1 Address-only operations are broadcast on the bus Affected instructions are <b>dcbi</b> , <b>dcbf</b> , and <b>dcbst</b> . Note that these cache control instruction broadcasts are not snooped by the e300 core. Refer to Section 4.3.3, “Data Cache Control,” for more information.
29–30	—	Reserved, should be cleared.
31	NOOPTI	No-op the data cache touch instructions 0 The <b>dcbt</b> and <b>dcbst</b> instructions are enabled 1 The <b>dcbt</b> and <b>dcbst</b> instructions are no-oped internal to the e300 core

Table 7-3 shows how HID0[ECLK] and HID0[SBCLK] are used to configure the *clk\_out* signal.

Table 7-3. Using HID0[ECLK] and HID0[SBCLK] to Configure *clk\_out*

$\overline{hreset}$	ECLK	SBCLK	<i>clk_out</i>
Asserted	x	x	Bus clock (small pulse for every rising edge of sysclk)
Negated	0	0	Clock output off
	0	1	Core clock/2
	1	0	Core clock
	1	1	Bus clock

Table 7-4 shows the bit definitions for HID1

Table 7-4. HID1 Bit Descriptions

Bits	Name	Description
0	PC0	PLL configuration bit 0 (read-only)
1	PC1	PLL configuration bit 1 (read-only)
2	PC2	PLL configuration bit 2 (read-only)
3	PC3	PLL configuration bit 3 (read-only)
4	PC4	PLL configuration bit 4 (read-only)
5	PC5	PLL configuration bit 5 (read-only)
6	PC6	PLL configuration bit 6 (read-only)

Table 7-4. HID1 Bit Descriptions (continued)

Bits	Name	Description
7–31	—	Reserved, should be cleared

**Note:** The clock configuration bits reflect the state of the *pll\_cfg[0:6]* signals.

Table 7-5 shows the bit definitions for HID2.

Table 7-5. e300HID2 Bit Descriptions

Bits	Name	Description
0–3	—	Reserved, should be cleared.
4	LET	True little-endian. This bit enables true little-endian mode operation for instruction and data accesses. This bit is set to reflect the state of the <i>tle</i> signal at the negation of <i>hreset</i> . This bit is used in conjunction with MSR[LE] to determine the endian mode of operation. 0 No function 1 True little-endian mode, when MSR[LE] = 1 Changing the value of this bit during normal operation is not recommended
5	IFEB	Instruction fetch burst extension. This bit enables the instruction fetch burst extension. 0 Instruction fetch burst extension disabled 1 Instruction fetch burst extension enabled
6	—	Reserved, should be cleared.
7	MESISTATE	MESI state enable. This bit enables the four-state MESI cache coherency protocol. 0 MESI disabled. The data cache uses a three-state MEI coherency protocol. 1 MESI enabled. The data cache uses a four-state MESI protocol.
8	IFEC	Instruction fetch cancel extension. This bit enables the instruction fetch cancel extension. 0 Instruction fetch cancel extension disabled 1 Instruction fetch cancel extension enabled
9	EBQS	Enable BIU queue sharing. This bit enables data cache queue sharing. 0 Data cache queue sharing disabled 1 Data cache queue sharing enabled
10	EBPX	Enable BIU pipeline extension. This bit enables the bus interface unit pipeline extension. 0 BIU pipeline extension disabled; 1 level pipeline 1 BIU pipeline extension enabled; 1-1/2 level pipeline
11–12	—	Reserved for e300c1, should be cleared.
11	ELRW	Enable weighted LRU. This bit enables the use of an adjusted (weighted) LRU. 0 Normal operation. 1 The <i>dcbt</i> , <i>dcbstst</i> , and <i>dcbz</i> instructions use an adjusted (weighted) LRU such that they always select and replace the lowest unlocked way in the data cache.
12	NOKS	No kill for snoop. This bit enables the forcing of kill-type snoops to flush data instead of killing it. 0 Normal operation. 1 Forces write-with-kill snoops to flush instead of kill (snoop can never kill data).
13	HBE	High BAT enable. Regardless of the setting of HID2[HBE], these BATs are accessible by <b>mf spr</b> and <b>mt spr</b> . 0 IBAT[4–7] and DBAT[4–7] are disabled 1 IBAT[4–7] and DBAT[4–7] are enabled
14–15	—	Reserved, should be cleared.

Table 7-5. e300HID2 Bit Descriptions (continued)

Bits	Name	Description
16–18	IWLCK[0–2]	Instruction cache way-lock. Useful for locking blocks of instructions into the instruction cache for time-critical applications that require deterministic behavior. 000 no ways locked 001 way 0 locked 010 way 0 through way 1 locked 011 way 0 through way 2 locked 100 way 0 through way 2 locked in e300c3. 101 way 0 through way 2 locked in e300c3. 110 way 0 through way 2 locked in e300c3. 111 way 0 through way 2 locked in e300c3. Setting HID0[ILOCK] will lock all ways.
19	ICWP	Instruction cache way protection. Used to protect locked ways in the instruction cache from being invalidated. 0 Instruction cache way protection disabled 1 Instruction cache way protection enabled
20–23	—	Reserved, should be cleared.
24–26	DWLCK[0–2]	Data cache way-lock. Useful for locking blocks of data into the data cache for time-critical applications where deterministic behavior is required. 000 no ways locked 001 way 0 locked 010 way 0 through way 1 locked 011 way 0 through way 2 locked 100 way 0 through way 2 locked in e300c3. 101 way 0 through way 2 locked in e300c3. 110 way 0 through way 2 locked in e300c3. 111 way 0 through way 2 locked in e300c3. Setting HID0[DLOCK] will lock all ways.
27–31	—	Reserved, should be cleared.

## 7.3.2 Instruction Set and Addressing Modes

The following sections describe the PowerPC instruction set and addressing modes in general.

### 7.3.2.1 PowerPC Instruction Set and Addressing Modes

All PowerPC instructions are encoded as single-word (32-bit) opcodes. Instruction formats are consistent among all instruction types, permitting efficient decoding to occur in parallel with operand accesses. This fixed instruction length and consistent format simplifies instruction pipelining.

The PowerPC instructions are divided into the following categories:

- Integer instructions—These include computational and logical instructions.
  - Integer arithmetic instructions
  - Integer compare instructions
  - Integer logical instructions
  - Integer rotate and shift instructions
- Floating-point instructions—These include floating-point computational instructions, as well as instructions that affect the FPSCR.
  - Floating-point arithmetic instructions
  - Floating-point multiply/add instructions
  - Floating-point rounding and conversion instructions
  - Floating-point compare instructions
  - Floating-point status and control instructions
- Load/store instructions—These include integer and floating-point load and store instructions.
  - Integer load and store instructions
  - Integer load and store multiple instructions
  - Floating-point load and store
  - Primitives used to construct atomic memory operations (**lwarx** and **stwx**. instructions)
- Flow control instructions—These include branching instructions, condition register logical instructions, trap instructions, and other instructions that affect the instruction flow.
  - Branch and trap instructions
  - Condition register logical instructions
- Processor control instructions—These instructions are used for synchronizing memory accesses and management of caches, TLBs, and the segment registers.
  - Move to/from SPR instructions
  - Move to/from MSR
  - Move to/from PMR
  - Synchronize
  - Instruction synchronize



- Memory control instructions—These instructions provide control of caches, TLBs, and segment registers.
  - Supervisor-level cache management instructions
  - Translation lookaside buffer management instructions. Note that there are additional implementation-specific instructions.
  - User-level cache instructions
  - Segment register manipulation instructions
- The e300 core implements the following instructions which are defined as optional by the PowerPC architecture:
  - Floating Select (**fsel**)
  - Floating Reciprocal Estimate Single-Precision (**fres**)
  - Floating Reciprocal Square Root Estimate (**frsqrte**)
  - Store Floating-Point as Integer Word (**stfiwx**)

Note that this grouping of instructions does not indicate the execution unit that executes a particular instruction or group of instructions.

Integer instructions operate on byte, half-word, and word operands. Floating-point instructions operate on single-precision (one word) and double-precision (one double word) floating-point operands. The PowerPC architecture uses instructions that are 4 bytes long and word-aligned. It provides for byte, half-word, and word operand loads and stores between memory and a set of 32 GPRs. It also provides for word and double-word operand loads and stores between memory and a set of 32 FPRs.

Computational instructions do not modify memory. To use a memory operand in a computation and then modify the same or another memory location, the memory contents must be loaded into a register, modified, and then written back to the target location with distinct instructions.

The core follows the program flow when it is in the normal execution state. However, the flow of instructions can be interrupted directly by the execution of an instruction or by an asynchronous event. Either kind of interrupt may cause one of several components of the system software to be invoked.

### 7.3.2.2 Implementation-Specific Instruction Set

The e300 core instruction set is defined as follows:

- The core provides hardware support for all 32-bit PowerPC instructions.
- The core provides two implementation-specific instructions used for software table search operations following TLB misses:
  - Load Data TLB Entry (**tlbld**)
  - Load Instruction TLB Entry (**tlbli**)
- The core implements the following instruction which is added to support critical interrupts (also supported on the G2\_LE). This is a supervisor-level, context synchronizing instruction.
  - Return from Critical Interrupt (**rftci**)

- The core implements the following instruction which is added to support easy start-up initialization or reloading of the instruction cache.
  - Instruction Cache Block Touch (**icbt**)
- The core provides the following performance monitor instructions:
  - Move to Performance Monitor Register (**mtpmr**)
  - Move from Performance Monitor Register (**mfpmr**)

### 7.3.3 Cache Implementation

The following sections describe the general cache characteristics as implemented in the PowerPC architecture and the core implementation.

#### 7.3.3.1 PowerPC Cache Characteristics

The PowerPC architecture does not define hardware aspects of cache implementations. The e300 core controls the following memory access modes on a page or block basis:

- Write-back/write-through mode
- Caching-inhibited mode
- Memory coherency

Note that in the core, a cache block is defined as eight words. The VEA defines cache management instructions that provide a means by which the application programmer can affect the cache contents.


#### 7.3.3.2 Implementation-Specific Cache Organization

The e300c3 provides 16-Kbyte, four-way set-associative instruction and data caches. The caches are physically addressed, and the data cache can operate in either write-back or write-through mode as specified by the PowerPC architecture.

The data cache is configured as 128 sets of 4 blocks each on the e300c3. Each block consists of 32 bytes, 2 state bits, and an address tag. The two state bits implement the three-state MEI (modified/exclusive/invalid) protocol. Each block contains eight 32-bit words. Note that the PowerPC architecture defines the term ‘block’ as the cacheable unit. For the core, the block size is equivalent to a cache line. A block diagram of the data cache organization is shown in [Figure 7-3](#).

The instruction cache is configured as 128 sets of 4 blocks each on the e300c3. Each block consists of 32 bytes, an address tag, and a valid bit. The instruction cache may not be written to, except through a block fill operation. In the e300 core, the instruction cache is blocked only until the critical load completes. The e300 core supports instruction fetching from other instruction cache lines following the forwarding of the critical-first-double-word of a cache line load operation. Successive instruction fetches from the cache line being loaded are forwarded, and accesses to other instruction cache lines can proceed during the cache line load operation. The instruction cache is not snooped, and cache coherency must be maintained by software. A fast hardware invalidation capability is provided to support cache maintenance.

The e300c3 data cache is configured as 128 sets of four blocks per set. The organization of the data cache is shown in [Figure 7-3](#).



**Figure 7-3. e300c3 Data Cache Organization**

Each cache block contains eight contiguous words from memory that are loaded from an 8-word boundary (that is, bits A[27–31] of the effective addresses are zero); thus, a cache block never crosses a page boundary. Misaligned accesses across a page boundary can incur a performance penalty.

The e300 core cache blocks are loaded in four beats of 64 bits each on the 64-bit data bus. The burst load is performed as critical-double-word-first. The data cache is blocked to internal accesses until the load completes; the instruction cache allows sequential fetching during a cache block load. In the core, the critical-double-word is simultaneously written to the cache and forwarded to the requesting unit, thus minimizing stalls due to load delays.

To ensure coherency among caches in a multiprocessor (or multiple caching-device) implementation, the core implements the MEI protocol during normal operation of the data cache. The new data cache MESI extension supports the additional fourth cache coherency shared state for the data cache. To support this feature, the shared signal, *shd*, has been added to the bus interface. Although the MESI protocol is supported by the e300 core, it is not implemented on MPC8313E. The following four states indicate the state of the cache block:

- **Modified**—The cache block is modified with respect to system memory; that is, data for this address is valid only in the cache and not in system memory.
- **Exclusive**—This cache block holds valid data that is identical to the data at this address in system memory. No other cache has this data.
- **Shared**—Only available if HID2[MESISTATE] register bit is set. The address block is valid in the cache and in at least one other cache. This block is always consistent with system memory. That is, the shared state is shared-unmodified; there is no shared-modified state. Although the MESI protocol is supported by the e300 core, it is not implemented on MPC8313E.
- **Invalid**—This cache block does not hold valid data.

Cache coherency is enforced by on-chip bus snooping logic. Because the e300 core data cache tags are single-ported, a simultaneous load/store and snoop access represents a resource contention. The snoop access is given first access to the tags. The load or store then occurs on the clock following the snoop.

Parity is now integrated into both instruction and data cache memory. A machine check interrupt is now taken upon the detection of an instruction or data cache parity error. Parity is checked whenever valid data is returned from the instruction or data cache for a cache hit or whenever valid data is read out of the cache for a castout or snoop-push operation.

### 7.3.3.3 Instruction and Data Cache Way-Locking

The e300 core implements instruction and data cache way-locking, which guarantees that certain memory accesses will hit in the cache. This provides deterministic access times for those accesses.

## 7.3.4 Interrupt Model

This section describes the PowerPC interrupt model and the e300 core implementation specifically.

### 7.3.4.1 PowerPC Interrupt Model

The PowerPC interrupt mechanism allows the core to change to supervisor state as a result of external signals, errors, or unusual conditions arising in the execution of instructions. The conditions that can cause interrupts are called exceptions. When interrupts occur, information about the state of the core is saved to certain registers and the core begins execution at an address (interrupt vector) predetermined for each interrupt type. Interrupts are processed in supervisor mode.

Some interrupts, such as program interrupts, can be triggered by a broad range of exception conditions. Other interrupts, such as the decremter interrupt, have only a single exception condition. Although multiple exception conditions can map to a single interrupt vector, a more specific condition may be determined by examining a register associated with the interrupt—for example, the DSISR and the FPSCR. Additionally, some exception conditions can be explicitly enabled or disabled by software.

The PowerPC architecture requires that interrupts be handled in program order; therefore, although a particular implementation may recognize exception conditions out of order, they are presented strictly in order. When an instruction-caused interrupt is recognized, any unexecuted instructions that appear earlier in the instruction stream, including any that have not yet entered the execute stage, are required to complete before the interrupt is taken. Any interrupts caused by those instructions are handled first. Likewise, asynchronous, precise interrupts are recognized when they occur, but are not handled until the instruction currently in the completion stage successfully completes execution or generates an interrupt, and the completed store queue is emptied.

Unless a catastrophic condition causes a system reset or machine check interrupt, only one interrupt is handled at a time. If, for example, a single instruction encounters multiple interrupt conditions, those conditions are handled sequentially. After the interrupt handler completes, the instruction execution continues until the next interrupt condition is encountered. However, in many cases there is no attempt to re-execute the instruction. This method of recognizing and handling interrupts sequentially guarantees that interrupts are recoverable.

To prevent the program state from being lost due to a system reset, a machine check interrupt, or an instruction-caused interrupt in the interrupt handler, interrupt handlers should save the information stored in SRR0 and SRR1 early and before enabling external interrupts.

The PowerPC architecture supports four types of interrupts:

- Synchronous, precise—These are caused by instructions. All instruction-caused interrupts are handled precisely; that is, the machine state at the time the interrupt occurs is known and can be completely restored. This means that (excluding the trap and system call interrupts) the address of the faulting instruction is provided to the interrupt handler and neither the faulting instruction nor subsequent instructions in the code stream will complete execution before the interrupt is taken. Once the interrupt is processed, execution resumes at the address of the faulting instruction (or at an alternate address provided by the interrupt handler). When an interrupt is taken due to a trap or system call instruction, execution resumes at an address provided by the handler.
- Synchronous, imprecise—The PowerPC architecture defines two imprecise floating-point exception modes: recoverable and nonrecoverable. Even though the core provides a means to enable the imprecise modes, it implements these modes identically to the precise mode (that is, all enabled floating-point exceptions are always precise on the core).
- Asynchronous, maskable—The external system management interrupt (SMI) and decremter interrupts are maskable, asynchronous interrupts. When these interrupts occur, their handling is postponed until the next instruction and any of its associated interrupts complete execution. If there are no instructions in the execution units, the interrupt is taken immediately upon determination of the correct restart address (for loading SRR0).
- Asynchronous, nonmaskable—The system reset and the machine check interrupt are nonmaskable, asynchronous interrupts. They may not be recoverable, or they may provide a limited degree of recoverability. All interrupts report recoverability through MSR[RI].

### 7.3.4.2 Implementation-Specific Interrupt Model

As specified by the PowerPC architecture, all interrupts can be described as either precise or imprecise and either synchronous or asynchronous. Asynchronous interrupts (some of which are maskable) are caused by events external to the processor's execution; synchronous interrupts, which are all handled precisely by the e300 core, are caused by instructions. A system management interrupt is an implementation-specific interrupt. The interrupt classes are shown in [Table 7-6](#).

**Table 7-6. Interrupt Classifications**

Synchronous/Asynchronous	Precise/Imprecise	Interrupt Type
Asynchronous, nonmaskable	Imprecise	Machine check System reset
Asynchronous, maskable	Precise	External interrupt Decrementer System management interrupt Critical interrupt
Synchronous	Precise	Instruction-caused interrupts

Although interrupts have other characteristics, such as whether they are maskable, the distinctions shown in [Table 7-6](#) define categories of interrupts that the core handles uniquely. Note that [Table 7-6](#) includes no synchronous, imprecise instructions. While the PowerPC architecture supports imprecise handling of floating-point exceptions, the core implements floating-point exception modes as precise.

The e300 core interrupts and exception conditions that cause them are listed in [Table 7-7](#).

**Table 7-7. Exceptions and Interrupts**

Interrupt Type	Vector Offset (hex)	Exception Conditions
Reserved	00000	—
System reset	00100	Caused by the assertion of either $\overline{hreset}$ .
Machine check	00200	Caused by the assertion of the $\overline{tea}$ signal during a data bus transaction, assertion of $\overline{mcp}$ , an address or data parity error, or an instruction or data cache parity error. Note that the e300 has SRR1 register values that are different from the G2/G2_LE cores' when a machine check occurs.
DSI	00300	Determined by the bit settings in the DSISR, listed as follows: <ul style="list-style-type: none"> <li>1 Set if the translation of an attempted access is not found in the primary hash table entry group (HTEG), or in the rehashed secondary HTEG, or in the range of a DBAT register; otherwise cleared</li> <li>4 Set if a memory access is not permitted by the page or DBAT protection mechanism; otherwise cleared</li> <li>6 Set for a store operation and cleared for a load operation</li> <li>9 Set if a data address breakpoint interrupt occurs when the data [0–28] in the DABR or DABR2 matches the next data access (load or store instruction) to complete in the completion unit. The different breakpoints are enabled as follows: <ul style="list-style-type: none"> <li>• Write breakpoints enabled when DABR[30] is set</li> <li>• Read breakpoints enabled when DABR[31] is set</li> </ul> </li> </ul>
ISI	00400	Caused when an instruction fetch cannot be performed for any of the following reasons: <ul style="list-style-type: none"> <li>• The effective (logical) address cannot be translated. That is, there is a page fault for this portion of the translation, so an ISI interrupt must be taken to load the PTE (and possibly the page) into memory.</li> <li>• The fetch access violates memory protection (indicated by SRR1[4] set). If the key bits (Ks and Kp) in the segment register and the PP bits in the PTE are set to prohibit read access, instructions cannot be fetched from this location.</li> </ul>
External interrupt	00500	Caused when MSR[EE] = 1 and the $\overline{int}$ signal is asserted.
Alignment	00600	Caused when the core cannot perform a memory access for any of the reasons described below: <ul style="list-style-type: none"> <li>• The operand of a floating-point load or store instruction is not word-aligned.</li> <li>• The operands of <b>lmw</b>, <b>stmw</b>, <b>lwarx</b>, and <b>stwcx</b>. instructions are not aligned.</li> <li>• The instruction is <b>lswi</b>, <b>lswx</b>, <b>stswi</b>, <b>stswx</b>, and the core is in little-endian mode. Note that PowerPC little-endian mode is not supported on the e300 core.</li> <li>• The operand of <b>dcbz</b> is in memory that is write-through-required or caching-inhibited.</li> </ul>

Table 7-7. Exceptions and Interrupts (continued)

Interrupt Type	Vector Offset (hex)	Exception Conditions
Program	00700	<p>Caused by one of the following exception conditions, which correspond to bit settings in SRR1 and arise during execution of an instruction.</p> <p>Floating-point enabled exception—A floating-point enabled exception condition is generated when the following condition is met: (MSR[FE0]   MSR[FE1]) and FPSCR[FEX] is 1.</p> <ul style="list-style-type: none"> <li>FPSCR[FEX] is set by the execution of a floating-point instruction that causes an enabled exception or by the execution of one of the Move to FPSCR instructions that results in both an exception condition bit and its corresponding enable bit being set in the FPSCR.</li> <li>Illegal instruction—An illegal instruction program interrupt is generated when execution of an instruction is attempted with an illegal opcode or illegal combination of opcode and extended opcode fields (including PowerPC instructions not implemented in the core), or when execution of an optional instruction not provided in the core is attempted (these do not include those optional instructions that are treated as no-ops).</li> <li>Privileged instruction—A privileged instruction program interrupt is generated when the execution of a privileged instruction is attempted and the MSR register user privilege bit, MSR[PR], is set. In the e300 core, this interrupt is generated for <b>mtspr</b> or <b>mfspir</b> with an invalid SPR field if SPR[0] = 1 and MSR[PR] = 1. This may not be true for all cores that implement the PowerPC architecture.</li> <li>Trap—A trap type program interrupt is generated when any of the conditions specified in a trap instruction are met.</li> </ul>
Floating-point unavailable	00800	Caused by an attempt to execute a floating-point instruction (including floating-point load, store, and move instructions) when the floating-point available bit (MSR[FP]) is cleared.
Decrementer	00900	Occurs when DEC[0] changes from 0 to 1. This interrupt is enabled with MSR[EE].
Critical interrupt	00A00	Taken when $\overline{cint}$ is asserted and MSR[CE] = 1.
Reserved	00B00–00BFF	—
System call	00C00	Occurs when a System Call ( <b>sc</b> ) instruction is executed.
Trace	00D00	Taken when MSR[SE] = 1 or when the currently completing instruction is a branch and MSR[BE] = 1.
Reserved	00E00	The e300 core does not generate an interrupt to this vector. Other devices may use this vector for floating-point assist interrupts.
Performance monitor	00F00	Caused when a configured PM counter using the pm_event_in to transition overflows.
Instruction translation miss	01000	Caused when the effective address for an instruction fetch cannot be translated by the ITLB.
Data load translation miss	01100	Caused when the effective address for a data load operation cannot be translated by the DTLB.
Data store translation miss	01200	Caused when the effective address for a data store operation cannot be translated by the DTLB, or when a DTLB hit occurs and the change bit in the PTE must be set due to a data store operation.
Instruction address breakpoint	01300	Occurs when the address (bits 0–29) in the IABR matches the next instruction to complete in the completion unit, and IABR[30] is set. Note that the e300 core also implements IABR2, which functions identically to IABR.

Table 7-7. Exceptions and Interrupts (continued)

Interrupt Type	Vector Offset (hex)	Exception Conditions
System management interrupt	01400	Caused when MSR[EE] = 1 and the $\overline{smi}$ input signal is asserted.
Reserved	01500–02FFF	—

## 7.3.5 Memory Management

The following sections describe the memory management features of the PowerPC architecture and the e300 core implementation, respectively.

### 7.3.5.1 PowerPC Memory Management

The primary functions of the MMU are to translate logical (effective) addresses to physical addresses for memory accesses and to provide access protection on blocks and pages of memory.

The core generates two types of accesses that require address translation: instruction accesses and data accesses to memory generated by load and store instructions.

The PowerPC MMU and interrupt model support demand-paged virtual memory. Virtual memory management permits execution of programs larger than the size of physical memory; demand-paged implies that individual pages are loaded into physical memory from system memory only when they are first accessed by an executing program.

The hashed page table is a variable-sized data structure that defines the mapping between virtual page numbers and physical page numbers. The page table size is a power of two, and its starting address is a multiple of its size.

The page table contains a number of page-table entry groups (PTEGs). A PTEG contains eight page-table entries (PTEs) of 8 bytes each; therefore, each PTEG is 64 bytes long. PTEG addresses are entry points for table search operations.

Address translations are enabled by setting bits in the MSR—MSR[IR] enables instruction address translations, and MSR[DR] enables data address translations.

### 7.3.5.2 Implementation-Specific Memory Management

The instruction and data memory management units in the e300 core provide 4 Gbytes of logical address space accessible to supervisor and user programs with a 4-Kbyte page size and 256-Mbyte segment size. Block sizes range from 128 Kbytes to 256 Mbytes and are software selectable. In addition, the core uses an interim 52-bit virtual address and hashed page tables for generating 32-bit physical addresses. The MMUs in the e300 core rely on the interrupt processing mechanism for the implementation of the paged virtual memory environment and for enforcing protection of designated memory areas.

Instruction and data TLBs provide address translation in parallel with the on-chip cache access, incurring no additional time penalty in the event of a TLB hit. A TLB is a cache of the most recently used page table



entries. Software is responsible for maintaining the consistency of the TLB with memory. The core TLBs are 64-entry, two-way, set-associative caches that contain instruction and data address translations. The core provides hardware assist for software table search operations through the hashed page table on TLB misses. Supervisor software can invalidate TLB entries selectively.

For instructions and data that correspond to block address translation, the e300 core provides independent eight-entry BAT arrays. These entries define blocks that can vary from 128 Kbytes to 256 Mbytes. The BAT arrays are maintained by system software. HID2[HBE] is added to the e300 for enabling or disabling the four additional pairs of BAT registers. However, regardless of the setting of HID2[HBE], these BATs are accessible by **mfspr** and **mtspr**.

As specified by the PowerPC architecture, the hashed page table is a variable-sized data structure that defines the mapping between virtual page numbers and physical page numbers. The page table size is a power of two, and its starting address is a multiple of its size.

Also as specified by the PowerPC architecture, the page table contains a number of PTEGs. A PTEG contains 8 PTEs of 8 bytes each; therefore, each PTEG is 64 bytes long. PTEG addresses are entry points for table search operations.

### 7.3.6 Instruction Timing

The e300 core is a pipelined superscalar processor core. Because instruction processing is reduced into a series of stages, an instruction does not require all of the resources of an execution unit at the same time. For example, after an instruction completes the decode stage, it can pass on to the next stage, while the subsequent instruction can advance into the decode stage. This improves the throughput of the instruction flow. For example, it may take three cycles for a single floating-point instruction to execute, but if there are no stalls in the floating-point pipeline, a series of floating-point instructions can have a throughput of one instruction per cycle.

The core instruction pipeline has four major pipeline stages, described as follows:

- The fetch pipeline stage primarily involves retrieving instructions from the memory system and determining the location of the next instruction fetch. Additionally, if possible, the BPU decodes branches during the fetch stage and folds out branch instructions before the dispatch stage.
- The dispatch pipeline stage is responsible for decoding the instructions supplied by the instruction fetch stage and determining which of the instructions are eligible to be dispatched in the current cycle. In addition, the source operands of the instructions are read from the appropriate register file and dispatched with the instruction to the execute pipeline stage. At the end of the dispatch pipeline stage, the dispatched instructions and their operands are latched by the appropriate execution unit.
- In the execute pipeline stage, each execution unit with an instruction executes the selected instruction (perhaps over multiple cycles), writes the instruction's result into the appropriate rename register, and notifies the completion stage when the execution has finished. In the case of an internal interrupt, the execution unit reports the interrupt to the completion/write-back pipeline stage and discontinues instruction execution until the interrupt is handled. The interrupt is not signaled until that instruction is the next to be completed. Execution of most floating-point instructions is pipelined within the FPU, allowing up to three instructions to execute in the FPU concurrently. The FPU pipeline stages are multiply, add, and round-convert. The LSU has two

pipeline stages: the first stage, for effective address calculation and MMU translation, and the second, for accessing data in the cache.

- The complete/write-back pipeline stage maintains the correct architectural machine state and transfers the contents of the rename registers to the GPRs and FPRs as instructions are retired. If the completion logic detects an instruction causing an interrupt, all subsequent instructions are canceled, their execution results in rename registers are discarded, and instructions are fetched from the correct instruction stream.

A superscalar processor core issues multiple, independent instructions into multiple pipelines, allowing instructions to execute in parallel. The e300c1 core has independent execution units for: integer instructions, floating-point instructions, branch instructions, load/store instructions, and system register instructions. The e300c3 provides two IUs, which improves the throughput of integer instructions. The e300c3 provides two integer units for greater integer instruction throughput along with enhanced multipliers in each IU that reduce the multiply instruction latency to a maximum of two cycles. The IU and the FPU each have dedicated register files for maintaining operands (GPRs and FPRs, respectively), allowing integer and floating-point calculations to occur simultaneously without interference.

The core provides support for single-cycle store, and it provides an adder/comparator in the system register unit that allows the dispatch and execution of multiple integer add and compare instructions on each cycle.

Because the PowerPC architecture can be applied to such a wide variety of implementations, instruction timing among processor cores varies accordingly.

### 7.3.7 Core Interface

The core interface is specific for each processor core implementation.

The MPC8313E contains an internal coherent system bus (CSB) that interfaces the processor core to the peripheral logic. This internal bus is very similar in function to the external 60x bus interface on the MPC603e. In the case of the MPC8313E, the CSB system logic decodes e300-initiated transactions and directs all accesses to the appropriate interface.

The e300 core can operate at a variety of frequencies allowing the designer to trade off performance for power consumption. The processor core is clocked from a separate PLL, which is referenced to the CSB frequency. This allows the processor core and the peripheral logic to operate at different frequencies.

The e300 core provides a versatile core interface that allows for a wide range of implementations. The interface includes a 32-bit address bus, a 64-bit data bus, and 56 control and information signals (see [Figure 7-4](#)). The core interface allows for address-only transactions, as well as address and data transactions. The core control and information signals include the address arbitration, address start, address transfer, transfer attribute, address termination, data arbitration, data transfer, data termination, and core state signals. Test and control signals provide diagnostics for selected internal circuits.




Figure 7-4. Core Interface

The core interface supports bus pipelining, allowing the address tenure of one transaction to overlap the data tenure of another. The extent of the pipelining depends on external arbitration and control circuitry. Similarly, the core supports split-bus transactions for systems with multiple potential bus masters—one device can have mastership of the address bus while another has mastership of the data bus. Allowing multiple bus transactions to occur simultaneously increases the available bus bandwidth for other activity and, as a result, improves performance.

The core clocking structure allows the bus to operate at integer multiples of the core cycle time.

The following sections describe the core bus support for memory operations. Note that some signals perform different functions depending on the addressing protocol used.

### 7.3.7.1 Memory Accesses

The e300 core CSB is a 64-bit data bus.

With a 64-bit CSB, memory accesses allow transfer sizes of 8, 16, 24, 32, 40, 48, 56, or 64 bits in one bus clock cycle. Data transfers occur in either single-beat transactions or four-beat burst transactions. Single-beat transactions are caused by noncached accesses that access memory directly (that is, reads and writes when caching is disabled, caching-inhibited accesses, and stores in write-through mode). Four-beat burst transactions, which always transfer an entire cache block (32 bytes), are initiated when a line is read from or written to memory.

### 7.3.7.2 Signals

The e300 core signals are grouped as follows:

- **Interrupts/Resets**—These signals include the external interrupt signal ( $\overline{int}$ ), critical interrupt signal ( $\overline{cint}$ ), checkstop signals, performance monitor signal ( $pm\_event\_in$ ) via the PM counters, and both soft reset and hard reset signals. They are used to interrupt and, under various conditions, to reset the core.

- JTAG/debug interface signals—The JTAG (based on the IEEE 1149.1 standard) interface and debug unit provides a serial interface to the system for performing monitoring and boundary tests. Two additional signals are added to the e300 core to allow observation of the internal clock state of the core (*stopped*) and to allow the external input to force the core into a halted state (*ext\_halt*).
- Core status and control—These signals include the memory reservation signal, machine quiesce control signals, time base/decrementer clock base enable signal, and the *tlbisynd* signal.
- Clock control—These signals provide for system clock input and frequency control.
- Test interface signals—Signals like address matching, combinational matching, and watchpoint are used in the core for production testing.
- Transfer attribute signals—These signals provide information about the type of transfer, such as the transfer size and whether the transaction is bursted, write-through, or cache-inhibited.

### 7.3.8 Debug Features

Some new debug features are specific to the e300 core. Accesses to the debug facilities are available only in supervisor mode by using the **mtspr** and **mfspr** instructions. The e300 provides the following additional feature in the JTAG/debug interface: Inclusion of breakpoint status and control pins: *stopped* and *ext\_halt*.

#### 7.3.8.1 Breakpoint Signaling

The breakpoint signaling provided on the e300 core allows observability of breakpoint matches external to the core. The *iabr*, *iabr2*, *dabr*, and *dabr2* breakpoint signals are asserted for at least one bus clock cycle when the respective breakpoint occurs. The status of the run state of the e300 core is indicated by the *stopped* pin. An asynchronous external breakpoint can be asserted to the e300 core using the *ext\_halt* pin:

- When DBCR and IBCR are configured for an OR combinational signal type, the breakpoint signals *iabr*, *iabr2* and *dabr*, *dabr2* reflect their respective breakpoints.
- When the DBCR and IBCR are configured for AND combinational signal type, only the *iabr2* and *dabr2* breakpoint signals are asserted after the AND condition is met (that is, both instruction breakpoints occurred or both data breakpoints occurred).
- When the core\_stopped pin is asserted, the e300 core has entered a stopped state and all internal clocking has stopped, indicating that a hardware debug event has occurred.
- The *ext\_halt* input pin can be used to force the core into halted state. The halted state may be a hardstop, conditional upon the HARDSTOP condition being set through the JTAG/debug interface

## 7.4 Differences Between Cores

The e300 core has similar functionality to the G2\_LE core. [Table 7-8](#) describes the differences between the G2\_LE and the e300.

**Table 7-8. Differences Between e300 and G2\_LE Cores**

e300 Core	G2_LE Core	Impact
New HID0 bits	—	The e300 core has a new HID0 bit defined to enable cache parity error reporting (ECPE).
New HID1 bits	—	The e300 core has new HID1 bits defined to extend the number of PLL configuration signals to seven (PC5, PC6).
New HID2 bits	—	The e300 core has new HID2 bits defined to support instruction fetch bursting (IFEB), MESI coherency protocol (MESI), instruction fetch cancels (IFEC), data cache queue sharing (EBQS), pipelining extension (EBPX), additional cache way locking (IWLCK and DWLCK), and instruction cache way protection (ICWP).
New PVR register value	—	The processor version register values differ.
New IBCR and DBCR bits	—	The e300 core has new IBCR[IABRSTAT, IABR2STAT] and DBCR[DABR1STAT, DABR2STAT] fields to provide instruction and data address breakpoint status.
—	16-Kbyte, four-way, set-associative, instruction and data caches	Some e300 cores may have different cache sizes than the G2_LE. See the <i>e300 PowerPC Core Reference Manual</i> for detailed information.
L1 cache parity	—	The e300 core supports parity for both instruction and data caches; the G2_LE does not support cache parity.
MEI or MESI coherency protocols	MEI protocol only	The e300 supports two coherency protocols: MEI and MESI; the G2_LE only supports the MEI protocol.
Instruction cancel extension	—	The e300 instruction cancel mechanism improves utilization of instruction cache by supporting 'hits-under-cancels' and 'misses-under-cancels'; the G2_LE requires the cancel to complete before new instruction fetches can begin.
Instruction fetch bursts to caching-inhibited space	Single-beat instruction fetches to caching-inhibited space	The e300's instruction fetch burst extension allows all caching-inhibited instruction fetches to be performed on the bus as burst transactions, even though the instructions are not cached. This improves performance for instruction space that is caching-inhibited, because up to eight instructions are returned with one bus operation. The G2_LE core must use single-beat instruction fetches for caching-inhibited space, returning only two instructions per bus operation.
Instruction cache way protection	—	The e300 core can protect locked ways in the instruction cache from invalidation; the G2_LE does not support instruction cache way protection.

Table 7-8. Differences Between e300 and G2\_LE Cores (continued)

e300 Core	G2_LE Core	Impact
Data cache queue sharing	—	The e300 has a new data cache queue sharing extension that allows the two burst-write queues in the bus unit to be used interchangeably for cache replacements and snoop pushes. Thus, the data cache can support two outstanding cache replacements or two outstanding snoop push operations on the bus at any given time.
<b>icbt</b> instruction	—	The e300 supports a new instruction cache block touch instruction that facilitates preloading the instruction cache before locking; the G2_LE core requires speculatively fetching instructions before locking the instruction cache.
1-1/2-level bus pipelining	1-level bus pipelining	For the e300, a new transaction can complete an address tenure when the previous transaction has been granted the data bus; for the G2_LE, a new transaction must wait until the previous data tenure has completed before completing its address tenure.
PowerPC little-endian not supported	PowerPC little-endian supported	PowerPC little-endian will not be supported in the e300 core, although true little-endian will be fully supported.
Data retry mode removed	Data retry mode available	$\overline{drtry}$ and $drtrymode$ will no longer be supported on the e300 and future versions.
External control instructions removed	External control instructions available	The <b>eciwx</b> and <b>ecowx</b> instruction pair will not be supported on the e300 core. These are optional instructions in the PowerPC architecture.
Reduced pin mode removed	Reduced pin mode available	Reduced pinout mode and the signal $redpinmode$ will not be supported in the e300 core.

# Chapter 8

## Integrated Programmable Interrupt Controller (IPIC)

This chapter describes the integrated programmable interrupt controller (IPIC), including a definition of the external signals and their functions. Also, the configuration, control, and status registers are described in this chapter. Note that individual chapters in this reference manual describe specific initialization aspects for each individual block.

### 8.1 Introduction

This chapter describes the IPIC interrupt protocol, various types of interrupt sources controlled by the IPIC unit, and the IPIC registers with some programming guidelines. The programming model is similar to the interrupt controller of the MPC8260. The interrupt controller provides interrupt management that is responsible for receiving hardware-generated interrupts from different sources (both internal and external). It also prioritizes and delivers the interrupts to the CPU for servicing. The IPIC prioritizes and manages interrupts from the following controller units:

- DDR memory controller (DDR)
- Enhanced local bus memory controller (eLBC)
- PCI
- Four-channel DMA controller (DMA)
- Message unit (MU)
- Dual three-speed Ethernet controllers (eTSEC1 and eTSEC2)
- DUART communication module (DUART)
- USB 2.0 dual role controller (USB DR)
- Security engine (SEC)
- System bus arbiter (SBA)
- Periodic interval timer (PIT)
- Real time clock timer (RTC ALR and RTC SEC)
- Eight global timers (GTM1–GTM8)
- Software watchdog timer (WDT)
- I<sup>2</sup>C controllers (I<sup>2</sup>C1 and I<sup>2</sup>C2)
- SPI controller (SPI)
- Power management controller (PMC)
- General-purpose I/O controller (GPIO)
- External pins ( $\overline{\text{IRQ}}[0:4]$ )

The interrupt sources controlled by the IPIC unit cause exceptions in the processor core. The internal interrupt ( $\overline{int}$ ) signal is the main interrupt output from the IPIC to the core and it causes the regular interrupt exception. The  $\overline{cint}$  signal is the critical interrupt output from the IPIC to the processor core and causes the critical interrupt exception. The  $\overline{smi}$  signal is the system management interrupt output from the IPIC to the processor core and causes the system management interrupt exception. The machine check exception is caused by the internal  $\overline{mcp}$  signal generated by the IPIC, informing the host processor of error conditions, assertion of the external  $\overline{IRQ0}$  machine-check request (enabled when  $SEMSR[SIRQ0] = 1$ ), and other conditions.

Table 8-1 shows the relationship of the various functional blocks and external signals of the device to the IPIC unit.

The IPIC receives interrupt request signals from the following two sources:

- External to the integrated device
- Internal to the integrated device

The unit selects the highest priority interrupt from all current interrupts and forwards it to the internal processor core, or off-chip for external servicing.

The IPIC also manages an internal non-maskable machine-check processor ( $\overline{mcp}$ ) signal and the interrupt generated by the off-chip interrupt sources ( $\overline{IRQ}[0:4]$ ).

The interrupt router of the IPIC monitors the outputs of the internal configuration registers. When the priority is highest in one of the received interrupt signals, the IPIC sets the corresponding bit in one of the interrupt pending registers—system internal interrupt pending register (SIPNR)/system external interrupt pending register (SEPNR). If the interrupt is not masked, the IPIC asserts the  $\overline{int}$  signal to indicate an interrupt request to the processor. When the processor is running the specific  $\overline{int}$ ,  $\overline{cint}$ , or  $\overline{smi}$  interrupt handler code, the processor must vectorize the external interrupt handler by explicitly (in software) reading the corresponding interrupt vector register (SIVCR, SCVCR, or SMVCR). In response to this read, the IPIC unit returns the vector (associated with the interrupt source) to the interrupt handler routine. In addition, the handler can vectorize different branches of interrupt handling.






Figure 8-1. Interrupt Sources Block Diagram

The IPIC receives the following types of interrupts:

- External interrupt—triggered by the off-chip signals ( $\overline{IRQn}$ ) listed in [Table 8-1](#)
- Internal interrupts—on-chip interrupts, triggered by the sources listed in [Table 8-7](#) and [Table 8-9](#)
- External and internal non-maskable machine check conditions, signaled by the sources listed in [Table 8-21](#) through  $\overline{mcp}$

The interrupt controller provides the ability to mask each interrupt source. Any source that can be caused by multiple events are also maskable.

When the IPIC receives an internal or external interrupt, its configuration register is checked to determine if it should be routed off-chip (to the external  $\overline{PCI\_INTA}$ ) or serviced as a normal external interrupt by the processor core (through the  $\overline{int}$  signal). As a third alternative, if the incoming interrupt has been configured as a critical or system management interrupt, the IPIC completes the processing of the interrupt by asserting  $\overline{cint}$  or  $\overline{smi}$  to the core. The assertion of the  $\overline{cint}$  or  $\overline{smi}$  signal to the core causes the interrupt to be serviced as a critical or a system management interrupt, respectively.

## 8.2 Features

The IPIC unit implements the following features:

- Functional and programming compatibility with the MPC8260 interrupt controller
- Support for external and internal discrete vectorized interrupt sources
- Support for external and internal non-maskable machine check conditions, signaled by  $\overline{mcp}$
- Programmable highest priority request (can be programmed to support a critical ( $\overline{cint}$ ) or system management interrupt ( $\overline{smi}$ ) type)
- Two programmable priority mixed groups of four on-chip and four external interrupt signals with two priority schemes for each group: grouped and spread
- Two programmable priority internal groups of eight on-chip interrupt signals with two priority schemes for each group: grouped and spread
- Two highest priority interrupts from each group can be programmed to support a critical or system management interrupt type
- External and internal interrupts directed to host processor
- Unique vector number for each interrupt source

## 8.3 Modes of Operation

The IPIC unit can operate in the core enable or core disable mode.

### 8.3.1 Core Enable Mode

In core enable mode, all internal interrupts (including those from the PCI block) are routed to and from the IPIC; the interrupts are sent to the PowerPC core. The DMA controller can optionally (depending on the programming of the DMA registers) steer its interrupt to the PCI host through the  $\overline{PCI\_INTA}$  signal.

In this mode all machine check interrupts are gathered by the IPIC unit and sent to the PowerPC core. If the device performs as a PCI host, the interrupts of the other PCI agents should be connected to the implementation's  $\overline{\text{IRQ}}_x$  signals and treated like normal external interrupts (sent to the core).

### 8.3.2 Core Disable Mode

In core disable mode, all internal interrupts (including those from the PCI block) are routed to and from the IPIC, the interrupts are then sent through the  $\overline{\text{PCI\_INTA}}$  signal to the PCI host CPU. Note that the core interrupt signal is masked. The user should use in this mode only the *int* output interrupt type (should not use *cint* or *smi* output interrupt types) to read an updated SIVCR. (See Section 8.5.7, “System Internal Interrupt Control Register (SICNR),” and Section 8.5.12, “System External Interrupt Control Register (SECNR).”)

In this mode, machine check interrupts are driven either on  $\overline{\text{PCI\_INTA}}$  or on  $\overline{\text{MCP\_OUT}}$  as level-sensitive interrupts. SERCR[MCPR] (see Section 8.5.15, “System Error Control Register (SERCR)”) controls which external signal is used.

## 8.4 External Signal Description

The following sections provide an overview and detailed descriptions of the IPIC signals.

### 8.4.1 Overview

The device has 5 distinct external interrupt request input signals ( $\overline{\text{IRQ}}[0:4]$ ) and one interrupt request output signal ( $\overline{\text{PCI\_INTA}}$ ). The IPIC interface signals are defined in Table 8-1.

Table 8-1. IPIC Signal Properties

Name	Port	Function	I/O	Reset	Requires Pull Up
$\overline{\text{IRQ}}[0:4]$	$\overline{\text{IRQ}}[0:4]$	External interrupts	I	—	Yes
$\overline{\text{PCI\_INTA}}$	$\overline{\text{PCI\_INTA}}$	Interrupt request output	O	Z	Yes
$\overline{\text{MCP\_OUT}}$	$\overline{\text{MCP\_OUT}}$	Interrupt request output	O	Z	Yes

### 8.4.2 Detailed Signal Descriptions

Table 8-2 provides detailed descriptions of the external IPIC signals.

Table 8-2. IPIC External Signals—Detailed Signal Descriptions

Signal	I/O	Description
$\overline{\text{IRQ}}[0:4]$	I	Interrupt request 0–4. The sense (level or edge) of each of these signals is programmable. All of these inputs can be driven completely asynchronously.
		<b>State Meaning</b> Asserted—When an external interrupt request signal is asserted, the priority is checked by the IPIC unit, and the interrupt is conditionally passed to the processor. Negated—There is no incoming interrupt from that source.
		<b>Timing</b> Assertion—All of these inputs can be asserted completely asynchronously. Negation—Interrupts programmed as level-sensitive must remain asserted until serviced.

**Table 8-2. IPIC External Signals—Detailed Signal Descriptions (continued)**

Signal	I/O	Description
$\overline{\text{PCI\_INTA}}$	OD	Interrupt request out. Active-low, open drain. When the IPIC is programmed in core disable mode, this output reflects the raw interrupts generated by on-chip sources. See <a href="#">Section 8.3, “Modes of Operation,”</a> for details.
		<b>State Meaning</b> Asserted—At least one interrupt is currently being signalled to the external system. Negated—Indicates no interrupt source currently routed to $\overline{\text{IRQ\_OUT}}$ .
		<b>Timing</b> Because external interrupts are asynchronous with respect to the system clock, both assertion and negation of $\overline{\text{IRQ\_OUT}}$ occur asynchronously with respect to the interrupt source. All timing given here is approximate. Assertion—Internal interrupt source: 3 system bus clock cycles after the interrupt occurs. External interrupt source: 4 cycles after the interrupt occurs. Negation—Follows interrupt source negation with the following delay: Internal interrupt: 3 system bus clock cycles. External interrupt: 4 cycles.
$\overline{\text{MCP\_OUT}}$	OD	Non-maskable Interrupt (machine check) request out. Active-low, open drain. When the IPIC is programmed in core disable mode, this output reflects the <i>mcp</i> interrupts generated by on-chip sources. See <a href="#">Section 8.3, “Modes of Operation.”</a>
		<b>State Meaning</b> Asserted—At least one machine check interrupt is currently being signalled to the external system. Negated—Indicates no interrupt source currently routed to $\overline{\text{MCP\_OUT}}$ .
		<b>Timing</b> Because external interrupts are asynchronous with respect to the system clock, both assertion and negation of $\overline{\text{MCP\_OUT}}$ occurs asynchronously with respect to the interrupt source. All timing given here is approximate. Assertion—Internal interrupt source: 2 system bus clock cycles after interrupt occurs. External interrupt source: 4 cycles after interrupt occurs. Negation—Follows interrupt source negation with the following delay: Internal interrupt: 2 system bus clock cycles. External interrupt: 4 cycles.

## 8.5 Memory Map/Register Definition

The IPIC programmable register map occupies 256 bytes of memory-mapped space. Reading undefined portions of the memory map returns all zeros; writing has no effect.

All IPIC registers are 32 bits wide and they are located on 32-bit address boundaries. Software can perform byte, half-word or word accesses to any IPIC registers. All addresses used in this chapter are offsets from the IPIC base, as defined in [Chapter 2, “Memory Map.”](#)

[Table 8-3](#) shows the memory map of the IPIC unit.

**Table 8-3. IPIC Register Address Map**

Offset	Register	Access	Reset Value	Section/ Page
0x00	System global interrupt configuration register (SICFR)	R/W	0x0000_0000	<a href="#">8.5.1/8-7</a>
0x04	System regular interrupt vector register (SIVCR)	R	0x0000_0000	<a href="#">8.5.2/8-9</a>
0x08	System internal interrupt pending register (SIPNR_H)	R	0x0000_0000	<a href="#">8.5.3/8-11</a>
0x0C	System internal interrupt pending register (SIPNR_L)	R	0x0000_0000	<a href="#">8.5.3/8-11</a>
0x10	System internal interrupt group A priority register (SIPRR_A)	R/W	0x0530_9770	<a href="#">8.5.4/8-13</a>
0x1C	System internal interrupt group D priority register (SIPRR_D)	R/W	0x0530_9770	<a href="#">8.5.5/8-14</a>

Table 8-3. IPIC Register Address Map (continued)

Offset	Register	Access	Reset Value	Section/ Page
0x20	System internal interrupt mask register (SIMSR_H)	R/W	0x0000_0000	8.5.6/8-15
0x24	System internal interrupt mask register (SIMSR_L)	R/W	0x0000_0000	8.5.6/8-15
0x28	System internal interrupt control register (SICNR)	R/W	0x0000_0000	8.5.7/8-16
0x2C	System external interrupt pending register (SEPNR)	R/W	Special	8.5.8/8-18
0x30	System mixed interrupt group A priority register (SMPRR_A)	R/W	0x0530_9770	8.5.9/8-18
0x34	System mixed interrupt group B priority register (SMPRR_B)	R/W	0x0530_9770	8.5.10/8-19
0x38	System external interrupt mask register (SEMSR)	R/W	0x0000_0000	8.5.11/8-20
0x3C	System external interrupt control register (SECNR)	R/W	0x0000_0000	8.5.12/8-21
0x40	System error status register (SERSR)	R/W	0x0000_0000	8.5.13/8-22
0x44	System error mask register (SERMR)	R/W		8.5.14/8-23
0x48	System error control register (SERCR)	R/W	0x0000_0000	8.5.15/8-24
0x4C–0x4F	Reserved	—	—	—
0x50	System internal interrupt force register (SIFCR_H)	R/W	0x0000_0000	8.5.16/8-25
0x54	System internal interrupt force register (SIFCR_L)	R/W	0x0000_0000	8.5.16/8-25
0x58	System external interrupt force register (SEFCR)	R/W	0x0000_0000	8.5.17/8-26
0x5C	System error force register (SERFR)	R/W	0x0000_0000	8.5.18/8-26
0x60	System critical interrupt vector register (SCVCR)	R	0x0000_0000	8.5.19/8-27
0x64	System management interrupt vector register (SMVCR)	R	0x0000_0000	8.5.20/8-27
0x68–0xFF	Reserved	—	—	—

### 8.5.1 System Global Interrupt Configuration Register (SICFR)

SICFR, shown in Figure 8-2, defines the highest priority interrupt and whether interrupts are grouped or spread in the priority table. See Table 8-4 for more information.




Figure 8-2. System Global Interrupt Configuration Register (SICFR)

Table 8-4 defines the bit fields of SICFR.

Table 8-4. SICFR Field Descriptions


Bits	Name	Description
0	—	Write ignored, read = 0
1–7	HPI	Highest priority interrupt. Specifies the 7-bit unique interrupt number/vector (see Table 8-6) of the single interrupt controller interrupt source that is advanced to the highest priority in the IPIC priority table (see Table 8-31). HPI can be modified dynamically.
8	—	Write ignored, read = 0
9	MPSB	Mixed interrupts priority scheme for group B. Selects the relative MIXB priority scheme. It cannot be changed dynamically. 0 Grouped. The MIXBs are grouped by priority at the top of the table. 1 Spread. The MIXBs are spread by priority in the table.
10	MPSA	Mixed interrupts priority scheme for group A. Selects the relative MIXA priority scheme. It cannot be changed dynamically. 0 Grouped. The MIXAs are grouped by priority at the top of the table. 1 Spread. The MIXAs are spread by priority in the table.
11	—	Write ignored, read = 0
12	IPSD	Internal interrupts priority scheme for group D. Selects the relative SYSD priority scheme. It cannot be changed dynamically. 0 Grouped. The SYSDs are grouped by priority at the top of the table. 1 Spread. The SYSDs are spread by priority in the table.
13	IPSC	Internal interrupts priority scheme for group C. Selects the relative SYSC priority scheme. It cannot be changed dynamically. 0 Grouped. The SYSCs are grouped by priority at the top of the table. 1 Spread. The SYSCs are spread by priority in the table.
14	IPSB	Internal interrupts priority scheme for group B. Selects the relative SYSB priority scheme. It cannot be changed dynamically. 0 Grouped. The SYSBs are grouped by priority at the top of the table. 1 Spread. The SYSBs are spread by priority in the table.
15	IPSA	Internal interrupts priority scheme for group A. Selects the relative SYSA priority scheme. It cannot be changed dynamically. 0 Grouped. The SYSAs are grouped by priority at the top of the table. 1 Spread. The SYSAs are spread by priority in the table.
16–21	—	Write ignored, read = 0
22–23	HPIT	HPI priority position IPIC output interrupt type. Defines which type of IPIC output interrupt signal ( $\overline{int}$ , $\overline{cint}$ , or $\overline{smi}$ ) asserts its request to the core in the HPI priority position. These bits cannot be changed dynamically. (If software really wants to change it, it has to make sure the corresponding interrupt source is masked or it won't happen during the change). The definition of HPIT is as follows: 00 $\overline{int}$ request is asserted to the core for HPI. 01 $\overline{smi}$ request is asserted to the core for HPI. 10 $\overline{cint}$ request is asserted to the core for HPI. 11 Reserved.
24–31	—	Write ignored, read = 0

### 8.5.2 System Global Interrupt Vector Register (SIVCR)

SIVCR, shown in Figure 8-3, contains a 7-bit code (Table 8-5) representing the regular unmasked interrupt source (INT) of the highest priority level.

**NOTE**

Note that in core disabled mode the user should use SIVCR only in order to read an updated interrupt vector (SCVCR and SMVCR should not be used).



**Figure 8-3. System Global Interrupt Vector Register (SIVCR)**

Table 8-5 defines the bit fields of SIVCR.

**Table 8-5. SIVCR Field Descriptions**

Bits	Name	Description
0–5	IVECx	Backward (MPC8260) compatible regular interrupt vector. Specifies a 6-bit unique number of the IPIC’s highest priority regular interrupt source, pending to the core. When a regular interrupt request occurs, SIVCR can be read. If there are multiple regular interrupt sources, SIVCR latches the highest priority regular interrupt. Note that IVECx field correctly reflects only the first 64 interrupt vectors (See Table 8-6 for details). The value of SIVEC cannot change while it is being read.
6–24	—	Write ignored, read = 0
25–31	IVEC	Regular interrupt vector. Specifies a 7-bit unique number of the IPIC’s highest priority regular interrupt source, pending to the core. Note that the when a regular interrupt request occurs, SIVCR can be read. If there are multiple regular interrupt sources, SIVCR latches the highest priority regular interrupt. Note that the IVEC field correctly reflects all interrupt vectors (see Table 8-6 for details). The value of SIVCR cannot change while it is being read.

Table 8-6 shows the definition of IVEC.

**Table 8-6. IVEC/CVEC/MVEC Field Definition**

Interrupt ID Number	Interrupt Meaning	Interrupt Vector
0	Error (no interrupt)	0b000_0000
1–8	Reserved	0b000_0001–0b000_1000
9	UART1	0b000_1001
10	UART2	0b000_1010
11	SEC	0b000_1011
12	eTSEC1 1588 timer	0b000_1100
13	eTSEC2 1588 timer	0b000_1101

Table 8-6. IVEC/CVEC/MVEC Field Definition (continued)

Interrupt ID Number	Interrupt Meaning	Interrupt Vector
14	I2C1	0b000_1110
15	I2C2	0b000_1111
16	SPI	0b001_0000
17	IRQ1	0b001_0001
18	IRQ2	0b001_0010
19	IRQ3	0b001_0011
20	IRQ4	0b001_0100
21–31	Reserved	0b001_0101–0b001_1111
32	TSEC1 Tx	0b010_0000
33	TSEC1 Rx	0b010_0001
34	TSEC1 Err	0b010_0010
35	TSEC2 TX	0b010_0011
36	TSEC2 Rx	0b010_0100
37	TSEC2 Err	0b010_0101
38	USB DR	0b010_0110
39–47	Reserved	0b010_0111–0b010_1111
48	IRQ0	0b011_0000
49–63	Reserved	0b011_0001–0b011_1111
64	RTC SEC	0b100_0000
65	PIT	0b100_0001
66	PCI	0b100_0010
67	Reserved	0b100_0011
68	RTC ALR	0b100_0100
69	MU	0b100_0101
70	SBA	0b100_0110
71	DMA	0b100_0111
72	GTM4	0b100_1000
73	GTM8	0b100_1001
74	GPIO	0b100_1010
75	Reserved	0b100_1011
76	DDR	0b100_1100
77	LBC	0b100_1101
78	GTM2	0b100_1110
79	GTM6	0b100_1111
80	PMC	0b101_0000



Table 8-6. IVEC/CVEC/MVEC Field Definition (continued)

Interrupt ID Number	Interrupt Meaning	Interrupt Vector
81–83	Reserved	0b101_0001–0b101_0011
84	GTM3	0b101_0100
85	GTM7	0b101_0101
86–89	Reserved	0b101_0110–0b101_1001
90	GTM1	0b101_1010
91	GTM5	0b101_1011
92–127	Reserved	0b101_1100–0b111_1111

### 8.5.3 System Internal Interrupt Pending Registers (SIPNR\_H and SIPNR\_L)

Each bit in SIPNR\_H and SIPNR\_L, shown in [Figure 8-4](#) and [Figure 8-5](#), may be assigned an internal interrupt source. (Implemented bits are listed in [Table 8-7](#).) When an interrupt request is received, the interrupt controller sets the corresponding SIPNR bit. When a pending interrupt is handled, the user clears the SIPNR bit by clearing the corresponding event register bit.

Note that SIPNR bit positions are not changed according to relative priority.




Figure 8-4. System Internal Interrupt Pending Register (SIPNR\_H)

[Table 8-7](#) lists implemented SIPNR\_H fields. Note that these field descriptions are also valid for SIFCR\_H and SIMSR\_H.

Table 8-7. SIPNR\_H/SIFCR\_H/SIMSR\_H Bit Assignments

Bits	Field
0	TSEC1 Tx
1	TSEC1 Rx
2	TSEC1 Err
3	TSEC2 Tx
4	TSEC2 Rx
5	TSEC2 Err
6	USB DR
7–23	—
24	UART1
25	UART2

**Table 8-7. SIPNR\_H/SIFCR\_H/SIMSR\_H Bit Assignments (continued)**


Bits	Field
26	SEC
27	eTSEC1 1588 timer
28	eTSEC2 1588 timer
29	I2C1
30	I2C2
31	SPI

Table 8-8 defines the bit fields of SIPNR\_H.

**Table 8-8. SIPNR\_H Field Descriptions**

Bits	Name	Description
0–31	INT $n$	Each implemented bit (listed in Table 8-7) corresponds to an internal interrupt source. When an interrupt is received, the interrupt controller sets the corresponding SIPNR bit. When a pending interrupt is handled, the user clears the SIPNR bit by clearing the corresponding event register bit. SIPNR bits are read only. Writing to this register has no effect. Note that the SIPNR bit positions are not changed according to their relative priority. For unimplemented bits, writes are ignored, read = 0.

SIPNR\_L is shown in Figure 8-4.



**Figure 8-5. System Internal Interrupt Pending Register (SIPNR\_L)**

Table 8-9 lists implemented SIPNR\_L fields. Note that these field assignments are also valid for SIFCR\_L and SIMSR\_L.

**Table 8-9. SIPNR\_L/SIFCR\_L/SIMSR\_L Bit Assignments**

Bits	Field
0	RTC SEC
1	PIT
2	PCI
3	—
4	RTC ALR
5	MU
6	SBA
7	DMA
8	GTM4
9	GTM8

**Table 8-9. SIPNR\_L/SIFCR\_L/SIMSR\_L Bit Assignments (continued)**

Bits	Field
10	GPIO
11	—
12	DDR
13	LBC
14	GTM2
15	GTM6
16	PMC
17–19	—
20	GTM3
21	GTM7
22–25	—
26	GTM1
27	GTM5
28–30	—
31	—

Table 8-10 defines the bit fields of SIPNR\_L.

**Table 8-10. SIPNR\_L Field Descriptions**

Bits	Name	Description
0–31	INT <sub>n</sub>	Each implemented bit (listed in Table 8-9) corresponds to an internal interrupt source. When an interrupt is received, the interrupt controller sets the corresponding SIPNR bit. When a pending interrupt is handled, the user clears the SIPNR bit by clearing the corresponding event register bit. SIPNR bits are read only. Writing to this register has no effect. Note that the SIPNR bit positions are not changed according to their relative priority. For unimplemented bits, writes are ignored, read = 0.

#### 8.5.4 System Internal Interrupt Group A Priority Register (SIPRR\_A)

SIPRR\_A, shown in Figure 8-6, defines the priority between TSEC1 transmit request (TSEC1 Tx), TSEC1 receive request (TSEC1 Rx), TSEC1 transmit/receive error (TSEC1 Err), TSEC2 transmit request (TSEC2 Tx), TSEC2 receive request (TSEC2 Rx) TSEC2 transmit/receive error (TSEC2 Err), USB DR, internal interrupt signals.

For more information, see Section 8.6.3, “Internal Interrupts Group Relative Priority.”

Offset 0x10

Access: Read/write




Figure 8-6. System Internal Interrupt Group A Priority Register (SIPRR\_A)

Table 8-11 defines the bit fields of SIPRR\_A.

Table 8-11. SIPRR\_A Field Descriptions

Bits	Name	Description
0–2	SYSA0P	SYSA0 priority order. Defines which interrupt source asserts its request in the SYSA0 priority position. The user should not program the same code to multiple priority positions (0–7). These bits can be changed dynamically. The definition of SYSA0P is as follows: 000 TSEC1 Tx asserts its request in the SYSA0 position. 001 TSEC1 Rx asserts its request in the SYSA0 position. 010 TSEC1 Err asserts its request in the SYSA0 position. 011 TSEC2 Tx asserts its request in the SYSA0 position. 100 TSEC2 Rx asserts its request in the SYSA0 position. 101 TSEC2 Err asserts its request in the SYSA0 position. 110 USB DR asserts its request in the SYSA0 position. 111 Reserved
3–11, 16–27	SYSA1P–SYSA7P	Same as SYSA0P, but for SYSA1P–SYSA7P.
12–15, 28–31	—	Write ignored, read = 0

### 8.5.5 System Internal Interrupt Group D Priority Register (SIPRR\_D)

SIPRR\_D, shown in Figure 8-7, defines the priority among the interrupt sources listed in Table 8-12.

Offset 0x1C

Access: Read/write




Figure 8-7. System Internal Interrupt Group D Priority Register (SIPRR\_D)

Table 8-12 defines the bit fields of SIPRR\_D.

**Table 8-12. SIPRR\_D Field Descriptions**

Bits	Name	Description
0–2	SYSD0P	SYSD0 priority order. Defines which interrupt source asserts its request in the SYSD0 priority position. The user should not program the same code to more than one priority position (0–7). These bits can be changed dynamically. SYSD0P is defined as follows: 000 UART1 asserts its request in the SYSD0 position. 001 UART2 asserts its request in the SYSD0 position. 010 SEC asserts its request in the SYSD0 position. 011 eTSEC1 1588 timer asserts its request in the SYSD0 position. 100 eTSEC2 1588 timer asserts its request in the SYSD0 position. 101 I2C1 asserts its request in the SYSD0 position. 110 I2C2 asserts its request in the SYSD0 position. 111 SPI asserts its request in the SYSD0 position.
3–11, 16–27	SYSD1P– SYSD7P	Same as SYSD0P, but for SYSD1P–SYSD7P.
12–15, 28–31	—	Write ignored, read = 0

### 8.5.6 System Internal Interrupt Mask Register (SIMSR\_H and SIMSR\_L)

Each implemented bit in SIMSR\_H and SIMSR\_L, shown in [Figure 8-8](#) and [Figure 8-9](#), corresponds to an internal interrupt source. The user masks an interrupt by clearing the corresponding SIMSR bit. When an interrupt request occurs, the corresponding SIPNR bit is set, regardless of the SIMSR bit. However, if the corresponding SIMSR bit is cleared, no interrupt request is passed to the core.

When an SIMSR bit is cleared by the user at the same time corresponding interrupt source requests an interrupt service, the request stops. If the user sets the SIMSR bit later, the core processes any pending corresponding interrupt requests according to its priority.



**Figure 8-8. System Internal Interrupt Mask Register (SIMSR\_H)**

Table 8-13 defines the bit fields of SIMSR\_H.

**Table 8-13. SIMSR\_H Field Descriptions**

Bits	Name	Description
0–31	INT <sub>n</sub>	<p>Each implemented bit (listed in Table 8-7) corresponds to an external interrupt source. The user masks an interrupt by clearing the corresponding SIMSR bit. An interrupt is unmasked (enable) by setting the corresponding SIMSR bit. The SIMSR can be read by the user at any time.</p> <p><b>Note:</b></p> <ul style="list-style-type: none"> <li>• SIMSR bit positions do not change according to their relative priority.</li> <li>• The user can clear pending register bits that were set by multiple interrupt events only by clearing all unmasked events in the corresponding event register.</li> <li>• If an SIMSR bit is masked at the same time that the corresponding SIPNR bit causes an interrupt request to the core, the error vector is issued (if no other interrupts are pending). Thus, the user should always include an error vector routine, even if it contains only an rfi instruction. The error vector cannot be masked.</li> </ul> <p>Unimplemented bits, shown as reserved in Table 8-7, are ignored on writes; read = 0.</p>

Figure 8-9 shows SIMSR\_L.



**Figure 8-9. System Internal Interrupt Mask Register (SIMSR\_L)**

Table 8-14 defines the bit fields of SIMSR\_L.

**Table 8-14. SIMSR\_L Field Descriptions**


Bits	Name	Description
0–31	INT <sub>n</sub>	<p>Each implemented bit (listed in Table 8-9) corresponds to an external interrupt source. The user masks an interrupt by clearing the corresponding SIMSR bit. An interrupt is unmasked (enabled) by setting the corresponding SIMSR bit. The SIMSR can be read by the user at any time.</p> <p><b>Note:</b></p> <ul style="list-style-type: none"> <li>• SIMSR bit positions are not changed according to their relative priority.</li> <li>• The user can clear pending register bits that were set by multiple interrupt events only by clearing all unmasked events in the corresponding event register.</li> <li>• If an SIMSR bit is masked at the same time that the corresponding SIPNR bit causes an interrupt request to the core, the error vector is issued (if no other interrupts are pending). Thus, the user should always include an error</li> </ul> <p>Unimplemented bits, shown as reserved in Figure 8-9, are ignored on writes; read = 0.</p>

## 8.5.7 System Internal Interrupt Control Register (SICNR)

SICNR, shown in Figure 8-10, defines the IPIC output interrupt type ( $\overline{\text{int}}$ ,  $\overline{\text{cint}}$ , or  $\overline{\text{smi}}$ ) in the SYSA0–SYSA1 and SYSD0–SYSD1 priority positions. All other priority positions assert  $\overline{\text{int}}$  to the core.

Offset 0x28

Access: Read write



**Figure 8-10. System Internal Interrupt Control Register (SICNR)**

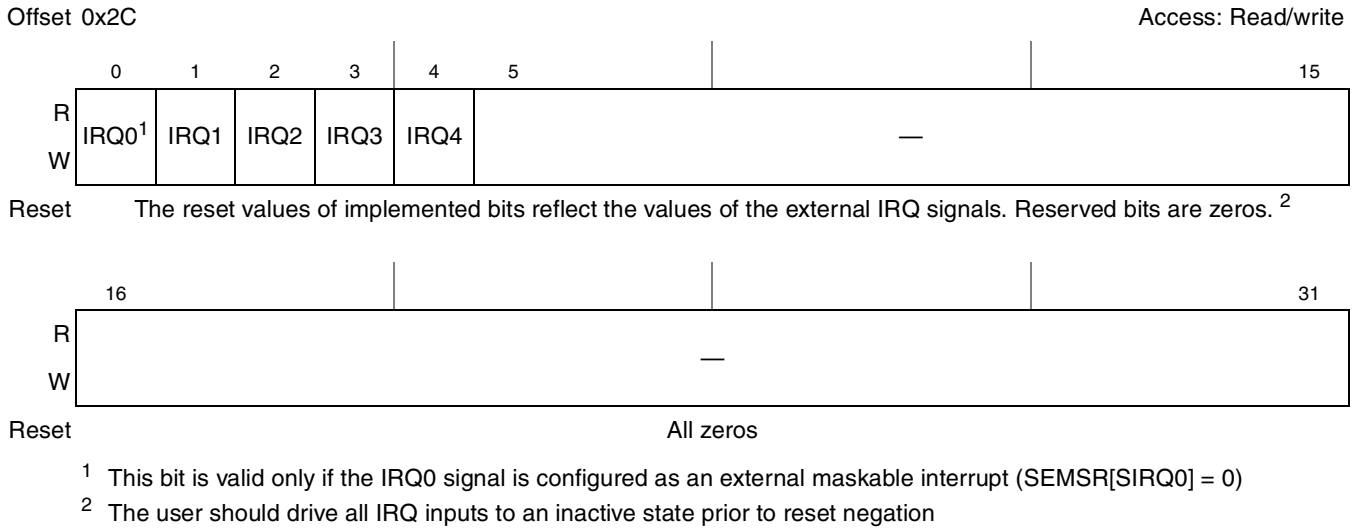
Table 8-15 defines the bit fields of SICNR.

**Table 8-15. SICNR Field Descriptions**

Bits	Name	Description
0–1	SYSD0T	SYSD0 priority position IPIC output interrupt type. Defines which type of the IPIC output interrupt signal ( $\overline{int}$ , $\overline{cint}$ , or $\overline{smi}$ ) asserts its request to the core in the SYSD0 priority position. These bits cannot be changed dynamically. (to change it, software must make sure the corresponding interrupt source is masked or it cannot happen during the change). The definition of SYSD0T is as follows: 00 $\overline{int}$ request is asserted to the core for SYSD0. 01 $\overline{smi}$ request is asserted to the core for SYSD0. 10 $\overline{cint}$ request is asserted to the core for SYSD0. 11 Reserved
2–3	SYSD1T	Same as SYSD0T, but for SYSD1T.
4–23	—	Write ignored, read = 0
24–25	SYSA0T	SYSA0 priority position IPIC output interrupt type. Defines which type of the IPIC output interrupt signal ( $\overline{int}$ , $\overline{smi}$ , or $\overline{cint}$ ) asserts its request to the core in the SYSA0 priority position. These bits can not be changed dynamically. (If s/w really wants to change it, it has to make sure the corresponding interrupt source is masked or it won't happen during the change). The definition of SYSA0T is as follows: 00 $\overline{int}$ request is asserted to the core for SYSA0. 01 $\overline{smi}$ request is asserted to the core for SYSA0. 10 $\overline{cint}$ request is asserted to the core for SYSA0. 11 Reserved.
26–27	SYSA1T	Same as SYSA0T, but for SYSA1T
28–31	—	Write ignored, read = 0

### 8.5.8 System External Interrupt Pending Register (SEPNR)

Each bit in the SEPNR, shown in [Figure 8-11](#), corresponds to an external interrupt source. When an interrupt is received, the interrupt controller sets the corresponding SEPNR bit.



**Figure 8-11. System External Interrupt Pending Register (SEPNR)**


[Table 8-16](#) defines the bit fields of SEPNR.

**Table 8-16. SEPNR Field Descriptions**

Bits	Name	Description
0–4	IRQ <sub>n</sub>	Each bit corresponds to an external interrupt source. When an external interrupt is received, the interrupt controller sets the corresponding SEPNR bit. When a pending interrupt is handled, the user must clear the corresponding SEPNR bit. For level triggered cases, the software needs to cause the $\overline{\text{IRQ}}_n$ to negate which automatically clears the bit in SEPNR. For edge-triggered cases, the software needs to clear the corresponding bit in SEPNR. SEPNR bits are cleared by writing ones to them. Because the user can only clear bits in this register, writing zeros to this register has no effect. Note that the SEPNR bit positions are not changed according to their relative priority.
5–31	—	Write ignored, read = 0

### 8.5.9 System Mixed Interrupt Group A Priority Register (SMPRR\_A)

The SMPRR\_A, shown in [Figure 8-12](#), defines the priority among the sources listed in [Table 8-17](#).



**Figure 8-12. System Mixed Interrupt Group A Priority Register (SMPRR\_A)**



Table 8-17 defines the bit fields of SMPRR\_A.

**Table 8-17. SMPRR\_A Field Descriptions**

Bits	Name	Description
0–2	MIXA0P	MIXA0 priority order. Defines which interrupt source asserts its request in the MIXA0 priority position. The user must not program the same code to more than one priority position (0–7). These bits can be changed dynamically. The definition of MIXA0P is as follows: 000 RTC SEC asserts its request to the MIXA0 position. 001 PIT asserts its request to the MIXA0 position. 010 PCI asserts its request to the MIXA0 position. 011 Reserved. 100 IRQ0 asserts its request to the MIXA0 position. This field for MIXA0 position is valid (must not be ignored) if IRQ0 signal configured as an external maskable interrupt (SEMSR[SIRQ0] = 0). 101 IRQ1 asserts its request to the MIXA0 position. 110 IRQ2 asserts its request to the MIXA0 position. 111 IRQ3 asserts its request to the MIXA0 position.
3–11, 16–27	MIXA1P– MIXA7P	Same as MIXA0P, but for MIXA1P–MIXA7P.
12–15, 28–31	—	Write ignored, read = 0

### 8.5.10 System Mixed Interrupt Group B Priority Register (SMPRR\_B)

SMPRR\_B, shown in Figure 8-13, defines the priority among the sources listed in Table 8-18.

Offset 0x34

Access: Read/write

	0	2	3	5	6	8	9	11	12	15	16	18	19	21	22	24	25	27	28	31
R	MIXB0P	MIXB1P	MIXB2P	MIXB3P	—	MIXB4P	MIXB5P	MIXB6P	MIXB7P	—										
W																				
Reset	0	0	0	0	0	1	0	1	0	0	1	1	0	0	0	1	0	1	1	1
	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0

**Figure 8-13. System Mixed Interrupt Group B Priority Register (SMPRR\_B)**

Table 8-18 defines the bit fields of SMPRR\_B.

**Table 8-18. SMPRR\_B Field Descriptions**

Bits	Name	Description
0–2 3–11, 16–27	MIXB $n$ P	MIXB $n$ priority order. Defines which interrupt source asserts its request in the MIXB $n$ priority position. The user must not program the same code to more than one priority position (0–7). These bits can be changed dynamically. The definition of MIXB $n$ P is as follows: 000 RTC ALR asserts its request to the MIXB $n$ position. 001 MU asserts its request to the MIXB $n$ position. 010 SBA asserts its request to the MIXB $n$ position. 011 DMA asserts its request to the MIXB $n$ position. 100 IRQ4 asserts its request to the MIXB $n$ position. 101–111 Reserved
12–15, 28–31	—	Write ignored, read = 0

### 8.5.11 System External Interrupt Mask Register (SEMSR)


Each bit in the system external interrupt mask register (SEMSR), shown in Figure 8-11, corresponds to an external interrupt source. The user masks an interrupt by clearing the corresponding SEMSR bit. An interrupt is unmasked (enabled) by setting the corresponding SEMSR bit.

When an external interrupt request occurs, the corresponding SEP $n$ R bit is set regardless of the setting of the corresponding SEMSR bit. However, if the corresponding SEMSR bit is cleared, no interrupt request is passed to the core.

When an SEMSR bit is cleared by the user at the same time that an interrupt source requests an interrupt service, the request stops. If the user sets the SEMSR bit later, a previously pending interrupt request is processed by the core according to its assigned priority. SEMSR can be read by the user at any time.

Offset 0x38

Access: Read/write



Reset The reset values of implemented bits reflect the values of the external IRQ signals. Reserved bits are zeros.<sup>2</sup>



Reset All zeros

<sup>1</sup> This bit is valid only if the IRQ0 signal is configured as an external maskable interrupt (SEMSR[SIRQ0] = 0)

<sup>2</sup> The user should drive all IRQ inputs to an inactive state prior to reset negation

**Figure 8-14. System External Interrupt Mask Register (SEMSR)**

Table 8-19 defines the bit fields of SEMSR.

**Table 8-19. SEMSR Field Descriptions**

Bits	Name	Description
0–4	—	Each bit corresponds to an external interrupt source. The user masks an interrupt by clearing the SEMSR bit. An interrupt can be enabled by setting the corresponding SEMSR bit. SEMSR can be read by the user at any time. <b>Note:</b> <ul style="list-style-type: none"> <li>SEMSR bit positions are not affected by their relative priority.</li> <li>The user can clear pending register bits that were set by multiple interrupt events only by clearing all unmasked events in the corresponding event register.</li> <li>If an SEMSR bit is masked at the same time that the corresponding SEPNR bit causes an interrupt request to the core, the error vector is issued (if no other interrupts pending). Thus, the user must always include an error vector routine, even if it contains only an <code>rfi</code> instruction. The error vector cannot be masked.</li> </ul>
5–15	—	Write ignored, read = 0
16	SIRQ0	Steer IRQ0. 0 IRQ0 is used as external interrupt request 1 IRQ0 is used as external MCP request
17–31	—	Write ignored, read = 0

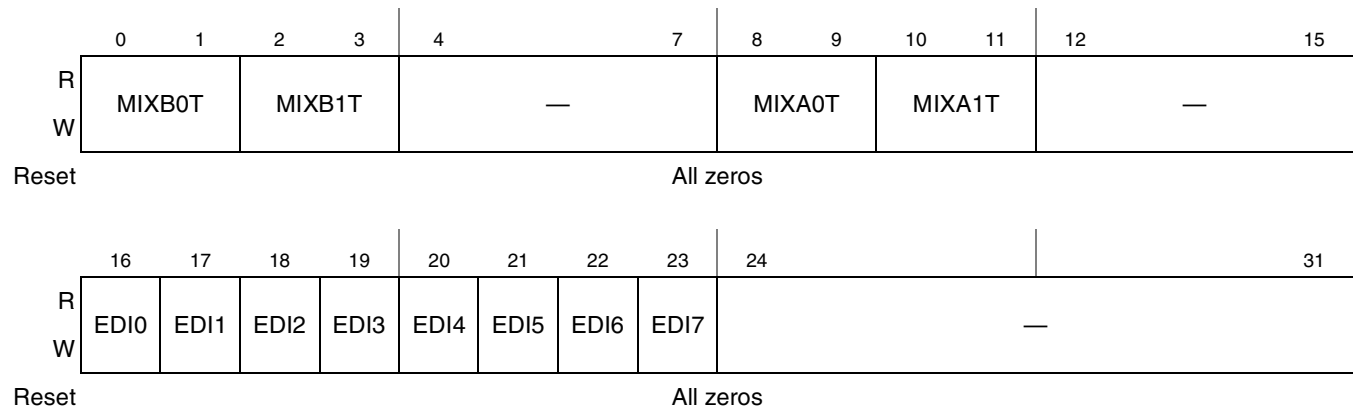
### 8.5.12 System External Interrupt Control Register (SECNR)

SECNR, shown in Figure 8-15, defines the edge detect mode for external  $\overline{IRQ}_n$  interrupt signals and determines whether the corresponding  $\overline{IRQ}_n$  signal asserts an interrupt request upon either a high-to-low change or assertion on the pin. It also defines the IPIC output interrupt type ( $\overline{int}$ ,  $\overline{cint}$ , or  $\overline{smi}$ ) in the MIXA0–MIXA1 and MIXB0–MIXB1 priority positions.

Note that in core disabled mode of operation the user should use the  $\overline{int}$  output interrupt type (should not use  $\overline{cint}$  or  $\overline{smi}$  output interrupt types) in order to read an updated SIVCR.

Offset 0x3C

Access: Read/write



**Figure 8-15. System External Interrupt Control Register (SECNR)**


Table 8-20 defines the bit fields of SECNR.

**Table 8-20. SECNR Field Descriptions**

Bits	Name	Description
0–1	MIXB0T	MIXB0 priority position IPIC output interrupt type. Defines which type of the IPIC output interrupt signal ( $\overline{int}$ , $\overline{cint}$ , or $\overline{smi}$ ) asserts its request to the core in the MIXB0 priority position. These bits can be changed dynamically. The definition of MIXB0T is as follows: 00 $\overline{int}$ request is asserted to the core for MIXB0. 01 $\overline{smi}$ request is asserted to the core for MIXB0. 10 $\overline{cint}$ request is asserted to the core for MIXB0. 11 Reserved
2–3	MIXB1T	Same as MIXB0T, but for MIXB1T.
4–7	—	Write ignored, read = 0
8–9	MIXA0T	MIXA0 priority position IPIC output interrupt Type. Defines which type of the IPIC output interrupt signal ( $\overline{int}$ , $\overline{cint}$ , or $\overline{smi}$ ) asserts its request to the core in the MIXA0 priority position. These bits can be changed dynamically. The definition of MIXA0T is as follows: 00 $\overline{int}$ request is asserted to the core for MIXA0. 01 $\overline{smi}$ request is asserted to the core for MIXA0. 10 $\overline{cint}$ request is asserted to the core for MIXA0. 11 Reserved
10–11	MIXA1T	Same as MIXA0T, but for MIXA1T.
12–15	—	Write ignored, read = 0
16–23	EDIx	Each bit defines the edge detect mode for the external $\overline{IRQn}$ interrupt signals, determines whether the corresponding $\overline{IRQn}$ signal asserts an interrupt request upon either a high-to-low change or low assertion on the pin. The corresponding $\overline{IRQn}$ signal asserts an interrupt request as follows: 0 Low assertion on $\overline{IRQn}$ generates an interrupt request (level sensitive). 1 High-to-low change on $\overline{IRQn}$ generates an interrupt request (edge sensitive).
24–31	—	Write ignored, read = 0

### 8.5.13 System Error Status Register (SERSR)

The bits in the SERSR, shown in Figure 8-16, correspond to the external and internal non-maskable error source machine check (mcp) conditions listed in Table 8-21. When an error interrupt signal is received, the interrupt controller sets the corresponding SERSR bit.



**Figure 8-16. System Error Status Register (SERSR)**

Table 8-21 lists the implemented SERSR bits. Note that these field assignments are valid for SERMR and SERFR.

**Table 8-21. SERSR/SERMR/SERFR Bit Assignments**

Bits	Field
0	IRQ0 <sup>1</sup>
1	WDT
2	SBA
3	—
4	—
5	PCI
6	—
7	MU
8–14	—
15	—
16–31	—

<sup>1</sup> This bit is valid only if the IRQ0 signal is configured as an external MCP interrupt (SEMSR[SIRQ0] = 1)

Table 8-22 defines the bit fields of SERSR.

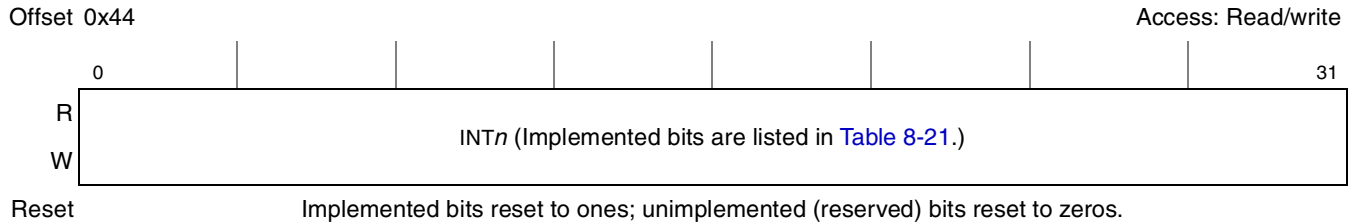
**Table 8-22. SERSR Field Descriptions**

Bits	Name	Description
0–31	INT <sub>n</sub>	Each implemented bit in the SERSR, listed in Table 8-21, corresponds to an external and an internal error source (mcp). When an error interrupt signal is received, the interrupt controller sets the corresponding SERSR bit. SERSR bits are cleared by writing ones to them. Unmasked event register bits should be cleared before clearing SERSR bits. Because the user can only clear bits in this register, writing zeros to this register has no effect. SERSR bits are cleared by power-on reset. Subsequent soft and hard resets do not affect SERSR bit states. For unimplemented bits (listed as reserved in Table 8-21), writes are ignored, read = 0

### 8.5.14 System Error Mask Register (SERMR)

Each implemented bit in SERMR, shown in Figure 8-17, corresponds to an external and an internal  $\overline{mcp}$  source (MCP). The user masks an MCP by clearing and enables an interrupt by setting the corresponding SERMR bit. When a masked MCP occurs, the corresponding SERSR bit is set, regardless of the setting of the corresponding SERMR bit although no MCP request is passed to the core in this case. The SERMR can be read by the user at any time.

## Integrated Programmable Interrupt Controller (IPIC)



**Figure 8-17. System Error Mask Register (SERMR)**


[Table 8-23](#) defines the bit fields of SERMR.

**Table 8-23. SERMR Field Descriptions**

Bits	Name	Description
0–31	INT $n$	Each implemented SERMR bit, listed in <a href="#">Table 8-21</a> , corresponds to an external and an internal MCP source. The user masks an MCP by clearing and enables an interrupt by setting the corresponding SERMR bit. When a masked MCP occurs, the corresponding SERSR bit is set, regardless of the setting of the SERMR bit although no MCP request is passed to the core. The SERMR can be read by the user at any time. Writes to unimplemented (reserved) bits are ignored; read = 0

### 8.5.15 System Error Control Register (SERCR)

SERCR, shown in [Figure 8-18](#), defines the control bits that route MCP requests in core disable mode to either `MCP_OUT` or `PCI_INTA` in core-disable mode.



**Figure 8-18. System Error Control Register (SERCR)**

[Table 8-24](#) defines the bit fields of SERCR.

**Table 8-24. SERCR Field Descriptions**

Bits	Name	Description
0–30	—	Write ignored, read = 0
31	MCPR	MCP route. Route MCP request to either <code>MCP_OUT</code> or <code>PCI_INTA</code> (in core disable mode). 0 MCP routed to <code>PCI_INTA</code> (in core disable mode). 1 MCP routed to <code>MCP_OUT</code> (in core disable mode).

## 8.5.16 System Internal Interrupt Force Registers (SIFCR\_H and SIFCR\_L)

Each implemented bit SIFCR\_H and SIFCR\_L, shown in [Figure 8-19](#) and [Figure 8-20](#), corresponds to an internal interrupt source. When a bit is set, the interrupt controller generates the corresponding interrupt (sets the corresponding SIPNR bit). The SIFCR can be read by the user at any time.



**Figure 8-19. System Internal Interrupt Force Register (SIFCR\_H)**

[Table 8-25](#) defines the bit fields of SIFCR\_H.

**Table 8-25. SIFCR\_H Field Descriptions**

Bits	Name	Description
0–31	INT $n$	Each implemented bit, listed in <a href="#">Table 8-7</a> , corresponds to an internal interrupt source. The user forces an interrupt by setting the corresponding SIFCR $x$ bit. SIFCR $n$ bit positions are not changed according to their relative priority. Writes to unimplemented (reserved) bits are ignored; read = 0

SIFCR\_L is shown in [Figure 8-20](#).



**Figure 8-20. System Internal Interrupt Force Register (SIFCR\_L)**

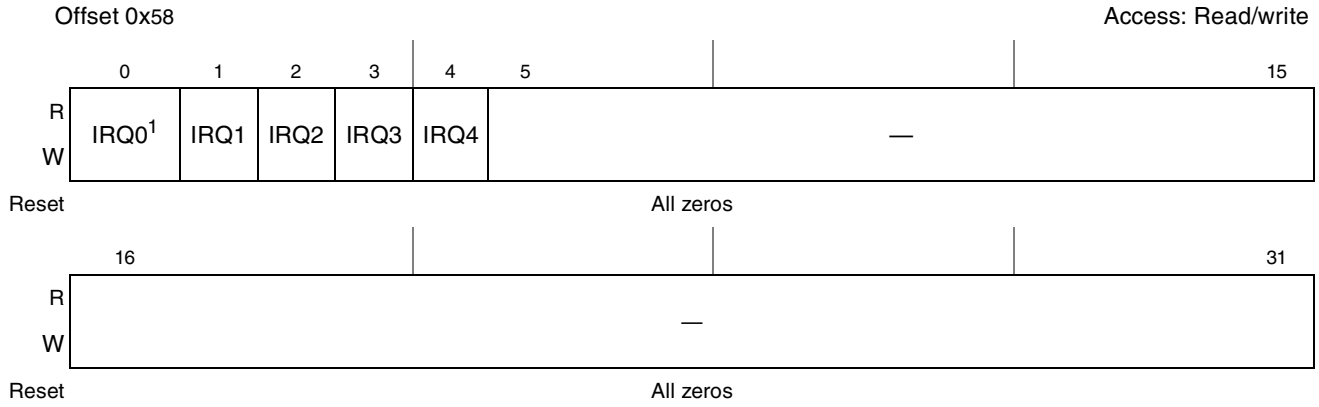
[Table 8-26](#) defines the bit fields of SIFCR\_L.

**Table 8-26. SIFCR\_L Field Descriptions**

Bits	Name	Description
0–31	INT $n$	Each implemented bit, listed in <a href="#">Table 8-9</a> , corresponds to an internal interrupt source. The user forces an interrupt by setting the corresponding SIFCR $x$ bit. SIFCR $x$ bit positions are not changed according to their relative priority. Writes to unimplemented (reserved) bits are ignored; read = 0

### 8.5.17 System External Interrupt Force Register (SEFCR)

Each implemented bit in SEFCR, shown in Figure 8-21, corresponds to an external interrupt source. When a bit is set, the interrupt controller generates the corresponding external interrupt (sets the corresponding SEPNR bit). SEFCR can be read by the user at any time.



<sup>1</sup> This bit is valid only if IRQ0 is configured as an external maskable interrupt (SEMSR[SIRQ0] = 0)

Figure 8-21. System External Interrupt Force Register (SEFCR)

Table 8-27 defines the bit fields of SEFCR.

Table 8-27. SEFCR Field Descriptions

Bits	Name	Description
0–4	IRQ <sub>n</sub>	Each bit corresponds to an external interrupt source. The user force an interrupt by setting the SEFCR bit. <b>Note:</b> SEFCR bit positions are not affected by their relative priority.
5–31	—	Write ignored, read = 0

### 8.5.18 System Error Force Register (SERFR)

Each bit in the system error force register (SERFR), shown in Figure 8-22, corresponds to an external MCP source. When a bit is set, the interrupt controller generates the corresponding MCP interrupt (sets the corresponding SERSR bit). The SERFR can be read by the user at any time.



Figure 8-22. System Error Status Register (SERFR)



Table 8-28 defines the bit fields of SERFR.


**Table 8-28. SERFR Field Descriptions**

Bits	Name	Description
0–31	INT <sub>n</sub>	Each implemented bit, listed in Table 8-21, corresponds to an external MCP source. The user forces an MCP by setting the SERFR bit. SERFR bit positions are not affected by their relative priority. Attempts to write to unimplemented (reserved) bits are ignored; read = 0

### 8.5.19 System Critical Interrupt Vector Register (SCVCR)

SCVCR, shown in Figure 8-23, contains a 7-bit code (Table 8-29) representing the unmasked critical interrupt (CINT) source of the highest priority level.

Note that in core-disabled mode the user should use SIVCR only to read an updated interrupt vector (SCVCR should not be used).



**Figure 8-23. System Critical Interrupt Vector Register (SCVCR)**

Table 8-29 defines SCVCR bit fields.


**Table 8-29. SCVCR Field Descriptions**

Bits	Name	Description
0–5	CVECx	Backward (MPC8260) compatible critical interrupt vector. Specifies a 6-bit unique number of the IPIC's highest priority critical interrupt source, pending to the core. When a critical interrupt request occurs, SCVCR can be read. If there are multiple critical interrupt sources, SCVCR latches the highest priority critical interrupt. Note that CVECx field will correctly reflect only first 64 interrupt vectors (See Table 8-6 for details). The value of SCVEC cannot change while it is being read.
6–24	—	Write ignored, read = 0
25–31	CVEC	Critical interrupt vector. Specifies a 7-bit unique number of the IPIC's highest priority critical interrupt source, pending to the core. When a critical interrupt request occurs, SCVCR can be read. If there are multiple critical interrupt sources, SCVCR latches the highest priority critical interrupt. Note that CVEC field correctly reflects all of the interrupt vectors (See Table 8-6 for details). The value of SCVEC cannot change while it is being read.

### 8.5.20 System Management Interrupt Vector Register (SMVCR)

SMVCR, shown in Figure 8-24, contains a 7-bit code (Table 8-30) representing the unmasked system management interrupt (SMI) source of the highest priority level.

Note that in core disabled mode the user should use SIVCR only read an updated interrupt vector (SMVCR should not be used).



**Figure 8-24. System Management Interrupt Vector Register (SMVCR)**

Table 8-30 defines the bit fields of SMVCR.

**Table 8-30. SMVCR Field Descriptions**

Bits	Name	Description
0–5	MVECx	Backward (MPC8260) compatible system management interrupt vector. Specifies a 6-bit unique number of the IPIC’s highest priority system management interrupt source, pending to the core. When a system management interrupt request occurs, SMVCR can be read. If there are multiple system management interrupt sources, SMVCR latches the highest priority system management interrupt. Note that MVECx f correctly reflects only the first 64 interrupt vectors (See Table 8-6 for details). The value of SMVEC cannot change while it is being read.
6–24	—	Write ignored, read = 0
25–31	MVEC	System management interrupt vector. Specifies a 7-bit unique number of the IPIC’s highest priority system management interrupt source, pending to the core. When a system management interrupt request occurs, SMVCR can be read. If there are multiple system management interrupt sources, SMVCR latches the highest priority system management interrupt. Note that MVEC field will correctly reflect all interrupt vectors (See Table 8-6 for details). The value of SMVEC cannot change while it is being read.

## 8.6 Functional Description

The following sections describe the types of interrupts, interrupt configurations, and their priorities.

### 8.6.1 Interrupt Types

The IPIC is responsible for receiving hardware-generated interrupts from different sources (both internal and external) along with prioritizing and delivering them to the CPU for servicing. The interrupt sources are controlled by the IPIC unit and may cause three types of exceptions in the processor core. The *int* signal is the main interrupt output from the IPIC to the processor core and causes the external interrupt exception. The *cint* signal is the critical interrupt output from the IPIC to the processor core and causes the critical external interrupt exception. The *smi* signal is the system management interrupt output from the IPIC to the processor core and causes the system management interrupt exception. The machine check exception is caused by the internal *mcp* signal generated by the IPIC, informing the processor of error conditions, assertion of the external MCP request, and other conditions.

## 8.6.2 Interrupt Configuration

Figure 8-25 shows the interrupt configuration.




Figure 8-25. Interrupt Structure

The interrupt controller allows masking of each interrupt source. When an unmasked interrupt source is pending in the SIPNR register, the interrupt controller sends an interrupt request to the core. When an interrupt is taken, the interrupt mask bit in the machine state register is cleared to disable further interrupt requests to the PowerPC core until software can handle them.

All interrupt sources are prioritized and bits are set in the system interrupt pending register (SIPNR, SEPNR) as interrupts occur regardless of whether they are masked in the IPIC. The prioritization of the interrupt sources is flexible within the following groups:

- The relative priority of the eTSEC1 Tx, eTSEC1 Rx, eTSEC1 Err, eTSEC2 TX, eTSEC2 Rx and USB DR internal interrupt signals can be modified.
- The relative priority of the UART1, UART2, SPI, I2C1, I2C2, eTSEC1 1588 timer, eTSEC2 1588 timer, and SEC internal interrupt signals can be modified.
- The relative priority of the IRQ0, IRQ1, IRQ2, and IRQ3 external interrupts, and RTC SEC and PCI internal interrupts can be modified.
- The relative priority of the IRQ4 external interrupts, and RTC ALR, MU, SBA, and DMA internal interrupts can be modified.
- One interrupt source can be assigned to be the programmable highest priority.

The SIVCR is updated with a 7-bit vector corresponding to the sub-block with the highest current priority.

### 8.6.3 Internal Interrupts Group Relative Priority

The relative priority in each internal group is programmable and can be changed dynamically. The group priorities are programmed in the IPIC internal interrupt priority registers (SIPRR<sub>x</sub>) and can be changed dynamically to implement a rotating priority.

In addition, the grouping of the locations of the interrupt entries has the following two options:

- **Grouped.** In the group scheme, all interrupts are grouped together at the top of [Table 8-31](#), ahead of most other interrupt sources. This scheme is ideal for applications where all interrupt sources function at a very high data rate and interrupt latency is very important.
- **Spread.** In the spread scheme, priorities are spread over [Table 8-31](#) so other sources can have lower interrupt latencies. This scheme is also programmed but cannot be changed dynamically.

### 8.6.4 Mixed Interrupts Group Relative Priority

The relative priority between up to four internal and four external interrupts in each group is programmable and can be changed dynamically. The group priorities are programmed in the IPIC mixed interrupt priority registers (SMPRR<sub>x</sub>) and can be changed dynamically to implement a rotating priority.

In addition, the grouping of the locations of the mixed interrupt entries has the following two options:

- **Grouped.** In the group scheme, all interrupts are grouped together at the top of the priority table, ahead of most other interrupt sources. See [Table 8-31](#) for more information. This scheme is ideal for applications where all interrupt sources function at a very high data rate and interrupt latency is very important.
- **Spread.** In the spread scheme, priorities are spread over the table so other sources can have lower interrupt latencies. This scheme is also programmed but cannot be changed dynamically.

## 8.6.5 Highest Priority Interrupt

In addition to the group relative priority option, SICFR[HPI] can be used to specify one interrupt source as having the highest priority. This interrupt remains within the same interrupt level as the other interrupt controller interrupts, but is serviced before any other interrupt in [Table 8-31](#).

If the highest priority feature is not used, the IPIC selects the interrupt request in MIXA0 to be the highest priority interrupt and the standard interrupt priority order is used from [Table 8-31](#). SICFR[HPI] can be updated dynamically to allow the user to change a normally low priority source into a high priority-source for a period as needed.

## 8.6.6 Interrupt Source Priorities

Each of the IPIC's internal and external interrupt sources can independently assert one interrupt request to the core. [Table 8-31](#) shows the prioritization of these interrupt sources. As described in previous sections, flexibility exists in the relative ordering of the interrupts, but, in general, relative priorities are as shown. A single interrupt priority number is associated with each table entry.

**Table 8-31. 7Interrupt Source Priority Levels**

Priority	Interrupt Source Description
1	Highest
2	MIXA0 (Grouped/Spread)
3	MIXA1 (Grouped)
4	MIXA2 (Grouped)
5	MIXA3 (Grouped)
6	MIXB0 (Spread)
7–10	Reserved
11	MIXA1 (Spread)
12–15	Reserved
16	MIXB0 (Grouped)
17	MIXB1 (Grouped)
18	MIXB2 (Grouped)
19	MIXB3 (Grouped)
20	MIXB1 (Spread)
21	SYSA0 (Grouped)
22	SYSA1 (Grouped)
23	SYSA2 (Grouped)
24	SYSA3 (Grouped)
25	MIXA2 (Spread)
26	SYSA4 (Grouped)
27	SYSA5 (Grouped)
28	SYSA6 (Grouped)
29	SYSA7 (Grouped)

**Table 8-31. 7Interrupt Source Priority Levels (continued)**

Priority	Interrupt Source Description
30	MIXA4 (Grouped)
31	MIXA5 (Grouped)
32	MIXA6 (Grouped)
33	MIXA7 (Grouped)
34	MIXB2 (Spread)
35–38	Reserved
39	MIXA3 (Spread)
40–43	Reserved
44	MIXB4 (Grouped)
45	MIXB5 (Grouped)
46	MIXB6 (Grouped)
47	MIXB7 (Grouped)
48	MIXB3 (Spread)
49	SYSD0 (Grouped)
50	SYSD1 (Grouped)
51	SYSD2 (Grouped)
52	SYSD3 (Grouped)
53	MIXA4 (Spread)
54	SYSD4 (Grouped)
55	SYSD5 (Grouped)
56	SYSD6 (Grouped)
57	SYSD7 (Grouped)
58	MIXB4 (Spread)
59	GTM4
60	Reserved
61	SYSA0 (Spread)
62	GTM8
63	Reserved
64	SYSD0 (Spread)
65	Reserved
66	GPIO
67	MIXA5 (Spread)
68	Reserved
69	Reserved
70	SYSA1 (Spread)

**Table 8-31. 7Interrupt Source Priority Levels (continued)**

Priority	Interrupt Source Description
71	DDR
72	Reserved
73	SYSD1 (Spread)
74	Reserved
75	LBC
76	MIXB5 (Spread)
77	GTM2
78	Reserved
79	SYSA2 (Spread)
80	GTM6
81	Reserved
82	SYSD2 (Spread)
83	Reserved
84	PMC
85	MIXA6 (Spread)
86	Reserved
87	Reserved
88	SYSA3 (Spread)
89	Reserved
90	Reserved
91	SYSD3 (Spread)
92	Reserved
93	Reserved
94	MIXB6 (Spread)
95	GTM3
96	Reserved
97	SYSA4 (Spread)
98	GTM7
99	Reserved
100	SYSD4 (Spread)
101	Reserved
102	Reserved
103	MIXA7 (Spread)
104	Reserved
105	Reserved
106	SYSA5 (Spread)

**Table 8-31. 7Interrupt Source Priority Levels (continued)**

Priority	Interrupt Source Description
107	Reserved
108	Reserved
109	SYSD5 (Spread)
110	Reserved
111	Reserved
112	MIXB7 (Spread)
113	GTM1
114	Reserved
115	SYSA6 (Spread)
116	GTM5
117	Reserved
118	SYSD6 (Spread)
119	Reserved
120	Reserved
121	Reserved
122	Reserved
123	SYSA7 (Spread)
124	Reserved
125	Reserved
126	SYSD7 (Spread)
127	Reserved
128	Reserved

### 8.6.7 Masking Interrupt Sources

By programming the system interrupt mask registers, SIMSR $x$  and SEMSR, the user can mask interrupt requests to the core. Each SIMSR $x$  and SEMSR bit corresponds to an interrupt source. To enable an interrupt, set the corresponding SIMSR or SEMSR bit. When a masked interrupt source has a pending interrupt request, the corresponding SIPNR $x$  or SEMSR bit is set, even though the interrupt is not generated to the core. The user can mask all interrupt sources to implement a polling interrupt servicing scheme.

When an interrupt source has multiple interrupting events, the user can individually mask these events by programming a mask register within that particular block. [Table 8-31](#) shows which interrupt sources have multiple interrupting events.



Figure 8-26 shows an example of how the masking occurs, using a DDR block as an example.

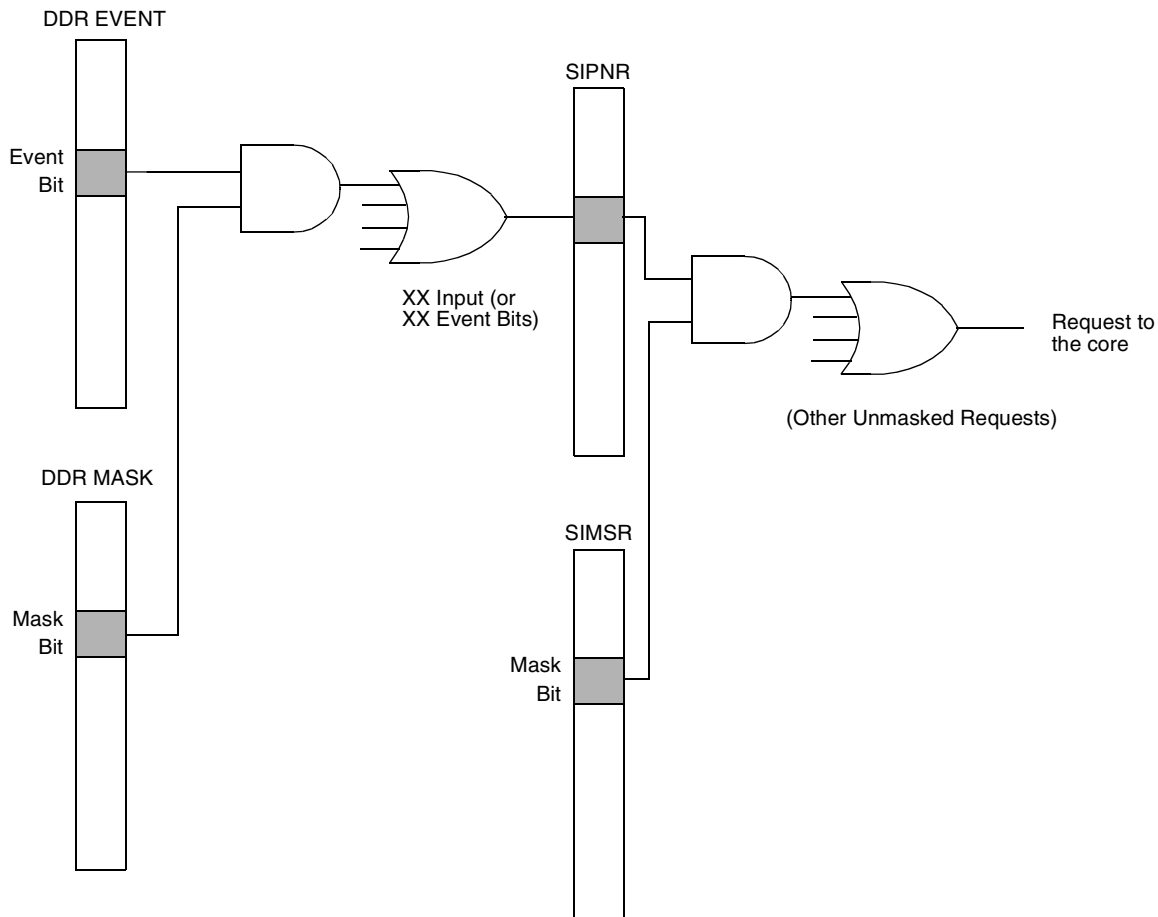


Figure 8-26. DDR Interrupt Request Masking

### 8.6.8 Interrupt Vector Generation and Calculation

Pending unmasked interrupts are presented to the core in order of priority according to [Table 8-31](#). The interrupt vector that allows the core to locate the interrupt service routine is made available to the core by interrupt handler software reading SIVCR. The interrupt controller passes an interrupt vector corresponding to the highest-priority, unmasked, pending interrupt in response to a read of SIVCR. [Table 8-5](#) lists the encodings for the seven low-order bits of the interrupt vector.

### 8.6.9 Machine Check Interrupts

The PIC supports the non-maskable machine check interrupts. When an error interrupt signal is received, the interrupt controller indicates the source by setting the corresponding SERSR bit. These sources are listed in [Table 8-21](#).



# Chapter 9

## DDR Memory Controller

### 9.1 Introduction

The fully programmable DDR SDRAM controller supports most JEDEC standard x8, x16, or x32 DDR2 and DDR memories available. In addition, unbuffered and registered DRAM modules are supported. However, mixing different memory types or unbuffered and registered DRAM modules in the same system is not supported. Dynamic power management and auto-precharge modes simplify memory system design. A large set of special features support rapid system debug.

#### NOTE

In this chapter, the word ‘bank’ refers to a physical bank specified by a chip select; ‘logical bank’ refers to one of the four or eight sub-banks in each SDRAM chip. A sub-bank is specified by the 2 or 3 bits on the bank address (MBA) pins during a memory access.

Figure 9-1 is a high-level block diagram of the DDR memory controller with its associated interfaces. Section 9.5, “Functional Description,” contains detailed figures of the controller.




Figure 9-1. DDR Memory Controller Simplified Block Diagram

## 9.2 Features

The DDR memory controller includes these distinctive features:

- Support for DDR2 and DDR SDRAM
- 32-bit SDRAM, 16-bit SDRAM for DDR and DDR2
- Programmable settings for meeting all SDRAM timing parameters
- The following SDRAM configurations are supported:
  - As many as two physical banks (chip selects), each bank independently addressable
  - 64-Mbit to 4-Gbit devices depending on internal device configuration with x8/x16/x32 data ports (no direct x4 support)
  - Unbuffered and registered DRAM modules
- Open page management (dedicated entry for each logical bank)
- Automatic DRAM initialization sequence or software-controlled initialization sequence
- Automatic DRAM data initialization

- Support for up to eight posted refreshes
- Memory controller clock frequency of two times the SDRAM clock with support for sleep power management

## 9.2.1 Modes of Operation

The DDR memory controller supports the following modes:

- 32-byte cache line wrap mode
- Dynamic power management mode. The DDR memory controller can reduce power consumption by negating the SDRAM CKE signal when no transactions are pending to the SDRAM.
- Auto-precharge mode. Clearing DDR\_SDRAM\_INTERVAL[BSTOPRE] causes the memory controller to issue an auto-precharge command with every read or write transaction. Auto-precharge mode can be enabled for separate chip selects by setting CS<sub>n</sub>\_CONFIG[AP\_n\_EN].

## 9.3 External Signal Descriptions

This section provides descriptions of the DDR memory controller's external signals. It describes each signal's behavior when the signal is asserted or negated and when the signal is an input or an output.

### 9.3.1 Signals Overview

Memory controller signals are grouped as follows:

- Memory interface signals
- Clock signals
- Debug signals

Table 9-1 shows how DDR memory controller external signals are grouped. The device hardware specification has a pinout diagram showing pin numbers. It also lists all electrical and mechanical specifications.

**Table 9-1. DDR Memory Interface Signal Summary**

Name	Function/Description	Reset	Pins	I/O
MDQ[0:31]	Data bus	All zeros	32	I/O
MDQS[0:3]	Data strobes	All zeros	4	I/O
$\overline{\text{MCAS}}$	Column address strobe	One	1	O
MA[14:0]	Address bus	All zeros	15	O
MBA[2:0]	Logical bank address	All zeros	3	O
$\overline{\text{MCS}}$ [0:1]	Chip selects	All ones	2	O
$\overline{\text{MWE}}$	Write enable	One	1	O
$\overline{\text{MRAS}}$	Row address strobe	One	1	O

Table 9-1. DDR Memory Interface Signal Summary (continued)

Name	Function/Description	Reset	Pins	I/O
MDM[0:3]	Data mask	All zeros	4	O
MCK	DRAM clock outputs	Zero	1	O
$\overline{\text{MCK}}$	DRAM clock outputs (complement)	One	1	O
MCKE	DRAM clock enable	Zero	1	O
MODT[0:1]	DRAM on-die termination external control.	All zeros	2	O
MDVAL	Memory debug data valid	Zero	1	O
MSRCID[0:4]	Memory debug source ID	All zeros	5	O

Table 9-2 shows the memory address signal mappings.

Table 9-2. Memory Address Signal Mappings

Signal Name (Outputs)		JEDEC DDR DIMM Signals (Inputs)
msb	MA14	A14
	MA13	A13
	MA12	A12
	MA11	A11
	MA10	A10 (AP for DDR) <sup>1</sup>
	MA9	A9
	MA8	A8 (alternate AP for DDR) <sup>2</sup>
	MA7	A7
	MA6	A6
	MA5	A5
	MA4	A4
	MA3	A3
	MA2	A2
	MA1	A1
lsb	MA0	A0
msb	MBA2	MBA2
	MBA1	MBA1
lsb	MBA0	MBA0

<sup>1</sup> Auto-precharge for DDR signaled on A10 when DDR\_SDRAM\_CFG[PCHB8] = 0.

<sup>2</sup> Auto-precharge for DDR signaled on A8 when DDR\_SDRAM\_CFG[PCHB8] = 1.

## 9.3.2 Detailed Signal Descriptions

The following sections describe the DDR SDRAM controller input and output signals, the meaning of their different states, and relative timing information for assertion and negation.

### 9.3.2.1 Memory Interface Signals

Table 9-3 describes the DDR controller memory interface signals.

**Table 9-3. Memory Interface Signals—Detailed Signal Descriptions**

Signal	I/O	Description	
MDQ[0:31]	I/O	Data bus. Both input and output signals on the DDR memory controller.	
	O	As outputs for the bidirectional data bus, these signals operate as described below.	
		<b>State Meaning</b>	Asserted/Negated—Represent the value of data being driven by the DDR memory controller.
		<b>Timing</b>	Assertion/Negation—Driven coincident with corresponding data strobes (MDQS) signal. High impedance—No READ or WRITE command is in progress; data is not being driven by the memory controller or the DRAM.
	I	As inputs for the bidirectional data bus, these signals operate as described below.	
		<b>State Meaning</b>	Asserted/Negated—Represents the state of data being driven by the external DDR SDRAMs.
<b>Timing</b>		Assertion/Negation—The DDR SDRAM drives data during a READ transaction. High impedance—No READ or WRITE command in progress; data is not being driven by the memory controller or the DRAM.	
MDQS[0:3]	I/O	Data strobes. Inputs with read data, outputs with write data.	
		O	As outputs, the data strobes are driven by the DDR memory controller during a write transaction. The memory controller always drives these signals low unless a read has been issued and incoming data strobes are expected. This keeps the data strobes from floating high when there are no transactions on the DRAM interface.
			<b>State Meaning</b>
	<b>Timing</b>	Assertion/Negation—If a WRITE command is registered at clock edge $n$ , data strobes at the DRAM assert centered in the data eye on clock edge $n + 1$ . See the JEDEC DDR SDRAM specification for more information.	
	I	As inputs, the data strobes are driven by the external DDR SDRAMs during a read transaction. The data strobes are used by the memory controller to synchronize data latching.	
		<b>State Meaning</b>	Asserted/Negated—Driven high when positive capture data is received and driven low when negative capture data is received. Centered in the data eye for writes; coincident with the data eye for reads. Treated as a clock. Data is valid when signals toggle. See Table 9-24 for byte lane assignments.
<b>Timing</b>		Assertion/Negation—If a READ command is registered at clock edge $n$ , and the latency is programmed in TIMING_CFG_1[CASLAT] to be $m$ clocks, data strobes at the DRAM assert coincident with the data on clock edge $n + m$ . See the JEDEC DDR SDRAM specification for more information.	

**Table 9-3. Memory Interface Signals—Detailed Signal Descriptions (continued)**

Signal	I/O	Description	
MA[14:0]	O	Address bus. Memory controller outputs for the address to the DRAM. MA[14:0] carry 15 of the address bits for the DDR memory interface corresponding to the row and column address bits. MA0 is the lsb of the address output from the memory controller.	
		<b>State Meaning</b>	Asserted/Negated—Represents the address driven by the DDR memory controller. Contains different portions of the address depending on the memory size and the DRAM command being issued by the memory controller. See <a href="#">Table 9-27</a> and <a href="#">Table 9-28</a> for a complete description of the mapping of these signals.
		<b>Timing</b>	Assertion/Negation—The address is always driven when the memory controller is enabled. It is valid when a transaction is driven to DRAM (when $\overline{MCSn}$ is active). High impedance—When the memory controller is disabled
MBA[2:0]	O	Logical bank address. Outputs that drive the logical (or internal) bank address pins of the SDRAM. Each SDRAM supports four or eight addressable logical sub-banks. Bit zero of the memory controller's output bank address must be connected to bit zero of the SDRAM's input bank address. MBA0, the least-significant bit of the three bank address signals, is asserted during the mode register set command to specify the extended mode register.	
		<b>State Meaning</b>	Asserted/Negated—Selects the DDR SDRAM logical (or internal) bank to be activated during the row address phase and selects the SDRAM internal bank for the read or write operation during the column address phase of the memory access. <a href="#">Table 9-27</a> and <a href="#">Table 9-28</a> describes the mapping of these signals in all cases.
		<b>Timing</b>	Assertion/Negation—Same timing as $MA_n$ High impedance—Same timing as $MA_n$
$\overline{MCAS}$	O	Column address strobe. Active-low SDRAM address multiplexing signal. $\overline{MCAS}$ is asserted for read or write transactions and for mode register set, refresh, and precharge commands.	
		<b>State Meaning</b>	Asserted—Indicates that a valid SDRAM column address is on the address bus for read and write transactions. See <a href="#">Table 9-30</a> for more information on the states required on $\overline{MCAS}$ for various other SDRAM commands. Negated—The column address is not guaranteed to be valid.
		<b>Timing</b>	Assertion/Negation—Assertion and negation timing is directed by the values described in <a href="#">Section 9.4.1.4, "DDR SDRAM Timing Configuration 0 (TIMING_CFG_0),"</a> <a href="#">Section 9.4.1.5, "DDR SDRAM Timing Configuration 1 (TIMING_CFG_1),"</a> <a href="#">Section 9.4.1.6, "DDR SDRAM Timing Configuration 2 (TIMING_CFG_2),"</a> and <a href="#">Section 9.4.1.3, "DDR SDRAM Timing Configuration 3 (TIMING_CFG_3)." </a> High impedance— $\overline{MCAS}$ is always driven unless the memory controller is disabled.
$\overline{MRAS}$	O	Row address strobe. Active-low SDRAM address multiplexing signal. Asserted for activate commands. In addition; used for mode register set commands and refresh commands.	
		<b>State Meaning</b>	Asserted—Indicates that a valid SDRAM row address is on the address bus for read and write transactions. See <a href="#">Table 9-30</a> for more information on the states required on $\overline{MRAS}$ for various other SDRAM commands. Negated—The row address is not guaranteed to be valid.
		<b>Timing</b>	Assertion/Negation—Assertion and negation timing is directed by the values described in <a href="#">Section 9.4.1.4, "DDR SDRAM Timing Configuration 0 (TIMING_CFG_0),"</a> <a href="#">Section 9.4.1.5, "DDR SDRAM Timing Configuration 1 (TIMING_CFG_1),"</a> <a href="#">Section 9.4.1.6, "DDR SDRAM Timing Configuration 2 (TIMING_CFG_2),"</a> and <a href="#">Section 9.4.1.3, "DDR SDRAM Timing Configuration 3 (TIMING_CFG_3)." </a> High impedance— $\overline{MRAS}$ is always driven unless the memory controller is disabled.



Table 9-3. Memory Interface Signals—Detailed Signal Descriptions (continued)

Signal	I/O	Description	
$\overline{\text{MCS}}[0:1]$	O	Chip selects. Two chip selects supported by the memory controller.	
		<b>State Meaning</b>	Asserted—Selects a physical SDRAM bank to perform a memory operation as described in Section 9.4.1.1, “Chip Select Memory Bounds (CS <sub>n</sub> _BNDS),” and Section 9.4.1.2, “Chip Select Configuration (CS <sub>n</sub> _CONFIG).” The DDR controller asserts one of the $\overline{\text{MCS}}[0:1]$ signals to begin a memory cycle. Negated—Indicates no SDRAM action during the current cycle.
		<b>Timing</b>	Assertion/Negation—Asserted to signal any new transaction to the SDRAM. The transaction must adhere to the timing constraints set in TIMING_CFG_0–TIMING_CFG_3. High impedance—Always driven unless the memory controller is disabled.
$\overline{\text{MWE}}$	O	Write enable. Asserted when a write transaction is issued to the SDRAM. This is also used for mode registers set commands and precharge commands.	
		<b>State Meaning</b>	Asserted—Indicates a memory write operation. See Table 9-30 for more information on the states required on $\overline{\text{MWE}}$ for various other SDRAM commands. Negated—Indicates a memory read operation.
		<b>Timing</b>	Assertion/Negation—Similar timing as $\overline{\text{MRA}}\overline{\text{S}}$ and $\overline{\text{MCAS}}$ . Used for write commands. High impedance— $\overline{\text{MWE}}$ is always driven unless the memory controller is disabled.
MDM[0:3]	O	DDR SDRAM data output mask. Masks unwanted bytes of data transferred during a write. They are needed to support sub-burst-size transactions (such as single-byte writes) on SDRAM where all I/O occurs in multi-byte bursts. MDM <sub>0</sub> corresponds to the most significant byte (MSB). Table 9-24 shows byte lane encodings.	
		<b>State Meaning</b>	Asserted—Prevents writing to DDR SDRAM. Asserted when data is written to DRAM if the corresponding byte(s) should be masked for the write. Note that the MDM <sub>n</sub> signals are active-high for the DDR controller. MDM <sub>n</sub> is part of the DDR command encoding. Negated—Allows the corresponding byte to be read from or written to the SDRAM.
		<b>Timing</b>	Assertion/Negation—Same timing as MDQx as outputs. High-impedance—Always driven unless the memory controller is disabled.
MODT[0:1]	O	On-Die termination. Memory controller outputs for the ODT to the DRAM. MODT[0:1] represents the on-die termination for the associated data, data masks, and data strobes.	
		<b>State Meaning</b>	Asserted/Negated—Represents the ODT driven by the DDR memory controller.
		<b>Timing</b>	Assertion/Negation—Driven in accordance with JEDEC DRAM specifications for on-die termination timings. It is configured through the CS <sub>n</sub> _CONFIG[ODT_RD_CFG] and CS <sub>n</sub> _CONFIG[ODT_WR_CFG] fields. High impedance—Always driven.

### 9.3.2.2 Clock Interface Signals

Table 9-4 contains the detailed descriptions of the clock signals of the DDR controller.

Table 9-4. Clock Signals—Detailed Signal Descriptions

Signal	I/O	Description	
$\overline{\text{MCK}}$ , MCK	O	DRAM clock output and its complement. See Section 9.5.4.1, “Clock Distribution.”	
		<b>State Meaning</b>	Asserted/Negated—The JEDEC DDR SDRAM specifications require true and complement clocks. A clock edge is seen by the SDRAM when the true and complement cross.
		<b>Timing</b>	Assertion/Negation—Timing is controlled by the DDR_CLK_CNTL register at offset 0x130.

**Table 9-4. Clock Signals—Detailed Signal Descriptions (continued)**

Signal	I/O	Description	
MCKE	O	Clock enable. Output signals used as the clock enables to the SDRAM. MCKE can be negated to stop clocking the DDR SDRAM. The MCKE signals should be connected to the same rank of memory as the corresponding $\overline{MCS}$ and $\overline{MODT}$ signals. For example, MCKE[0] should be connected to the same rank of memory as $\overline{MCS}[0]$ and $\overline{MODT}[0]$ .	
		<b>State Meaning</b>	Asserted—Clocking to the SDRAM is enabled. Negated—Clocking to the SDRAM is disabled and the SDRAM should ignore signal transitions on MCK or $\overline{MCK}$ . MCK/ $\overline{MCK}$ are don't cares while MCKE is negated.
		<b>Timing</b>	Assertion/Negation—Asserted when DDR_SDRAM_CFG[MEM_EN] is set. Can be negated when entering dynamic power management or self refresh. Are asserted again when exiting dynamic power management or self refresh. High impedance—Always driven.

### 9.3.2.3 Debug Signals

The debug signals MSRCID[0:4] and MDVAL have no function in normal DDR controller operation. A detailed description of these signals can be found in [Section 5.4.3.8, “Debug Configuration.”](#)

## 9.4 Memory Map/Register Definition

[Table 9-5](#) shows the register memory map for the DDR memory controller.

In this table and in the register figures and field descriptions, the following access definitions apply:

- Reserved fields are always ignored for the purposes of determining access type.
- R/W, R, and W (read/write, read only, and write only) indicate that all the non-reserved fields in a register have the same access type.
- w1c indicates that all of the non-reserved fields in a register are cleared by writing ones to them.
- Mixed indicates a combination of access types.
- Special is used when no other category applies. In this case the register figure and field description table should be read carefully.

**Table 9-5. DDR Memory Controller Memory Map**

Offset	Register	Access	Reset	Section/Page
<b>DDR Memory Controller—Block Base Address 0x0_2000</b>				
0x000	CS0_BNDS—Chip select 0 memory bounds	R/W	0x0000_0000	<a href="#">9.4.1.1/9-9</a>
0x008	CS1_BNDS—Chip select 1 memory bounds	R/W	0x0000_0000	<a href="#">9.4.1.1/9-9</a>
0x080	CS0_CONFIG—Chip select 0 configuration	R/W	0x0000_0000	<a href="#">9.4.1.2/9-10</a>
0x084	CS1_CONFIG—Chip select 1 configuration	R/W	0x0000_0000	<a href="#">9.4.1.2/9-10</a>
0x100	TIMING_CFG_3—DDR SDRAM timing configuration 3	R/W	0x0000_0000	<a href="#">9.4.1.3/9-11</a>
0x104	TIMING_CFG_0—DDR SDRAM timing configuration 0	R/W	0x0011_0105	<a href="#">9.4.1.4/9-12</a>
0x108	TIMING_CFG_1—DDR SDRAM timing configuration 1	R/W	0x0000_0000	<a href="#">9.4.1.5/9-14</a>
0x10C	TIMING_CFG_2—DDR SDRAM timing configuration 2	R/W	0x0000_0000	<a href="#">9.4.1.6/9-16</a>

Table 9-5. DDR Memory Controller Memory Map (continued)

Offset	Register	Access	Reset	Section/Page
0x110	DDR_SDRAM_CFG—DDR SDRAM control configuration	R/W	0x0200_0000	9.4.1.7/9-18
0x114	DDR_SDRAM_CFG_2—DDR SDRAM control configuration 2	R/W	0x0000_0000	9.4.1.8/9-21
0x118	DDR_SDRAM_MODE—DDR SDRAM mode configuration	R/W	0x0000_0000	9.4.1.9/9-22
0x11C	DDR_SDRAM_MODE_2—DDR SDRAM mode configuration 2	R/W	0x0000_0000	9.4.1.10/9-23
0x120	DDR_SDRAM_MD_CNTL—DDR SDRAM mode control	R/W	0x0000_0000	9.4.1.11/9-24
0x124	DDR_SDRAM_INTERVAL—DDR SDRAM interval configuration	R/W	0x0000_0000	9.4.1.12/9-26
0x128	DDR_DATA_INIT—DDR SDRAM data initialization	R/W	0x0000_0000	9.4.1.13/9-27
0x130	DDR_SDRAM_CLK_CNTL—DDR SDRAM clock control	R/W	0x0200_0000	9.4.1.14/9-27
0x140– 0x144	Reserved	—	—	—
0x148	DDR_INIT_ADDR—DDR training initialization address	R/W	0x0000_0000	9.4.1.15/9-28
0x150– 0xBF4	Reserved	—	—	—
0xBF8	DDR_IP_REV1—DDR IP block revision 1	R	0xn <sub>nnn</sub> _n <sub>nnn</sub> <sup>1</sup>	9.4.1.16/9-28
0xBFC	DDR_IP_REV2—DDR IP block revision 2	R	0x00n <sub>n</sub> _00n <sub>n</sub> <sup>1</sup>	9.4.1.17/9-29

<sup>1</sup> Implementation-dependent reset values are listed in specified section/page.

## 9.4.1 Register Descriptions

This section describes the DDR memory controller registers. Shading indicates reserved fields that should not be written.

### 9.4.1.1 Chip Select Memory Bounds (CS<sub>n</sub>\_BNDS)

The chip select bounds registers (CS<sub>n</sub>\_BNDS) define the starting and ending address of the memory space that corresponds to the individual chip selects. Note that the size specified in CS<sub>n</sub>\_BNDS should equal the size of physical DRAM. Also, note that EAn must be greater than or equal to SAn.

If chip select interleaving is enabled, all fields in the lower interleaved chip select are used, and the other chip selects' bounds registers are unused. For example, if chip selects 0 and 1 are interleaved, all fields in CS<sub>0</sub>\_BNDS are used, and all fields in CS<sub>1</sub>\_BNDS are unused.

CS<sub>n</sub>\_BNDS are shown in Figure 9-2.




Figure 9-2. Chip Select Bounds Registers (CS<sub>n</sub>\_BNDS)

Table 9-6 describes the CS<sub>n</sub>\_BNDS register fields.

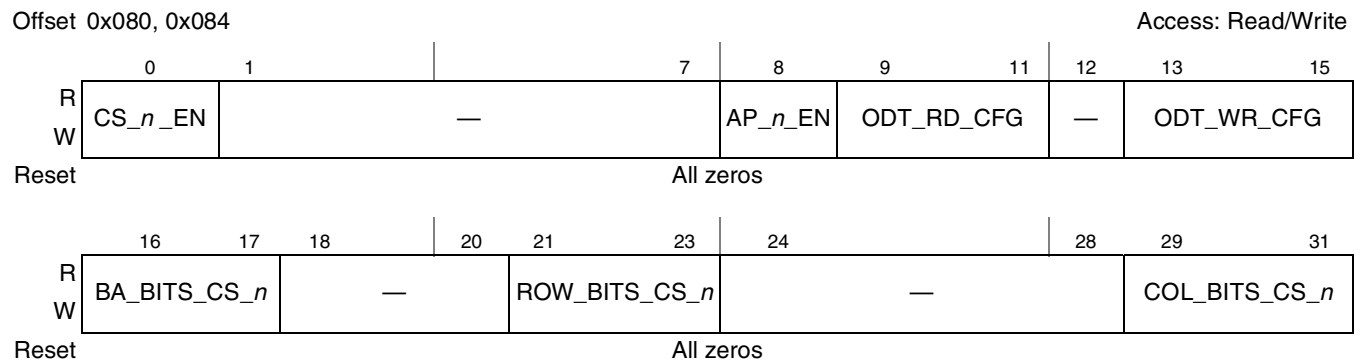
**Table 9-6. CS<sub>n</sub>\_BNDS Field Descriptions**

Bits	Name	Description
0–7	—	Reserved
8–15	SA <sub>n</sub>	Starting address for chip select (bank) <i>n</i> . This value is compared against the 8 msbs of the 32-bit address.
16–23	—	Reserved
24–31	EA <sub>n</sub>	Ending address for chip select (bank) <i>n</i> . This value is compared against the 8 msbs of the 32-bit address.

### 9.4.1.2 Chip Select Configuration (CS<sub>n</sub>\_CONFIG)

The chip select configuration (CS<sub>n</sub>\_CONFIG) registers shown in Figure 9-3 enable the DDR chip selects and set the number of row and column bits used for each chip select. These registers should be loaded with the correct number of row and column bits for each SDRAM. Because CS<sub>n</sub>\_CONFIG[ROW\_BITS\_CS<sub>n</sub>, COL\_BITS\_CS<sub>n</sub>] establish address multiplexing, the user should take great care to set these values correctly.

If chip select interleaving is enabled, then all fields in the lower interleaved chip select are used, and the other registers' fields are unused, with the exception of the ODT\_RD\_CFG and ODT\_WR\_CFG fields. For example, if chip selects 0 and 1 are interleaved, all fields in CS<sub>0</sub>\_CONFIG are used, but only the ODT\_RD\_CFG and ODT\_WR\_CFG fields in CS<sub>1</sub>\_CONFIG are used.



**Figure 9-3. Chip Select Configuration Register (CS<sub>n</sub>\_CONFIG)**

Table 9-7 describes the CS<sub>n</sub>\_CONFIG register fields.

**Table 9-7. CS<sub>n</sub>\_CONFIG Field Descriptions**

Bits	Name	Description
0	CS <sub>n</sub> _EN	Chip select <i>n</i> enable 0 Chip select <i>n</i> is not active 1 Chip select <i>n</i> is active and assumes the state set in CS <sub>n</sub> _BNDS.
1–7	—	Reserved


Table 9-7. CS<sub>n</sub>\_CONFIG Field Descriptions (continued)

Bits	Name	Description
8	AP <sub>n</sub> _EN	Chip select <i>n</i> auto-precharge enable 0 Chip select <i>n</i> is only auto-precharged if global auto-precharge mode is enabled (DDR_SDRAM_INTERVAL[BSTOPRE] = 0). 1 Chip select <i>n</i> always issues an auto-precharge for read and write transactions.
9–11	ODT_RD_CFG	ODT for reads configuration. Note that CAS latency plus additive latency must be at least 3 cycles for ODT_RD_CFG to be enabled. ODT should only be used with DDR2 memories. 000 Never assert ODT for reads 001 Assert ODT only during reads to CS <sub>n</sub> 010 Assert ODT only during reads to other chip selects 011 Reserved 100 Assert ODT for all reads 101–111 Reserved
12	—	Reserved
13–15	ODT_WR_CFG	ODT for writes configuration. Note that write latency plus additive latency must be at least 3 cycles for ODT_WR_CFG to be enabled. ODT should only be used with DDR2 memories. 000 Never assert ODT for writes 001 Assert ODT only during writes to CS <sub>n</sub> 010 Assert ODT only during writes to other chip selects 011 Reserved 100 Assert ODT for all writes 101–111 Reserved
16–17	BA_BITS_CS <sub>n</sub>	Number of bank bits for SDRAM on chip select <i>n</i> . These bits correspond to the sub-bank bits driven on MBA <sub>n</sub> in Table 9-27 and Table 9-28. 00 2 logical bank bits 01 3 logical bank bits 10–11 Reserved
18–20	—	Reserved
21–23	ROW_BITS_CS <sub>n</sub>	Number of row bits for SDRAM on chip select <i>n</i> . See Table 9-27 and Table 9-28 for details. 000 12 row bits 001 13 row bits 010 14 row bits 011 15 row bits 000–111 Reserved
24–28	—	Reserved
29–31	COL_BITS_CS <sub>n</sub>	Number of column bits for SDRAM on chip select <i>n</i> . For DDR, the decoding is as follows: 000 8 column bits 001 9 column bits 010 10 column bits 011 11 column bits 100–111 Reserved

### 9.4.1.3 DDR SDRAM Timing Configuration 3 (TIMING\_CFG\_3)

DDR SDRAM timing configuration register 3, shown in Figure 9-4, sets the extended refresh recovery time, which is combined with TIMING\_CFG\_1[REFREC] to determine the full refresh recovery time.

## DDR Memory Controller



**Figure 9-4. DDR SDRAM Timing Configuration 3 (TIMING\_CFG\_3)**


Table 9-8 describes TIMING\_CFG\_3 fields.

**Table 9-8. TIMING\_CFG\_3 Field Descriptions**

Bits	Name	Description
0–12	—	Reserved, should be cleared.
13–15	EXT_REFREC	Extended refresh recovery time ( $t_{RFC}$ ). Controls the number of clock cycles from a refresh command until an activate command is allowed. This field is concatenated with TIMING_CFG_1[REFREC] to obtain a 7-bit value for the total refresh recovery. Note that hardware adds an additional 8 clock cycles to the final, 7-bit value of the refresh recovery. $t_{RFC} = \{EXT\_REFREC \parallel REFREC\} + 8$ , such that $t_{RFC}$ is calculated as follows:  000 0 clocks 001 16 clocks 010 32 clocks 011 48 clocks 100 64 clocks 101 80 clocks 110 96 clocks 111 112 clocks
16–31	—	Reserved, should be cleared.

### 9.4.1.4 DDR SDRAM Timing Configuration 0 (TIMING\_CFG\_0)

DDR SDRAM timing configuration register 0, shown in Figure 9-5, sets the number of clock cycles between various SDRAM control commands.



**Figure 9-5. DDR SDRAM Timing Configuration 0 (TIMING\_CFG\_0)**

Table 9-9 describes TIMING\_CFG\_0 fields.

**Table 9-9. TIMING\_CFG\_0 Field Descriptions**

Bits	Name	Description
0–1	RWT	Read-to-write turnaround ( $t_{RTW}$ ). Specifies how many extra cycles are added between a read to write turnaround. If 0 clocks is chosen, then the DDR controller uses a fixed number based on the CAS latency and write latency. Choosing a value other than 0 adds extra cycles past this default calculation. As a default the DDR controller determines the read-to-write turnaround as $CL - WL + BL/2 + 2$ . In this equation, CL is the CAS latency rounded up to the next integer, WL is the programmed write latency, and BL is the burst length. 00 0 clocks 01 1 clock 10 2 clocks 11 3 clocks
2–3	WRT	Write-to-read turnaround. Specifies how many extra cycles are added between a write to read turnaround. If 0 clocks is chosen, then the DDR controller uses a fixed number based on the, read latency, and write latency. Choosing a value other than 0 adds extra cycles past this default calculation. As a default, the DDR controller determines the write-to-read turnaround as $WL - CL + BL/2 + 1$ . In this equation, CL is the CAS latency rounded down to the next integer, WL is the programmed write latency, and BL is the burst length. 00 0 clocks 01 1 clock 10 2 clocks 11 3 clocks
4–5	RRT	Read-to-read turnaround. Specifies how many extra cycles are added between reads to different chip selects. As a default, 3 cycles are required between read commands to different chip selects. Extra cycles may be added with this field. Note: If 8-beat bursts are enabled, then 5 cycles are the default. Note that DDR2 does not support 8-beat bursts. 00 0 clocks 01 1 clock 10 2 clocks 11 3 clocks
6–7	WWT	Write-to-write turnaround. Specifies how many extra cycles are added between writes to different chip selects. As a default, 2 cycles are required between write commands to different chip selects. Extra cycles may be added with this field. Note: If 8-beat bursts are enabled, then 4 cycles are the default. Note that DDR2 does not support 8-beat bursts. 00 0 clocks 01 1 clock 10 2 clocks 11 3 clocks
8	—	Reserved, should be cleared.
9–11	ACT_PD_EXIT	Active powerdown exit timing ( $t_{XARD}$ and $t_{XARDS}$ ). Specifies how many clock cycles to wait after exiting active powerdown before issuing any command. 000 Reserved 001 1 clock 010 2 clocks 011 3 clocks 100 4 clocks 101 5 clocks 110 6 clocks 111 7 clocks
12	—	Reserved, should be cleared.
13–15	PRE_PD_EXIT	Precharge powerdown exit timing ( $t_{XP}$ ). Specifies how many clock cycles to wait after exiting precharge powerdown before issuing any command. 000 Reserved 001 1 clock 010 2 clocks 011 3 clocks 100 4 clocks 101 5 clocks 110 6 clocks 111 7 clocks

**Table 9-9. TIMING\_CFG\_0 Field Descriptions (continued)**


Bits	Name	Description
16–19	—	Reserved, should be cleared.
20–23	ODT_PD_EXIT	ODT powerdown exit timing ( $t_{AXPD}$ ). Specifies how many clocks must pass after exiting powerdown before ODT may be asserted. 0000 0 clock 0001 1 clock 0010 2 clocks 0011 3 clocks 0100 4 clocks 0101 5 clocks 0110 6 clocks 0111 7 clocks 1000 8 clocks 1001 9 clocks 1010 10 clocks 1011 11 clocks 1100 12 clocks 1101 13 clocks 1110 14 clocks 1111 15 clocks
24–27	—	Reserved, should be cleared.
28–31	MRS_CYC	Mode register set cycle time ( $t_{MRD}$ ). Specifies the number of cycles that must pass after a Mode Register Set command until any other command. 0000 Reserved 0001 1 clock 0010 2 clocks 0011 3 clocks 0100 4 clocks 0101 5 clocks 0110 6 clocks 0111 7 clocks 1000–1111 Reserved

### 9.4.1.5 DDR SDRAM Timing Configuration 1 (TIMING\_CFG\_1)

DDR SDRAM timing configuration register 1, shown in Figure 9-6, sets the number of clock cycles between various SDRAM control commands.

Offset 0x108

Access: Read/Write



**Figure 9-6. DDR SDRAM Timing Configuration 1 (TIMING\_CFG\_1)**

Table 9-10 describes TIMING\_CFG\_1 fields.

**Table 9-10. TIMING\_CFG\_1 Field Descriptions**

Bits	Name	Description
0	—	Reserved, should be cleared.



Table 9-10. TIMING\_CFG\_1 Field Descriptions (continued)


Bits	Name	Description
1–3	PRETOACT	Precharge-to-activate interval ( $t_{RP}$ ). Determines the number of clock cycles from a precharge command until an activate or refresh command is allowed.  000 Reserved 001 1 clock 010 2 clocks 011 3 clocks 100 4 clocks 101 5 clocks 110 6 clocks 111 7 clocks
4–7	ACTTOPRE	Activate to precharge interval ( $t_{RAS}$ ). Determines the number of clock cycles from an activate command until a precharge command is allowed.  0000 16 clocks                      0101 5 clocks 0001 17 clocks                      0110 6 clocks 0010 18 clocks                      0111 7 clocks 0011 19 clocks                      ... 0100 4 clocks                        1111 15 clocks
8	—	Reserved, should be cleared.
9–11	ACTTORW	Activate to read/write interval for SDRAM ( $t_{RCD}$ ). Controls the number of clock cycles from an activate command until a read or write command is allowed.  000 Reserved 001 1 clock 010 2 clocks 011 3 clocks 100 4 clocks 101 5 clocks 110 6 clocks 111 7 clocks
12–15	CASLAT	$\overline{MCAS}$ latency from READ command. Number of clock cycles between registration of a READ command by the SDRAM and the availability of the first output data. If a READ command is registered at clock edge $n$ and the latency is $m$ clocks, data is available nominally coincident with clock edge $n + m$ . This value must be programmed at initialization as described in <a href="#">Section 9.4.1.8, “DDR SDRAM Control Configuration 2 (DDR_SDRAM_CFG_2).”</a>  0000 Reserved                      1000 4.5 clocks 0001 1 clock                        1001 5 clocks 0010 1.5 clocks                    1010 5.5 clocks 0011 2 clocks                      1011 6 clocks 0100 2.5 clocks                    1100 6.5 clocks 0101 3 clocks                      1101 7 clocks 0110 3.5 clocks                    1110 7.5 clocks 0111 4 clocks                      1111 8 clocks

**Table 9-10. TIMING\_CFG\_1 Field Descriptions (continued)**

Bits	Name	Description
16–19	REFREC	Refresh recovery time ( $t_{RFC}$ ). Controls the number of clock cycles from a refresh command until an activate command is allowed. This field is concatenated with TIMING_CFG_3[EXTREFREC] to obtain a 7-bit value for the total refresh recovery. Note that hardware adds an additional 8 clock cycles to the final, 7-bit value of the refresh recovery, such that $t_{RFC}$ is calculated as follows: $t_{RFC} = \{EXT\_REFREC \parallel REFREC\} + 8$ .  0000 8 clocks                      0011 11 clocks 0001 9 clocks                      ... 0010 10 clocks                     1111 23 clocks
20	—	Reserved, should be cleared.
21–23	WRREC	Last data to precharge minimum interval ( $t_{WR}$ ). Determines the number of clock cycles from the last data associated with a write command until a precharge command is allowed.  000 Reserved 001 1 clock 010 2 clocks 011 3 clocks 100 4 clocks 101 5 clocks 110 6 clocks 111 7 clocks
24	—	Reserved, should be cleared.
25–27	ACTTOACT	Activate-to-activate interval ( $t_{RRD}$ ). Number of clock cycles from an activate command until another activate command is allowed for a different logical bank in the same physical bank (chip select).  000 Reserved                      100 4 clocks 001 1 clock                        101 5 clocks 010 2 clocks                       110 6 clocks 011 3 clocks                       111 7 clocks
28	—	Reserved, should be cleared.
29–31	WRTORD	Last write data pair to read command issue interval ( $t_{WTR}$ ). Number of clock cycles between the last write data pair and the subsequent read command to the same physical bank.  000 Reserved                      100 4 clocks 001 1 clock                        101 5 clocks 010 2 clocks                       110 6 clocks 011 3 clocks                       111 7 clocks

**9.4.1.6 DDR SDRAM Timing Configuration 2 (TIMING\_CFG\_2)**

DDR SDRAM timing configuration 2, shown in Figure 9-7, sets the clock delay to data for writes.



**Figure 9-7. DDR SDRAM Timing Configuration 2 Register (TIMING\_CFG\_2)**

Table 9-11 describes the TIMING\_CFG\_2 fields.

**Table 9-11. TIMING\_CFG\_2 Field Descriptions**

Bits	Name	Description
0	—	Reserved
1–3	ADD_LAT	Additive latency. The additive latency must be set to a value less than TIMING_CFG_1[ACTTORW]. (DDR2-specific) 000 0 clocks 001 1 clock 010 2 clocks 011 3 clocks 100 4 clocks 101 5 clocks 110 Reserved 111 Reserved
4–8	CPO <sup>1</sup>	MCAS-to-preamble override. Defines the number of DRAM cycles between when a read is issued and when the corresponding DQS preamble is valid for the memory controller. For these decodings, “READ_LAT” is equal to the CAS latency plus the additive latency. 00000 READ_LAT + 1                      01100 READ_LAT + 5/2 00001 Reserved                            01101 READ_LAT + 11/4 00010 READ_LAT                            01110 READ_LAT + 3 00011 READ_LAT + 1/4                    01111 READ_LAT + 13/4 00100 READ_LAT + 1/2                    10000 READ_LAT + 7/2 00101 READ_LAT + 3/4                    10001 READ_LAT + 15/4 00110 READ_LAT + 1                       10010 READ_LAT + 4 00111 READ_LAT + 5/4                    10011 READ_LAT + 17/4 01000 READ_LAT + 3/2                    10100 READ_LAT + 9/2 01001 READ_LAT + 7/4                    10101 READ_LAT + 19/4 01010 READ_LAT + 2                       10110–11111 Reserved 01011 READ_LAT + 9/4
9	—	Reserved
10–12	WR_LAT	Write latency. Note that the total write latency for DDR2 is equal to WR_LAT + ADD_LAT; the write latency for DDR1 is 1. 000 Reserved 001 1 clock 010 2 clocks 011 3 clocks 100 4 clocks 101 5 clocks 110 6 clocks 111 7 clocks
13–15	—	Reserved
16–18	RD_TO_PRE	Read to precharge (t <sub>RTP</sub> ). For DDR2, with a non-zero ADD_LAT value, takes a minimum of ADD_LAT + t <sub>RTP</sub> cycles between read and precharge. For DDR1 with burst length of 4, must be set to 010; for DDR1 with burst length of 8, must be set to 100. 000 Reserved                                100 4 cycles 001 1 cycle                                    101–111 Reserved 010 2 cycles 011 3 cycles

Table 9-11. TIMING\_CFG\_2 Field Descriptions (continued)

Bits	Name	Description
19–21	WR_DATA_DELAY	Write command to write data strobe timing adjustment. Controls the amount of delay applied to the data and data strobes for writes. See Section 9.5.7, “DDR SDRAM Write Timing Adjustments,” for details. 000 0 clock delay                   100 1 clock delay 001 1/4 clock delay                101 5/4 clock delay 010 1/2 clock delay                110 3/2 clock delay 011 3/4 clock delay                111 Reserved
22	—	Reserved
23–25	CKE_PLS	Minimum CKE pulse width ( $t_{CKE}$ ) Can be set to 001 for DDR1. 000 Reserved                        011 3 cycles 001 1 cycle                            100 4 cycles 010 2 cycles                          101–111 Reserved
26–31	FOUR_ACT	Window for four activates ( $t_{FAW}$ ). This is applied to DDR2 with eight logical banks only. Must be set to 000001 for DDR1. 000000 Reserved                    ... 000001 1 cycle                        01001119 cycles 000010 2 cycles                      010100 20 cycles 000011 3 cycles                      010101–111111 Reserved 000100 4 cycles

<sup>1</sup> For CPO decodings other than 00000 and 11111, ‘READ\_LAT’ is rounded up to the next integer value.

### 9.4.1.7 DDR SDRAM Control Configuration (DDR\_SDRAM\_CFG)

The DDR SDRAM control configuration register, shown in Figure 9-8, enables the interface logic and specifies certain operating features such as self refreshing, error checking and correcting, registered DIMMs, and dynamic power management.

Offset 0x110

Access: Read/Write




Figure 9-8. DDR SDRAM Control Configuration Register (DDR\_SDRAM\_CFG)

Table 9-12 describes the DDR\_SDRAM\_CFG fields.

**Table 9-12. DDR\_SDRAM\_CFG Field Descriptions**

Bits	Name	Description
0	MEM_EN	DDR SDRAM interface logic enable. 0 SDRAM interface logic is disabled. 1 SDRAM interface logic is enabled. Must not be set until all other memory configuration parameters have been appropriately configured by initialization code.
1	SREN	Self refresh enable (during sleep). 0 SDRAM self refresh is disabled during sleep. Whenever self-refresh is disabled, the system is responsible for preserving the integrity of SDRAM during sleep. 1 SDRAM self refresh is enabled during sleep.
2	—	Reserved. Must be cleared.
3	RD_EN	Registered DRAM module enable. Specifies the type of DRAM module used in the system. 0 Indicates unbuffered DRAM modules. 1 Indicates registered DRAM modules. <b>Note:</b> RD_EN and 2T_EN must not both be set at the same time.
4	—	Reserved
5–7	SDRAM_TYPE	Type of SDRAM device to be used. This field is used when issuing the automatic hardware initialization sequence to DRAM through Mode Register Set and Extended Mode Register Set commands. Default value is 010 designating DDR1 SDRAM. 000–001 Reserved 010 DDR1 SDRAM 011 DDR2 SDRAM 100 Reserved 101 Reserved 110 Reserved 111 Reserved
8–9	—	Reserved
10	DYN_PWR	Dynamic power management mode 0 Dynamic power management mode is disabled. 1 Dynamic power management mode is enabled. If there is no ongoing memory activity, the SDRAM CKE signal is negated.
11–12	DBW	DRAM data bus width. 00 Reserved 01 32-bit bus is used 10 16-bit bus is used 11 Reserved
13	8_BE	8-beat burst enable. 0 4-beat bursts are used on the DRAM interface. 1 8-beat bursts are used on the DRAM interface. <b>Note:</b> DDR1 (SDRAM_TYPE = 010) must use 8-beat bursts when using 32-bit bus mode (32_BE = 1) and 4-beat bursts when using 64-bit bus mode; DDR2 (SDRAM_TYPE = 011) must use 4-beat bursts, even when using 32-bit bus mode
14	NCAP	Non-concurrent auto-precharge. Some older DDR DRAMs do not support concurrent auto precharge. If one of these devices is used, then this bit needs to be set if auto precharge is used. 0 DRAMs in system support concurrent auto-precharge. 1 DRAMs in system do not support concurrent auto-precharge.

Table 9-12. DDR\_SDRAM\_CFG Field Descriptions (continued)


Bits	Name	Description
15	—	Reserved
16	2T_EN	Enable 2T timing. 0 1T timing is enabled. The DRAM command/address are held for only 1 cycle on the DRAM bus. 1 2T timing is enabled. The DRAM command/address are held for 2 full cycles on the DRAM bus for every DRAM transaction. However, the chip select is only held for the second cycle. <b>Note:</b> RD_EN and 2T_EN must not both be set at the same time.
17–23	BA_INTLV_CTL	Bank (chip select) interleaving control. Set this field only if you wish to use bank interleaving. (All unlisted field values are reserved.) 0000000 No external memory banks are interleaved 1000000 External memory banks 0 and 1 are interleaved
24–25	—	Reserved
26	x32_EN	x32 enable. 0 Either x8 or x16 discrete DRAM chips are used. In this mode, each data byte has a dedicated corresponding data strobe. 1 x32 discrete DRAM chips are used. In this mode, DQS0 is used to capture DQ[0:31].
27	PCHB8	Precharge bit 8 enable. 0 MA[10] is used to indicate the auto-precharge and precharge all commands. 1 MA[8] is used to indicate the auto-precharge and precharge all commands. If x32_EN is cleared, then PCHB8 should be cleared as well.
28	HSE	Global half-strength override Sets I/O driver impedance to half strength. This impedance is used by the MDIC, address/command, data, and clock impedance values, but only if automatic hardware calibration is disabled and the corresponding group's software override is disabled in the DDR control driver register(s) described in <a href="#">Section 5.3.2.8, “DDR Control Driver Register (DDRCDR).”</a> This bit should be cleared if using automatic hardware calibration. 0 I/O driver impedance is configured to full strength. 1 I/O driver impedance is configured to half strength.
29	—	Reserved
30	MEM_HALT	DDR memory controller halt. When this bit is set, the memory controller does not accept any new data read/write transactions to DDR SDRAM until the bit is cleared again. This can be used when bypassing initialization and forcing MODE REGISTER SET commands through software. 0 DDR controller accepts new transactions. 1 DDR controller finishes any remaining transactions, and then it remains halted until this bit is cleared by software.
31	BI	Bypass initialization 0 DDR controller cycles through initialization routine based on SDRAM_TYPE 1 Initialization routine is bypassed. Software is responsible for initializing memory through DDR_SDRAM_MODE2 register. If software is initializing memory, then the MEM_HALT bit can be set to prevent the DDR controller from issuing transactions during the initialization sequence. Note that the DDR controller does not issue a DLL reset to the DRAMs when bypassing the initialization routine, regardless of the value of DDR_SDRAM_CFG[DLL_RST_DIS]. If a DLL reset is required, then the controller should be forced to enter and exit self refresh after the controller is enabled.

### 9.4.1.8 DDR SDRAM Control Configuration 2 (DDR\_SDRAM\_CFG\_2)

The DDR SDRAM control configuration register 2, shown in [Figure 9-9](#), provides more control configuration for the DDR controller.

Offset 0x114

Access: Read/Write



**Figure 9-9. DDR SDRAM Control Configuration Register 2 (DDR\_SDRAM\_CFG\_2)**

[Table 9-13](#) describes the DDR\_SDRAM\_CFG\_2 fields.

**Table 9-13. DDR\_SDRAM\_CFG\_2 Field Descriptions**


Bits	Name	Description
0	FRC_SR	Force self refresh 0 DDR controller operates in normal mode. 1 DDR controller enters self-refresh mode.
1	—	Reserved. Should be cleared.
2	DLL_RST_DIS	DLL reset disable. The DDR controller typically issues a DLL reset to the DRAMs when exiting self refresh. However, this function may be disabled by setting this bit during initialization. 0 DDR controller issues a DLL reset to the DRAMs when exiting self refresh. 1 DDR controller does not issue a DLL reset to the DRAMs when exiting self refresh.
3	—	Reserved
4–5	DQS_CFG	DQS configuration 00 Only true DQS signals are used. 01 Reserved 10 Reserved 11 Reserved
6–8	—	Reserved
9–10	ODT_CFG	ODT configuration. This field defines how ODT is driven to the on-chip IOs. See <a href="#">Section 5.4.4.12, “DDR Control Driver Register (DDRCDR),”</a> which defines the termination value that is used. (DDR2-specific, must be cleared for DDR1) 00 Never assert ODT to internal IOs 01 Assert ODT to internal IOs only during writes to DRAM 10 Assert ODT to internal IOs only during reads to DRAM 11 Always keep ODT asserted to internal IOs
11–15	—	Reserved.

**Table 9-13. DDR\_SDRAM\_CFG\_2 Field Descriptions (continued)**

Bits	Name	Description
16–19	NUM_PR	Number of posted refreshes. This determines how many posted refreshes, if any, can be issued at one time. Note that if posted refreshes are used, then this field, along with DDR_SDRAM_INTERVAL[REFINT], must be programmed such that the maximum $t_{ras}$ specification cannot be violated. For example, some DDR1 SDRAMs are not able to use more than 3 posted refreshes because the required refresh interval could then exceed the maximum constraint for $t_{ras}$ . 0000 Reserved 0001 1 refresh is issued at a time 0010 2 refreshes is issued at a time 0011 3 refreshes is issued at a time ... 1000 8 refreshes is issued at a time 1001–1111 Reserved
20–26	—	Reserved, should be cleared.
27	D_INIT	DRAM data initialization This bit is set by software, and it is cleared by hardware. If software sets this bit before the memory controller is enabled, the controller automatically initializes DRAM after it is enabled. This bit is automatically cleared by hardware once the initialization is completed. This data initialization bit should only be set when the controller is idle. 0 There is not data initialization in progress, and no data initialization is scheduled 1 The memory controller initializes memory once it is enabled. This bit remains asserted until the initialization is complete. The value in DDR_DATA_INIT register is used to initialize memory.
28–31	—	Reserved

### 9.4.1.9 DDR SDRAM Mode Configuration (DDR\_SDRAM\_MODE)

The DDR SDRAM mode configuration register, shown in Figure 9-10, sets the values loaded into the DDR’s mode registers.



**Figure 9-10. DDR SDRAM Mode Configuration Register (DDR\_SDRAM\_MODE)**



Table 9-14 describes the DDR\_SDRAM\_MODE fields.

**Table 9-14. DDR\_SDRAM\_MODE Field Descriptions**

Bits	Name	Description
0–15	ESDMODE	Extended SDRAM mode. Specifies the initial value loaded into the DDR SDRAM extended mode register. The range and meaning of legal values is specified by the DDR SDRAM manufacturer. When this value is driven onto the address bus (during the DDR SDRAM initialization sequence), MA[0] presents the lsb of ESDMODE, which, in the big-endian convention shown in Figure 9-10, corresponds to ESDMODE[15]. The msb of the SDRAM extended mode register value must be stored at ESDMODE[0].
16–31	SDMODE	SDRAM mode. Specifies the initial value loaded into the DDR SDRAM mode register. The range of legal values is specified by the DDR SDRAM manufacturer. When this value is driven onto the address bus (during DDR SDRAM initialization), MA[0] presents the lsb of SDMODE, which, in the big-endian convention shown in Figure 9-10, corresponds to SDMODE[15]. The msb of the SDRAM mode register value must be stored at SDMODE[0]. Because the memory controller forces SDMODE[7] to certain values depending on the state of the initialization sequence, the corresponding bits of this field are ignored by the memory controller. Note that SDMODE[7] is mapped to MA[8].

#### 9.4.1.10 DDR SDRAM Mode 2 Configuration (DDR\_SDRAM\_MODE\_2)

The DDR SDRAM mode 2 configuration register, shown in Figure 9-11, sets the values loaded into the DDR's extended mode 2 and 3 registers (for DDR2).



**Figure 9-11. DDR SDRAM Mode 2 Configuration Register (DDR\_SDRAM\_MODE\_2)**

Table 9-15 describes the DDR\_SDRAM\_MODE\_2 fields.

**Table 9-15. DDR\_SDRAM\_MODE\_2 Field Descriptions**


Bits	Name	Description
0–15	ESDMODE2	Extended SDRAM mode 2. Specifies the initial value loaded into the DDR SDRAM extended 2 mode register. The range and meaning of legal values is specified by the DDR SDRAM manufacturer. When this value is driven onto the address bus (during the DDR SDRAM initialization sequence), MA[0] presents the lsb bit of ESDMODE2, which, in the big-endian convention shown in Figure 9-11, corresponds to ESDMODE2[15]. The msb of the SDRAM extended mode 2 register value must be stored at ESDMODE2[0].
16–31	ESDMODE3	Extended SDRAM mode 3. Specifies the initial value loaded into the DDR SDRAM extended 3 mode register. The range of legal values of legal values is specified by the DDR SDRAM manufacturer. When this value is driven onto the address bus (during DDR SDRAM initialization), MA[0] presents the lsb of ESDMODE3, which, in the big-endian convention shown in Figure 9-11, corresponds to ESDMODE3[15]. The msb of the SDRAM extended mode 3 register value must be stored at ESDMODE3[0].

### 9.4.1.11 DDR SDRAM Mode Control Register (DDR\_SDRAM\_MD\_CNTL)

The DDR SDRAM mode control register, shown in Figure 9-12, allows the user to carry out the following tasks:

- Issue a mode register set command to a particular chip select
- Issue an immediate refresh to a particular chip select
- Issue an immediate precharge or precharge all command to a particular chip select
- Force the CKE signals to a specific value

Table 9-16 describes the fields of this register. Table 9-17 shows the user how to set the fields of this register to accomplish the above tasks.



**Figure 9-12. DDR SDRAM Mode Control Register (DDR\_SDRAM\_MD\_CNTL)**

Table 9-16 describes the DDR\_SDRAM\_MD\_CNTL fields.

**NOTE**

Note that MD\_EN, SET\_REF, and SET\_PRE are mutually exclusive; only one of these fields can be set at a time.

**Table 9-16. DDR\_SDRAM\_MD\_CNTL Field Descriptions**

Bits	Name	Description
0	MD_EN	Mode enable. Setting this bit specifies that valid data in MD_VALUE is ready to be written to DRAM as one of the following commands: <ul style="list-style-type: none"> <li>• MODE REGISTER SET</li> <li>• EXTENDED MODE REGISTER SET</li> <li>• EXTENDED MODE REGISTER SET 2</li> <li>• EXTENDED MODE REGISTER SET 3</li> </ul> The specific command to be executed is selected by setting MD_SEL. In addition, the chip select must be chosen by setting CS_SEL. MD_EN is set by software and cleared by hardware once the command has been issued. 0 Indicates that no mode register set command needs to be issued. 1 Indicates that valid data contained in the register is ready to be issued as a mode register set command.
1	—	Reserved
2-3	CS_SEL	Select chip select. Specifies the chip select that is driven active due to any command forced by software in DDR_SDRAM_MD_CNTL. 00 Chip select 0 is active 01 Chip select 1 is active 10 Reserved 11 Reserved
4	—	Reserved

Table 9-16. DDR\_SDRAM\_MD\_CNTL Field Descriptions (continued)

Bits	Name	Description
5–7	MD_SEL	Mode register select. MD_SEL specifies one of the following: <ul style="list-style-type: none"> <li>During a mode select command, selects the SDRAM mode register to be changed</li> <li>During a precharge command, selects the SDRAM logical bank to be precharged. A precharge all command ignores this field.</li> <li>During a refresh command, this field is ignored.</li> </ul> Note that MD_SEL contains the value that is presented onto the memory bank address pins (MBA <sub>n</sub> ) of the DDR controller. 000 MR 001 EMR 010 EMR2 011 EMR3
8	SET_REF	Set refresh. Forces an immediate refresh to be issued to the chip select specified by DDR_SDRAM_MD_CNTL[CS_SEL]. This bit is set by software and cleared by hardware once the command has been issued. 0 Indicates that no refresh command needs to be issued. 1 Indicates that a refresh command is ready to be issued.
9	SET_PRE	Set precharge. Forces a precharge or precharge all to be issued to the chip select specified by DDR_SDRAM_MD_CNTL[CS_SEL]. This bit is set by software and cleared by hardware once the command has been issued. 0 Indicates that no precharge all command needs to be issued. 1 Indicates that a precharge all command is ready to be issued.
10–11	CKE_CNTL	Clock enable control. Allows software to globally clear or set all CKE signals issued to DRAM. Once software has forced the value driven on CKE, that value continues to be forced until software clears the CKE_CNTL bits. At that time, the DDR controller continues to drive the CKE signals to the same value forced by software until another event causes the CKE signals to change (such as, self refresh entry/exit, power down entry/exit). 00 CKE signals are not forced by software. 01 CKE signals are forced to a low value by software. 10 CKE signals are forced to a high value by software. 11 Reserved
12–15	—	Reserved
16–31	MD_VALUE	Mode register value. This field, which specifies the value that is presented on the memory address pins of the DDR controller during a mode register set command, is significant only when this register is used to issue a mode register set command or a precharge or precharge all command. For a mode register set command, this field contains the data to be written to the selected mode register. For a precharge command, only bit five is significant: 0 Issue a precharge command; MD_SEL selects the logical bank to be precharged 1 Issue a precharge all command; all logical banks are precharged

Table 9-17 shows how DDR\_SDRAM\_MD\_CNTL fields should be set for each of the tasks described above.

Table 9-17. Settings of DDR\_SDRAM\_MD\_CNTL Fields


Field	Mode Register Set	Refresh	Precharge	Clock Enable Signals Control
MD_EN	1	0	0	—
SET_REF	0	1	0	—
SET_PRE	0	0	1	—

**Table 9-17. Settings of DDR\_SDRAM\_MD\_CNTL Fields (continued)**

Field	Mode Register Set	Refresh	Precharge	Clock Enable Signals Control
CS_SEL	Chooses chip select (CS)			—
MD_SEL	Select mode register. See <a href="#">Table 9-16</a> .	—	Selects logical bank	—
MD_VALUE	Value written to mode register	—	Only bit five is significant. See <a href="#">Table 9-16</a> .	—
CKE_CNTL	0	0	0	See <a href="#">Table 9-16</a> .

### 9.4.1.12 DDR SDRAM Interval Configuration (DDR\_SDRAM\_INTERVAL)

The DDR SDRAM interval configuration register, shown in [Figure 9-13](#), sets the number of DRAM clock cycles between bank refreshes issued to the DDR SDRAMs. In addition, the number of DRAM cycles that a page is maintained after it is accessed is provided here.



**Figure 9-13. DDR SDRAM Interval Configuration Register (DDR\_SDRAM\_INTERVAL)**


[Table 9-18](#) describes the DDR\_SDRAM\_INTERVAL fields.

**Table 9-18. DDR\_SDRAM\_INTERVAL Field Descriptions**

Bits	Name	Description
0–15	REFINT	Refresh interval. Represents the number of memory bus clock cycles between refresh cycles. Depending on DDR_SDRAM_CFG_2[NUM_PR], some number of rows are refreshed in each DDR SDRAM physical bank during each refresh cycle. The value for REFINT depends on the specific SDRAMs used and the interface clock frequency. Refreshes are not issued when the REFINT is set to all 0s.
16–17	—	Reserved
18–31	BSTOPRE	Precharge interval. Sets the duration (in memory bus clocks) that a page is retained after a DDR SDRAM access. If BSTOPRE is zero, the DDR memory controller uses auto-precharge read and write commands rather than operating in page mode. This is called global auto-precharge mode.

### 9.4.1.13 DDR SDRAM Data Initialization (DDR\_DATA\_INIT)

The DDR SDRAM data initialization register, shown in [Figure 9-14](#), provides the value that is used to initialize memory if DDR\_SDRAM\_CFG2[D\_INIT] is set.



**Figure 9-14. DDR SDRAM Data Initialization Configuration Register (DDR\_DATA\_INIT)**


[Table 9-19](#) describes the DDR\_DATA\_INIT fields.

**Table 9-19. DDR\_DATA\_INIT Field Descriptions**

Bits	Name	Description
0–31	INIT_VALUE	Initialization value. Represents the value that DRAM is initialized with if DDR_SDRAM_CFG2[D_INIT] is set.

### 9.4.1.14 DDR SDRAM Clock Control (DDR\_SDRAM\_CLK\_CNTL)

The DDR SDRAM clock control configuration register, shown in [Figure 9-15](#), provides a 1/4-cycle clock adjustment.



**Figure 9-15. DDR SDRAM Clock Control Configuration Register (DDR\_SDRAM\_CLK\_CNTL)**


[Table 9-20](#) describes the DDR\_SDRAM\_CLK\_CNTL fields.

**Table 9-20. DDR\_SDRAM\_CLK\_CNTL Field Descriptions**

Bits	Name	Description
0–4	—	Reserved
5–7	CLK_ADJUST	Clock adjust. 000 Clock is launched aligned with address/command 001 Clock is launched 1/4 applied cycle after address/command 010 Clock is launched 1/2 applied cycle after address/command 011 Clock is launched 3/4 applied cycle after address/command 100 Clock is launched 1 applied cycle after address/command 101–111 Reserved
8	—	Reserved, should be cleared.
9–31	—	Reserved

### 9.4.1.15 DDR Initialization Address (DDR\_INIT\_ADDR)

The DDR SDRAM initialization address register, shown in [Figure 9-16](#), provides the address that is used for the automatic  $\overline{\text{CAS}}$  to preamble calibration after POR.



**Figure 9-16. DDR Initialization Address Configuration Register (DDR\_INIT\_ADDR)**


[Table 9-21](#) describes the DDR\_INIT\_ADDR fields.

**Table 9-21. DDR\_INIT\_ADDR Field Descriptions**

Bits	Name	Description
0–31	INIT_ADDR	Initialization address. Represents the address that is used for the automatic CAS to preamble calibration at POR.

### 9.4.1.16 DDR IP Block Revision 1 (DDR\_IP\_REV1)

The DDR IP block revision 1 register, shown in [Figure 9-17](#), provides read-only fields with the IP block ID, along with major and minor revision information.



**Figure 9-17. DDR IP Block Revision 1 (DDR\_IP\_REV1)**

<sup>1</sup> For reset values, see [Table 9-22](#).

[Table 9-22](#) describes the DDR\_IP\_REV1 fields.

**Table 9-22. DDR\_IP\_REV1 Field Descriptions**

Bits	Name	Description
0–15	IP_ID	IP block ID. For the DDR controller, this value is 0x0002.
16–23	IP_MJ	Major revision. This is currently set to 0x02.
24–31	IP_MN	Minor revision. This is currently set to 0x01.

### 9.4.1.17 DDR IP Block Revision 2 (DDR\_IP\_REV2)

The DDR IP block revision 2 register, shown in Figure 9-18, provides read-only fields with the IP block integration and configuration options.




Figure 9-18. DDR IP Block Revision 2 (DDR\_IP\_REV2)

Table 9-23 describes the DDR\_IP\_REV2 fields.

Table 9-23. DDR\_IP\_REV2 Field Descriptions

Bits	Name	Description
0–7	—	Reserved
8–15	IP_INT	IP block integration options
16–23	—	Reserved
24–31	IP_CFG	IP block configuration options

## 9.5 Functional Description

The DDR SDRAM controller controls processor and I/O interactions with system memory. It provides support for JEDEC-compliant DDR2 and DDR SDRAMs. The memory system allows a wide range of memory devices to be mapped to any arbitrary chip select, and support is provided for registered DRAM modules and unbuffered DRAM modules. However, registered DRAM modules cannot be mixed with unbuffered DRAM modules.

Figure 9-19 is a high-level block diagram of the DDR memory controller. Requests are received from the internal mastering device and the address is decoded to generate the physical bank, logical bank, row, and column addresses. The transaction is compared with values in the row open table to determine if the address maps to an open page. If the transaction does not map to an open page, an active command is issued.

The memory interface supports as many as two physical banks of 32-bit wide memory. Bank sizes up to 512 Mbytes are supported, providing up to a maximum of 512 Mbytes of DDR main memory.

Programmable parameters allow for a variety of memory organizations and timings. The controller allows as many as 16 pages to be open simultaneously. The amount of time (in clock cycles) the pages remain open is programmable with DDR\_SDRAM\_INTERVAL[BSTOPRE].




Figure 9-19. DDR Memory Controller Block Diagram

Read and write accesses to memory are burst oriented; accesses start at a selected location and continue for a programmed number of higher locations (4 or 8) in a programmed sequence. Accesses to closed pages start with the registration of an ACTIVE command followed by a READ or WRITE. (Accessing open pages does not require an ACTIVE command.) The address bits registered coincident with the activate command specifies the logical bank and row to be accessed. The address coincident with the READ or WRITE command specify the logical bank and starting column for the burst access.

The data interface is source synchronous, meaning whatever sources the data also provides a clocking signal to synchronize data reception. These bidirectional data strobes (MDQS[0:3]) are inputs to the controller during reads and outputs during writes. The DDR SDRAM specification requires the data strobe signals to be centered within the data tenure during writes and to be offset by the controller to the center of the data tenure during reads. This delay is implemented in the controller for both reads and writes.

The address and command interface is also source synchronous, although 1/4 cycle adjustments are provided for adjusting the clock alignment.



Figure 9-20 shows an example DDR SDRAM configuration with four logical banks.

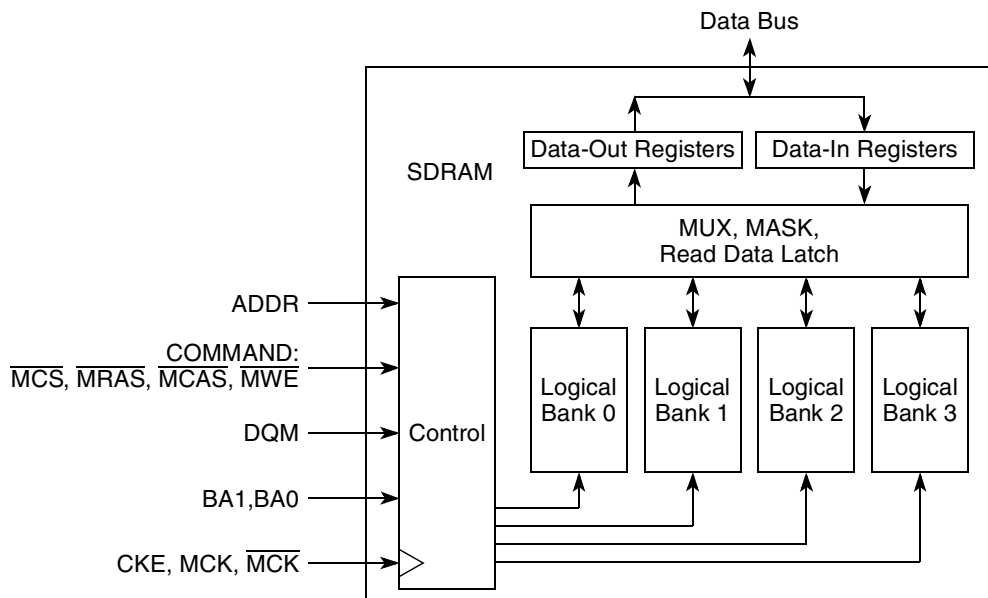


Figure 9-20. Typical Dual Data Rate SDRAM Internal Organization

Figure 9-21 shows some typical signal connections.

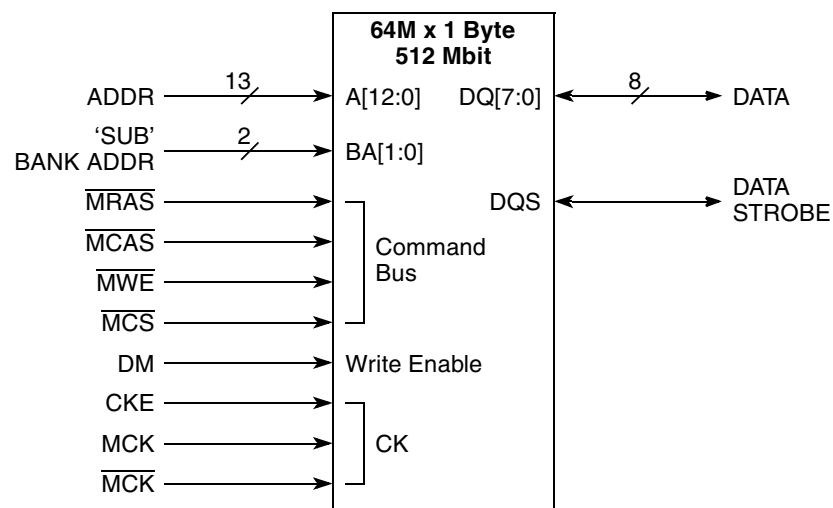
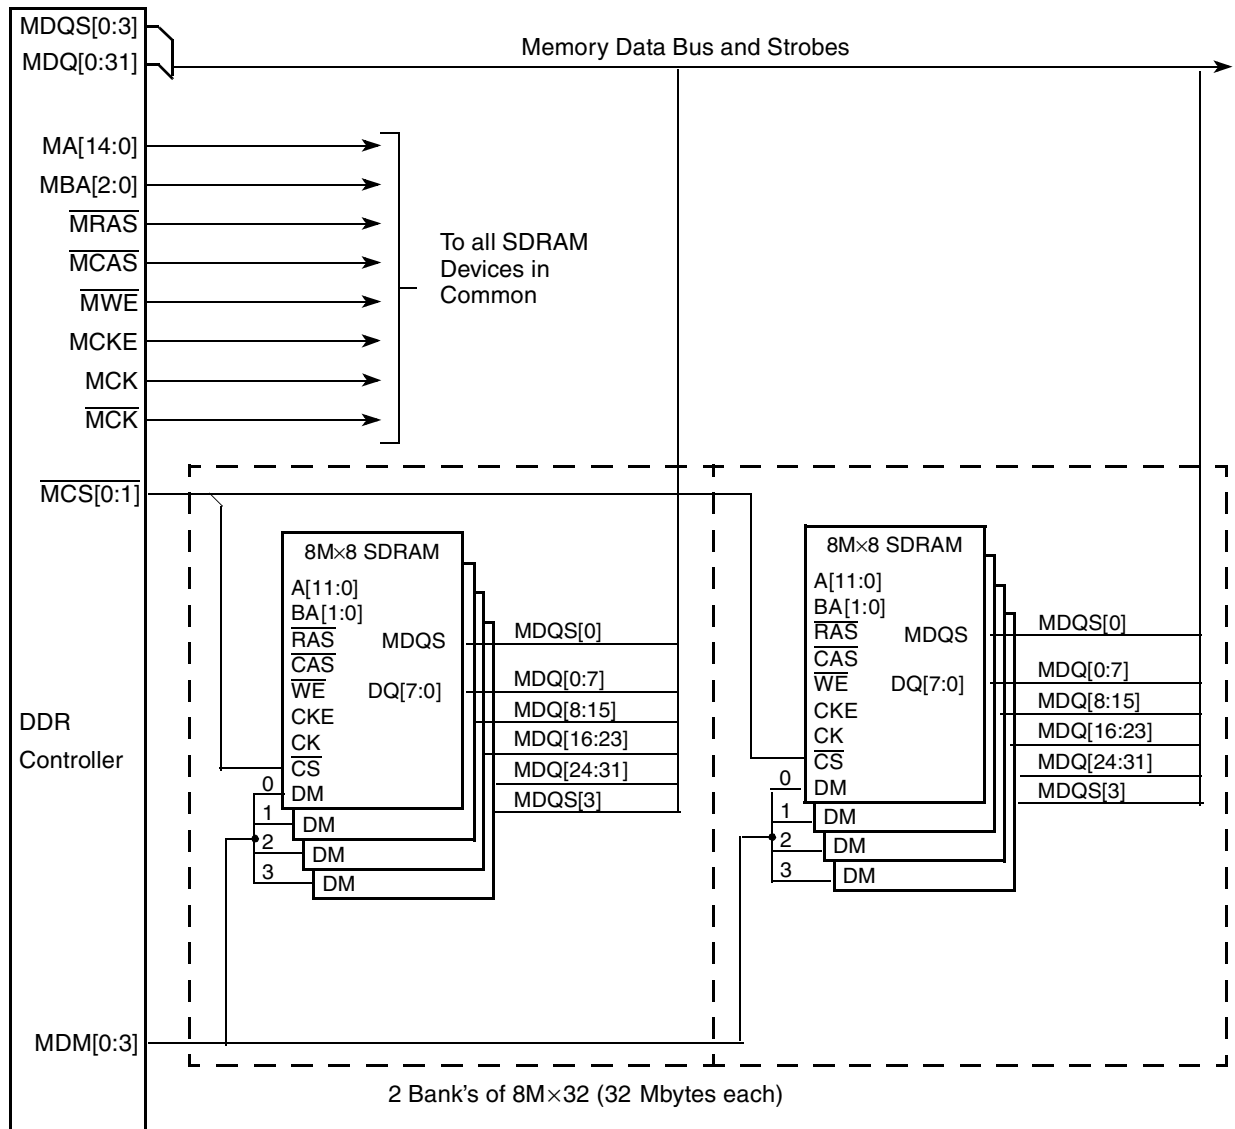


Figure 9-21. Typical DDR SDRAM Interface Signals

Figure 9-22 shows an example DDR SDRAM configuration with two physical banks each comprised of four  $8\text{M} \times 8$  DDR modules for a total of 128 Mbytes of system memory. Certain address and control lines may require buffering. Analysis of the device's AC timing specifications, desired memory operating frequency, capacitive loads, and board routing loads can assist the system designer in deciding signal buffering requirements. The DDR memory controller drives 15 address pins, but in this example the DDR SDRAM devices use only 12 bits.



1. All signals are connected in common (in parallel) except for  $\overline{MCS}[0:1]$ , MCK, MDM[0:3], and the data bus signals.
2. Each of the  $\overline{MCS}[0:1]$  signals correspond with a separate physical bank of memory.
3. Buffering may be needed if large memory arrays are used.

Figure 9-22. Example 64-Mbyte DDR SDRAM Configuration

### 9.5.1 DDR SDRAM Interface Operation

The DDR memory controller supports many different DDR SDRAM configurations. SDRAMs with different sizes can be used in the same system. Fifteen multiplexed address signals and three logical bank select signals support device densities from 64 Mbits to 2 Gbits. Two chip select ( $\overline{CS}$ ) signals support one bank of memory. The DDR SDRAM physical banks can be built from standard memory modules or directly-attached memory devices. The data path to individual physical banks is bits wide. The DDR memory controller supports physical bank sizes from 16 Mbytes to 512 Mbytes. The physical banks can

be constructed using x8, x16, or x32 memory devices. The memory technologies supported are 64 Mbits, 128 Mbits, 256 Mbits, 512 Mbits, 1 Gbit, 2 Gbits, and 4 Gbits. Four data qualifier (DQM) signals provide byte selection for memory accesses.

#### NOTE

An 8-bit DDR SDRAM device has a DQM signal and eight data signals (DQ[0:7]). A 16-bit DDR SDRAM device has two DQM signals associated with specific halves of the 16 data signals (DQ[0:7] and DQ[8:15]).

Table 9-24 shows the DDR memory controller's relationships between data byte lane0–3, MDM[0:3], MDQS[0:3], and MDQ[0:31] when DDR SDRAM memories are used with x8 or x16 devices.

**Table 9-24. Byte Lane to Data Relationship**

Data Byte Lane	Data Bus Mask	Data Bus Strobe	Data Bus
0 (MSB)	MDM[0]	MDQS[0]	MDQ[0:7]
1	MDM[1]	MDQS[1]	MDQ[8:15]
2	MDM[2]	MDQS[2]	MDQ[16:23]
3	MDM[3]	MDQS[3]	MDQ[24:31]

#### 9.5.1.1 Supported DDR SDRAM Organizations

Although the DDR memory controller multiplexes row and column address bits onto 15 memory address signals and 3 logical bank select signals, a physical bank may be implemented with memory devices requiring fewer than 30 address bits. The physical bank may be configured to provide from 12 to 15 row address bits, plus 2 or 3 logical bank-select bits and from 8–11 column address bits.

Table 9-25 and Table 9-26 describe DDR SDRAM device configurations supported by the DDR memory controller.

#### NOTE

DDR SDRAM is limited to 29 total address bits.

**Table 9-25. Supported DDR1 SDRAM Device Configurations**

SDRAM Device	Device Configuration	Row x Column x Sub-bank Bits	32-Bit Bank Size	Two Banks of Memory
64 Mbits	8 Mbits x 8	12 x 9 x 2	32 Mbytes	64 Mbytes
64 Mbits <sup>1</sup>	4 Mbits x 16	12 x 8 x 2	16 Mbytes	32 Mbytes
128 Mbits	16 Mbits x 8	12 x 10 x 2	64 Mbytes	128 Mbytes
128 Mbits	8 Mbits x 16	12 x 9 x 2	32 Mbytes	64 Mbytes
256 Mbits	32 Mbits x 8	13 x 10 x 2	128 Mbytes	256 Mbytes
256 Mbits	16 Mbits x 16	13 x 9 x 2	64 Mbytes	128 Mbytes
512 Mbits	64 Mbits x 8	13 x 11 x 2	256 Mbytes	512 Mbytes
512 Mbits	32 Mbits x 16	13 x 10 x 2	128 Mbytes	256 Mbytes
1 Gbits	128 Mbits x 8	14 x 11 x 2	512 Mbytes	1 Gbyte

**Table 9-25. Supported DDR1 SDRAM Device Configurations (continued)**

SDRAM Device	Device Configuration	Row x Column x Sub-bank Bits	32-Bit Bank Size	Two Banks of Memory
1 Gbits	64 Mbits x 16	14 x 10 x 2	256 Mbytes	512 Mbytes
2 Gbits	256 Mbits x 8	15 x 11 x 2	1 Gbyte	2 Gbytes
2 Gbits	128 Mbits x 16	15 x 10 x 2	512 Mbytes	1 Gbyte

<sup>1</sup> This configuration is not supported in 16-bit bus mode.

**Table 9-26. Supported DDR2 SDRAM Device Configurations**

SDRAM Device	Device Configuration	Row x Column x Sub-bank Bits	32-Bit Bank Size	Two Banks of Memory
256 Mbits	32 Mbits x 8	13 x 10 x 2	128 Mbytes	256 Mbytes
256 Mbits	16 Mbits x 16	13 x 9 x 2	64 Mbytes	128 Mbytes
512 Mbits	64 Mbits x 8	14 x 10 x 2	256 Mbytes	512 Mbytes
512 Mbits	32 Mbits x 16	13 x 10 x 2	128 Mbytes	256 Mbytes
1 Gbits	128 Mbits x 8	14 x 10 x 3	512 Mbytes	1 Gbyte
1 Gbits	64 Mbits x 16	13 x 10 x 3	256 Mbytes	512 Mbytes
2 Gbits	256 Mbits x 8	15 x 10 x 3	1 Gbytes	2 Gbytes
2 Gbits	128 Mbits x 16	14 x 10 x 3	512 Mbytes	1 Gbytes
4 Gbits	512 Mbits x 8	15 x 11 x 3	2 Gbyte	4 Gbytes
4 Gbits	256 Mbits x 16	15 x 10 x 3	1 Gbyte	2 Gbytes

If a transaction request is issued to the DDR memory controller and the address does not lie within any of the programmed address ranges for an enabled chip select, a memory select error is flagged.

Using a memory-polling algorithm at power-on reset or by querying the JEDEC serial presence detect capability of memory modules, system firmware uses the memory-boundary registers to configure the DDR memory controller to map the size of each bank in memory. The memory controller uses its bank map to assert the appropriate  $\overline{MCS}_n$  signal for memory accesses according to the provided bank starting and ending addresses. The memory banks are not required to be mapped to a contiguous address space.

## 9.5.2 DDR SDRAM Address Multiplexing

Table 9-27 and Table 9-28 show the address bit encodings for each DDR SDRAM configuration. The address presented at the memory controller signals MA[14:0] use MA[14] as the msb and MA[0] as the lsb. Also, MA[10] is used as the auto-precharge bit in DDR1/DDR2 modes for reads and writes, so the column address can never use MA[10].

Table 9-27. DDR1 Address Multiplexing for 32-Bit Data Bus with Interleaving Disabled

Row x Col	msb	Address from Core Master																												lsb			
		0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20	21	22	23	24	25	26	27		28	29	30-31
15 x 11 x 2	$\overline{\text{MRAS}}$			14	13	12	11	10	9	8	7	6	5	4	3	2	1	0															
	MBA																		1	0													
	$\overline{\text{MCAS}}$																					11	9	8	7	6	5	4	3	2	1	0	
15 x 10 x 2	$\overline{\text{MRAS}}$			14	13	12	11	10	9	8	7	6	5	4	3	2	1	0															
	MBA																			1	0												
	$\overline{\text{MCAS}}$																					9	8	7	6	5	4	3	2	1	0		
14 x 11 x 2	$\overline{\text{MRAS}}$			13	12	11	10	9	8	7	6	5	4	3	2	1	0																
	MBA																			1	0												
	$\overline{\text{MCAS}}$																					11	9	8	7	6	5	4	3	2	1	0	
14 x 10 x 2	$\overline{\text{MRAS}}$			13	12	11	10	9	8	7	6	5	4	3	2	1	0																
	MBA																				1	0											
	$\overline{\text{MCAS}}$																					9	8	7	6	5	4	3	2	1	0		
13 x 11 x 2	$\overline{\text{MRAS}}$			12	11	10	9	8	7	6	5	4	3	2	1	0																	
	MBA																				1	0											
	$\overline{\text{MCAS}}$																					11	9	8	7	6	5	4	3	2	1	0	
13 x 10 x 2	$\overline{\text{MRAS}}$			12	11	10	9	8	7	6	5	4	3	2	1	0																	
	MBA																				1	0											
	$\overline{\text{MCAS}}$																					9	8	7	6	5	4	3	2	1	0		
13 x 9 x 2	$\overline{\text{MRAS}}$					12	11	10	9	8	7	6	5	4	3	2	1	0															
	MBA																					1	0										
	$\overline{\text{MCAS}}$																						8	7	6	5	4	3	2	1	0		
12 x 10 x 2	$\overline{\text{MRAS}}$					11	10	9	8	7	6	5	4	3	2	1	0																
	MBA																					1	0										
	$\overline{\text{MCAS}}$																						9	8	7	6	5	4	3	2	1	0	
12 x 9 x 2	$\overline{\text{MRAS}}$					11	10	9	8	7	6	5	4	3	2	1	0																
	MBA																						1	0									
	$\overline{\text{MCAS}}$																							8	7	6	5	4	3	2	1	0	
12 x 8 x 2	$\overline{\text{MRAS}}$						11	10	9	8	7	6	5	4	3	2	1	0															
	MBA																							1	0								
	$\overline{\text{MCAS}}$																								7	6	5	4	3	2	1	0	

**Table 9-28. DDR2 Address Multiplexing**

Row x Col	msb	Address from Core Master																												lsb						
		0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20	21	22	23	24	25	26	27		28	29	30-31			
15 x 11 x 3	MRAS		14	13	12	11	10	9	8	7	6	5	4	3	2	1	0																			
	MBA																	2	1	0																
	MCAS																				11	9	8	7	6	5	4	3	2	1	0					
15 x 10 x 3	MRAS			14	13	12	11	10	9	8	7	6	5	4	3	2	1	0																		
	MBA																		2	1	0															
	MCAS																				9	8	7	6	5	4	3	2	1	0						
14 x 10 x 2	MRAS					13	12	11	10	9	8	7	6	5	4	3	2	1	0																	
	MBA																			1	0															
	MCAS																					9	8	7	6	5	4	3	2	1	0					
13 x 10 x 3	MRAS					12	11	10	9	8	7	6	5	4	3	2	1	0																		
	MBA																			2	1	0														
	MCAS																					9	8	7	6	5	4	3	2	1	0					
13 x 10 x 2	MRAS					12	11	10	9	8	7	6	5	4	3	2	1	0																		
	MBA																				1	0														
	MCAS																					9	8	7	6	5	4	3	2	1	0					
13 x 9 x 2	MRAS						12	11	10	9	8	7	6	5	4	3	2	1	0																	
	MBA																					1	0													
	MCAS																						8	7	6	5	4	3	2	1	0					

Chip select interleaving is supported for the memory controller, and is programmed in DDR\_SDRAM\_CFG[BA\_INTLV\_CTL]. Interleaving is supported between chip selects 0 and 1. When interleaving is enabled, the chip selects being interleaved must use the same size of memory. One extra bit in the address decode is used for the interleaving to determine which chip select to access.

Table 9-29 illustrates examples of address decode when interleaving between two chip selects.

**Table 9-29. Example of Address Multiplexing for 32-Bit Data Bus Interleaving between Two Banks with Partial Array Self Refresh Disabled**

Row x Col	msb	Address from Core Master																												lsb						
		1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20	21	22	23	24	25	26	27	28		29	30-31				
14 x 10 x 3	MRAS		13	12	11	10	9	8	7	6	5	4	3	2	1	0																				
	MBA																				CS SEL	2	1	0												
	MCAS																							9	8	7	6	5	4	3	2	1	0			

**Table 9-29. Example of Address Multiplexing for 32-Bit Data Bus Interleaving between Two Banks with Partial Array Self Refresh Disabled (continued)**

Row x Col	msb	Address from Core Master																												lsb								
		1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20	21	22	23	24	25	26	27	28		29	30–31						
14 x 10 x 2	MRAS			13	12	11	10	9	8	7	6	5	4	3	2	1	0	CS SEL																				
	MBA																			1	0																	
	MCAS																					9	8	7	6	5	4	3	2	1	0							
13 x 10 x 3	MRAS			12	11	10	9	8	7	6	5	4	3	2	1	0	CS SEL																					
	MBA																		2	1	0																	
	MCAS																				9	8	7	6	5	4	3	2	1	0								
13 x 10 x 2	MRAS			12	11	10	9	8	7	6	5	4	3	2	1	0	CS SEL																					
	MBA																	1	0																			
	MCAS																			9	8	7	6	5	4	3	2	1	0									

### 9.5.3 JEDEC Standard DDR SDRAM Interface Commands

The following section describes the commands and timings the controller uses when operating in DDR2 or DDR modes.

All read or write accesses to DDR SDRAM are performed by the DDR memory controller using JEDEC standard DDR SDRAM interface commands. The SDRAM device samples command and address inputs on rising edges of the memory clock; data is sampled using both the rising and falling edges of DQS. Data read from the DDR SDRAM is also sampled on both edges of DQS.

The following DDR SDRAM interface commands (summarized in [Table 9-30](#)) are provided by the DDR controller. All actions for these commands are described from the perspective of the SDRAM device.

- Row activate—Latches row address and initiates memory read of that row. Row data is latched in SDRAM sense amplifiers and must be restored by a precharge command before another row activate occurs.
- Precharge—Restores data from the sense amplifiers to the appropriate row. Also initializes the sense amplifiers in preparation for reading another row in the memory array (performing another activate command). Precharge must occur after read or write, if the row address changes on the next open page mode access.
- Read—Latches column address and transfers data from the selected sense amplifier to the output buffer as determined by the column address. During each succeeding clock edge, additional data is driven without additional read commands. The amount of data transferred is determined by the burst size which defaults to 4.
- Write—Latches column address and transfers data from the data pins to the selected sense amplifier as determined by the column address. During each succeeding clock edge, additional data is transferred to the sense amplifiers from the data pins without additional write commands. The

amount of data transferred is determined by the data masks and the burst size, which is set to four by the DDR memory controller.

- Refresh (similar to  $\overline{MCAS}$  before  $\overline{MRAS}$ )—Causes a row to be read in all logical banks (JEDEC SDRAM) as determined by the refresh row address counter. This refresh row address counter is internal to the SDRAM. After being read, the row is automatically rewritten in the memory array. All logical banks must be in a precharged state before executing a refresh. The memory controller also supports posted refreshes, where several refreshes may be executed at once, and the refresh interval may be extended.
- Mode register set (for configuration)—Allows setting of DDR SDRAM options. These options are:  $\overline{MCAS}$  latency, additive latency (for DDR2), write recovery (for DDR2), burst type, and burst length.  $\overline{MCAS}$  latency may be chosen as provided by the preferred SDRAM (some SDRAMs provide  $\overline{MCAS}$  latency {1,2,3}, some provide  $\overline{MCAS}$  latency {1,2,3,4,5}, and so on). Burst type is always sequential. Although some SDRAMs provide burst lengths of 1, 2, 4, 8, and page size, this memory controller supports a burst length of 4. A burst length of 8 is supported for DDR1 memory only. For DDR2 in 32-bit bus mode, all 32-byte burst accesses from the platform are split into two 16-byte (that is, 4-beat) accesses to the SDRAMs in the memory controller. The mode register set command is performed by the DDR memory controller during system initialization. Parameters such as mode register data,  $\overline{MCAS}$  latency, burst length, and burst type, are set by software in DDR\_SDRAM\_MODE[SDMODE] and transferred to the SDRAM array by the DDR memory controller after DDR\_SDRAM\_CFG[MEM\_EN] is set. If DDR\_SDRAM\_CFG[BI] is set to bypass the automatic initialization, then the MODE registers can be configured through software through use of the DDR\_SDRAM\_MD\_CNTL register.
- Self refresh (for long periods of standby)—Used when the device is in standby for very long periods of time. Automatically generates internal refresh cycles to keep the data in all memory banks refreshed. Before execution of this command, the DDR controller places all logical banks in a precharged state.

Table 9-30. DDR SDRAM Command Table

Operation	CKE Prev.	CKE Current	$\overline{MCS}$	$\overline{MRAS}$	$\overline{MCAS}$	$\overline{MWE}$	MBA	MA10	MA
Activate	H	H	L	L	H	H	Logical bank select	Row	Row
Precharge select logical bank	H	H	L	L	H	L	Logical bank select	L	X
Precharge all logical banks	H	H	L	L	H	L	X	H	X
Read	H	H	L	H	L	H	Logical bank select	L	Column
Read with auto-precharge	H	H	L	H	L	H	Logical bank select	H	Column
Write	H	H	L	H	L	L	Logical bank select	L	Column
Write with auto-precharge	H	H	L	H	L	L	Logical bank select	H	Column
Mode register set	H	H	L	L	L	L	Opcode	Opcode	Opcode and mode
Auto refresh	H	H	L	L	L	H	X	X	X
Self refresh	H	L	L	L	L	H	X	X	X



## 9.5.4 DDR SDRAM Interface Timing

The DDR memory controller supports four-beat bursts to SDRAM. For single-beat reads, the DDR memory controller performs a four- (or eight-) beat burst read, but ignores the last three (or seven) beats. Single-beat writes are performed by masking the last three (or seven) beats of the four- (or eight-) beat burst using the data mask MDM[0:3].

### NOTE

If a second read or write is pending, reads shorter than four beats are not terminated early even if some data is irrelevant.

To accommodate available memory technologies across a wide spectrum of operating frequencies, the DDR memory controller allows the setting of the intervals defined in [Table 9-31](#) with granularity of one memory clock cycle, except for CASLAT, which can be programmed with 1/2 clock granularity.

**Table 9-31. DDR SDRAM Interface Timing Intervals**

Timing Intervals	Definition
ACTTOACT	The number of clock cycles from a bank-activate command until another bank-activate command within a physical bank. This interval is listed in the AC specifications of the SDRAM as $t_{RRD}$ .
ACTTOPRE	The number of clock cycles from an activate command until a precharge command is allowed. This interval is listed in the AC specifications of the SDRAM as $t_{RAS}$ .
ACTTORW	The number of clock cycles from an activate command until a read or write command is allowed. This interval is listed in the AC specifications of the SDRAM as $t_{RCD}$ .
BSTOPRE	The number of clock cycles to maintain a page open after an access. The page open duration counter is reloaded with BSTOPRE each time the page is accessed (including page hits). When the counter expires, the open page is closed with an SDRAM precharge bank command as soon as possible.
CASLAT	Used in conjunction with additive latency to obtain the READ latency. The number of clock cycles between the registration of a READ command by the SDRAM and the availability of the first piece of output data. If a READ command is registered at clock edge $n$ , and the read latency is $m$ clocks, the data is available nominally coincident with clock edge $n + m$ .
PRETOACT	The number of clock cycles from a precharge command until an activate or a refresh command is allowed. This interval is listed in the AC specifications of the SDRAM as $t_{RP}$ .
REFINT	Refresh interval. Represents the number of memory bus clock cycles between refresh cycles. Depending on DDR_SDRAM_CFG_2[ <code>NUM_PR</code> ], some number of rows are refreshed in each SDRAM bank during each refresh cycle. The value of REFINT depends on the specific SDRAMs used and the frequency of the interface as $t_{RP}$ .
REFREC	The number of clock cycles from the refresh command until an activate command is allowed. This can be calculated by referring to the AC specification of the SDRAM device. The AC specification indicates a maximum refresh-to-activate interval in nanoseconds.
WR_DATA_DELAY	Provides different options for the timing between a write command and the write data strobe. This allows write data to be sent later than the nominal time to meet the SDRAM timing requirement between the registration of a write command and the reception of a data strobe associated with the write command. The specification dictates that the data strobe may not be received earlier than 75% of a cycle, or later than 125% of a cycle, from the registration of a write command. This parameter is not defined in the SDRAM specification. It is implementation-specific, defined for the DDR memory controller in TIMING_CFG_2.

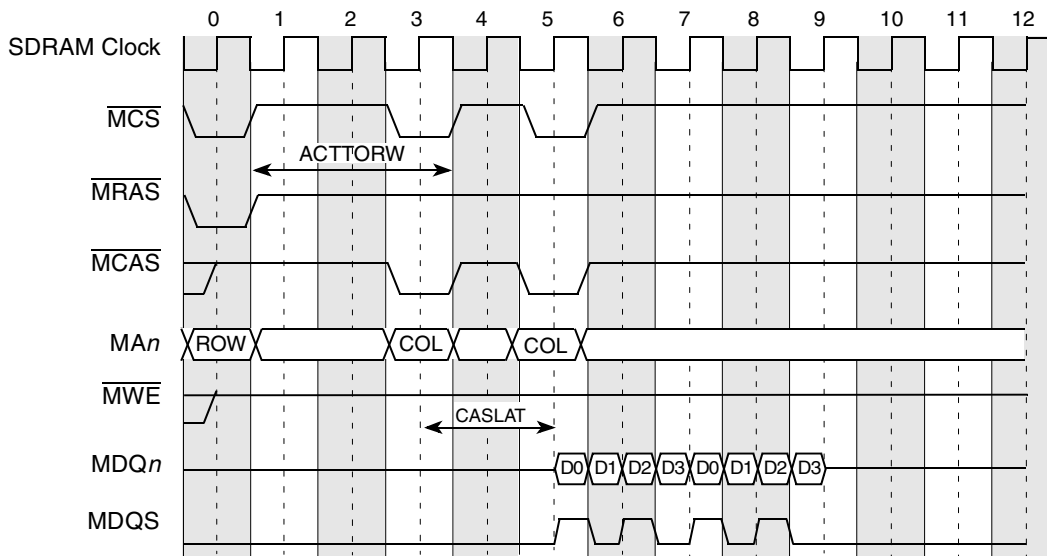
**Table 9-31. DDR SDRAM Interface Timing Intervals (continued)**

Timing Intervals	Definition
WRREC	The number of clock cycles from the last beat of a write until a precharge command is allowed. This interval, write recovery time, is listed in the AC specifications of the SDRAM as $t_{WR}$ .
WRTORD	Last write pair to read command. Controls the number of clock cycles from the last write data pair to the subsequent read command to the same bank as $t_{WTR}$ .

The value of the above parameters (in whole clock cycles) must be set by boot code at system start-up (in the TIMING\_CFG\_0, TIMING\_CFG\_1, TIMING\_CFG\_2, and TIMING\_CFG\_3 registers as described in Section 9.4.1.4, “DDR SDRAM Timing Configuration 0 (TIMING\_CFG\_0),” Section 9.4.1.5, “DDR SDRAM Timing Configuration 1 (TIMING\_CFG\_1),” Section 9.4.1.6, “DDR SDRAM Timing Configuration 2 (TIMING\_CFG\_2),” and Section 9.4.1.3, “DDR SDRAM Timing Configuration 3 (TIMING\_CFG\_3)”) and be kept in the DDR memory controller configuration register space.

The following figures show SDRAM timing for various types of accesses. System software is responsible (at reset) for optimally configuring SDRAM timing parameters. The programmable timing parameters apply to both read and write timing configuration. The configuration process must be completed and the DDR SDRAM initialized before any accesses to SDRAM are attempted.

Figure 9-23 through Figure 9-25 show DDR SDRAM timing for various types of accesses; see Figure 9-23 for a single-beat read operation, Figure 9-24 for a single-beat write operation, and Figure 9-25 for a burst-write operation. Note that all signal transitions occur on the rising edge of the memory bus clock and that single-beat read operations are identical to burst-reads. These figures assume the CLK\_ADJUST is set to 1/2 DRAM cycle, an additive latency of 0 DRAM cycles is used, and the write latency is 1 DRAM cycle (for DDR1).



**Figure 9-23. DDR SDRAM Burst Read Timing—ACTTORW = 3, MCAS Latency = 2**




Figure 9-24. DDR SDRAM Single-Beat (Double Word) Write Timing—ACTTOR




Figure 9-25. DDR SDRAM Single-Beat (Double Word) Write Timing—ACTTORW = 3

### 9.5.4.1 Clock Distribution

- If running with many devices, zero-delay PLL clock buffers, JEDEC-JESD82 standard, should be used. These buffers were designed for DDR applications.
- PCB traces for DDR clock signals should be short, all on the same layer, and of equal length and loading.
- DDR SDRAM manufacturers provide detailed information on PCB layout and termination issues.




Figure 9-26. DDR SDRAM Clock Distribution Example for x8 DDR SDRAMs

### 9.5.5 DDR SDRAM Mode-Set Command Timing

The DDR memory controller transfers the mode register set commands to the SDRAM array, and it uses the setting of TIMING\_CFG\_0[MRS\_CYC] for the Mode Register Set cycle time.

Figure 9-27 shows the timing of the mode-set command. The first transfer corresponds to the ESDMODE code; the second corresponds to SDMODE. The Mode Register Set cycle time is set to 2 DRAM cycles.




Figure 9-27. DDR SDRAM Mode-Set Command Timing

## 9.5.6 DDR SDRAM Registered DIMM Mode

To reduce loading, registered DRAM modules latch the DDR SDRAM control signals internally before using them to access the array. Setting `DDR_SDRAM_CFG[RD_EN]` compensates for this delay on the DRAM modules' control bus by delaying the data and data mask writes (on SDRAM buses) by an extra SDRAM clock cycle.

### NOTE

Application system board must assert the reset signal on DDR memory devices until software is able to program the DDR memory controller configuration registers, and must deassert the reset signal on DDR memory devices before `DDR_SDRAM_CFG[MEM_EN]` is set. This ensures that the DDR memory devices are held in reset until a stable clock is provided and, further, that a stable clock is provided before memory devices are released from reset.

Figure 9-28 shows the registered DDR SDRAM DIMM single-beat write timing.




Figure 9-28. Registered DDR SDRAM DIMM Burst Write Timing

## 9.5.7 DDR SDRAM Write Timing Adjustments

The DDR memory controller facilitates system design flexibility by providing a write timing adjustment parameter, write data delay, (`TIMING_CFG_2[WR_DATA_DELAY]`) for data and DQS. The DDR SDRAM specification requires DQS be received no sooner than 75% of an SDRAM clock period—and no later than 125% of a clock period—from the capturing clock edge of the command/address at the SDRAM. `TIMING_CFG_2[WR_DATA_DELAY]` specifies how much to delay the launching of DQS and data from the first clock edge occurring one SDRAM clock cycle after the command is launched. The delay increment step sizes are in 1/4 SDRAM clock periods starting with the default value of 0.

Figure 9-29 shows the use of the WR\_DATA\_DELAY parameter.




Figure 9-29. Write Timing Adjustments Example for Write Latency = 1

### 9.5.8 DDR SDRAM Refresh

The DDR memory controller supports auto-refresh and self-refresh. Auto refresh is used during normal operation and is controlled by the `DDR_SDRAM_INTERVAL[REFINT]` value; self-refresh is used only when the DDR memory controller is set to enter a sleep power management state. The REFINT value, which represents the number of memory bus clock cycles between refresh cycles, must allow for possible outstanding transactions to complete before a refresh request is sent to the memory after the REFINT value is reached. If a memory transaction is in progress when the refresh interval is reached, the refresh cycle waits for the transaction to complete. In the worst case, the refresh cycle must wait the number of bus clock cycles required by the longest programmed access. To ensure that the latency caused by a memory transaction does not violate the device refresh period, it is recommended that the programmed value of REFINT be less than that required by the SDRAM.

When a refresh cycle is required, the DDR memory controller does the following:

1. Completes all current memory requests.
2. Closes all open pages with a PRECHARGE-ALL command to each DDR SDRAM bank with an open page (as indicated by the row open table).
3. Issues one or more auto-refresh commands to each DDR SDRAM bank (as identified by its chip select) to refresh one row in each logical bank of the selected physical bank.

The auto-refresh commands are staggered across the two possible banks to reduce the system's instantaneous power requirements. Three sets of auto refresh commands are issued on consecutive cycles when the memory is populated with one DIMMs. The initial PRECHARGE-ALL commands are also staggered in three groups for convenience. It is important to note that when entering self-refresh mode, only one refresh command is issued simultaneously to all physical banks. For this entire refresh sequence, no cycle optimization occurs for the usual case where fewer than two banks are installed. After the refresh sequence completes, any pending memory request is initiated after an inactive period specified by `TIMING_CFG_1 [REFREC]` and `TIMING_CFG_3[EXT_REFREC]`. In addition, posted refreshes are supported to allow the refresh interval to be set to a larger value.

### 9.5.8.1 DDR SDRAM Refresh Timing

Refresh timing for the DDR SDRAM is controlled by the programmable timing parameter `TIMING_CFG_1 [REFREC]`, which specifies the number of memory bus clock cycles from the refresh command until a logical bank activate command is allowed. The DDR memory controller implements bank staggering for refreshes, as shown in [Figure 9-30](#) (`TIMING_CFG_1 [REFREC] = 10` in this example).




Figure 9-30. DDR SDRAM Bank Staggered Auto Refresh Timing

System software is responsible for optimal configuration of `TIMING_CFG_1 [REFREC]` and `TIMING_CFG_3[EXT_REFREC]` at reset. Configuration must be completed before DDR SDRAM accesses are attempted.

### 9.5.8.2 DDR SDRAM Refresh and Power-Saving Modes

In full-on mode, the DDR memory controller supplies the normal auto refresh to SDRAM. In sleep mode, the DDR memory controller can be configured to take advantage of self-refreshing SDRAMs or to provide no refresh support. Self-refresh support is enabled with the `SREN` memory control parameter.

Table 9-32 summarizes the refresh types available in each power-saving mode.

**Table 9-32. DDR SDRAM Power-Saving Modes Refresh Configuration**

Power Saving Mode	Refresh Type	SREN
Sleep	Self	1
	None	—


Note that in the absence of refresh support, system software must preserve DDR SDRAM data (such as by copying the data to disk) before entering the power-saving mode.

The dynamic power-saving mode uses the CKE DDR SDRAM pin to dynamically power down when there is no system memory activity. The CKE pin is negated when both of the following conditions are met:

- No memory refreshes are scheduled
- No memory accesses are scheduled

CKE is reasserted when a new access or refresh is scheduled or the dynamic power mode is disabled. This mode is controlled with DDR\_SDRAM\_CFG[DYN\_PWR\_MGMT].

Dynamic power management mode offers tight control of the memory system’s power consumption by trading power for performance through the use of CKE. Powering up the DDR SDRAM when a new memory reference is scheduled causes an access latency penalty, depending on whether active or precharge powerdown is used, along with the settings of TIMING\_CFG\_0[ACT\_PD\_EXIT] and TIMING\_CFG\_0[PRE\_PD\_EXIT]. A penalty of 1 cycle is shown in Figure 9-31.



**Figure 9-31. DDR SDRAM Power-Down Mode**



### 9.5.8.2.1 Self-Refresh in Sleep Mode

The entry and exit timing for self-refreshing SDRAMs is shown in Figure 9-32 and Figure 9-33.




Figure 9-32. DDR SDRAM Self-Refresh Entry Timing




Figure 9-33. DDR SDRAM Self-Refresh Exit Timing

## 9.5.9 DDR Data Beat Ordering

Transfers to and from memory are always performed in four- or eight-beat bursts. For transfer sizes other than four or eight beats, the data transfers are still operated as four- or eight-beat bursts. The DDR memory controller uses data masks to prevent all unintended full double words from writing to SDRAM. For example, if a write transaction is desired with a size of one double word (8 bytes), then the second, third, and fourth beats of data are not written to DRAM.

Table 9-33 lists the data beat sequencing to and from the DDR SDRAM and the data queues for each of the possible transfer sizes with each of the possible starting double-word offsets. All underlined double-word offsets are valid for the transaction.

**Table 9-33. Memory Controller–Data Beat Ordering**

Transfer Size	Starting Double-Word Offset	Double-Word Sequence <sup>1</sup> to/from DRAM and Queues
1 double word	0	<u>0</u> - 1 - 2 - 3
	1	<u>1</u> - 2 - 3 - 0
	2	<u>2</u> - 3 - 0 - 1
	3	<u>3</u> - 0 - 1 - 2
2 double words	0	<u>0</u> - <u>1</u> - 2 - 3
	1	<u>1</u> - <u>2</u> - 3 - 0
	2	<u>2</u> - <u>3</u> - 0 - 1
3 double words	0	<u>0</u> - <u>1</u> - <u>2</u> - 3
	1	<u>1</u> - <u>2</u> - <u>3</u> - 0

<sup>1</sup> All underlined **double**-word offsets are valid for the transaction.

## 9.5.10 Page Mode and Logical Bank Retention

The DDR memory controller supports an open/closed page mode with an allowable open page for each logical bank of DRAM used. In closed page mode for DDR SDRAMs, the DDR memory controller uses the SDRAM auto-precharge feature, which allows the controller to indicate that the page must be automatically closed by the DDR SDRAM after the READ or WRITE access. This is performed using MA[10] of the address during the COMMAND phase of the access to enable auto-precharge. Auto-precharge is non-persistent in that it is either enabled or disabled for each individual READ or WRITE command. It can, however, be enabled or disabled separately for each chip select.

When the DDR memory controller operates in open page mode, it retains the currently active SDRAM page by not issuing a precharge command. The page remains open until one of the following conditions occurs:

- Refresh interval is met.
- The user-programmable DDR\_SDRAM\_INTERVAL[BSTOPRE] value is exceeded.
- There is a logical bank row collision with another transaction that must be issued.

Page mode can dramatically reduce access latencies for page hits. Depending on the memory system design and timing parameters, using page mode can save two to three clock cycles for subsequent burst accesses that hit in an active page. Also, better performance can be obtained using more banks, especially in systems which use many different channels. Page mode is disabled by clearing DDR\_SDRAM\_INTERVAL[BSTOPRE] or setting CS<sub>n</sub>\_CONFIG[AP\_nEN].

## 9.6 Initialization/Application Information

System software must configure the DDR memory controller, using a memory polling algorithm at system start-up, to correctly map the size of each bank in memory. Then, the DDR memory controller uses its bank map to assert the appropriate  $\overline{MCS}_n$  signal for memory accesses according to the provided bank depths. System software must also configure the DDR memory controller at system start-up to appropriately multiplex the row and column address bits for each bank. Refer to row-address configuration in [Section 9.4.1.2, “Chip Select Configuration \(CS<sub>n</sub>\\_CONFIG\).”](#) Address multiplexing occurs according to these configuration bits.

At system reset, initialization software (boot code) must set up the programmable parameters in the memory interface configuration registers. See [Section 9.4.1, “Register Descriptions,”](#) for more detailed descriptions of the configuration registers. These parameters are shown in [Table 9-34.](#)

**Table 9-34. Memory Interface Configuration Register Initialization Parameters**

Name	Description	Parameter	Section/Page
CS <sub>n</sub> _BNDS	Chip select memory bounds	SA <sub>n</sub> EA <sub>n</sub>	<a href="#">9.4.1.1/9-9</a>
CS <sub>n</sub> _CONFIG	Chip select configuration	CS <sub>n</sub> _EN      BA_BITS_CS <sub>n</sub> AP <sub>n</sub> _EN      ROW_BITS_CS <sub>n</sub> ODT_RD_CFG    COL_BITS_CS <sub>n</sub> ODT_WR_CFG	<a href="#">9.4.1.2/9-10</a>
TIMING_CFG_3	Extended timing parameters for fields in TIMING_CFG_1	EXT_REFREC	<a href="#">9.4.1.3/9-11</a>
TIMING_CFG_0	Timing configuration	RWT          ACT_PD_EXIT WRT          PRE_PD_EXIT RRT          ODT_PD_EXIT WWT          MRS_CYC	<a href="#">9.4.1.4/9-12</a>
TIMING_CFG_1	Timing configuration	PRETOACT      REFREC ACTTOPRE      WRREC ACTTORW      ACTTOACT CASLAT        WRTORD	<a href="#">9.4.1.5/9-14</a>
TIMING_CFG_2	Timing configuration	ADD_LAT        WR_DATA_DELAY CPO            CKE_PLS WR_LAT        FOUR_ACT RD_TO_PRE	<a href="#">9.4.1.6/9-16</a>
DDR_SDRAM_CFG	Control configuration	SREN            NCAP RD_EN          2T_EN SDRAM_TYPE    BA_INTLV_CTL DYN_PWR        x32_EN 32_BE           HSE 8_BE            BI DBW	<a href="#">9.4.1.7/9-18</a>
DDR_SDRAM_CFG_2	Control configuration	DQS_CFG        NUM_PR ODT_CFG        D_INIT	<a href="#">9.4.1.8/9-21</a>
DDR_SDRAM_MODE	Mode configuration	ESDMODE SDMODE	<a href="#">9.4.1.9/9-22</a>

**Table 9-34. Memory Interface Configuration Register Initialization Parameters (continued)**

Name	Description	Parameter	Section/Page
DDR_SDRAM_MODE_2	Mode configuration	ESDMODE2 ESDMODE3	9.4.1.10/9-23
DDR_SDRAM_INTERVAL	Interval configuration	REFINT BSTOPRE	9.4.1.12/9-26
DDR_DATA_INIT	Data initialization configuration register	INIT_VALUE	9.4.1.13/9-27
DDR_SDRAM_CLK_CNTL	Clock adjust	CLK_ADJUST	9.4.1.14/9-27
DDR_INIT_ADDR	Initialization address	INIT_ADDR	9.4.1.15/9-28

### 9.6.1 Programming Differences Between Memory Types

Depending on the memory type used, certain fields must be programmed differently. Table 9-35 illustrates the differences in certain fields for different memory types. Note that this table does not list all fields that must be programmed.

**Table 9-35. Programming Differences Between Memory Types**

Parameter	Description	Differences		Section/Page
AP <sub>n</sub> _EN	Chip Select <i>n</i> Auto Precharge Enable	DDR1	Can be used to place chip select <i>n</i> in auto precharge mode	9.4.1.2/9-10
		DDR2	Can be used to place chip select <i>n</i> in auto precharge mode	
ODT_RD_CFG	Chip Select ODT Read Configuration	DDR1	Should always be set to 000	9.4.1.2/9-10
		DDR2	Can be enabled to assert ODT if desired. This could be set differently depending on system topology. However, systems with only 1 chip select typically not uses ODT when issuing reads to the memory.	
ODT_WR_CFG	Chip Select ODT Write Configuration	DDR1	Should always be set to 000	9.4.1.2/9-10
		DDR2	Can be enabled to assert ODT if desired. This could be set differently depending on system topology. However, ODT typically is set to assert for the chip select that is getting written to (value would be set to 001).	
ODT_PD_EXIT	ODT Powerdown Exit	DDR1	Should be set to 0001	9.4.1.4/9-12
		DDR2	Should be set according to the DDR2 specifications for the memory used. The JEDEC parameter this applies to is $t_{AXPD}$ .	
PRETOACT	Precharge to Activate Timing	DDR1	Should be set according to the specifications for the memory used ( $t_{RP}$ )	9.4.1.5/9-14
		DDR2	Should be set according to the specifications for the memory used ( $t_{RP}$ )	

Table 9-35. Programming Differences Between Memory Types (continued)

Parameter	Description	Differences		Section/Page
ACTTOPRE	Activate to Precharge Timing	DDR1	Should be set, along with the Extended Activate to Precharge Timing, according to the specifications for the memory used ( $t_{RAS}$ )	9.4.1.5/9-14
		DDR2	Should be set, along with the Extended Activate to Precharge Timing, according to the specifications for the memory used ( $t_{RAS}$ )	
ACTTORW	Activate to Read/Write Timing	DDR1	Should be set according to the specifications for the memory used ( $t_{RCD}$ )	9.4.1.5/9-14
		DDR2	Should be set according to the specifications for the memory used ( $t_{RCD}$ )	
CASLAT	CAS Latency	DDR1	Should be set, along with the Extended CAS Latency, to the desired CAS latency	9.4.1.5/9-14
		DDR2	Should be set, along with the Extended CAS Latency, to the desired CAS latency	
REFREC	Refresh Recovery	DDR1	Should be set, along with the Extended Refresh Recovery, to the specifications for the memory used ( $t_{RFC}$ )	9.4.1.5/9-14
		DDR2	Should be set, along with the Extended Refresh Recovery, to the specifications for the memory used ( $T_{RFC}$ )	
WRREC	Write Recovery	DDR1	Should be set according to the specifications for the memory used ( $t_{WR}$ )	9.4.1.5/9-14
		DDR2	Should be set according to the specifications for the memory used ( $t_{WR}$ )	
ACTTOACT	Activate <i>A</i> to Activate <i>B</i>	DDR1	Should be set according to the specifications for the memory used ( $t_{RRD}$ )	9.4.1.5/9-14
		DDR2	Should be set according to the specifications for the memory used ( $t_{RRD}$ )	
WRTORD	Write to Read Timing	DDR1	Should be set according to the specifications for the memory used ( $t_{WTR}$ )	9.4.1.5/9-14
		DDR2	Should be set according to the specifications for the memory used ( $t_{WTR}$ )	
ADD_LAT	Additive Latency	DDR1	Should be set to 000	9.4.1.6/9-16
		DDR2	Should be set to the desired additive latency. This must be set to a value less than TIMING_CFG_1[ACTTORW]	
WR_LAT	Write Latency	DDR1	Should be set to 001	9.4.1.6/9-16
		DDR2	Should be set to CAS latency – 1 cycle. For example, if the CAS latency is 5 cycles, then this field should be set to 100 (4 cycles).	

Table 9-35. Programming Differences Between Memory Types (continued)

Parameter	Description	Differences		Section/Page
RD_TO_PRE	Read to Precharge Timing	DDR1	Should be set to 010 if burst length is 4 and 100 if burst length is 8	9.4.1.6/9-16
		DDR2	Should be set according to the specifications for the memory used ( $t_{RTP}$ ). Time between read and precharge for non-zero value of additive latency (AL) is a minimum of $AL + t_{RTP}$ cycles.	
CKE_PLS	Minimum CKE Pulse Width	DDR1	Can be set to 001	9.4.1.6/9-16
		DDR2	Should be set according to the specifications for the memory used ( $t_{CKE}$ )	
FOUR_ACT	Four Activate Window	DDR1	Should be set to 00001	9.4.1.6/9-16
		DDR2	Should be set according to the specifications for the memory used ( $t_{FAW}$ ). Only applies to eight logical banks.	
RD_EN	Registered DIMM Enable	DDR1	If registered DRAM modules are used, then this field should be set to 1	9.4.1.7/9-18
		DDR2	If registered DRAM modules are used, then this field should be set to 1	
8_BE	8-beat burst enable	DDR1	If 8-beat bursts are desired, then this field should be set to 1	9.4.1.7/9-18
		DDR2	Should be set to 0	
2T_EN	2T Timing Enable	DDR1	In heavily loaded systems, this can be set to 1 to gain extra timing margin on the interface at the cost of address/command bandwidth.	9.4.1.7/9-18
		DDR2	In heavily loaded systems, this can be set to 1 to gain extra timing margin on the interface at the cost of address/command bandwidth.	
ODT_CFG	ODT Configuration	DDR1	Should be set to 00	9.4.1.8/9-21
		DDR2	Can be set for termination at the IOs according to system topology. Typically, if ODT is enabled, then the internal IOs should be set up for termination only during reads to DRAM.	
BSTOPR	Burst To Precharge Interval	DDR1	Can be set to any value, depending on the application. Auto precharge can be enabled by setting this field to all 0s.	9.4.1.12/9-26
		DDR2	Can be set to any value, depending on the application. Auto precharge can be enabled by setting this field to all 0s.	

## 9.6.2 DDR SDRAM Initialization Sequence

After configuration of all parameters is complete, system software must set `DDR_SDRAM_CFG[MEM_EN]` to enable the memory interface. Note that 200  $\mu$ s must elapse after DRAM clocks are stable (`DDR_SDRAM_CLK_CNTL[CLK_ADJUST]` is set and any chip select is enabled) before `MEM_EN` can be set, so a delay loop in the initialization code may be necessary if software is enabling the memory controller. If `DDR_SDRAM_CFG[BI]` is not set, the DDR memory controller conducts an automatic initialization sequence to the memory, which follows the memory specifications. If the bypass initialization mode is used, then software can initialize the memory through the `DDR_SDRAM_MD_CNTL` register.





# Chapter 10

## Enhanced Local Bus Controller

This chapter describes the enhanced local bus controller (eLBC) block. It describes the external signals and the memory-mapped registers as well as a functional description of the general-purpose chip-select machine (GPCM), NAND Flash control machine (FCM), and user-programmable machines (UPMs) of the eLBC. Finally, it includes an initialization and applications information section with many specific examples of its use.

### 10.1 Introduction

Figure 10-1 is a functional block diagram of the eLBC, which supports three interfaces: GPCM, FCM, and UPM controllers.




Figure 10-1. Enhanced Local Bus Controller Block Diagram

## 10.1.1 Overview

The main component of the eLBC is its memory controller, which provides a seamless interface to many types of memory devices and peripherals. The memory controller is responsible for controlling four memory banks shared by a GPCM, an FCM, and up to three UPMs. As such, it supports a minimal glue logic interface to SRAM, EPROM, NOR Flash EEPROM, NAND Flash EEPROM, burstable RAM, regular DRAM devices, extended data output DRAM devices, and other peripherals. The external address latch signal (LALE) allows multiplexing of addresses with data signals to reduce the device pin count. The eLBC also includes a number of data checking and protection features such as write protection and a bus monitor to ensure that each bus cycle is terminated within a user-specified period.

## 10.1.2 Features

The eLBC main features are as follows:

- Memory controller with four memory banks
  - 32-bit address decoding with mask
  - Variable memory block sizes (32 Kbytes to 4 Gbytes in FCM mode, 32 Kbytes to 64 Mbytes in GPCM and UPM modes)
  - Selection of control signal generation on a per-bank basis
  - Data buffer controls activated on a per-bank basis
  - Automatic segmentation of large transactions into memory accesses optimized for bus width and addressing capability
  - Write-protection capability
  - Atomic operation
- General-purpose chip-select machine (GPCM)
  - Compatible with SRAM, EPROM, NOR Flash EEPROM, and peripherals
  - Global (boot) chip-select available at system reset
  - Boot chip-select support for 8- and 16-bit devices
  - Minimum three-clock access to external devices
  - Two byte-write-enable signals ( $\overline{\text{LWE}}[0:1]$ )
  - Output enable signal ( $\overline{\text{LOE}}$ )
  - External access termination signal ( $\overline{\text{LGTA}}$ )
- NAND Flash control machine (FCM)
  - Compatible with small (512+16 bytes) and large (2048+64 bytes) page parallel NAND Flash EEPROM
  - Global (boot) chip-select available at system reset, with 4-Kbyte boot block buffer for execute-in-place boot loading
  - Read-only ECC registers to verify after write operation
  - Boot chip-select support for 8-bit devices
  - Dual 2-Kbyte/eight 512-byte buffers allow simultaneous data transfer during flash reads and programming

- Interrupt-driven block transfer for reads and writes
- Programmable command and data transfer sequences of up to eight steps supported
- Generic command and address registers support proprietary flash interfaces
- Block write locking to ensure system security and integrity
- Three user-programmable machines (UPMs)
  - Programmable-array-based machine controls external signal timing with a granularity of up to one quarter of an external bus clock period
  - User-specified control-signal patterns run when an internal master requests a single-beat or burst read or write access.
  - UPM refresh timer runs a user-specified control signal pattern to support refresh
  - User-specified control-signal patterns can be initiated by software
  - Each UPM can be defined to support DRAM devices with depths of 64, 128, 256, and 512 Kbytes, and 1, 2, 4, 8, 16, 32, 64, 128, and 256 Mbytes
  - Support for 8- and 16-bit devices
  - Page mode support for successive transfers within a burst
  - Internal address multiplexing supporting 64-, 128-, 256-, and 512-Kbyte, and 1-, 2-, 4-, 8-, 16-, 32-, 64-, 128-, and 256-Mbyte page banks
- Optional monitoring of transfers between local bus internal masters and local bus slaves (local bus error reporting on interrupt and status registers)

### 10.1.3 Modes of Operation

The eLBC provides one GPCM, one FCM, and three UPMs for the local bus, with no restriction on how many of the four banks (chip selects) can be programmed to operate with any given machine. The internal transaction address is limited to 32 bits, so all chip selects must fall within the 4-Gbyte window addressed by the internal transaction address. When a memory transaction is dispatched to the eLBC, the internal transaction address is compared with the address information of each bank (chip select). The corresponding machine assigned to that bank (GPCM, FCM, or UPM) then takes ownership of the external signals that control the access and maintains control until the transaction ends. Thus, with the eLBC in GPCM or FCM, or UPM mode, only one of the four chip selects is active at any time for the duration of the transaction except in the case of UPM refresh where all UPM machines that are enabled for refresh have concurrent chip select assertion.

#### 10.1.3.1 eLBC Bus Clock and Clock Ratios

The eLBC supports ratios of 2, 4, and 8 between the faster internal (system) clock and slower external bus clock (LCLK[0:1]). This ratio is software programmable through the clock ratio register (LCRR[CLKDIV]). This ratio affects the resolution of signal timing shifts in GPCM and FCM modes and the interpretation of UPM array words in UPM mode. The bus clock is driven identically onto pins, LCLK[0:1], to allow the clock load to be shared equally across a set of signal nets, thereby enhancing the edge rates of the bus clock.

### 10.1.3.2 Source ID Debug Mode

The eLBC provides the ID of a transaction source on external device pins. When those pins are selected, the 5-bit internal ID of the current transaction source appears on LSRCID[0:4] whenever valid address or data is available on the eLBC external pins. The reserved value of 0x1F, which indicates invalid address or data, appears on the source ID pins at all other times. The combination of a valid source ID (any value except 0x1F) and the value of external address latch enable (LALE) and data valid (LDVAL) facilitate capturing useful debug data as follows:

- If a valid source ID is detected on LSRCID[0:4] and LALE is asserted, a valid full 26-bit address may be latched from LAD[0:15] and combined with LA[16:25].
- If a valid source ID is detected on LSRCID[0:4] and LDVAL is asserted, valid data may be latched from LAD.

The LSRCID[0:4] and LDVAL signals are multiplexed with other functions sharing the same external pins. Refer to [Chapter 3, “Signal Descriptions,”](#) and [Chapter 4, “Reset, Clocking, and Initialization,”](#) to learn how to enable the LSRCID/LDVAL pins.

## 10.2 External Signal Descriptions

[Table 10-1](#) contains a list of external signals related to the eLBC and summarizes their function. The table also shows the reset state of all external signals during assertion of  $\overline{\text{HRESET}}$ . For more information on the use of some of these signals as reset configuration signals on the device, see “Power on Reset Flow.”

**Table 10-1. Signal Properties—Summary**

Name	Alternate Function(s)	Mode	Descriptions	No. of Signals	I/O	Reset State (Outputs)
LALE	—	—	External address latch enable	1	O	Reset_cfg
$\overline{\text{CS}}[0:3]$	—	—	Chip selects 0–3	4	O	Reset_cfg
$\overline{\text{LWE}}0/$ $\overline{\text{LWE}}/$ $\overline{\text{LBS}}0$	$\overline{\text{LWE}}0$	GPCM	Write enable 0	1	O	Reset_cfg
	$\overline{\text{LWE}}$	FCM	Write enable	1		
	$\overline{\text{LBS}}0$	UPM	Byte (lane) select 0	1		
$\overline{\text{LWE}}1/$ $\overline{\text{LBS}}1$	$\overline{\text{LWE}}$	GPCM	Write enable 1	1	O	Reset_cfg
	$\overline{\text{LBS}}1$	UPM	Byte (lane) select 1	1		
LGPL0/ LFCLE	LGPL0	UPM	General purpose line 0	1	O	Reset_cfg
	LFCLE	FCM	Flash command latch enable	1		
LGPL1/ LFALE	LGPL1	UPM	General purpose line 1	1	O	Reset_cfg
	LFALE	FCM	Flash address latch enable	1		
$\overline{\text{LOE}}/$ LGPL2/ $\overline{\text{LFRE}}$	$\overline{\text{LOE}}$	GPCM	Output enable	1	O	
	$\overline{\text{LFRE}}$	FCM	Flash read enable	1		
	LGPL2	UPM	General purpose line 2	1		

Table 10-1. Signal Properties—Summary (continued)

Name	Alternate Function(s)	Mode	Descriptions	No. of Signals	I/O	Reset State (Outputs)
LGPL3/ LFWP	LGPL3	UPM	General purpose line 3	1	O	Reset_cfg
	LFWP	FCM	Flash write protect	1		
LGTA/ LFRB/ LGPL4/ LUPWAIT	LGTA	GPCM	Transaction termination	1	I	High-Z
	LFRB	FCM	Flash ready/busy, open-drain shared pin	1	I	
	LGPL4	UPM	General purpose line 4	1	O	
	LUPWAIT	UPM	External device wait	1	I	
LGPL5	—	UPM	General purpose line 5	1	O	Reset_cfg
LBCTL	—	—	Data buffer control	1	O	
LA[0:25]	—	—	Non-multiplexed address bus	26	O	
LAD[0:15]	—	—	Multiplexed address/data bus	16	I/O	
LCLK[0:1]	—	—	Local bus clocks	2	O	Driven
LDVAL	—	eLBC debug	Local bus data valid	1	O	
LSRCID[0:4]	—	eLBC debug	Local bus source ID	5	O	

Table 10-2 shows the detailed external signal descriptions for the eLBC.

Table 10-2. Enhanced Local Bus Controller Detailed Signal Descriptions

Signal	I/O	Description	
LALE	O	External address latch enable. The local bus memory controller provides control for an external address latch, which allows address and data to be multiplexed on the device pins.	
		<b>State Meaning</b>	Asserted/Negated—LALE is asserted with the address at the beginning of each memory controller transaction. The number of cycles for which it is asserted is governed by the ORn[EAD] and LCRR[EADC] fields. Note that no other control signals are asserted during the assertion of LALE.
LCS[0:3]	O	Chip selects. Four chip selects are provided that are mutually exclusive.	
		<b>State Meaning</b>	Asserted/Negated—Used to enable specific memory devices or peripherals connected to the eLBC. LCS[0:3] are provided on a per-bank basis with LCS0 corresponding to the chip select for memory bank 0, which has the memory type and attributes defined by BR0 and OR0.

Table 10-2. Enhanced Local Bus Controller Detailed Signal Descriptions (continued)

Signal	I/O	Description	
$\overline{\text{LWE0}}$ / $\overline{\text{LWE1}}$ / $\overline{\text{LBS0}}$ / $\overline{\text{LBS1}}$	O	GPCM write enable 0/FCM write enable/UPM byte select 0. These signals select or validate each byte lane of the data bus. For an 8-bit port size, bit 0 is the only defined signal. The least-significant address bits of each access also determine which byte lanes are considered valid for a given data transfer.	
		<b>State Meaning</b>	Asserted/Negated—For GPCM operation, $\overline{\text{LWE}}[0:1]$ assert for each byte lane enabled for writing. $\overline{\text{LWE}}$ enables command, address, and data writes to NAND Flash EEPROMs controlled by FCM. $\overline{\text{LBS}}[0:1]$ are programmable byte-select signals in UPM mode. See <a href="#">Section 10.4.4.4, “RAM Array,”</a> for programming details about $\overline{\text{LBS}}[0:1]$ .
		<b>Timing</b>	Assertion/Negation—See <a href="#">Section 10.4.2, “General-Purpose Chip-Select Machine (GPCM),”</a> for details regarding the timing of $\overline{\text{LWE}}[0:1]$ .
LGPL0/ LFCLE	O	General purpose line 0/FCM command latch enable.	
		<b>State Meaning</b>	Asserted/Negated—In UPM mode, LGPL0 is one of six general purpose signals; it is driven with a value programmed into the UPM array. In FCM mode, LFCLE enables command cycles to NAND Flash EEPROMs.
LGPL1/ LFALE	O	General-purpose line 1/FCM address latch enable.	
		<b>State Meaning</b>	Asserted/Negated—In UPM mode, LGPL1 is one of six general purpose signals; it is driven with a value programmed into the UPM array. In FCM mode, LFALE enables address cycles to NAND Flash EEPROMs.
$\overline{\text{LOE}}$ /LGPL2/ LFRE	O	GPCM output enable/General-purpose line 2/FCM read enable.	
		<b>State Meaning</b>	Asserted/Negated—Controls the output buffer of memory when accessing memory/devices in GPCM mode. In UPM mode, LGPL2 is one of six general purpose signals; it is driven with a value programmed into the UPM array. $\overline{\text{LFRE}}$ enables data read cycles from NAND Flash EEPROMs controlled by FCM.
LGPL3/ $\overline{\text{LFWP}}$	O	General-purpose line 3/FCM write protect.	
		<b>State Meaning</b>	Asserted/Negated—In UPM mode, LGPL3 is one of six general purpose signals; it is driven with a value programmed into the UPM array. In FCM mode $\overline{\text{LFWP}}$ protects NAND Flash EEPROMs from accidental erasure and programming when $\overline{\text{LFWP}}$ is asserted low—see <a href="#">Section 10.3.1.17, “Flash Mode Register (FMR),”</a> for programming of FCM operations to control $\overline{\text{LFWP}}$ .
$\overline{\text{LGT}}$ /LGPL4/ $\overline{\text{LFRB}}$ / LUPWAIT	I/O	GPCM transfer acknowledge/General-purpose line 4/FCM Flash ready-busy/UPM wait.	
		<b>State Meaning</b>	Asserted/Negated—Input in GPCM or FCM modes used for transaction termination. It may also be configured as one of six general-purpose output signals when in UPM mode or as an input to force the UPM controller to wait for the memory/device. FCM uses $\overline{\text{LFRB}}$ to stall during long-latency read and programming operations, continuing once $\overline{\text{LFRB}}$ returns high.
LGPL5	O	General-purpose line 5	
		<b>State Meaning</b>	Asserted/Negated—One of six general purpose signals when in UPM mode, and drives a value programmed in the UPM array.

Table 10-2. Enhanced Local Bus Controller Detailed Signal Descriptions (continued)

Signal	I/O	Description
LBCTL	O	Data buffer control. The memory controller activates LBCTL for the local bus when a GPCM-, UPM-, or FCM-controlled bank is accessed. Buffer control is disabled by setting ORn[BCTLD].
		<b>State Meaning</b> Asserted/Negated—The LBCTL pin normally functions as a write/read control for a bus transceiver connected to the LAD lines. Note that an external data buffer must not drive the LAD lines in conflict with the eLBC when LBCTL is high, because LBCTL remains high after reset and during address phases.
LA[0:25]	O	Nonmultiplexed address bus. All bits driven on LA[0:25] are defined for 8-bit port sizes. For 16-bit port sizes LA[25] is a don't care.
		<b>State Meaning</b> Asserted/Negated—LA is the address bus used to transmit addresses to external RAM devices. Refer to <a href="#">Section 10.5, "Initialization/Application Information,"</a> for address signal multiplexing.
LAD[0:15]	I/O	Multiplexed address/data bus. For a port size of 16 bits, LAD[0:7] connect to the most-significant byte lane (at address offset 0), while LAD[8:15] connect to the least-significant byte lane (at address offset 1). For a port size of 8 bits, only LAD[0:7] are connected to the external RAM.
		<b>State Meaning</b> Asserted/Negated—LAD is the shared 16-bit address/data bus through which external RAM devices transfer data and receive addresses.
		<b>Timing</b> Assertion/Negation—During assertion of LALE, LAD are driven with the RAM address for the access to follow. External logic should propagate the address on LAD while LALE is asserted, and latch the address upon negation of LALE. After LALE is negated, LAD are either driven by write data or are made high-impedance by the eLBC in order to sample read data driven by an external device. Following the last data transfer of a write access, LAD are again taken into a high-impedance state.
LCLK[0:1]	O	Local bus clocks
		<b>State Meaning</b> Asserted/Negated—LCLK[0:1] drive an identical bus clock signal for distributed loads.
LDVAL	O	Local bus data valid (eLBC debug mode only)
		<b>State Meaning</b> Asserted/Negated—For a read, LDVAL asserts for one bus cycle in the cycle immediately preceding the sampling of read data on LAD. For a write, LDVAL asserts for one bus cycle during the final cycle for which the current write data on LAD is valid. During burst transfers, LDVAL asserts for each data beat.
		<b>Timing</b> Assertion/Negation—Valid only while the eLBC is in system debug mode. In debug mode, LDVAL asserts when the eLBC generates a data transfer acknowledge.
LSRCID[0:4]	O	Local bus source ID (eLBC debug mode only). In debug mode, all LSRCID[0:4] pins are driven high unless LSRCID[0:4] is driving a debug source ID for identifying the internal system device controlling the eLBC.
		<b>State Meaning</b> Asserted/Negated—Remain high until the last bus cycle of the assertion of LALE, in which case the source ID of the address is indicated, or until LDVAL is asserted, in which case the source ID relating to the data transfer is indicated. In case of address debug, LSRCID[0:4] is valid only when the address on LAD consists of all physical address bits—with optional padding—for reconstructing the system address presented to the eLBC.

## 10.3 Memory Map/Register Definition

Table 10-3 shows the memory mapped registers of the eLBC. Undefined 4-byte address spaces within offset 0x000–0xFFF are reserved.

**Table 10-3. Enhanced Local Bus Controller Registers**

Enhanced Local Bus Controller—Block Base Address 0x0_5000				
Offset	Register	Access	Reset	Section/Page
0x000	BR0—Base register 0	R/W	0x0000_0000	<a href="#">10.3.1.1/10-10</a>
0x008	BR1—Base register 1	R/W	0x0000_0000	<a href="#">10.3.1.1/10-10</a>
0x010	BR2—Base register 2	R/W	0x0000_0000	<a href="#">10.3.1.1/10-10</a>
0x018	BR3—Base register 3	R/W	0x0000_0000	<a href="#">10.3.1.1/10-10</a>
0x020–0x038	Reserved	R/W	0x0000_0000	<a href="#">10.3.1.1/10-10</a>
0x004	OR0—Options register 0	R/W	0x0000_0FF7	<a href="#">10.3.1.2/10-11</a>
0x00C	OR1—Options register 1	R/W	0x0000_0000	<a href="#">10.3.1.2/10-11</a>
0x014	OR2—Options register 2	R/W	0x0000_0000	<a href="#">10.3.1.2/10-11</a>
0x01C	OR3—Options register 3	R/W	0x0000_0000	<a href="#">10.3.1.2/10-11</a>
0x024–0x064	Reserved	—	—	—
0x068	MAR—UPM address register	R/W	0x0000_0000	<a href="#">10.3.1.3/10-19</a>
0x06C	Reserved	—	—	—
0x070	MAMR—UPMA mode register	R/W	0x0000_0000	<a href="#">10.3.1.4/10-20</a>
0x074	MBMR—UPMB mode register	R/W	0x0000_0000	<a href="#">10.3.1.4/10-20</a>
0x078	MCMR—UPMC mode register	R/W	0x0000_0000	<a href="#">10.3.1.4/10-20</a>
0x07C–0x080	Reserved	—	—	—
0x084	MRTPR—Memory refresh timer prescaler register	R/W	0x0000_0000	<a href="#">10.3.1.5/10-22</a>
0x088	MDR—UPM/FCM data register	R/W	0x0000_0000	<a href="#">10.3.1.6/10-22</a>
0x08C	Reserved	—	—	—
0x090	LSOR—Special operation initiation register	R/W	0x0000_0000	<a href="#">10.3.1.7/10-23</a>
0x094–0x09C	Reserved	—	—	—
0x0A0	LURT—UPM refresh timer	R/W	0x0000_0000	<a href="#">10.3.1.4/10-20</a>
0x0A4–0x0AC	Reserved	—	—	—
0x0B0	LTESR—Transfer error status register	w1c	0x0000_0000	<a href="#">10.3.1.9/10-25</a>
0x0B4	LTEDR—Transfer error disable register	R/W	0x0000_0000	<a href="#">10.3.1.10/10-27</a>



Table 10-3. Enhanced Local Bus Controller Registers (continued)

Enhanced Local Bus Controller—Block Base Address 0x0_5000				
Offset	Register	Access	Reset	Section/Page
0x0B8	LTEIR—Transfer error interrupt register	R/W	0x0000_0000	10.3.1.11/10-28
0x0BC	LTEATR—Transfer error attributes register	R/W	0x0000_0000	10.3.1.12/10-29
0x0C0	LTEAR—Transfer error address register	R/W	0x0000_0000	10.3.1.13/10-30
0x0C4	LTECCR—Transfer error ECC register	w1c	0x0000_0000	10.3.1.14/10-30
0x0C8– 0x0CC	Reserved	—	—	—
0x0D0	LBCR—Configuration register	R/W	0x0004_0000	10.3.1.15/10-31
0x0D4	LCRR—Clock ratio register	R/W	0x8000_0008	10.3.1.16/10-33
0x0D8– 0x0DC	Reserved	—	—	—
0x0E0	FMR—Flash mode register	R/W	0x0000_0n00	10.3.1.17/10-34
0x0E4	FIR—Flash instruction register	R/W	0x0000_0000	10.3.1.18/10-35
0x0E8	FCR—Flash command register	R/W	0x0000_0000	10.3.1.19/10-36
0x0EC	FBAR—Flash block address register	R/W	0x0000_0000	10.3.1.20/10-37
0x0F0	FPAR—Flash page address register	R/W	0x0000_0000	10.3.1.21/10-37
0x0F4	FBCR—Flash byte count register	R/W	0x0000_0000	10.3.1.22/10-39
0x0F8– 0x0FC	Reserved	—	—	—
0x100	FECC0—Flash ECC block 0 register	R	0x0000_0000	10.3.1.23/10-39
0x104	FECC1—Flash ECC block 1 register	R	0x0000_0000	10.3.1.23/10-39
0x108	FECC2—Flash ECC block 2 register	R	0x0000_0000	10.3.1.23/10-39
0x10C	FECC3—Flash ECC block 3 register	R	0x0000_0000	10.3.1.23/10-39

### 10.3.1 Register Descriptions

This section provides a detailed description of the eLBC configuration, status, and control registers with detailed bit and field descriptions.

Address offsets in the eLBC address range that are not defined in Table 10-3 should not be accessed for reading or writing. Similarly, only zero should be written to reserved bits of defined registers, as writing ones can have unpredictable results in some cases.


Bits designated as write-one-to-clear are cleared only by writing ones to them. Writing zeros to them has no effect.

### 10.3.1.1 Base Registers (BR0–BR3)

The base registers ( $BR_n$ ), shown in Figure 10-2, contain the base address and address types for each memory bank. The memory controller uses this information to compare the address bus value with the current address accessed. Each register (bank) includes a memory attribute and selects the machine for memory operation handling. Note that after system reset,  $BR0[V]$  is set,  $BR1[V]$ – $BR3[V]$  are cleared, and the value of  $BR0[PS]$  reflects the initial port size configured by the boot ROM location field of the reset configuration word.

Offset BR0: 0x0\_5000  
 BR1: 0x0\_5008  
 BR2: 0x0\_5010  
 BR3: 0x0\_5018

Access: Read/Write



<sup>1</sup> BR0 has its valid bit (V) set for  $RCWH[ROMLOC] = LBC$ . Thus bank 0 is valid with the port size (PS) configured from  $RCWH[ROMLOC]$  as loaded during reset. M = 0 for MSEL of GPCM, 1 for MSEL of FCM at boot. All other base registers have all bits cleared to zero during reset.

Figure 10-2. Base Registers ( $BR_n$ )

Table 10-4 describes  $BR_n$  fields.

Table 10-4.  $BR_n$  Field Descriptions

Bits	Name	Description
0–16	BA	Base address. The upper 17 bits of each base register are compared to the address on the address bus to determine if the bus master is accessing a memory bank controlled by the memory controller. Used with the address mask bits $OR_n[AM]$ .
17–18	—	Reserved
19–20	PS	Port size. Specifies the port size of this memory region. For BR0, PS is configured from the field in reset configuration word as loaded during reset. For all other banks the value is reset to 00 (port size not defined). 00 Reserved 01 8-bit 10 16-bit (not supported for FCM) 11 Reserved
21–22	DECC	Specifies the method for data error checking. 00 Data error checking disabled. No ECC generation for FCM. 01 ECC checking is enabled, but ECC generation is disabled, for FCM on full-page transfers. 10 ECC checking and generation are enabled for FCM on full-page transfers. 11 Reserved
23	WP	Write protect. 0 Read and write accesses are allowed. 1 Only read accesses are allowed. The memory controller does not assert $\overline{LCS}_n$ on write cycles to this memory bank. $LTESR[WP]$ is set (if WP is set) if a write to this memory bank is attempted, and a local bus error interrupt is generated (if enabled), terminating the cycle.

Table 10-4. BR<sub>n</sub> Field Descriptions (continued)

Bits	Name	Description
24–26	MSEL	Machine select. Specifies the machine to use for handling memory operations. 000 GPCM (possible reset value) 001 FCM (possible reset value) 010 Reserved 011 Reserved 100 UPMA 101 UPMB 110 UPMC 111 Reserved
27	—	Reserved
28–29	ATOM	Atomic operation. Writes (reads) to the address space handled by the memory controller bank reserve the selected memory bank for the exclusive use of the accessing device. The reservation is released when the device performs a read (write) operation to this memory controller bank. If a subsequent read (write) request to this memory controller bank is not detected within 256 bus clock cycles of the last write (read), the reservation is released and an atomic error is reported (if enabled). 00 The address space controlled by this bank is not used for atomic operations. 01 Read-after-write-atomic (RAWA). 10 Write-after-read-atomic (WARA). 11 Reserved
30	—	Reserved
31	V	Valid bit. Indicates that the contents of the BR <sub>n</sub> and OR <sub>n</sub> pair are valid. $\overline{LCSn}$ does not assert unless V is set (an access to a region that has no valid bit set may cause a bus time-out). After a system reset, only BR0[V] is set. 0 This bank is invalid. 1 This bank is valid.

### 10.3.1.2 Option Registers (OR0–OR3)

The OR<sub>n</sub> registers define the sizes of memory banks and access attributes. The OR<sub>n</sub> attribute bits support the following three modes of operation as defined by BR<sub>n</sub>[MSEL]:

- GPCM mode
- FCM mode
- UPM mode

The OR<sub>n</sub> registers are interpreted differently depending on which of the three machine types is selected for that bank. Because bank 0 can be used to boot, the reset value of OR0 may be different depending on power-on configuration options. Table 10-5 shows the reset values for OR0.

Table 10-5. Reset value of OR0 Register

Boot Source	OR0 Reset Value
FCM (small page NAND Flash)	0000_03AE
FCM (large page NAND Flash)	0000_07AE
GPCM	0000_0FF7
eLBC not used as a boot source	0000_0F07

### 10.3.1.2.1 Address Mask

The address mask field of the option registers (OR<sub>n</sub>[AM]) masks up to 17 corresponding BR<sub>n</sub>[BA] fields. The 15 LSBs of the 32-bit internal transaction address do not participate in bank address matching in selecting a bank for access. Masking address bits independently allows external devices of different size address ranges to be used. Address mask bits can be set or cleared in any order in the field, allowing a resource to reside in more than one area of the address map. Table 10-6 shows memory bank sizes from 32 Kbytes to 4 Gbytes. Memory block sizes vary from 32 Kbytes to 4 Gbytes in FCM mode, and 32 Kbytes to 64 Mbytes in GPCM and UPM modes.

**Table 10-6. Memory Bank Sizes in Relation to Address Mask**

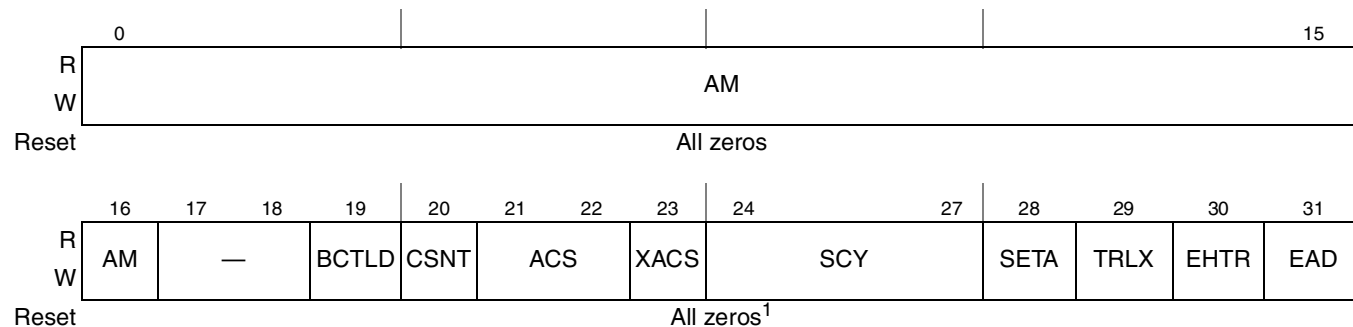
AM	Memory Bank Size
0000_0000_0000_0000_0	4 Gbytes
1000_0000_0000_0000_0	2 Gbytes
1100_0000_0000_0000_0	1 Gbyte
1110_0000_0000_0000_0	512 Mbytes
1111_0000_0000_0000_0	256 Mbytes
1111_1000_0000_0000_0	128 Mbytes
1111_1100_0000_0000_0	64 Mbytes
1111_1110_0000_0000_0	32 Mbytes
1111_1111_0000_0000_0	16 Mbytes
1111_1111_1000_0000_0	8 Mbytes
1111_1111_1100_0000_0	4 Mbytes
1111_1111_1110_0000_0	2 Mbytes
1111_1111_1111_0000_0	1 Mbyte
1111_1111_1111_1000_0	512 Kbytes
1111_1111_1111_1100_0	256 Kbytes
1111_1111_1111_1110_0	128 Kbytes
1111_1111_1111_1111_0	64 Kbytes
1111_1111_1111_1111_1	32 Kbytes

### 10.3.1.2.2 Option Registers (OR<sub>n</sub>)—GPCM Mode

Figure 10-3 shows the bit fields for OR<sub>n</sub> when the corresponding BR<sub>n</sub>[MSEL] selects the GPCM machine.

Offset OR0: 0x0\_5004  
 OR1: 0x0\_500c  
 OR2: 0x0\_5014  
 OR3: 0x0\_501c

Access: Read/Write



<sup>1</sup> Refer to Table 10-5 for the OR0 reset value. All other option registers have all bits cleared.

**Figure 10-3. Option Registers (OR<sub>n</sub>) in GPCM Mode**

Table 10-7 describes OR<sub>n</sub> fields for GPCM mode.

**Table 10-7. OR<sub>n</sub>—GPCM Field Descriptions**

Bits	Name	Description												
0–16	AM	GPCM address mask. Masks corresponding BR <sub>n</sub> bits. Masking address bits independently allows external devices of different size address ranges to be used. Address mask bits can be set or cleared in any order in the field, allowing a resource to reside in more than one area of the address map. 0 Corresponding address bits are masked and therefore don't care for address checking. 1 Corresponding address bits are used in the comparison between base and transaction addresses.												
17–18	—	Reserved												
19	BCTLD	Buffer control disable. Disables assertion of LBCTL during access to the current memory bank. 0 LBCTL is asserted upon access to the current memory bank. 1 LBCTL is not asserted upon access to the current memory bank.												
20	CSNT	Chip select negation time. Determines when $\overline{LCSn}$ and $\overline{LWE}$ are negated during an external memory write access handled by the GPCM, provided that ACS ≠ 00 (when ACS = 00, only $\overline{LWE}$ is affected by the setting of CSNT). This helps meet address/data hold times for slow memories and peripherals. 0 $\overline{LCSn}$ and $\overline{LWE}$ are negated normally. 1 $\overline{LCSn}$ and $\overline{LWE}$ are negated earlier depending on the value of LCRR[CLKDIV].												
		<table border="1"> <thead> <tr> <th>LCRR [CLKDIV]</th> <th>CSNT</th> <th>Meaning</th> </tr> </thead> <tbody> <tr> <td>x</td> <td>0</td> <td><math>\overline{LCSn}</math> and <math>\overline{LWE}</math> are negated normally.</td> </tr> <tr> <td>2</td> <td>1</td> <td><math>\overline{LCSn}</math> and <math>\overline{LWE}</math> are negated normally.</td> </tr> <tr> <td>4 or 8</td> <td>1</td> <td><math>\overline{LCSn}</math> and <math>\overline{LWE}</math> are negated one quarter bus clock cycle earlier.</td> </tr> </tbody> </table>	LCRR [CLKDIV]	CSNT	Meaning	x	0	$\overline{LCSn}$ and $\overline{LWE}$ are negated normally.	2	1	$\overline{LCSn}$ and $\overline{LWE}$ are negated normally.	4 or 8	1	$\overline{LCSn}$ and $\overline{LWE}$ are negated one quarter bus clock cycle earlier.
LCRR [CLKDIV]	CSNT	Meaning												
x	0	$\overline{LCSn}$ and $\overline{LWE}$ are negated normally.												
2	1	$\overline{LCSn}$ and $\overline{LWE}$ are negated normally.												
4 or 8	1	$\overline{LCSn}$ and $\overline{LWE}$ are negated one quarter bus clock cycle earlier.												

Table 10-7. OR<sub>n</sub>—GPCM Field Descriptions (continued)

Bits	Name	Description																		
21–22	ACS	<p>Address to chip-select setup. Determines the delay of the <math>\overline{\text{LCSn}}</math> assertion relative to the address change when the external memory access is handled by the GPCM. At system reset, OR0[ACS] = 11</p> <table border="1"> <thead> <tr> <th>LCRR [CLKDIV]</th> <th>Value</th> <th>Meaning</th> </tr> </thead> <tbody> <tr> <td rowspan="2">x</td> <td>00</td> <td><math>\overline{\text{LCSn}}</math> is output at the same time as the address lines. Note that this overrides the value of CSNT such that CSNT = 0.</td> </tr> <tr> <td>01</td> <td>Reserved.</td> </tr> <tr> <td rowspan="2">2</td> <td>10</td> <td><math>\overline{\text{LCSn}}</math> is output one half bus clock cycle after the address lines.</td> </tr> <tr> <td>11</td> <td><math>\overline{\text{LCSn}}</math> is output one half bus clock cycle after the address lines.</td> </tr> <tr> <td rowspan="2">4 or 8</td> <td>10</td> <td><math>\overline{\text{LCSn}}</math> is output one quarter bus clock cycle after the address lines.</td> </tr> <tr> <td>11</td> <td><math>\overline{\text{LCSn}}</math> is output one half bus clock cycle after the address lines.</td> </tr> </tbody> </table>	LCRR [CLKDIV]	Value	Meaning	x	00	$\overline{\text{LCSn}}$ is output at the same time as the address lines. Note that this overrides the value of CSNT such that CSNT = 0.	01	Reserved.	2	10	$\overline{\text{LCSn}}$ is output one half bus clock cycle after the address lines.	11	$\overline{\text{LCSn}}$ is output one half bus clock cycle after the address lines.	4 or 8	10	$\overline{\text{LCSn}}$ is output one quarter bus clock cycle after the address lines.	11	$\overline{\text{LCSn}}$ is output one half bus clock cycle after the address lines.
LCRR [CLKDIV]	Value	Meaning																		
x	00	$\overline{\text{LCSn}}$ is output at the same time as the address lines. Note that this overrides the value of CSNT such that CSNT = 0.																		
	01	Reserved.																		
2	10	$\overline{\text{LCSn}}$ is output one half bus clock cycle after the address lines.																		
	11	$\overline{\text{LCSn}}$ is output one half bus clock cycle after the address lines.																		
4 or 8	10	$\overline{\text{LCSn}}$ is output one quarter bus clock cycle after the address lines.																		
	11	$\overline{\text{LCSn}}$ is output one half bus clock cycle after the address lines.																		
23	XACS	<p>Extra address to chip-select setup. Setting this bit increases the delay of the <math>\overline{\text{LCSn}}</math> assertion relative to the address change when the external memory access is handled by the GPCM. After a system reset, OR0[XACS] = 1.</p> <p>0 Address to chip-select setup is determined by ORx[ACS] and LCRR[CLKDIV].  1 Address to chip-select setup is extended (see Table 10-32 and Table 10-33).</p>																		
24–27	SCY	<p>Cycle length in bus clocks. Determines the number of wait states inserted in the bus cycle, when the GPCM handles the external memory access. Thus it is the main parameter for determining cycle length. The total cycle length depends on other timing attribute settings. After a system reset, OR0[SCY] = 1111.</p> <p>0000 No wait states  0001 1 bus clock cycle wait state  ...  1111 15 bus clock cycle wait states</p>																		
28	SETA	<p>External address termination.</p> <p>0 Access is terminated internally by the memory controller unless the external device asserts <math>\overline{\text{LGTA}}</math> earlier to terminate the access.  1 Access is terminated externally by asserting the <math>\overline{\text{LGTA}}</math> external pin. (Only <math>\overline{\text{LGTA}}</math> can terminate the access).</p>																		
29	TRLX	<p>Timing relaxed. Modifies the settings of timing parameters for slow memories or peripherals.</p> <p>0 Normal timing is generated by the GPCM.  1 Relaxed timing on the following parameters:</p> <ul style="list-style-type: none"> <li>• Adds an additional cycle between the address and control signals (only if ACS is not equal to 00).</li> <li>• Doubles the number of wait states specified by SCY, providing up to 30 wait states.</li> <li>• Works in conjunction with EHTR to extend hold time on read accesses.</li> <li>• <math>\overline{\text{LCSn}}</math> (only if ACS is not equal to 00) and <math>\overline{\text{LWE}}</math> signals are negated one cycle earlier during writes.</li> </ul>																		

Table 10-7. OR<sub>n</sub>—GPCM Field Descriptions (continued)

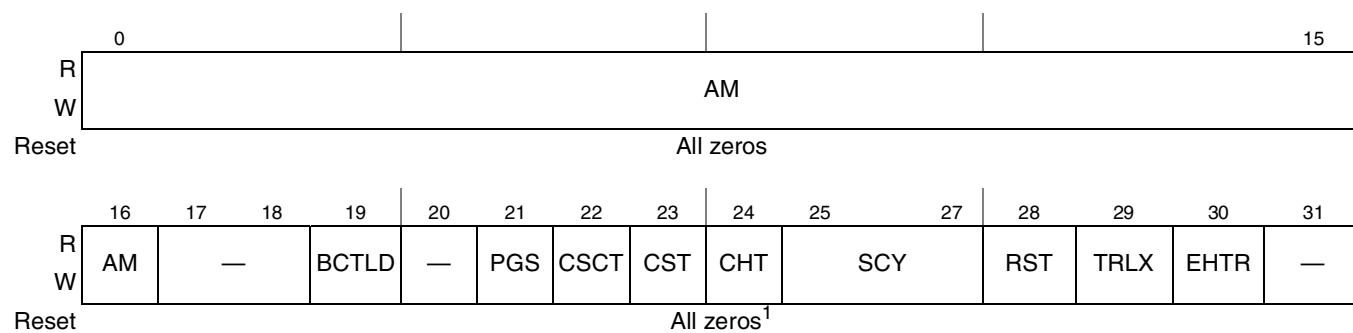
Bits	Name	Description															
30	EHTR	Extended hold time on read accesses. Indicates with TRLX how many cycles are inserted between a read access from the current bank and the next access. <table border="1" data-bbox="409 382 1403 651"> <thead> <tr> <th>TRLX</th> <th>EHTR</th> <th>Meaning</th> </tr> </thead> <tbody> <tr> <td>0</td> <td>0</td> <td>The memory controller generates normal timing. No additional cycles are inserted.</td> </tr> <tr> <td>0</td> <td>1</td> <td>1 idle clock cycle is inserted.</td> </tr> <tr> <td>1</td> <td>0</td> <td>4 idle clock cycles are inserted.</td> </tr> <tr> <td>1</td> <td>1</td> <td>8 idle clock cycles are inserted.</td> </tr> </tbody> </table>	TRLX	EHTR	Meaning	0	0	The memory controller generates normal timing. No additional cycles are inserted.	0	1	1 idle clock cycle is inserted.	1	0	4 idle clock cycles are inserted.	1	1	8 idle clock cycles are inserted.
TRLX	EHTR	Meaning															
0	0	The memory controller generates normal timing. No additional cycles are inserted.															
0	1	1 idle clock cycle is inserted.															
1	0	4 idle clock cycles are inserted.															
1	1	8 idle clock cycles are inserted.															
31	EAD	External address latch delay. Allow extra bus clock cycles when using external address latch (LALE). 0 No additional bus clock cycles (LALE asserted for one bus clock cycle only) 1 Extra bus clock cycles are added (LALE is asserted for the number of bus clock cycles specified by LCRR[EADC]).															

### 10.3.1.2.3 Option Registers (OR<sub>n</sub>)—FCM Mode

Figure 10-4 shows the bit fields for OR<sub>n</sub> when the corresponding BR<sub>n</sub>[MSEL] selects the FCM machine.

Offset OR0: 0x0\_5004  
 OR1: 0x0\_500c  
 OR2: 0x0\_5014  
 OR3: 0x0\_501c

Access: Read/Write



<sup>1</sup> Refer to Table 10-5 for the OR0 reset value. All other option registers have all bits cleared.

Figure 10-4. Option Registers (OR<sub>n</sub>) in FCM Mode

Table 10-8 describes  $OR_n$  fields for FCM mode.

**Table 10-8.  $OR_n$ —FCM Field Descriptions**

Bits	Name	Description															
0–16	AM	FCM address mask. Masks corresponding $BR_n$ bits. Masking address bits independently allows external devices of different size address ranges to be used. Address mask bits can be set or cleared in any order in the field, allowing a resource to reside in more than one area of the address map. 0 Corresponding address bits are masked. 1 Corresponding address bits are used in the comparison between base and transaction addresses.															
17–18	—	Reserved															
19	BCTLD	Buffer control disable. Disables assertion of LBCTL during access to the current memory bank. 0 LBCTL is asserted upon access to the current memory bank. 1 LBCTL is not asserted upon access to the current memory bank.															
20	—	Reserved															
21	PGS	NAND Flash EEPROM page size, buffer size, and block size. 0 Page size of 512 main area bytes plus 16 spare area bytes (small page devices); FCM RAM buffers are 1 Kbyte each; Flash block size of 16 Kbytes. 1 Page size of 2048 main area bytes plus 64 spare area bytes (large page devices); FCM RAM buffers are 4 Kbytes each; Flash block size of 128 Kbytes.															
22	CSCT	Chip select to command time. Determines how far in advance $\overline{LCS_n}$ is asserted prior to any bus activity during a NAND Flash access handled by the FCM. This helps meet chip-select setup times for slow memories. <table border="1" data-bbox="391 963 1442 1203"> <thead> <tr> <th>TRLX</th> <th>CSCT</th> <th>Meaning</th> </tr> </thead> <tbody> <tr> <td>0</td> <td>0</td> <td>The chip-select is asserted 1 clock cycle before any command.</td> </tr> <tr> <td>0</td> <td>1</td> <td>The chip-select is asserted 4 clock cycles before any command.</td> </tr> <tr> <td>1</td> <td>0</td> <td>The chip-select is asserted 2 clock cycles before any command.</td> </tr> <tr> <td>1</td> <td>1</td> <td>The chip-select is asserted 8 clock cycles before any command.</td> </tr> </tbody> </table>	TRLX	CSCT	Meaning	0	0	The chip-select is asserted 1 clock cycle before any command.	0	1	The chip-select is asserted 4 clock cycles before any command.	1	0	The chip-select is asserted 2 clock cycles before any command.	1	1	The chip-select is asserted 8 clock cycles before any command.
TRLX	CSCT	Meaning															
0	0	The chip-select is asserted 1 clock cycle before any command.															
0	1	The chip-select is asserted 4 clock cycles before any command.															
1	0	The chip-select is asserted 2 clock cycles before any command.															
1	1	The chip-select is asserted 8 clock cycles before any command.															
23	CST	Command setup time. Determines the delay of $\overline{LFW\overline{E}}$ assertion relative to the command, address, or data change when the external memory access is handled by the FCM. <table border="1" data-bbox="391 1310 1442 1606"> <thead> <tr> <th>TRLX</th> <th>CST</th> <th>Meaning</th> </tr> </thead> <tbody> <tr> <td>0</td> <td>0</td> <td>The write-enable is asserted coincident with any command.</td> </tr> <tr> <td>0</td> <td>1</td> <td>The write-enable is asserted 0.25 clock cycles after any command, address, or data.</td> </tr> <tr> <td>1</td> <td>0</td> <td>The write-enable is asserted 0.5 clock cycles after any command, address, or data.</td> </tr> <tr> <td>1</td> <td>1</td> <td>The write-enable is asserted 1 clock cycle after any command, address, or data.</td> </tr> </tbody> </table>	TRLX	CST	Meaning	0	0	The write-enable is asserted coincident with any command.	0	1	The write-enable is asserted 0.25 clock cycles after any command, address, or data.	1	0	The write-enable is asserted 0.5 clock cycles after any command, address, or data.	1	1	The write-enable is asserted 1 clock cycle after any command, address, or data.
TRLX	CST	Meaning															
0	0	The write-enable is asserted coincident with any command.															
0	1	The write-enable is asserted 0.25 clock cycles after any command, address, or data.															
1	0	The write-enable is asserted 0.5 clock cycles after any command, address, or data.															
1	1	The write-enable is asserted 1 clock cycle after any command, address, or data.															



Table 10-8. OR<sub>n</sub>—FCM Field Descriptions (continued)

Bits	Name	Description															
24	CHT	<p>Command hold time. Determines the <math>\overline{\text{LFW}}\overline{\text{E}}</math> negation prior to the command, address, or data change when the external memory access is handled by the FCM.</p> <table border="1"> <thead> <tr> <th>TRLX</th> <th>CHT</th> <th>Meaning</th> </tr> </thead> <tbody> <tr> <td>0</td> <td>0</td> <td>The write-enable is negated 0.5 clock cycles before any command, address, or data change.</td> </tr> <tr> <td>0</td> <td>1</td> <td>The write-enable is negated 1 clock cycle before any command, address, or data change.</td> </tr> <tr> <td>1</td> <td>0</td> <td>The write-enable is negated 1.5 clock cycles before any command, address, or data change.</td> </tr> <tr> <td>1</td> <td>1</td> <td>The write-enable is negated 2 clock cycles before any command, address, or data change.</td> </tr> </tbody> </table>	TRLX	CHT	Meaning	0	0	The write-enable is negated 0.5 clock cycles before any command, address, or data change.	0	1	The write-enable is negated 1 clock cycle before any command, address, or data change.	1	0	The write-enable is negated 1.5 clock cycles before any command, address, or data change.	1	1	The write-enable is negated 2 clock cycles before any command, address, or data change.
TRLX	CHT	Meaning															
0	0	The write-enable is negated 0.5 clock cycles before any command, address, or data change.															
0	1	The write-enable is negated 1 clock cycle before any command, address, or data change.															
1	0	The write-enable is negated 1.5 clock cycles before any command, address, or data change.															
1	1	The write-enable is negated 2 clock cycles before any command, address, or data change.															
25–27	SCY	<p>Cycle length in bus clocks. Determines:</p> <ul style="list-style-type: none"> <li>the number of wait states inserted in command, address, or data transfer bus cycles, when the FCM handles the external memory access. Thus it is the main parameter for determining cycle length. The total cycle length depends on other timing attribute settings.</li> <li>the delay between command/address writes and data write cycles, or the delay between write cycles and read cycles from NAND Flash EEPROM. A delay of <math>4 \times (2 + \text{SCY})</math> clock cycles (TRLX = 0) or <math>8 \times (2 + \text{SCY})</math> clock cycles (TRLX = 1) is inserted between the last write and the first data transfer to/from NAND Flash devices.</li> <li>the delay between a command write and the first sample point of the RDY/<math>\overline{\text{BSY}}</math> pin (connected to LFR<math>\overline{\text{B}}</math>). LFR<math>\overline{\text{B}}</math> is not sampled until <math>8 \times (2 + \text{SCY})</math> clock cycles (TRLX = 0) or <math>16 \times (2 + \text{SCY})</math> clock cycles (TRLX = 1) have elapsed following the command.</li> </ul> <p>000 No extra wait states  001 1 bus clock cycle wait state  ...  111 7 bus clock cycle wait states</p>															
28	RST	<p>Read setup time. Determines the delay of <math>\overline{\text{LFR}}\overline{\text{E}}</math> assertion relative to sampling of read data when the external memory access is handled by the FCM.</p> <table border="1"> <thead> <tr> <th>TRLX</th> <th>RST</th> <th>Meaning</th> </tr> </thead> <tbody> <tr> <td>0</td> <td>0</td> <td>The read-enable is asserted 0.75 clock cycles prior to any wait states.</td> </tr> <tr> <td>0</td> <td>1</td> <td>The read-enable is asserted 1 clock cycle prior to any wait states.</td> </tr> <tr> <td>1</td> <td>0</td> <td>The read-enable is asserted 0.5 clock cycles prior to any wait states.</td> </tr> <tr> <td>1</td> <td>1</td> <td>The read-enable is asserted 1 clock cycle prior to any wait states.</td> </tr> </tbody> </table>	TRLX	RST	Meaning	0	0	The read-enable is asserted 0.75 clock cycles prior to any wait states.	0	1	The read-enable is asserted 1 clock cycle prior to any wait states.	1	0	The read-enable is asserted 0.5 clock cycles prior to any wait states.	1	1	The read-enable is asserted 1 clock cycle prior to any wait states.
TRLX	RST	Meaning															
0	0	The read-enable is asserted 0.75 clock cycles prior to any wait states.															
0	1	The read-enable is asserted 1 clock cycle prior to any wait states.															
1	0	The read-enable is asserted 0.5 clock cycles prior to any wait states.															
1	1	The read-enable is asserted 1 clock cycle prior to any wait states.															
29	TRLX	<p>Timing relaxed. Modifies the settings of timing parameters for slow memories.</p> <p>0 Normal timing is generated by the FCM.</p> <p>1 Relaxed timing on the following parameters:</p> <ul style="list-style-type: none"> <li>Doubles the number of clock cycles between <math>\overline{\text{LCS}}\overline{\text{n}}</math> assertion and commands.</li> <li>Doubles the number of wait states specified by SCY, providing up to 14 wait states.</li> <li>Works in conjunction with CST and RST to extend command/address/data setup times.</li> <li>Adds one clock cycle to the command/address/data hold times.</li> <li>Works in conjunction with CBT to extend the wait time for read/busy status sampling by 16 clock cycles.</li> <li>Works in conjunction with EHTR to double hold time on read accesses.</li> </ul>															

Table 10-8. OR<sub>n</sub>—FCM Field Descriptions (continued)


Bits	Name	Description															
30	EHTR	Extended hold time on read accesses. Indicates with TRLX how many cycles are inserted between a read access from the current bank and the next access. <table border="1" style="margin-left: 20px;"> <thead> <tr> <th>TRLX</th> <th>EHTR</th> <th>Meaning</th> </tr> </thead> <tbody> <tr> <td>0</td> <td>0</td> <td>1 idle clock cycle is inserted.</td> </tr> <tr> <td>0</td> <td>1</td> <td>2 idle clock cycles are inserted.</td> </tr> <tr> <td>1</td> <td>0</td> <td>4 idle clock cycles are inserted.</td> </tr> <tr> <td>1</td> <td>1</td> <td>8 idle clock cycles are inserted.</td> </tr> </tbody> </table>	TRLX	EHTR	Meaning	0	0	1 idle clock cycle is inserted.	0	1	2 idle clock cycles are inserted.	1	0	4 idle clock cycles are inserted.	1	1	8 idle clock cycles are inserted.
TRLX	EHTR	Meaning															
0	0	1 idle clock cycle is inserted.															
0	1	2 idle clock cycles are inserted.															
1	0	4 idle clock cycles are inserted.															
1	1	8 idle clock cycles are inserted.															
31	—	Reserved															

### 10.3.1.2.4 Option Registers (OR<sub>n</sub>)—UPM Mode

Figure 10-5 shows the bit fields for OR<sub>n</sub> when the corresponding BR<sub>n</sub>[MSEL] selects a UPM machine.

Offset OR0: 0x0\_5004  
OR1: 0x0\_500c  
OR2: 0x0\_5014  
OR3: 0x0\_501c

Access: Read/Write



<sup>1</sup> Refer to Table 10-5 for the OR0 reset value. All other option registers have all bits cleared.

Figure 10-5. Option Registers (OR<sub>n</sub>) in UPM Mode

Table 10-9 describes BR<sub>n</sub> fields for UPM mode.

Table 10-9. OR<sub>n</sub>—UPM Field Descriptions

Bits	Name	Description
0–16	AM	UPM address mask. Masks corresponding BR <sub>n</sub> bits. Masking address bits independently allows external devices of different size address ranges to be used. Address mask bits can be set or cleared in any order in the field, allowing a resource to reside in more than one area of the address map. 0 Corresponding address bits are masked. 1 The corresponding address bits are used in the comparison with address pins.
17–18	—	Reserved

**Table 10-9. OR<sub>n</sub>—UPM Field Descriptions (continued)**

Bits	Name	Description															
19	BCTLD	Buffer control disable. Disables assertion of LBCTL during access to the current memory bank. 0 LBCTL is asserted upon access to the current memory bank. 1 LBCTL is not asserted upon access to the current memory bank.															
20–22	—	Reserved															
23	BI	Burst inhibit. Indicates if this memory bank supports burst accesses. 0 The bank supports burst accesses. 1 The bank does not support burst accesses. The selected UPM executes burst accesses as a series of single accesses.															
24–28	—	Reserved															
29	TRLX	Timing relaxed. Works in conjunction with EHTR to extend hold time on read accesses.															
30	EHTR	Extended hold time on read accesses. Indicates with TRLX how many cycles are inserted between a read access from the current bank and the next access. <table border="1" data-bbox="365 737 1446 976"> <thead> <tr> <th>TRLX</th> <th>EHTR</th> <th>Meaning</th> </tr> </thead> <tbody> <tr> <td>0</td> <td>0</td> <td>The memory controller generates normal timing. No additional cycles are inserted.</td> </tr> <tr> <td>0</td> <td>1</td> <td>1 idle clock cycle is inserted.</td> </tr> <tr> <td>1</td> <td>0</td> <td>4 idle clock cycles are inserted.</td> </tr> <tr> <td>1</td> <td>1</td> <td>8 idle clock cycles are inserted.</td> </tr> </tbody> </table>	TRLX	EHTR	Meaning	0	0	The memory controller generates normal timing. No additional cycles are inserted.	0	1	1 idle clock cycle is inserted.	1	0	4 idle clock cycles are inserted.	1	1	8 idle clock cycles are inserted.
TRLX	EHTR	Meaning															
0	0	The memory controller generates normal timing. No additional cycles are inserted.															
0	1	1 idle clock cycle is inserted.															
1	0	4 idle clock cycles are inserted.															
1	1	8 idle clock cycles are inserted.															
31	EAD	External address latch delay. Allow extra bus clock cycles when using external address latch (LALE). 0 No additional bus clock cycles (LALE asserted for one bus clock cycle only) 1 Extra bus clock cycles are added (LALE is asserted for the number of bus clock cycles specified by LCRR[EADC]).															

### 10.3.1.3 UPM Memory Address Register (MAR)

Figure 10-6 shows the fields of the UPM memory address register (MAR).


**Figure 10-6. UPM Memory Address Register (MAR)**

Table 10-10 describes the MAR fields.

**Table 10-10. MAR Field Descriptions**

Bits	Name	Description
0–31	A	Address that can be output to the address signals under control of the AMX bits in the UPM RAM word.

### 10.3.1.4 UPM Mode Registers (M<sub>x</sub>MR)

The UPM machine mode registers (MAMR, MBMR and MCMR), shown in Figure 10-7, contain the configuration for the three UPMs.

Offset MAMR: 0x0\_5070  
 MBMR: 0x0\_5074  
 MCMR: 0x0\_5078

Access: Read/Write



**Figure 10-7. UPM Mode Registers (M<sub>x</sub>MR)**

Table 10-11 describes UPM mode fields.

**Table 10-11. M<sub>x</sub>MR Field Descriptions**

Bits	Name	Description
0	—	Reserved
1	RFEN	Refresh enable. Indicates that the UPM needs refresh services. This bit must be set for UPMA (refresh executor) if refresh services are required on any UPM assigned chip selects. If MAMR[RFEN] = 0, no refresh services can be provided, even if UPMB and/or UPMC have their RFEN bit set. 0 Refresh services are not required 1 Refresh services are required
2–3	OP	Command opcode. Determines the command executed by the UPM <sub>n</sub> when a memory access hits a UPM assigned bank. 00 Normal operation 01 Write to UPM array. On the next memory access that hits a UPM assigned bank, write the contents of the MDR into the RAM location pointed to by MAD. After the access, MAD is automatically incremented. 10 Read from UPM array. On the next memory access that hits a UPM assigned bank, read the contents of the RAM location pointed to by MAD into the MDR. After the access, MAD is automatically incremented. 11 Run pattern. On the next memory access that hits a UPM assigned bank, run the pattern written in the RAM array. The pattern run starts at the location pointed to by MAD and continues until the LAST bit is set in the RAM word.
4	UWPL	LUPWAIT polarity active low. Sets the polarity of the LUPWAIT pin when in UPM mode. 0 LUPWAIT is active high. 1 LUPWAIT is active low.

Table 10-11. MxMR Field Descriptions (continued)

Bits	Name	Description														
5–7	AM	<p>Address multiplex size. Determines how the address of the current memory cycle can be output on the address pins. This field is needed when interfacing with devices requiring row and column addresses multiplexed on the same pins. See <a href="#">Section 10.4.4.4.7, “Address Multiplexing (AMX)”</a> for more information.</p> <p>000 Internal transaction address a[8:23] driven on LA[10:25]; LAD[0:15] driven low.            001 Internal transaction address a[7:22] driven on LA[10:25]; LAD[0:15] driven low.            010 Internal transaction address a[6:21] driven on LA[10:25]; LAD[0:15] driven low.            011 Internal transaction address a[5:20] driven on LA[10:25]; LAD[0:15] driven low.            100 Internal transaction address a[4:19] driven on LA[10:25]; LAD[0:15] driven low.            101 Internal transaction address a[3:18] driven on LA[10:25]; LAD[0:15] driven low.            110 Reserved            111 Reserved</p>														
8–9	DS	<p>Disable timer period. Guarantees a minimum time between accesses to the same memory bank controlled by UPM<math>n</math>. The disable timer is turned on by the TODT bit in the RAM array word, and when expired, the UPM<math>n</math> allows the machine access to handle a memory pattern to the same bank. Accesses to a different bank by the same UPM<math>n</math> is also allowed. To avoid conflicts between successive accesses to different banks, the minimum pattern in the RAM array for a request serviced, should not be shorter than the period established by DS.</p> <p>00 1-bus clock cycle disable period            01 2-bus clock cycle disable period            10 3-bus clock cycle disable period            11 4-bus clock cycle disable period</p>														
10–12	G0CL	<p>General line 0 control. Determines which logical address line can be output to the LGPL0 pin when the UPM<math>n</math> is selected to control the memory access.</p> <p>000 A12            001 A11            010 A10            011 A9            100 A8            101 A7            110 A6            111 A5</p>														
13	GPL4	<p>LGPL4 output line disable. Determines how the LGPL4/LUPWAIT pin is controlled by the corresponding bits in the UPM<math>n</math> array. See <a href="#">Table 10-40 on page 10-79</a>.</p> <table border="1" data-bbox="386 1325 1182 1518"> <thead> <tr> <th rowspan="2">Value</th> <th rowspan="2">LGPL4/LUPWAIT Pin Function</th> <th colspan="2">Interpretation of UPM Word Bits</th> </tr> <tr> <th>G4T1/DLT3</th> <th>G4T3/WAEN</th> </tr> </thead> <tbody> <tr> <td>0</td> <td>LGPL4 (output)</td> <td>G4T1</td> <td>G4T3</td> </tr> <tr> <td>1</td> <td>LUPWAIT (input)</td> <td>DLT3</td> <td>WAEN</td> </tr> </tbody> </table>	Value	LGPL4/LUPWAIT Pin Function	Interpretation of UPM Word Bits		G4T1/DLT3	G4T3/WAEN	0	LGPL4 (output)	G4T1	G4T3	1	LUPWAIT (input)	DLT3	WAEN
Value	LGPL4/LUPWAIT Pin Function	Interpretation of UPM Word Bits														
		G4T1/DLT3	G4T3/WAEN													
0	LGPL4 (output)	G4T1	G4T3													
1	LUPWAIT (input)	DLT3	WAEN													
14–17	RLF	<p>Read loop field. Determines the number of times a loop defined in the UPM<math>n</math> will be executed for a burst- or single-beat read pattern or when MxMR[OP] = 11 (RUN command)</p> <p>0000 16            0001 1            0010 2            0011 3            ...            1110 14            1111 15</p>														

Table 10-11. MxMR Field Descriptions (continued)

Bits	Name	Description
18–21	WLF	Write loop field. Determines the number of times a loop defined in the UPM $n$ will be executed for a burst- or single-beat write pattern. 0000 16 0001 1 0010 2 0011 3 ... 1110 14 1111 15
22–25	TLF	Refresh loop field. Determines the number of times a loop defined in the UPM $n$ will be executed for a refresh service pattern. 0000 16 0001 1 0010 2 0011 3 ... 1110 14 1111 15
26–31	MAD	Machine address. RAM address pointer for the command executed. This field is incremented by 1, each time the UPM is accessed and the OP field is set to WRITE or READ. Address range is 64 words per UPM $n$ .

### 10.3.1.5 Memory Refresh Timer Prescaler Register (MRTPR)

The refresh timer prescaler register (MRTPR), shown in [Figure 10-8](#), is used to divide the system clock to provide the UPM refresh timers clock.




Figure 10-8. Memory Refresh Timer Prescaler Register (MRTPR)

[Table 10-12](#) describes MRTPR fields.


Table 10-12. MRTPR Field Descriptions

Bits	Name	Description
0–7	PTP	Refresh timers prescaler. Determines the period of the refresh timers input clock. The system clock is divided by PTP except when the value is 00000_0000, which represents the maximum divider of 256.
8–31	—	Reserved


### 10.3.1.6 UPM/FCM Data Register (MDR)

The memory data register (MDR), shown in [Figure 10-9](#) and [Figure 10-10](#), contains data written to or read from the RAM array for UPM read or write commands. MDR also contains data written to or read from an external NAND Flash EEPROM for FCM write address, write data, and read status commands. MDR

must be set up before issuing a write command to the UPM, or before issuing a FCM operation sequence that uses MDR to source address or data bytes.



**Figure 10-9. UPM Data Register in UPM Mode (MDR)**



**Figure 10-10. FCM Data Register in FCM Mode (MDR)**

Table 10-13 describes MDR[D].


**Table 10-13. MDR Field Description**

Bits	Name	Description
0–31	D	In UPM mode, D is the data to be read or written into the RAM array when a write or read command is supplied to the UPM (MxMR[OP] = 01 or MxMR[OP] = 10).
0–7	AS3	In FCM mode, AS3 is the fourth byte of address sent by a custom address write operation, or the fourth byte of data read from a read status operation.
8–15	AS2	In FCM mode, AS2 is the third byte of address sent by a custom address write operation, or the third byte of data read from a read status operation.
16–23	AS1	In FCM mode, AS1 is the second byte of address sent by a custom address write operation, or the second byte of data read from a read status operation.
24–31	AS0	In FCM mode, AS0 is the first byte of address sent by a custom address write operation, or the first byte of data read from a read status operation.

### 10.3.1.7 Special Operation Initiation Register (LSOR)

The special operation initiation register (LSOR), shown in Figure 10-11, is used by software to trigger a special operation on the indicated bank. Writing to LSOR activates a special operation on bank LSOR[BANK] provided that the bank is valid and controlled by a memory controller whose mode OP field is set to a value other than ‘normal operation.’ If eLBC is currently busy with a memory transaction, writing LSOR completes immediately, but the special operation request is queued until eLBC can service it. To avoid race conditions between software and a busy eLBC, registers that affect currently running special operation and LSOR must not be re-written before a pending special operation has been completed. The UPM and FCM have different indications of when such special operations are completed. The behavior of eLBC is unpredictable if special operation modes are altered between LSOR being written and the relevant memory controller completing that access.

UPM special operation modes are set in registers MxMR[OP], see [Section 10.3.1.4, “UPM Mode Registers \(MxMR\).”](#) FCM special operation modes are set in FMR[OP], see [Section 10.3.1.17, “Flash Mode Register \(FMR\).”](#) Writing LSOR has the same effect as setting a special controller mode and performing a dummy access to a bank associated with the controller in question, but use of LSOR avoids changing settings for the address space occupied by the bank. More details of special operation sequences appear in [Section 10.4.4.2.1, “UPM Programming Example \(Two Sequential Writes to the RAM Array\).”](#)



**Figure 10-11. Special Operation Initiation Register (LSOR)**


[Table 10-14](#) describes LSOR.

**Table 10-14. LSOR Field Description**

Bits	Name	Description
0–28	—	Reserved
29–31	BANK	Bank on which a special operation is initiated. If the bank identified by BANK is marked valid (BRn[V] set) and the bank is controlled by a memory controller whose current mode OP is non-zero—or a special operation—eLBC will request the special operation to be activated on the selected bank when this field is written. Otherwise, writing this field has no effect. 000 Bank 0 is triggered for special operation ... 011 Bank 3 is triggered for special operation 100–111 Reserved

### 10.3.1.8 UPM Refresh Timer (LURT)

The UPM refresh timer (LURT), shown in [Figure 10-12](#), generates a refresh request for all valid banks that selected a UPM machine and are refresh-enabled (MxMR[RFEN] = 1). Each time the timer expires, a qualified bank generates a refresh request using the selected UPM. The qualified banks rotate their requests.



**Figure 10-12. UPM Refresh Timer (LURT)**



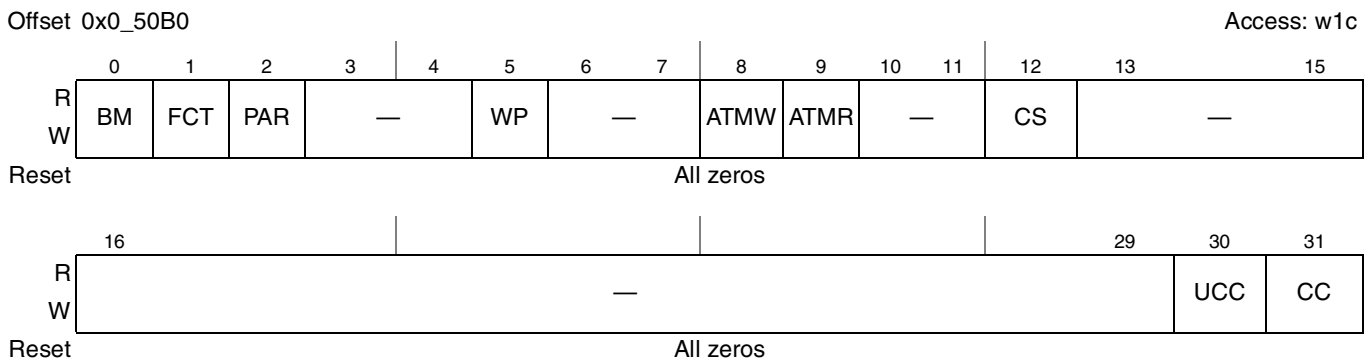
Table 10-15 describes LURT fields.

**Table 10-15. LURT Field Descriptions**

Bits	Name	Description
0–7	LURT	<p>UPM refresh timer period. Determines, along with the timer prescaler (MRTPR), the timer period according to the following equation:</p> $\text{TimerPeriod} = \frac{\text{LURT}}{\left(\frac{F_{\text{systemclock}}}{\text{MRTPR}[\text{PTP}]}\right)}$ <p>Example: For a 266-MHz system clock and a required service rate of 15.6 <math>\mu\text{s}</math>, given MRTPR[PTP] = 32, the LURT value should be 128 decimal. <math>128/(266 \text{ MHz}/32) = 15.4 \mu\text{s}</math>, which is less than the required service period of 15.6 <math>\mu\text{s}</math>. Note that the reset value (0x00) sets the maximum period to 256 x MRTPR[PTP] system clock cycles.</p>
8–31	—	Reserved

### 10.3.1.9 Transfer Error Status Register (LTESR)

The transfer error status register (LTESR) indicates the cause of an error or event. LTESR, shown in Figure 10-13, is a write-1-to-clear register. Reading LTESR occurs normally; however, write operations can clear but not set bits. A bit is cleared whenever the register is written, and the data in the corresponding bit location is a 1. For example, to clear only the write protect error bit (LTESR[WP]) without affecting other LTESR bits, 0x0400\_0000 should be written to the register. After any error/event reported by LTESR, LTEATR[V] must be cleared for LTESR to updated again.



**Figure 10-13. Transfer Error Status Register (LTESR)**

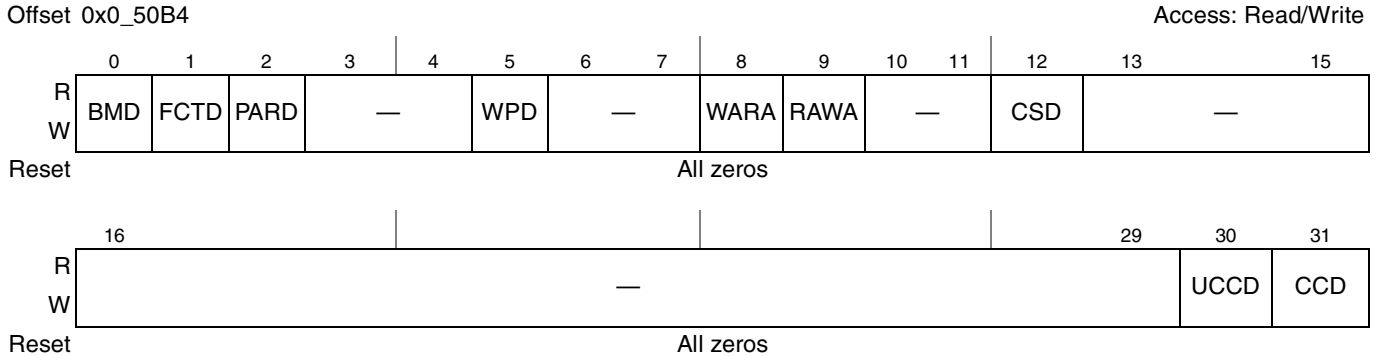
Table 10-16 describes LTESR fields.

**Table 10-16. LTESR Field Descriptions**

Bits	Name	Description
0	BM	Bus monitor time-out 0 No local bus monitor time-out occurred. 1 Local bus monitor time-out occurred. No data beat was acknowledged on the bus within LBCR[BMT] x LBCR[BMTPS] bus clock cycles from the start of a transaction.
1	FCT	FCM command time-out 0 No FCM command time-out occurred. 1 A CW0, CW1, CW2, or CW3 command issued to FCM timed-out with respect to the timer configured by FMR[CWTO].
2	PAR	ECC error for FCM mode 0 No local bus ECC error 1 Uncorrectable ECC error (FCM). LTEATR[PB] indicates the block that caused the error and LTEATR[BNK] indicates which memory controller bank was accessed.
3–4	—	Reserved
5	WP	Write protect error 0 No write protect error occurred. 1 A write was attempted to a local bus memory region that was defined as read-only in the memory controller. Usually, in this case, a bus monitor time-out will occur (as the cycle is not automatically terminated).
6–7	—	Reserved
8	ATMW	Atomic error write 0 No atomic write error occurred. 1 The subsequent write (WARA) to a memory bank did not occur within 256 bus clock cycles.
9	ATMR	Atomic error read 0 No atomic read error occurred. 1 The subsequent read (RAWA) to a memory bank did not occur within 256 bus clock cycles.
10–11	—	Reserved
12	CS	Chip select error 0 No chip select error occurred. 1 A transaction was sent to the eLBC that did not hit any memory bank.
13–29	—	Reserved
30	UCC	UPM Run pattern (MxMR[OP]=11) command completion event 0 No UPM Run pattern operation in progress, or operation pending. 1 UPM Run pattern operation has completed, allowing software to continue processing of results.
31	CC	FCM command completion event 0 No FCM operation in progress, or operation pending. 1 FCM operation has completed, allowing software to continue processing of results.

### 10.3.1.10 Transfer Error Check Disable Register (LTEDR)

The transfer error check disable register (LTEDR), shown in Figure 10-14, is used to disable error/event checking. Note that control of error/event checking is independent of control of reporting of errors/events (LTEIR) through the interrupt mechanism.



**Figure 10-14. Transfer Error Check Disable Register (LTEDR)**

Table 10-17 describes LTEDR fields.

**Table 10-17. LTEDR Field Descriptions**


Bits	Name	Description
0	BMD	Bus monitor disable 0 Bus monitor is enabled. 1 Bus monitor is disabled, but internal bus time-outs can still occur.
1	FCTD	FCM command time-out disable 0 FCM command timer is enabled. 1 FCM command time-out is disabled, but internal FCM command timer can terminate command waits.
2	PARD	ECC error checking disabled. 0 ECC error checking is enabled. 1 ECC error checking is disabled.
3–4	—	Reserved
5	WPD	Write protect error checking disable. 0 Write protect error checking is enabled. 1 Write protect error checking is disabled.
6–7	—	Reserved
8	WARA	Write after read atomic (WARA) error checking disable. 0 WARA error checking is enabled. 1 WARA error checking is disabled.
9	RAWA	Read after write atomic (RAWA) error checking disable. 0 RAWA error checking is enabled. 1 RAWA error checking is disabled.
10–11	—	Reserved
12	CSD	Chip select error checking disable. 0 Chip select error checking is enabled. 1 Chip select error checking is disabled.

**Table 10-17. LTEDR Field Descriptions (continued)**

Bits	Name	Description
13–29	—	Reserved
30	UCCD	UPM Run pattern command completion checking disable. 0 UPM Run pattern command completion checking is enabled. 1 UPM Run pattern command completion checking is disabled.
31	CCD	FCM command completion checking disable. 0 Command completion checking is enabled. 1 Command completion checking is disabled.

### 10.3.1.11 Transfer Error Interrupt Enable Register (LTEIR)

The transfer error interrupt enable register (LTEIR), shown in [Figure 10-15](#), is used to send or block error/event reporting through the eLBC internal interrupt mechanism. Software should clear pending errors/events in LTESR before enabling interrupts. After an interrupt has occurred, clearing relevant LTESR error/event bits negates the interrupt.

**Figure 10-15. Transfer Error Interrupt Enable Register (LTEIR)**

[Table 10-18](#) describes LTEIR fields.

**Table 10-18. LTEIR Field Descriptions**

Bits	Name	Description
0	BMI	Bus monitor error interrupt enable. 0 Bus monitor error reporting is disabled. 1 Bus monitor error reporting is enabled.
1	FCTI	FCM command time-out interrupt enable. 0 FCM command time-out error reporting is disabled. 1 FCM command time-out error reporting is enabled.
2	PARI	ECC error interrupt enable. 0 ECC error reporting is disabled. 1 ECC error reporting is enabled.
3–4	—	Reserved

Table 10-18. LTEIR Field Descriptions (continued)

Bits	Name	Description
5	WPI	Write protect error interrupt enable. 0 Write protect error reporting is disabled. 1 Write protect error reporting is enabled.
6–7	—	Reserved
8	WARA	Write after read atomic (WARA) error interrupt enable. 0 WARA error reporting is disabled. 1 WARA error reporting is enabled.
9	RAWA	Read after write atomic (RAWA) error interrupt enable. 0 RAWA error reporting is disabled. 1 RAWA error reporting is enabled.
10–11	—	Reserved
12	CSI	Chip select error interrupt enable. 0 Chip select error reporting is disabled. 1 Chip select error reporting is enabled.
13–29	—	Reserved
30	UCCI	UPM Run pattern command completion Event interrupt enable. 0 UPM Run pattern command completion reporting is disabled. 1 UPM Run pattern command completion reporting is enabled.
31	CCI	FCM command completion Event interrupt enable. 0 Command completion reporting is disabled. 1 Command completion reporting is enabled.

### 10.3.1.12 Transfer Error Attributes Register (LTEATR)

The transfer error attributes register (LTEATR) captures source attributes of an error/event. Figure 10-16 shows the LTEATR. After LTEATR[V] has been set, software must clear this bit to allow LTESR, LTEATR, and LTEAR to update following any subsequent events/errors.




Figure 10-16. Transfer Error Attributes Register (LTEATR)


Table 10-19 describes LTEATR fields.

**Table 10-19. LTEATR Field Descriptions**

Bits	Name	Description
0–2	—	Reserved
3	RWB	Transaction type for the error: 0 The transaction for the error was a write transaction. 1 The transaction for the error was a read transaction.
4–10	—	Reserved
11–15	SRCID	Captures the source of the transaction when this information is provided on the internal interface to the eLBC.
16–19	PB	Error on block for FCM. For FCM, there are at most four 512-byte page blocks (for a large page device) checked by ECC. A bit is set for the 512-byte block that had an uncorrectable ECC error on read (bit 16 represents block 0, the first 512 bytes of a page; if ORx[PGS] = 0, bits 17–19 are always 0).
20–23	BNK	Memory controller bank. There is one error status bit per memory controller bank (bit 20 represents bank 0). A bit is set for the local bus memory controller bank that had an error.
24–30	—	Reserved
31	V	Error attribute capture is valid. Indicates that the captured error information is valid. 0 Captured error attributes and address are not valid. 1 Captured error attributes and address are valid.

### 10.3.1.13 Transfer Error Address Register (LTEAR)

The transfer error address register (LTEAR) captures the address of a transaction that caused an error/event. The transfer error address register (LTEAR) is shown in Figure 10-17.



**Figure 10-17. Transfer Error Address Register (LTEAR)**

Table 10-20 describes LTEAR fields.

**Table 10-20. LTEAR Field Descriptions**

Bits	Name	Description
0–31	A	Transaction address for the error. For GPCM and UPM, holds the 32-bit address of the transaction resulting in an error. For FCM, this register is undefined.

### 10.3.1.14 Transfer Error ECC Register (LTECCR)

The transfer error ECC register (LTECCR) captures single bit and multibit errors per 512-byte sector in FCM mode. LTECCR, shown in Figure 10-18, is a write-1-to-clear register. Write operations can clear but not set bits. It captures the errors during full page read transfers on FCM command completion event, provided ECC check is enabled in BRx[DECC].




Figure 10-18. Transfer Error ECC Register (LTECCR)

Table 10-21. LTECCR Field Descriptions

Bits	Name	Description
0–11	—	Reserved
12–15	SBCE	Single bit correctable error There are at most four 512-byte page blocks (for a large page device) checked by ECC. A bit is set for the 512-byte block that had a single bit correctable ECC error on read (bit 12 represents block 0, the first 512 bytes of a page; if ORx[PGS] = 0, bits 13–15 are always 0).
16–27	—	Reserved
28–31	MBUE	Multi bit uncorrectable error There are at most four 512-byte page blocks (for a large page device) checked by ECC. A bit is set for the 512-byte block that had an uncorrectable ECC error on read (bit 28 represents block 0, the first 512 bytes of a page; if ORx[PGS] = 0, bits 29–31 are always 0).

### 10.3.1.15 Local Bus Configuration Register (LBCR)

The local bus configuration register (LBCR) is shown in [Figure 10-19](#).




Figure 10-19. Local Bus Configuration Register

Table 10-22 describes LBCR fields.

**Table 10-22. LBCR Field Descriptions**

Bits	Name	Description
0	LDIS	Local bus disable 0 Local bus is enabled. 1 Local bus is disabled. No internal transactions will be acknowledged.
1–7	—	Reserved
8–9	BCTLC	Defines the use of LBCTL 00 LBCTL is used as $\overline{W/R}$ control for GPCM or UPM accesses (buffer control). 01 LBCTL is used as $\overline{LOE}$ for GPCM accesses only. 10 LBCTL is used as $\overline{LWE}$ for GPCM accesses only. 11 Reserved.
10	AHD	Address hold disable. Removes part of the hold time for LAD with respect to LALE in order to lengthen the LALE pulse. 0 During address phases on the local bus, the LALE signal negates 2 platform clock period prior to the address being invalidated. At 66.6 MHz, this provides 3 ns of address hold time at the external address latch. 1 During address phases on the local bus, the LALE signal negates 1 platform clock period prior to the address being invalidated. This halves the address hold time, but extends the latch enable duration. This may be necessary for very high frequency designs.
11–15	—	Reserved. Reads to bit 13 return 1.
16–23	BMT	Bus monitor timing. Defines the bus monitor time-out period. Clearing BMT (reset value) selects the maximum count of bus clock cycles. For non-zero values of BMT, the number of LCLK clock cycles to count down before a time-out error is generated is given by: bus cycles = BMT × PS, where PS is set according to LBCR[BMTPS]. The value of BMT × PS must not be less than 40 bus cycles for reliable operation.
24–27	—	Reserved
28–31	BMTPS	Bus monitor timer prescale. Defines the multiplier, PS, to scale LBCR[BMT] for determining bus time-outs. 0000 PS = 8 0001 PS = 16 0010 PS = 32 0011 PS = 64 0100 PS = 128 0101 PS = 256 0110 PS = 512 0111 PS = 1024 1000 PS = 2048 1001 PS = 4096 1010 PS = 8192 1011 PS = 16,384 1100 PS = 32,768 1101 PS = 65,536 1110 PS = 131,072 1111 PS = 262,144




### 10.3.1.16 Clock Ratio Register (LCRR)

The clock ratio register, shown in [Figure 10-20](#), sets the system clock to eLBC bus frequency ratio. It also provides configuration bits for extra delay cycles for address and control signals.

#### NOTE

For proper operation of the system, it is required that this register setting will not be altered while local bus memories or devices are being accessed. Special care needs to be taken when running instructions from an eLBC memory.



**Figure 10-20. Clock Ratio Register (LCRR)**

[Table 10-23](#) describes LCRR fields.

**Table 10-23. LCRR Field Descriptions**

Bits	Name	Description
0–13	—	Reserved
14–15	EADC	External address delay cycles of LCLK. Defines the number of cycles for the assertion of LALE. 00 4 01 1 10 2 11 3
16–26	—	Reserved
27–31	CLKDIV	System clock divider. Sets the frequency ratio between the system clock and the local bus clock. The system clock is equivalent to <code>csb_clk</code> or twice <code>csb_clk</code> (if <code>RCWL[LBIUCM]</code> is set). Only the values shown below are allowed. <b>Note:</b> It is critical that no transactions are being executed via the local bus while <code>CLKDIV</code> is being modified. As such, prior to modification, the user must ensure that code is not executing out of the local bus. Once <code>LCRR[CLKDIV]</code> is written, the register should be read, and then an <code>isync</code> should be executed. 00000–00001 Reserved 00010 2 00011 Reserved 00100 4 00101–00111 Reserved 01000 8 01001–11111 Reserved

### 10.3.1.17 Flash Mode Register (FMR)

The local bus Flash mode register (FMR), shown in Figure 10-21, controls global operation of the FCM.




Figure 10-21. Flash Mode Register

<sup>1</sup> Bit R (field BOOT) is set if power-on-reset configuration selects FCM as the boot ROM target.

Table 10-24 describes FMR fields.

Table 10-24. FMR Field Descriptions

Bits	Name	Description
0–15	—	Reserved
16–19	CWTO	Command wait time-out. For FCM commands that wait on LFR $\bar{B}$ being sampled high (CW0, CW1, RBW and RSW), FCM pauses execution of the instruction sequence until either LFR $\bar{B}$ is sampled high, or a timer controlled by CTO expires, whichever occurs first. The time-out in the latter case is: 0000 256 cycles of LCLK 0001 512 cycles of LCLK 0010 1024 cycles of LCLK 0011 2048 cycles of LCLK 0100 4096 cycles of LCLK 0101 8192 cycles of LCLK 0110 16,384 cycles of LCLK 0111 32,768 cycles of LCLK 1000 65,536 cycles of LCLK 1001 131,072 cycles of LCLK 1010 262,144 cycles of LCLK 1011 524,288 cycles of LCLK 1100 1,048,576 cycles of LCLK 1101 2,097,152 cycles of LCLK 1110 4,194,304 cycles of LCLK 1111 8,388,608 cycles of LCLK
20	BOOT	Flash auto-boot load mode. During system boot from NAND Flash EEPROM, this bit remains set to alter the use of the FCM buffer RAM. Software should clear BOOT once FCM is to be restored to normal operation. Setting BOOT without auto-boot in progress only alters the mapping of the buffer RAM. 0 FCM is operating in normal functional mode, with an 8 Kbyte FCM buffer RAM. 1 eLBC has been configured—either from reset or by a special operation OP = 01—to auto-load a 4-Kbyte boot block into the FCM buffer RAM, which maps only the 4 Kbytes of NAND flash main data region comprising the boot block. Any access to the buffer RAM is delayed until the entire boot block has been loaded.
21–22	—	Reserved

Table 10-24. FMR Field Descriptions (continued)

Bits	Name	Description
23	ECCM	ECC mode. When hardware checking and/or generation of error correcting codes (ECC) is enabled (that is, when $BR_n[DECC]$ is 01 or 10, and full page transfers are specified with $FBCR[BC] = 0$ ), ECCM sets the ECC block size and position of the ECC code word(s) in the NAND Flash spare region for both checking and generation functions. The format of the ECC code word conforms with the Samsung/Toshiba spare region assignment specifications. 0 ECC is checked/calculated over 512-Byte blocks. A 24-bit ECC is assigned to spare region bytes at offsets $(N \times 16) + 6$ through $(N \times 16) + 8$ for spare region $N$ , $N = 0-3$ . 1 ECC is checked/calculated over 512-Byte blocks. A 24-bit ECC is assigned to spare region bytes at offsets $(N \times 16) + 8$ through $(N \times 16) + 10$ for spare region $N$ , $N = 0-3$ .
24–25	—	Reserved
26–27	AL	Address length. AL sets the number of address bytes issued during page address (PA) operations. However, the number of address bytes issued for column address (CA) operations is determined by the device page size (for $OR_n[PGS] = 0$ , 1 CA byte is issued; for $OR_n[PGS] = 1$ , 2 CA bytes are issued). 00 2 bytes are issued for page addresses, thus a total of 3 ( $OR_n[PGS] = 0$ ) or 4 ( $OR_n[PGS] = 1$ ) address bytes are issued for a {CA,PA} sequence 01 3 bytes are issued for page addresses, thus a total of 4 ( $OR_n[PGS] = 0$ ) or 5 ( $OR_n[PGS] = 1$ ) address bytes are issued for a {CA,PA} sequence 10 4 bytes are issued for page addresses, thus a total of 5 ( $OR_n[PGS] = 0$ ) or 6 ( $OR_n[PGS] = 1$ ) address bytes are issued for a {CA,PA} sequence 11 —
28–29	—	Reserved
30–31	OP	Flash operation. For OP not equal to 00, a special operation is triggered on the next write to LSOR or dummy access to a bank controlled by FCM. Once a special operation has commenced, OP is automatically reset to 00 by FCM. Individual blocks may be temporarily unlocked for erase and reprogramming operations. 00 Normal operation. All read and write accesses to banks controlled by FCM access the shared FCM buffer RAM. No bus activity is caused by this operation. 01 Simulate auto-boot block loading, and set $FMR[BOOT]$ . Boot block loading occurs from the bank triggered on the special operation, therefore the appropriate bank configuration must be initialized prior to issuing this operation. 10 Execute the command sequence contained in FIR, but with write protection enabled (pin $\overline{LFWP}$ asserted low) so that all Flash blocks are protected from accidental erasure and reprogramming. 11 Execute the command sequence contained in FIR, but permit the single block identified by $FBAR[BLK]$ to be erased or reprogrammed, with pin $\overline{LFWP}$ remaining high during the access.

### 10.3.1.18 Flash Instruction Register (FIR)

The local bus Flash instruction register (FIR), shown in Figure 10-22, holds a sequence of up to eight instructions for issue by the FCM. Setting  $FMR[OP]$  non-zero and writing LSOR or accessing a bank controlled by FCM causes FCM to read FIR 4 bits at a time, starting at bit 0 and continuing with adjacent 4-bit opcodes, until only NOP opcodes remain. The programmed instruction sequence of OP0, OP1, ..., OP7 is performed on the activated bank, using the data buffer addressed by FPAR. If  $LTEIR[CCI] = 1$  and  $LTEDR[CCD] = 0$ , eLBC will generate an interrupt once the entire sequence has completed, and software should examine LTEATR and clear its V bit.

Software must not alter the contents of the addressed FCM buffer, FIR, MDR, FCR, FBAR, FPAR, or FBCR while an operation is in progress—or eLBC will behave unpredictably—but software can freely

modify the contents of any currently unused FCM RAM buffer in preparation for the next operation.

Offset 0x0\_50E4

Access: Read/Write




Figure 10-22. Flash Instruction Register

Table 10-25 describes FIR fields.

Table 10-25. FIR Field Descriptions

Bits	Name	Description
0–3	OP0	FCM operation codes. OP0 is executed first, followed by OP1, through to OP7.
4–7	OP1	0000 NOP—No-operation and end of operation sequence
8–11	OP2	0001 CA—Issue current column address as set in FPAR, with length set by ORx[PGS]
		0010 PA—Issue current block+page address as set in FBAR and FPAR, with length set by FMR[AL]
		0011 UA—Issue user-defined address byte from next AS field in MDR
12–15	OP3	0100 CM0—Issue command from FCR[CMD0]
		0101 CM1—Issue command from FCR[CMD1]
16–19	OP4	0110 CM2—Issue command from FCR[CMD2]
		0111 CM3—Issue command from FCR[CMD3]
20–23	OP5	1000 WB—Write FBCR bytes of data from current FCM buffer to Flash device
24–27	OP6	1001 WS—Write one byte (8b port) of data from next AS field of MDR to Flash device
		1010 RB—Read FBCR bytes of data from Flash device into current FCM RAM buffer
28–31	OP7	1011 RS—Read one byte (8b port) of data from Flash device into next AS field of MDR
		1100 CW0—Wait for LFR $\bar{B}$ to return high or time-out, then issue command from FCR[CMD0]
		1101 CW1—Wait for LFR $\bar{B}$ to return high or time-out, then issue command from FCR[CMD1]
		1110 RBW—Wait for LFR $\bar{B}$ to return high or time-out, then read FBCR bytes of data from Flash device into current FCM RAM buffer
		1111 RSW—Wait for LFR $\bar{B}$ to return high or time-out, then read one byte (8b port) of data from Flash device into next AS field of MDR

### 10.3.1.19 Flash Command Register (FCR)

The local bus Flash command register (FCR), shown in Figure 10-23, holds up to four NAND Flash EEPROM command bytes that may be referenced by opcodes in FIR during FCM operation. The values of the commands should follow the manufacturer’s datasheet for the relevant NAND Flash device.

Offset 0x0\_50E8

Access: Read/Write




Figure 10-23. Flash Command Register

Table 10-26 describes FCR fields.

**Table 10-26. FCR Field Descriptions**


Bits	Name	Description
0–7	CMD0	General purpose FCM Flash command byte 0. Opcodes in FIR that issue command index 0 write CMD0 to the NAND Flash command/data bus.
8–15	CMD1	General purpose FCM Flash command byte 1. Opcodes in FIR that issue command index 1 write CMD1 to the NAND Flash command/data bus.
16–23	CMD2	General purpose FCM Flash command byte 2. Opcodes in FIR that issue command index 2 write CMD2 to the NAND Flash command/data bus.
24–31	CMD3	General purpose FCM Flash command byte 3. Opcodes in FIR that issue command index 3 write CMD3 to the NAND Flash command/data bus.

### 10.3.1.20 Flash Block Address Register (FBAR)

The local bus Flash block address register (FBAR), shown in Figure 10-24, locates the NAND Flash block index for the page currently accessed.

Offset 0x0\_50EC

Access: Read/Write



**Figure 10-24. Flash Block Address Register**

Table 10-27 describes FBAR fields.

**Table 10-27. FBAR Field Descriptions**


Bits	Name	Description
0–7	—	Reserved
8–31	BLK	Flash block address. The size of the NAND Flash, as configured in ORn[PGS] and FMR[AL], determines the number of bits of BLK that are issued to the EEPROM during block address phases.

### 10.3.1.21 Flash Page Address Register (FPAR)


The local bus Flash page address register (FPAR), shown in Figure 10-25 and Figure 10-26, locates the current NAND Flash page in both the external NAND Flash device and FCM buffer RAM.

Offset 0x0\_50F0

Access: Read/Write



**Figure 10-25. Flash Page Address Register, Small Page Device (ORx[PGS] = 0)**



**Figure 10-26. Flash Page Address Register, Large Page Device (ORx[PGS] = 1)**

Table 10-28 describes FPAR fields for small page devices.

**Table 10-28. FPAR Field Descriptions, Small Page Device (ORx[PGS] = 0)**

Bits	Name	Description
0–16	—	Reserved
17–21	PI	Page index. PI indexes the page in NAND Flash EEPROM at the current block defined by FBAR, and locates the corresponding transfer buffer in the FCM buffer RAM. The 3 LSBs of PI index one of the eight 1 Kbyte buffers in the FCM buffer RAM as follows: 000 The page is transferred to/from FCM buffer 0, address offsets 0x0000–0x03FF 001 The page is transferred to/from FCM buffer 1, address offsets 0x0400–0x07FF 010 The page is transferred to/from FCM buffer 2, address offsets 0x0800–0x0BFF 011 The page is transferred to/from FCM buffer 3, address offsets 0x0C00–0x0FFF 100 The page is transferred to/from FCM buffer 4, address offsets 0x1000–0x13FF 101 The page is transferred to/from FCM buffer 5, address offsets 0x1400–0x17FF 110 The page is transferred to/from FCM buffer 6, address offsets 0x1800–0x1BFF 111 The page is transferred to/from FCM buffer 7, address offsets 0x1C00–0x1FFF
22	MS	Main/spare region locator. In the case that FBCR[BC] = 0, MS is treated as 0. 0 Data is transferred to/from the main region of the FCM buffer; that is, the first 512 bytes of the buffer are used as the starting address. 1 Data is transferred to/from the spare region of the FCM buffer; that is, the second 512 bytes of the buffer are used as the starting address, but only an initial 16 bytes of spare region are defined.
23–31	CI	Column index. CI indexes the first byte to transfer to/from the main or spare region of the NAND Flash EEPROM and corresponding transfer buffer. In the case that FBCR[BC] = 0, CI is treated as 0. For MS = 0, CI can range 0x000–0x1FFF; for MS = 1, CI can range 0x000–0x00F.

Table 10-29 describes FPAR fields for large page devices.

**Table 10-29. FPAR Field Descriptions, Large Page Device (ORx[PGS] = 1)**


Bits	Name	Description
0–13	—	Reserved
14–19	PI	Page index. PA indexes the page in NAND Flash EEPROM at the current block defined by FBAR, and locates the corresponding transfer buffer in the FCM buffer RAM. The LSB of PI indexes one of the two 4 Kbyte buffers in the FCM buffer RAM as follows: 0 The page is transferred to/from FCM buffer 0, address offsets 0x0000–0x0FFF 1 The page is transferred to/from FCM buffer 1, address offsets 0x1000–0x1FFF

**Table 10-29. FPAR Field Descriptions, Large Page Device (ORx[PGS] = 1) (continued)**

Bits	Name	Description
20	MS	Main/spare region locator. In the case that FBCR[BC] = 0, MS is treated as 0. 0 Data is transferred to/from the main region of the FCM buffer; that is, the first 2048 bytes of the buffer are used as the starting address. 1 Data is transferred to/from the spare region of the FCM buffer; that is, the second 2048 bytes of the buffer are used as the starting address, but only an initial 64 bytes of spare region are defined.
21–31	CI	Column index. CI indexes the first byte to transfer to/from the main or spare region of the NAND Flash EEPROM and corresponding transfer buffer. In the case that FBCR[BC] = 0, CI is treated as 0. For MS = 0, CI can range 0x000–0x7FF; for MS = 1, CI can range 0x000–0x03F.

### 10.3.1.22 Flash Byte Count Register (FBCR)

The local bus Flash byte count register (FBCR), shown in [Figure 10-27](#), defines the size of FCM block transfers for reads and writes to the NAND Flash EEPROM.

**Figure 10-27. Flash Byte Count Register**


[Table 10-30](#) describes FBCR fields.

**Table 10-30. FBCR Field Descriptions**

Bits	Name	Description
0–19	—	Reserved
20–31	BC	Byte count determines how many bytes are transferred by the FCM during data read (RB) or data write (WB) opcodes. The first byte accessed in the NAND Flash EEPROM is located by the FPAR register, and successive bytes are transferred until either BC bytes have been counted, or the end of the spare region of the currently addressed Flash page has been reached. If BC = 0, an entire Flash page and its spare region will be transferred by FCM, in which case FPAR[MS] and FPAR[CI] are treated as zero regardless of their values. BC = 0 is the only setting that permits FCM to generate and check ECC.

### 10.3.1.23 Flash ECC Block $n$ Register (FECC0–FECC3)

The local bus Flash ECC block $n$  register (FECC $n$ ), shown in [Figure 10-28](#), specifies the ECC value calculated during writes or reads by eLBC. It can be used for verify after write feature in software. Note that the valid bit sets before the command completion event and hence, the correct ECC could be read before actual completion of writes/reads.

Figure 10-28. Flash ECC Block $n$  Register (FECC0–FECC3)Table 10-31. FECC $n$  Field Descriptions

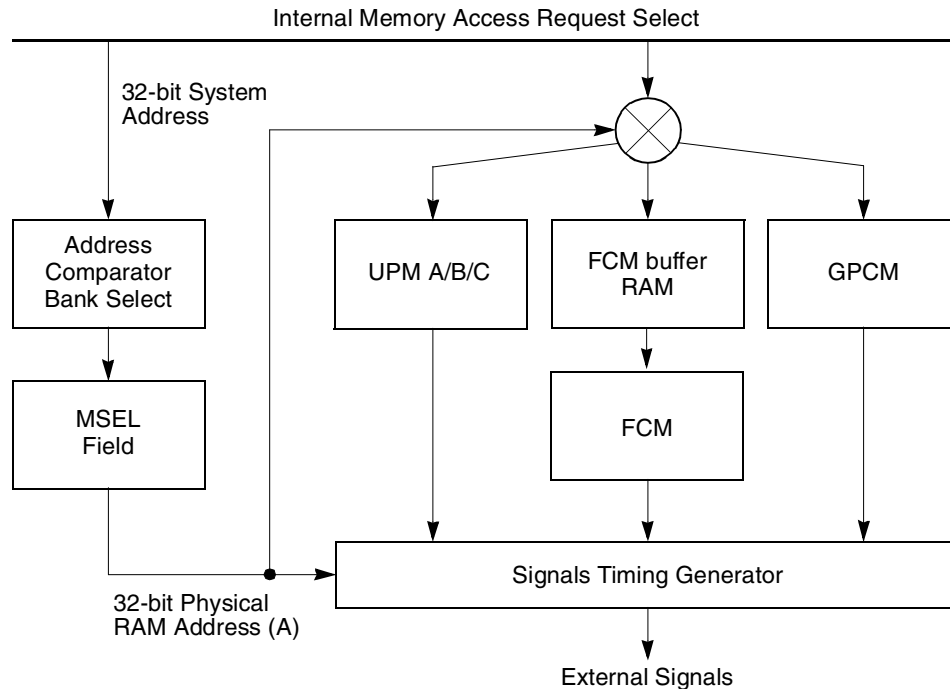
Bits	Name	Description
0	V	Valid bit. This bit denotes that the ECC stored in this register is valid. It is set for full page write/read transfers if ECC generation/checking is enabled in BR $n$ [DECC].
1–7	—	Reserved
8–31	ECC	24 bit ECC; For $n^{\text{th}}$ 512 bytes of a page in case of large page or for $(4k + n)^{\text{th}}$ 512 byte page for small page where $k = 0,1,2,\dots$ . It stores calculated ECC value during writes/reads.

## 10.4 Functional Description

The eLBC allows the implementation of memory systems with very specific timing requirements.

- The GPCM provides interfacing for simpler, lower-performance memories and memory-mapped devices. It has inherently lower performance because it does not support bursting. For this reason, GPCM-controlled banks are used primarily for boot-loading from NVRAM or NOR Flash, and access to low-performance memory-mapped peripherals.
- The FCM interfaces the eLBC to NAND Flash EEPROMs with 8-bit data bus. The FCM has an automatic boot-loading feature that allows the CPU to boot from high density EEPROM, loading the boot block into 4 Kbytes of RAM for execution of the first level boot code. Following boot, FCM provides a flexible instruction sequencer that allows a user-defined command, address, and data transfer sequence of up to 8 steps to be executed against a memory-mapped buffer RAM. Programmable set-up time, hold time, and wait states permit the FCM to maximize the performance of NAND Flash block transfers, which can proceed in parallel with software processing of the multiple RAM buffers. A single-pass ECC engine in the FCM permits zero-overhead error checking, reporting, and correction in both boot blocks and page data transfers if enabled.
- The UPM supports refresh timers, address multiplexing of the external bus, and generation of programmable control signals for row address and column address strobes, to allow for a minimal glue logic interface to DRAMs, burstable SRAMs, and almost any other kind of peripheral with asynchronous timing or single data rate clocking. The UPM can be used to generate flexible, user-defined timing patterns for control signals that govern a memory device. These patterns define how the external control signals behave during a read, write, burst-read, or burst-write access. Refresh timers are also available to periodically initiate user-defined refresh patterns.





**Figure 10-29. Basic Operation of Memory Controllers in the eLBC**

Each memory bank (chip select) can be assigned to any of these three types of machines through the machine select bits of the base register for that bank ( $BR_n[MSEL]$ ), as illustrated in Figure 10-29. If a bank match occurs, the corresponding machine (GPCM, FCM, or UPM) then takes ownership of the external signals that control the access and maintains control until the transaction ends.

## 10.4.1 Basic Architecture

The following subsections describe the basic architecture of the eLBC.

### 10.4.1.1 Address and Address Space Checking

The defined base addresses are written to the  $BR_n$  registers, while the corresponding address masks are written to the  $OR_n$  registers. Each time a local bus access is requested, the internal transaction address is compared with each bank. Addresses are decoded by comparing the 17 MSBs of the address, masked by  $OR_n[AM]$ , with the base address for each bank ( $BR_n[BA]$ ). If a match is found on a memory controller bank, the attributes defined in the  $BR_n$  and  $OR_n$  for that bank are used to control the memory access. If a match is found in more than one bank, the lowest-numbered bank handles the memory access (that is, bank 0 has priority over bank 1).

### 10.4.1.2 External Address Latch Enable Signal (LALE)

The local bus uses a multiplexed address/data bus. Therefore the eLBC must distinguish between address and data phases, which take place on the same bus (LAD pins). The LALE signal, when asserted, signifies an address phase during which the eLBC drives the memory address on the LAD pins. An external address

latch uses this signal to capture the address and provide it to the address pins of the memory or peripheral device. If the eLBC is not used in multiplexed mode, the 26 LSBs of the address are provided by LA[0:25], unlatched, to the address pins of the memory or peripheral device. When LALE is negated, LAD then serves as the (bi-directional) data bus for the access. Any address phase initiates the assertion of LALE, which has a programmable duration of between 1 and 4 bus clock cycles.

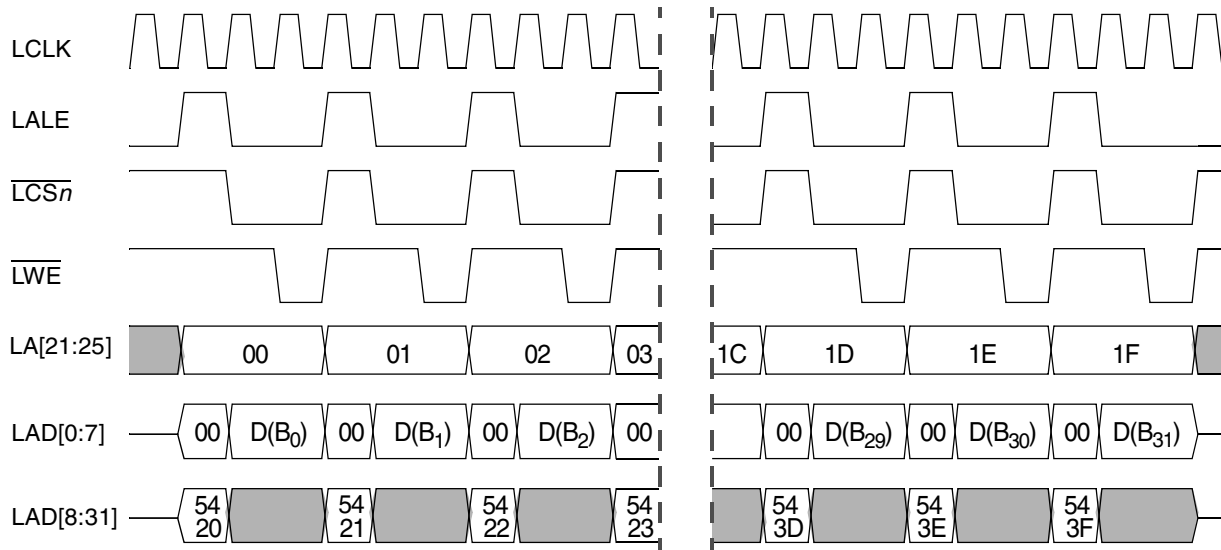
To ensure adequate hold time on the external address latch, LALE negates earlier than the address changes on LAD during address phases. By default, LALE negates earlier by 2 platform clock period. For example, if the platform clock is operating at 66.6 MHz, then 3 ns of address hold time is introduced. However, at higher frequencies, the duration of the shortened LALE pulse may not meet the minimum latch enable pulse width specifications of some latches. In such cases, setting LBCR[AHD] = 1 increases the LALE pulse width by 1 platform clock cycle, but decreases the address hold time by the same amount. If both longer hold time and longer LALE pulse duration are needed, then the address phase can be extended using the ORn[EAD] and LCRR[EADC] fields, and the LBCR[AHD] bit can be left at 0. However, this will add latency to all address tenures.

The frequency of LALE assertion varies across the three memory controllers:

- For GPCM, every assertion of  $\overline{\text{LCS}}_n$  is considered an independent access, and accordingly, LALE asserts prior to each such access. For example, GPCM driving an 8-bit port would assert LALE and  $\overline{\text{LCS}}_n$  32 times in order to satisfy a 32-byte cache line transfer.
- For FCM, LALE asserts prior to each multi-command operation sequence, but LALE can be ignored on NAND Flash EEPROM accesses as the signal does *not* enable address latching in such devices. The value on the LAD and LA pins during LALE assertion is driven low-impedance, but otherwise not defined for FCM banks.
- In the case of UPM, the frequency of LALE assertion depends on how the UPM RAM is programmed. UPM single accesses typically assert LALE once, upon commencement, but it is possible to program UPM to assert LALE several times, and to change the values of LA<sub>n</sub> with and without LALE being involved.

In general, when using the GPCM controller it is not necessary to use LA if a sufficiently wide latch is used to capture the entire address during LALE phases. The UPMs may require LA if the eLBC is generating its own burst address sequence.

To illustrate how a large transaction is handled by the eLBC, [Figure 10-30](#) shows eLBC signals for the GPCM performing a 32-byte write starting at address 0x5420. Note that during each of the 32 assertions of LALE, LA[21:25] exactly mirror LAD[27:31], but during data phases, only LAD[0:7] is driven with valid data.




**Note:** All address and signal values are shown in hexadecimal.  
 $D(B_k) = k^{\text{th}}$  of 32 data bytes

**Figure 10-30. Example of 8-Bit GPCM Writing 32 Bytes to Address 0x5420 (LCRR[PBYP] = 0)**

### 10.4.1.3 Data Transfer Acknowledge (TA)

The three memory controllers in the eLBC generate an internal transfer acknowledge signal, TA, to allow data on LAD to be either sampled (for reads) or changed (on writes). The data sampling/data change always occurs at the end of the bus cycle in which the eLBC asserts TA internally. In eLBC debug mode, TA is also visible externally on the LDVAL pin. The GPCM controller automatically generates TA according to the timing parameters programmed for them in the option and mode registers; FCM generates TA whenever data read and write instructions are executed out of register FIR; a UPM generates TA only when a UPM pattern has the UTA RAM word bit set. Figure 10-31 shows LALE, TA (internal), and  $\overline{\text{LCSn}}$ . Note that TA and LALE are never asserted together, and that for the duration of LALE,  $\overline{\text{LCSn}}$  (or any other control signal) remains negated or frozen.



**Figure 10-31. Basic eLBC Bus Cycle with LALE, TA, and  $\overline{\text{LCSn}}$**

### 10.4.1.4 Data Buffer Control (LBCTL)

The memory controller provides a data buffer control signal for the local bus (LBCTL). This signal is activated when a GPCM-, FCM-, or UPM-controlled bank is accessed. LBCTL can be disabled by setting  $OR_n[BCTL D]$ . LBCTL can be further configured by  $LBCR[BCTL C]$  to act as an extra  $\overline{LWE}$  or an extra  $\overline{LOE}$  signal when in GPCM mode.

If LBCTL is configured as a data buffer control ( $LBCR[BCTL C] = 00$ ), the signal is asserted (high) on the rising edge of the bus clock on the first cycle of the memory controller operation, coincident with LALE. If the access is a write, LBCTL remains high for the whole duration. However, if the access is a read, LBCTL is negated (low) with the negation of LALE so that the memory device is able to drive the bus. If back-to-back read accesses are pending, LBCTL is asserted (high) one bus clock cycle before the next transaction starts (that is, one bus clock cycle before LALE) to allow a whole bus cycle for the bus to turn around before the next address is driven.

### 10.4.1.5 Atomic Operation

The eLBC supports the following kinds of atomic bus operations (set by  $BR_n[ATOM]$ ):

- Read-after-write atomic (RAWA). When a write access hits a memory bank in which  $ATOM = 01$ , the eLBC reserves the selected memory bank for the exclusive use of the accessing master. While the bank is reserved, no other device can be granted access to this bank. The reservation is released when the master that created it accesses the same bank with a read transaction. Additional write transactions prior to the releasing read do not change reservation status, but are otherwise processed normally. If the master fails to release the reservation within 256 bus clock cycles, the reservation is released and an atomic error is reported (if enabled); additional write transactions prior to the releasing read restart the reservation timer. This feature is intended for CAM operations.
- Write-after-read atomic (WARA). When a read access hits a memory bank in which  $ATOM = 10$ , the eLBC reserves the bus for the exclusive use of the accessing master. During the reservation period, no other device can be granted access to the atomic bank. The reservation is released when the device that created it accesses the same bank with a write transaction. Additional read transactions prior to the releasing write are otherwise processed normally and do not change the reservation status. If the device fails to release the reservation within 256 bus clock cycles, the reservation is released and an atomic error is reported (if enabled); additional read transactions prior to the releasing write restart the reservation timer.

### 10.4.1.6 Bus Monitor

A bus monitor is provided to ensure that each bus cycle is terminated within a reasonable (user defined) period. When a transaction starts, the bus monitor starts counting down from the time-out value ( $LBCR[BMT] \times LBCR[BMTPS]$ ) until a data beat is acknowledged on the bus. It then reloads the time-out value and resumes the countdown until the data tenure completes and then idles if there is no pending transaction. Setting  $LTEDR[BMD]$  disables bus monitor error checking (i.e. the  $LTESR[BM]$  bit is not set by a bus monitor time-out); however, the bus monitor is still active and can generate a UPM exception (as noted in [Section 10.4.4.1.4, “Exception Requests,”](#)) or terminate a GPCM access.

It is very important to ensure that the value of LBCR[BMT] is not set too low; otherwise spurious bus time-outs may occur during normal operation—resulting in incomplete data transfers. Accordingly, the time-out value represented by the LBCR[BMT], LBCR[BMTPS] pair must not be set below 40 bus cycles for time-out under any circumstances.

## 10.4.2 General-Purpose Chip-Select Machine (GPCM)

The GPCM allows a minimal glue logic and flexible interface to SRAM, EPROM, FEPROM, ROM devices, and external peripherals. The GPCM contains two basic configuration register groups—BR<sub>n</sub> and OR<sub>n</sub>.

Figure 10-32 shows a simple connection between an 8-bit port size SRAM device and the eLBC in GPCM mode. Byte-write enable signals ( $\overline{\text{LWE}}$ ) are available for each byte written to memory. Also, the output enable signal ( $\overline{\text{LOE}}$ ) is provided to minimize external glue logic. On system reset, a global (boot) chip-select is available that provides a boot ROM chip-select ( $\overline{\text{LCS0}}$ ) prior to the system being fully configured.




Figure 10-32. Enhanced Local Bus to GPCM Device Interface

Figure 10-33 shows  $\overline{\text{LCS}}$  as defined by the setup time required between the address lines and  $\overline{\text{CE}}$ . The user can configure OR<sub>n</sub>[ACS] to specify  $\overline{\text{LCS}}$  to meet this requirement. Generally, the attributes for the memory cycle are taken from OR<sub>n</sub>. These attributes include the CSNT, ACS, XACS, SCY, TRLX, EHTR and SETA fields.





Figure 10-33. GPCM Basic Read Timing (XACS = 0, ACS = 1x, TRLX = 0, CLKDIV = 4,8)

### 10.4.2.1 GPCM Read Signal Timing

The basic GPCM read timing parameters that may be set by the  $OR_n$  attributes are shown in Figure 10-34. The read access cycle commences upon latching of the memory address (LALE negated), and concludes when LBCTL returns high to turn the local bus around for a subsequent address phase. Read data is captured by eLBC on the falling edge of TA.  $\overline{LOE}$  and  $\overline{LCSn}$  negate high simultaneously, in some cases before the end of the read access to provide additional hold time for the external memory.



**Notes:**  
 $t_{RC}$  = Read cycle time.  $t_{CSRP}$  = Read chip-select assertion period.  
 $t_{ARCS}$  = Address valid to read chip-select time.  $t_{OEN}$  = Output enable negated time.  
 $t_{AOE}$  = Address valid to output enable time.

Figure 10-34. GPCM General Read Timing Parameters

Table 10-32 lists the signal timing parameters for a GPCM read access as the option register attributes are varied.

Table 10-32. GPCM Read Control Signal Timing


Option Register Attributes				Signal Timing (LCLK clock cycles) <sup>1</sup>				
TRLX	EHTR	XACS	ACS	t <sub>ARCS</sub>	t <sub>CSRP</sub>	t <sub>AOE</sub>	t <sub>OEN</sub>	t <sub>RC</sub>
0	0	0	0X	0	2+SCY	1	0	2+SCY
0	0	0	10	¼ (½)	1¾+SCY (2+SCY)	1	0	2+SCY
0	0	0	11	½	1½+SCY	1	0	2+SCY
0	0	1	0X	0	2+SCY	1	0	2+SCY
0	0	1	10	1	1+SCY	1	0	2+SCY
0	0	1	11	2	1+SCY	2	0	3+SCY
0	1	0	0X	0	2+SCY	1	1	3+SCY
0	1	0	10	¼ (½)	1¾+SCY (1½+SCY)	1	1	3+SCY
0	1	0	11	½	1½+SCY	1	1	3+SCY
0	1	1	0X	0	2+SCY	1	1	3+SCY
0	1	1	10	1	1+SCY	1	1	3+SCY
0	1	1	11	2	1+SCY	2	1	4+SCY
1	0	0	0X	0	2+2xSCY	1	4	6+2xSCY
1	0	0	10	¼ (½)	1¾+2xSCY (1½+2xSCY)	2	4	7+2xSCY
1	0	0	11	½	1½+2xSCY	2	4	7+2xSCY
1	0	1	0X	0	2+2xSCY	1	4	6+2xSCY
1	0	1	10	2	1+2xSCY	2	4	7+2xSCY
1	0	1	11	3	1+2xSCY	3	4	8+2xSCY
1	1	0	0X	0	2+2xSCY	1	8	10+2xSCY
1	1	0	10	¼ (½)	1¾+2xSCY (1½+2xSCY)	2	8	11+2xSCY
1	1	0	11	½	1½+2xSCY	2	8	11+2xSCY
1	1	1	0X	0	2+2xSCY	1	8	10+2xSCY
1	1	1	10	2	1+2xSCY	2	8	11+2xSCY
1	1	1	11	3	1+2xSCY	3	8	12+2xSCY

<sup>1</sup> Times in parentheses are specific for the case LCRR[CLKDIV] = 2; other times apply to all CLKDIV values.

### 10.4.2.2 GPCM Write Signal Timing

The basic GPCM write timing parameters that may be set by the OR<sub>n</sub> attributes are shown in [Figure 10-35](#). The write access cycle commences upon latching of the memory address (LAL<sub>E</sub> negated), and concludes when LCS<sub>n</sub> returns high. LBCTL remains stable for the entire cycle to drive data onto any secondary data

bus. Write data becomes invalid following the falling edge of TA.  $\overline{\text{LWE}}$  may, in some cases, negate high before the end of the write access to provide additional hold time for the external memory.



**Notes:**  
 $t_{WC}$  = Write cycle time.  $t_{CSWP}$  = Write chip-select assertion period.  
 $t_{AWCS}$  = Address valid to write chip-select time.  $t_{WEN}$  = Write enable negated time wrt chip-select.  
 $t_{AWE}$  = Address valid to write enable time.

**Figure 10-35. GPCM General Write Timing Parameters**

Table 10-33 lists the signal timing parameters for a GPCM write access as the option register attributes are varied.

**Table 10-33. GPCM Write Control Signal Timing**

Option Register Attributes				Signal Timing (LCLK clock cycles) <sup>1</sup>				
TRLX	XACS	ACS	CSNT	$t_{AWCS}$	$t_{CSWP}$	$t_{AWE}$	$t_{WEN}$	$t_{WC}$
0	0	00	0	0	2+SCY	1	0	2+SCY
0	0	10	0	$\frac{1}{4}$ ( $\frac{1}{2}$ )	$1\frac{3}{4}$ +SCY (2+SCY)	1	0	2+SCY
0	0	11	0	$\frac{1}{2}$	$1\frac{1}{2}$ +SCY	1	0	2+SCY
0	1	00	0	0	2+SCY	1	0	2+SCY
0	1	10	0	1	1+SCY	1	0	2+SCY
0	1	11	0	2	1+SCY	2	0	3+SCY
0	0	00	1	0	2+SCY	1	$\frac{1}{4}$ (0)	2+SCY
0	0	10	1	$\frac{1}{4}$ ( $\frac{1}{2}$ )	$1\frac{1}{2}$ +SCY	1	0	$1\frac{3}{4}$ +SCY ( $1\frac{1}{2}$ +SCY)
0	0	11	1	$\frac{1}{2}$	$1\frac{1}{4}$ +SCY (1+SCY)	1	0	$1\frac{3}{4}$ +SCY ( $1\frac{1}{2}$ +SCY)



Table 10-33. GPCM Write Control Signal Timing (continued)

Option Register Attributes				Signal Timing (LCLK clock cycles) <sup>1</sup>				
TRLX	XACS	ACS	CSNT	t <sub>AWCS</sub>	t <sub>CSWP</sub>	t <sub>AWE</sub>	t <sub>WEN</sub>	t <sub>WC</sub>
0	1	00	1	0	2+SCY	1	¼ (0)	2+SCY
0	1	10	1	1	¾+SCY (½+SCY)	1	0	1¾+SCY (1½+SCY)
0	1	11	1	2	¾+SCY (½+SCY)	2	0	2¾+SCY (2½+SCY)
1	0	00	0	0	2+2×SCY	1	0	2+2×SCY
1	0	10	0	1¼ (1½)	1¾+2×SCY (2+2×SCY)	2	0	3+2×SCY
1	0	11	0	1½	1½+2×SCY	2	0	3+2×SCY
1	1	00	0	0	2+2×SCY	1	0	2+2×SCY
1	1	10	0	2	1+2×SCY	2	0	3+2×SCY
1	1	11	0	3	1+2×SCY	3	0	4+2×SCY
1	0	00	1	0	3+2×SCY	1	¼ (1)	3+2×SCY
1	0	10	1	1¼ (1½)	1½+2×SCY	2	0	2¾+2×SCY (2½+2×SCY)
1	0	11	1	1½	1¼+2×SCY (1+2×SCY)	2	0	2¾+2×SCY (2½+2×SCY)
1	1	00	1	0	3+2×SCY	1	¼ (1)	3+2×SCY
1	1	10	1	2	¾+2×SCY (½+2×SCY)	2	0	2¾+2×SCY (2½+2×SCY)
1	1	11	1	3	¾+2×SCY (½+2×SCY)	3	0	3¾+2×SCY (3½+2×SCY)

<sup>1</sup> Times in parentheses are specific for the case LCRR[CLKDIV] = 2; other times apply to all CLKDIV values.

### 10.4.2.3 Chip-Select Assertion Timing

The banks selected to work with the GPCM support an option to drive the  $\overline{LCSn}$  signal with different timings (with respect to the external address/data bus).  $\overline{LCSn}$  can be driven in any of the following ways:

- Simultaneous with the latched memory address. (This refers to the externally latched address and not the address timing on LAD. That is, the chip select does not assert during LALE).
- One quarter of a clock cycle later (for LCRR[CLKDIV] = 4, 8).
- One half of a clock cycle later (for LCRR[CLKDIV] = 2, 4, or 8).
- One clock cycle later (for LCRR[CLKDIV] = 4), when OR<sub>n</sub>[XACS] = 1.
- Two clock cycles later (for LCRR[CLKDIV] = 2, 4, or 8), when OR<sub>n</sub>[XACS] = 1.

- Three clock cycles later (for  $\text{LCRR}[\text{CLKDIV}] = 2, 4, \text{ or } 8$ ), when  $\text{OR}_n[\text{XACS}] = 1$  and  $\text{OR}_n[\text{TRLX}] = 1$ .

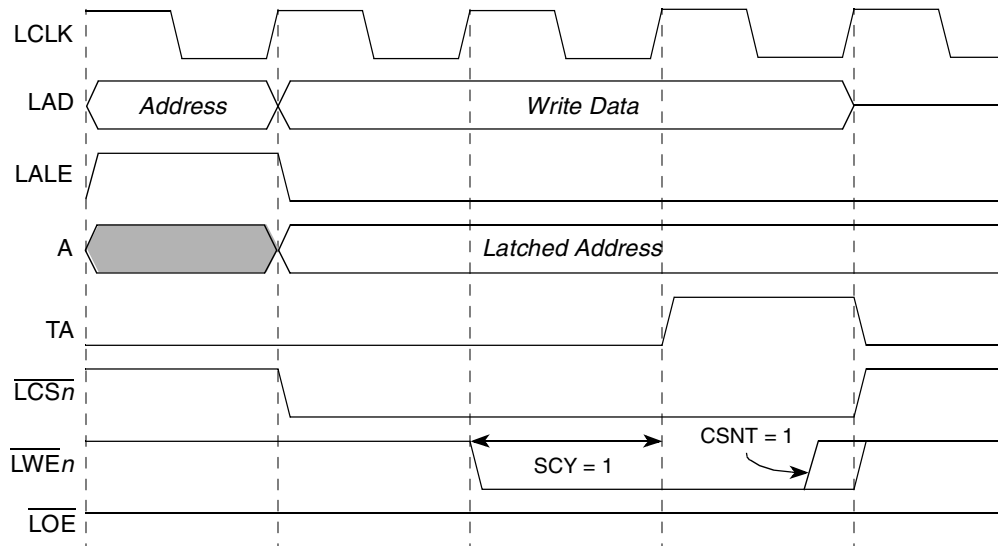
The timing diagram in [Figure 10-33](#) shows two chip-select assertion timings for the case  $\text{LCRR}[\text{CLKDIV}] = 4$  or  $8$ . If  $\text{LCRR}[\text{CLKDIV}] = 2$ ,  $\overline{\text{LCS}}_n$  asserts identically for  $\text{OR}_n[\text{ACS}] = 10$  or  $11$ .

#### 10.4.2.3.1 Programmable Wait State Configuration

The GPCM supports internal generation of transfer acknowledge. It allows between zero and 30 wait states to be added to an access by programming  $\text{OR}_n[\text{SCY}]$  and  $\text{OR}_n[\text{TRLX}]$ . Internal generation of transfer acknowledge is enabled if  $\text{OR}_n[\text{SETA}] = 0$ . If  $\overline{\text{LGTA}}$  is asserted externally two bus clock cycles or more before the wait state counter has expired (to allow for synchronization latency), the current memory cycle is terminated by  $\overline{\text{LGTA}}$ ; otherwise it is terminated by the expiration of the wait state counter. Regardless of the setting of  $\text{OR}_n[\text{SETA}]$ , wait states prolong the assertion duration of both  $\overline{\text{LOE}}$  and  $\overline{\text{LWE}}_n$  in the same manner. When  $\text{TRLX} = 1$ , the number of wait states inserted by the memory controller is doubled from  $\text{OR}_n[\text{SCY}]$  cycles to  $2 \times \text{OR}_n[\text{SCY}]$  cycles, allowing a maximum of 30 wait states.

#### 10.4.2.3.2 Chip-Select and Write Enable Negation Timing

[Figure 10-32](#) shows a basic connection between the local bus and a static memory device. In this case,  $\overline{\text{LCS}}_n$  is connected directly to  $\overline{\text{CE}}$  of the memory device. The  $\overline{\text{LWE}}[0:1]$  signals are connected to the respective  $\overline{\text{WE}}[1:0]$  signals on the memory device where each  $\overline{\text{LWE}}[0:1]$  signal corresponds to a different data byte.



**Figure 10-36. GPCM Basic Write Timing**  
( $\text{XACS} = 0, \text{ACS} = 00, \text{CSNT} = 1, \text{SCY} = 1, \text{TRLX} = 0, \text{CLKDIV} = 4, 8$ )

As [Figure 10-36](#) shows, the timing for  $\overline{\text{LCS}}_n$  is the same as for the latched address. The strobes for the transaction are supplied by  $\overline{\text{LOE}}$  or  $\overline{\text{LWE}}_n$ , depending on the transaction direction—read or write (write case shown in the figure).  $\text{OR}_n[\text{CSNT}]$ , along with  $\text{OR}_n[\text{TRLX}]$ , control the timing for the appropriate strobe negation in write cycles. When this attribute is asserted, the strobe is negated one quarter of a clock before the normal case provided that  $\text{LCRR}[\text{CLKDIV}] = 4$  or  $8$ . For example, when  $\text{ACS} = 00$  and

CSNT = 1,  $\overline{LWE}n$  is negated one quarter of a clock earlier, as shown in Figure 10-36. If LCRR[CLDIV] = 2,  $\overline{LWE}n$  is negated either coincident with  $\overline{LCS}n$  or one cycle earlier.

1.  $\overline{LCS}n$  is affected by CSNT and TRLX only if ACS[0] is non zero. However,  $\overline{LWE}n$  is affected independent of ACS.
2. When CSNT attribute is asserted, the strobe is negated one quarter of a clock before the normal case provided that LCRR[CLDIV] = 4 or 8.
3. TRLX = 1 in conjunction with CSNT = 1, negates the  $\overline{LCS}n$  and  $\overline{LWE}n$  1+1/4 cycle earlier if LCRR[CLKDIV] = 4 or 8.

If LCRR[CLKDIV] = 2,  $\overline{LCS}n$  and  $\overline{LWE}n$  are negated either normally or one cycle earlier if TRLX = 1.

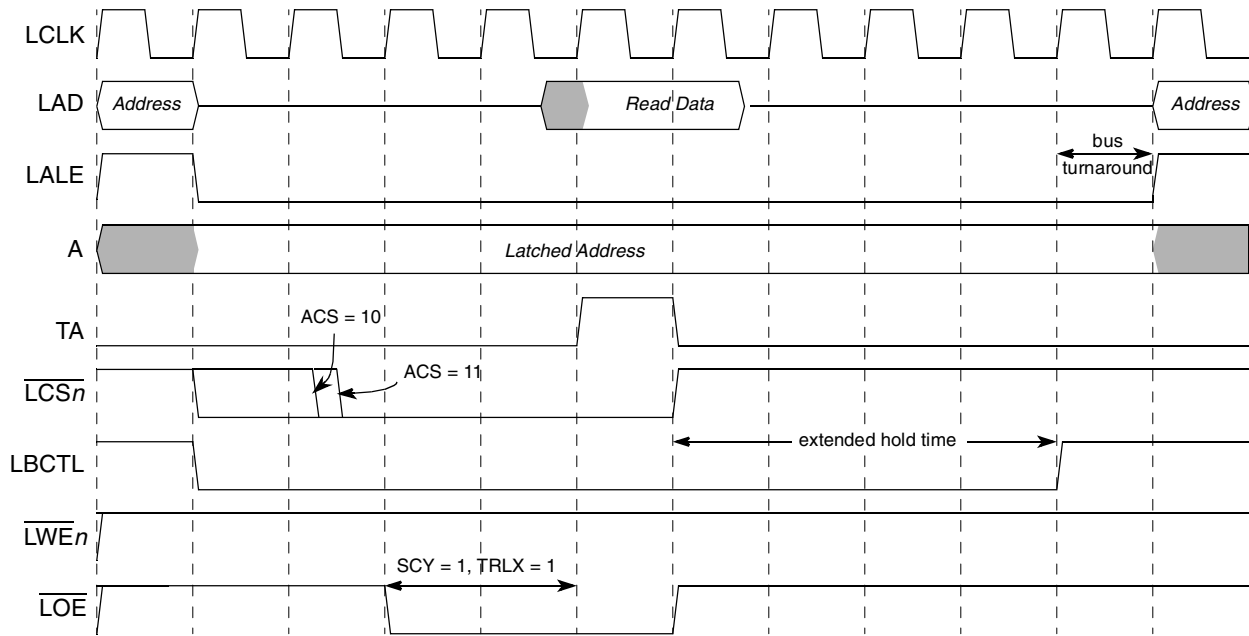
For example, when ACS = 00, CSNT = 1 and TRLX = 0,  $\overline{LWE}n$  is negated one quarter of a clock earlier and  $\overline{LCS}n$  is negated normally as shown in Figure 10-36.

### 10.4.2.3.3 Relaxed Timing

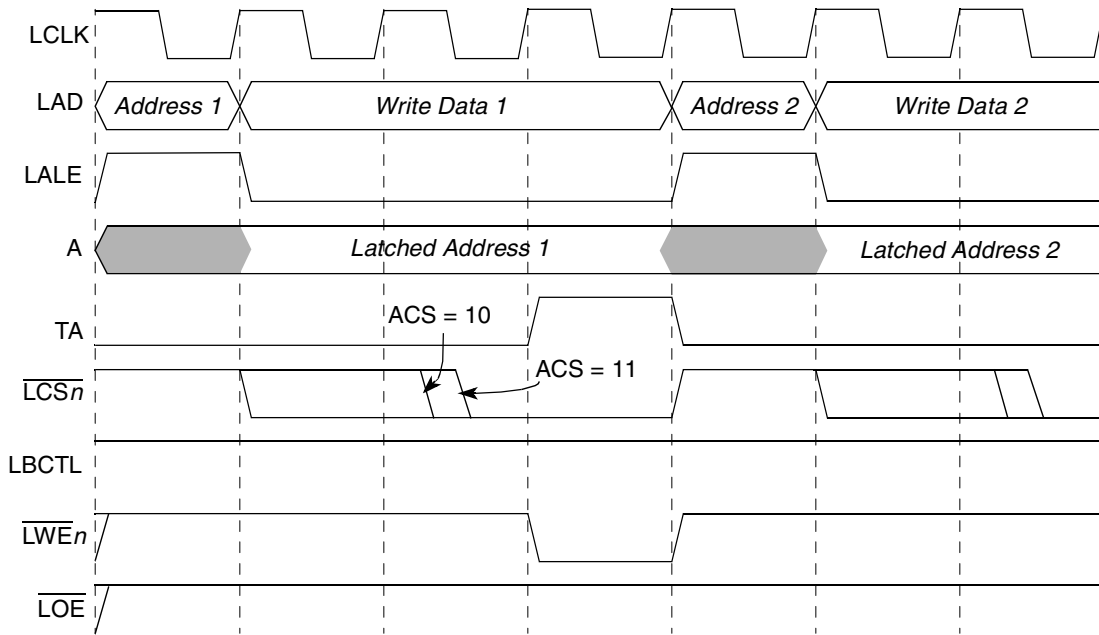
ORx[TRLX] is provided for memory systems that require more relaxed timing between signals. Setting TRLX = 1 has the following effect on timing:

- An additional bus cycle is added between the address and control signals (but only if ACS is not equal to 00).
- The number of wait states specified by SCY is doubled, providing up to 30 wait states.
- The extended hold time on read accesses (EHTR) is extended further.
- $\overline{LCS}n$  signals are negated one cycle earlier during writes (but only if ACS is not equal to 00).
- $\overline{LWE}[0:1]$  signals are negated one cycle earlier during writes.

Figure 10-37 and Figure 10-38 show relaxed timing read and write transactions. The effect of LCRR[CLKDIV] = 2 for these examples is only to delay the assertion of  $\overline{LCS}n$  in the ACS = 10 case to the ACS = 11 case. The example in Figure 10-38 also shows address and data multiplexing on LAD for a pair of writes issued consecutively.




**Figure 10-37. GPCM Relaxed Timing Back-to-Back Reads**  
 (XACS = 0, ACS = 1x, SCY = 1, CSNT = 0, TRLX = 1, EHTR = 0, CLKDIV = 4, 8)




**Figure 10-38. GPCM Relaxed Timing Back-to-Back Writes**  
 (XACS = 0, ACS = 1x, SCY = 0, CSNT = 0, TRLX = 1, CLKDIV = 4, 8)

When  $\overline{\text{TRLX}}$  and  $\overline{\text{CSNT}}$  are set in a write access, the  $\overline{\text{LWE}}[0:1]$  strobe signals are negated one clock earlier than in the normal case, as shown in Figure 10-39 and Figure 10-40. If  $\overline{\text{ACS}} \neq 00$ ,  $\overline{\text{LCSn}}$  is also negated one clock earlier.



**Figure 10-39. GPCM Relaxed Timing Write**  
(XACS = 0, ACS = 10, SCY = 0, CSNT = 1,  $\overline{\text{TRLX}}$  = 1, CLKDIV = 4, 8)



**Figure 10-40. GPCM Relaxed Timing Write**  
(XACS = 0, ACS = 00, SCY = 1, CSNT = 1,  $\overline{\text{TRLX}}$  = 1, CLKDIV = 4, 8)

#### 10.4.2.3.4 Output Enable ( $\overline{\text{LOE}}$ ) Timing

The timing of the  $\overline{\text{LOE}}$  is affected only by  $\text{TRLX}$ . It always asserts and negates on the rising edge of the bus clock.  $\overline{\text{LOE}}$  asserts either on the rising edge of the bus clock after  $\overline{\text{LCSn}}$  is asserted or coinciding with  $\overline{\text{LCSn}}$  (if  $\text{XACS} = 1$  and  $\text{ACS} = 10$  or  $\text{ACS} = 11$ ). Accordingly, assertion of  $\overline{\text{LOE}}$  can be delayed (along with the assertion of  $\overline{\text{LCSn}}$ ) by programming  $\text{TRLX} = 1$ .  $\overline{\text{LOE}}$  negates on the rising clock edge coinciding with  $\overline{\text{LCSn}}$  negation

#### 10.4.2.3.5 Extended Hold Time on Read Accesses

Slow memory devices that take a long time to disable their data bus drivers on read accesses should choose some combination of  $\text{OR}_n[\text{TRLX}, \text{EHTR}]$ . Any access following a read access to the slower memory bank is delayed by the number of clock cycles specified in Table 10-7 in addition to any existing bus turnaround cycle. The final bus turnaround cycle is automatically inserted by the eLBC for reads, regardless of the setting of  $\text{OR}_n[\text{EHTR}]$ .




Figure 10-41. GPCM Read Followed by Read ( $\text{TRLX} = 0$ ,  $\text{EHTR} = 0$ , Fastest Timing)

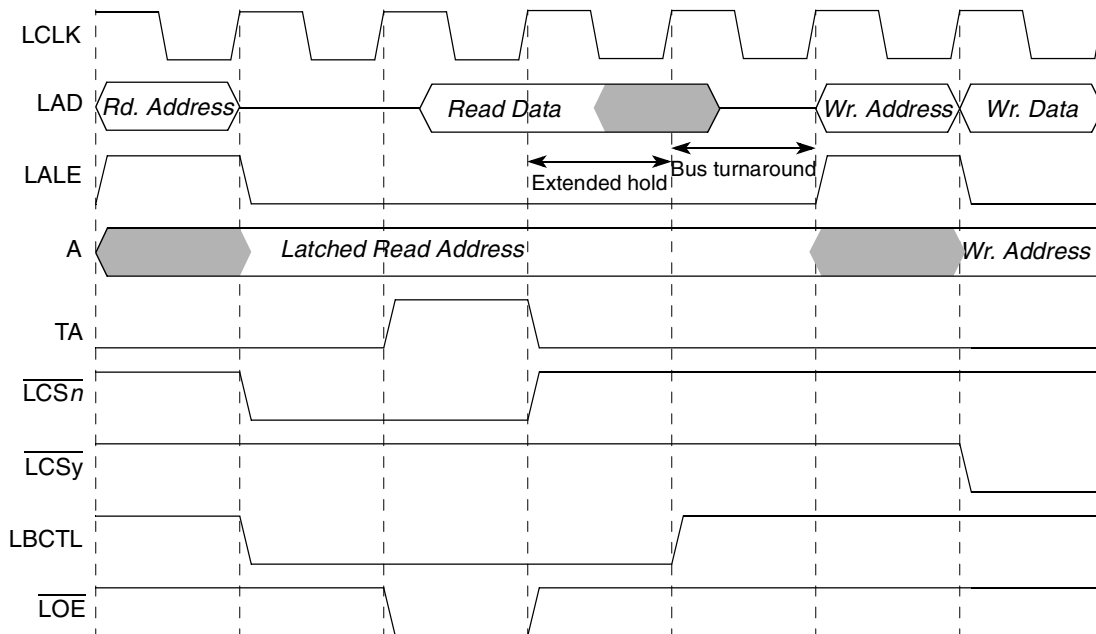


Figure 10-42. GPCM Read Followed by Write  
(TRLX = 0, EHTR = 1, One-Cycle Extended Hold Time on Reads)

#### 10.4.2.4 External Access Termination ( $\overline{LGTA}$ )

External access termination is supported by the GPCM using the asynchronous  $\overline{LGTA}$  input signal, which is synchronized and sampled internally by the local bus. If, during assertion of  $\overline{LCSn}$ , the sampled  $\overline{LGTA}$  signal is asserted, it is converted to an internal generation of transfer acknowledge, which terminates the current GPCM access (regardless of the setting of  $OR_n[SETA]$ ).  $\overline{LGTA}$  should be asserted for at least one bus cycle to be effective. Note that because  $\overline{LGTA}$  is synchronized, bus termination occurs two cycles after  $\overline{LGTA}$  assertion, so in case of read cycle, the device still must drive data as long as  $\overline{LOE}$  is asserted.

The user selects whether transfer acknowledge is generated internally or externally ( $\overline{LGTA}$ ) by programming  $OR_n[SETA]$ . Asserting  $\overline{LGTA}$  always terminates an access, even if  $OR_n[SETA] = 0$  (internal transfer acknowledge generation), but it is the only means by which an access can be terminated if  $OR_n[SETA] = 1$ . The timing of  $\overline{LGTA}$  is illustrated by the example in Figure 10-43.




Figure 10-43. External Termination of GPCM Access

### 10.4.2.5 GPCM Boot Chip-Select Operation

Boot chip-select operation allows address decoding for a boot ROM before system initialization.  $\overline{\text{LCS0}}$  is the boot chip-select output; its operation differs from other external chip-select outputs after a system reset. When the core begins accessing memory after system reset,  $\overline{\text{LCS0}}$  is asserted for every local bus access until BR0 or OR0 is reconfigured.

The boot chip-select also provides a programmable port size, which is configured during reset. The boot chip-select does not provide write protection.  $\overline{\text{LCS0}}$  operates this way until the first write to OR0 and it can be used as any other chip-select register after the preferred address range is loaded into BR0. After the first write to OR0, the boot chip-select can be restarted only with a hardware reset. Table 10-34 describes the initial values of the boot bank in the memory controller.

Table 10-34. Boot Bank Field Values after Reset for GPCM as Boot Controller

Register	Field	Setting
BR0	BA	0000_0000_0000_0000_0
	PS	From RCWH[ROMLOC]
	DECC	00
	WP	0
	MSEL	000
	ATOM	00
	V	1



Table 10-34. Boot Bank Field Values after Reset for GPCM as Boot Controller (continued)

Register	Field	Setting
OR0	AM	0000_0000_0000_0000_0
	BCTLD	0
	CSNT	1
	ACS	11
	XACS	1
	SCY	1111
	SETA	0
	TRLX	1
	EHTR	1
	EAD	1

### 10.4.3 Flash Control Machine (FCM)

The FCM provides a glueless interface to parallel-bus NAND Flash EEPROM devices. The FCM contains three basic configuration register groups— $BR_n$ ,  $OR_n$ , and FMR.

Figure 10-44 shows a simple connection between an 8-bit port size NAND Flash EEPROM and the eLBC in FCM mode. Commands, address bytes, and data are all transferred on  $LAD[0:7]$ <sup>1</sup>, with  $\overline{LFW\overline{E}}$  asserted for transfers written to the device, or  $\overline{LFR\overline{E}}$  asserted for transfers read from the device. eLBC signals LFCLE and LFALE determine whether writes are of type command (only LFCLE asserted), address (only LFALE asserted), or write data (neither LFCLE nor LFALE asserted). The NAND Flash RDY/ $\overline{BSY}$  pin is normally open-drain, and should be pulled high by a 4.7-K $\Omega$  resistor. On system reset, a global (boot) chip-select is available that provides a boot ROM chip-select ( $\overline{LCS0}$ ) prior to the system being fully configured.

1. Note bit numbering reversal: LAD[0] (msb) connects to Flash IO[7], while LAD[7] (lsb) connects to IO[0].

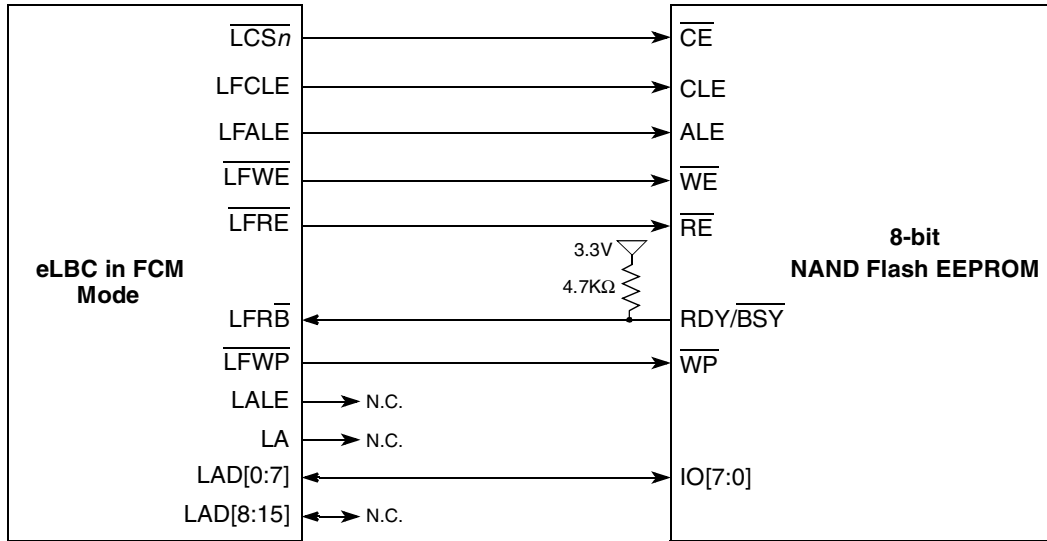


Figure 10-44. Local Bus to 8-bit FCM Device Interface

Basic read access timing for FCM is shown in Figure 10-45. Although LCLK is shown for reference, NAND Flash EEPROMs do not make use of the clock.




Figure 10-45. FCM Basic Page Read Timing  
(PGS = 1, CSCT = 0, CST = 0, CHT = 1, RST = 1, SCY = 0, TRLX = 0, EHTR = 1)

Following the assertion of LALE, FCM asserts  $\overline{\text{LCSn}}$  to commence a command sequence to the Flash device. After a delay of  $t_{\text{CSCT}}$ , the first command can be written to the device on assertion of  $\overline{\text{LFW\overline{E}}}$ , followed by any parameters (typically address bytes and data), and concluded with a secondary command. In many cases, the second command initiates a long-running operation inside the Flash device, which pulls the wired-OR pin  $\text{LFR}\overline{\text{B}}$  low to indicate that the device is busy. Since in [Figure 10-45](#) FCM is now expecting a read response, it takes LBCTL low to turnaround any bus transceivers that are present. Upon  $\text{LFR}\overline{\text{B}}$  indicating ready status, FCM asserts  $\overline{\text{LFR\overline{E}}}$  repeatedly to recover bytes of read data, and the bytes are stored in eLBC's FCM buffer RAM while an ECC is optionally computed on the bytes transferred. Finally, FCM negates  $\overline{\text{LCSn}}$  and delays eLBC by  $t_{\text{EHTR}}$  before any subsequent memory access occurs.

### 10.4.3.1 FCM Buffer RAM

Read and write accesses to eLBC banks controlled by FCM do not access attached NAND Flash EEPROMs directly. Rather, these accesses read and write the FCM buffer RAM—a single, shared 8-Kbyte space internal to eLBC and mapped by the base address of every FCM bank. Even though each FCM-controlled bank will have a different base address to differentiate it, all accesses to such banks will access the same buffer space. External eLBC signals, such as LALE and  $\overline{\text{LCSn}}$ , will not assert upon accesses to the buffer RAM. The FCM buffer RAM is logically divided into two or more buffers, depending on the setting of  $\text{ORn}[\text{PGS}]$ , with different buffers being accessible concurrently by software and FCM.

To perform a page read operation from a NAND Flash device, software initializes the FCM command, mode, and address registers, before issuing a special operation (FMR[OP] set non-zero) to a particular FCM-controlled bank. FCM will execute the sequence of op-codes held in FIR, reading data from the Flash device into the shared buffer RAM. While this read is taking place, software is free to access any data stored in other, currently inactive buffers of the FCM buffer RAM through reads or writes to any bank controlled by FCM. If command completion interrupts are enabled, an interrupt will be generated once FCM has completed the read. When FCM has completed its last command, software can switch to the newly read buffer and issue further commands.

To perform a page write operation, software first prepares data to be written in a fresh buffer. Then, the FCM command, mode, and address registers are initialized, and a special operation (FMR[OP] set non-zero) is issued to a particular FCM-controlled bank. FCM will execute the sequence of op-codes held in FIR, writing data from shared buffer RAM to the Flash device. To ensure that the device is enabled for programming, software must initialize  $\text{FMR}[\text{OP}] = 11$ , which prevents assertion of  $\overline{\text{LFW\overline{P}}}$  during the write. While this write is taking place, software is free to access any data stored in other, currently inactive buffers of the FCM buffer RAM through reads or writes to any bank controlled by FCM. When FCM has completed its last command, software can re-use the previously written buffer and issue further commands.

See [Section 10.4.3.4.2, “Boot Block Loading into the FCM Buffer RAM,”](#) for a description of the shared buffer RAM layout during boot.

#### 10.4.3.1.1 Buffer Layout and Page Mapping for Small-Page NAND Flash Devices

The FCM buffer space is divided into eight 1-Kbyte buffers for small-page devices ( $\text{ORn}[\text{PGS}] = 0$ ), mapped as shown in [Figure 10-46](#). Each page in a small-page NAND Flash comprises 528 bytes, where

512 bytes appear as main region data, and 16 bytes appear as spare region data. The EEPROM's page numbered  $P$  is associated with buffer number  $(P \bmod 8)$ , where  $P = \text{FPAR}[\text{PI}]$ . Since the bank size set by  $\text{ORn}[\text{AM}]$  will be greater than 8 Kbytes, an identical image of the FCM buffer RAM appears replicated every 8 Kbytes throughout the bank address space. It is recommended that the bank size be set to 32 Kbytes, which covers a single NAND Flash block for small-page devices.

For FCM commands, register FPAR sets the page address and, therefore, also the buffer number. In the case that  $\text{FBCR}[\text{BC}] = 0$ , FCM transfers an entire page, comprising the 512-byte main region followed by the 16-byte spare region; the 496-byte reserved region is not accessed, and remains undefined for software. However, for commands given a specific byte count in  $\text{FBCR}[\text{BC}]$ ,  $\text{FPAR}[\text{MS}]$  locates the starting address in either the main region ( $\text{MS} = 0$ ) or the spare region ( $\text{MS} = 1$ ). Where different eLBC banks control both small and large-page devices, a large-page 4-Kbyte buffer must be assigned to either the first 4 or last 4 small-page buffers.




Figure 10-46. FCM Buffer RAM Memory Map for Small-Page (512-byte page) NAND Flash Devices

#### 10.4.3.1.2 Buffer Layout and Page Mapping for Large-Page NAND Flash Devices

The FCM buffer space is divided into two 4 Kbyte buffers for large-page devices ( $\text{ORn}[\text{PGS}] = 1$ ), mapped as shown in Figure 10-47. Each page in a large-page NAND Flash comprises 2112 bytes, where 2048 bytes appear as main region data, and 64 bytes appear as spare region data. The EEPROM's page numbered  $P$  is associated with buffer number  $(P \bmod 2)$ , where  $P = \text{FPAR}[\text{PI}]$ . Since the bank size set by  $\text{ORn}[\text{AM}]$  will be greater than 8 Kbytes, an identical image of the FCM buffer RAM appears replicated every 8 Kbytes throughout the bank address space. It is recommended that the bank size be set to 256 Kbytes, which covers a single NAND Flash block for large-page devices.

For FCM commands, register FPAR sets the page address and, therefore, also the buffer number. In the case that  $\text{FBCR}[\text{BC}] = 0$ , FCM transfers an entire page, comprising the 2048-byte main region followed by the 64-byte spare region; the 1984-byte reserved region is not accessed, and remains undefined for

software. However, for commands given a specific byte count in FBCR[BC], FPAR[MS] locates the starting address in either the main region (MS = 0) or the spare region (MS = 1). Where different eLBC banks control both small and large-page devices, a large-page 4 Kbyte buffer must be assigned to either the first 4 or last 4 small-page buffers.




Figure 10-47. FCM Buffer RAM Memory Map for Large-Page (2-Kbyte page) NAND Flash Devices

### 10.4.3.1.3 Error Correcting Codes and the Spare Region

The FCM's ECC engine makes use of data in the NAND Flash spare region to store pre-computed ECC code words. ECC is calculated in a single pass over blocks of 512 bytes of data in the main region. The setting of FMR[ECCM] determines the location of the 24-bit ECC in the spare region.

The basic ECC algorithm is depicted in Figure 10-48. The stream of data bytes is considered to form a matrix having 8 columns (corresponding with the device bus IO[7:0] or IO[15:8]) and 512 rows (corresponding with each byte in the ECC block).




Figure 10-48. FCM ECC Calculation

The placement of ECC code words in relation to FMR[ECCM] is shown in Figure 10-49. For small-page devices, only a single 512-byte main region is ECC-protected. For large-page devices, there are four adjacent main regions, and each has a 16-byte spare region—of which only one is shown in the figure. If eLBC is configured to generate ECC ( $BR_n[DECC] = 10$ ), FCM will substitute on full-page write transfers the three code word bytes in place of the spare region data originally provided at the locations shown in Figure 10-49 and write the same 24-bit ECC code in the appropriate FECC $_n$  register for software reference. Transfers shorter than a full page, however, require software to prepare the appropriate ECC in the spare region. Similarly, FCM can check and correct bit errors on full-page reads if  $BR_n[DECC] = 01$  or 10. A correctable error is a single bit error in any 512-byte block of main region data, as judged by comparison of a regenerated ECC with the ECC retrieved from the spare region, or a single bit error in the retrieved ECC only. Bit errors in the main region are corrected before FCM completes its final read transfer and signals an event in LTESR[CC]. The bit vector in LTECCR[SBCE] can be checked on FCM CC event to find out if any 512-byte block or the corresponding ECC have single bit correctable errors. Errors that appear more complex (two or more bits in error per 512-byte block) are not corrected, but are flagged as parity errors by FCM. The bit vector in LTEATR[PB] or LTECCR[MBUE] can be checked to determine which 512-byte blocks in a large-page NAND Flash main region were found to be uncorrectable.

ECCM	Byte 0	Byte 511	Other Mains	Spare 0	5	6	7	8	9	10	11	12	13	14	15
0	Main Region			—	EC0	EC1	EC2						—		
1	Main Region			—				EC0	EC1	EC2			—		

Figure 10-49. ECC Placement in NAND Flash Spare Regions in Relation to FMR[ECCM]

### 10.4.3.2 Programming FCM

FCM has a fully general command and data transfer sequencer that caters for both common and specific/proprietary NAND Flash command sequences. The command sequencer reads a program out of the FIR register, which can hold up to 8 instructions, each represented by a 4-bit op-code, as illustrated in Figure 10-50. The first instruction executed is read from FIR[OP0], the next is read from FIR[OP1], and likewise to subsequent instructions, ending at FIR[OP7] or until the only instructions remaining are NOPs. If FIR contains nothing but NOP instructions, FCM will not assert  $\overline{LCS}_n$ , otherwise,  $\overline{LCS}_n$  is asserted prior to the first instruction and remains asserted until the last instruction has completed. If LTESR[CC] is enabled, completion of the last instruction will trigger a command completion event interrupt from eLBC.

Prior to executing a sequence, necessary operands for the instructions will need to be set in the FMR, FCR, MDR, FBCR, FBAR, and FPAR registers. The AS0–AS3 address and data pointers associated with FCM's use of MDR all reset to select AS0 at the start of the instruction sequence. A complete list of op-codes can be found in Section 10.3.1.18, "Flash Instruction Register (FIR)."




Figure 10-50. FCM Instruction Sequencer Mechanism

#### 10.4.3.2.1 FCM Command Instructions

There are two kinds of command instruction:

- Commands that issue immediately—CM0, CM1, CM2, and CM3. These commands write a single command byte by asserting LFCLE and  $\overline{\text{LFWE}}$  while driving an 8-bit command onto LAD[0:7]. Op-code CM $n$  sources its command byte from field FCR[CMD $n$ ], therefore up to four different commands can be issued in any FCM instruction sequence.
- Commands that wait for LFRB to be sampled high (EEPROM in ready state) before issuing—CW0, and CW1. These commands first poll the LFRB pin, waiting for it to go high, before writing a single command byte onto LAD[0:7], sourced from FCR[CMD $n$ ] for op-code CW $n$ . It is necessary to use CW $n$  op-codes whenever the EEPROM is expected to be in a busy state (such as following a page read, block erase, or program operation) and therefore initially unresponsive to commands. To avoid deadlock in cases where the device is already available, FCM does not expect a transition on LFRB. Rather, FCM waits for  $8 \times (2 + \text{OR}_n[\text{SCY}])$  clock cycles (when  $\text{OR}_n[\text{TRLX}] = 0$ ) or  $16 \times (2 + \text{OR}_n[\text{SCY}])$  clock cycles (when  $\text{OR}_n[\text{TRLX}] = 1$ ) before sampling the level of LFRB. If the level of LFRB does not return high before a time-out set by FMR[CWTO] occurs, FCM proceeds to issue the command normally, and a FCT event is issued to LTESR.

The manufacturer's datasheet should be consulted to determine values for programming into the FCR register, and whether a given command in the sequence is expected to initiate busy device behavior.

### 10.4.3.2.2 FCM No-Operation Instruction

A NOP instruction that appears in FIR ahead of the last instruction is executed with the timing of a regular command instruction, but neither LFCLE nor  $\overline{\text{LFW}}\overline{\text{E}}$  are asserted. Thus a NOP instruction may be used to insert a pause matching the time taken for a regular command write.

### 10.4.3.2.3 FCM Address Instructions

Address instructions are used to issue addresses to the NAND Flash EEPROM. A complete device address is formed from a sequence of one or more bytes, each written onto LAD[0:7] with LFALE and  $\overline{\text{LFW}}\overline{\text{E}}$  asserted together. There are three kinds of address generation provided:

- Column address—CA. A column address comprises one byte ( $\text{OR}_n[\text{PGS}] = 0$ ) or two bytes ( $\text{OR}_n[\text{PGS}] = 1$ ) locating the starting byte or word to be transferred in the next page read or write sequence. FPAR[CI] sets the value of the column index provided that FBCR[BC] is non-zero. In the case that FBCR[BC] = 0, a column index of zero is issued to the device, regardless of the value in FPAR[CI].
- Page address—PA. A page address comprises 2, 3, or 4 bytes, depending on the setting of FMR[AL], and locates the data page in the NAND Flash address space. The complete page address is the concatenation of the block index, read from FBAR[BLK], with the page-in-block index, read from FPAR[PI]. The page address length set in FMR[AL] should correspond with the size of EEPROM being accessed. Similarly, the block index in FBAR[BLK] must not exceed the maximum block index for the device, as most devices require reserved address bits to be written as zero.
- User-defined address—UA. This instruction allows the FCM to write a user-defined address byte, which is read from the next AS field in MDR, starting at MDR[AS0]. Each subsequent UA instruction reads an adjacent AS field in MDR, until all four AS bytes (MDR[AS0], MDR[AS1], MDR[AS2], MDR[AS3]) have been sent; a fifth and any following UA instructions send zero as the address byte. Note that each UA instruction advances the MDR pointer for writes by one byte, and therefore a mix of UA and WS instructions can consume adjacent bytes from MDR.

### 10.4.3.2.4 FCM Data Read Instructions

Data read instructions assert  $\overline{\text{LFR}}\overline{\text{E}}$  repeatedly to transfer one or more bytes of read data from the NAND Flash EEPROM. Data read instructions are distinguished by their data destination:

- Read data to buffer RAM immediately—RB. This instruction reads FBCR[BC] bytes of data into the current FCM RAM buffer addressed by FPAR. If FBCR[BC] = 0, an entire page (including spare region) is transferred in a burst, starting at the page boundary, and the ECC calculation is checked against the ECC stored in the spare region. Correctable ECC errors are corrected and reported in LTECCR[SBCE]; other errors may cause an interrupt if enabled. If the value of FBCR[BC] takes the read pointer beyond the end of the spare region in the buffer, FCM discards any excess bytes read.
- Read data/status to MDR immediately—RS. This instruction asserts  $\overline{\text{LFR}}\overline{\text{E}}$  exactly once to read one byte (8-bit port size) of data into the next AS field of MDR. Reads beyond the fourth byte of MDR are discarded. The MDR read pointer is independent of the MDR write pointer used by UA and WS instructions.



- Read data to buffer RAM once waited on ready—RBW. This instruction first polls the  $\overline{\text{LFRB}}$  pin, waiting for it to go high, before proceeding with a read to buffer as described for the RB instruction. Sampling and time-outs for polling the  $\overline{\text{LFRB}}$  pin follow the behavior of  $\text{CW}_n$  instructions.
- Read data/status to MDR once waited on ready—RSW. This instruction first polls the  $\overline{\text{LFRB}}$  pin, waiting for it to go high, before proceeding with a status read to MDR as described for the RS instruction. Sampling and time-outs for polling the  $\overline{\text{LFRB}}$  pin follow the behavior of  $\text{CW}_n$  instructions.

#### 10.4.3.2.5 FCM Data Write Instructions

Data write instructions assert  $\overline{\text{LFW}}_E$  repeatedly (with  $\overline{\text{LFCLE}}$  and  $\overline{\text{LFALE}}$  both negated) to transfer one or more bytes of write data to the NAND Flash EEPROM. Data write instructions are distinguished by their data source:

- Write data from FCM buffer RAM—WB. This instruction writes  $\text{FBCR}[\text{BC}]$  bytes of data from the current FCM RAM buffer addressed by  $\text{FPAR}$ . If  $\text{FBCR}[\text{BC}] = 0$ , an entire page (including spare region) is transferred in a burst, starting at the page boundary, and the ECC calculation is stored in the appropriate  $\text{FECC}_n$  registers and spare region in accordance with the setting of  $\text{FMR}[\text{ECCM}]$ . If the value of  $\text{FBCR}[\text{BC}]$  takes the write pointer beyond the end of the spare region in the buffer, the value of data written by FCM is undefined.
- Write data/status from MDR—WS. This instruction asserts  $\overline{\text{LFW}}_E$  exactly once to write one byte (8-bit port size) of data taken from the next AS field of MDR. Attempts to write beyond four bytes of MDR has the effect of writing zeros. The MDR write pointer is independent of the MDR read pointer used by RS and RSW instructions.

#### 10.4.3.3 FCM Signal Timing

If  $\text{BR}_n[\text{MSEL}]$  selects the FCM, the attributes for the memory cycle are taken from  $\text{OR}_n$ . These attributes include the CSCT, CST, CHT, RST, SCY, TRLX, and EHTR fields.

##### 10.4.3.3.1 FCM Chip-Select Timing

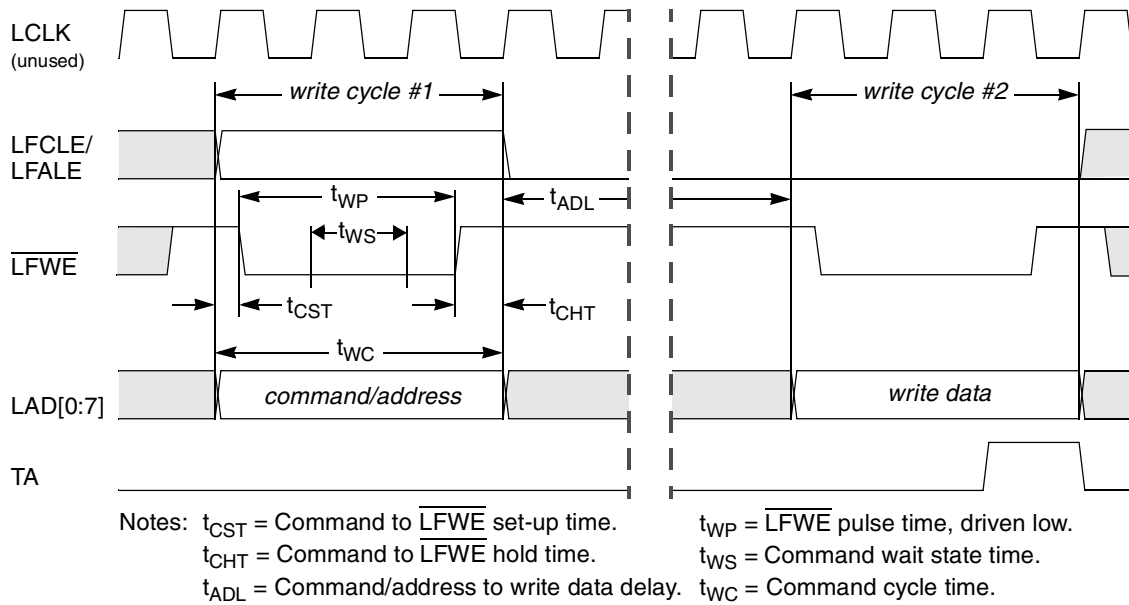
The timing of  $\overline{\text{LCS}}_n$  assertion in FCM mode is illustrated by the timing diagram in [Figure 10-45](#).  $\overline{\text{LCS}}_n$  is asserted immediately following  $\overline{\text{LALE}}$  negation, and remains asserted until the last instruction in FIR has completed. The delay,  $t_{\text{CSCT}}$ , between  $\overline{\text{LCS}}_n$  assertion and commencement of the first NAND Flash instruction is controlled by  $\text{OR}_n[\text{CSCT}]$  and  $\text{OR}_n[\text{TRLX}]$ , as shown in [Table 10-35](#).  $\text{OR}_n[\text{CSCT}]$  should be set in accordance with the NAND Flash EEPROM chip-select to  $\overline{\text{WE}}$  set-up time specification.

**Table 10-35. FCM Chip-Select to First Command Timing**

$\text{OR}_n[\text{TRLX}]$	$\text{OR}_n[\text{CSCT}]$	$\overline{\text{LCS}}_n$ to First Command Delay
0	0	1 LCLK clock cycle
0	1	4 LCLK clock cycles
1	0	2 LCLK clock cycles
1	1	8 LCLK clock cycles

### 10.4.3.3.2 FCM Command, Address, and Write Data Timing

The FCM command (CM0–CM3, CW0, CW1), address (CA, PA, UA), and data write (WB, WS) instructions all share the same basic timing attributes. Assertion of  $\overline{\text{LFWE}}$  initiates transfer via LAD[0:7], and the options in OR $n$  for FCM mode establish the set-up, hold, and wait state timings with respect to  $\overline{\text{LFWE}}$ , as shown in Figure 10-51.



**Figure 10-51. Timing of FCM Command/Address and Write Data Cycles (for TRLX = 0, CHT = 0, CST = 1, SCY = 1, CLKDIV = 4\*N)**

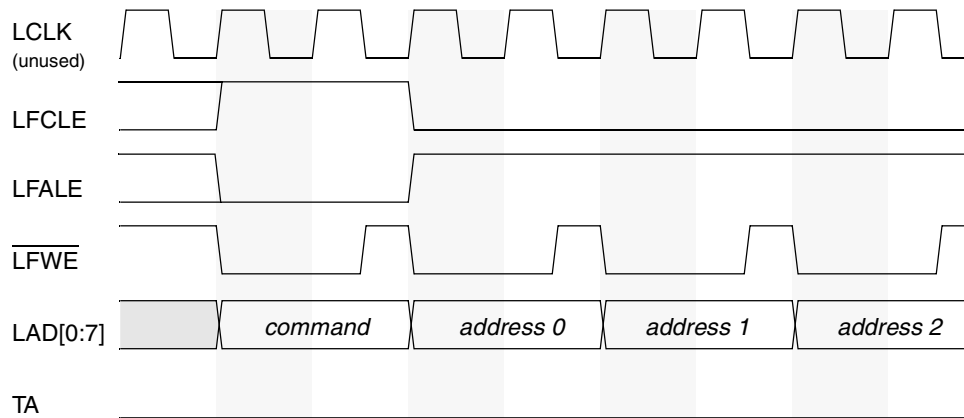
The timing parameters are summarized in Table 10-36.

**Table 10-36. FCM Command, Address, and Write Data Timing Parameters**

Option Register Attributes			Timing Parameter (LCLK Clock Cycles) <sup>1</sup>					
TRLX	CHT	CST	$t_{\text{CST}}$	$t_{\text{CHT}}$	$t_{\text{WS}}$	$t_{\text{WP}}$	$t_{\text{WC}}$	$t_{\text{ADL}}$
0	0	0	0	½	SCY	1½+SCY	2+SCY	4x(2+SCY)
0	0	1	¼	½	SCY	1¼+SCY	2+SCY	4x(2+SCY)
0	1	0	0	1	SCY	1+SCY	2+SCY	4x(2+SCY)
0	1	1	¼	1	SCY	¾+SCY	2+SCY	4x(2+SCY)
1	0	0	½	1½	2xSCY	1+2xSCY	3+2xSCY	8x(2+SCY)
1	0	1	1	1½	2xSCY	½+2xSCY	3+2xSCY	8x(2+SCY)
1	1	0	½	2	2xSCY	½+2xSCY	3+2xSCY	8x(2+SCY)
1	1	1	1	2	2xSCY	2xSCY	3+2xSCY	8x(2+SCY)


<sup>1</sup> In the parameters, SCY refers to a delay of OR $n$ [SCY] clock cycles.

An example of minimum delay command timing appears in Figure 10-52. Note that the set-up, wait-state, and hold timing of command, address, and write data cycles with respect to  $\overline{\text{LFW}}\overline{\text{E}}$  assertion are all identical, and that the minimum cycle extends for two LCLK clock cycles.



**Figure 10-52. Example of FCM Command and Address Timing with Minimum Delay Parameters**  
(for  $\text{TRLX} = 0$ ,  $\text{CHT} = 0$ ,  $\text{CST} = 0$ ,  $\text{SCY} = 0$ ,  $\text{CLKDIV} = 4 \cdot \text{N}$ )

An example of relaxed command timing is shown in Figure 10-53.



**Figure 10-53. Example of FCM Command and Address Timing with Relaxed Parameters**  
(for  $\text{TRLX} = 1$ ,  $\text{CHT} = 0$ ,  $\text{CST} = 1$ ,  $\text{SCY} = 2$ ,  $\text{CLKDIV} = 4 \cdot \text{N}$ )

#### 10.4.3.3.3 FCM Ready/Busy Timing

Instructions CW0, CW1, RBW, and RSW force FCM to observe the state of the  $\text{LFR}\overline{\text{B}}$  pin, which may be driven low by a long-latency NAND Flash operation, such as a page read. Following the issue of such commands, FCM waits as shown in Figure 10-54 before sampling the state of  $\text{LFR}\overline{\text{B}}$ . This guards against observing  $\text{LFR}\overline{\text{B}}$  before it has been properly driven low by the device, but does not preclude  $\text{LFR}\overline{\text{B}}$  from remaining high after a command. In addition, FCM samples and compares the state of  $\text{LFR}\overline{\text{B}}$  on two consecutive cycles of LCLK to filter out noise on this signal as it rises to the ready state ( $\text{LFR}\overline{\text{B}} = 1$ ).

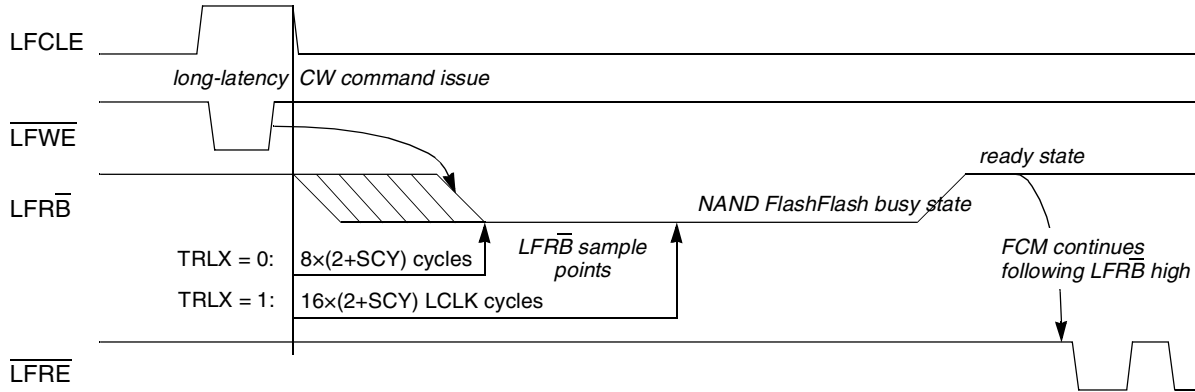
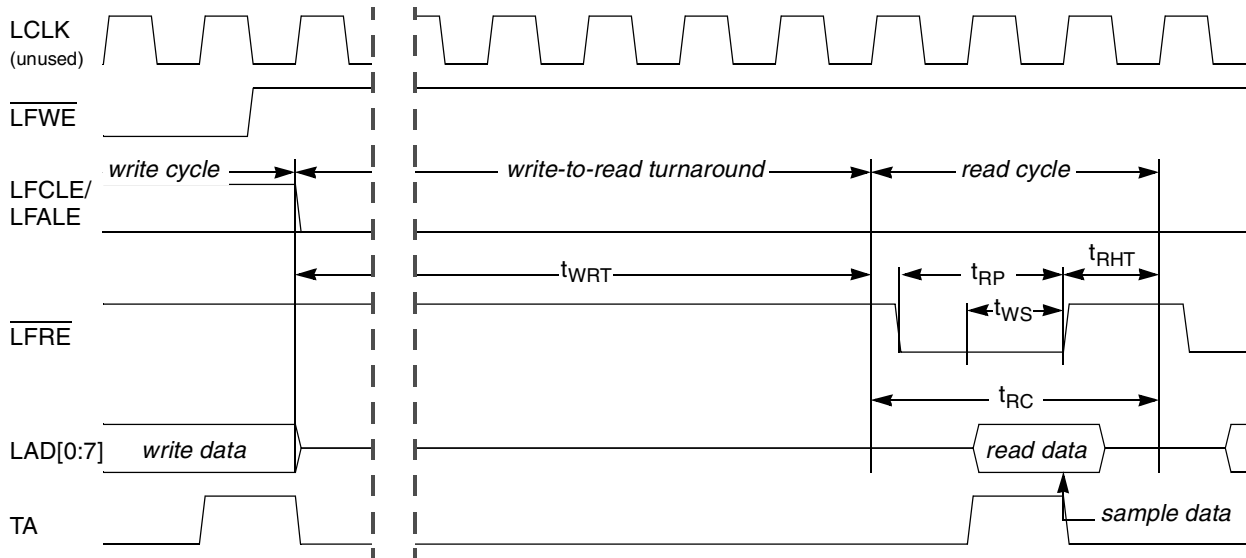


Figure 10-54. FCM Delay Prior to Sampling LFRB State

### 10.4.3.3.4 FCM Read Data Timing

The timing for read data transfers is shown in Figure 10-55. Upon assertion of  $\overline{\text{LFRE}}$ , the Flash device will enable its output drivers and drive valid read data while  $\overline{\text{LFRE}}$  is held low. FCM samples read data on the rising edge of  $\overline{\text{LFRE}}$ , which follows an optional number of wait states. Note that FCM will delay the first read if a RBW or RSW instruction is issued, in which case LFRB sample timing takes effect (see Section 10.4.3.3.3, “FCM Ready/Busy Timing”).



Notes:  $t_{RP}$  =  $\overline{\text{LFRE}}$  pulse time, read period.  $t_{WS}$  = Read wait state time.  
 $t_{RHT}$  =  $\overline{\text{LFRE}}$  hold time.  $t_{RC}$  = Read data cycle time.  
 $t_{WRT}$  = Write to read turnaround time.

Figure 10-55. FCM Read Data Timing  
 (for  $\text{TRLX} = 0, \text{RST} = 0, \text{SCY} = 1, \text{CLKDIV} = 4 \cdot N$ )

The timing parameters are summarized in [Table 10-37](#).

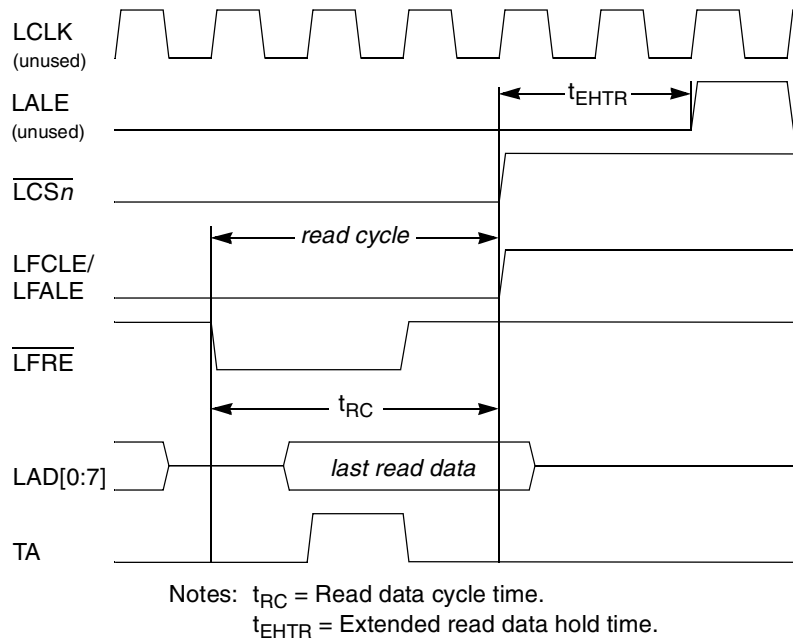
**Table 10-37. FCM Read Data Timing Parameters**

Option Register Attributes		Timing Parameter (LCLK Clock Cycles) <sup>1</sup>				
TRLX	RST	$t_{RP}$	$t_{RHT}$	$t_{WS}$	$t_{RC}$	$t_{WRT}$
0	0	$\frac{3}{4}+SCY$	1	SCY	$2+SCY$	$4 \times (2+SCY)$
0	1	$1+SCY$	1	SCY	$2+SCY$	$4 \times (2+SCY)$
1	0	$\frac{1}{2}+2 \times SCY$	2	$2 \times SCY$	$3+2 \times SCY$	$8 \times (2+SCY)$
1	1	$1+2 \times SCY$	2	$2 \times SCY$	$3+2 \times SCY$	$8 \times (2+SCY)$

<sup>1</sup> In the parameters, SCY refers to a delay of  $OR_n[SCY]$  clock cycles.

#### 10.4.3.3.5 FCM Extended Read Hold Timing

Allowance for slow output driver turn-off when reading NAND Flash EEPROMs is made via setting of  $OR_n[EHTR]$  and  $OR_n[TRLX]$ . The extended read data hold time, shown at  $t_{EHTR}$  in [Figure 10-45](#) and [Figure 10-56](#), is a delay inserted by FCM between the last data read and another eLBC memory access (requiring LALE assertion).  $\overline{LCSn}$  is negated during  $t_{EHTR}$  to allow external devices and bus transceivers time to disable their drivers.



**Figure 10-56. FCM Read Data Timing with Extended Hold Time**  
 (for  $TRLX = 0$ ,  $EHTR = 1$ ,  $RST = 1$ ,  $SCY = 1$ ,  $CLKDIV = 4 \times N$ )

#### 10.4.3.4 FCM Boot Chip-Select Operation

Boot chip-select operation allows address decoding for a boot ROM before system initialization.  $\overline{LCS0}$  is the boot chip-select output; its operation differs from other external chip-select outputs after a system reset.

When the core begins accessing memory after system reset,  $\overline{LCS0}$  is asserted initially to load a 4-Kbyte boot block into the FCM buffer RAM, but core instruction fetches occur from the buffer RAM.

#### 10.4.3.4.1 FCM Bank 0 Reset Initialization

The boot chip-select also provides a programmable port size, which is configured during reset. The boot chip-select does not provide write protection.  $\overline{LCS0}$  operates this way until the first write to OR0 and it can be used as any other chip-select register after the preferred address range is loaded into BR0. After the first write to OR0, the boot chip-select can be restarted only with a hardware reset. [Table 10-38](#) describes the initial values of the boot bank in the memory controller.


**Table 10-38. Boot Bank Field Values after Reset for FCM as Boot Controller**

Register	Field	Setting
BR0	BA	0000_0000_0000_0000_0
	PS	From 01
	DECC	00
	WP	0
	MSEL	001
	ATOM	00
	V	0
OR0	AM	0000_0000_0000_0000_0
	BCTLD	0
	PGS	From RCWH[ROMLOC]
	CSCT	1
	CST	1
	CHT	1
	RST	1
	SCY	From por_cfg_scy[1:3]
	TRLX	1
	EHTR	1

#### 10.4.3.4.2 Boot Block Loading into the FCM Buffer RAM

If FCM is selected as the boot ROM controller from power-on-reset configuration, eLBC will automatically load from bank 0 a single 4-Kbyte page of boot code into the FCM buffer RAM during  $\overline{HRESET}$  (see [Section 4.3.2.2.4, “Boot ROM Location”](#)). The CPU can execute boot code directly from the FCM buffer RAM, but must ensure that any further data read from the NAND Flash EEPROM is transferred under software control in order to continue the bootstrap process.

Since OR0[AM] is initially cleared during reset, all CPU fetches to eLBC will access the FCM buffer RAM, which appears in the memory map as a 4-Kbyte RAM. No NAND Flash spare regions are mapped during boot, therefore only 4 Kbytes of contiguous, main region data, loaded from the first pages of the boot block, are accessible in eLBC bank 0, as indicated in [Figure 10-57](#).



**Figure 10-57. FCM Buffer RAM Memory Map During Boot Loading**

The process for booting is as follows:

1. Following negation of  $\overline{\text{PORESET}}$ , eLBC is released from reset and commences automatic boot block loading if FCM is selected as the boot ROM location. Small-page or large-page, 8-bit NAND Flash devices can be used for boot loading when enabled with  $\overline{\text{LCS0}}$ . eLBC drives  $\overline{\text{LFWP}}$  low during boot accesses to prevent accidental erasure of the NAND Flash boot ROM.
2. FCM starts searching for a valid boot block at block index 0.
3. FCM reads the spare regions of the first two pages of the current block, checking the bad block indication (BI) bytes to validate the block for reading. BI bytes must all hold the value 0xFF for the page to be considered readable.
  - For small-page devices, BI is a single byte read from spare region byte offset 5.
  - For large-page devices, BI is a single byte read from spare region byte offset 0.

If either of the first two pages of the current block are marked invalid, then the boot block index is incremented by 1, and FCM repeats step 3. eLBC will continue searching for a bootable block indefinitely, therefore at least one block must be marked valid for boot loading to proceed. At the conclusion of the boot block search, the value of  $\text{FBAR}[\text{BLK}]$  points to the boot block.

4. If ECC checking is enabled, the FCM recovers from the spare region the stored ECC for each 512-byte block of boot data. The boot block must be prepared with ECC protection. During ECC generation, software should use  $\text{FMR}[\text{ECCM}] = 0$  for small-page devices, and  $\text{FMR}[\text{ECCM}] = 1$  for large-page devices.

If RCW initialization is required, the first 64 bytes of the boot block must be prepared in accordance with the layout described in [Section 4.3.3.1.1, “Local Bus Controller Setting.”](#)

5. FCM performs a sequence of random-access page reads, reading entire pages from the boot block until 4 Kbytes have been saved to the FCM buffer RAM. If ECC checking is enabled, the ECC of each 512-byte region is verified and single-bit errors are corrected if possible. If FCM is unable to correct ECC errors, eLBC halts the boot process and signals an unrecoverable error by asserting the  $\overline{\text{hreset\_req}}$  signal.

- The CPU now commences fetching instructions, in random order, from the FCM buffer RAM. This first-level boot loader typically copies a secondary boot loader into system memory, and continues booting from there. Boot software must clear FMR[BOOT] to enable normal operation of FCM.

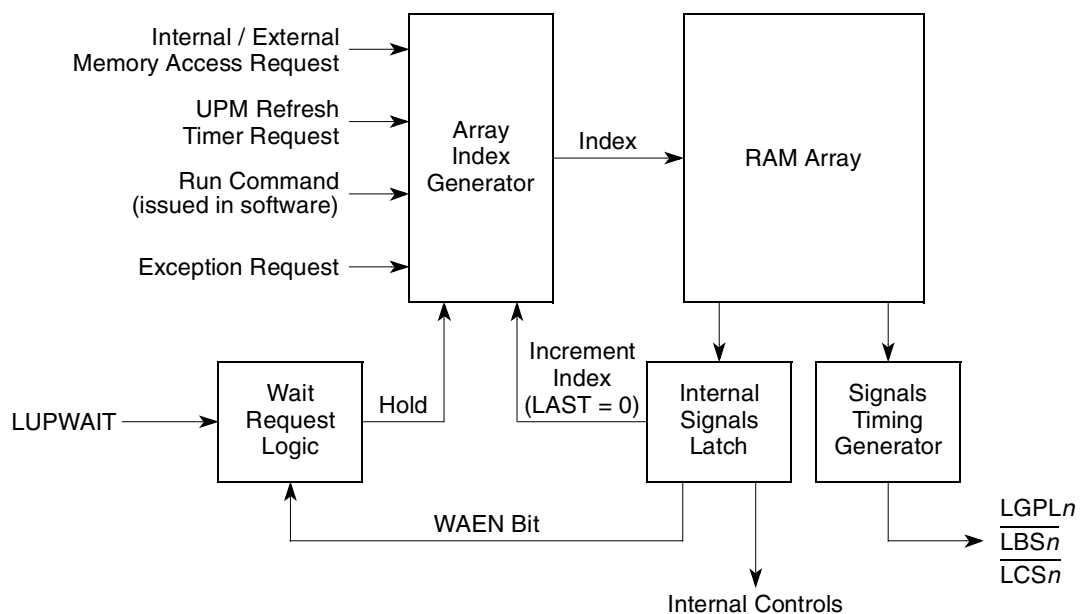
#### 10.4.4 User-Programmable Machines (UPMs)

UPMs are flexible interfaces that connect to a wide range of memory devices. At the heart of each UPM is an internal RAM array that specifies the logical value driven on the external memory control signals ( $\overline{LCSn}$ ,  $\overline{LBS}[0:1]$  and  $\overline{LGPL}[0:5]$ ) for a given clock cycle. Each word in the RAM array provides bits that allow a memory access to be controlled with a resolution of up to one quarter of the external bus clock period on the byte-select and chip-select lines. A gap of 2 dead LCLK cycles is present on the UPM interface between UPM transactions.

#### NOTE

If the  $\overline{LGPL4}/\overline{LGTA}/\overline{LFRB}/\overline{LUPWAIT}$  signal is used as both an input and an output, a weak pull-up is required. Refer to the hardware specification for details regarding termination options.

Figure 10-58 shows the basic operation of each UPM.



**Figure 10-58. User-Programmable Machine Functional Block Diagram**

The following events initiate a UPM cycle:

- Any internal device requests an external memory access to an address space mapped to a chip-select serviced by the UPM
- A UPM refresh timer expires and requests a transaction, such as a DRAM refresh
- A bus monitor time-out error during a normal UPM cycle redirects the UPM to execute an exception sequence



The RAM array contains 64 words of 32-bits each. The signal timing generator loads the RAM word from the RAM array to drive the general-purpose lines, byte-selects, and chip-selects. If the UPM reads a RAM word with WAEN set, the external LUPWAIT signal is sampled and synchronized by the memory controller and the current request is frozen.

#### 10.4.4.1 UPM Requests

A special pattern location in the RAM array is associated with each of the possible UPM requests. An internal device's request for a memory access initiates one of the following patterns ( $MxMR[OP] = 00$ ):

- Read single-beat pattern (RSS)
- Read burst cycle pattern (RBS)
- Write single-beat pattern (WSS)
- Write burst cycle pattern (WBS)

A UPM refresh timer request pattern initiates a refresh timer pattern (RTS).

An exception (caused by a bus monitor time-out error) occurring while another UPM pattern is running initiates an exception condition pattern (EXS).

Figure 10-59 and Table 10-39 show the start addresses of these patterns in the UPM RAM, according to cycle type. RUN commands ( $MxMR[OP] = 11$ ), however, can initiate patterns starting at any of the 64 UPM RAM words.

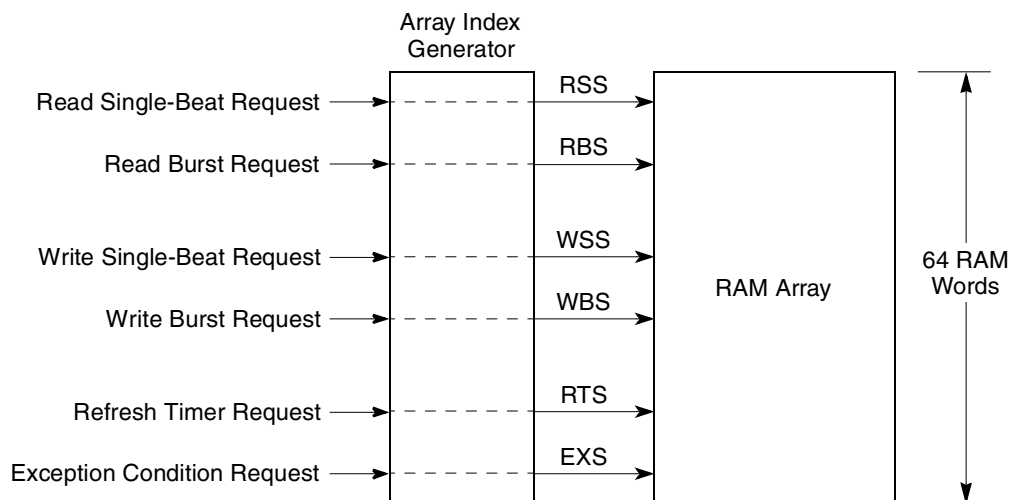


Figure 10-59. RAM Array Indexing

Table 10-39. UPM Routines Start Addresses

UPM Routine	Routine Start Address
Read single-beat (RSS)	0x00
Read burst (RBS)	0x08
Write single-beat (WSS)	0x18
Write burst (WBS)	0x20

**Table 10-39. UPM Routines Start Addresses (continued)**

UPM Routine	Routine Start Address
Refresh timer (RTS)	0x30
Exception condition (EXS)	0x3C

#### 10.4.4.1.1 Memory Access Requests

The user must ensure that the UPM is appropriately initialized before a request occurs.

The UPM supports two types of memory reads and writes:


- A single-beat transfer transfers one operand consisting of up to a single word (dependent on port size). A single-beat cycle starts with one transfer start and ends with one transfer acknowledge.
- A burst transfer transfers exactly 4 double words regardless of port size. For 32-bit accesses, the burst cycle starts with one transfer start but ends after eight transfer acknowledges, whereas an 8-bit device requires 32 transfer acknowledges.

The user must ensure that patterns for single-beat transfers contain one, and only one, transfer acknowledge (UTA bit in RAM word set high) and for a burst transfer, contain the exact number of transfer acknowledges required.

Any transfers that do not naturally fit single or burst transfers are synthesized as a series of single transfers. These accesses are treated by the UPM as back-to-back, single-beat transfers. Burst transfers can also be inhibited by setting  $OR_n[BI]$ . Burst performance can be achieved by ensuring that UPM transactions are 32-byte aligned with a transaction size being some multiple of 32-bytes, which is a natural fit for cache-line transfers, for example.

#### 10.4.4.1.2 UPM Refresh Timer Requests

Each UPM contains a refresh timer that can be programmed to generate refresh service requests of a particular pattern in the RAM array. Figure 10-60 shows the clock division hardware associated with memory refresh timer request generation. The UPM refresh timer register (LURT) defines the period for the timers associated with all three UPMs.

**Figure 10-60. Memory Refresh Timer Request Block Diagram**

By default, all local bus refreshes are performed using the refresh pattern of UPMA. This means that if refresh is required,  $MAMR[RFEN]$  must be set. It also means that only one refresh routine should be programmed and be placed in UPMA, which serves as the refresh executor. Any banks assigned to a UPM are provided with the common UPMA refresh pattern if the  $RFEN$  bit of the corresponding UPM is set, concurrently. UPMA assigned banks, therefore, always receive refresh services when  $MAMR[RFEN]$  is set, while UPMB and UPMC assigned banks also receive (the same) refresh services if the corresponding  $MxMR[RFEN]$  bits are set. In this scenario, more than one chip select may assert at the same time, as refresh pattern runs for all banks assigned to UPM with  $RFEN$  bit set.

### 10.4.4.1.3 Software Requests—RUN Command

Software can start a request to the UPM by issuing a RUN command to the UPM. Some memory devices have their own signal handshaking protocol to put them into special modes, such as self-refresh mode.

For these special cycles, the user creates a special RAM pattern that can be stored in any unused areas in the UPM RAM. Then a RUN command is used to run the cycle. The UPM runs the pattern beginning at the specified RAM location until it encounters a RAM word with its LAST bit set. The RUN command is issued by setting  $M_xMR[OP] = 11$  and accessing UPM $_n$  memory region with any write transaction that hits the corresponding UPM machine.  $M_xMR[MAD]$  determines the starting address in the RAM array for the pattern.

Note that transfer acknowledges (UTA bit in the RAM word) are ignored for software (RUN command) requests, and hence the LAD signals remain high-impedance unless the normal initial LALE occurs or the RUN pattern causes assertion of LALE to occur on changes to the RAM word AMX field.

### 10.4.4.1.4 Exception Requests

When the eLBC under UPM control initiates an access to a memory device and an exception occurs (bus monitor time-out), the UPM provides a mechanism by which memory control signals can meet the device's timing requirements without losing data. The mechanism is the exception pattern that defines how the UPM negates its signals in a controlled manner.

### 10.4.4.2 Programming the UPMs

The UPM is a micro sequencer that requires microinstructions or RAM words to generate signal timings for different memory cycles. Follow these steps to program the UPMs:

1. Set up  $BR_n$  and  $OR_n$  registers.
2. Write patterns into the RAM array.
3. Program MRTPR, LURT and MAMR[RFEN] if refresh is required.
4. Program  $M_xMR$ .

Patterns are written to the RAM array by setting  $M_xMR[OP] = 01$  and accessing the UPM with any write transaction that hits the relevant chip select. The entire array is thus programmed by an alternating series of writes: to MDR (RAM word to be written) each time followed by a read from MDR and then followed by a (dummy) write transaction to the relevant UPM assigned bank. A read from MDR is required to ensure that the MDR update has occurred prior to the (dummy) write transaction.

RAM array contents may also be read for debug purposes, for example, by alternating dummy read transactions, each time followed by reads of MDR (when  $M_xMR[OP] = 10$ ).

#### NOTE

$M_xMR$  / MDR registers should not be updated while dummy read/write access is still in progress. If the  $M_xMR[MAD]$  is incremented then the previous dummy transaction is already completed.

In order to enforce proper ordering between updates to the MxMR/MDR register and the dummy accesses to the UPM memory region, two rules must be followed:

1. Since the result of any update to the MxMR/MDR register must be in effect before the dummy read or write to the UPM region, a write to MxMR/MDR should be followed immediately by a read of MxMR/MDR.
2. The UPM memory region should have the same MMU settings as the memory region containing the MxMR configuration register; both should be mapped by the MMU as cache-inhibited and guarded. This prevents the CPU from re-ordering a read of the UPM memory around the read of MxMR. Once the programming of the UPM array is complete the MMU setting for the associated address range can be set to the proper mode for normal operation, such as cacheable and copyback.

For proper signalling, the following guidelines must be followed while programming UPM RAM words:

- For UPM reads, program UTA and LAST in the same or consecutive RAM words.
- For UPM burst reads, program last UTA and LAST in the same or consecutive RAM words.
- For UPM writes, program UTA and LAST in the same RAM word.
- For UPM burst writes, program last UTA and LAST in the same RAM word.

#### 10.4.4.2.1 UPM Programming Example (Two Sequential Writes to the RAM Array)

The following example further illustrates the steps required to perform two writes to the RAM array at non-sequential addresses assuming that the relevant BR<sub>n</sub> and OR<sub>n</sub> registers have been previously set up:

1. Program MxMR for the first write (with the desired RAM array address).
2. Write pattern/data to MDR to ensure that the MxMR has already been updated with the desired configuration.
3. Read MDR to ensure that the MDR has already been updated with the desired pattern. (Or, read MxMR register if step 2 is not performed.)
4. Perform a dummy write transaction.
5. Read/check MxMR[MAD]. If incremented, the previous dummy write transaction is completed; proceed to step 6. Repeat step 5 until incremented.
6. Program MxMR for the second write with the desired RAM array address.
7. Write pattern/data to MDR to ensure that the MxMR has already been updated with the desired configuration.
8. Read MDR to ensure that the MDR has already been updated with the desired pattern.
9. Perform a dummy write transaction.
10. Read/check MxMR[MAD]. If incremented, the previous dummy write transaction is completed.

Note that if step 1 (or 6) and 2 (or 7) are reversed, step 3 (or 8) is replaced by the following:

- Read MxMR to ensure that the MxMR has already been updated with the desired configuration.

### 10.4.4.2 UPM Programming Example (Two Sequential Reads from the RAM Array)

RAM array contents may also be read for debug purposes, for example, by alternating dummy read transactions, each time followed by reads of MDR ( $MxMR[OP] = 0b10$ ). The following example further illustrates the steps required to perform two reads from the RAM array at non-sequential addresses assuming that the relevant  $BR_n$  and  $OR_n$  registers have been previously set up:

1. Program  $MxMR$  for the first read with the desired RAM array address.
2. Read  $MxMR$  to ensure that the  $MxMR$  has already been updated with the desired configuration, such as RAM array address.
3. Perform a dummy read transaction.
4. Read/check  $MxMR[MAD]$ . If incremented, the previous dummy read transaction is completed; proceed to step 5. Repeat step 4 until incremented.
5. Read MDR.
6. Program  $MxMR$  for the second read with the desired RAM array address.
7. Read  $MxMR$  to ensure that the  $MxMR$  has already been updated with the desired configuration, such as RAM array address.
8. Perform a dummy read transaction.
9. Read/check  $MxMR[MAD]$ . If incremented, the previous dummy read transaction is completed; proceed to step 10. Repeat step 9 until incremented.
10. Read MDR.

### 10.4.4.3 UPM Signal Timing

RAM word fields specify the value of the various external signals at a granularity of up to four values for each bus clock cycle. The signal timing generator causes external signals to behave according to timing specified in the current RAM word. For  $LCRR[CLKDIV] = 4$  or  $8$ , each bit in the RAM word relating to  $\overline{LCS}_n$  and  $\overline{LBS}$  timing specifies the value of the corresponding external signal at each quarter phase of the bus clock. If  $LCRR[CLKDIV] = 2$ , the external signal can change value only on each half phase of the bus clock. If the RAM word in this case ( $LCRR[CLKDIV] = 2$ ) specifies a quarter phase signal change, the signal timing generator interprets this as a half cycle change.

The division of UPM bus cycles into phases is shown in [Figure 10-61](#) and [Figure 10-62](#). If  $LCRR[CLKDIV] = 2$ , the bus cycle comprises only two active phases, T1 and T3, which correspond with the first and second halves of the bus clock cycle, respectively. However, if  $LCRR[CLKDIV] = 4$  or  $8$ , four phases, T1–T4, define four quarters of the bus clock cycle. Because T2 and T4 are inactive when  $LCRR[CLKDIV] = 2$ , UPM ignores signal timing programmed for assertion in either of these phases in the case  $LCRR[CLKDIV] = 2$ .




Figure 10-61. UPM Clock Scheme for LCRR[CLKDIV] = 2




Figure 10-62. UPM Clock Scheme for LCRR[CLKDIV] = 4 or 8

### 10.4.4.4 RAM Array

The RAM array for each UPM is 64 locations deep and 32 bits wide, as shown in Figure 10-63. The signals at the bottom of the figure are UPM outputs. The selected  $\overline{LCS}_n$  is for the bank that matches the current address. The selected  $\overline{LBS}$  is for the byte lanes read or written by the access.

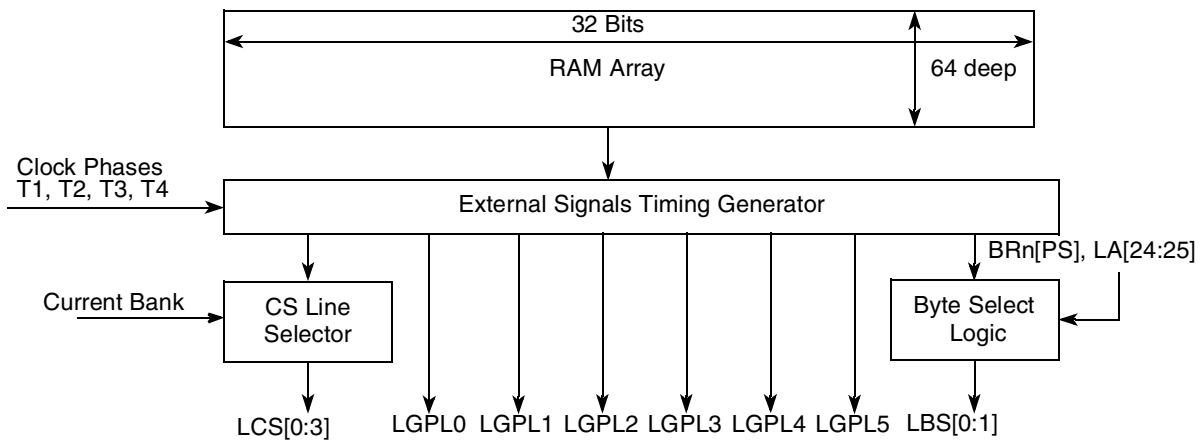


Figure 10-63. RAM Array and Signal Generation

### 10.4.4.4.1 RAM Words

The RAM word is a 32-bit microinstruction stored in one of 64 locations in the RAM array. It specifies timing for external signals controlled by the UPM. Figure 10-37 shows the RAM word fields. When  $LCRR[CLKDIV] = 4$  or  $8$ , the  $CST_n$  and  $BST_n$  bits determine the state of UPM signals  $\overline{LCS}_n$  and  $\overline{LBS}[0:1]$  at each quarter phase of the bus clock. When  $LCRR[CLKDIV] = 2$ ,  $CST_2$  and  $CST_4$  are ignored and the external has the values defined by  $CST_1$  and  $CST_3$  but extended to half the clock cycle in duration. The same interpretation occurs for the  $BST_n$  bits when  $LCRR[CLKDIV] = 2$ .

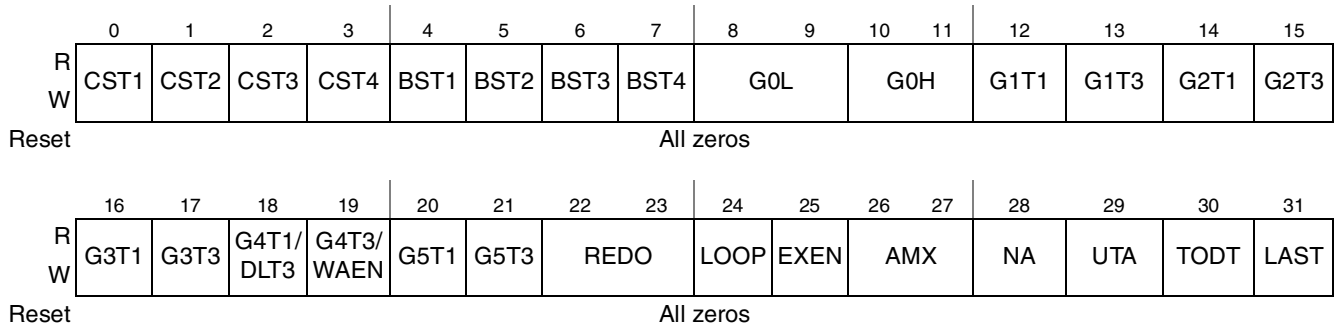


Figure 10-64. RAM Word Fields

Table 10-40 contains descriptions of the RAM word fields.

Table 10-40. RAM Word Field Descriptions

Bits	Name	Description
0	CST1	Chip select timing 1. Defines the state (0 or 1) of $\overline{LCS}_n$ during bus clock quarter phase 1 if $LCRR[CLKDIV] = 4$ or $8$ . Defines the state (0 or 1) of $\overline{LCS}_n$ during bus clock half phase 1 if $LCRR[CLKDIV] = 2$ .
1	CST2	Chip select timing 2. Defines the state (0 or 1) of $\overline{LCS}_n$ during bus clock quarter phase 2 if $LCRR[CLKDIV] = 4$ or $8$ . Ignored when $LCRR[CLKDIV] = 2$ .
2	CST3	Chip select timing 3. Defines the state (0 or 1) of $\overline{LCS}_n$ during bus clock quarter phase 3 if $LCRR[CLKDIV] = 4$ or $8$ . Defines the state (0 or 1) of $\overline{LCS}_n$ during bus clock half phase 2 if $LCRR[CLKDIV] = 2$ .
3	CST4	Chip select timing 4. Defines the state (0 or 1) of $\overline{LCS}_n$ during bus clock quarter phase 4 if $LCRR[CLKDIV] = 4$ or $8$ . Ignored when $LCRR[CLKDIV] = 2$ .
4	BST1	Byte select timing 1. Defines the state (0 or 1) of $\overline{LBS}$ during bus clock quarter phase 1 ( $LCRR[CLKDIV] = 4$ or $8$ ) or bus clock half phase 1 ( $LCRR[CLKDIV] = 2$ ), in conjunction with $BR_n[PS]$ and $LA[24:25]$ .
5	BST2	Byte select timing 2. Defines the state (0 or 1) of $\overline{LBS}$ during bus clock quarter phase 2 ( $LCRR[CLKDIV] = 4$ or $8$ ), in conjunction with $BR_n[PS]$ and $LA[24:25]$ . Ignored when $LCRR[CLKDIV] = 2$ .
6	BST3	Byte select timing 3. Defines the state (0 or 1) of $\overline{LBS}$ during bus clock quarter phase 3 ( $LCRR[CLKDIV] = 4$ or $8$ ) or bus clock half phase 2 ( $LCRR[CLKDIV] = 2$ ), in conjunction with $BR_n[PS]$ and $LA[24:25]$ .
7	BST4	Byte select timing 4. Defines the state (0 or 1) of $\overline{LBS}$ during bus clock quarter phase 4 ( $LCRR[CLKDIV] = 4$ or $8$ ), in conjunction with $BR_n[PS]$ and $LA[24:25]$ . Ignored when $LCRR[CLKDIV] = 2$ .

Table 10-40. RAM Word Field Descriptions (continued)

Bits	Name	Description
8–9	G0L	General purpose line 0 lower. Defines the state of LGPL0 during the bus clock quarter phases 1 and 2 (first half phase). 00 Value defined by MxMR[G0CL] 01 Reserved 10 0 11 1
10–11	G0H	General purpose line 0 higher. Defines the state of LGPL0 during the bus clock quarter phases 3 and 4 (second half phase). 00 Value defined by MxMR[G0CL] 01 Reserved 10 0 11 1
12	G1T1	General purpose line 1 timing 1. Defines the state (0 or 1) of LGPL1 during bus clock quarter phases 1 and 2 (first half phase).
13	G1T3	General purpose line 1 timing 3. Defines the state (0 or 1) of LGPL1 during bus clock quarter phases 3 and 4 (second half phase)
14	G2T1	General purpose line 2 timing 1. Defines state (0 or 1) of LGPL2 during bus clock quarter phases 1 and 2 (first half phase).
15	G2T3	General purpose line 2 timing 3. Defines the state (0 or 1) of LGPL2 during bus clock quarter phases 3 and 4 (second half phase).
16	G3T1	General purpose line 3 timing 1. Defines the state (0 or 1) of LGPL3 during bus clock quarter phases 1 and 2 (first half phase).
17	G3T3	General purpose line 3 timing 3. Defines the state (0 or 1) of LGPL3 during bus clock quarter phases 3 and 4 (second half phase).
18	G4T1/DLT3	General purpose line 4 timing 1/delay time 3. The function of this bit is determined by MxMR[GPL4]. If MxMR[GPL4] = 0 and LGPL4/LUPWAIT pin functions as an output (LGPL4), G4T1/DLT3 defines the state (0 or 1) of LGPL4 during bus clock quarter phases 1 and 2 (first half phase). If MxMR[GPL4] = 1 and LGPL4/LUPWAIT functions as an input (LUPWAIT), if a read burst or single read is executed, G4T1/DLT3 defines the sampling of the data bus as follows: 0 In the current word, the data bus should be sampled at the start of bus clock quarter phase 1 of the next bus clock cycle. 1 In the current word, the data bus should be sampled at the start of bus clock quarter phase 3 of the current bus clock cycle.
19	G4T3/WAEN	General purpose line 4 timing 3/wait enable. Bit function is determined by MxMR[GPL4]. If MxMR[GPL4] = 0 and LGPL4/LUPWAIT pin functions as an output (LGPL4), G4T3/WAEN defines the state (0 or 1) of LGPL4 during bus clock quarter phases 3 and 4 (second half phase). If MxMR[GPL4] = 1 and LGPL4/LUPWAIT functions as an input (LUPWAIT), G4T3/WAEN is used to enable the wait mechanism: 0 LUPWAIT detection is disabled. 1 LUPWAIT is enabled. If LUPWAIT is detected as being asserted, a freeze in the external signals logical values occurs until LUPWAIT is detected as being negated.
20	G5T1	General purpose line 5 timing 1. Defines the state (0 or 1) of LGPL5 during bus clock quarter phases 1 and 2 (first half phase).
21	G5T3	General purpose line 5 timing 3. Defines the state (0 or 1) of LGPL5 during bus clock quarter phases 3 and 4 (second half phase).



Table 10-40. RAM Word Field Descriptions (continued)

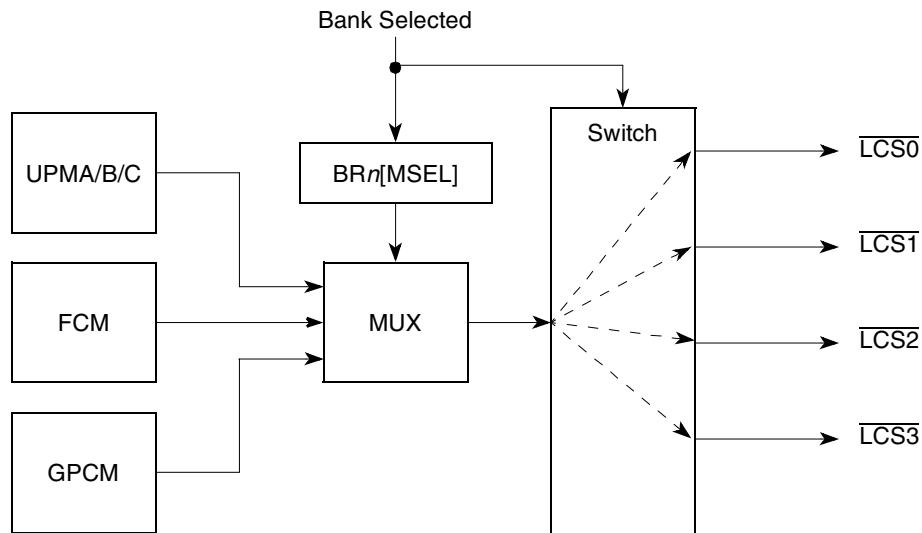
Bits	Name	Description
22–23	REDO	Redo current RAM word. Defines the number of times to execute the current RAM word. 00 Once (normal operation) 01 Twice 10 Three times 11 Four times
24	LOOP	Loop start/end. The first RAM word in the RAM array where LOOP is 1 is recognized as the loop start word. The next RAM word where LOOP is 1 is the loop end word. RAM words between, and including the start and end words, are defined as part of the loop. The number of times the UPM executes this loop is defined in the corresponding loop fields of the MxMR. 0 The current RAM word is not the loop start word or loop end word. 1 The current RAM word is the start or end of a loop. <b>Note:</b> AMX must not change values in any RAM word which begins a loop
25	EXEN	Exception enable. Allows branching to an exception pattern at the exception start address (EXS). When an internal bus monitor time-out exception is recognized and EXEN in the RAM word is set, the UPM branches to the special exception start address (EXS) and begins operating as the pattern defined there specifies. The user should provide an exception pattern to negate signals controlled by the UPM in a controlled fashion. For DRAM control, a handler should negate RAS and CAS to prevent data corruption. If EXEN = 0, exceptions are ignored by UPM (but not by local bus) and execution continues. After the UPM branches to the exception start address, it continues reading until the LAST bit is set in the RAM word. 0 The UPM continues executing the remaining RAM words, ignoring any internal bus monitor time-out. 1 The current RAM word allows a branch to the exception pattern after the current cycle if an exception condition is detected.
26–27	AMX	Address multiplexing. Determines the source of LAD during an LALE phase. Any change in the AMX field initiates a new LALE (address) phase. 00 LAD (and/or in conjunction with LA) is the non-multiplexed address. For example, column address. 01 Reserved 10 LAD (and/or in conjunction with LA) is driven with the multiplexed address according to MxMR[AM]. For example, row address. See <a href="#">Section 10.4.4.4.7, “Address Multiplexing (AMX)”</a> for more information. 11 LAD (and/or in conjunction with LA) is driven with the contents of MAR. Used, for example, to initialize a mode. <b>Note:</b> AMX must not change values in any RAM word which begins a loop. <b>Note:</b> Source ID debug mode is only supported for the AMX = 00 setting.
28	NA	Next burst address. Determines when the address is incremented during a burst access. 0 The address increment function is disabled. 1 The address is incremented in the next cycle. In conjunction with the BRn[PS], the increment value of LAN is 1 or 2 for port sizes of 8 and 16 bits, respectively.
29	UTA	UPM transfer acknowledge. Indicates assertion of transfer acknowledge in the current cycle. 0 Transfer acknowledge is not asserted in the current cycle. 1 Transfer acknowledge is asserted in the current cycle. In case of UPM writes, program UTA and LAST in same RAM word. In case of UPM reads, program UTA and LAST in consecutive or same RAM words.

Table 10-40. RAM Word Field Descriptions (continued)

Bits	Name	Description
30	TODT	Turn-on disable timer. The disable timer associated with each UPM allows a minimum time to be guaranteed between two successive accesses to the same memory bank. This feature is critical when DRAM requires a RAS precharge time. TODT turns the timer on to prevent another UPM access to the same bank until the timer expires. The disable timer period is determined in $MxMR[DSn]$ . The disable timer does not affect memory accesses to different banks. Note that TODT must be set together with LAST, otherwise it is ignored. 0 The disable timer is turned off. 1 The disable timer for the current bank is activated preventing a new access to the same bank (when controlled by the UPMs) until the disable timer expires. For example, precharge time.
31	LAST	Last word. When LAST is read in a RAM word, the current UPM pattern terminates and control signal timing set in the RAM word is applied to the current (and last) cycle. However, if the disable timer is activated and the next access is to the same bank, execution of the next UPM pattern is held off and the control signal values specified in the last word are extended in duration for the number of clock cycles specified in $MxMR[DSn]$ . 0 The UPM continues executing RAM words. 1 Indicates the last RAM word in the program. The service to the UPM request is done after this cycle concludes. In case of UPM writes, program UTA and LAST in same RAM word. In case of UPM reads, program UTA and LAST in consecutive or same RAM words.

#### 10.4.4.4.2 Chip-Select Signal Timing ( $CST_n$ )

If  $BR_n[MSEL]$  of the accessed bank selects a UPM on the currently requested cycle, the UPM manipulates the  $\overline{LCS}_n$  for that bank with timing as specified in the UPM RAM word  $CST_n$  fields. The selected UPM affects only the assertion and negation of the appropriate  $\overline{LCS}_n$  signal. The state of the selected  $\overline{LCS}_n$  signal of the corresponding bank depends on the value of each  $CST_n$  bit. Figure 10-65 shows how UPMs control  $\overline{LCS}_n$  signals.

Figure 10-65.  $\overline{LCS}_n$  Signal Selection

### 10.4.4.4.3 Byte Select Signal Timing ( $BST_n$ )

If  $BR_n[MSEL]$  of the accessed memory bank selects a UPM on the currently requested cycle, the selected UPM affects the assertion and negation of the appropriate  $\overline{LBS}[0:1]$  signal. The timing of both byte-select signals is specified in the RAM word. However,  $\overline{LBS}[0:1]$  are also controlled by the port size of the accessed bank, the number of bytes to transfer, and the address accessed. Figure 10-66 shows how UPMs control  $\overline{LBS}[0:1]$ .

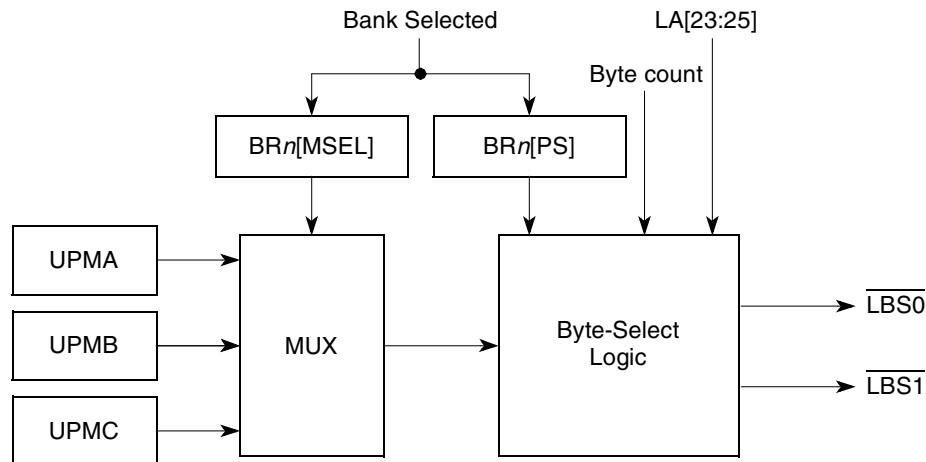


Figure 10-66.  $\overline{LBS}$  Signal Selection

The uppermost byte select ( $\overline{LBS0}$ ), when asserted, indicates that  $LAD[0:7]$  contains valid data during a cycle. Likewise,  $\overline{LBS1}$  indicates that  $LAD[8:15]$  contain valid data. For a UPM refresh timer request, all  $\overline{LBS}[0:1]$  signals are asserted/negated by the UPM according to the refresh pattern only. Following any internal bus monitor exception, the  $\overline{LBS}[0:1]$  signals are negated regardless of the exception handling provided by any UPM exception pattern to prevent spurious writes to external RAM.

### 10.4.4.4.4 General-Purpose Signals ( $G_nT_n$ , $GO_n$ )

The general-purpose signals ( $LGPL[0:5]$ ) each have two bits in the RAM word that define the logical value of the signal to be changed at the rising edge of the bus clock and/or at the falling edge of the bus clock.  $LGPL0$  offers enhancements beyond the other  $LGPL_n$  lines.

$LGPL0$  can be controlled by an address line specified in  $M_xMR[G0CL]$ . To use this feature,  $G0H$  and  $G0L$  should be set in the RAM word. For example, for a SIMM with multiple banks, this address line can be used to switch between internal memory device banks.

### 10.4.4.4.5 Loop Control (LOOP)

The LOOP bit in the RAM word specifies the beginning and end of a set of UPM RAM words that are to be repeated. The first time  $LOOP = 1$ , the memory controller recognizes it as a loop start word and loads the memory loop counter with the corresponding contents of the loop field shown in Table 10-41. The next RAM word for which  $LOOP = 1$  is recognized as a loop end word. When it is reached, the loop counter is decremented by one.

Continued loop execution depends on the loop counter. If the counter is not zero, the next RAM word executed is the loop start word. Otherwise, the next RAM word executed is the one after the loop end word. Loops can be executed sequentially but cannot be nested. Also, special care must be taken:

- LAST and LOOP must not be set together.
- Loop start word should not have an AMX change with regard to the previous word.

**Table 10-41. MxMR Loop Field Use**

Request Serviced	Loop Field
Read single-beat cycle	RLF
Read burst cycle	RLF
Write single-beat cycle	WLF
Write burst cycle	WLF
Refresh timer expired	TLF
RUN command	RLF

#### 10.4.4.4.6 Repeat Execution of Current RAM Word (REDO)

The REDO function is useful for wait-state insertion in a long UPM routine that would otherwise need too many RAM words. Setting the REDO bits of the RAM word to a nonzero value causes the UPM to re-execute the current RAM word up to three more times, as defined in the REDO field of the current RAM word.

Special care must be taken in the following cases:

- When UTA and REDO are set together, TA is asserted the number of times specified by the REDO function.
- When NA and REDO are set together, the address is incremented the number of times specified by the REDO function.
- When LOOP and REDO are set together, the loop mechanism works as usual and the line is repeated according to the REDO function.
- LAST and REDO must not be set together.
- REDO should not be used within the exception routine.

#### 10.4.4.4.7 Address Multiplexing (AMX)

Address lines can be controlled by the user-provided pattern in the UPM. The address multiplex (AMX) bits in the RAM word can choose between driving the transaction address (AMX = 00), driving it according to the multiplexing specified by the MxMR[AM] field (AMX = 10), or driving the contents of MAR (AMX = 11) on the address signals. The next address (NA) bit of the RAM word does not affect LA signals, unless AMX = 00 and chooses the column address for NA = 1.

In all cases, LA[21:25] of the eLBC are driven by the five lsbs of the address selected by AMX, regardless of whether the next address (NA) bit of the RAM word is used to increment the current address. The effect of NA = 1 is visible only when AMX = 00 chooses the column address.

Table 10-42 shows how the RAM word AMX bits and MxMR[AM] settings can be used to affect row × column address multiplexing on the LA[10:25] signals.

**Table 10-42. UPM Address Multiplexing**

	msb		Internal Transaction Address																												lsb															
	0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20	21	22	23	24	25	26	27	28	29		30	31													
AMX = 10 MxMR[AM] = 000 (Row)											LAD										LA																									
AMX = 00 (Col)																												LA																		
AMX = 10 MxMR[AM] = 001 (Row)												10	11	12	13	14	15	16	17	18	19	20	21	22	23	24	25																			
AMX = 00 (Col)																													LA																	
AMX = 10 MxMR[AM] = 010 (Row)												10	11	12	13	14	15	16	17	18	19	20	21	22	23	24	25																			
AMX = 00 (Col)																													LA																	
AMX = 10 MxMR[AM] = 011 (Row)												10	11	12	13	14	15	16	17	18	19	20	21	22	23	24	25																			
AMX = 00 (Col)																													LAD						LA											
AMX = 10 MxMR[AM] = 100 (Row)												10	11	12	13	14	15	16	17	18	19	20	21	22	23	24	25																			
AMX = 00 (Col)																												LAD		LA																
AMX = 10 MxMR[AM] = 101 (Row)												10	11	12	13	14	15	16	17	18	19	20	21	22	23	24	25																			
AMX = 00 (Col)																												LAD		LA																
AMX = 10 MxMR[AM] = 110	Reserved																																													
AMX = 10 MxMR[AM] = 111	Reserved																																													

Note that any change to the AMX field from one RAM word to the next RAM word executed results in an address phase on the {LAD<sub>n</sub>, LA<sub>n</sub>} bus with the assertion of LALE for the number of cycles set for LALE in the OR<sub>n</sub> and LCRR registers. The LGPL[0:5] signals maintain the value specified in the RAM word during the LALE phase.

#### NOTE

AMX must not change values in any RAM word which begins a loop.

#### 10.4.4.4.8 Data Valid and Data Sample Control (UTA)

When a read access is handled by the UPM, and the UTA bit is 1 (data is to be sampled by the eLBC), the value of the DLT3 bit in the same RAM word, in conjunction with MxMR[GPL4], determines when the data input is sampled by the eLBC as follows:

- If  $MxMR[GPL4] = 1$  (G4T4/DLT3 functions as DLT3) and  $DLT3 = 1$  in the RAM word, data is latched on the falling edge of the bus clock instead of the rising edge. The eLBC samples the data on the next falling edge of the bus clock, which is during the middle of the current bus cycle. This feature should be used only in systems without external synchronous bus devices that require mid-cycle sampling.
- If  $MxMR[GPL4] = 0$  (G4T4/DLT3 functions as G4T4), or if  $MxMR[GPL4] = 1$  but  $DLT3 = 0$  in the RAM word, data is latched on the rising edge of the bus clock, which occurs at the end of the current bus clock cycle (normal operation).

Figure 10-67 shows how data sampling is controlled by the UPM.




Figure 10-67. UPM Read Access Data Sampling

#### 10.4.4.4.9 LGPL[0:5] Signal Negation (LAST)

When the LAST bit is read in a RAM word, the current UPM pattern is terminated at the end of the current cycle. On the next cycle (following LAST) all the UPM signals are negated unconditionally (driven to logic 1), unless there is a back-to-back UPM request pending. In this case, the signal values for the cycle following the one in which the LAST bit was set are taken from the first RAM word of the pending UPM routine.

In case of UPM writes, program UTA and LAST in same RAM word. In case of UPM reads, program UTA and LAST in consecutive or same RAM words.

#### 10.4.4.4.10 Wait Mechanism (WAEN)

The WAEN bit in the RAM array word can be used to enable the UPM wait mechanism in selected UPM RAM words. If the UPM reads a RAM word with WAEN set, the external LUPWAIT signal is sampled and synchronized by the memory controller as if it were an asynchronous signal. The WAEN bit is ignored if  $LAST = 1$  in the same RAM word.

Synchronization of LUPWAIT starts at the rising edge of the bus clock and takes at least 1 bus cycle to complete. If LUPWAIT is asserted and  $WAEN = 1$  in the current UPM word, the UPM is frozen until LUPWAIT is negated. The value of external signals driven by the UPM remains as indicated in the previous RAM word. When LUPWAIT is negated, the UPM continues normal functions. Note that during WAIT cycles, the UPM does not handle data.

Figure 10-68 shows how the WAEN bit in the word read by the UPM and the LUPWAIT signal are used to hold the UPM in a particular state until LUPWAIT is negated. As the example shows, the  $\overline{LCSn}$  and LGPL1 states and the WAEN value are frozen until LUPWAIT is recognized as negated. WAEN is typically set before the line that contains  $UTA = 1$ . Note that if WAEN and NA are both set in the same RAM word, NA causes the burst address to increment once as normal regardless of whether the UPM freezes.

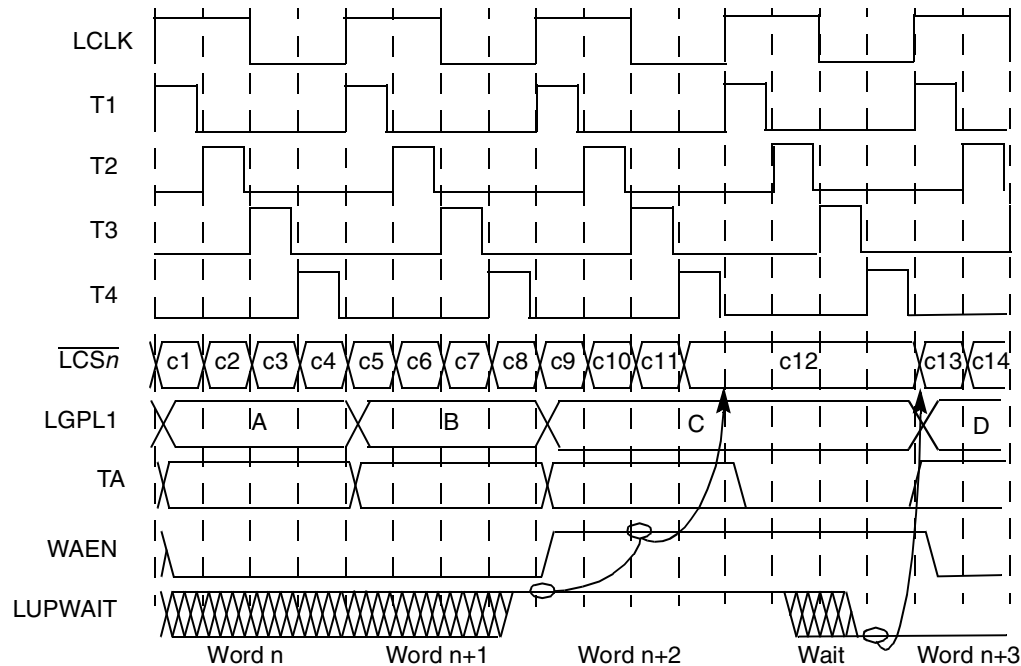


Figure 10-68. Effect of LUPWAIT Signal

#### 10.4.4.5 Synchronous Sampling of LUPWAIT for Early Transfer Acknowledge

If LUPWAIT is to be considered an asynchronous signal, which can be asserted/negated at any time, no UPM RAM word must contain both  $WAEN = 1$  and  $UTA = 1$  simultaneously.

However, programming  $WAEN = 1$  and  $UTA = 1$  in the same RAM word allows the UPM to treat LUPWAIT as a synchronous signal, which must meet set-up and hold times in relation to the rising edge of the bus clock. In this mode, as soon as UPM samples LUPWAIT negated on the rising edge of the bus clock, it immediately generates an internal transfer acknowledge, which allows a data transfer one bus clock cycle later. The generation of transfer acknowledge is early because LUPWAIT is not re-synchronized. The acknowledge occurs early or normally depending on whether the UPM was already frozen in WAIT cycles or not. This feature allows the synchronous negation of LUPWAIT to affect a data transfer, even if UTA, WAEN, and LAST are set simultaneously.

#### 10.4.4.6 Extended Hold Time on Read Accesses

Slow memory devices that take a long time to turn off their data bus drivers on read accesses should choose some non-zero combination of  $OR_n[TRLX]$  and  $OR_n[EHTR]$ . The next accesses after a read access to the

slow memory device is delayed by the number of clock cycles specified in the  $OR_n$  register in addition to any existing bus turnaround cycle.

## 10.5 Initialization/Application Information

### 10.5.1 Interfacing to Peripherals in Different Address Modes

This section provides guidelines for interfacing to peripherals.

#### 10.5.1.1 Multiplexed Address/Data Bus for 26-Bit Addressing

In this mode, the eLBC is used with port sizes of 8 and 16 bits; address bits  $A[6:21]$  will appear on  $LAD[0:15]$  (with zero bits on  $LAD[16:31]$ ) during address phases, while the lower 10 bits of the address,  $A[22:31]$ , are driven permanently on  $LA[16:25]$ . The connection is illustrated in [Figure 10-69](#).

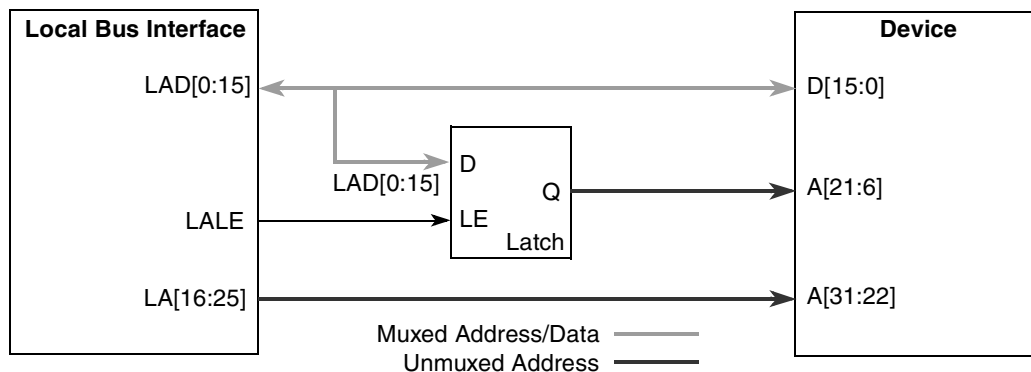


Figure 10-69. Multiplexed Address/Data Bus for 26-Bit Addressing

#### 10.5.1.2 Non-Multiplexed Address and Data Buses

For small address space applications the address latch may be eliminated entirely if the local bus address is taken entirely from  $LA[0:25]$ , in which case addresses driven onto  $LAD$  during address phases are simply ignored. The connection is illustrated in [Figure 10-70](#). In non-multiplexed mode, waveforms etc remain the same except that few things need not be taken care of like  $ASHIFT$  parameter,  $LAD$  bus turnaround time,  $LALE$  timings etc.






Figure 10-70. Non-Multiplexed Address and Data Buses

### 10.5.1.3 Multiplexed Address and Data to Save Maximum Pins in 8- to 16-Bit Addressing

With the use of a feature called address byte swap by setting `LBCR[ABSWP]`, data and address muxing can be swapped from the default available. Currently, `LAD[0:31]` carries `A[0:31]`. In case of 8-bit interface with 8-bit addressing we do not get benefit of pin reduction with the available muxing. This is because, the MSB of data is muxed with MSB of address. While 8-bit data required is `LAD[0:7]` and address required is lower order bits `A[24:31]` we need to pull out all the 16 bits out of the device. For pin limited devices, this feature can be used where `LAD[0:7]` is mapped to the lsb's of `A[24:31]` and `LAD[8:15]` carries `lsb+1 [16:23]`. As a result while interfacing with 8 bit address and 8-bit data only `LAD[0:7]` is suffice with `LALE` pin. The only drawback of this feature is that it does not support burst as all the address is latched from `LAD` bus.

### 10.5.1.4 Peripheral Hierarchy on the Local Bus for High Bus Speeds

To achieve the highest possible bus speeds on the local bus, it is recommended to reduce the number of devices connected directly to the bus. For best results, only one bank of synchronous SRAMs should have a direct connection, and a bus demultiplexer should be used to replace separate latch and separate bus transceiver combinations. Figure 10-71 shows an example of such a hierarchy. This section is only a guideline, and the board designer must simulate the electric characteristics of the scenario to determine the maximum operating frequency.

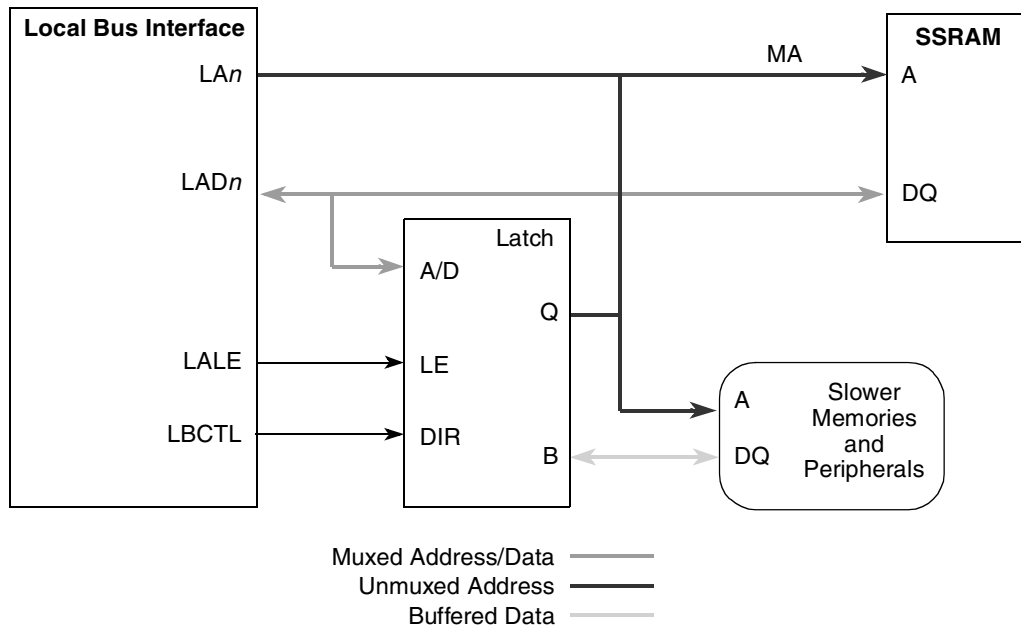


Figure 10-71. Local Bus Peripheral Hierarchy for High Bus Speeds

### 10.5.1.5 GPCM Timings

In case a system contains a memory hierarchy with high speed synchronous memories (synchronous SRAM) and lower speed asynchronous memories (for example, FLASH EPROM and peripherals) the GPCM-controlled memories should be decoupled by buffers to reduce capacitive loading on the bus. Those buffers have to be taken into account for the timing calculations.




Figure 10-72. GPCM Address Timings

To calculate address setup timing for a slower peripheral/memory device, several parameters have to be added: propagation delay for the address latch, propagation delay for the buffer and the address setup for the actual peripheral. Typical values for the two propagation delays are in the order of 3–6 ns, so for a 133-MHz bus frequency,  $\overline{LCS}$  should arrive on the order of 3 bus clocks later.

For data timings, only the propagation delay of one buffer plus the actual data setup time has to be considered.




Figure 10-73. GPCM Data Timings

## 10.5.2 Bus Turnaround

Because the local bus uses multiplexed address and data, special consideration must be given to avoid bus contention at bus turnaround. The following cases must be examined:

- Address phase after previous read
- Read data phase after address phase
- Read-modify-write cycle for parity protected memory banks
- UPM cycles with additional address phases

The bus does not change direction for the following cases so they need no special attention:

- Continued burst after the first beat
- Write data phase after address phase
- Address phase after previous write

### 10.5.2.1 Address Phase after Previous Read

During a read cycle, the memory/peripheral drives the bus and the bus transceiver drives LAD. After the data has been sampled, the output drivers of the external device must be disabled. This can take some time; for slow devices the EHTR feature of the GPCM or the programmability of the UPM should be used to guarantee that those devices have stopped driving the bus when the eLBC memory controller ends the bus cycle.

In this case, after the previous cycle ends, LBCTL goes high and changes the direction of the bus transceiver. The eLBC then inserts a bus turnaround cycle to avoid contention. The external device has now already placed its data signals in high impedance and no bus contention will occur.

### 10.5.2.2 Read Data Phase after Address Phase

During the address phase, LAD actively drives the address and LBCTL is high, driving the bus transceivers in the same direction as during a write. After the end of the address phase, LBCTL goes low and changes the direction of the bus transceiver. The eLBC places the LAD signals in high impedance after its  $t_{dis}(LB)$ . The LBCTL will have its new state after  $t_{en}(LB)$  and, because this is an asynchronous input, the transceiver starts to drive those signals after its  $t_{en}(\text{transceiver})$  time. The system designer has to ensure, that  $[t_{en}(LB) + t_{en}(\text{transceiver})]$  is larger than  $t_{dis}(LB)$  to avoid bus contention.

### 10.5.2.3 Read-Modify-Write Cycle for Parity Protected Memory Banks

Principally, a read-modify-write cycle is a read cycle immediately followed by a write cycle. Because the write cycle will have a new address phase in any case, this basically is the same case as an address phase after a previous read.

### 10.5.2.4 UPM Cycles with Additional Address Phases

The flexibility of the UPM allows the user to insert additional address phases during read cycles by changing the AMX field, therefore turning around the bus during one pattern. The eLBC automatically inserts a single bus turnaround cycle if the bus (LAD) was previously high impedance for any reason, such as a read, before LALE is driven and LAD is driven with the new address. The turnaround cycle is not inserted on a write, because the bus was already driven to begin with.

However, bus contention could potentially still occur on the far side of a bus transceiver. It is the responsibility of the designer of the UPM pattern to guarantee that enough idle cycles are inserted in the UPM pattern to avoid this.

## 10.5.3 Interface to Different Port-Size Devices

The eLBC supports 8- and 16-bit data port sizes. However, the bus requires that the portion of the data bus used for a transfer to or from a particular port size be fixed. A 16-bit port must reside on LAD[0:15], and an 8-bit port must reside on LAD[0:7]. The local bus always tries to transfer the maximum amount of data on all bus cycles. [Figure 10-74](#) shows the device connections on the data bus.




Figure 10-74. Interface to Different Port-Size Devices

Table 10-43 lists the bytes required on the data bus for read cycles.

**Table 10-43. Data Bus Drive Requirements For Read Cycles**

Transfer Size	Address State <sup>1</sup> 3 lsbs	Port Size/LAD Data Bus Assignments							
		16-Bit				8-Bit			
		0-7	8-15	16-23	24-31	0-7	8-15	16-23	24-31
Byte	000	OP0	—			OP0			
	001	—	OP1			OP1			
	010	OP2	—			OP2			
	011	—	OP3			OP3			
	100	OP4	—			OP4			
	101	—	OP5			OP5			
	110	OP6	—			OP6			
	111	—	OP7			OP7			
Half Word	000	OP0	OP1			OP0			
	001	—	OP1			OP1			
	010	OP2	OP3			OP2			
	100	OP4	OP5			OP4			
	101	—	OP5			OP5			
	110	OP6	OP7			OP6			
Word	000	OP0	OP1			OP0			
	100	OP4	OP5			OP4			

<sup>1</sup> Address state is the calculated address for port size.

## 10.5.4 Command Sequence Examples for NAND Flash EEPROM

In order to program the eLBC and FCM for executing NAND Flash command sequences, command codes and pause states should be obtained from the relevant NAND Flash device data sheet and programmed into FCM configuration registers. This section illustrates some common sequences for large-page, multi-gigabit NAND Flash EEPROMs; however, details should be verified against manufacturers' specific programming data.

Throughout these examples it is assumed that one or more banks of eLBC has been configured under FCM control ( $BR_n[MSEL] = 001$ ), with base address, port size, ECC mode, and timing parameters configured in accordance with the device's hardware specifications.

### 10.5.4.1 NAND Flash Soft Reset Command Sequence Example

An example of configuring FCM to execute a soft reset command to large-page NAND Flash is shown in [Table 10-44](#). This sequence does not require use of the shared FCM buffer RAM. The sequence is initiated by writing FMR[OP] = 10, and issuing a special operation to the bank. At the conclusion of the sequence, eLBC will issue a command complete interrupt (LTESR[CC]) if interrupts are enabled.

**Table 10-44. FCM Register Settings for Soft Reset (ORn[PGS] = 1)**

Register	Initial Contents	Description
FCR	0xFF000000	CMD0 = 0xFF = reset command; other commands unused
FBAR	—	unused
FPAR	—	unused
FBCR	—	unused
MDR	—	unused
FIR	0x40000000	OP0 = CM0 = command 0; OP1–OP7 = NOP

### 10.5.4.2 NAND Flash Read Status Command Sequence Example

An example of configuring FCM to execute a status read command to large-page NAND Flash is shown in [Table 10-45](#). This sequence does not require use of the shared FCM buffer RAM, but reads the NAND Flash status into register MDR[AS0]. The sequence is initiated by writing FMR[OP] = 10 and issuing a special operation to the bank. At the conclusion of the sequence, eLBC will issue a command complete interrupt (LTESR[CC]) if interrupts are enabled.

**Table 10-45. FCM Register Settings for Status Read (ORn[PGS] = 1)**

Register	Initial Contents	Description
FCR	0x70000000	CMD0 = 0x70 = read status command; other commands unused
FBAR	—	unused
FPAR	—	unused
FBCR	—	unused
MDR	—	Status returned in AS0
FIR	0x4B000000	OP0 = CM0 = command 0; OP1 = RS = read status to MDR; OP2–OP7 = NOP

### 10.5.4.3 NAND Flash Read Identification Command Sequence Example

An example of configuring FCM to execute a status ID command to large-page NAND Flash is shown in [Table 10-46](#). This sequence does not require use of the shared FCM buffer RAM, but uses MDR to set up a dummy address prior to the sequence, and then to receive the first 4 bytes of ID during the sequence. The sequence is initiated by writing FMR[OP] = 10, and issuing a special operation to the bank. At the

conclusion of the sequence, eLBC will issue a command complete interrupt (LTESR[CC]) if interrupts are enabled. MDR[AS3–AS0] then can be read to obtain the first 4 bytes of NAND Flash ID.

**Table 10-46. FCM Register Settings for ID Read (ORn[PGS] = 1)**

Register	Initial Contents	Description
FCR	0x90000000	CMD0 = 0x90 = read ID command; other commands unused
FBAR	—	unused
FPAR	—	unused
FBCR	—	unused
MDR	0x00000000	AS0 = 0x00 = dummy address for read ID command; AS0–AS3 return with first 4 bytes of ID code
FIR	0x43BBBBB0	OP0 = CM0 = command 0; OP1 = UA = user address from MDR; OP2–OP6 = RS = read 4 bytes ID into MDR[AS3–AS0]; OP7 = NOP

#### 10.5.4.4 NAND Flash Page Read Command Sequence Example

An example of configuring FCM to execute a random page read command to large-page NAND Flash is shown in Table 10-47. This sequence reads an entire page (main and spare region) into the shared FCM buffer RAM, checking ECC as it proceeds. The sequence is initiated by writing FMR[OP] = 10, and issuing a special operation to the bank. A few cycles before completion itself, FECCn gets updated with the ECC bytes for the main region validated by FECCn[0]. At the conclusion of the sequence, eLBC will issue a command complete interrupt (LTESR[CC]) if interrupts are enabled. Once the sequence has completed, the shared buffer (buffer 1 for page index 5) and transfer error registers (LTECCR that reports the 512 blocks with unibit /multibit errors if any) are valid.

**Table 10-47. FCM Register Settings for Page Read (ORn[PGS] = 1)**

Register	Initial Contents	Description
FCR	0x00300000	CMD0 = 0x00 = random read address entry; CMD1 = 0x30 = read page
FBAR	block index (e.g. block 0x00010ab4)	BLK locates index of 128-Kbyte block
FPAR	page offset (e.g. 0x00005000 locates page 5, buffer 1)	PI locates page index in BLK; PI mod 2 indexes FCM buffer RAM; MS = 0 and CI = 0
FBCR	0x00000000	BC = 0 to read entire 2112-byte page with ECC check
MDR	—	unused
FIR	0x4125E000	OP0 = CM0 = command 0; OP1 = CA = column address; OP2 = PA = page address; OP3 = CM1 = command 1; OP4 = RBW = wait on Flash ready and read data into FCM buffer; OP5–OP7 = NOP

### 10.5.4.5 NAND Flash Block Erase Command Sequence Example

An example of configuring FCM to execute a block erase command to large-page NAND Flash is shown in [Table 10-48](#). This sequence does not require use of the shared FCM buffer RAM, but returns with the erase status in MDR[AS0]. The sequence is initiated by writing FMR[OP] = 10, and issuing a special operation to the bank. At the conclusion of the sequence, eLBC will issue a command complete interrupt (LTER[CC]) if interrupts are enabled.

Note that operations specified by OP3 and OP4 (status read) should never be skipped while erasing a NAND Flash device, because, in case that happens, contention may arise on LGPL4. A possible case is that the next transaction from eLBC may try to use that pin as an output and since the NAND Flash device might already be driving it, contention will occur. In case OP3 and OP4 operations are skipped, it may also happen that a new command is issued to the NAND Flash device even when the device has not yet finished processing the previous request. This may also result in unpredictable behavior.

**Table 10-48. FCM Register Settings for Block Erase (ORn[PGS] = 1)**

Register	Initial Contents	Description
FCR	0x6070D000	CMD0 = 0x60 = block address entry; CMD1 = 0x70 = read status CMD2 = 0xD0 = erase block;
FBAR	block index (e.g. block 0x00010AB4)	BLK locates index of 128-Kbyte block
FPAR	0x00000000	PI = 0 to locate block boundary
FBCR	—	unused
MDR	—	returns with AS0 holding erase status
FIR	0x426DB000	OP0 = CM0 = command 0; OP1 = PA = page address; OP2 = CM2 = command 2; OP3 = CW1 = wait on Flash ready and issue command 1; OP4 = RS = read erase status into MDR[AS0]; OP5–OP7 = NOP

### 10.5.4.6 NAND Flash Program Command Sequence Example

An example of configuring FCM to execute a program command to large-page NAND Flash is shown in [Table 10-49](#). This sequence writes an entire page (main and spare region) from the shared FCM buffer RAM, generating ECC as it proceeds. The shared buffer (buffer 1 for page index 5) must be initialized by software prior to starting the sequence. The sequence is initiated by writing FMR[OP] = 10, and issuing a special operation to the bank. At the conclusion of the sequence, eLBC will issue a command complete interrupt (LTER[CC]) if interrupts are enabled. The status of the programming operation is returned in MDR[AS0].

Note that operations specified by OP5 and OP6 (status read) should never be skipped while programming a NAND Flash device, because, in case that happens, contention may arise on LGPL4. A possible case is that the next transaction from eLBC may try to use that pin as an output and since the NAND Flash device might already be driving it, contention will occur. In case OP5 and OP6 operations are skipped, it may also



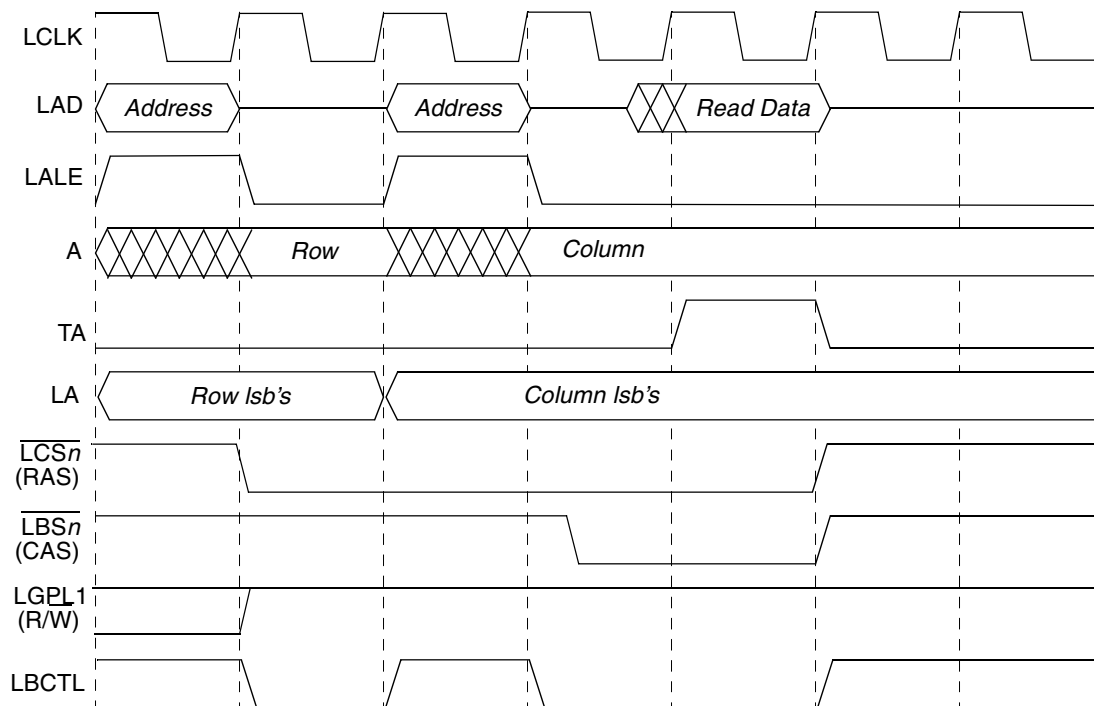
happen that a new command is issued to the NAND Flash device even when the device has not yet finished processing the previous request. This may also result in unpredictable behavior.

**Table 10-49. FCM Register Settings for Page Program (OR<sub>n</sub>[PGS] = 1)**

Register	Initial Contents	Description
FCR	0x80701000	CMD0 = 0x80 = page address and data entry; CMD1 = 0x70 = read status CMD2 = 0x10 = program page;
FBAR	block index (e.g. block 0x00010AB4)	BLK locates index of 128-Kbyte block
FPAR	page offset (e.g. 0x00005000 locates page 5, buffer 1)	PI locates page index in BLK; PI mod 2 indexes FCM buffer RAM; MS = 0 and CI = 0
FBCR	0x00000000	BC = 0 to write entire 2112-Byte page with ECC generation
MDR	—	returns with AS0 holding program status
FIR	0x41286DB0	OP0 = CM0 = command 0; OP1 = CA = column address; OP2 = PA = page address; OP3 = WB = write data from buffer; OP4 = CM2 = command 2; OP5 = CW1 = wait on Flash ready and issue command 1; OP6 = RS = read erase status into MDR[AS0]; OP7 = NOP

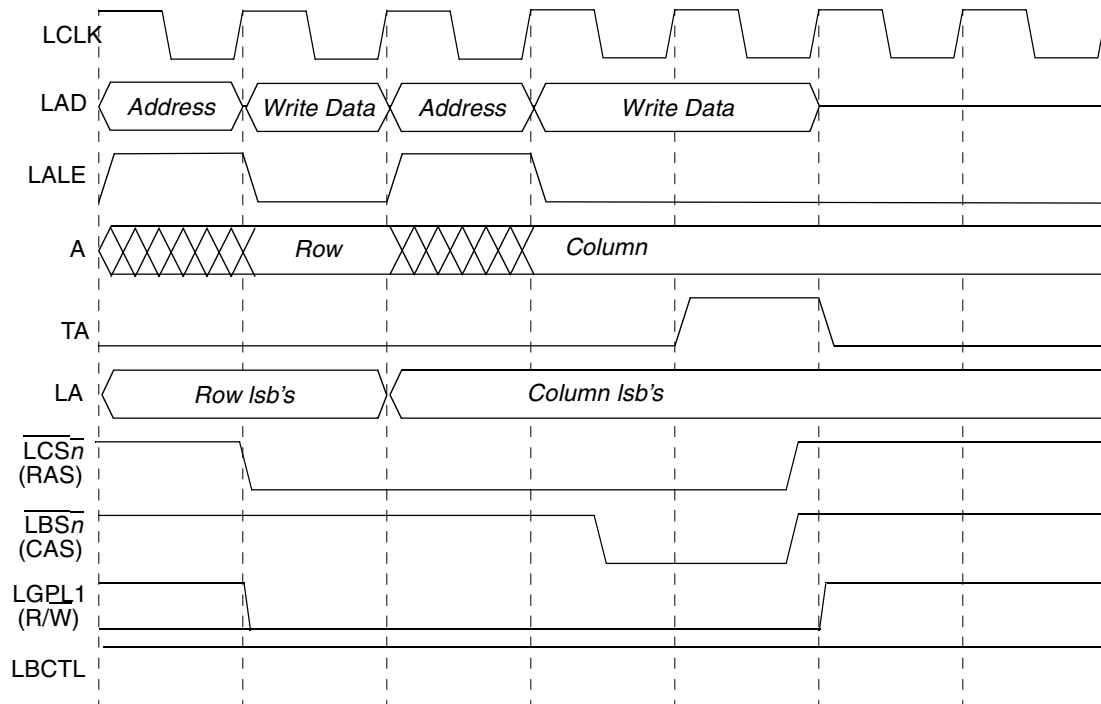
### 10.5.5 Interfacing to Fast-Page Mode DRAM Using UPM

Connecting the local bus UPM controller to a DRAM device requires a detailed examination of the timing diagrams representing the possible memory cycles that must be performed when accessing this device. This section describes timing diagrams for various UPM configurations for fast-page mode DRAM, with LCRR[CLKDIV] = 4 or 8. These illustrative examples may not represent the timing necessary for any specific device used with the eLBC. Here, LGPL1 is programmed to drive  $\overline{R/W}$  of the DRAM, although any LGPL<sub>n</sub> signal may be used for this purpose.



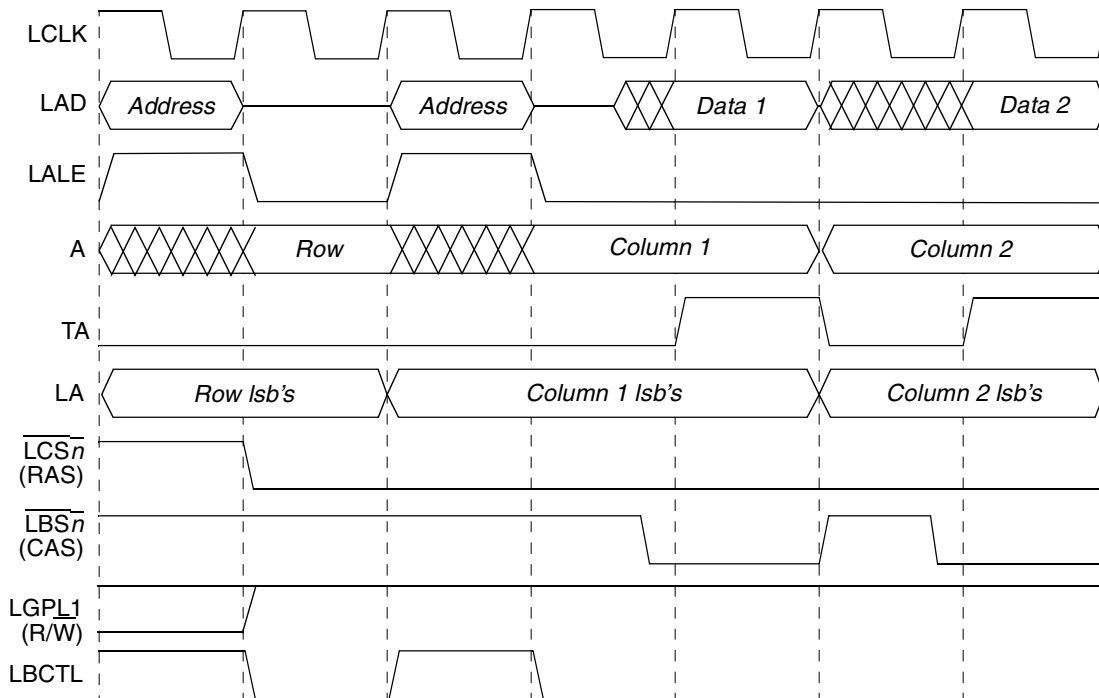
cst1	0	LALE pause (due to change in AMX)	0	0	Bit 0
cst2	0		0	0	Bit 1
cst3	0		0	0	Bit 2
cst4	0		0	0	Bit 3
bst1	1		1	0	Bit 4
bst2	1		0	0	Bit 5
bst3	1		0	0	Bit 6
bst4	1		0	0	Bit 7
g0l0					Bit 8
g0l1					Bit 9
g0h0					Bit 10
g0h1					Bit 11
g1t1	1		1	1	Bit 12
g1t3	1		1	1	Bit 13
g2t1					Bit 14
g2t3					Bit 15
g3t1					Bit 16
g3t3					Bit 17
g4t1					Bit 18
g4t3					Bit 19
g5t1					Bit 20
g5t3					Bit 21
redo[0]					Bit 22
redo[1]					Bit 23
loop	0		0	0	Bit 24
exen	0		0	0	Bit 25
amx0	1		0	0	Bit 26
amx1	0		0	0	Bit 27
na	0		0	0	Bit 28
uta	0		0	1	Bit 29
todt	0		0	1	Bit 30
last	0		0	1	Bit 31
	RSS	RSS+1	RSS+1	RSS+2	

Figure 10-75. Single-Beat Read Access to FPM DRAM




cst1	0	LALE pause (due to change in AMX)	0	0	Bit 0	
cst2	0		0	0	Bit 1	
cst3	0		0	0	Bit 2	
cst4	0		0	1	Bit 3	
bst1	1		1	0	Bit 4	
bst2	1		1	0	Bit 5	
bst3	1		0	0	Bit 6	
bst4	1		0	1	Bit 7	
g0i0					Bit 8	
g0i1					Bit 9	
g0h0					Bit 10	
g0h1					Bit 11	
g1t1	0			0	0	Bit 12
g1t3	0			0	0	Bit 13
g2t1						Bit 14
g2t3						Bit 15
g3t1					Bit 16	
g3t3					Bit 17	
g4t1					Bit 18	
g4t3					Bit 19	
g5t1					Bit 20	
g5t3					Bit 21	
redo[0]					Bit 22	
redo[1]					Bit 23	
loop	0		0	0	Bit 24	
exen	0		0	0	Bit 25	
amx0	1		0	0	Bit 26	
amx1	0		0	0	Bit 27	
na	0		0	0	Bit 28	
uta	0		0	1	Bit 29	
todt	0		0	1	Bit 30	
last	0		0	1	Bit 31	
	WSS	WSS+1	WSS+1	WSS+2		

Figure 10-76. Single-Beat Write Access to FPM DRAM




cst1	0	LALE pause (due to change in AMX)	0	0	1	Bit 0
cst2	0		0	0	1	Bit 1
cst3	0		0	0	1	Bit 2
cst4	0		0	0	1	Bit 3
bst1	1		1	0	1	Bit 4
bst2	1		1	0	1	Bit 5
bst3	1		1	0	1	Bit 6
bst4	1		0	0	1	Bit 7
g0i0						Bit 8
g0i1						Bit 9
g0h0						Bit 10
g0h1						Bit 11
g1i1	1		1	1	1	Bit 12
g1i3	1		1	1	1	Bit 13
g2i1						Bit 14
g2i3						Bit 15
g3i1						Bit 16
g3i3						Bit 17
g4i1						Bit 18
g4i3						Bit 19
g5i1						Bit 20
g5i3						Bit 21
redo[0]						Bit 22
redo[1]						Bit 23
loop	0		1	1	0	Bit 24
exen	0		0	1	0	Bit 25
amx0	1		0	0	0	Bit 26
amx1	0		0	0	0	Bit 27
na	0		0	1	0	Bit 28
uta	0		0	1	0	Bit 29
todt	0		0	0	1	Bit 30
last	0		0	0	1	Bit 31
	<b>RBS</b>		<b>RBS+1</b>	<b>RBS+2</b>	<b>RBS+3</b>	

Figure 10-77. Burst Read Access to FPM DRAM Using LOOP (Two Beats Shown)



cst1	1	0	0	Bit 0
cst2	1	0	0	Bit 1
cst3	1	0	1	Bit 2
cst4	1	0	1	Bit 3
bst1	1	0	0	Bit 4
bst2	0	0	0	Bit 5
bst3	0	0	1	Bit 6
bst4	0	0	1	Bit 7
g0i0				Bit 8
g0i1				Bit 9
g0h0				Bit 10
g0h1				Bit 11
g1i1				Bit 12
g1i3				Bit 13
g2i1				Bit 14
g2i3				Bit 15
g3i1				Bit 16
g3i3				Bit 17
g4i1				Bit 18
g4i3				Bit 19
g5i1				Bit 20
g5i3				Bit 21
redo[0]				Bit 22
redo[1]				Bit 23
loop	0	0	0	Bit 24
exen	0	0	0	Bit 25
amx0	0	0	0	Bit 26
amx1	0	0	0	Bit 27
na	0	0	0	Bit 28
uta	0	0	0	Bit 29
todt	0	0	1	Bit 30
last	0	0	1	Bit 31
	<b>PTS</b>	<b>PTS+1</b>	<b>PTS+2</b>	

Figure 10-78. Refresh Cycle (CBR) to FPM DRAM



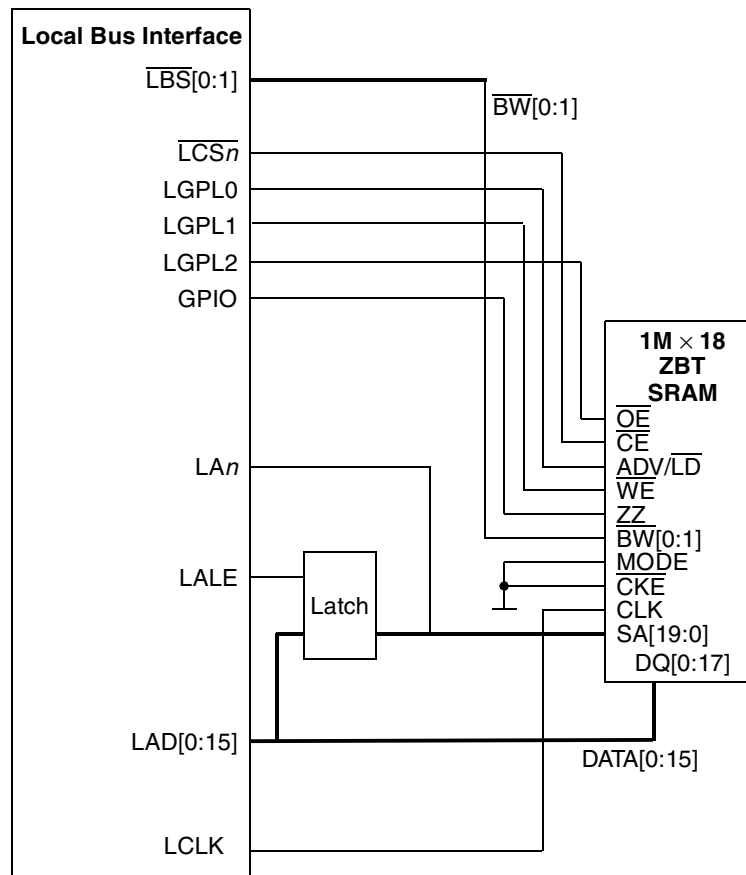
cst1	1	Bit 0
cst2	1	Bit 1
cst3	1	Bit 2
cst4	1	Bit 3
bst1	1	Bit 4
bst2	1	Bit 5
bst3	1	Bit 6
bst4	1	Bit 7
g0l0		Bit 8
g0l1		Bit 9
g0h0		Bit 10
g0h1		Bit 11
g1l1		Bit 12
g1t3		Bit 13
g2t1		Bit 14
g2t3		Bit 15
g3t1		Bit 16
g3t3		Bit 17
g4t1		Bit 18
g4t3		Bit 19
g5t1		Bit 20
g5t3		Bit 21
redo[0]		Bit 22
redo[1]		Bit 23
loop	0	Bit 24
exen	0	Bit 25
amx0	0	Bit 26
amx1	0	Bit 27
na	0	Bit 28
uta	0	Bit 29
todt	1	Bit 30
last	1	Bit 31
EXS		

Figure 10-79. Exception Cycle

### 10.5.6 Interfacing to ZBT SRAM Using UPM

ZBT SRAMs have been designed to optimize the performance of table access in networking applications. This section describes how to interface to ZBT SRAMs. Figure 10-80 shows the connections. The UPM is used to generate control signals. The same interfacing is used for pipelined and flow-through versions of ZBT SRAMs. However different UPM patterns must be generated for those cases. Because ZBT

SRAMs will mostly be used by performance-critical applications, we assume here that, typically, the maximum width of the local bus of 16 bits will be used.



**Figure 10-80. Interface to ZBT SRAM**

ZBT SRAMs allow different configurations. For the local bus, the burst order should be set to linear burst order by tying the mode pin to GND.  $\overline{\text{CKE}}$  should also be tied to ground.

ZBT SRAMs perform four-beat bursts. Because the eLBC generates sixteen-beat transactions (for 16-bit ports) the UPM breaks down each burst into four consecutive four-beat bursts. The internal address generator of the eLBC generates the new  $\{A21, A22\}$  for the second, third, and fourth burst. In other words, because linear burst is used on the SRAM, the device itself bursts with the burst addresses of  $[0:1:2:3]$ . The local bus always generates linear bursts and expects  $[0:1:2:3:4:5:6:7:\dots:15]$ . Therefore, four consecutive linear bursts of the ZBT SRAM with  $\{A21, A22\} = \{0,0\}$  for the first burst,  $\{A21, A22\} = \{0,1\}$  for the second burst,  $\{A21, A22\} = \{1,0\}$  for the third burst, and  $\{A21, A22\} = \{1,1\}$  for the fourth burst give the desired burst pattern.

The UPM also supports single beat accesses. Because the ZBT SRAM does not support this and always responds with a burst, the UPM pattern has to take care that data for the critical beat is provided (for write) or sampled (for read), and that the rest of the burst is ignored (by negating  $\overline{\text{WE}}$ ). The UPM controller basically has to wait for the end of the SRAM burst to avoid bus contention with further bus activities.

If a UPM device has  $\overline{OE}$ , it should not be asserted in the same RAM word as the TA signal. If  $\overline{OE}$  and TA are both asserted in the same RAM word, then the eLBC may not be able to sample the correct data during reads. Therefore,  $\overline{OE}$  must be asserted earlier than TA.



# Chapter 11 Sequencer

## 11.1 Overview

The I/O sequencer switches transactions among its ports, using a buffer pool to minimize blocking. [Figure 11-1](#) is a block diagram of the I/O sequencer (IOS).




Figure 11-1. I/O Sequencer Block Diagram

## 11.1.1 Features

The I/O sequencer includes the following features:

- Switches transactions among its ports
- Contains 8 cache-line (32-byte) buffers to allow streaming of PCI transactions
- Performs address translation on outbound PCI transactions

Note that the number of CSB masters allowed to access to PCI outbound windows is restricted to no more than four non-CPU masters or no more than two non-CPU plus the CPU. If this number is exceeded, deadlock can occur on CSB arbitration.

## 11.2 External Signal Description

The I/O sequencer has no external signals.

## 11.3 Memory Map/Register Definition

Table 11-1 shows the I/O sequencer memory map.

**Table 11-1. Sequencer Memory Map**

Offset	Register	Access	Reset	Section/Page
0x00	POTAR0—PCI outbound translation address register 0	R/W	0x0000_0000	<a href="#">11.4.1/11-3</a>
0x08	POBAR0—PCI outbound base address register 0	R/W	0x0000_0000	<a href="#">11.4.2/11-3</a>
0x10	POCMR0—PCI outbound comparison mask register 0	R/W	0x0000_0000	<a href="#">11.4.3/11-4</a>
0x18	POTAR1—PCI outbound translation address register 1	R/W	0x0000_0000	<a href="#">11.4.1/11-3</a>
0x20	POBAR1—PCI outbound base address register 1	R/W	0x0000_0000	<a href="#">11.4.2/11-3</a>
0x28	POCMR1—PCI outbound comparison mask register 1	R/W	0x0000_0000	<a href="#">11.4.3/11-4</a>
0x30	POTAR2—PCI outbound translation address register 2	R/W	0x0000_0000	<a href="#">11.4.1/11-3</a>
0x38	POBAR2—PCI outbound base address register 2	R/W	0x0000_0000	<a href="#">11.4.2/11-3</a>
0x40	POCMR2—PCI outbound comparison mask register 2	R/W	0x0000_0000	<a href="#">11.4.3/11-4</a>
0x48	POTAR3—PCI outbound translation address register 3	R/W	0x0000_0000	<a href="#">11.4.1/11-3</a>
0x50	POBAR3—PCI outbound base address register 3	R/W	0x0000_0000	<a href="#">11.4.2/11-3</a>
0x58	POCMR3—PCI outbound comparison mask register 3	R/W	0x0000_0000	<a href="#">11.4.3/11-4</a>
0x60	POTAR4—PCI outbound translation address register 4	R/W	0x0000_0000	<a href="#">11.4.1/11-3</a>
0x68	POBAR4—PCI outbound base address register 4	R/W	0x0000_0000	<a href="#">11.4.2/11-3</a>
0x70	POCMR4—PCI outbound comparison mask register 4	R/W	0x0000_0000	<a href="#">11.4.3/11-4</a>
0x78	POTAR5—PCI outbound translation address register 5	R/W	0x0000_0000	<a href="#">11.4.1/11-3</a>
0x80	POBAR5—PCI outbound base address register 5	R/W	0x0000_0000	<a href="#">11.4.2/11-3</a>
0x88	POCMR5—PCI outbound comparison mask register 5	R/W	0x0000_0000	<a href="#">11.4.3/11-4</a>

Table 11-1. Sequencer Memory Map (continued)

Offset	Register	Access	Reset	Section/Page
0xF0	PMCR—Power management control register	R/W	0x0000_0000	11.4.4/11-5
0xF8	DTCR—Discard timer control register	R/W	0x0000_0000	11.4.5/11-6

## 11.4 Register Descriptions

### 11.4.1 PCI Outbound Translation Address Registers (POTAR<sub>n</sub>)

The PCI outbound translation address register defines the location of the outbound translation window in the PCI (translated) address space. Figure 11-2 shows the POTAR<sub>n</sub> register fields.


Figure 11-2. PCI Outbound Translation Address Registers (POTAR<sub>n</sub>)

Table 11-2 describes POTAR<sub>n</sub> fields.

Table 11-2. POTAR<sub>n</sub> Field Descriptions

Bits	Name	Description
0–11	—	Reserved
12–31	TA	Translation address. Contains the starting address of the outbound translated address. It also corresponds to the most-significant 20 bits of a 32-bit address. The translation address must be aligned based on the window's size.

### 11.4.2 PCI Outbound Base Address Registers (POBAR<sub>n</sub>)

The PCI outbound base address register (POBAR<sub>n</sub>) defines the location of the outbound translation window in the local (source) memory space. Figure 11-3 shows the POBAR<sub>n</sub> register fields.


Figure 11-3. PCI Outbound Base Address Registers (POBAR<sub>n</sub>)


Table 11-3 describes POBAR<sub>n</sub> fields.

**Table 11-3. POBAR<sub>n</sub> Field Descriptions**

Bits	Name	Description
0–11	—	Reserved
12–31	BA	Base address. This field contains the starting address of the outbound translated window. This field corresponds to the most-significant 20 bits of a 32-bit address.

### 11.4.3 PCI Outbound Comparison Mask Registers (POCMR<sub>n</sub>)

The PCI outbound comparison mask register (POCMR<sub>n</sub>) defines the size and destination of the outbound translation window. It also defines some properties of the window in the PCI address space. See Section 11.5.1, “Transaction Forwarding,” for more information. Figure 11-4 shows the POCMR<sub>n</sub> register fields.



**Figure 11-4. PCI Outbound Comparison Mask Registers (POCMR<sub>n</sub>)**

Table 11-4 describes the bit settings of the POCMR<sub>n</sub> register.

**Table 11-4. POCMR<sub>n</sub> Field Descriptions**

Bits	Name	Description
0	EN	Enable. Enables the address translation window. 0 Address translation is disabled for this window. 1 Address translation is enabled for this window. Local addresses that match the definition of the window will be recognized by the device and translated to the PCI memory space.
1	IO	I/O space. Determines whether the window is mapped to the PCI memory space or PCI I/O space. 0 Memory space 1 I/O space
2–11	—	Reserved, should be cleared.

Table 11-4. POCMR<sub>n</sub> Field Descriptions (continued)

Bits	Name	Description																																																
12–31	CM	Comparison mask. Contains the size of the translation window. The bits that are 1 in this field indicate bits of the transaction address that should be matched to the value in the PCI outbound base address register and translated in case of a match. This field corresponds to the most-significant 20 bits of a 32-bit address.																																																
		<table border="1"> <thead> <tr> <th>Value</th> <th>Meaning</th> <th>Value</th> <th>Meaning</th> </tr> </thead> <tbody> <tr> <td>0000_0000_0000_0000_0000</td> <td>4 GB</td> <td>1111_1111_1110_0000_0000</td> <td>2 MB</td> </tr> <tr> <td>1000_0000_0000_0000_0000</td> <td>2 GB</td> <td>1111_1111_1111_0000_0000</td> <td>1 MB</td> </tr> <tr> <td>1100_0000_0000_0000_0000</td> <td>1 GB</td> <td>1111_1111_1111_1000_0000</td> <td>512 KB</td> </tr> <tr> <td>1110_0000_0000_0000_0000</td> <td>512 MB</td> <td>1111_1111_1111_1100_0000</td> <td>256 KB</td> </tr> <tr> <td>1111_0000_0000_0000_0000</td> <td>256 MB</td> <td>1111_1111_1111_1110_0000</td> <td>128 KB</td> </tr> <tr> <td>1111_1000_0000_0000_0000</td> <td>128 MB</td> <td>1111_1111_1111_1111_0000</td> <td>64 KB</td> </tr> <tr> <td>1111_1100_0000_0000_0000</td> <td>64 MB</td> <td>1111_1111_1111_1111_1000</td> <td>32 KB</td> </tr> <tr> <td>1111_1110_0000_0000_0000</td> <td>32 MB</td> <td>1111_1111_1111_1111_1100</td> <td>16 KB</td> </tr> <tr> <td>1111_1111_0000_0000_0000</td> <td>16 MB</td> <td>1111_1111_1111_1111_1110</td> <td>8 KB</td> </tr> <tr> <td>1111_1111_1000_0000_0000</td> <td>8 MB</td> <td>1111_1111_1111_1111_1111</td> <td>4 KB</td> </tr> <tr> <td>1111_1111_1100_0000_0000</td> <td>4 MB</td> <td></td> <td></td> </tr> </tbody> </table>	Value	Meaning	Value	Meaning	0000_0000_0000_0000_0000	4 GB	1111_1111_1110_0000_0000	2 MB	1000_0000_0000_0000_0000	2 GB	1111_1111_1111_0000_0000	1 MB	1100_0000_0000_0000_0000	1 GB	1111_1111_1111_1000_0000	512 KB	1110_0000_0000_0000_0000	512 MB	1111_1111_1111_1100_0000	256 KB	1111_0000_0000_0000_0000	256 MB	1111_1111_1111_1110_0000	128 KB	1111_1000_0000_0000_0000	128 MB	1111_1111_1111_1111_0000	64 KB	1111_1100_0000_0000_0000	64 MB	1111_1111_1111_1111_1000	32 KB	1111_1110_0000_0000_0000	32 MB	1111_1111_1111_1111_1100	16 KB	1111_1111_0000_0000_0000	16 MB	1111_1111_1111_1111_1110	8 KB	1111_1111_1000_0000_0000	8 MB	1111_1111_1111_1111_1111	4 KB	1111_1111_1100_0000_0000	4 MB		
Value	Meaning	Value	Meaning																																															
0000_0000_0000_0000_0000	4 GB	1111_1111_1110_0000_0000	2 MB																																															
1000_0000_0000_0000_0000	2 GB	1111_1111_1111_0000_0000	1 MB																																															
1100_0000_0000_0000_0000	1 GB	1111_1111_1111_1000_0000	512 KB																																															
1110_0000_0000_0000_0000	512 MB	1111_1111_1111_1100_0000	256 KB																																															
1111_0000_0000_0000_0000	256 MB	1111_1111_1111_1110_0000	128 KB																																															
1111_1000_0000_0000_0000	128 MB	1111_1111_1111_1111_0000	64 KB																																															
1111_1100_0000_0000_0000	64 MB	1111_1111_1111_1111_1000	32 KB																																															
1111_1110_0000_0000_0000	32 MB	1111_1111_1111_1111_1100	16 KB																																															
1111_1111_0000_0000_0000	16 MB	1111_1111_1111_1111_1110	8 KB																																															
1111_1111_1000_0000_0000	8 MB	1111_1111_1111_1111_1111	4 KB																																															
1111_1111_1100_0000_0000	4 MB																																																	
		All others values are Reserved																																																

#### 11.4.4 Power Management Control Register (PMCR)

PMCR provides control of system-level low-power modes. Figure 11-5 shows the PMCR register fields.




Figure 11-5. Power Management Control Register (PMCR)

Table 11-5 describes PMCR fields.

Table 11-5. PMCR Field Descriptions

Bits	Name	Description
0–30	—	Reserved
31	IDLE	This read-only bit indicates that the logic controlled by the power management function is idle. This bit could be used as an indication that it is okay to disable the clocks to this complex. 0 Logic is active. 1 Logic is idle. There are no outstanding transactions in the IOS, the DMA, or the PCI port.

## 11.4.5 Discard Timer Control Register (DTCR)

DTCR configures the discard timer, which is used to place a time limit on PCI delayed read transactions from non-prefetchable memory. Although prefetched reads may be discarded whenever the IOS is full and needs to allocate another buffer, other delayed reads must not be discarded until the originator actually receives the data. The DTCR is used to release stuck buffers in case of malfunctioning or disconnected masters that never come back to read the data they requested

Figure 11-6 shows the DTCR register fields.




Figure 11-6. Discard Timer Control Register (DTCR)

Table 11-6 describes DTCR fields.

Table 11-6. DTCR Field Descriptions

Bits	Name	Description
0	EN	Enable. This bit enables the discard timer. 0 Disabled 1 Enabled
1–7	—	Reserved
8–31	PTV	Preload timer value (PTV). This field contains the preload value for the discard timer. PCI delayed reads from non-prefetchable address space are discarded after $(2^{24} - \text{PTV})$ internal clock cycles if the master has not repeated the transaction. 0xFFFFFFFF is not valid for PTV. For example, to discard a delayed completion if the PCI master has not repeated the transaction in $2^{15}$ PCI clocks, <ul style="list-style-type: none"> <li>Assuming the internal frequency is twice the PCI frequency</li> <li>The PTV should equal <math>2^{24} - 2^{16}</math> (0xFF0000).</li> </ul>

## 11.5 Functional Description

The IOS is a four-port switch with buffering. Each port has master and slave interfaces. When a port masters a transaction, the transaction attributes are stored in a buffer and the IOS generates a transaction to the slave interface of the destination port. The data is also buffered between the ports. The IOS contains 8 cache line (32-byte) transaction buffers, some of which are reserved for specific types of transactions.

The address and data phases of the transactions are independent. The data phases of the transactions are not required to be in order.

### 11.5.1 Transaction Forwarding

Although the ports use a similar interface, the I/O sequencer is not actually symmetrical. The transaction forwarding from each source is explained in the following sections.

### 11.5.1.1 Transactions from the Coherency System Bus (CSB) Port

Transactions from the CSB port are forwarded as follows:

- If the address matches the 12-byte PCI controller software configuration memory space of the PCI controller, the transaction is forwarded to the PCI port. These address values are configuration options of the I/O sequencer. See [Table 13-3](#) for more information on PCI controller software configuration memory space.
- If the address matches the DMA register memory space, the transaction is forwarded to the DMA port.
- If the address hits any of the outbound translation windows, the transaction is forwarded to the PCI port, with the address translated. See [Section 11.5.2, “PCI Outbound Address Translation,”](#) for more information.

### 11.5.1.2 Transactions from the PCI Port

Transactions from the PCI port are forwarded as follows:

- If the address matches the DMA register memory space, the transaction is forwarded to the DMA port.
- All other transactions are forwarded to the CSB port.

### 11.5.1.3 Transactions from the DMA Port

Transactions from the DMA port are forwarded as follows:

- If the address hits any of the outbound translation windows, the transaction is forwarded to the PCI port, with the address translated. See [Section 11.5.2, “PCI Outbound Address Translation,”](#) for more information.
- All other transactions are forwarded to the CSB port.

## 11.5.2 PCI Outbound Address Translation

Outbound address translation is provided to allow the outbound transactions to access any address over the PCI memory or I/O space. Translation window base addresses are defined in the PCI outbound base address registers. See [Section 11.4.2, “PCI Outbound Base Address Registers \(POBARn\),”](#) for more information. Transactions to these address ranges are issued on the PCI bus with a translated address. The translation addresses are defined in the associated PCI outbound translation address registers (POTARs).

Figure 11-7 shows an example translation window for outbound memory accesses.




Figure 11-7. Outbound PCI Memory Address Translation

The six sets of outbound translation registers allow six simultaneous translation windows to the PCI port. Software can move and adjust the memory window translations and sizes during run-time. This allows software to access different PCI memory/IO spaces on-the-fly, but the PCI outbound translation source windows must not overlap. However, outbound translation destination windows can be overlapped.

### 11.5.3 Transaction Ordering

The following rules are applied to maintain proper ordering of transactions:

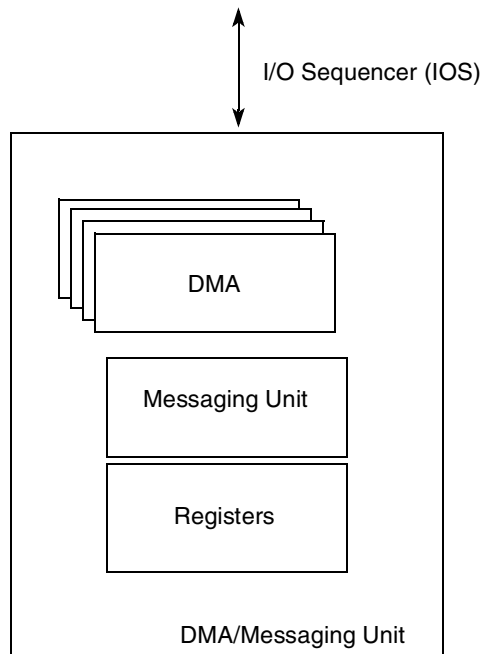
- The transactions arriving from each port are dispatched to the destination port in the order of arrival. The dispatch order of transactions arriving on different ports is not necessarily maintained.
- A read transaction that originates at the CSB port and reads from the PCI port pulls out of the IOS any posted writes that originated on the PCI port and were posted before the read data arrives from the PCI.
- The IOS can always accept a write from the PCI port without forcing the PCI port to first accept a read.



## Chapter 12

# DMA/Messaging Unit

The DMA/messaging unit supports communication between two processors on different buses, for example, a local processor and a processor on a PCI bus. This unit operates with generic messages and doorbell registers. [Figure 12-1](#) is a block diagram of the DMA/messaging unit.



**Figure 12-1. DMA/Messaging Unit Block Diagram**

This block also provides a DMA controller, which transfers blocks of data independent of the local processor or PCI hosts. The DMA module has four high-speed DMA channels, which share buffer space in the I/O sequencer (IOS) to facilitate the gathering and sending of data.

## 12.1 Features

The DMA/messaging unit includes the following features:

- Message and doorbell registers for inter-processor communication
- DMA controller
  - Four DMA channels
  - Concurrent execution across multiple channels with programmable bandwidth control
  - Misaligned transfer capability

- Data chaining and direct mode
- Interrupt on completed segment, chain, and error

## 12.2 Memory Map/Register Definition

Table 12-1 lists the address and access of the memory map module.

**Table 12-1. Module Memory Map**

Offset	Register	Access	Reset	Section/Page
0x0_8030	OMISR—Outbound message interrupt status register	Mixed	0x0000_0000	<a href="#">12.3.1/12-3</a>
0x0_8034	OMIMR—Outbound message interrupt mask register	R/W	0x0000_0000	<a href="#">12.3.2/12-4</a>
0x0_8050	IMR0—Inbound message register 0	R/W	0x0000_0000	<a href="#">12.3.3/12-5</a>
0x0_8054	IMR1—Inbound message register 1	R/W	0x0000_0000	<a href="#">12.3.3/12-5</a>
0x0_8058	OMR0—Outbound message register 0	R/W	0x0000_0000	<a href="#">12.3.4/12-5</a>
0x0_805C	OMR1—Outbound message register 1	R/W	0x0000_0000	<a href="#">12.3.4/12-5</a>
0x0_8060	ODR—Outbound doorbell register	R/W	0x0000_0000	<a href="#">12.3.5/12-6</a>
0x0_8068	IDR—Inbound doorbell register	R/W	0x0000_0000	<a href="#">12.3.5/12-6</a>
0x0_8080	IMISR—Inbound message interrupt status register	Mixed	0x0000_0000	<a href="#">12.3.6/12-7</a>
0x0_8084	IMIMR—Inbound message interrupt mask register	R/W	0x0000_0000	<a href="#">12.3.7/12-8</a>
0x0_8100	DMAMR0—DMA 0 mode register	R/W	0x0000_0000	<a href="#">12.3.8.1/12-9</a>
0x0_8104	DMASR0—DMA 0 status register	R/W	0x0000_0000	<a href="#">12.3.8.2/12-11</a>
0x0_8108	DMACDAR0—DMA 0 current descriptor address register	R/W	0x0000_0000	<a href="#">12.3.8.3/12-12</a>
0x0_8110	DMASAR0—DMA 0 source address register	R/W	0x0000_0000	<a href="#">12.3.8.4/12-13</a>
0x0_8118	DMADAR0—DMA 0 destination address register	R/W	0x0000_0000	<a href="#">12.3.8.5/12-13</a>
0x0_8120	DMABCR0—DMA 0 byte count register	R/W	0x0000_0000	<a href="#">12.3.8.6/12-14</a>
0x0_8124	DMANDAR0—DMA 0 next descriptor address register	R/W	0x0000_0000	<a href="#">12.3.8.7/12-14</a>
0x0_8180	DMAMR1—DMA 1 mode register	R/W	0x0000_0000	<a href="#">12.3.8.1/12-9</a>
0x0_8184	DMASR1—DMA 1 status register	R/W	0x0000_0000	<a href="#">12.3.8.2/12-11</a>
0x0_8188	DMACDAR1—DMA 1 current descriptor address register	R/W	0x0000_0000	<a href="#">12.3.8.3/12-12</a>
0x0_8190	DMASAR1—DMA 1 source address register	R/W	0x0000_0000	<a href="#">12.3.8.4/12-13</a>
0x0_8198	DMADAR1—DMA 1 destination address register	R/W	0x0000_0000	<a href="#">12.3.8.5/12-13</a>
0x0_81A0	DMABCR1—DMA 1 byte count register	R/W	0x0000_0000	<a href="#">12.3.8.6/12-14</a>
0x0_81A4	DMANDAR1—DMA 1 next descriptor address register	R/W	0x0000_0000	<a href="#">12.3.8.7/12-14</a>
0x0_8200	DMAMR2—DMA 2 mode register	R/W	0x0000_0000	<a href="#">12.3.8.1/12-9</a>
0x0_8204	DMASR2—DMA 2 status register	R/W	0x0000_0000	<a href="#">12.3.8.2/12-11</a>
0x0_8208	DMACDAR2—DMA 2 current descriptor address register	R/W	0x0000_0000	<a href="#">12.3.8.3/12-12</a>
0x0_8210	DMASAR2—DMA 2 source address register	R/W	0x0000_0000	<a href="#">12.3.8.4/12-13</a>

Table 12-1. Module Memory Map (continued)

Offset	Register	Access	Reset	Section/Page
0x0_8218	DMADAR2—DMA 2 destination address register	R/W	0x0000_0000	<a href="#">12.3.8.5/12-13</a>
0x0_8220	DMABCR2—DMA 2 byte count register	R/W	0x0000_0000	<a href="#">12.3.8.6/12-14</a>
0x0_8224	DMANDAR2—DMA 2 next descriptor address register	R/W	0x0000_0000	<a href="#">12.3.8.7/12-14</a>
0x0_8280	DMAMR3—DMA 3 mode register	R/W	0x0000_0000	<a href="#">12.3.8.1/12-9</a>
0x0_8284	DMASR3—DMA 3 status register	R/W	0x0000_0000	<a href="#">12.3.8.2/12-11</a>
0x0_8288	DMACDAR3—DMA 3 current descriptor address register	R/W	0x0000_0000	<a href="#">12.3.8.3/12-12</a>
0x0_8290	DMASAR3—DMA 3 source address register	R/W	0x0000_0000	<a href="#">12.3.8.4/12-13</a>
0x0_8298	DMADAR3—DMA 3 destination address register	R/W	0x0000_0000	<a href="#">12.3.8.5/12-13</a>
0x0_82A0	DMABCR3—DMA 3 byte count register	R/W	0x0000_0000	<a href="#">12.3.8.6/12-14</a>
0x0_82A4	DMANDAR3—DMA 3 next descriptor address register	R/W	0x0000_0000	<a href="#">12.3.8.7/12-14</a>
0x0_82A8	DMAGSR—DMA general status register	R	0x0000_0000	<a href="#">12.3.8.8/12-15</a>
0x0_82B0– 0x0_82FF	Reserved	—	—	—

## 12.3 Register Descriptions

The following sections describe the DMA/messaging unit configuration, control, and status registers.

### NOTE

The registers described in this section use little-endian byte ordering. Software running on the local processor in big-endian mode must byte-swap the data. No byte swapping occurs when the registers are accessed from the PCI bus.

### 12.3.1 Outbound Message Interrupt Status Register (OMISR)

OMISR contains the interrupt status of the doorbell and outbound message registers. A PCI device acknowledges the outbound message interrupt by writing a 1 to the appropriate status bit: OMISR[OM1I] or OMISR[OM0I]. Setting one of these bits clears both the interrupt and the corresponding status bit. The local processor provokes an outbound message interrupt by writing to either of the two outbound message registers: OMR0 or OMR1. OMISR can be accessed from the CSB or the PCI bus, but it is normally accessed only from the PCI bus. [Figure 12-2](#) shows the OMISR fields.




Figure 12-2. Outbound Message Interrupt Status Register (OMISR)


Table 12-2 describes the OMISR register.

**Table 12-2. OMISR Field Descriptions**

Bits	Name	Description
31–4	—	Reserved
3	ODI	Outbound doorbell interrupt. This read-only bit indicates the status of the ODR bits. It is masked by OMIMR[ODIM]. 0 No outbound doorbell interrupt. 1 There is an outbound doorbell interrupt.
2	—	Reserved
1	OM1I	Outbound message 1 interrupt. When set, indicates that there is an outbound message 1 interrupt. Write 1 to this position to clear this bit. 0 No outbound message 1 interrupt. 1 There is an outbound message 1 interrupt.
0	OM0I	Outbound message 0 interrupt. When set, indicates that there is an outbound message 0 interrupt. Write 1 to this position to clear this bit. 0 No outbound message 0 interrupt. 1 There is an outbound message 0 interrupt.

### 12.3.2 Outbound Message Interrupt Mask Register (OMIMR)

OMIMR contains the interrupt mask of the doorbell and message register events generated by the local processor. OMIMR can be read from the CSB or the PCI bus, but it can be written only from the PCI bus. Figure 12-3 shows the OMIMR.



**Figure 12-3. Outbound Message Interrupt Mask Register (OMIMR)**

Table 12-3 describes the OMIMR register.

**Table 12-3. OMIMR Field Descriptions**

Bits	Name	Description
31–4	—	Reserved
3	ODIM	Outbound doorbell interrupt mask. 0 Outbound doorbell interrupt is allowed 1 Outbound doorbell interrupt is masked
2	—	Reserved

Table 12-3. OMIMR Field Descriptions (continued)

Bits	Name	Description
1	OM1IM	Outbound message 1 interrupt mask. 0 Outbound message 1 interrupt is allowed 1 Outbound message 1 interrupt is masked
0	OM0IM	Outbound message 0 interrupt mask. 0 Outbound message 0 interrupt is allowed 1 Outbound message 0 interrupt is masked

### 12.3.3 Inbound Message Registers (IMR0–IMR1)

The inbound message registers can be read from the PCI bus and the CSB in both host and agent modes. They can be written only from the PCI bus. Figure 12-4 shows the IMR0 and IMR1 fields.



Figure 12-4. Inbound Message Registers (IMR0, IMR1)

Table 12-4 describes the IMR<sub>n</sub> register.

Table 12-4. IMR0 and IMR1 Field Descriptions

Bits	Name	Description
31–0	IMSG <sub>n</sub>	Inbound message <i>n</i> . Contains generic data to be passed between the local processor and external hosts.

### 12.3.4 Outbound Message Registers (OMR0–OMR1)

The outbound message registers can be read from the PCI bus and the CSB in both host and agent modes. They can be written only from the CSB.




Figure 12-5. Outbound Message Registers (OMR0–OMR1)

Table 12-5 describes the OMR<sub>n</sub> registers.

Table 12-5. OMR0 and OMR1 Field Descriptions

Bits	Name	Description
31–0	OMSG <sub>n</sub>	Outbound message <i>n</i> . Contains generic data to be passed between the local processor and external hosts.

## 12.3.5 Doorbell Registers

The following sections describe the outbound and inbound doorbell registers.

### 12.3.5.1 Outbound Doorbell Register (ODR)

ODR is accessible from the PCI bus and the CSB in both host and agent modes. Figure 12-6 shows the ODR<sub>n</sub> fields.




Figure 12-6. Outbound Doorbell Register (ODR)

Table 12-6 describes the ODR registers.

Table 12-6. ODR Field Descriptions

Bits	Name	Description
31–29	—	Reserved
28–0	ODR <sub>n</sub>	Outbound doorbell <i>n</i> . Write 1 from the CSB to set. Write 1 from the PCI bus to clear. Writing 0 has no effect. (Writing a bit in this register from the CSB causes an interrupt ( $\overline{\text{PCI\_INTA}}$ ) to be generated.)

### 12.3.5.2 Inbound Doorbell Register (IDR)

IDR is accessible from the PCI bus and the CSB in both host and agent modes. Figure 12-7 shows the IDR fields.

Offset: 0x068

Access: User read/write




Figure 12-7. Inbound Doorbell Register (IDR)

Table 12-7 describes the IDR registers.

Table 12-7. IDR Field Descriptions

Bits	Name	Descriptions
31	IMC	Inbound machine check. Write 1 from the PCI bus to set. Write 1 from the CSB to clear. Writing 0 has no effect. Writing this bit from the PCI bus causes a machine check interrupt to be generated to the local processor.
30–0	IDR $n$	Inbound doorbell $n$ . Write 1 from the PCI bus to set. Write 1 from the CSB to clear. Writing 0 has no effect. Writing a bit in this register from the PCI bus causes an interrupt to be generated to the local processor.

### 12.3.6 Inbound Message Interrupt Status Register (IMISR)

The IMISR contains the interrupt status of the doorbell and message register events. Writing a 1 to IM1I clears the bit. The events are generated by the PCI masters.

Figure 12-8 shows the IMISR fields.

Offset: 0x080

Access: User Mixed




Figure 12-8. Inbound Message Interrupt Status Register (IMISR)


Table 12-8 describes the IMISR register.

**Table 12-8. IMISR Field Descriptions**

Bits	Name	Descriptions
31–5	—	Reserved
4	MCI	Machine check interrupt. Indicates whether a machine check interrupt condition was generated by setting the IDR[31]. The interrupt is cleared by writing a 1 to IDR[IMC] from the CSB. 0 No machine check interrupt 1 There is a machine check interrupt
3	IDI	Inbound doorbell interrupt. Indicates whether an inbound doorbell interrupt occurred. 0 No inbound doorbell interrupt 1 There is an inbound doorbell interrupt
2	—	Reserved
1	IM1I	Inbound message 1 interrupt. Indicates whether an inbound message 1 interrupt occurred. Write 1 to this position to clear this bit. 0 No inbound message 1 interrupt. 1 There is an inbound message 1 interrupt.
0	IM0I	Inbound message 0 interrupt. Indicates whether an inbound message 0 interrupt occurred. Write 1 to this position to clear this bit. 0 No inbound message 0 interrupt. 1 There is an inbound message 0 interrupt.

### 12.3.7 Inbound Message Interrupt Mask Register (IMIMR)

This register contains the interrupt mask of the doorbell and message register events generated by the PCI master. Figure 12-9 shows the IMIMR fields.



**Figure 12-9. Inbound Message Interrupt Mask Register (IMIMR)**

Table 12-9 describes the IMISR register.

**Table 12-9. IMIMR Field Descriptions**

Bits	Name	Description
31–5	—	Reserved
4	MCIM	Machine check interrupt mask. 0 Machine check interrupt from the IDR is allowed 1 Machine check interrupt is masked. IMISR[MC1] is cleared
3	IDIM	Inbound doorbell interrupt mask. 0 Inbound doorbell interrupt is allowed 1 Inbound doorbell interrupt is masked. IMISR[IDI] is cleared.
2	—	Reserved



Table 12-9. IMIMR Field Descriptions (continued)

Bits	Name	Description
1	IM1IM	Inbound message 1 interrupt mask. 0 Inbound message 1 interrupt is allowed 1 Inbound message 1 interrupt is masked. IMISR[IM1] is cleared
0	IM0IM	Inbound message 0 interrupt mask. 0 Inbound message 0 interrupt is allowed 1 Inbound message 0 interrupt is masked. IMISR[IM0] is cleared

## 12.3.8 DMA Registers

Each DMA channel has a set of seven 32-bit registers (mode, status, current descriptor address, next descriptor address, source address, destination address, and byte count) to support transactions. The following sections describe the format of the DMA support registers.

### 12.3.8.1 DMA Mode Register (DMAMR<sub>n</sub>)

This section describes the DMA mode register. The mode register allows software to start the DMA transfer and to control various DMA transfer characteristics. Figure 12-10 shows the DMAMR<sub>n</sub> fields.

Offset: 0x100, 0x180, 0x200, 0x280

Access: User read/write




Figure 12-10. DMA Mode Register (DMAMR<sub>n</sub>)

Table 12-10 describes the DMAMR<sub>n</sub> register.

Table 12-10. DMAMR<sub>n</sub> Field Descriptions

Bits	Name	Description
31–24	—	Reserved
23–21	BWC	Bandwidth control. Only applies when multiple channels are executing transfers concurrently. The field determines how many cache lines a given channel is allowed to transfer after it is granted access to the IOS interface and before it releases the interface to the next channel. This allows the user to prioritize the DMA channels. The BWC values are listed as follows: 000 1 cache line 001 2 cache lines 010 4 cache lines 011 8 cache lines 100 16 cache lines Others Reserved

Table 12-10. DMAMR $n$  Field Descriptions (continued)

Bits	Name	Description
20	DMSEN	Direct mode snoop enable. This bit controls snooping of direct mode DMA transactions. 0 Snooping is disabled 1 Snooping is enabled
19	IRQS	Interrupt steer. This bit determines the destination of the DMA interrupts. 0 All DMA interrupts are routed to the on-chip interrupt controller 1 All DMA interrupts are routed to the PCI bus through $\overline{\text{PCI\_INTA}}$
18	—	Reserved
17–16	DAHTS	Destination address hold transfer size. This field indicates the transfer size used for each transaction when DAHE is 1. The byte count register must be in multiples of the size, and the destination address register must be aligned based on the size. 00 1 byte 01 2 bytes 10 4 bytes 11 8 bytes
15–14	SAHTS	Source address hold transfer size. This field indicates the transfer size used for each transaction when SAHE is 1. The byte count register must be in multiples of the size, and the source address register must be aligned based on the size. 00 1 byte 01 2 bytes 10 4 bytes 11 8 bytes
13	DAHE	Destination address hold enable. This bit allows the DMA controller to hold the destination address constant for every transfer. The size used for transfer is indicated by DAHTS. Note that hardware supports only aligned transfers for this feature. 0 Do not hold the destination address constant 1 Hold the destination address constant <b>Note:</b> The DMA does not support address hold for both the source and the destination at the same transfer.
12	SAHE	Source address hold enable. This bit allows the DMA controller to hold the source address constant for every transfer. The size used for transfer is indicated by SAHTS. Note that hardware supports only aligned transfers for this feature. 0 Do not hold the source address constant 1 Hold the source address constant <b>Note:</b> The DMA does not support address hold for both the source and the destination at the same transfer.
11–10	PRC	PCI read command. This field indicates the type of PCI read command to use. 00 Reserved 01 PCI read line 10 PCI read multiple 11 Reserved
9–8	—	Reserved
7	EOTIE	End-of-transfer interrupt enable. This bit determines whether an interrupt is generated at the completion of a DMA transfer. End-of-transfer is defined as the end of a direct mode transfer or in chaining mode, as the end of the transfer of the last segment of a chain. 0 No EOT interrupt is generated 1 EOT interrupt is generated
6–4	—	Reserved

Table 12-10. DMAMR $n$  Field Descriptions (continued)

Bits	Name	Description
3	TEM	Transfer error mask. This bit determines the DMA response in the event of a transfer error. 0 The DMA will halt when a transfer error occurs. 1 The DMA will complete the transfer regardless of whether a transfer error occurs. <b>Note:</b> Regardless of the setting of TEM, if an error condition was detected during the DMA transfer, it will cause DMASR $n$ [TE] to be set.
2	CTM	Channel transfer mode. 0 Chaining mode 1 Direct mode
1	CC	Channel continue. This bit applies only to chaining mode. Setting this bit indicates that the current descriptor segment should be repeated. CC is cleared by the DMA once the repeat takes effect, so it only causes a single repeat. 0 Normal chaining 1 DMACDAR is not loaded from DMANDAR, causing a repeat of the current descriptor segment
0	CS	Channel start. A 0-to-1 transition occurring on this bit when the channel is not busy (SR[CB] bit is 0) will start the DMA process. If the channel is busy and a 0-to-1 transition occurs, the DMA channel will restart from a previous halt condition. A 1-to-0 transition when the channel is busy (CB bit is 1) will halt the DMA process. Nothing happens if the channel is not busy and a 1-to-0 transition occurs. This bit is cleared by the DMA at the end of a transfer.

### 12.3.8.2 DMA Status Register (DMASR $n$ )

This section describes the DMA status register. The status register reports various DMA conditions during and after the DMA transfer. Writing a 1 to a specific set bit clears the bit. Figure 12-11 shows the DMASR $n$  fields.


Figure 12-11. DMA Status Register (DMASR $n$ )

Table 12-11 describes the DMASR $n$  register.

Table 12-11. DMASR $n$  Field Descriptions

Bits	Name	Description
31–8	—	Reserved
7	TE	Transfer error. Set when there is an error condition during the DMA transfer.
6–3	—	Reserved
2	CB	Channel busy. This bit indicates whether the channel is busy. It is cleared as a result of any of the following conditions: an error or completion of the DMA transfer. 0 No DMA transfer is currently in progress 1 A DMA transfer is currently in progress

Table 12-11. DMASR<sub>n</sub> Field Descriptions (continued)

Bits	Name	Description
1	EOSI	End-of-segment interrupt. After transferring a segment of data, if the DMACDAR <sub>n</sub> [EOSIE] bit in the current descriptor address register is set, this bit is set and an interrupt is generated.
0	EOCDI	End-of-chain/direct interrupt. When the last DMA transfer is finished, either in chaining or direct mode, if DMAMR[EOTIE] is set, this bit is set and an interrupt is generated.

### 12.3.8.3 DMA Current Descriptor Address Register (DMACDAR<sub>n</sub>)

DMACDAR<sub>n</sub> contains the address of the current segment descriptor being transferred. In chaining mode, software must initialize this register to point to the first descriptor in the chain. After processing the first descriptor, the DMA controller moves the contents of the next descriptor address register into DMACDAR, loads the following descriptor into DMANDAR, and executes the current transfer.

Figure 12-12 shows the DMACDAR<sub>n</sub> fields.


Figure 12-12. DMA Current Descriptor Address Register (DMACDAR<sub>n</sub>)


Table 12-12 describes the DMACDAR<sub>n</sub> register.

Table 12-12. DMACDAR<sub>n</sub> Field Descriptions

Bits	Name	Description
31–5	CDA	Current descriptor address. This field contains the current descriptor address of the segment descriptor in memory. It must be aligned on an 8-word boundary.
4	SNEN	Snoop enable. 0 Snooping is disabled on DMA transactions of the current segment. 1 Snooping is enabled on DMA transactions of the current segment.
3	EOSIE	End-of-segment interrupt enable 0 No end-of-segment interrupt is generated. 1 An interrupt is generated when the current DMA transfer for the current descriptor is finished.
2–0	—	Reserved

### 12.3.8.4 DMA Source Address Register (DMASAR<sub>n</sub>)

DMASAR<sub>n</sub> indicates the address from which the DMA controller will be reading data. The software must ensure that this is a valid memory address. Figure 12-13 shows the DMASAR<sub>n</sub>.



**Figure 12-13. DMA Source Address Register (DMASAR<sub>n</sub>)**


Table 12-13 describes the DMASAR<sub>n</sub> register.

**Table 12-13. DMASAR<sub>n</sub> Field Descriptions**

Bits	Name	Description
31–0	SA	Source address of DMA transfer. The content of this field is updated after each DMA read operation.

### 12.3.8.5 DMA Destination Address Register (DMADAR<sub>n</sub>)

DMADAR<sub>n</sub> indicates the address to which the DMA controller will be writing data. The software must ensure that this is a valid memory address. Figure 12-14 shows the DMADAR<sub>n</sub> fields.



**Figure 12-14. DMA Destination Address Register (DMADAR<sub>n</sub>)**

Table 12-14 describes the DMADAR<sub>n</sub> register.

**Table 12-14. DMASAR<sub>n</sub> Field Descriptions**

Bits	Name	Description
31–0	DA	Destination address of DMA transfer. Updated after each DMA write operation.

### 12.3.8.6 DMA Byte Count Register (DMABCR $n$ )

DMABCR $n$  contains the number of bytes per transfer (maximum transfer size is 64 Mbytes). Figure 12-15 shows the DMABCR $n$ .




Figure 12-15. DMA Byte Count Register (DMABCR $n$ )

Table 12-15 describes the DMABCR $n$  register.

Table 12-15. DMABCR $n$  Field Descriptions

Bits	Name	Description
31–26	—	Reserved
25–0	BC	Byte count. This field contains the number of bytes to transfer. The value in this register is decremented after each DMA read operation. Maximum transfer size is 64 Mbytes.

### 12.3.8.7 DMA Next Descriptor Address Register (DMANDAR $n$ )

DMANDAR $n$  contains the address for the next segment descriptor in the chain. In chaining mode, this register is loaded from the ‘next descriptor’ field of the descriptor to which the current descriptor address register is pointing. Figure 12-16 shows the DMANDAR $n$ .




Figure 12-16. DMA Next Descriptor Address Register (DMANDAR $n$ )

Table 12-16 describes the DMANDAR $n$  register.

Table 12-16. DMANDAR $n$  Field Descriptions

Bits	Name	Descriptions
31–5	NDA	Next descriptor address. This field contains the next descriptor address of the next segment descriptor in memory. It must be aligned on an 8-word boundary.
4	NSNEN	Next snoop enable. 0 Snooping is disabled on DMA transactions. 1 Snooping is enabled on DMA transactions.
3	NEOSIE	Next end-of-segment interrupt enable. 0 No end-of-segment interrupt is generated. 1 An interrupt is generated when the DMA transfer for the next descriptor is finished.

Table 12-16. DMANDAR<sub>n</sub> Field Descriptions (continued)

Bits	Name	Descriptions
2–1	—	Reserved
0	EOTD	End-of-transfer descriptor. 0 This descriptor contains a link to another descriptor. 1 This descriptor is the last to be executed.

### 12.3.8.8 DMA General Status Register (DMAGSR)

DMAGSR provides faster access to the status bits by combining the status bits of all of the DMA channels into one register. Each byte of this register provides the value of bits 7–0 of a channel’s DMA status register. These bits are cleared by writing to the individual DMA status registers. Figure 12-17 shows the DMAGSR fields.




Figure 12-17. DMA General Status Register (DMAGSR)

## 12.4 Functional Description

### 12.4.1 Message Unit

An embedded processor is often part of a larger system containing many processors and distributed memory. These processors tend to work on tasks independent of the host and other peripheral processors in the system. Because of the independent nature of the tasks, it is necessary to provide a communication mechanism between the peripheral processors and the rest of the system. One such method is the use of messages. This block provides a messaging unit to further facilitate communications between host and peripheral. The message unit uses generic messages and doorbell registers.

#### 12.4.1.1 Messaging Registers (IMR0–IMR1, OMR0–OMR1)

There are two 32-bit inbound message registers (IMR0–IMR1) and two 32-bit outbound message registers (OMR0–OMR1). IMR0 and IMR1 allow a remote host or PCI master to write a 32-bit value that, in turn, causes an interrupt request to the on-chip interrupt controller that drives an interrupt line to the local processor. OMR0 and OMR1 allow the local processor to write an outbound message which, in turn, causes the outbound interrupt signal  $\overline{\text{PCI\_INTA}}$  to assert.

The interrupt to the local processor is cleared by writing 1 to the appropriate IMISR bit. The interrupt to PCI ( $\overline{\text{PCI\_INTA}}$ ) is cleared by writing 1 to the appropriate OMISR bit.

### 12.4.1.2 Doorbell Registers (IDR and ODR)

This block contains the inbound doorbell register (IDR) and the outbound doorbell register (ODR). The inbound doorbell allows a remote processor to set a bit in the register from the PCI bus. This, in turn, generates an interrupt request to the on-chip interrupt controller that drives an interrupt line to the local processor. The local processor can write to the ODR, which causes the outbound interrupt signal `PCI_INTA` to assert, thus interrupting the remote processor on the PCI bus.

The interrupt to the local processor is cleared by writing 1 to the appropriate IDR bit. The interrupt to PCI (`PCI_INTA`) is cleared by writing 1 to the appropriate ODR bit.

## 12.4.2 DMA Controller

The DMA controller transfers blocks of data independent of the local processor or PCI hosts. Data movement occurs on the PCI bus and/or CSB. The DMA module has four high-speed DMA channels, which share buffer space in the IOS to facilitate the gathering and sending of data. Both the local processor and PCI masters can initiate a DMA transfer.

Features of the DMA controller include the following:

- Four channels
- Concurrent execution across multiple channels with programmable bandwidth control
- All channels are accessible by local processor and remote PCI masters
- Unaligned transfer capability
- Data chaining and direct mode
- Interrupt on completed segment, chain, and error

Figure 12-18 shows a diagram of the DMA controller in the integrated device.




Figure 12-18. DMA Controller Block Diagram

## 12.4.3 DMA Operation

The DMA controller operates in the following two modes:

- Direct mode, in direct mode, the DMA controller does not read a chain of descriptors from memory but instead uses the current parameters in the DMA registers to start a DMA transfer. The DMA



transfer finishes after all the bytes specified in the byte count register have been transferred. See [Section 12.5.1, “Initialization Steps in Direct Mode,”](#) for more details on initialization steps.

- Chaining mode, in chaining mode, the DMA controller loads descriptors from memory prior to a DMA transfer. The DMA controller begins the transfer according to the descriptor information loaded for each segment. Once the current segment is finished, the DMA controller reads the next descriptor from memory and begins another DMA transfer. The process is finished if the current descriptor is the last one in the chain. See [Section 12.5.2, “Initialization Steps in Chaining Mode,”](#) for more details on initialization steps.

In both modes, setting the start bit in the DMA mode register begins the DMA transfer.

The DMA controller supports unaligned transfers for both the source and destination addresses. It gathers data beginning at the source address and aligns the data accordingly before sending it to the destination address. The DMA controller assumes that the source and destination addresses are valid PCI or CSB memory addresses.

Accesses to CSB memory depend on the alignment of the source and destination addresses and the size of the transfer. The DMA controller transfers a full cache line whenever possible. Misaligned destination addresses result in sub-transfers of less than a cache line on the initial and final beats of the transfer; intermediate beats transfer full cache lines. Configuring a DMA channel for address hold mode  $DMAMR_n$  precludes cache line transfers.

PCI memory read operations depend on the PRC (PCI read command) field in the mode register, the alignment of the source address, and the size of the transfer. The DMA controller attempts to read a full cache line whenever possible. Writing to PCI memory depends on the alignment of the destination address and the size of the transfer.

### 12.4.3.1 DMA Coherency

The four DMA channels use up to four cache lines (128 bytes) of buffer space in the IOS in addition to 16 bytes of local buffer space. Because no address snooping occurs in these internal queues, data posted in these queues is not visible to the rest of the system while a DMA transfer is in progress. It is the responsibility of application software to ensure the coherency of the region being transferred during the DMA process.

Snooping of the CPU or processor data cache is selectable during DMA transactions. A snoop bit is provided in the DMA current descriptor address register ( $DMACDAR_n$ ) and the DMA next descriptor address register ( $DMANDAR_n$ ) that allows software to control when the cache is snooped on a per segment basis.

### 12.4.3.2 Halt and Error Conditions

DMA transfers are halted either by clearing the CS (channel start) bit in the DMA mode register ( $DMAMR_n$ ) or when encountering an error condition. In either case, the application software can do one of the following:

- Continue the DMA transfer
- Reconfigure the DMA for a new transfer

- Leave the channel in the halted state

When a DMA channel is halted, its programming model is completely accessible. If the DMA is halted due to an error condition, the TE (transfer error) bit in the DMA status register (DMASR $n$ ) must be cleared before the transfer can be resumed or a new transfer initiated. Note that the TE bit is not cleared automatically by hardware.

## 12.4.4 DMA Segment Descriptors

DMA segment descriptors contain the source and destination addresses of the data segment, the segment byte count, and a link to the next descriptor. Segment descriptors are built on cache-line (32-byte) boundaries in either CSB or PCI memory and are linked together into chains using the next-descriptor-address field.

**Table 12-17. DMA Segment Descriptor Fields**

Descriptor Field	Description
Source address	Contains the source address of the DMA transfer. After the DMA controller reads the descriptor from memory, this field will be loaded into the DMA source address register (DMASAR $n$ ).
Destination address	Contains the destination address of the DMA transfer. After the DMA controller reads the descriptor from memory, this field will be loaded into the DMA destination address register (DMADAR $n$ ).
Next descriptor address	Points to the next descriptor in memory. After the DMA controller reads the descriptor from memory, this field will be loaded into the DMA next descriptor address register (DMANDAR $n$ ).
Byte count	Contains the number of bytes to transfer. After the DMA controller reads the descriptor from memory, this field will be loaded into the DMA byte count register (DMABCR $n$ ).

Application software initializes the current DMA current descriptor address register (DMACDAR $n$ ) to point to the first descriptor in the chain. For each descriptor in the chain, the DMA controller starts a new DMA transfer with the control parameters specified by the descriptor. The DMA controller traverses the descriptor chain until reaching the last descriptor (with its EOTD bit set).

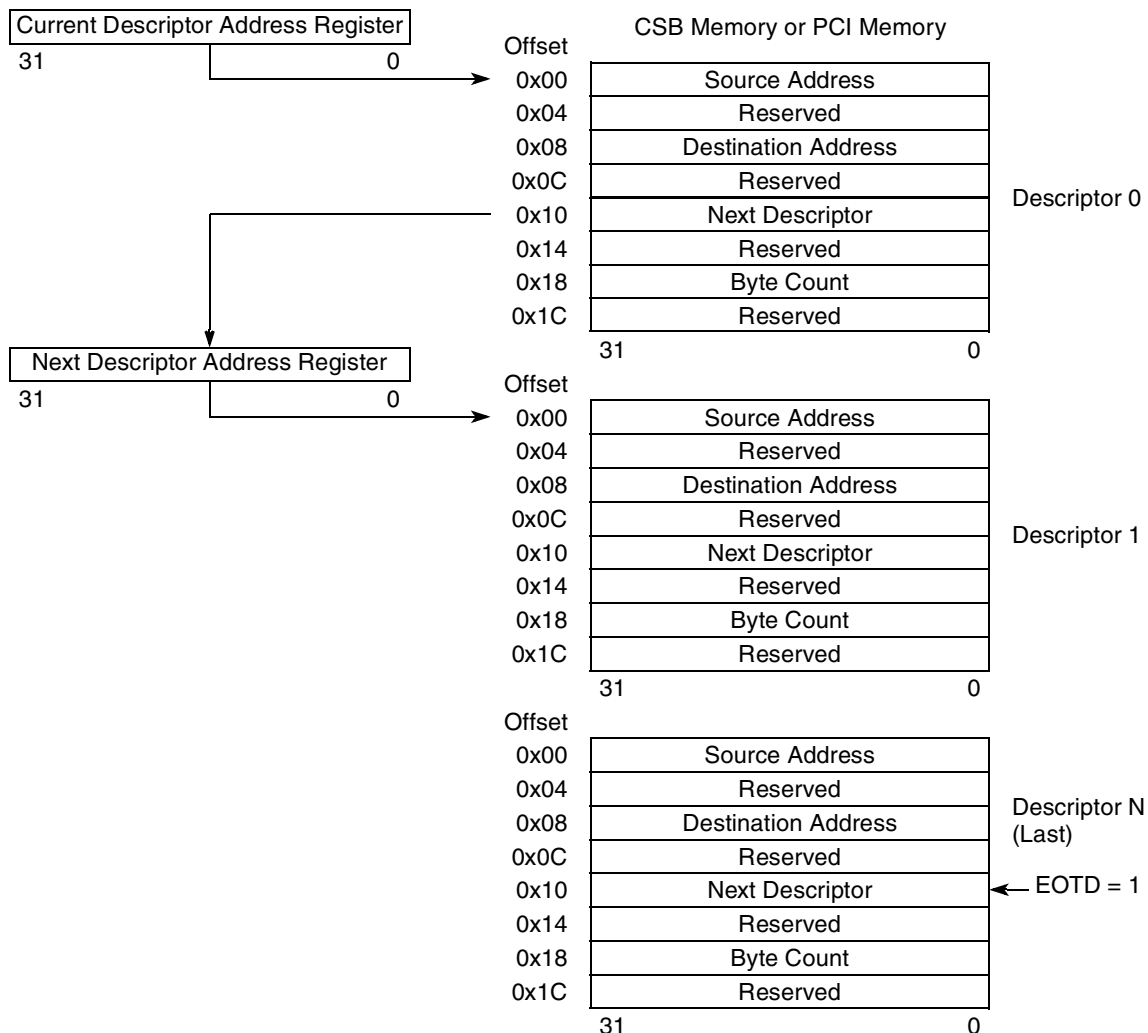


Figure 12-19. DMA Chain of Segment Descriptors

#### 12.4.4.1 Descriptor in Big-Endian Mode

In big-endian mode, the descriptor in CSB memory should be programmed such that data appears in ascending significant-byte order. If segment descriptors are written to memory located in the CSB, they should be treated like they are translated from big-endian to little-endian mode.

**Example:** Big-endian mode descriptor's data structure. Note that the descriptor structure must be aligned on an 8-word boundary.

```
struct {
    double a;          /* 0x1122334455667788 double word*/
    double b;          /* 0x55667788aabbccdd double word*/
    double c;          /* 0x8765432101234567 double word */
    double d;          /* 0x0123456789abcdef double word */
} Descriptor;
Results: Source Address = 0x44332211 <MSB..LSB>
        Destination Address = 0x88776655 <MSB..LSB>
        Next Descriptor Address = 0x21436587 <MSB..LSB>
        Byte Count = 0x67452301 <MSB..LSB>
```

### 12.4.4.2 Descriptor in Little-Endian Mode

In little-endian mode, each segment descriptor should be programmed in descending significant-byte order.

**Example:** Little-endian mode descriptor's data structure. Note that the descriptor structure must be aligned on an 8-word boundary.

```
struct {
    double a;          /* 0x8877665544332211 double word*/
    double b;          /* 0x1122334488776655 double word*/
    double c;          /* 0x7654321012345678 double word */
    double d;          /* 0x0123456776543210 double word */
} Descriptor;
Results: Source Address = 0x44332211 <MSB..LSB>
        Destination Address = 0x88776655 <MSB..LSB>
        Next Descriptor Address = 0x12345678 <MSB..LSB>
        Byte Count = 0x76543210 <MSB..LSB>
```

## 12.5 Initialization/Application Information

### 12.5.1 Initialization Steps in Direct Mode

The initialization steps of a DMA transfer in direct mode are described as follows:

1. Poll the CB (channel busy) bit in the DMA status register (DMASR $n$ ) to make sure the DMA channel is idle.
2. Initialize the DMASAR $n$ , DMADAR $n$ , and the DMABCR $n$ .
3. Initialize DMAMR $n$ [CTM]) to indicate direct mode. Other control parameters in the mode register can also be initialized here if necessary.
4. First clear then set the DMAMR $n$ [CS] to start the DMA transfer.

### 12.5.2 Initialization Steps in Chaining Mode

The initialization steps of a DMA transfer in chaining mode are described as follows:

1. Build a chain of descriptor segments in memory. Refer to [Section 12.4.4, "DMA Segment Descriptors."](#)

2. Poll the  $DMASR_n[CB]$  to make sure the DMA channel is idle.
3. Initialize the  $DMACDAR_n$  to point to the first descriptor in the chain.
4. Initialize the  $DMAMR_n[CTM]$  to indicate chaining mode. Other control parameters in the mode register can also be initialized here if necessary.
5. First clear then set the  $DMAMR_n[CS]$  to start the DMA transfer.



---

## Chapter 13

# PCI Bus Interface

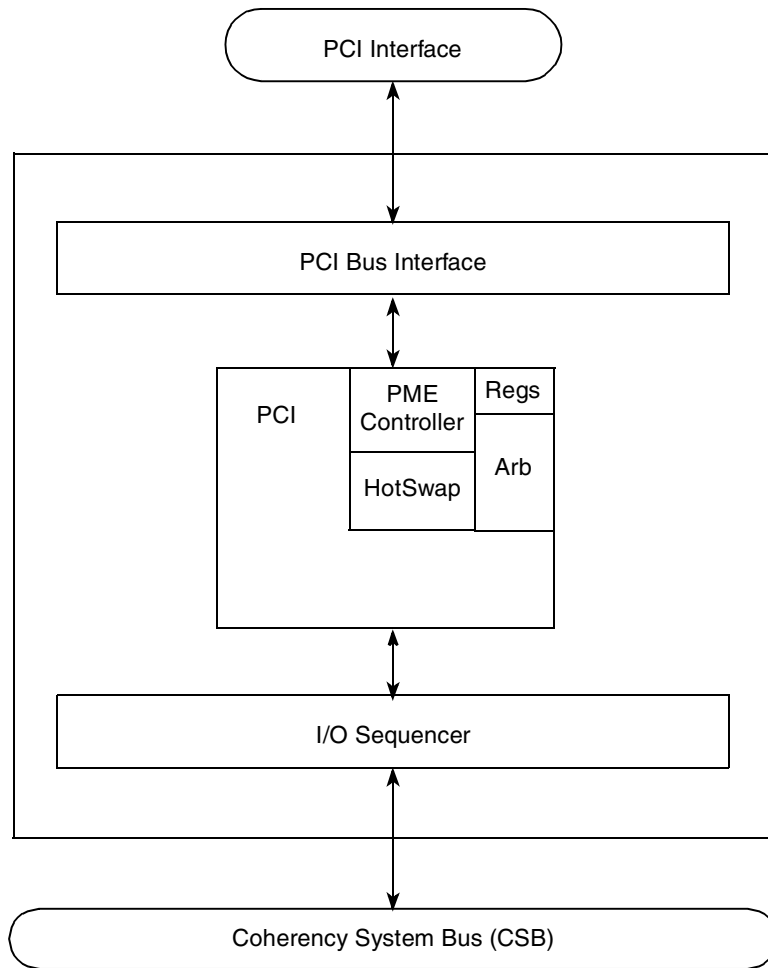
The PCI interface is compatible with the *PCI Local Bus Specification*, Rev. 2.3. It is beyond the scope of this manual to document the intricacies of PCI. This chapter describes the PCI controller and provides a basic description of the PCI bus operations. The specific emphasis is directed at how this device implements the PCI specification. Designers of systems incorporating PCI devices should refer to the respective specifications for a thorough description of the PCI buses.

### NOTE

Much of the available PCI literature refers to a 16-bit quantity as a WORD and a 32-bit quantity as a DWORD. Because this is inconsistent with the terminology in this manual, the terms ‘word’ and ‘double word’ are not used in this chapter. Instead, the number of bits or bytes indicates the exact quantity.

## 13.1 Introduction

The PCI controller acts as a bridge between the PCI interface and the CSB. The I/O sequencer buffers the data. [Figure 13-1](#) is a high-level block diagram of the PCI controller.



**Figure 13-1. PCI Controller Block Diagram**

The PCI controller connects the processor and memory system to the I/O components through the PCI system bus. This interface acts as both initiator (master) and target (slave) device. The PCI controller uses a 32-bit multiplexed, address/data bus that can run at frequencies up to 66-MHz. The interface provides address and data parity with error checking and reporting. The interface provides for three physical address spaces—64-bit address memory, 32-bit address I/O, and PCI configuration space.

Note that PCI supports up to three external masters.

The PCI interface can function as either a PCI host bridge referred to as host mode or a peripheral device on the PCI bus referred to as agent mode. See [Section 13.4.4.4, “Host Mode Configuration Access,”](#) for more information. Note that the PCI controller can be configured from the PCI bus while in agent mode. An address translation mechanism is provided to map PCI memory windows between the PCI bus and the internal bus.

The PCI interface does not flush pending outbound writes as a result of an inbound read command. Systems must not rely on inbound reads to ensure all pending outbound writes have completed. For example, consider the case where a core writes data to a PCI device and then updates a flag in the local



DDR memory indicating the write to PCI has completed. An external PCI master may misread the flag ahead of the actual write transaction's completion on the PCI bus.

### 13.1.1 Features

The PCI controller includes the following features:

- PCI specification revision 2.3 compliant
- 32-bit PCI interface support
- Host and agent mode support
- PCI bus power management unit
- Supports accesses to all PCI address spaces
- 64-bit dual-address cycle (DAC) support (as a target only)
- Internal configuration registers accessible from PCI
- On-chip arbitration supporting three masters on PCI
- Arbiter supports two-level priority request/grant signal pairs
- Supports PCI-to-memory and memory-to-PCI streaming
- Memory prefetching of PCI read accesses and support for delayed read transactions
- Supports posting of processor-to-PCI and PCI-to-memory writes
- Supports selectable snooping for inbound transactions
- Address translation units for address mapping between host and peripheral
- Supports parity
- PCI 3.3-V compatible

### 13.1.2 Modes of Operation

PCI controller modes of operation are determined at reset by the reset configuration word high (RCWH) as described in [Section 4.3.2, “Reset Configuration Words.”](#) [Table 13-1](#) summarizes these modes.

**Table 13-1. PCI Controller Modes**

Parameter	Description	Section/Page
Host/agent configuration	Selects between host and agent mode for the PCI interface.	<a href="#">4.3.2.2.1/4-16</a>
PCI arbiter enable	Enables the on-chip PCI bus arbiter	<a href="#">4.3.2.2/4-14</a>

#### 13.1.2.1 Host/Agent Mode Configuration

The PCI controller can function as either a PCI host bridge (referred to as host mode) or a peripheral device on the PCI bus (referred to as agent mode). Note that host/agent mode selection is determined at power-up as summarized in [Section 4.3.2.2.1, “PCI Host/Agent Configuration.”](#)

When the device powers up in host mode, all inbound configuration accesses are ignored (and thus master aborted). When the device powers up in agent mode, it acknowledges inbound configuration accesses. Note that in PCI agent mode, the PCI controller ignores all PCI memory accesses except those to the

memory-mapped registers until inbound address translation is enabled. In agent mode, configuration cycles are acknowledged if CFG\_LOCK is 0 (see Section 13.3.3.24, “PCI Function Configuration Register”), either from reset configuration or after being cleared by software.

### 13.1.2.2 PCI Arbiter Configuration

The interface can be configured to use an on-chip or off-chip PCI arbiter. Arbitration for PCI is determined by the value in RCWH[PCIARB]. See Section 4.3.2.2, “Reset Configuration Word High Register (RCWHR),” for more information.

## 13.2 External Signal Description

Table 13-2 shows the properties of the PCI signals.

Table 13-2. Signal Properties

Name	Function	Reset State	Pull Up
CPCI_HS_ENUM	CompactPCI hot swap enumerator	High impedance	Required
CPCI_HS_ES	CompactPCI hot swap ejector switch	—	—
CPCI_HS_LED	CompactPCI hot swap LED	Asserted	—
PCI_AD[31:0]	PCIn address / data	High impedance	—
PCI_C/BE[3:0]	PCI bus command / byte enable	High impedance	—
PCI_DEVSEL	PCI device select	High impedance	Required
PCI_FRAME	PCI cycle frame	High impedance	Required
PCI_REQ[0:2]	PCI arbiter requests	Configuration-dependent	Required on inputs
PCI_GNT[0:2]	PCI arbiter grants	Configuration-dependent	—
PCI_IDSEL	PCI initialization device select	—	—
PCI_INTA	PCI interrupt A	High impedance	Required
PCI_IRDY	PCI initiator ready	High impedance	Required
PCI_PAR	PCI parity	High impedance	—
PCI_PERR	PCI parity error	High impedance	Required
PCI_RESET_OUT	PCI reset output	Asserted	
PCI_SERR	PCI system error	High impedance	Required
PCI_STOP	PCI stop	High impedance	Required
PCI_TRDY	PCI target ready	High impedance	Required
PCI_PME	PCI PME assertion request	High impedance	Required

Figure 13-2 shows the external PCI signals.




Figure 13-2. PCI Interface External Signals

Table 13-3 contains detailed descriptions of the external PCI interface signals.

Table 13-3. PCI Interface Signals—Detailed Signal Descriptions

Signal	I/O	Description
CPCI_HS_ENUM	O	CompactPCI hot swap enumerator. Used for the hot swap interface to connect to the host as the enumeration request in a compact PCI system. This signal is used for agent mode only.
		<b>State Meaning</b> Asserted—This card was inserted and needs to be configured, or this card is about to be extracted and needs to be removed from the system resources list. Negated—No action is needed.
		<b>Timing</b> Assertion/Negation—No timing is specified
CPCI_HS_ES	I	CompactPCI hot swap ejector switch. Used for agent mode only. In a compact PCI system this input signal is used for the hot swap interface to connect to the ejector switch logic.
		<b>State Meaning</b> Asserted—The switch is open. Negated—The switch is closed.
		<b>Timing</b> Assertion/Negation—No timing is specified
CPCI_HS_LED	O	CompactPCI hot swap LED. Used for the hot swap interface to connect to the hot swap LED in a CompactPCI system. This signal is used for agent mode only.
		<b>State Meaning</b> Asserted—Output is driving logic 1 to illuminate the hot swap LED. Negated—Output is driving logic 0 to turn off the hot swap LED.
		<b>Timing</b> Assertion/Negation—No timing is specified
M66EN	I	66-MHz enable. Determines the AC timing of the PCI interface.
		<b>State Meaning</b> Asserted—The PCI interface signals use the 66-MHz PCI AC timing parameters. Negated—The PCI interface signals use the 33-MHz PCI AC timing parameters.
		<b>Timing</b> Assertion/Negation—Constant

**Table 13-3. PCI Interface Signals—Detailed Signal Descriptions (continued)**

Signal	I/O	Description	
PCI_AD[31:0]	I/O	PCI address/data bus. During an address phase, these signals contain a physical address. During a data phase, these signals contain the data bytes.	
	O	Outputs for the bi-directional PCI address/data bus.	
		<b>State Meaning</b>	Asserted/Negated—Represents the physical address during the address phase of a PCI transaction. During the data phase(s) of a PCI transaction, the PCI address/data bus contain the data being written. PCI_AD[7:0] define the LSB and, PCI_AD[31:24] define the MSB.
		<b>Timing</b>	Assertion/Negation—As specified by <i>PCI Local Bus Specification Rev 2.3</i>
	I	Inputs for the bi-directional PCI address/data bus.	
		<b>State Meaning</b>	Asserted/Negated—Represents the address to be decoded as a check for device select during the address phase of a PCI transaction or the data being received during the data phase(s) of a PCI transaction. PCI_AD[7:0] define the LSB and, PCI_AD[31:24] define the MSB.
<b>Timing</b>		Assertion/Negation—As specified by <i>PCI Local Bus Specification Rev 2.3</i>	
PCI_C/BE[3:0]	I/O	PCI bus command/byte enable.	
	O	Outputs for the bi-directional command/byte enable.	
		<b>State Meaning</b>	Asserted/Negated—During the address phase, PCI_CBE[3:0], define the bus command. Byte enables determine which byte lanes carry meaningful data for PCI bus data phases. The PCI_CBE[0] signal applies to the LSB.
		<b>Timing</b>	Assertion/Negation—As specified by <i>PCI Local Bus Specification Rev 2.3</i>
	I	Inputs for the bi-directional command/byte enable.	
		<b>State Meaning</b>	Asserted/Negated—During the address phase, PCI_CBE[3:0], indicate the command that another master is sending. During the PCI bus data phase, PCI_CBE[3:0], indicate which byte lanes are valid.
<b>Timing</b>		Assertion/Negation—As specified by <i>PCI Local Bus Specification Rev 2.3</i>	
PCI_DEVSEL	I/O	PCI device select.	
	O	Outputs for the bi-directional device select.	
		<b>State Meaning</b>	Asserted—The PCI controller has decoded the address and is the target of the current access. Negated—The PCI controller has decoded the address and is not the target of the current access.
		<b>Timing</b>	Assertion/Negation—As specified by <i>PCI Local Bus Specification Rev 2.3</i>
	I	Inputs for the bi-directional device select.	
		<b>State Meaning</b>	Asserted—Some PCI agents (other than this PCI controller) have decoded its address as the target of the current access. Negated—No PCI agent has been selected.
<b>Timing</b>		Assertion/Negation—As specified by <i>PCI Local Bus Specification Rev 2.3</i>	

Table 13-3. PCI Interface Signals—Detailed Signal Descriptions (continued)

Signal	I/O	Description	
$\overline{\text{PCI\_FRAME}}$	I/O	PCI cycle frame. Used by the current PCI master to indicate the beginning and duration of an access.	
	O	Outputs for the bi-directional frame.	
		<b>State Meaning</b>	Asserted—The PCI controller acting as a PCI master which is initiating a bus transaction. While $\overline{\text{PCI\_FRAME}}$ is asserted, data transfers may continue. Negated—If $\overline{\text{PCI\_IRDY}}$ is asserted, indicates that the PCI transaction is in the final data phase; if $\overline{\text{PCI\_IRDY}}$ is negated, indicates that the PCI bus is idle.
		<b>Timing</b>	Assertion/Negation—As specified by <i>PCI Local Bus Specification Rev 2.3</i>
	I	Inputs for the bi-directional frame.	
		<b>State Meaning</b>	Asserted—Another PCI master is initiating a bus transaction. Negated—The transaction is in the final data phase or that the bus is idle.
<b>Timing</b>		Assertion/Negation—As specified by <i>PCI Local Bus Specification Rev 2.3</i>	
$\overline{\text{PCI\_GNT0}}$	I/O	PCI arbiter grants. Output signal on this PCI controller when the arbiter is enabled. Input signal when the arbiter is disabled. <b>Note:</b> $\overline{\text{PCI\_GNT}}[0]$ is a point-to-point signal. Every master has its own bus grant signal.	
	O	Outputs for the bi-directional arbiter grants.	
		<b>State Meaning</b>	Asserted—The PCI controller granted control of the PCI bus to agent 0. Negated—The PCI controller did not grant control of the PCI bus to agent 0.
		<b>Timing</b>	Assertion/Negation—As specified by <i>PCI Local Bus Specification Rev 2.3</i>
	I	Inputs for the bi-directional arbiter grants.	
		<b>State Meaning</b>	Asserted—The PCI controller has been granted control of the PCI bus by an external arbiter. Negated—The PCI controller has not been granted control of the PCI bus by an external arbiter.
<b>Timing</b>		Assertion/Negation—As specified by <i>PCI Local Bus Specification Rev 2.3</i>	
$\overline{\text{PCI\_GNT}}[1:2]$	O	PCI arbiter grants. Output signals on this PCI controller when the arbiter is enabled. Note that $\overline{\text{PCI\_GNT}}n$ is a point-to-point signal. Every master has its own bus grant signal.	
		<b>State Meaning</b>	Asserted—The PCI controller granted control of the PCI bus to agent $n$ . Negated—The PCI controller did not grant control of the PCI bus to agent $n$ .
		<b>Timing</b>	Assertion/Negation—As specified by <i>PCI Local Bus Specification Rev 2.3</i>
$\overline{\text{PCI\_IDSEL}}$	I	PCI initialization device select. Used as a chip select during a PCI configuration cycle in agent mode. This signal should be tied low in host mode.	
		<b>State Meaning</b>	Asserted—The PCI controller is being selected as a target of a configuration read or write transactions. Negated—The PCI controller is not being selected as a target of configuration read or write transactions.
		<b>Timing</b>	Assertion/Negation—As specified by <i>PCI Local Bus Specification Rev 2.3</i>

**Table 13-3. PCI Interface Signals—Detailed Signal Descriptions (continued)**

Signal	I/O	Description
$\overline{\text{PCI\_INTA}}$	O	PCI interrupt A.
	<b>State Meaning</b>	Asserted—The PCI controller signals an interrupt to the PCI host. Negated—The PCI controller is not currently signalling an interrupt.
	<b>Timing</b>	Assertion/Negation—As specified by <i>PCI Local Bus Specification Rev 2.3</i>
$\overline{\text{PCI\_IRDY}}$	I/O	PCI initiator ready. This signal is driven by the PCI controller when it is the initiator of a PCI transfer.
	O	Outputs for the bi-directional initiator ready.
	<b>State Meaning</b>	Asserted—The PCI controller, acting as a PCI master, can complete the current data phase of a PCI transaction. During a write, this PCI controller asserts $\overline{\text{PCI\_IRDY}}$ to indicate that valid data is present on $\text{PCI\_AD}[31:0]$ . During a read, this PCI controller asserts $\overline{\text{PCI\_IRDY}}$ to indicate that it is prepared to accept data. Negated—The PCI target needs to wait before this PCI controller, acting as a PCI master, can complete the current data phase. During a write, this PCI controller negates $\overline{\text{PCI\_IRDY}}$ to insert a wait cycle when it cannot provide valid data to the target. During a read, this PCI controller negates $\overline{\text{PCI\_IRDY}}$ to insert a wait cycle when it cannot accept data from the target.
	<b>Timing</b>	Assertion/Negation—As specified by <i>PCI Local Bus Specification Rev 2.3</i>
	I	Inputs for the bi-directional initiator ready.
	<b>State Meaning</b>	Asserted—Another PCI master can complete the current data phase of a transaction. Negated—If $\overline{\text{PCI\_FRAME}}$ is asserted, indicates a wait cycle from another master. If $\overline{\text{PCI\_FRAME}}$ is negated, indicates that the PCI bus is idle.
$\text{PCI\_PAR}$	I/O	PCI parity.
	O	Outputs for the bi-directional parity.
	<b>State Meaning</b>	Asserted—Odd parity across $\text{PCI\_AD}[31:0]$ and $\text{PCI\_CBE}[3:0]$ during address and data phases. Negated—Even parity across $\text{PCI\_AD}[31:0]$ and $\text{PCI\_AD}[31:0]$ during address and data phases.
	<b>Timing</b>	Assertion/Negation—As specified by <i>PCI Local Bus Specification Rev 2.3</i>
	I	Inputs for the bi-directional parity.
	<b>State Meaning</b>	Asserted—Odd parity driven by another PCI master or the PCI target during address and data phases. Negated—Even parity driven by another PCI master or the PCI target during address and data phases.
<b>Timing</b>	Assertion/Negation—As specified by <i>PCI Local Bus Specification Rev 2.3</i>	

Table 13-3. PCI Interface Signals—Detailed Signal Descriptions (continued)

Signal	I/O	Description	
$\overline{\text{PCI\_PERR}}$	I/O	PCI parity error	
	O	Outputs for the bi-directional parity error.	
		<b>State Meaning</b>	Asserted—The PCI controller, acting as a PCI agent, detected a data parity error. (driven by the PCI initiator on reads; driven by the PCI target on writes.) Negated—No error.
		<b>Timing</b>	Assertion/Negation—As specified by <i>PCI Local Bus Specification Rev 2.3</i>
	I	Inputs for the bi-directional parity error.	
		<b>State Meaning</b>	Asserted—Another PCI agent detects a data parity error while this PCI controller is sourcing data (this PCI controller was acting as the PCI initiator during a write, or is acting as the PCI target during a read). Negated—No error.
<b>Timing</b>		Assertion/Negation—As specified by <i>PCI Local Bus Specification Rev 2.3</i>	
$\overline{\text{PCI\_PME}}$	I/O	PCI PME signal	
	O	Outputs for the bi-directional $\overline{\text{PCI\_PME}}$ signal. This is an open-drain signal.	
		<b>State Meaning</b>	Asserted—Indicates that a power management event has occurred Negated—Indicates that no power management event has occurred
		<b>Timing</b>	Assertion/Negation—As specified by <i>PCI Local Bus Specification Rev 2.3</i>
	I	Inputs for the bi-directional $\overline{\text{PCI\_PME}}$ signal.	
		<b>State Meaning</b>	Asserted—Indicates an agent is requesting a power state change Negated—Indicates no power state change request
<b>Timing</b>		Assertion/Negation—As specified by <i>PCI Local Bus Specification Rev 2.3</i>	
$\overline{\text{PCI\_REQ0}}$	I/O	PCI bus request. Input signal on this PCI controller when the arbiter is enabled. Output signal when the arbiter is disabled. Note that $\overline{\text{PCI\_REQ}n}$ is a point-to-point signal. Every master has its own bus request signal.	
	O	Outputs for the bi-directional bus request.	
		<b>State Meaning</b>	Asserted—The PCI controller is requesting control of the PCI bus to perform a transaction. Negated—The PCI controller does not require use of the PCI bus.
		<b>Timing</b>	Assertion/Negation—As specified by <i>PCI Local Bus Specification Rev 2.3</i>
	I	Input for the bi-directional bus request.	
		<b>State Meaning</b>	Asserted—Agent 0 is requesting control of the PCI bus to perform a transaction. Negated—Agent 0 does not require use of the PCI bus.
<b>Timing</b>		Assertion/Negation—As specified by <i>PCI Local Bus Specification Rev 2.3</i>	
$\overline{\text{PCI\_REQ}}[1:2]$	I	PCI bus request. Input signals on this PCI controller when the arbiter is enabled. Note that $\overline{\text{PCI\_REQ}}[n]$ is a point-to-point signal. Every master has its own bus request signal. Following is the state meaning for the $\overline{\text{PCI\_REQ}}[n]$ input.	
		<b>State Meaning</b>	Asserted—An agent $n$ is requesting control of the PCI bus to perform a transaction. Negated—An agent $n$ does not require use of the PCI bus.
		<b>Timing</b>	Assertion/Negation—As specified by <i>PCI Local Bus Specification Rev 2.3</i>

Table 13-3. PCI Interface Signals—Detailed Signal Descriptions (continued)

Signal	I/O	Description
$\overline{\text{PCI\_RESET\_OUT}}$	O	PCI reset. This signal is used only in host mode. It should be left unconnected in agent mode.
		<b>State Meaning</b> Asserted—Devices on the PCI bus are in reset. Negated—Devices on the PCI bus operate normally.
		<b>Timing</b> Assertion/Negation—As specified by <i>PCI Local Bus Specification Rev 2.3</i>
$\overline{\text{PCI\_SERR}}$	O	PCI system error
		Outputs for the bi-directional system error.
		<b>State Meaning</b> Asserted—An address parity error, a target-abort (when this PCI controller is acting as the initiator), or some other system error (where the result is a catastrophic error) was detected. Negated—No error.
	<b>Timing</b> Assertion/Negation—As specified by <i>PCI Local Bus Specification Rev 2.3</i>	
	I	Inputs for the bi-directional system error.
		<b>State Meaning</b> Asserted—A device (other than this PCI controller) has detected a catastrophic error. Negated—No error.
<b>Timing</b> Assertion/Negation—As specified by <i>PCI Local Bus Specification Rev 2.3</i>		
$\overline{\text{PCI\_STOP}}$	O	PCI stop.
		Outputs for the bi-directional stop.
		<b>State Meaning</b> Asserted—The PCI controller, acting as a PCI target, is requesting that the initiator stop the current transaction. Negated—The current transaction can continue.
	<b>Timing</b> Assertion/Negation—As specified by <i>PCI Local Bus Specification Rev 2.3</i>	
	I	Inputs for the bi-directional stop.
		<b>State Meaning</b> Asserted—A target is requesting that this PCI controller, as the initiator, stop the current transaction. Negated—The current transaction can continue.
<b>Timing</b> Assertion/Negation—As specified by <i>PCI Local Bus Specification Rev 2.3</i>		



Table 13-3. PCI Interface Signals—Detailed Signal Descriptions (continued)

Signal	I/O	Description	
$\overline{\text{PCI\_TRDY}}$	I/O	PCI target ready.	
	O	Outputs for the bi-directional target ready.	
		<b>State Meaning</b>	Asserted—The PCI controller, acting as a PCI target, can complete the current data phase of a PCI transaction. During a read, this PCI controller asserts $\overline{\text{PCI\_TRDY}}$ to indicate that valid data is present on PCI_AD[31:0]. During a write, this PCI controller asserts $\overline{\text{PCI\_TRDY}}$ to indicate that it is prepared to accept data. Negated—The PCI initiator needs to wait before this PCI controller, acting as a PCI target, can complete the current data phase. During a read, this PCI controller negates $\overline{\text{PCI\_TRDY}}$ to insert a wait cycle when it cannot provide valid data to the initiator. During a write, this PCI controller negates $\overline{\text{PCI\_TRDY}}$ to insert a wait cycle when it cannot accept data from the initiator.
		<b>Timing</b>	Assertion/Negation—As specified by <i>PCI Local Bus Specification Rev 2.3</i>
	I	Inputs for the bi-directional target ready.	
		<b>State Meaning</b>	Asserted—Another PCI target is able to complete the current data phase of a transaction. Negated—A wait cycle from another target.
<b>Timing</b>		Assertion/Negation—As specified by <i>PCI Local Bus Specification Rev 2.3</i>	

### 13.3 Memory Map/Register Definitions

The PCI controller has the following types of registers:

- The PCI configuration access registers. Used for generating PCI configuration accesses from the CSB. These registers, listed in [Table 13-4](#), are memory-mapped on the CSB and accessed through the IMMR window.
- The PCI memory-mapped registers. Used to manage error functions, general control and status, and address translation control for the inbound path. These registers are shown in [Table 13-5](#). They can be accessed by PCI masters via the PCI controller to the CSB through the PIMMR inbound window. Note that [Table 13-5](#) does not list outbound address translation registers; these are contained in the I/O sequencer (IOS) memory-mapped registers. See [Chapter 12](#), “[DMA/Messaging Unit](#),” for more information.
- The PCI configuration space registers. Defined by the PCI specification. These registers are accessed by PCI masters using configuration accesses and are described in [Section 13.3.3](#), “[PCI Configuration Space Registers](#).”

Table 13-4. PCI Configuration Access Registers

Offset	Register	Access	Reset	Section/Page
<b>PCI Configuration Access Registers</b>				
0x0	PCI_CONFIG_ADDRESS	W	0x0000_0000	<a href="#">13.3.1.1/13-13</a>
0x4	PCI_CONFIG_DATA	R/W	0x0000_0000	<a href="#">13.3.1.2/13-14</a>
0x8	PCI_INT_ACK	R	0x0000_0000	<a href="#">13.3.1.3/13-15</a>

Table 13-5. PCI Memory-Mapped Registers

Offset	Register	Access	Reset	Section/Page
<b>PCI Error Management Registers</b>				
0x00	PCI error status register (PCI_ESR)	w1c	0x0000_0000	13.3.2.1/13-15
0x04	PCI error capture disable register (PCI_ECDR)	R/W	0x0000_0000	13.3.2.2/13-16
0x08	PCI error enable register (PCI_EER)	R/W	0x0000_0000	13.3.2.3/13-17
0x0C	PCI error attributes capture register (PCI_EATCR)	R/W	0x0000_0000	13.3.2.4/13-18
0x10	PCI error address capture register (PCI_EACR)	R	0x0000_0000	13.3.2.5/13-19
0x14	PCI error extended address capture register (PCI_EEACR)	R	0x0000_0000	13.3.2.6/13-20
0x18	PCI error data capture register (PCI_EDCR)	R/W	0x0000_0000	13.3.2.7/13-20
<b>PCI Control and Status Registers</b>				
0x20	PCI general control register (PCI_GCR)	R/W	0x0000_0000	13.3.2.8/13-20
0x24	PCI error control register (PCI_ECR)	R/W	0x0000_0000	13.3.2.9/13-21
0x28	PCI general status register (PCI_GSR)	R	0x0000_0000	13.3.2.10/13-22
<b>PCI Inbound ATU Registers</b>				
0x38	PCI inbound translation address register 2 (PITAR2)	R/W	0x0000_0000	13.3.2.11/13-22
0x3C	Reserved	—	—	—
0x40	PCI inbound base address register 2 (PIBAR2)	R/W	0x0000_0000	13.3.2.12/13-23
0x44	PCI inbound extended base address register 2 (PIEBAR2)	R/W	0x0000_0000	13.3.2.13/13-24
0x48	PCI inbound window attributes register 2 (PIWAR2)	R/W	0x0000_0000	13.3.2.14/13-24
0x50	PCI inbound translation address register 1 (PITAR1)	R/W	0x0000_0000	13.3.2.11/13-22
0x54	Reserved	—	—	—
0x58	PCI inbound base address register 1 (PIBAR1)	R/W	0x0000_0000	13.3.2.12/13-23
0x5C	PCI inbound extended base address register 1 (PIEBAR1)	R/W	0x0000_0000	13.3.2.13/13-24
0x60	PCI inbound window attributes register 1 (PIWAR1)	R/W	0x0000_0000	13.3.2.14/13-24
0x68	PCI inbound translation address register 0 (PITAR0)	R/W	0x0000_0000	13.3.2.11/13-22
0x6C	Reserved	—	—	—
0x70	PCI inbound base address register 0 (PIBAR0)	R/W	0x0000_0000	13.3.2.12/13-23
0x78	PCI inbound window attributes register 0 (PIWAR0)	R/W	0x0000_0000	13.3.2.13/13-24
0x7C– 0xFF	Reserved	—	—	—


### 13.3.1 PCI Configuration Access Registers

This section describes the registers used to allow a local bus master to access the PCI configuration space, and generate special cycle or interrupt acknowledge transactions on the PCI bus. A special case provides

access to the PCI controller's internal PCI configuration registers. The PCI registers, PCI\_CONFIG\_ADDRESS, PCI\_CONFIG\_DATA, and PCI\_INT\_ACK, are little endian registers.

### 13.3.1.1 PCI\_CONFIG\_ADDRESS

Figure 13-3 shows the PCI\_CONFIG\_ADDRESS register fields.



**Figure 13-3. PCI\_CONFIG\_ADDRESS Register**

The PCI\_CONFIG\_ADDRESS register holds the address for an access to the PCI configuration space from the local bus. This register must be programmed before accessing PCI\_CONFIG\_DATA to perform the transaction. Only 32-bit accesses are permitted.

If EN=1, BN=0, and DN=0, the access is to the internal PCI configuration registers, so no transaction is generated on the PCI bus.

If EN=1, BN=0, DN=31, FN=7, and RN=0, writing to PCI\_CONFIG\_DATA generates a special cycle transaction and reading from PCI\_CONFIG\_DATA generates an interrupt acknowledge transaction.

Table 13-6 shows the bit settings of the PCI\_CONFIG\_ADDRESS register.

**Table 13-6. PCI\_CONFIG\_ADDRESS Field Descriptions**

Bits	Name	Description
31	EN	Enable configuration transaction. Determines the type of transaction to be generated. 0 No configuration transaction will be generated by accessing the CONFIG_DATA register. Such an access will be passed through to the PCI bus as an I/O transaction. Since this is generally not desirable, the user should not access CONFIG_DATA when the EN bit is 0. 1 A configuration transaction will be generated by accessing the CONFIG_DATA register if BN and DN are not both zero.
30–24	—	Reserved
23–16	BN	Bus number. Specifies the bus segment to which a configuration transaction is directed. If this field is 0, a Type 0 configuration transaction is generated. Otherwise, a Type 1 configuration transaction is generated.

**Table 13-6. PCI\_CONFIG\_ADDRESS Field Descriptions (continued)**

Bits	Name	Description																																																		
15–11	DN	Device number. Specifies the device to which a configuration transaction is directed. For a Type 0 configuration transaction, this field is decoded to individual PCI1_IDSEL signals for the address phase according to the following values. For a Type 1 configuration transaction, this field is used directly for the address phase.																																																		
		<table border="0"> <thead> <tr> <th>Value</th> <th>AD Signal that is Driving High</th> </tr> </thead> <tbody> <tr><td>01010</td><td>31</td></tr> <tr><td>01011</td><td>11</td></tr> <tr><td>01100</td><td>12</td></tr> <tr><td>01101</td><td>13</td></tr> <tr><td>01110</td><td>14</td></tr> <tr><td>01111</td><td>15</td></tr> <tr><td>10000</td><td>16</td></tr> <tr><td>10001</td><td>17</td></tr> <tr><td>10010</td><td>18</td></tr> <tr><td>10011</td><td>19</td></tr> <tr><td>10100</td><td>20</td></tr> <tr><td>10101</td><td>21</td></tr> <tr><td>10110</td><td>22</td></tr> <tr><td>10111</td><td>23</td></tr> <tr><td>11000</td><td>24</td></tr> <tr><td>11001</td><td>25</td></tr> <tr><td>11010</td><td>26</td></tr> <tr><td>11011</td><td>27</td></tr> <tr><td>11100</td><td>28</td></tr> <tr><td>11101</td><td>29</td></tr> <tr><td>11110</td><td>30</td></tr> <tr><td>11111</td><td>Special cycle / interrupt acknowledge</td></tr> <tr><td>00000</td><td>Internal access</td></tr> <tr><td>Others</td><td>Reserved</td></tr> </tbody> </table>	Value	AD Signal that is Driving High	01010	31	01011	11	01100	12	01101	13	01110	14	01111	15	10000	16	10001	17	10010	18	10011	19	10100	20	10101	21	10110	22	10111	23	11000	24	11001	25	11010	26	11011	27	11100	28	11101	29	11110	30	11111	Special cycle / interrupt acknowledge	00000	Internal access	Others	Reserved
		Value	AD Signal that is Driving High																																																	
		01010	31																																																	
		01011	11																																																	
		01100	12																																																	
		01101	13																																																	
		01110	14																																																	
		01111	15																																																	
		10000	16																																																	
		10001	17																																																	
		10010	18																																																	
		10011	19																																																	
		10100	20																																																	
		10101	21																																																	
		10110	22																																																	
		10111	23																																																	
		11000	24																																																	
		11001	25																																																	
		11010	26																																																	
		11011	27																																																	
		11100	28																																																	
11101	29																																																			
11110	30																																																			
11111	Special cycle / interrupt acknowledge																																																			
00000	Internal access																																																			
Others	Reserved																																																			
10–8	FN	Function number. Specifies the function to which the configuration transaction is directed on a multi-function device. It is used directly in the address phase of the configuration transaction.																																																		
7–2	RN	Register number. Specifies the register being accessed in the PCI configuration space.																																																		
1–0	—	Reserved																																																		

### 13.3.1.2 PCI\_CONFIG\_DATA

An access to PCI\_CONFIG\_DATA usually generates a PCI configuration transaction if PCI\_CONFIG\_ADDRESS[EN] is set. There are some exceptions contained in the description of PCI\_CONFIG\_ADDRESS[EN].

This register may be accessed with an 8-, 16-, or 32-bit access, depending on the width of the register targeted by the configuration transaction.

Figure 13-4 shows the PCI\_CONFIG\_DATA register fields.




Figure 13-4. PCI\_CONFIG\_DATA

Table 13-7 shows the bit settings of the PCI\_CONFIG\_DATA register.

Table 13-7. PCI\_CONFIG\_DATA Field Descriptions

Bits	Name	Description
31–0	CFG_DATA	Configuration data. This field contains the data transferred on a PCI configuration transaction.

### 13.3.1.3 PCI Interrupt Acknowledge Register (PCI\_INT\_ACK)

Reading this register generates an interrupt acknowledge transaction on the PCI bus. The value that is read is undefined.

## 13.3.2 PCI Memory-Mapped Control and Status Registers

This section describes the control and status registers.

### 13.3.2.1 PCI Error Status Register (PCI\_ESR)

The PCI error status register (PCI\_ESR) contains status bits for various types of error conditions captured by the PCI controller. Each status bit is set when the corresponding error condition is captured. PCI\_ESR is a write-1-to-clear type register. A bit is cleared whenever the register is written and the data in the corresponding bit location is a 1. Figure 13-5 shows the PCI\_ESR fields.




Figure 13-5. PCI Error Status Register (PCI\_ESR)

Table 13-8 describes the bit settings of the PCI\_ESR register.

**Table 13-8. PCI\_ESR Field Descriptions**


Bits	Name	Description
0	MERR	Multiple errors. Set if any other bit of this register is 1 and the same error type occurs again.
1–20	—	Reserved
21	APAR	Address parity error. Set when there is an address parity error on a PCI access initiated by a device other than this PCI controller.
22	PCISERR	PCI system error. Set when the $\overline{\text{PCI\_SERR}}$ input signal is asserted. See Table 13-3 for more information on $\overline{\text{PCI\_SERR}}$ .
23	MPERR	Master parity error. Set when the $\overline{\text{PCI\_PERR}}$ input signal is asserted on a write access initiated by this PCI controller or when a data parity error is detected by this PCI controller on a read access that it initiated.
24	TPERR	Target parity error. Set when this PCI controller is the target of a transaction and the $\overline{\text{PCI\_PERR}}$ input signal is asserted on a read access or a data parity error is detected by this PCI controller on a write access.
25	NORSP	No response. Set when there is no response to a transaction initiated by this PCI controller on the PCI bus (no $\overline{\text{PCI\_DEVSEL}}$ assertion).
26	TABT	Target abort. Set when a PCI target abort occurs on a transaction initiated by this PCI controller.
27–31	—	Reserved

### 13.3.2.2 PCI Error Capture Disable Register (PCI\_ECDR)

PCI\_ECDR contains fields for controlling the capture of the transaction that caused an error. Each bit corresponds to the error condition reported in the PCI error status register (PCI\_ESR). Note that only the first error is captured, so disabling the capture of some error types may allow greater visibility of the significant errors.

- 1 = Do not capture the transaction that caused this error.
- 0 = Capture the transaction that caused this error.

Figure 13-6 shows the PCI\_ECDR fields.



**Figure 13-6. PCI Error Capture Disable Register (PCI\_ECDR)**

Table 13-9 describes the bit settings of the PCI\_ECDR register.

**Table 13-9. PCI\_ECDR Field Descriptions**


Bits	Name	Description
0–20	—	Reserved
21	APAR	Address parity error. Disable capture for address parity errors
22	PCISERR	PCI system error. Disable capture for received $\overline{\text{PCI\_SERR}}$ errors
23	MPERR	Master parity error. Disable capture for master $\overline{\text{PCI\_PERR}}$ errors
24	TPERR	Target parity error. Disable capture for target $\overline{\text{PCI\_PERR}}$ errors
25	NORSP	No response. Disable capture for master-abort errors
26	TABT	Target abort. Disable capture for target abort errors
27–31	—	Reserved

### 13.3.2.3 PCI Error Enable Register (PCI\_EER)

PCI\_EER contains fields for enabling the assertion of an interrupt for the error conditions reported in the PCI error status register (PCI\_ESR).

- 1 = The interrupt is enabled.
- 0 = The interrupt is disabled.

Figure 13-7 shows the PCI\_EER fields.



**Figure 13-7. PCI Error Enable Register (PCI\_EER)**

Table 13-10 describes the bit settings of the PCI\_EER register.

**Table 13-10. PCI\_EER Field Descriptions**

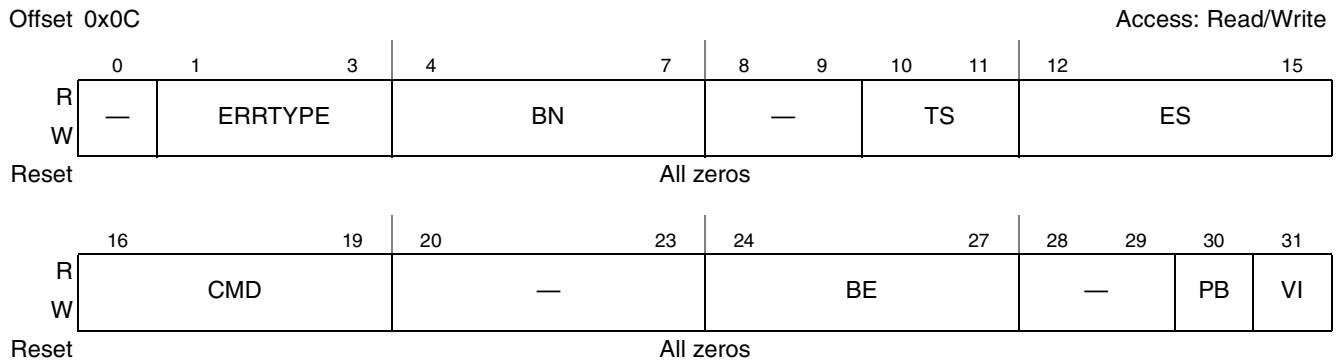
Bits	Name	Description
0–20	—	Reserved
21	APAR	Address parity error. Generate an interrupt when the corresponding bit of the PCI_ESR is 1.
22	PCISERR	PCI system error. Generate an interrupt when the corresponding bit of the PCI_ESR is 1.
23	MPERR	Master parity error. Generate an interrupt when the corresponding bit of the PCI_ESR is 1.
24	TPERR	Target parity error. Generate an interrupt when the corresponding bit of the PCI_ESR is 1.
25	NORSP	No response. Generate an interrupt when the corresponding bit of the PCI_ESR is 1.

**Table 13-10. PCI\_EER Field Descriptions (continued)**

Bits	Name	Description
26	TABT	Target abort. Generate an interrupt when the corresponding bit of the PCI_ESR is 1.
27–31	—	Reserved

### 13.3.2.4 PCI Error Attributes Capture Register (PCI\_EATCR)

PCI\_EATCR contains fields for storing information associated with the first PCI error captured. Figure 13-8 shows the PCI\_EATCR fields.



**Figure 13-8. PCI Error Attributes Capture Register (PCI\_EATCR)**

Table 13-11 describes the bit settings of the PCI\_EATCR register.

**Table 13-11. PCI\_EATCR Field Descriptions**

Bits	Name	Description
0	—	Reserved
1–3	ERRTYPE	First error type. This field is encoded to indicate the type of the first PCI error captured. 000 Address parity error 001 Write data parity error 010 Read data parity error 011 Master abort 100 Target abort 101 System error indication received 110 Parity error indication received on a read 111 Parity error indication received on a write
4–7	BN	Beat number. This field provides the data beat number on which the error occurred for data parity errors. The value of this field is undefined for other error types. The beat values are described as follows: 0000 1st beat 0001 2nd beat 0010 3rd beat 0011 4th beat 0100 5th beat 0101 6th beat 0110 7th beat 0111 8th beat 1000 9th beat or beyond (transaction larger than one cache line) Others Reserved



Table 13-11. PCI\_EATCR Field Descriptions (continued)

Bits	Name	Description
8–9	—	Reserved
10–11	TS	Transaction size. Indicates the size of the transaction in units of doublewords (8 bytes). If the transaction crossed a cache line (32-byte) boundary, this field indicates the number of actual double words in the cache line on which the error occurred. This field is valid only if the PCI controller was the master of the transaction. 00 4 double words 01 1 double word 10 2 double words 11 3 double words
12–15	ES	Error source. This field indicates the source of the PCI transaction. 0000 External master 0101 DMA Others reserved
16–19	CMD	PCI command. Contains the PCI command PCI_CBE[3:0] of the transaction.
20–23	—	Reserved
24–27	BE	PCI byte enables. Contains the PCI byte enables PCI_CBE[3:0] for the data word.
28–29	—	Reserved
30	PB	Parity bit. Contains the PCI parity bit for the captured data word.
31	VI	Error information valid. This bit indicates that the error information captured in this register, PCI_EACR, PCI_EEACR, and PCI_EDCR is valid. 0 No valid error information 1 Error information is valid

### 13.3.2.5 PCI Error Address Capture Register (PCI\_EACR)

PCI\_EACR contains fields for storing the low portion of the address associated with the first PCI error captured. Figure 13-9 shows the PCI\_EACR fields.




Figure 13-9. PCI Error Address Capture Register (PCI\_EACR)

Table 13-12 describes the bit settings of the PCI\_EACR register.

Table 13-12. PCI\_EACR Field Description

Bits	Name	Description
0–31	PCI_EA	PCI error address. Contains the low portion of the address associated with the first detected error. Read only.

### 13.3.2.6 PCI Error Extended Address Capture Register (PCI\_EEACR)

PCI\_EEACR contains fields for storing the high portion of the address associated with the first PCI error captured. Figure 13-10 shows the PCI\_EEACR fields.




Figure 13-10. PCI Error Extended Address Capture Register (PCI\_EEACR)

Table 13-13 describes the bit settings of the PCI\_EEACR register.

Table 13-13. PCI\_EEACR Field Description

Bits	Name	Description
0–31	PCI_EEA	PCI error extended address. Contains the high portion of the address associated with the first detected error.

### 13.3.2.7 PCI Error Data Low Capture Register (PCI\_EDLCR)

PCI\_EDLCR contains fields for storing the data associated with the first PCI error captured. Figure 13-11 shows the PCI\_EDLCR fields.




Figure 13-11. PCI Error Data Low Capture Register (PCI\_EDLCR)


Table 13-14 describes the bit settings of the PCI\_EDLCR register.

Table 13-14. PCI\_EDLCR Field Description

Bits	Name	Description
0–31	PCI_EDR	PCI error data. Contains the data associated with the first detected error.

### 13.3.2.8 PCI General Control Register (PCI\_GCR)

PCI\_GCR contains fields for controlling the behavior of the internal arbiter, the state of the bus signals, and the PCI reset signal for host mode. Figure 13-12 shows the PCI\_GCR fields.



**Figure 13-12. PCI General Control Register (PCI\_GCR)**

Table 13-15 shows the bit settings of PCI\_GCR. The bits that are not reserved have read and write capability.

**Table 13-15. PCI\_GCR Field Descriptions**


Bits	Name	Description
0–28	—	Reserved
29	BBR	Block bus requests. This bit could be used to prepare for entering a low-power mode by preventing transactions on the PCI bus. 0 External bus requests are treated normally. 1 Block external bus requests. When this bit is set, all bus requests from external devices to the PCI controller’s internal arbiter are blocked, and the bus is continuously granted to the PCI controller.
30	PPL	PCI pins low. This bit could be used to put the bus signals in a safe electrical state when the devices on the bus are powered down. This bit should never be set during normal operation of the PCI bus. 0 PCI pins function normally 1 PCI pins in the low state. Setting this bit forces all the output and bidirectional pins of the PCI bus to be driven low.
31	SPRST	Soft PCI reset. This bit provides software control of the $\overline{\text{PCI\_RESET\_OUT}}$ output signal. It is only valid in host mode. 0 $\overline{\text{PCI\_RESET\_OUT}}$ is driven low. 1 $\overline{\text{PCI\_RESET\_OUT}}$ is driven high.

**13.3.2.9 PCI Error Control Register (PCI\_ECR)**

PCI\_ECR contains fields for determining whether an interrupt or machine check is generated for the error conditions reported in the PCI error status register (PCI\_ESR). Note that if the corresponding bit in the PCI error enable register (PCI\_EER) is clear, the bit in the PCI error control register (PCI\_ECR) has no effect.

- 1 = A machine check is generated.
- 0 = An interrupt is generated.

Figure 13-13 shows the PCI\_ECR fields.



**Figure 13-13. PCI Error Control Register (PCI\_ECR)**


Table 13-16 describes the bit settings of the PCI\_ECR register.

**Table 13-16. PCI\_ECR Field Descriptions**

Bits	Name	Description
0–20	—	Reserved
21	APAR	0 An interrupt is generated if the corresponding bit of the PCI_ESR is 1. 1 A machine check is generated if the corresponding bit of the PCI_ESR is 1.
22	PCISERR	0 An interrupt is generated if the corresponding bit of the PCI_ESR is 1. 1 A machine check is generated if the corresponding bit of the PCI_ESR is 1.
23	MPERR	0 An interrupt is generated if the corresponding bit of the PCI_ESR is 1. 1 A machine check is generated if the corresponding bit of the PCI_ESR is 1.
24	TPERR	0 An interrupt is generated if the corresponding bit of the PCI_ESR is 1. 1 A machine check is generated if the corresponding bit of the PCI_ESR is 1.
25	NORSP	0 An interrupt is generated if the corresponding bit of the PCI_ESR is 1. 1 A machine check is generated if the corresponding bit of the PCI_ESR is 1.
26	TABT	0 An interrupt is generated if the corresponding bit of the PCI_ESR is 1. 1 A machine check is generated if the corresponding bit of the PCI_ESR is 1.
27–31	—	Reserved

### 13.3.2.10 PCI General Status Register (PCI\_GSR)

PCI\_GSR contains fields for providing status information, shown in Figure 13-14.



**Figure 13-14. PCI General Status Register (PCI\_GSR)**

Table 13-17 shows the bit settings of the PCI\_GSR register. All bits are read-only.


**Table 13-17. PCI\_GSR Field Descriptions**

Bits	Name	Description
0–30	—	Reserved
31	IDLE	PCI controller is idle. Indicates when the PCI bus is totally idle before setting PCI_GCR[PPL]. 0 The PCI controller is active. 1 The PCI controller is idle.

### 13.3.2.11 PCI Inbound Translation Address Registers (PITAR<sub>n</sub>)

PITAR<sub>n</sub> contains fields for defining the starting point of the inbound translation windows in the local memory space (see Section 13.4.6, “PCI Inbound Address Translation”; for more information on outbound address translation registers, see Chapter 11, “Sequencer”). Inbound and outbound windows for the same bus should not overlap. Therefore, situations where an inbound window translation points back into an

outbound window, or where an outbound translation window points back into an inbound window, are not allowed.



**Figure 13-15. PCI Inbound Translation Address Registers (PITAR $n$ )**


Table 13-18 shows the bit settings of PITAR $n$ .

**Table 13-18. PITAR $n$  Field Descriptions**

Bits	Name	Description
0–11	—	Reserved
12–31	TA	Translation address. Contains the starting address of the inbound translated address. TA corresponds to the 20 highest-order bits of a 32-bit local address. The specified address must be aligned to the window size, as defined by PIWAR $n$ [IWS].

### 13.3.2.12 PCI Inbound Base Address Registers (PIBAR $n$ )

PIBAR $n$  contains fields for defining the starting point of the inbound windows in the PCI memory space. A write to a PIBAR $n$  register also causes a change in the base address bits in the corresponding GPL base address register in the PCI configuration space. Figure 13-16 shows the PIBAR $x$  fields.



**Figure 13-16. PCI Inbound Base Address Registers (PIBAR $n$ )**

Table 13-19 shows the bit settings of PIBAR $n$ .

**Table 13-19. PIBAR $n$  Field Descriptions**

Bits	Name	Description
0–31	BA	Base address. Contains the starting address in the PCI memory space of the inbound window. This field corresponds to bits 43–12 of a 64-bit address. In PIBAR0, the upper 12 bits are reserved because only a 32-bit address is supported. The specified address must be aligned to the window size, as defined by PIWAR $n$ [IWS].

### 13.3.2.13 PCI Inbound Extended Base Address Registers (PIEBAR<sub>n</sub>)

PIEBAR<sub>n</sub> contains fields for defining the high portion of the starting point of the inbound windows in the PCI memory space. Figure 13-17 shows the PIEBAR<sub>n</sub> fields.




Figure 13-17. PCI Inbound Extended Base Address Registers (PIEBAR<sub>n</sub>)

Table 13-20 shows the bit settings of PIEBAR<sub>n</sub>.

Table 13-20. PIEBAR<sub>n</sub> Field Descriptions

Bits	Name	Description
0–11	—	Reserved
12–31	EBA	Extended base address. Contains the high portion of the starting address in the PCI memory space of the inbound base address. This 20-bit field corresponds to bits 63–44 of a 64-bit address.

### 13.3.2.14 PCI Inbound Window Attribute Registers (PIWAR<sub>n</sub>)

PIWAR<sub>n</sub> contains fields for defining the size of an inbound translation window. It also defines some properties of the window. Figure 13-18 shows the PIWAR<sub>n</sub> fields. See Section 4.3.1.1, “Reset Configuration Word Source,” for more information on reset configuration.




Figure 13-18. PCI Inbound Window Attribute Registers (PIWAR<sub>n</sub>)

Table 13-21 shows the bit settings of PIWAR<sub>n</sub>.

Table 13-21. PIWAR<sub>n</sub> Field Descriptions

Bits	Name	Description
0	EN	Enable. Used to enable the address translation window. 0 Address translation is disabled for this window. 1 Address translation is enabled for this window. PCI addresses that match the definition of the window will be recognized by the PCI controller and translated to the local memory space.
1	—	Reserved

Table 13-21. PIWAR<sub>n</sub> Field Descriptions (continued)

Bits	Name	Description
2	PF	Prefetchable. Defines whether the transactions that are translated through this window are prefetchable on the local bus. Streaming the transactions requires the memory space to be prefetchable. 0 Not prefetchable 1 Prefetchable
3–11	—	Reserved
12–15	RTT	Read transaction type. Determines the type of transaction performed on the local bus when the PCI transaction is a read. The RTT values are described as follows: 0100 Read without snoop on system bus 0101 Read with snoop on system bus Others reserved
16–19	WTT	Write transaction type. Determines the type of transaction performed on the local bus when the PCI transaction is a write. The WTT values are described as follows: 0100 Write without snoop of local processor 0101 Write with snoop of local processor Others reserved
20–25	—	Reserved
26–31	IWS	Inbound window size. Indicates the size of the inbound translation window. Inbound translation window size N which is the encoded $2^{(N+1)}$ bytes window size. The smallest window is 4 Kbytes (N = 11) 000000–001010 Reserved 001011 4-Kbyte window size 001100 8-Kbyte window size ... 011110 2-Gbyte window size 011111–111111 Reserved

### 13.3.3 PCI Configuration Space Registers

This section describes the PCI configuration space registers. These registers are shown with descending bit numbering to correspond to the PCI standard.

#### NOTE

The registers described in this section use little-endian byte ordering. Software running on the local processor in big-endian mode must byte-swap the data. No byte swapping occurs when the registers are accessed from the PCI bus.

Table 13-22 shows the PCI configuration registers that are mapped in PCI configuration space. Some fields are common to registers in both spaces to ensure consistency. These fields are discussed in the register definitions.

Table 13-22. PCI Configuration Space Registers

Address	Use	Access
00	Vendor ID configuration register	R
02	Device ID configuration register	R
04	PCI command configuration register	R/W
06	PCI status configuration register	Read/bit-reset
08	Revision ID configuration register	R
09	Standard programming interface	R
0A	Subclass code configuration register	R
0B	Base class code configuration register	R
0C	Cache line size configuration register	R/W
0D	Latency timer configuration register	R/W
0E	Header type configuration register	R
0F	BIST control configuration register	R
10	PIMMR base address register	R/W
14	GPL base address register 0	R/W
18	GPL base address register 1	R/W
1C	GPL extended base address register 1	R/W
20	GPL base address register 2	R/W
24	GPL extended base address register 2	R/W
2C	Subsystem vendor ID configuration register	R
2E	Subsystem device ID configuration register	R
34	Capabilities pointer configuration register	R
3C	Interrupt line configuration register	R/W
3D	Interrupt pin configuration register	R
3E	Minimum grant configuration register	R
3F	Maximum latency configuration register	R
44	PCI function configuration register	R/W
46	PCI arbiter control register (PCIACR)	R/W
48	Hot swap register block	R/W
80	PCI power management register 0	R
84	PCI power management register 1	R/W



### 13.3.3.1 Vendor ID Configuration Register

Figure 13-19 shows the vendor ID fields. This is a read-only register.




Figure 13-19. Vendor ID Configuration Register

Table 13-23 shows the bit settings of the vendor ID register.

Table 13-23. Vendor ID Configuration Register Field Descriptions

Bits	Name	Description
15–0	VID	Vendor ID. The read-only value 0x1957 specifies Freescale Semiconductor as the manufacturer of the device.

### 13.3.3.2 Device ID Configuration Register

Figure 13-20 shows the device ID fields. This is a read only register.




Figure 13-20. Device ID Configuration Register

Table 13-24 shows the bit settings of the device ID register.

Table 13-24. Device ID Configuration Register Field Descriptions

Bits	Name	Description
15–0	DID	Device ID. This field identifies the device. 00B0 MPC8313E 00B1 MPC8313 00B2 SC8311E 00B3 SC8311

### 13.3.3.3 PCI Command Configuration Register

Figure 13-21 shows the PCI command fields.

Offset 04

Access: Mixed




Figure 13-21. PCI Command Configuration Register

Table 13-25 shows the bit settings of the PCI command register.

Table 13-25. PCI Command Configuration Register Field Descriptions

Bits	Name	Description
15–11	—	Reserved
10	INTD	Interrupt Disable. Setting this bit masks the $\overline{\text{PCI\_INTA}}$ output. 0 $\overline{\text{PCI\_INTA}}$ provides the device interrupt status. 1 $\overline{\text{PCI\_INTA}}$ is always negated.
9	FB-B	Fast back-to-back. Hard-wired to 0.
8	SERREN	SERR enable. This bit is an enable bit for the SERR driver. Address parity errors are reported only if this bit and bit 6 are 1. 0 $\overline{\text{PCI\_SERR}}$ is never asserted. 1 $\overline{\text{PCI\_SERR}}$ may be asserted to indicate error conditions.
7	—	Reserved
6	PERRR	Parity error response. Controls the PCI controller's response to a parity error. 0 Parity errors are ignored and normal operation continues. 1 Standard parity error treatment.
5	—	Reserved
4	MWI	Memory-write-and-invalidate. Hard-wired to 0.
3	SC	Special cycles. Hard-wired to 0.
2	BMST	Bus master. Controls the PCI controller's ability to be a master on the PCI bus. At reset, this bit is cleared in Agent Mode and set in Host Mode. 0 The PCI controller does not generate PCI accesses. 1 The PCI controller behaves as a bus master.
1	MEM	Memory space. Controls the response to memory space accesses. 0 The PCI controller does not respond to Memory Space accesses. 1 The PCI controller as a target responds to Memory Space accesses.
0	I/O	I/O space. Hard-wired to 0.

### 13.3.3.4 PCI Status Configuration Register

This register is used to record status information for PCI bus-related events. Some of the bits are hard-wired to indicate the capabilities of the PCI controller. Other bits can be cleared by writing 1 to the bit location. Figure 13-22 shows the PCI status fields.

Offset 06													Access: Mixed			
	15	14	13	12	11	10	9	8	7	6	5	4	3	2	1	0
R	DPERR	SSERR	RMA	RTA	STA	DEVSEL_T		DPD	FB-BC	—	66M	CL	INTS	—		
W	w1c	w1c	w1c	w1c	w1c			w1c					w1c			
Reset	0	0	0	0	0	0	0	0	1	0	1	1	0	0	0	0

Figure 13-22. PCI Status Configuration Register


Table 13-26 shows the bit settings of the PCI status register.

Table 13-26. PCI Status Configuration Register Field Descriptions

Bits	Name	Description
15	DPERR	Detected parity error. Set whenever the PCI controller detects a parity error on the PCI bus, even if parity error handling is disabled (as controlled by bit 6 in the PCI Command register).
14	SSERR	Signaled system error. Set whenever $\overline{\text{PCI\_SERR}}$ is asserted.
13	RMA	Received master abort. Set whenever the PCI controller, acting as the PCI master on the PCI bus, terminates a transaction (except for a special-cycle) using master-abort.
12	RTA	Received target abort. Set whenever a transaction initiated by this PCI controller on the PCI bus is terminated by a target-abort.
11	STA	Signaled target abort. Set whenever the PCI controller, acting as the PCI target on the PCI bus, issues a target-abort to a PCI master.
10–9	DEVSEL_T	DEVSEL timing. Hard-wired to 00.
8	DPD	Master data parity error. Set when a data parity error is detected on the PCI bus, if the PCI controller is the master that initiated the transaction and bit 6 in the PCI command register is set.
7	FB-BC	Fast back-to-back capable. Hard-wired to 1.
6	—	Reserved
5	66M	66-MHz capable. Hard-wired to 1.
4	CL	Capabilities list. Hard-wired to 1.
3	INTS	Interrupt status. Contains the status of the device interrupt. The value of this bit is not affected by the INTD bit of the PCI command configuration register.
2–0	—	Reserved

### 13.3.3.5 Revision ID Configuration Register

Figure 13-23 shows the revision ID fields.



**Figure 13-23. Revision ID Configuration Register**


Table 13-27 shows the bit settings of the revision ID register.

**Table 13-27. Revision ID Configuration Register Field Descriptions**

Bits	Name	Description
7-0	RID	Revision ID. Specifies a revision code of the PCI controller. 8'h21 Revision 021.

### 13.3.3.6 Standard Programming Interface Configuration Register

Figure 13-24 shows the standard programming interface fields. This is the lower byte of the class code.



**Figure 13-24. Standard Programming Interface Configuration Register**

Table 13-28 shows the bit settings of the standard programming interface register.

**Table 13-28. Standard Programming Interface Configuration Register Field Descriptions**

Bits	Name	Description
7-0	PI	Programming interface. This field is hard-wired to 0x00.

### 13.3.3.7 Subclass Code Configuration Register

Figure 13-25 shows the subclass code fields. This is the middle byte of the class code.




Figure 13-25. Subclass Code Configuration Register

Table 13-29 shows the bit settings of the subclass code register.

Table 13-29. Subclass Code Configuration Register Field Descriptions

Bits	Name	Description
7–0	SC	Sub-class code. This field is hard-wired to 0x20, indicating a Power PC processor.

### 13.3.3.8 Base Class Code Configuration Register

Figure 13-26 shows the base class code fields. This is the upper byte of the class code.




Figure 13-26. Base Class Code Configuration Register


Table 13-30 shows the bit settings of the class code register.

Table 13-30. Class Code Configuration Register Field Descriptions

Bits	Name	Description
7–0	BCC	Base class code. This field is hard-wired to 0x0B, indicating a processor.

### 13.3.3.9 Cache Line Size Configuration Register

Figure 13-27 shows the cache line size fields.



**Figure 13-27. Cache Line Size Configuration Register**

Table 13-31 shows the bit settings of the cache line size register.

**Table 13-31. Cache Line Size Configuration Register Field Descriptions**

Bits	Name	Description
7–0	CLS	Cache line size. Cache-line in terms of 32-bit words. Although the register is writable, only the value 0x08 is legal.

### 13.3.3.10 Latency Timer Configuration Register

Figure 13-28 shows the latency timer fields.



**Figure 13-28. Latency Timer Configuration Register**

Table 13-32 shows the bit settings of the latency timer register.

**Table 13-32. Latency Timer Configuration Register Field Descriptions**

Bits	Name	Description
7–3	LT	Latency timer. Specifies a granularity of 8 PCI clocks, the length of time that the PCI controller, when mastering a transaction, may hold the bus as the result of a bus grant. Refer to the PCI 2.3 specification for the rules by which the PCI controller completes transactions when the timer has expired.
2–0	—	Reserved

### 13.3.3.11 Header Type Configuration Register

Figure 13-29 shows the read-only header type register, which is hard-wired to 0x00.




Figure 13-29. Header Type Configuration Register

### 13.3.3.12 BIST Control Configuration Register

Figure 13-30 shows the read-only BIST control register, which is hard-wired to 0x00.




Figure 13-30. BIST Control Configuration Register

### 13.3.3.13 PIMMR Base Address Configuration Register

Figure 13-31 shows the PIMMR base address register fields.




Figure 13-31. PIMMR Base Address Configuration Register

Table 13-33 shows the bit settings of the PIMMR base address register.

Table 13-33. PIMMR Base Address Configuration Register Field Descriptions

Bits	Name	Description
31–20	BA	Base address. Defines the base address for the internal (on-chip) memory-mapped register space. The size of this space is 1 MB.
19–4	—	Reserved
3	PRE	Prefetchable. Hard-wired to 0.


**Table 13-33. PIMMR Base Address Configuration Register Field Descriptions (continued)**

Bits	Name	Description
2–1	T	Type. Hard-wired to 00.
0	MSI	Memory space indicator. Hard-wired to 0

### 13.3.3.14 GPL Base Address Register 0

The GPL base address register 0 is provided to allow access to local memory space. This register is closely tied to PIBAR0 and PIWAR0 in the CSR memory space. A write to GPL base address register 0 also causes a change in the base address bits that are not masked according to the IWS field of PIWAR0 in PIBAR0. Note that this write operation will not change the bits that are masked by the IWS field. For read operation these masked bits will always return zeros.

Figure 13-32 shows the GPL base address register 0 fields.



**Figure 13-32. GPL Base Address Register 0**

Table 13-34 shows the bit settings of the GPL base address register 0.

**Table 13-34. GPL Base Address Register 0 Field Descriptions**

Bits	Name	Description
31–12	BA	Base address. Defines the base address for the inbound window. Bits 11–4 are hard-wired to 0 since the minimum window size is 4 Kbytes.
3	PRE	Prefetchable. This bit is read-only and contains the value of the PF bit in PIWAR0.
2–1	T	Type. Hard-wired to 00.
0	MSI	Memory space indicator. Hard-wired to 0

### 13.3.3.15 GPL Base Address Registers 1–2

The general purpose local access base address registers are provided to allow access to local memory space. These registers are closely tied to PIBAR<sub>n</sub> and PIWAR<sub>n</sub> in the CSR memory space. A write to a GPL base address register also causes a change in the base address bits that are not masked according to the IWS field of PIWAR<sub>n</sub> in the corresponding PIBAR<sub>n</sub>. Note that this write operation will not change the bits that are masked by the IWS field. For read operations, these masked bits always return zeros.

Figure 13-33 shows the GPL base address register 1–2 fields.






Figure 13-33. GPL Base Address Registers 1–2

Table 13-35 shows the bit settings of the GPL base address register 1–2.

Table 13-35. GPL Base Address Register 1,2 Field Descriptions

Bits	Name	Description
31–12	BA	Base address. Defines the two portion of the base address for the inbound window. Bits 11–4 are hard-wired to 0 since the minimum window size is 4 Kbytes.
11–4	—	Reserved
3	PRE	Prefetchable. This bit is read-only and contains the value of the PF bit in PIWAR $n$ .
2–1	T	Type. Hard-wired to 10.
0	MSI	Memory space indicator. Hard-wired to 0.

### 13.3.3.16 GPL Extended Base Address Registers 1–2

Two general-purpose local access base address registers are provided to allow access to local memory space. These registers are closely tied to PIBAR $n$ , PIEBAR $n$ , and PIWAR $n$  in the CSR memory space. A write to a GPL extended base address register also causes a change in the base address bits that are not masked according to the IWS field of PIWAR $x$  in the corresponding PIBAR $n$ /PIEBAR $n$ . Note that this write operation does not change bits that are masked by the IWS field. For read operations these masked bits will always return zeros.

Figure 13-34 shows the GPL extended base address registers 1–2 fields.



Figure 13-34. GPL Extended Base Address Registers 1–2

Table 13-36 shows the bit settings of the GPL extended base address register 1–2.

Table 13-36. GPL Extended Base Address Registers 1–2 Field Descriptions

Bits	Name	Description
31–0	EBA	Extended base address. Defines the high portion of the base address for the inbound window.

### 13.3.3.17 Subsystem Vendor ID Configuration Register

Figure 13-35 shows the subsystem vendor ID fields. The subsystem vendor ID configuration register is read-only from the PCI bus, but it can be programmed from the CSB.




Figure 13-35. Subsystem Vendor ID Configuration Register

Table 13-37 shows the bit settings of the subsystem vendor ID configuration register.

Table 13-37. Subsystem Vendor ID Configuration Register Field Descriptions

Bits	Name	Description
15–0	SVID	Subsystem vendor ID. Identifies the manufacturer of the board or subsystem that contains this device.

### 13.3.3.18 Subsystem Device ID Configuration Register

Figure 13-36 shows the subsystem device configuration register ID fields. The subsystem device ID configuration register is read-only from the PCI bus, but it can be programmed from the CSB.



Figure 13-36. Subsystem Device ID Configuration Register

Table 13-38 shows the bit settings of the subsystem device ID configuration register.

Table 13-38. Subsystem Device ID Configuration Register Field Descriptions

Bits	Name	Description
15–0	SDID	Subsystem device ID. This field identifies the board or subsystem that contains this device.

### 13.3.3.19 Capabilities Pointer Configuration Register

The capabilities pointer register specifies the byte offset in the PCI configuration space that contains the first item in the capabilities list. Figure 13-37 shows the capabilities pointer configuration register fields.




Figure 13-37. Capabilities Pointer Configuration Register

### 13.3.3.20 Interrupt Line Configuration Register

Figure 13-38 shows the interrupt line configuration register fields.




Figure 13-38. Interrupt Line Configuration Register

Table 13-39 shows the bit settings of the interrupt line configuration register.

Table 13-39. Interrupt Line Configuration Register Field Descriptions

Bits	Name	Description
7-0	IL	Interrupt line. Used to communicate interrupt line routing information. The value has no effect on the operation of the PCI controller.

### 13.3.3.21 Interrupt Pin Configuration Register

The interrupt pin configuration register tells which interrupt pin is used (0x01 means PCI\_INTA).

Figure 13-39 shows the interrupt pin configuration register fields.




Figure 13-39. Interrupt Pin Register

### 13.3.3.22 Minimum Grant Configuration Register

Figure 13-40 shows the minimum grant configuration register fields.

Offset 0x3E




Figure 13-40. Minimum Grant Configuration Register

### 13.3.3.23 Maximum Latency Configuration Register

Figure 13-41 shows the maximum latency configuration register fields.

Offset 0x3F




Figure 13-41. Maximum Latency Configuration Register

### 13.3.3.24 PCI Function Configuration Register

Figure 13-42 shows the PCI function configuration register fields.

Offset 0x44

Access: Read/Write




Figure 13-42. PCI Function Configuration Register

Table 13-40 shows the bit settings of the PCI function configuration register.

Table 13-40. PCI Function Configuration Register Field Descriptions

Bits	Name	Description
15–6	—	Reserved
5	CFG_LOCK	Configuration lock. Controls access to the PCI configuration space from the PCI port. In host mode the PCI configuration space is always inaccessible, so this bit is not used. Normally, this bit will be cleared in agent mode once the configuration of the PCI controller is complete to allow an external host to access the PCI configuration space. 0 Access to the configuration spaces is permitted. 1 Any inbound PCI access to the PCI configuration space is retried. See <a href="#">Section 4.3.1.1, “Reset Configuration Word Source,”</a> for more information on reset configuration.
3–4	—	Reserved

**Table 13-40. PCI Function Configuration Register Field Descriptions (continued)**

Bits	Name	Description
2	TLTD	Target latency timeout disable. Determines whether the PCI controller, while acting as a PCI target, times out when the first data phase of a transaction has not completed in 16 PCI cycles. 0 Target latency timeout enabled. 1 Target latency timeout disabled.
1	MLTD	Master latency timer disable. Determines whether the PCI controller, while acting as a PCI master, terminates a transaction upon the expiration of the master latency timer. 0 Master latency timer enabled. 1 Master latency timer disabled.
0	HA	Host/Agent. Indicates whether the PCI controller is in host mode or agent mode. It provides the value of the <code>PCI_HOST</code> —PCI host configuration bit is sampled at the end of the reset sequence. 0 Host mode 1 Agent mode

### 13.3.3.25 PCI Arbiter Control Register (PCIACR)

Figure 13-43 shows the PCI arbiter control register (PCIACR) fields.


**Figure 13-43. PCI Arbiter Control Register (PCIACR)**

Table 13-41 shows the bit settings of the PCIACR.

**Table 13-41. PCI Arbiter Control Register (PCIACR) Field Descriptions**

Bits	Name	Description
15	AD	Arbiter disable. Indicates whether the PCI controller functions as the arbiter for the PCI bus. It provides the value of the PCI arbiter enable configuration bit as sampled at the end of the reset sequence. See <a href="#">Chapter 4, “Reset, Clocking, and Initialization,”</a> for more information on reset configuration. 0 Arbiter enabled 1 Arbiter disabled
14	PM	Parking mode. Controls which device receives a bus grant when there are no outstanding bus requests and the bus is idle. 0 The bus is parked with the last device to use the bus. 1 The bus is parked with the PCI controller.
13	—	Reserved
12	PBMD	PCI broken master disable. Determines whether the PCI controller ignores the bus requests of an initiator that requests the bus for an excessive period without using it. 0 An initiator that requests the bus and receives the grant must begin using the bus within 16 PCI clock periods after the bus becomes idle or its request is subsequently ignored. 1 No requests are ignored.
11–7	—	Reserved


**Table 13-41. PCI Arbiter Control Register (PCIACR) Field Descriptions (continued)**

Bits	Name	Description
6–4	PRIn	Priority level for master <i>n</i> . When the PCI controller functions as the arbiter for the PCI bus, each PRIn bit determines the arbitration priority level for the PCI master connected to the REQn/GNTn pair. 0 Low priority 1 High priority
3–1	—	Reserved
0	MPRI	My priority. When the PCI controller functions as the arbiter for the PCI bus, this bit determines the arbitration priority level for the PCI controller when it acts as a PCI master. 0 Low priority 1 High priority

### 13.3.3.26 Hot Swap Register Block

Figure 13-44 shows the hot swap register block fields.

Offset 0x48



**Figure 13-44. Hot Swap Register Block**

Table 13-42 shows the bit settings of the Hot Swap register block.

**Table 13-42. Hot Swap Register Block Field Descriptions**

Bits	Name	Description
31–24	—	Reserved
23	INS	Insertion status. Indicates that a card has been inserted. Write 1 to clear this bit.
22	EXT	Extraction status. Indicates that a card has been extracted. Write 1 to clear this bit.
21–20	—	Reserved
19	LOO	LED On/Off. Controls the LED when the hardware is in state H2 0 LED off 1 LED on
18	—	Reserved
17	EIM	ENUM mask. This bit masks the CPCI_HS_ENUM input. 0 Enabled 1 Masked
16	—	Reserved
15–8	NXT_PTR	Next pointer—hardwired to 0x80 to point to the address of the power management capability in the PCI controller.
7–0	CAP_ID	Capability ID for hot swap (hardwired to 0x06)

### 13.3.3.27 PCI Power Management Register 0 (PCIPMR0)

The PCI power management register 0 (PCIPMR0), shown in Figure 13-45, indicates the power management policies to implement in the system.




Figure 13-45. PCI Power Management Register 0 (PCIPMR0)

Table 13-43 describes the PCIPMR0 fields.

Table 13-43. PCIPMR0 Field Descriptions

Bits	Name	Description
31–27	PME_Support	Indicates the power states in which the PCI Controller may assert PME#. PME_Support (bit 27): 0 PME# cannot be asserted from D0 1 PME# can be asserted from D0 PME_Support (bit 28): 0 PME# cannot be asserted from D1 1 PME# can be asserted from D1 PME_Support (bit 29): 0 PME# cannot be asserted from D2 1 PME# can be asserted from D2 PME_Support (bit 30): 0 PME# cannot be asserted from D3_hot 1 PME# can be asserted from D3_hot PME_Support (bit 31): 0 PME# cannot be asserted from D3_cold 1 PME# can be asserted from D3_cold
26	D2_Support	D2 power management state support 0 The PCI controller does not support D2 power management state. 1 The PCI controller supports the D2 power management state.
25	D1_Support	D1 power management state support 0 The PCI controller does not support D1 power management state. 1 The PCI controller supports the D1 power management state.
24–22	Aux_Current	Reports the 3.3 Vaux auxiliary current requirements
21	DSI	Device specific initialization. Indicates whether special initialization of this PCI controller is required.
20	—	Reserved

**Table 13-43. PCIPMR0 Field Descriptions (continued)**


Bits	Name	Description
19	PME_Clock	PME clock 0 Indicates that no PCI clock is required for the PCI controller to generate PME# 1 Indicates that PCI clock is required for the PCI controller to generate PME#
18–16	Version	PCI Power Management Interface Specification version. 011 Revision 1.2 of the PCI Power Management Interface Specification
15–8	NEXT_CAP_PTR	The next capability pointer points to the next item in the PCI controller's capability list. 0000_0000 The end of the capability list
7–0	CAP_ID	0000_0001 Indicates the power management support capability

### 13.3.3.28 PCI Power Management Register 1 (PCIPMR1)

The PCI power management register 1 (PCIPMR1), shown in Figure 13-46, contains the bit fields that software uses to manage the PCI controller's power management state, as well as to enable and monitor PMEs (power management events). This register can be accessed by the host in agent mode.

Offset 0x84

Access: Read / write



**Figure 13-46. PCI Power Management Register 1 (PCIPMR1)**

Table 13-44 describes the PCIPMR1 fields.

**Table 13-44. PCIPMR1 Field Descriptions**

Bits	Name	Description
31–24	Data	Reports the state dependent data requested by the Data_Select field. The value of this field is scaled by the value reported by the Data_Scale field.
23	BPCC_En	Bus power/Clock control enable 0 Disable the bus power/clock control policies defined in section 4.7.1 of the PCI Bus Power Management Interface Specification Revision 1.2 1 Enable the bus power/clock control policies defined in section 4.7.1 of the PCI Bus Power Management Interface Specification Revision 1.2 Note: This bit field is not implemented, only required for all PCI-to-PCI Bridge



Table 13-44. PCIPMR1 Field Descriptions (continued)

Bits	Name	Description
22	B2_B3	The state of this bit determines what will happen as a direct result of programming the function to D3_hot. 0 Indicates that when the bridge function is programmed to D3_hot, its secondary bus will have its power removed (B3). 1 Indicates that when the bridge function is programmed to D3_hot, its secondary bus's PCI clock will be stopped (B2). This bit only meaningful if bit 16 (BPCC_En) is set. Note: This bit field is not implemented, only required for all PCI-to-PCI Bridge
21–16	—	Reserved
15	PME_Status	This bit set when the PCI controller would normally assert PME# signal independent of the state of the PME_En bit. Writing a value of one to this bit will clear it and cause the PCI controller to stop asserting a PME# signal. 0 Default 1 If (Wake_Up & PME_En)
14–13	Data_Scale	The scale factor to be used when interpreting the value of the data register
12–9	Data_Select	Selects which data is to be reported through the data register and Data_Scale field
8	PME_En	Enables the function to assert PME# 0 Disables the function to assert PME# 1 Enables the function to assert PME#
7–4	—	Reserved
3	No_Soft_Reset	This bit field indicates whether an internal reset occurs during the transition from D3_hot to D0. 0 The Power_State command performs an internal reset. 1 The Power_State command does not perform an internal reset.
2	—	Reserved
1–0	Power_State	Determines the current power state of the PCI controller and sets the controller into a new power state. The power state definition is as follows: 00 D0 supports all PCI function. 01 D1 disables the inbound memory space, bus mastering and functional interrupt request. 10 D2 disables the inbound memory space, bus mastering and functional interrupt request. 11 D3_hot disables the inbound memory space, bus mastering and functional interrupt request.

## 13.4 Functional Description

The following sections discuss the operation of the PCI controller.

### 13.4.1 PCI Bus Arbitration

The PCI bus arbitration approach is access-based. Bus masters must arbitrate for each access performed on the bus. PCI uses a central arbitration scheme where each master has its own unique request ( $\overline{REQn}$ ) output and grant ( $GNTn$ ) input signal. A simple request-grant handshake is used to gain access to the bus. Arbitration for the bus occurs during the previous access so that no PCI bus cycles are consumed waiting for arbitration (except when the bus is idle).

The PCI internal arbiter supports three external masters (besides the PCI controller itself) by using the  $\overline{\text{REQ}}$  signals and generating the  $\overline{\text{GNT}}$  signals.

During reset, the PCI controller samples the reset configuration bit (and programs the PCI\_ARB\_DIS bit accordingly) to determine if the arbiter is enabled or disabled. See [Chapter 4, “Reset, Clocking, and Initialization,”](#) for more information. The arbiter can also be enabled or disabled by directly programming the PCI\_ARB\_DIS bit in the arbiter configuration register (see [Section 13.3.3.25, “PCI Arbiter Control Register \(PCIACR\),”](#) for more information). However, it is recommended to use the reset configuration bit to set the arbiter state because the arbiter state controls the direction of  $\overline{\text{REQ0}}$  and  $\overline{\text{GNT0}}$ .

If the arbiter is disabled, the PCI controller uses  $\overline{\text{REQ0}}$  to issue requests to an external arbiter, and uses  $\overline{\text{GNT0}}$  to receive grants from the external arbiter.

### 13.4.1.1 Bus Parking

When no devices are requesting the bus, the bus is granted, or parked, for a specified device to prevent the AD, PCI\_C/ $\overline{\text{BE}}$  and PCI\_PAR signals from floating. The PCI controller can be configured to either park on itself or park on the last master to use the bus (see [Section 13.3.3.25, “PCI Arbiter Control Register \(PCIACR\),”](#) for more information).

### 13.4.1.2 Arbitration Algorithm

The round-robin arbitration algorithm has two priority levels. Each of the external PCI bus masters, plus the PCI controller, are assigned either a high or a low priority level, as programmed in the arbiter configuration register (see [Section 13.3.3.25, “PCI Arbiter Control Register \(PCIACR\).”](#)) Within each priority group (high or low), the bus grant is given to the next requesting device in numerical order, with the PCI controller itself positioned before device 0.  $\overline{\text{GNT}}_n$  is asserted for device  $n$  as soon as the previously granted device begins a transaction. Conceptually, the lowest priority device at any given time is the current bus master and the highest priority device is the next one to follow the current master. This is considered to be a fair algorithm because a given device cannot prevent other devices from having access to the bus—a given device automatically becomes the lowest priority device as soon as it begins to use the bus. If a master is not requesting the bus, the transaction slot is given to the next requesting device within the priority group.

The grant given to one device may be taken away and whenever a higher priority device asserts its request. If the bus is idle when a new device is to receive a grant, no device receives a grant for one clock; in the next clock, the new winner of the arbitration receives a grant. This operation allows for a turnaround clock when a device is using address stepping or when the bus is parked.

The low priority group collectively receives one bus transaction request slot in the high priority group. Therefore, if there are  $N$  high-priority devices, each high-priority device is guaranteed to get at least one of  $(N+1)$  bus transactions, and the  $M$  low priority devices are guaranteed to each get at least one of  $(N+1) \times M$  bus transactions, with one of the low-priority devices receiving the grant in one of  $(N+1)$  bus transactions. If all devices are programmed to the same priority level or if there is only one device at the low priority, the algorithm provides each device an equal number of bus grants in a round-robin sequence.

An arbitration example with three masters in the high priority group and two in the low priority group is shown in [Figure 13-47](#). Noting that one position in the high priority group is actually a place-holder for

the low priority group, it can be seen that each high priority initiator is guaranteed at least 1 out of 3 transaction slots, and each low priority initiator is guaranteed at least 1 out of 6 slots. Assuming all devices are requesting the bus, the grant sequence (with device 1 being the current master) is as follows: 0, 2, the PCI controller, 0, 2, 1, 0, 2, the PCI controller, and so on. If, for example, device 2 is not requesting the bus, the grant sequence becomes 0, the PCI controller, 0, 1, 0, the PCI controller, and so on. If device 2 now requests the bus at a point in the sequence when device 0 is conducting a transaction and the PCI controller is the next grant, then the PCI controller's grant is removed, and the higher-priority device 2 is awarded the next grant.




Figure 13-47. PCI Arbitration Example

### 13.4.1.3 Broken Master Lock-Out

The broken master feature allows the arbiter to lock out any masters that are broken or ill-behaved. This feature is controlled by programming the PCI arbiter control register. When the broken master feature is enabled, a granted device that does not assert  $\overline{\text{PCI\_FRAME}}$  within 16 PCI clock cycles after the bus is idle, has its grant removed and subsequent requests are ignored until its  $\overline{\text{REQ}}$  is negated for at least one clock cycle. This prevents ill-behaved masters from monopolizing the bus. When the broken master feature is disabled, a device that requests the bus and receives a grant never loses its grant until and unless it begins a transaction or negates its  $\overline{\text{REQ}}$  signal. Note that disabling the broken master feature is not recommended.

### 13.4.1.4 Master Latency Timer

The PCI controller implements the master latency timer register (see [Section 13.3.3.10, “Latency Timer Configuration Register”](#)) to prevent itself from monopolizing the bus. When the master latency timer expires, the PCI controller checks the state of its  $\overline{\text{PCI\_GNT}}$  signals. If the  $\overline{\text{PCI\_GNT}}$  signal is not asserted, the PCI controller completes one more data phase and relinquishes the bus. The master latency timer can be disabled if needed (see [Section 13.3.3.24, “PCI Function Configuration Register,”](#) for more information).

## 13.4.2 Bus Commands

PCI bus commands indicate the type of transaction occurring on the bus. These commands are encoded on PCI\_C/BE[3:0] during the address phase of the transaction. PCI bus commands are described in [Table 13-45](#).

**Table 13-45. PCI Command Definitions**

PCI_C/ BE[3:0]	Command Type	Supported as:		Definition
		Initiator	Target	
0b0000	Interrupt acknowledge	Yes	No	A read implicitly addressed to the system interrupt controller. The size of the vector to be returned is indicated on the byte enables after the address phase.
0b0001	Special cycle	Yes	No	Provides a simple message broadcast mechanism. See <a href="#">Section 13.4.4.6, "Special Cycle Command,"</a> for more information.
0b0010	I/O read	Yes	No	Accesses agents mapped in I/O address space.
0b0011	I/O write	Yes	No	Accesses agents mapped in I/O address space.
0b010x	—	—	—	Reserved. No response occurs.
0b0110	Memory read	Yes	Yes	Accesses agents mapped in memory address space. A read from prefetchable space, when seen as a target, fetches a cache line of data (32 bytes) from the starting address, even though all 32 bytes may not actually be sent to the initiator.
0b0111	Memory write	Yes	Yes	Accesses agents mapped in memory address space. Note that for inbound writes less than 4-bytes, the PCI controller splits the transaction into single byte writes to the target. Thus, the PCI interface cannot be used to perform single beat writes to 16-bit devices on the local peripheral interfaces.
0b100x	—	—	—	Reserved. No response occurs.
0b1010	Configuration read	Yes	Yes	Accesses the configuration space of each agent. An agent is selected when its IDSEL signal is asserted. See <a href="#">Section 13.4.4.4, "Host Mode Configuration Access,"</a> for more information on configuration accesses. As a target, a configuration read is only accepted if the PCI controller is configured to be in agent mode.
0b1011	Configuration write	Yes	Yes	Accesses the configuration space of each agent. An agent is selected when its IDSEL signal is asserted. See <a href="#">Section 13.4.4.4, "Host Mode Configuration Access,"</a> for more information. As a target, a configuration write is only accepted if the PCI controller is configured to be in agent mode.
0b1100	Memory read multiple	Yes	Yes	Causes a prefetch of the next cache line.
0b1101	Dual address cycle	No	Yes	Transfers an 8-byte address to devices.
0b1110	Memory read line	Yes	Yes	Indicates that the initiator intends to transfer an entire cache line of data.
0b1111	Memory write and invalidate	No	Yes	Indicates that the initiator will transfer an entire cache line of data, and if PCI has any cacheable memory, this line needs to be invalidated.

### 13.4.3 PCI Protocol Fundamentals

The bus transfer mechanism on the PCI bus is called a burst. A burst is comprised of an address phase and one or more data phases.

All signals are sampled on the rising edge of the PCI clock. Each signal has a setup and hold window with respect to the rising clock edge, in which transitions are not allowed. Outside this aperture, signal values or transitions have no significance.

#### 13.4.3.1 Basic Transfer Control

PCI data transfers are controlled by the following signals:

- $\overline{\text{PCI\_FRAME}}$  is driven by an initiator to indicate the beginning and end of a transaction.
- $\overline{\text{PCI\_IRDY}}$  (initiator ready) is driven by an initiator, allowing it to force wait cycles.
- $\overline{\text{PCI\_TRDY}}$  (target ready) is driven by a target, allowing it to force wait cycles.

The bus is idle when both  $\overline{\text{PCI\_FRAME}}$  and  $\overline{\text{PCI\_IRDY}}$  are negated. The first clock cycle in which  $\overline{\text{PCI\_FRAME}}$  is asserted indicates the beginning of the address phase. The address and the bus command code are transferred in that cycle. The next cycle ends the address phase and begins the data phase.

During the data phase, data is transferred in each cycle that both  $\overline{\text{PCI\_IRDY}}$  and  $\overline{\text{PCI\_TRDY}}$  are asserted. Once the PCI controller, as an initiator, has asserted  $\overline{\text{PCI\_IRDY}}$ , it does not change  $\overline{\text{PCI\_IRDY}}$  or  $\overline{\text{PCI\_FRAME}}$  until the current data phase completes, regardless of the state of  $\overline{\text{PCI\_TRDY}}$ . Once the PCI controller, as a target, has asserted  $\overline{\text{PCI\_TRDY}}$  or  $\overline{\text{PCI\_STOP}}$  it does not change  $\overline{\text{PCI\_DEVSEL}}$ ,  $\overline{\text{PCI\_TRDY}}$ , or  $\overline{\text{PCI\_STOP}}$  until the current data phase completes.

When the PCI controller (as a master) intends to complete only one more data transfer,  $\overline{\text{PCI\_FRAME}}$  is negated and  $\overline{\text{PCI\_IRDY}}$  is asserted (or kept asserted) indicating the initiator is ready. After the target indicates it is ready ( $\overline{\text{PCI\_TRDY}}$  asserted) the bus returns to the idle state.

#### 13.4.3.2 Addressing

The PCI specification defines three physical address spaces—memory, I/O, and configuration. The memory and I/O address spaces are standard for all systems. The configuration address space supports the PCI hardware configuration. Each PCI device decodes the address for each PCI transaction with each agent responsible for its own address decode.

The information contained in the two lower address bits (AD1 and AD0) depends on the address space. In the I/O address space, all 32 address/data lines provide the full byte address. AD[1:0] are used for the generation of  $\overline{\text{PCI\_DEVSEL}}$  and indicate the least significant valid byte involved in the transfer. Once a target has claimed an I/O access, it first determines if it can complete the entire access as indicated by the byte enable signals. If all the selected bytes are not in the address range, the entire access should not be completed; that is, the target should not transfer any data and should terminate the transaction with a target-abort operation. See [Section 13.4.3.6, “Bus Transactions,”](#) for more information.

In the configuration address space, accesses are decoded to a 4-byte address using AD[7:2]. An agent determines if it is the target of the access when a configuration command is decoded, IDSEL is asserted, and AD[1:0] are 0b00; otherwise, the agent ignores the current transaction. The PCI controller determines

a configuration access is for a device on the PCI bus by decoding a configuration command. When in agent mode, the PCI controller responds to host-generated PCI configuration cycles when its IDSEL is asserted during a configuration cycle.

For memory accesses, the address is decoded using AD[31:2]; thereafter, the address is incremented internally by 4 bytes until the end of the burst transfer. Another initiator in a memory access should drive 0b00 on AD[1:0] during the address phase to indicate a linear incrementing burst order. The PCI controller checks AD[1:0] during a memory command access and provides the linear incrementing burst order. On reads, if AD[1:0] is 0b10, which represents a cache line wrap, the PCI controller linearly increments the burst order starting at the critical 64-bit address, wraps at the end of the cache line, and disconnects after reading one cache line. If AD[1:0] is 0bx1 (a reserved encoding) and the PCI\_C/BE[3:0] signals indicate a memory transaction, it executes a target disconnect after the first data phase is completed. Note that AD[1:0] are included in parity calculations.

### 13.4.3.3 Device Selection

As a target, the PCI controller drives  $\overline{\text{PCI\_DEVSEL}}$  one clock following the address phase as indicated in the configuration space status register; see [Section 13.3.3.4, “PCI Status Configuration Register,”](#) for more information. The PCI controller as a target qualifies the address/data lines with  $\overline{\text{PCI\_FRAME}}$  before asserting  $\overline{\text{PCI\_DEVSEL}}$ . The  $\overline{\text{PCI\_DEVSEL}}$  signal is asserted at or before the clock edge at which the PCI controller enables its  $\overline{\text{PCI\_TRDY}}$ ,  $\overline{\text{PCI\_STOP}}$ , or data (for a read). The  $\overline{\text{PCI\_DEVSEL}}$  signal is not negated until  $\overline{\text{PCI\_FRAME}}$  is negated, with  $\overline{\text{PCI\_IRDY}}$  asserted and either  $\overline{\text{PCI\_STOP}}$  or  $\overline{\text{PCI\_TRDY}}$  asserted. The exception to this is a target-abort; see [Section 13.4.3.8, “Transaction Termination,”](#) for more information.

As an initiator, if the PCI controller does not see the assertion of  $\overline{\text{PCI\_DEVSEL}}$  within 4 clocks of  $\overline{\text{PCI\_FRAME}}$ , it terminates the transaction with a master-abort as described in [Section 13.4.3.8, “Transaction Termination,”](#) for more information.

### 13.4.3.4 Byte Enable Signals

The byte enable signals ( $\overline{\text{BE}}[3:0]$ ) indicate which byte lanes carry valid data. The byte enable signals may enable different bytes for each of the data phases. The byte enable signals are valid on the edge of the clock that starts each data phase and remain valid for the entire data phase.

If the PCI controller, as a target, sees no byte enable signals asserted, it completes the current data phase with no permanent change. This implies that on a read transaction, the PCI controller expects the data not to be changed, and on a write transaction, the data is not stored.

### 13.4.3.5 Bus Driving and Turnaround


The turnaround-cycle is one clock cycle and is required to avoid contention. This cycle occurs at different times for different signals.  $\overline{\text{PCI\_IRDY}}$ ,  $\overline{\text{PCI\_TRDY}}$ , and  $\overline{\text{PCI\_DEVSEL}}$  use the address phase as their turnaround-cycle.  $\overline{\text{PCI\_FRAME}}$ ,  $\overline{\text{PCI\_C/BE}}[3:0]$ , and AD[31:0] use the idle cycle between transactions as their turnaround-cycle. (An idle cycle in PCI is when both  $\overline{\text{PCI\_FRAME}}$  and  $\overline{\text{PCI\_IRDY}}$  are negated).

Byte lanes not involved in the current data transfer are driven to a stable condition even though the data is not valid.

### 13.4.3.6 Bus Transactions

The timing diagrams in this section show the relationship of significant signals involved in bus transactions.

Note the following conventions:

- When a signal is drawn as a solid line, it is actively being driven by the current initiator or target.
- When a signal is drawn as a dashed line, no agent is actively driving it.
- Three-stated signals with slashes between the two rails have indeterminate values.
- The terms ‘edge’ and ‘clock edge’ refer to the rising edge of the clock.
- The terms ‘asserted’ and ‘negated’ refer to the globally visible state of the signal on the clock edge, and not to signal transitions.
- The symbol  represents a turnaround-cycle.

### 13.4.3.7 Read and Write Transactions

Both read and write transactions begin with an address phase followed by a data phase. The address phase occurs when  $\overline{\text{PCI\_FRAME}}$  is asserted for the first time, and the AD[31:0] signals contain a byte address and the PCI\_C/ $\overline{\text{BE}}$ [3:0] signals contain a bus command. The data phase consists of the actual data transfer and possible wait cycles; the byte enable signals remain actively driven from the first clock of the data phase through the end of the data transfer.

A read transaction starts when  $\overline{\text{PCI\_FRAME}}$  is asserted for the first time and the PCI\_C/ $\overline{\text{BE}}$ [3:0] signals indicate a read command. Figure 13-48 shows an example of a single beat read transaction.

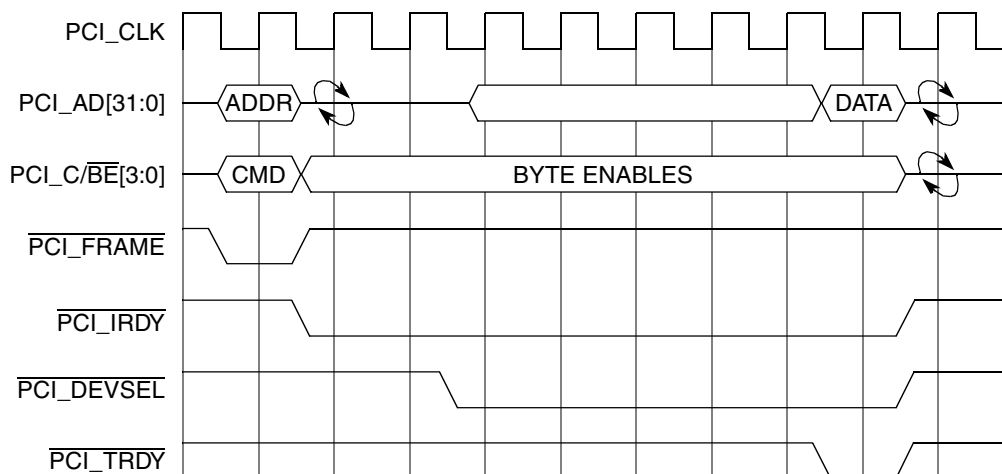


Figure 13-48. Single Beat Read Example

Figure 13-49 shows an example of a burst read transaction.

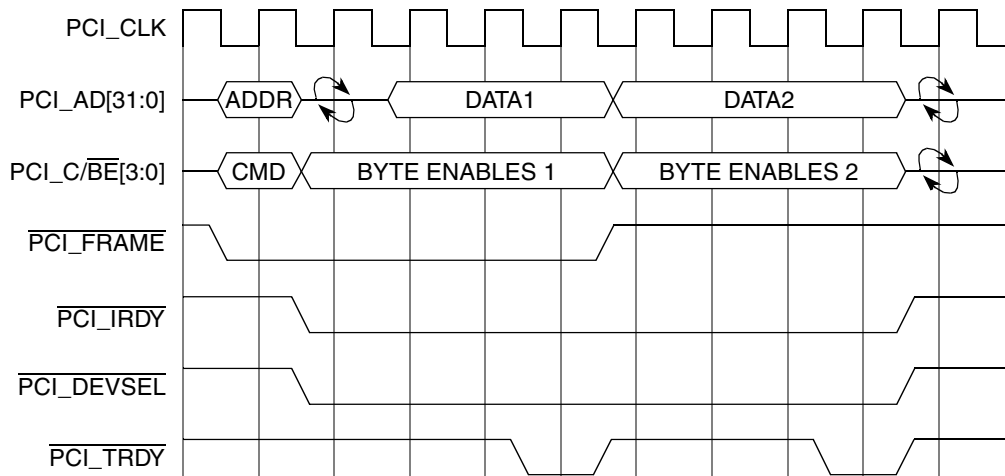


Figure 13-49. Burst Read Example

During the turnaround-cycle following the address phase, the PCI\_C/BE[3:0] signals indicate which byte lanes are involved in the data phase. The turnaround-cycle must be enforced by the target with the PCI\_TRDY signal if using fast PCI\_DEVSEL assertion. The earliest the target can provide valid data is one cycle after the turnaround cycle. The target must drive the AD[31:0] signals when PCI\_DEVSEL is asserted except during the turnaround cycle.

The data phase completes when data is transferred, which occurs when both PCI\_IRDY and PCI\_TRDY are asserted on the same clock edge. When either is negated, a wait cycle is inserted and no data is transferred. To indicate the last data phase PCI\_IRDY must be asserted when PCI\_FRAME is negated.

A write transaction starts when PCI\_FRAME is asserted for the first time and the PCI\_C/BE[3:0] signals indicate a write command. Figure 13-50 shows an example of a single-beat write transaction.

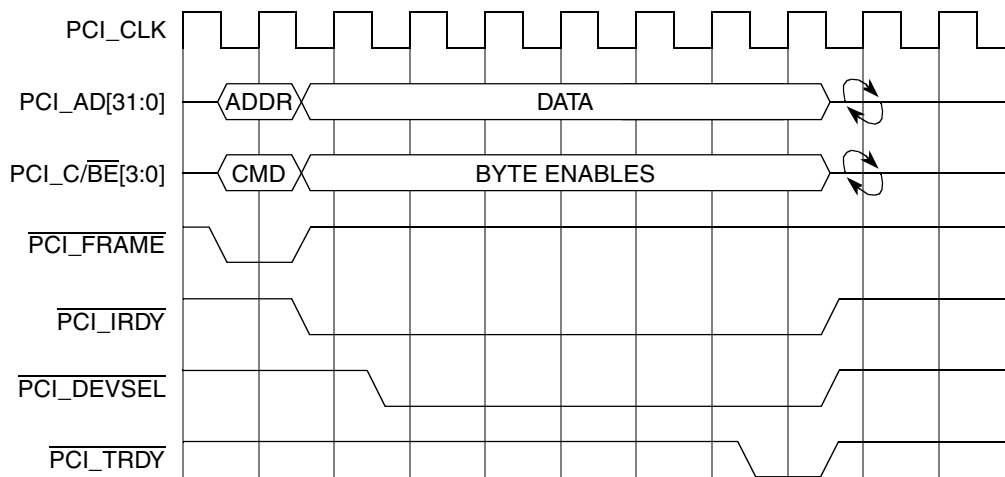


Figure 13-50. Single Beat Write Example



Figure 13-51 shows an example of a burst write transaction.

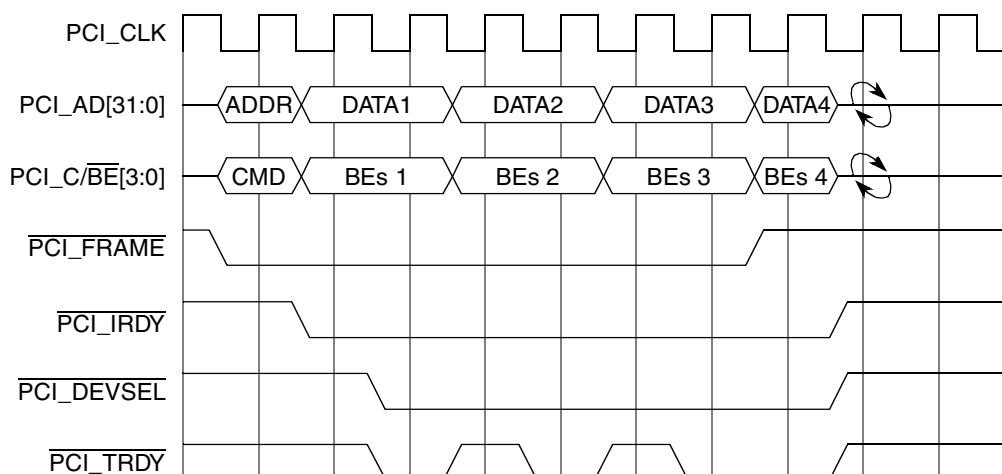


Figure 13-51. Burst Write Example

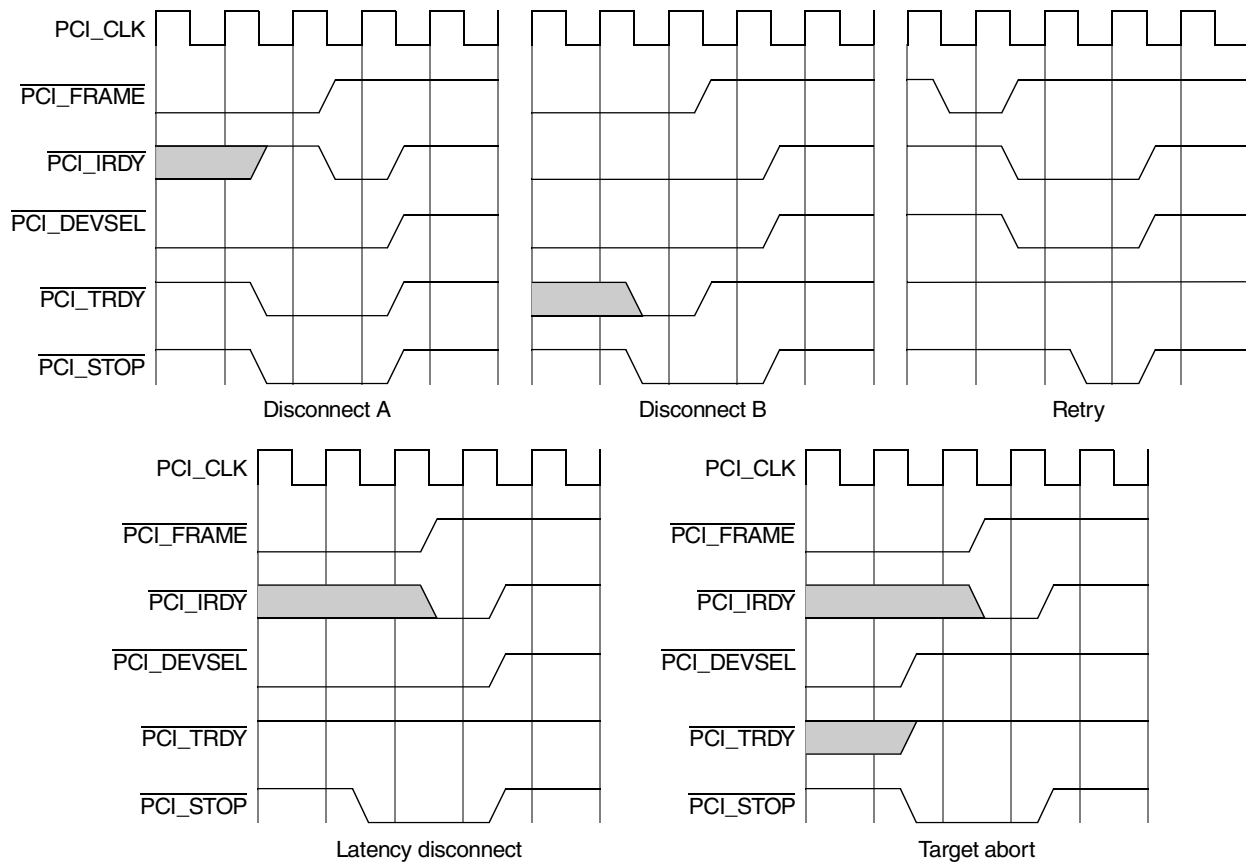
A write transaction is similar to a read transaction except no turnaround cycle is needed following the address phase because the initiator provides both address and data. Data phases are the same for both read and write transactions.

### 13.4.3.8 Transaction Termination

The termination of a PCI transaction is orderly and systematic, regardless of the cause of the termination. All transactions end when  $\overline{\text{PCI\_FRAME}}$  and  $\overline{\text{PCI\_IRDY}}$  are both negated, indicating the idle cycle.

The PCI controller as an initiator terminates a transaction when  $\overline{\text{PCI\_FRAME}}$  is negated and  $\overline{\text{PCI\_IRDY}}$  is asserted. This indicates that the final data phase is in progress. The final data transfer occurs when both  $\overline{\text{PCI\_TRDY}}$  and  $\overline{\text{PCI\_IRDY}}$  are asserted. A master-abort is an abnormal case of a master initiated termination. If the PCI controller detects that  $\overline{\text{PCI\_DEVSEL}}$  has remained negated for more than four clocks after the assertion of  $\overline{\text{PCI\_FRAME}}$ , it negates  $\overline{\text{PCI\_FRAME}}$  and then, on the next clock, negates  $\overline{\text{PCI\_IRDY}}$ . On aborted reads, the PCI controller returns 0xFFFF\_FFFF. The data is lost on aborted writes.

When the PCI controller as a target needs to suspend a transaction, it asserts  $\overline{\text{PCI\_STOP}}$ . Once asserted,  $\overline{\text{PCI\_STOP}}$  remains asserted until  $\overline{\text{PCI\_FRAME}}$  is negated. Depending on the circumstances, data may or may not be transferred during the request for termination. If  $\overline{\text{PCI\_TRDY}}$  and  $\overline{\text{PCI\_IRDY}}$  are asserted during the assertion of  $\overline{\text{PCI\_STOP}}$ , data is transferred. This type of target-initiated termination is called a disconnect B, shown in Figure 13-52. If  $\overline{\text{PCI\_TRDY}}$  is asserted when  $\overline{\text{PCI\_STOP}}$  is asserted but  $\overline{\text{PCI\_IRDY}}$  is not,  $\overline{\text{PCI\_TRDY}}$  must remain asserted until  $\overline{\text{PCI\_IRDY}}$  is asserted and the data is transferred. This is called a disconnect A target-initiated termination, also shown in Figure 13-52. However, if  $\overline{\text{PCI\_TRDY}}$  is negated when  $\overline{\text{PCI\_STOP}}$  is asserted, no more data is transferred, and the initiator therefore does not have to wait for a final data transfer (see the retry diagram in Figure 13-50).



**Figure 13-52. Target-Initiated Terminations**

Note that when an initiator is terminated by  $\overline{\text{PCI\_STOP}}$ , it must negate its  $\overline{\text{REQn}}$  signal for a minimum of two PCI clocks (of which one clock is needed for the bus to return to the idle state). If the initiator intends to complete the transaction, it should reassert its  $\overline{\text{REQn}}$  immediately following the two clocks or potential starvation may occur. If the initiator does not intend to complete the transaction, it can assert  $\overline{\text{REQn}}$  whenever it needs to use the PCI bus again.

The PCI controller terminates a transaction in the following cases:

- Eight PCI clock cycles have elapsed between data phases. This is a ‘latency disconnect’ (see [Figure 13-50](#)).
- AD[1:0] is 0bx1 (a reserved burst ordering encoding) during the address phase and one data phase has completed.
- The PCI command is a configuration command and one data phase has completed.
- A streaming transaction crosses a 4-Kbyte page boundary.
- A streaming transaction runs out of I/O sequencer buffer entries.
- A cache line wrap transaction has completed a cache line transfer.

Another target-initiated termination is the retry termination. Retry refers to termination requested because the target is currently in a state where it is unable to process the transaction. This can occur because no buffer entries are available in the I/O sequencer, or the sixteen clock latency timer has expired without

transfer of the first data. The target latency timer of the PCI controller can be optionally disabled. See [Section 13.3.3.24, “PCI Function Configuration Register,”](#) for more information.

When the PCI controller is in host mode it does not respond to any PCI configuration transactions. When the PCI controller is in agent mode and the CFG\_LOCK lock bit is set (see [Section 13.3.3.24, “PCI Function Configuration Register”](#)) the PCI controller retries all transactions to the PCI configuration space or the internal (on-chip) memory-mapped register space. Note that all retried accesses need to be completed. An example of a retry is shown in [Figure 13-50](#).

Note that because a target can determine whether or not data is transferred (when both  $\overline{\text{PCI\_IRDY}}$  and  $\overline{\text{PCI\_TRDY}}$  are asserted), if it wants to do only one more data transfer and then stop, it may assert  $\overline{\text{PCI\_TRDY}}$  and  $\overline{\text{PCI\_STOP}}$  at the same time.

Target-abort refers to the abnormal termination that is used when a fatal error has occurred, or when a target will never be able to respond. Target-abort is indicated when  $\overline{\text{PCI\_STOP}}$  is asserted and  $\overline{\text{PCI\_DEVSEL}}$  is negated. This indicates that the target requires the transaction to be terminated and does not want the transaction tried again. Note that any transferred data may have been corrupted.

The PCI controller terminates a transaction with target-abort in the case in which it is the intended target of a read transaction from system memory and the data from memory is corrupt. If the PCI controller is the intended target of a transaction and an address parity error occurs, or a data parity error occurs on a write transaction to system memory, it continues the transaction on the PCI bus but aborts internally. The PCI controller does not target-abort in this case.

If the PCI controller is mastering a transaction that terminates with a target-abort, undefined data is returned on a read and write data is lost. An example of a target-abort is shown in [Figure 13-50](#).

An initiator may retry any target disconnect accesses, except target-abort, at a later time starting with the address of the next non-transferred data. Retry is actually a special case of disconnect where no data transfer occurs at all and the initiator must start the entire transaction over again.

## 13.4.4 Other Bus Operations

The following sections provide information on additional PCI bus operations.

### 13.4.4.1 Fast Back-to-Back Transactions

In the two types of fast back-to-back transactions, the first type places the burden of avoiding contention on the initiator while the second places the burden on all potential targets. The PCI controller as a target supports both types of fast back-to-back transactions but does not support them as an initiator. The PCI controller as a target has the fast back-to-back enable bit hardwired to one, that is, enabled.

For the first type (governed by the initiator), the initiator may only run a fast back-to-back transaction to the same target. For the second type, when the PCI controller detects a fast-back-to-back operation and did not drive  $\overline{\text{PCI\_DEVSEL}}$  in the previous cycle, it delays the assertion of  $\overline{\text{PCI\_DEVSEL}}$  and  $\overline{\text{PCI\_TRDY}}$  for one cycle to allow the other target to get off the bus.

### 13.4.4.2 Dual Address Cycles

The PCI controller supports dual address cycle (DAC) commands (64-bit addressing on PCI bus) as a target only. DACs are different from single address cycles (SACs) in that the address phase takes two PCI beats instead of one PCI beat to transfer (64-bit vs. 32-bit addressing). Only PCI memory commands can use DAC cycles; I/O, configuration, interrupt acknowledge, and special cycle command cannot use DAC cycles. The PCI controller supports single-beat and burst DAC transactions.

### 13.4.4.3 Data Streaming

The PCI controller provides data streaming for PCI transactions to and from prefetchable memory. In other words, when the PCI controller is a target for a PCI initiated transaction, it supplies or accepts multiple cache lines of data without disconnecting. For PCI transactions to non-prefetchable space, the PCI controller disconnects after the first data phase so streaming cannot occur.

For PCI memory reads, streaming is achieved by performing speculative reads from memory in prefetchable space. A block of memory may be marked as prefetchable by setting the PCI configuration registers bit for the inbound address translation (see [Section 13.3.2.14, “PCI Inbound Window Attribute Registers \(PIWARn\),”](#) for more information) in the following cases:

- When reads do not alter the contents of memory (reads have no side effects)
- When reads return all bytes regardless of the byte enable signals
- When writes can be merged without causing errors

For a memory read command or a memory read line command, the PCI controller reads one cache line from memory. If the transaction crosses a cache line boundary, the PCI controller starts the read of a new cache line. For a memory read multiple command, the PCI controller reads two cache lines from memory. When the PCI transaction finishes the read for the first cache line, the PCI controller performs a speculative read of a third cache line. The PCI controller continues this prefetching until the end of the transaction.

For PCI writes to memory, streaming is achieved by buffering the transaction in the space available within the I/O sequencer. This allows PCI memory writes to execute with no wait states.

A disconnect occurs if the PCI controller runs out of buffer space on writes, or the PCI controller cannot supply consecutive data beats for reads within eight PCI bus clocks of each other. A disconnect also occurs if the transaction crosses a 4-Kbyte page boundary.

### 13.4.4.4 Host Mode Configuration Access

The PCI controller provides two types of configuration accesses to support hierarchical bridges. To access configuration space, a value is written to the CONFIG\_ADDR register specifying which PCI bus, which device, and which configuration register to be accessed.

When the PCI controller sees an access that falls inside the 4 bytes beginning at the CONFIG\_DATA address, it checks the enable bit, the device number and the bus number in the CONFIG\_ADDR register. If the enable bit is set and the device number is not equal to all ones, a configuration cycle translation is performed. When the device number field is equal to all ones, it has a special meaning (see [Section 13.4.4.6, “Special Cycle Command,”](#) for more information).

There are two types of translations supported:

- Type 0 translations—For when the device is on the PCI bus connected to the PCI controller.
- Type 1 translations—For when the device is on another bus somewhere behind the PCI controller.

For type 0 translations, the PCI controller decodes the device number field to assert the appropriate IDSEL line and perform a configuration cycle on the PCI bus with AD[1:0] as 0b00. All 21 IDSEL bits are decoded, starting with bit AD11. That is, if the device number field contains 0b01011, AD11 on the PCI bus is set. The IDSEL lines are bit-wise associated with increasing values for the device number such that AD12 corresponds to 0b01100, and so on up to bit 30 as shown in [Table 13-41](#). AD31 is selected with 0b01010. A device number of 0b11111 indicates a special cycle. Device number 0b00000 is used for configuring the PCI controller itself. Bits 10 through 8 are copied to the PCI bus as an encoded value for components which contain multiple functions. Bits 7 through 2 are also copied onto the PCI bus. The PCI controller implements address stepping on configuration cycles so that the target's PCI\_IDSEL, which is connected directly to one of the AD lines, reaches a stable value. This means that a valid address and command are driven on the AD and PCI\_C/ $\overline{\text{BE}}$  lines one cycle before the assertion of  $\overline{\text{PCI\_FRAME}}$ .

For type 1 translations, the PCI controller copies the contents of the CONFIG\_ADDR register directly onto the PCI address/data lines during the address phase of a configuration cycle, with the exception that AD[1-0] contains 0b01 (not 0b00 as in Type 0 translations).

When the PCI controller is configured as a host device, a local master sometimes needs to perform configuration reads from unpopulated PCI slots (as part of the system configuration). To avoid getting a machine check interrupt, the following steps should be taken:

1. Mask the NORSP bit in the error mask register. See [Section 13.3.2.9, “PCI Error Control Register \(PCI\\_ECR\).”](#)
2. Perform the PCI configuration reads.
3. Clear the NORSP bit in the error status register.
4. Unmask (write 1) the NORSP bit in the error mask register. See [Section 13.3.2.3, “PCI Error Enable Register \(PCI\\_EER\).”](#)

#### 13.4.4.5 Agent Mode Configuration Access

When the PCI controller is configured as an agent device, it responds to remote host generated PCI configuration accesses to the PCI interface. This is indicated by decoding the configuration command along with the PCI controller's IDSEL being asserted. A remote host can access the 256-byte PCI configuration area and the memory-mapped configuration registers within the PCI controller.

#### 13.4.4.6 Special Cycle Command

A special cycle command contains no explicit destination address but is broadcast to all PCI agents. Each receiving agent must determine whether the message is applicable to itself. No assertion of  $\overline{\text{PCI\_DEVSEL}}$  in response to a special cycle command is necessary.

A special cycle command is like any other bus command in that it has an address phase and a data phase. The address phase starts like all other commands with the assertion of  $\overline{\text{PCI\_FRAME}}$  and completes when

$\overline{\text{PCI\_FRAME}}$  and  $\overline{\text{PCI\_IRDY}}$  are negated. Special cycles terminate with a master-abort. (In the special cycle case, the received-master-abort bit in the configuration status register is not set.)

The address phase contains no valid information other than the command field. Even though there is no explicit address, the address/data lines are driven to a stable state and parity is generated. During the data phase, the address/data lines contain the message type and an optional data field. The message is encoded on the sixteen least-significant bits (AD[15:0]). The data field is encoded on AD[31:16]. When running a special cycle, the message and data are valid on the first clock  $\overline{\text{PCI\_IRDY}}$  is asserted.

When the `PCI_CONFIG_ADDRESS` register is written with a value so that the bus number matches the bridge bus, the device number is all ones, the function number is all ones, and the register number is zero. The next time the `PCI_CONFIG_DATA` register is accessed, the PCI controller executes either a special cycle or an interrupt acknowledge command. When the `PCI_CONFIG_DATA` register is written, the PCI controller generates a special cycle encoding on the command/byte enable lines during the address phase and drives the data from the `PCI_CONFIG_DATA` register onto the address/data lines during the first data phase.

If the bus number field of the `PCI_CONFIG_ADDRESS` does not match one of the PCI controller bus numbers, the PCI controller passes the write to `PCI_CONFIG_DATA` through to the PCI bus as a type 1 configuration cycle as it does any other time the bus number field does not match.

**Table 13-46. Special Cycle Commands**

Address (AD[15-0])	Message Type	Description
0x0000	SHUTDOWN (SLEEP)	Indicates the processor is entering its most power saving mode
0x0001	HALT (DOZE)	Indicates the processor is entering a power save mode where address decoding is still available
0x0002–0xFFFF	—	Reserved for future commands

### 13.4.4.7 Interrupt Acknowledge

When the `PCI_CONFIG_ADDRESS` register is written with a value such that the bus number is 0x00, the device number is all ones, the function number is all ones, and the register number is zero, the next time the `PCI_CONFIG_DATA` register is accessed the PCI controller does either a special cycle command or an interrupt acknowledge command. When the `PCI_CONFIG_DATA` register is read, the PCI controller generates an interrupt acknowledge command encoding on the command/byte enable lines during the address phase. During the address phase, AD[31:0] do not contain a valid address but are driven with stable data and valid parity ( $\overline{\text{PCI\_PAR}}$ ). During the data phase, the byte enable signals determine which bytes are involved in the transaction. The interrupt vector must be returned when  $\overline{\text{PCI\_TRDY}}$  is asserted.

An interrupt acknowledge transaction can also be issued on the PCI bus by reading from the `PCI_INT_ACK` register.

## 13.4.5 Error Functions

This section describes PCI bus errors.

### 13.4.5.1 Parity

During valid 32-bit address and data transfers, parity covers all 32 address/data lines and the 4 command/byte enable lines regardless of whether or not all lines carry meaningful information. Byte lanes not actually transferring data are driven with stable (albeit meaningless) data and are included in the parity calculation. During configuration, special cycle or interrupt acknowledge commands, some address lines are not defined but are still driven to stable values and included in the parity calculation.

Even parity is calculated for all PCI operations: the value of PCI\_PAR is generated such that the number of ones on PCI\_AD[31:0], PCI\_C/ $\overline{\text{BE}}$ [3:0] and PCI\_PAR equals an even number. The PCI\_PAR signal is driven when the address/data lines are driven and follow the corresponding address or data by one clock.

The PCI controller checks the parity after all valid address phases (the assertion of  $\overline{\text{PCI\_FRAME}}$ ) and for valid data transfers ( $\overline{\text{PCI\_IRDY}}$  and  $\overline{\text{PCI\_TRDY}}$  asserted) involving the PCI controller. When an address or data parity error is detected, the detected-parity-error bit in the configuration space status register is set (see [Section 13.3.3.4, “PCI Status Configuration Register.”](#))

### 13.4.5.2 Error Reporting

Except for setting the detected-parity-error bit, all parity error reporting and response is controlled by the parity-error-response bit (see [Section 13.3.3.3, “PCI Command Configuration Register,”](#) for more information). If the parity-error-response bit is cleared, the PCI controller completes all transactions regardless of parity errors (address or data). If the bit is set, the PCI controller asserts  $\overline{\text{PCI\_PERR}}$  two clocks after the actual data transfer in which a data parity error is detected, and keeps  $\overline{\text{PCI\_PERR}}$  asserted for one clock. When acting as an initiator during a read transaction or as a target involved in a write to system memory the PCI controller asserts  $\overline{\text{PCI\_PERR}}$ .

Figure 13-53 shows the possible assertion points for  $\overline{\text{PCI\_PERR}}$  if the PCI controller detects a data parity error.




Figure 13-53. PCI Parity Operation

As an initiator, the PCI controller attempts to complete the transaction on the PCI bus if a data parity error is detected and sets the data-parity-reported bit in the configuration space status register. If a data parity error occurs on a read transaction, the PCI controller aborts the transaction internally. As a target, the PCI controller completes the transaction on the PCI bus even if a data parity error occurs. If parity error occurs during a write to system memory, the transaction completes on the PCI bus, but is aborted internally, insuring that potentially corrupt data does not go to memory.

When the PCI controller asserts  $\overline{\text{PCI\_SERR}}$ , it sets the signaled-system-error bit in the configuration space status register. Additionally, if the error is an address parity error, the parity-error-detected bit is set; reporting an address parity error on  $\overline{\text{PCI\_SERR}}$  is conditioned on the parity-error-response bit being enabled in the command register.  $\overline{\text{PCI\_SERR}}$  is asserted when the PCI controller detects an address parity error while acting as a target. The system error is passed to the PCI controller's interrupt processing logic to assert  $\overline{\text{MCP}}$ . Figure 13-53 shows where the PCI controller could detect an address parity error and assert  $\overline{\text{PCI\_SERR}}$  or where the PCI controller, acting as an initiator, checks for the assertion of  $\overline{\text{PCI\_SERR}}$  signaled by the target detecting an address parity error.

As a target that asserts  $\overline{\text{PCI\_SERR}}$  on an address parity, the PCI controller completes the transaction on the PCI bus, aborting internally if the transaction is a write to system memory. If  $\overline{\text{PCI\_PERR}}$  is asserted during a PCI controller write to PCI, the PCI controller attempts to continue the transfer, allowing the target to abort/disconnect if desired. If the PCI controller detects a parity error on a read from PCI, the PCI controller aborts the transaction internally and continues the transfer on the PCI bus, allowing the target to abort/disconnect if desired.



In all cases of parity errors on the PCI bus, regardless of the parity-error-response bit, information about the transaction is logged in the PCI error control capture register, the PCI error address capture register and the PCI error data capture register;  $\overline{MCP}$  is also asserted to the core as an option.

### 13.4.6 PCI Inbound Address Translation

For inbound transactions (transactions generated by an external master on the PCI bus where the PCI controller responds as a slave device), the PCI controller only responds to PCI addresses within the windows mapped by the PCI inbound base address registers (PIBARs). If there is an address hit in one of the PIBARs, the PCI address is translated from PCI space to local memory space through the associated PCI inbound translation address registers (PITARs). This allows an external master to access local memory. Each PIBAR register is associated with a PITAR and PIWAR which are located in the PCI controller's PCI CSR space. Figure 13-54 shows an example translation window for inbound memory accesses.

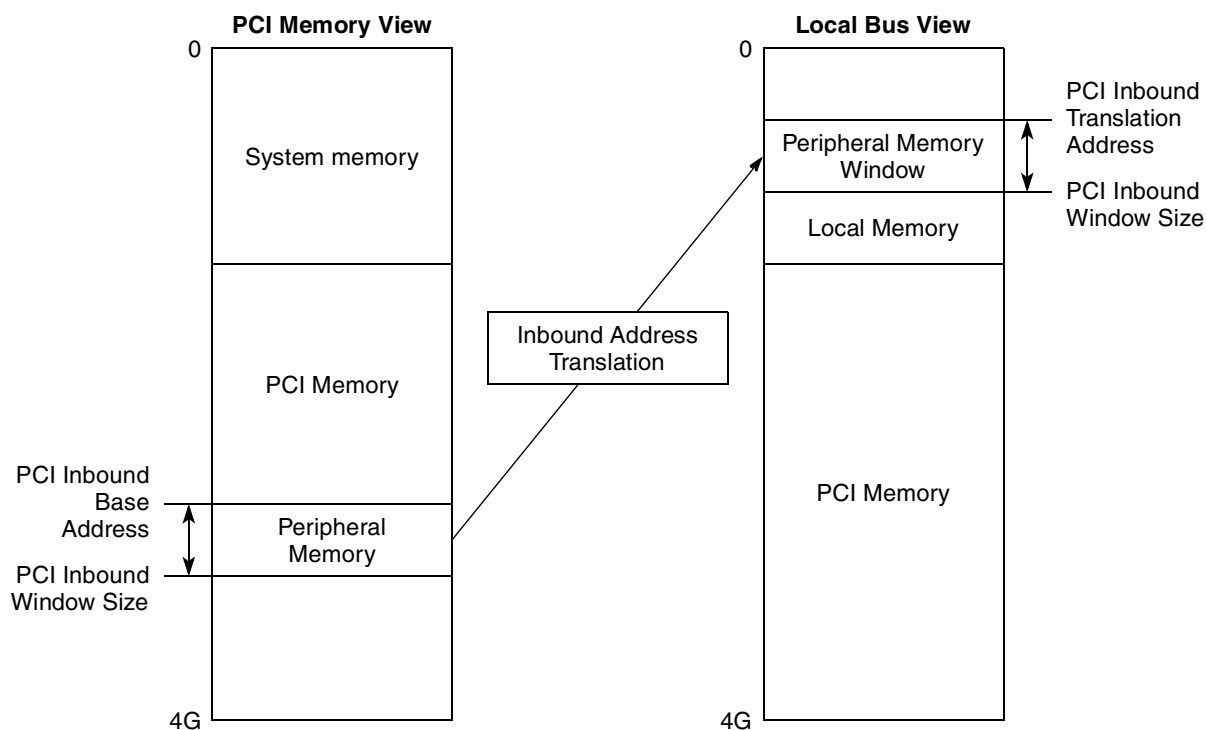


Figure 13-54. Inbound PCI Memory Address Translation

There are three full sets of inbound translation registers, in addition to the PIMMR base address register, allowing four simultaneous translation windows, one to a fixed destination and three programmable. Only two of the programmable windows can be mapped anywhere in the 64-bit PCI address space. Window 0 can only be mapped within the lowest 4-Gbyte space. Software can move the programmable translation base addresses during run-time to access different portions of local memory, but the PCI inbound translation windows may not overlap.

The translation windows are disabled after reset, that is, after reset, the PCI controller does not acknowledge externally mastered transactions on the PCI bus by asserting  $\overline{PCI\_DEVSEL}$  until the inbound translation windows are enabled.

## 13.4.7 CompactPCI Hot Swap Specification Support

CompactPCI is an open specification supported by the PCI Industrial Computer Manufacturers Group (PICMG) and is intended for embedded applications using PCI. CompactPCI Hot Swap is an extension of the CompactPCI specification and allows the insertion and extraction (or “hot swapping”) of boards without adversely affecting system operation. The hot swap specification defines the following levels of support:

- Hot swap capable
- Hot swap friendly
- Hot swap ready

The PCI controller is hot swap friendly, meaning that it supports the hardware and software connection processes as defined in the hot swap specification. This level of support allows the board and system designers to build full Hot Swap and high availability systems based on the PCI controller as a PCI target device. For details on the hot swap process, refer to the *Hot Swap Specification PICMG 2.1*, R1.0, August 3, 1998.

## 13.5 Initialization/Application Information

The following sections describe initialization sequences for host and agent modes.

### 13.5.1 Initialization Sequence for Host Mode

The following sequences must be followed in host mode:

1. Enable PCI output clocks and select desired frequency ratios. See [Section 4.4.1, “Clocking in PCI Host Mode.”](#)
2. Wait for at least 1 ms to enable stable clocks into agent devices
3. Deactivate PCI\_RESET\_OUT signal for PCI. See [Table 13-3](#) for more information on PCI\_RESET\_OUT signal.
4. Wait for at least 1 ms to enable devices to complete the powerup sequence.
5. Configure PCI internal registers and PCI agents to desired modes of operation

### 13.5.2 Initialization Sequence for Agent Mode

The following sequences must be followed in agent mode:

1. Optionally initialize subsystem vendor ID/device ID
  - a) Initialize PCI inbound window size in PIWAR[1:3] desired window size
  - b) Unlock configuration lock in PCI function configuration register

# Chapter 14

## Security Engine (SEC) 2.2

This chapter describes the functionality of the integrated security engine (SEC 2.2) in the MPC8313E. It addresses the following topics:

- Section 14.1, “SEC 2.2 Architecture Overview”
- Section 14.2, “Configuration of Internal Memory Space”
- Section 14.3, “Descriptor Overview”
- Section 14.4, “Execution Units”
- Section 14.5, “Channel”
- Section 14.6, “Controller”

### NOTE

The MPC8313 does not support a security engine.

The SEC 2.2 is designed to off-load computationally intensive security functions, such as authentication, and bulk encryption from the processor core of the MPC8313E. It is optimized to process all the algorithms associated with IPSec, SSL/TLS, iSCSI, SRTP, and 802.11i. The SEC 2.2 is derived from integrated security cores found in other members of the PowerQUICC family, including SEC 1.0, the version implemented in the MPC8272 and SEC 2.0, implemented on the MPC8555.

The security engine’s execution units (EUs) and primary features include the following:

- DEU—Data encryption standard execution unit
  - DES, 3DES
  - Two key (K1, K2, K1) or three key (K1, K2, K3)
  - ECB and CBC modes for both DES and 3DES
- AESU—Advanced encryption standard execution unit
  - Implements the Rijndael symmetric key cipher
  - ECB, CBC, CCM, and counter modes
  - 128-, 192-, 256-bit key lengths
- MDEU—Message digest execution unit
  - SHA with 160-, 224-, or 256-bit message digest
  - MD5 with 128-bit message digest
  - HMAC with either algorithm
- One channel, supporting a queue of commands (descriptor pointers)
  - Dynamic assignment of execution units through an integrated controller

- 256-byte buffer FIFOs on data input and output paths of each execution unit, with flow control for large data sizes. The input and output FIFOs are shared between AESU and DEU; MDEU has its own input FIFO.
- Master/slave logic, with DMA
  - 32-bit address/64-bit data
  - DMA blocks can be on any byte boundary
- Scatter/Gather capability
  - Gather capability enables the SEC 2.2 to concatenate multiple segments of memory when reading input data
  - Similarly, scatter capability enables the SEC 2.2 to write to multiple segments of memory when writing output data

## 14.1 SEC 2.2 Architecture Overview

The SEC 2.2 (referred to as SEC in this chapter) can act as a master on the internal system bus to allow the SEC to off-load the data movement bottleneck normally associated with slave-only cores. The host processor accesses the SEC through its device drivers using system memory for data storage. The SEC resides in the peripheral memory map of the processor, therefore when an application requires cryptographic functions, it simply creates descriptors for the SEC which define the cryptographic function to be performed and the location of the data. The SEC's bus-mastering capability permits the host processor to set up the channel with a few short register writes, leaving the SEC to perform reads and writes on system memory to complete the required task.

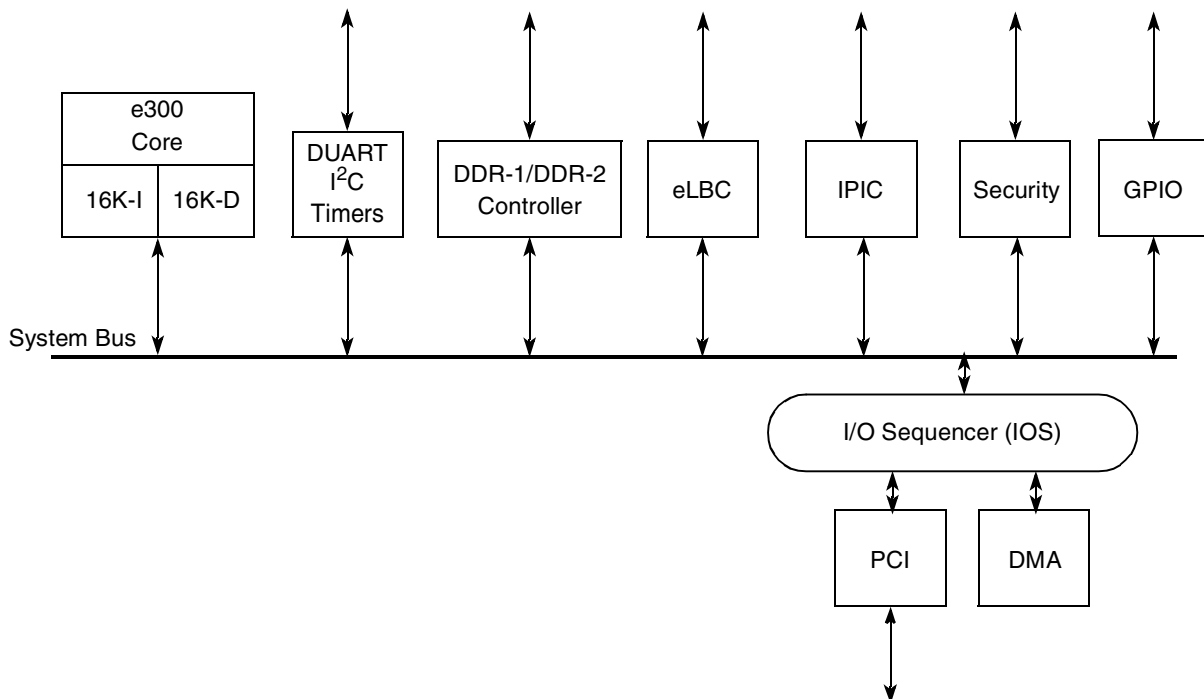
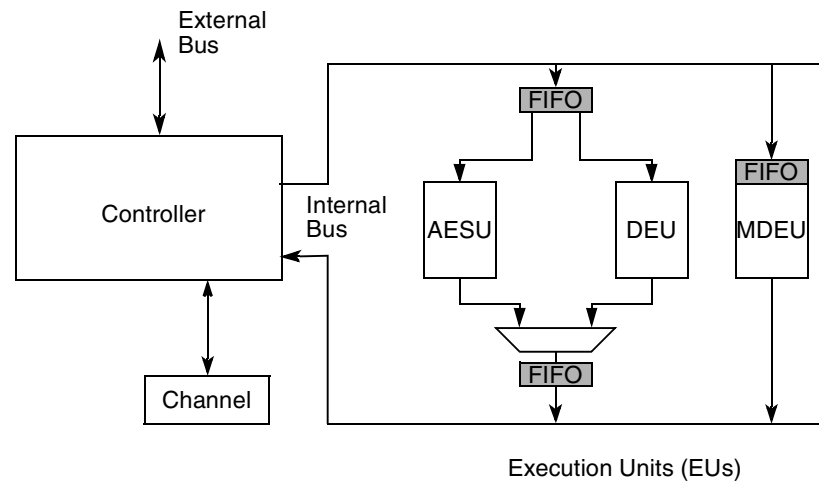


Figure 14-1. SEC Connected to MPC8313E System Bus

Figure 14-2 shows a simplified block diagram of the SEC internal architecture. The controller block is capable of transferring 64-bit words between the bus and any register inside the SEC.



**Figure 14-2. SEC Functional Modules**

An operation begins when the host writes a descriptor pointer to the fetch FIFO in the SEC channel. From this point on, the channel directs the sequence of operations. The channel uses the descriptor pointer to read the descriptor, then decodes the first word of the descriptor to determine the operation to be performed and the execution units needed to perform it. The channel requests the controller to assign the needed execution units. Next the channel requests that the controller fetch the keys, IVs and data from locations specified in the rest of the descriptor. The controller satisfies the requests by making requests to the master interface per the programmable priority scheme. Data is fed into the execution units through their registers and the proper input FIFOs. The execution units read from their input FIFOs and write processed data to their output FIFOs. The channel requests the controller to write data from the output FIFOs and registers back to system memory through the master/slave interface.

For most packets, the entire payload is too long to fit in the input or output FIFO. The SEC then uses a flow control scheme for reading and writing data. The channel directs the controller to read bursts of input as necessary to keep refilling the input FIFO, until the entire payload has been fetched. Similarly, the channel directs the controller to write bursts of output whenever enough accumulates in the output FIFO.

### 14.1.1 Descriptors

As a crypto acceleration block, the SEC controller has been designed for easy use and integration with existing systems and software. All cryptographic functions are accessible through descriptors. A descriptor specifies a cryptographic function to be performed, and contains pointers to all necessary input data and to the places where output data is to be written. Some descriptor types perform multiple functions to facilitate particular protocols. A descriptor is diagrammed in Table 14-1.

**Table 14-1. Example Descriptor**

Field Name	Value/Type	Description
DPD_DES_CTX_CRYPT	Variable	Representative header for DES using context to encrypt
LEN_CTXIN PTR_CTXIN	Length Pointer	Number of bytes to be written Pointer to context (IV) to be written into DES engine
LEN_KEY PTR_KEY	Length Pointer	Number of bytes in key Pointer to block cipher key
LEN_DATAIN PTR_DATAIN	Length Pointer	Number of bytes of data to be ciphered Pointer to data to perform cipher upon
LEN_DATAOUT PTR_DATAOUT	Length Pointer	Number of bytes of data after ciphering Pointer to location where cipher output is to be written
LEN_CTXOUT PTR_CTXOUT	Length Pointer	Length of output context (IV) Pointer to location where altered context is to be written
Null length Null pointer	Length Pointer	Zeros for fixed length descriptor Zeros for fixed length descriptor
Null length Null pointer	Length Pointer	Zeros for fixed length descriptor Zeros for fixed length descriptor

Each descriptor contains eight long-words (64 bits each), consisting of the following:

- One long-word of header—The header describes the required services and encodes information that indicates which EUs to use and which modes to set. It also indicates whether notification should be sent to the host when the descriptor operation is complete.
- Seven long-words containing pointers and lengths used to locate input or output data.

Upon completion of the current descriptor, the channel checks the next entry in its fetch FIFO, and, if non-zero, the channel is instructed to request a burst read of the next descriptor.

Processing of the next descriptor (and whether or not a done signal is generated) is determined by the programming of channel's configuration register. Two modes of operation are supported:

- Signal done at end of every descriptor
- Signal done at end of a selected descriptor

The channel can signal done through an interrupt, or by a write-back of to descriptor header in non-SEC memory after processing a descriptor. Either the value written back is identical to that of the header, with the exception that a DONE field is set, or special status fields are written back. That status writeback field can be reserved for descriptors performing ICV checking.

Occasionally, a descriptor field may not be applicable to the requested service. For example, if using DES in ECB mode, the contents of the IV field do not affect the result of the DES computation. Therefore, when processing descriptors, the channel skips any pointer that has an associated length of zero.

For more information, refer to [Section 14.3, “Descriptor Overview.”](#)

## 14.1.2 Execution Units (EUs)

'Execution unit' (EU) is the generic term for a functional block that performs the mathematical manipulations required by protocols used in cryptographic processing. The EUs are compatible with IPsec, SSL/TLS, iSCSI, SRTP, and 802.11i processing, and can work together to perform high-level cryptographic tasks. The SEC's execution units are as follows:

- DEU for performing block cipher, symmetric key cryptography using DES and 3DES
- AESU for performing the advanced encryption standard algorithm
- MDEU for performing security hashing using MD-5, SHA-1, SHA-224, or SHA-256

Each EU is described in detail in [Section 14.4, "Execution Units."](#)

### 14.1.2.1 Data Encryption Standard Execution Unit (DEU)

The DES execution unit (DEU) performs bulk data encryption/decryption, in compliance with the data encryption standard algorithm (ANSI x3.92). The DEU can also compute 3DES, an extension of the DES algorithm in which each 64-bit input block is processed three times. The SEC supports 2-key (K1 = K3) or 3-key 3DES.

The DEU operates by permuting 64-bit data blocks with a shared 56-bit key and an initialization vector (IV). The SEC supports two modes of operation: electronic code book (ECB) and cipher clock chaining (CBC).

For more information, refer to [Section 14.4.1, "Data Encryption Standard Execution Unit \(DEU\)."](#)

### 14.1.2.2 Advanced Encryption Standard Execution Unit (AESU)

The AESU is used to accelerate bulk data encryption/decryption in compliance with the advanced encryption standard algorithm Rijndael. The AESU executes on 128 bit blocks with a choice of key sizes: 128, 192, or 256 bits.

AESU is a symmetric-key algorithm; the sender and receiver use the same key for both encryption and decryption. The session key and IV are supplied to the AESU module prior to encryption. The processor supplies data to the module that is processed as 128-bit input. The AESU operates in ECB, CBC, CTR, and CCM modes.

For more information, refer to [Section 14.4.3, "Advanced Encryption Standard Execution Unit \(AESU\)."](#)

### 14.1.2.3 Message Digest Execution Unit (MDEU)

The MDEU computes a single message digest (or hash or integrity check) value of all the data presented on the input bus, using either the MD5, SHA-1, SHA-224, or SHA-256 algorithms for bulk data hashing. With any hash algorithm, the larger message is mapped onto a smaller output space, therefore collisions are potential, albeit not probable. The 160-bit hash value is a sufficiently large space such that collisions are extremely rare. The security of the hash function is based on the difficulty of locating collisions. That is, it is computation infeasible to construct two distinct but similar messages that produce the same hash output.

- The MD5 generates a 128-bit hash, and the algorithm is specified in RFC 1321.

- SHA-1 is a 160-bit hash function, specified by the ANSI X9.30-2 and FIPS 180-2 standards.
- SHA-224 is a 224-bit hash function specified by FIPS 180-2 that provides 224 bits of security against collision attacks.
- SHA-256 is a 256-bit hash function specified by FIPS 180-2 that provides 256 bits of security against collision attacks.
- The MDEU also supports HMAC computations, as specified in RFC 2104.

For more information, refer to [Section 14.4.2, “Message Digest Execution Unit \(MDEU\).”](#)

### 14.1.3 Channel

The SEC includes one channel that manages data and EU function by using the following:

- A fetch FIFO, which holds a queue of pointers to descriptors waiting to be serviced
- A configuration register, which allows the user a number of options for SEC event signaling.
- Control registers containing information about the transaction in process
- A status register containing an indication of the last unfulfilled bus request
- A descriptor buffer memory used to store the active descriptor

Whenever the channel is idle and its fetch FIFO is non-empty, the channel reads the next descriptor pointer from the fetch FIFO. Using this pointer, the channel fetches the descriptor and places it in its descriptor buffer. To service this descriptor, the channel directs execution of the following steps.

1. Analyze the descriptor header to determine the cryptographic services required, and request use of the appropriate EUs from the controller.
2. Wait for the controller to grant access to the required EUs.
3. Set the appropriate mode bits in the EU(s) for the required service.
4. Fetch ‘data parcels’ using pointers from the descriptor buffer, and place them in either a EU input FIFO or EU registers (as appropriate). The term ‘data parcel’ refers here to any input or output of a cryptographic process, such as a key, hash result, input context, output context, or text-data. ‘Context’ refers to either an IV (initialization vector) or other internal EU state that can be read out or loaded in. ‘Text-data’ refers to plaintext or ciphertext to be operated on.
5. If the data size is greater than EU FIFO size, continue fetching input data, and writing output data to memory.
6. Wait for EU(s) to complete processing.
7. Upon completion, unload results from output FIFOs and context registers and write them to external memory using pointers in the descriptor buffer.
8. If multiple services are requested, go back to step 3.
9. Release the EUs.
10. If ‘done notification’ is enabled in the descriptor header, perform this notification.

The channel can generate two types of done notification signals when it completes operation on a descriptor. It can signal done through an interrupt or by a writeback of the descriptor header after processing a descriptor. Two values can be written back: the first is identical to that of the header, with the



exception that a 'done' field is set. The second is a status field writeback. That status field writeback occurs to report the result of ICV checking if ICV check writeback is enabled.

Many security protocols involve both encryption and hashing of packet payloads. To accomplish this without requiring two passes through the data, the channel can configure data flows through more than one EU. In such cases, one EU is designated the primary EU, and the other as the secondary EU. The primary EU receives its data from memory through the controller, and the secondary EU receives its data by 'snooping' SEC buses.

There are two types of snooping:

- Input data can be fed to the primary EU and the same input data snooped by the secondary EU. This is called 'in-snooping.'
- Output data from the primary EU can be snooped by the secondary EU. This is called 'out-snooping.'

In the SEC the secondary EU is always the MDEU.

For more information, refer to [Section 14.5, "Channel."](#)

### 14.1.4 SEC Controller

The SEC controller manages on-chip resources, including the individual execution units (EUs), FIFOs, the master/slave interface to the MPC8313E system bus, and the internal buses that connect all the various modules. The controller receives service requests from the master/slave interface and from the channel, and schedules the required activities. The controller provides for two ways of operating the execution units:

- Channel-controlled access—The channel can request a particular service from any available execution unit. This is the normal operating condition.
- Host-controlled access—The host can move data into and out of any execution unit directly through memory-mapped EU registers. This is typically only used for debug.

The system bus interface and access to system memory are critical factors in performance, and the 64-bit master/slave interface of the SEC controller allows it to achieve performance unattainable on secondary buses.

#### 14.1.4.1 Channel-Controlled Access

Processing begins when a descriptor pointer is written to the fetch FIFO of the channel. Based on the services requested by the descriptor header, the channel asks the controller to assign the necessary EUs to the channel. Once the required EU has been reserved, the channel requests that the controller fetch and load the appropriate data. The controller acts as a master on the system bus, reading and writing on byte boundaries. The channel operates the EU, and makes further requests to the controller to write output data to system memory. When the descriptor processing is complete, the channel asks the controller to release the EU.

### 14.1.4.2 Host-Controlled Access

All execution units (EUs) are memory-mapped, and can be used entirely through register read/write access. The SEC operates as a slave, and the host must write the information typically provided through the descriptor into the appropriate registers and FIFOs of the SEC. This method is more CPU intensive, and requires a great deal of familiarity with SEC registers. It is recommended that host-controlled access be used only for operations using a single EU, and for debug purposes.

For more information, refer to [Section 14.6, “Controller.”](#)

## 14.2 Configuration of Internal Memory Space

[Table 14-2](#) shows the base address map, while [Table 14-3](#) provides the address map, including all registers in the execution units. The 18-bit SEC address bus value is shown. These address values are offsets from IMMRBAR. See [Section 5.2.4.1, “Internal Memory Map Registers Base Address Register \(IMMRBAR\),”](#) for more information.

Note that these tables show addresses of 64-bit words; the three least-significant address bits that are used to select bytes within 64-bit words are not shown.

**Table 14-2. SEC Address Map**

Address Offset (AD 17–0)	Module	Description	Type	Reference
0x3_0000–0x3_0FFF	—	Reserved	—	—
0x3_1000–0x3_10FF	Controller	Arbiter/controller control register space	Resource control	<a href="#">Section 14.6, “Controller”</a>
0x3_1100–0x3_11FF	Channel	Channel	Data control	<a href="#">Section 14.5, “Channel”</a>
0x3_2000–0x3_2FFF	DEU	DES/3DES execution unit	EU	<a href="#">Section 14.4.1, “Data Encryption Standard Execution Unit (DEU)”</a>
0x3_4000–0x3_4FFF	AESU	AES execution unit		<a href="#">Section 14.4.3, “Advanced Encryption Standard Execution Unit (AESU)”</a>
0x3_6000–0x3_6FFF	MDEU	Message digest execution unit		<a href="#">Section 14.4.2, “Message Digest Execution Unit (MDEU)”</a>

[Table 14-3](#) shows the system address map showing all functional registers.

**Table 14-3. SEC Address Map**

Address Offset (AD 17–0)	Register	Access	Access By	Section/Page
<b>Controller</b>				
0x3_1008	IMR—Interrupt mask register	R/W		<a href="#">14.6.4.2/14-68</a>
0x3_1010	ISR—Interrupt status register	R		<a href="#">14.6.4.3/14-70</a>
0x3_1018	ICR—Interrupt clear register	W		<a href="#">14.6.4.4/14-71</a>
0x3_1020	ID—Identification register	R		<a href="#">14.6.4.5/14-73</a>

Table 14-3. SEC Address Map (continued)

Address Offset (AD 17-0)	Register	Access	Access By	Section/Page
0x3_1028	EUASR—EU assignment status register	R		14.6.4.1/14-67
0x3_1030	MCR—Master control register	R/W		14.6.4.7/14-74
<b>Channel</b>				
0x3_1108	CCCR—Crypto-channel configuration register	R/W		14.5.1.1/14-55
0x3_1110	CCPSR—Crypto-channel pointer status register	R		14.5.1.2/14-57
0x3_1140	CDPR—Crypto-channel current descriptor pointer register	R		14.5.1.3/14-62
0x3_1148	FF—Crypto-channel fetch FIFO register	W		14.5.1.4/14-62
0x3_1180–0x3_11BF	DB $n$ —Descriptor buffers [0–7]	R		14.5.1.5/14-63
<b>Controller</b>				
0x3_1BF8	IP block revision register	R	Word	14.6.4.6/14-73
<b>DEU</b>				
0x3_2000	DEUMR—DEU mode register	R/W		14.4.1.1/14-19
0x3_2008	DEUKSR—DEU key size register (bytes)	R/W		14.4.1.2/14-20
0x3_2010	DEUDSR—DEU data size register (bits)	R/W		14.4.1.3/14-21
0x3_2018	DEURCR—DEU reset control register	R/W	Word	14.4.1.4/14-22
0x3_2028	DEUSR—DEU status register	R	Word	14.4.1.5/14-23
0x3_2030	DEUI SR—DEU interrupt status register	R	Word	14.4.1.6/14-24
0x3_2038	DEUI CR—DEU interrupt control register	R/W	Word	14.4.1.7/14-25
0x3_2050	DEUEMR—DEU end-of-message register	W	Word	14.4.1.8/14-27
0x3_2100	DEUIV—DEU IV register	R/W		14.4.1.9/14-27
0x3_2400	DEUK1—DEU key 1 register	W		14.4.1.10/14-28
0x3_2408	DEUK2—DEU key 2 register	W		14.4.1.10/14-28
0x3_2410	DEUK3—DEU key 3 register	W		14.4.1.10/14-28
0x3_2800–0x3_2FFF	DEU FIFO	R/W		14.4.1.11/14-28
<b>AESU</b>				
0x3_4000	AESUMR—AESU mode register	R/W		14.4.3.1/14-40
0x3_4008	AESUKSR—AESU key size register (bytes)	R/W		14.4.3.2/14-42
0x3_4010	AESUDSR—AESU data size register (bits)	R/W		14.4.3.3/14-43
0x3_4018	AESURCR—AESU reset control register	R/W	Word	14.4.3.4/14-43
0x3_4028	AESUSR—AESU status register	R	Word	14.4.3.5/14-44
0x3_4030	AESUI SR—AESU interrupt status register	R	Word	14.4.3.6/14-45
0x3_4038	AESUI CR—AESU interrupt control register	R/W	Word	14.4.3.7/14-47

Table 14-3. SEC Address Map (continued)

Address Offset (AD 17-0)	Register	Access	Access By	Section/Page
0x3_4050	AESUEMR—AESU end-of-message register	W	Word	14.4.3.8/14-48
0x3_4100–0x3_4108	AESU context memory registers	R/W		14.4.3.9/14-49
0x3_4400–0x3_4408	AESU key memory registers	R/W		14.4.3.9.5/14-53
0x3_4800–0x3_4FFF	AESU FIFO	R/W		14.4.3.9.6/14-53
<b>MDEU</b>				
0x3_6000	MDEUMR—MDEU mode register	R/W		14.4.2.1/14-28
0x3_6008	MDEUKSR—MDEU key size register (bytes)	R/W		14.4.2.3/14-32
0x3_6010	MDEUDSR—MDEU data size register (bits)	R/W		14.4.2.4/14-32
0x3_6018	MDEURCR—MDEU reset control register	R/W	Word	14.4.2.5/14-33
0x3_6028	MDEUSR—MDEU status register	R	Word	14.4.2.6/14-34
0x3_6030	MDEUIR—MDEU interrupt status register	R	Word	14.4.2.7/14-35
0x3_6038	MDEUICR—MDEU interrupt control register	R/W	Word	14.4.2.8/14-36
0x3_6040	MDEU ICV size register	W		14.4.2.9/14-37
0x3_6050	MDEUEMR—MDEU end-of-message register	W	Word	14.4.2.10/14-38
0x3_6100–0x3_6120	MDEU context memory registers	R/W		14.4.2.11/14-38
0x3_6400–0x3_647F	MDEU key memory registers	W		14.4.2.12/14-39
0x3_6800–0x3_6FFF	MDEU FIFO	W		14.4.2.13/14-40

## 14.3 Descriptor Overview

The host processor maintains a record of current secure sessions and the corresponding keys and contexts of those sessions. Once the host has determined that a security operation is required, it creates a ‘descriptor’ containing all the information the SEC needs to perform the security operation. The host creates the descriptor in main memory, then writes a pointer to the descriptor into the fetch FIFO of the SEC channel. The channel uses this pointer to read the descriptor into its descriptor buffer. Once it obtains the descriptor, the SEC uses its bus mastering capability to obtain inputs and write results, thus off-loading data movement and encryption operations from the host processor.

For test purposes, it is also possible for the host to write keys, context, and text-data directly to execution units, using the SEC’s host-controlled access. This method avoids use of descriptors.

### 14.3.1 Descriptor Structure

SEC descriptors are conceptually similar to descriptors used by most devices with DMA capability. The descriptors have a fixed length of 64 bytes, that is, eight long-words, consisting of one ‘header dword’ and seven ‘pointer dwords.’ See [Figure 14-3](#) for the descriptor format.

	0	15	16	17	23	24	31	32	63
Header Dword	Header							Reserved	
Pointer Dword 0	Length0	J0	Extent0	—			Pointer0		
Pointer Dword 1	Length1	J1	Extent1	—			Pointer1		
Pointer Dword 2	Length2	J2	Extent2	—			Pointer2		
Pointer Dword 3	Length3	J3	Extent3	—			Pointer3		
Pointer Dword 4	Length4	J4	Extent4	—			Pointer4		
Pointer Dword 5	Length5	J5	Extent5	—			Pointer5		
Pointer Dword 6	Length6	J6	Extent6	—			Pointer6		

**Figure 14-3. Descriptor Format**

The header dword specifies the security operation to be performed, the execution unit(s) needed, and the modes for each execution unit. The pointer dwords, all of which have the same format, contain pointer and length information for locating input or output data parcels (such as keys, context, or text-data). The large number of pointers provided in the descriptor allows for multi-algorithm operations that require fetching of multiple keys, as well as fetch and return of contexts. Any pointer dword that is not needed can be given a length of zero, and the channel will skip over the corresponding operations.

SEC descriptors include scatter/gather capability, which means that each pointer in a descriptor can be either a direct pointer to a contiguous parcel of data, or can be a pointer to a link table, which is a list of pointers and lengths used to assemble the data parcel. When a link table is used to read input data, this is referred to as a ‘gather’ operation; when used to write output data, it is referred to as a ‘scatter’ operation.

### 14.3.2 Descriptor Format: Header Dword

Descriptors are created by the host to guide the SEC through required cryptographic operations. The header dword defines the operations to be performed, the mode for each operation, and internal addressing used by the controller and channel for internal data movement. The fields that must be supplied to the SEC are shown in the ‘Field’ rows of [Figure 14-4](#), and described in [Table 14-4](#). The SEC device drivers allow the host to create proper headers for each cryptographic operation.

In processing a descriptor, the SEC can also write back certain fields to the header dword. These are shown in the ‘Writeback’ rows of [Figure 14-4](#), and described in [Table 14-5](#).




Figure 14-4. Header Dword

Table 14-4. Header Dword Bit Definitions

Bits	Name	Description
0–3	OP_0: EU_SEL0	Primary EU select. See <a href="#">Section 14.3.2.1, “Selecting Execution Units—EU_SEL0 and EU_SEL1,”</a> for possible values.
4–11	MODE0	Primary mode. Mode data used to program the primary EU. The mode data is to the chosen EU. This field is passed directly to bits 56–63 of the mode register in the selected EU.
12–15	OP_1: EU_SEL1	Secondary EU select. See <a href="#">Section 14.3.2.1, “Selecting Execution Units—EU_SEL0 and EU_SEL1,”</a> for possible values.
16–23	MODE1	Secondary mode. Mode data used to program the primary EU. The mode data is to the chosen EU. This field is passed directly to bits 56–63 of the mode register in the selected EU.
24–28	DESC_TYPE	Descriptor type. This, along with the DIR field, determines the sequence of actions to be performed by the channel and selected EUs using the blocks of data listed in the rest of the descriptor. The attributes determined include the direction of data flow for each data block, which EU (primary or secondary) is accessed, what snooping options are used, and which internal EU addresses are accessed. See <a href="#">Section 14.3.2.2, “Selecting Descriptor Type—DESC_TYPE,”</a> for possible values.
29	—	Reserved.
30	DIR	Direction. Direction of overall data flow 0 Outbound 1 Inbound This, along with the DESC_TYPE field, helps determine the sequence of actions to be performed by the channel and selected EUs.
31	DN	Done notification. 0 No done notification. 1 Signal ‘done’ to the host on completion of this descriptor. This enables done notification if the NT field is 1 in the channel configuration register (see <a href="#">Table 14-31</a> ). The done notification can take the form of an interrupt, a writeback in the DONE field of this header dword (see <a href="#">Table 14-5</a> ), or both, depending upon the states of the CDIE (Channel Done Interrupt Enable) and CDWE (Channel Done Writeback Enable) bits in the channel configuration register.

**Table 14-5. Header Dword Writeback Bit Definitions**

Bits	Name	Description
0–7	DONE	When done writeback is used, then at the completion of descriptor processing this byte is written with the value 0xFF. To determine when done writeback is used, see the CDWE, NT, and CDIE fields in the channel configuration register (Table 14-31).
8–34	—	Reserved.
35–36	ICR0	Integrity check result from primary. These bits are supplied by the primary EU when descriptor processing is complete. 00 No integrity check was performed. 01 The integrity check passed. 10 The integrity check failed. 11 Reserved
37–42	—	Reserved.
43–44	ICR1	Integrity check result from secondary. These bits are supplied by the secondary EU (if any) when descriptor processing is complete. 00 No integrity check was performed. 01 The integrity check passed. 10 The integrity check failed. 11 Reserved
45–63	—	Reserved.

### 14.3.2.1 Selecting Execution Units—EU\_SEL0 and EU\_SEL1

Table 14-6 shows the values for EU\_SEL0 and EU\_SEL1 in the descriptor header. The following rules govern the choices for these fields:

1. EU\_SEL0 values of ‘No EU selected’ or ‘Reserved’ will result in an ‘Unrecognized Header Error’ condition during processing of the descriptor header.
2. The only valid choices for EU\_SEL1 are ‘No EU selected’ or MDEU. Any other choice will result in an ‘Unrecognized Header’ error condition.
3. If EU\_SEL1 is MDEU, then EU\_SEL0 must be DEU or AESU. All other values of EU\_SEL0 will result in an ‘Unrecognized header’ error condition.

**Table 14-6. EU\_SEL0 and EU\_SEL1 Values**

Value (Binary)	Selected EU
0000	No EU selected
0010	DEU
0011	MDEU
0110	AESU
Others	Reserved
1111	Reserved for header writeback

### 14.3.2.2 Selecting Descriptor Type—DESC\_TYPE

Table 14-7 shows the permissible values for the DESC\_TYPE field in the descriptor header. Descriptor types from SEC 1.0, which have ‘0’ in the last bit, are listed first, followed by SEC 2.x types, which have ‘1’ in the last bit.

**Table 14-7. Descriptor Types**

Value (Binary)	Descriptor Type	Notes
0000_0	aesu_ctr_nonsnoop	AESU CTR non-snooping <sup>1</sup>
0001_0	common_nonsnoop	Common, non-snooping
0010_0	hmac_snoop_no_afeu	Snooping, HMAC
0011_0	—	Reserved
0100_0	—	Reserved
0101_0	—	Reserved
0110_0	—	Reserved
0111_0	—	Reserved
1000_0	—	Reserved
1001_0	—	Reserved
1010_0	—	Reserved
1011_0	—	Reserved
1100_0	hmac_snoop_aesu_ctr	AESU CTR hmac snooping <sup>2</sup>
1101_0	—	Reserved
1110_0	—	Reserved
1111_0	—	Reserved
0000_1	ipsec_esp	IPsec ESP mode encryption and hashing
0001_1	802.11i AES ccmp	CCMP encryption and hashing, suitable for IEEE Std. 802.11i
0010_1	srtplib	SRTP encryption and hashing
0011_1	—	Reserved
0100_1	—	Reserved
0101_1	—	Reserved
0110_1	—	Reserved
0111_1	—	Reserved
1000_1	tls_ssl_block	TLS/SSL generic block cipher
1001_1	—	Reserved



Table 14-7. Descriptor Types (continued)

Value (Binary)	Descriptor Type	Notes
1010_1	raid_xor	XOR 3 sources together
Others	—	Reserved

<sup>1</sup> Type 0000\_0 is for AES-CTR operations. Type 0001\_0 also supports AES-CTR, however, to use AES-CTR with 0001\_0, the user must pre-pend zeros to the AES-Ctx before loading the AES context registers.

<sup>2</sup> Type 1100\_0 is for AES-CTR operations with HMAC. Type 0010\_0 also supports AES-CTR with HMAC, however, to use AES-CTR with 0010\_0, the user must pre-pend zeros to the AES-Ctx before loading the AES context registers.

For more about descriptor types and the data used for each type, see [Section 14.3.5, “Descriptor Types.”](#)

### 14.3.3 Descriptor Format: Pointer Dwords

The descriptor contains seven ‘pointer dwords,’ which define where in memory the SEC should access its input and output data parcels. The pointer dwords are numbered 0–6 as shown in [Figure 14-3](#). The channel determines how it will use each of the pointer dwords based on the ‘Descriptor Type’ and ‘Direction’ fields in the header. The channel accesses the first data parcel by starting at a location given by a POINTER value, and accessing a number of bytes given by a LENGTH or EXTENT value. Subsequent data parcels may be accessed by starting where a previous data parcel ended, or by starting at a different POINTER. The LENGTH or EXTENT used with any POINTER may be from the same pointer dword or from a different pointer dword in the same descriptor. Although the EXTENT field exists in each pointer dword of the SEC descriptor, only the EXTENTS in pointer dwords 3, 4, and 5 are currently in use.




Figure 14-5. Pointer Dword

Table 14-8. Pointer Dword Field Definitions

Bits	Name	Description
0–15	LENGTH	Length. A number of bytes in the range 0–65535. The use of this field depends on the ‘Descriptor Type’ and ‘Direction’ in the header dword. A value of zero causes the channel to skip this dword.
16	J	Jump. Determines whether to ‘jump’ to a link table whenever the POINTER field in this same lword is used. 0 The POINTER field points to data. 1 The POINTER field points to a link table, and scatter/gather is enabled.
17–23	EXTENT	Extent. A number of bytes in the range 0–127. The use of this field depends on the ‘Descriptor Type’ and ‘Direction’ in the header dword.
24–31	—	Reserved
32–63	POINTER	Pointer: A memory address.

On occasion, a descriptor field may not be applicable to the requested service. With seven pointer dwords, it is possible that not all these dwords will be required to specify the input and output parameters. (Some operations, for example, do not require context.) Where a particular dword is not used, all fields should be set to 0.

Some descriptors involve more than seven parcels of input and output data. In these cases, it is necessary to use one POINTER field to address a sequence of data parcels.

LENGTH and EXTENT fields normally specify the sizes of data parcels. In some cases, however, the POINTER field is zero, and the LENGTH and/or EXTENT fields simply specify values to be written to an EU.

The J bit in each pointer dword is used to enable the scatter/gather feature. If a data parcel to be read or written by the SEC is in one contiguous block of memory locations, then the scatter/gather feature is not needed. In this case the POINTER should be set to point directly at the first byte of the parcel, and the J bit should be 0. On the other hand, if the data parcel is stored in several separate segments of memory, then the scatter/gather capability is needed to assemble or distribute the complete parcel. In this case, the POINTER should be set to point to a link table, and the J bit should be 1. For link table format, see [Section 14.3.4, “Link Table Format.”](#)

### 14.3.4 Link Table Format

Link tables implement scatter/gather capability. For gather operations, a link table specifies a list of memory segments that are to be concatenated in the process of assembling data parcels. For scatter operations, a link table specifies a list of memory segments into which the output data should be written. Scatter or gather of a data parcel may be specified by a single link table or by a chain of link tables that are linked together with pointers.

The link table or chain of link tables accessed through some descriptor POINTER must specify enough memory segments to hold all the data that will be accessed through that pointer. In most cases, only a single data parcel is accessed through a given POINTER, and the chain of link tables specifies just that parcel. In other cases, the descriptor POINTER is used multiple times to access a sequence of data parcels, and the chain of link tables must supply data for the entire sequence. If a link table is used to access a sequence of data parcels, the end of each parcel must also be at the end of a memory segment. In other words, a single memory segment must not straddle two data parcels.

A link table may contain any number of long word entries. There are two kinds of entries—regular entries and next entries. Each ‘regular entry’ specifies a memory segment by means of a 32-bit starting address (SegAdr) and a 16-bit length (SegLen). A ‘next entry’ is used at the end of a link table to specify that the list of memory segments is continued in another link table. In a next entry, the N bit is set, the SegAdr field gives the address of the next link table, and the SEGLen field must be 0. A chain of link tables may contain any number of link tables.

Whether the list of memory segments is in a single link table or split into several link tables, the last entry in the last link table is a regular entry with the R (return) bit set. The R bit signifies the end of link table operations so that the channel returns to the descriptor for its next pointer (if any). A single link table entry is shown in [Figure 14-6](#).




Figure 14-6. Link Table Entry

Table 14-9. Link Table Field Definitions

Bits	Name	Description
0–15	SEGLen	Length. When N=0, a number in the range 1–65535, specifying the number of bytes in the memory segment. pointed to by SEGADR. A value of 0 will cause an error state to be set in the channel pointer status register—G-STATE for a gather operation or S-STATE for a scatter operation (see <a href="#">Section 14.5.1.2, “Crypto-Channel Pointer Status Register (CCPSR)”</a> ). When N = 1, must be 0.
16–21	—	Reserved
22	R	Return. When N=0: 0 No special action. 1 This is the last entry in the chain of link tables. If this entry does not specify the right number of bytes to complete the last data parcel, a G-STATE or S-STATE error will be set in the channel pointer status register (see <a href="#">Section 14.5.1.2, “Crypto-Channel Pointer Status Register (CCPSR)”</a> ). When N = 1, ignored.
23	N	Next 0 No special action. 1 This is the last long word in the current link table. The SEGADR field is the address of the next link table in the chain.
24–31	—	Reserved
32–63	SEGADR	Segment address. A memory address.

For any sequence of data parcels accessed by a link table or chain of link tables, the combined lengths of the parcels (the sum of their LENGTH and/or EXTENT fields) must equal the combined lengths of the link table memory segments (SEGLen fields). Otherwise the channel sets the appropriate error state in the channel pointer status register—G-STATE for gather error or S-STATE for scatter error (see [Section 14.5.1.2, “Crypto-Channel Pointer Status Register \(CCPSR\)”](#)).

Example (from [Figure 14-6](#)): To demonstrate use of a link table, assume that the current descriptor type calls for the channel to read a data parcel using Pointer3 and Extent3 fields, and assume that J3 = 1. Due to the J3 value, Pointer3 is not used as a data address but instead used as the address of a link table. The channel begins by reading the first four long words starting at Pointer3 into an internal ‘gather table buffer.’

Using the first entry of the gather table buffer, the channel starts accessing the data parcel by reading SEGLen bytes beginning at SEGADR. If the required data parcel size (Extent3) is greater than this first SegLen, the channel moves on to the next entry of the gather table buffer, and reads SEGLen bytes starting at SEGADR. While there are more bytes to be read in the data parcel, this process continues. If the channel’s gather table buffer is exhausted, the channel reads the next four long words of the link table into its gather table buffer. If a gather table buffer entry is encountered in which the N bit is set, the channel uses the SEGADR field in that word to find the next link table in the chain. The last byte of the required parcel size (Extent3) must coincide with the last byte of a memory segment, or unpredictable results may occur.

Now assume that the channel accesses its next data parcel using Pointer3 again, this time with length given by Length3. In this case the channel continues to the next line of the link table, and begins reading the memory segment specified there. As before, the channel concatenates memory segments from as many link table entries as necessary to obtain the required number of bytes (Length3).

Similarly, the next data parcel is obtained by using Pointer3 yet again, this time with length given by Extent4.

Assume that for the current descriptor type, the Extent4 data parcel is the last one to be accessed through Pointer3. Then the link table entry that supplies the last memory segment for Extent4 has the R bit set, signifying that this is the last entry in the chain of link tables.

### 14.3.5 Descriptor Types

An example of how the pointer dwords should be used with the various descriptor types to load keys, context, and text-data into the execution units, and how the required outputs should be unloaded is shown below.

**Table 14-10. Descriptor Format by Type**

Descriptor Type	field type	Pointer Dword1	Pointer Dword2	Pointer Dword 3	Pointer Dword4	Pointer Dword 5	Pointer Dword 6	Pointer Dword 7
0000_0 aesu_ctr_ nosnoop	Length	nil	Cipher IV	Cipher Key	In FIFO	Out FIFO	Cipher IV Out	nil
	Extent	undefined	undefined	undefined	nil	nil	nil	undefined
0001_0 common_ nosnoop	Length	nil	Cipher IV	Cipher Key	In FIFO	Out FIFO	Cipher IV Out	nil
	Extent	undefined	undefined	undefined	nil	nil	nil	undefined
1100_0 hmac_snoop_ aesu_ctr	Length	HMAC Key	HMAC Data	AES Key	AES Ctx	In FIFO	Out FIFO	HMAC Out
	Extent	undefined	undefined	undefined	nil	nil	nil	undefined
0000_1 ipsec_esp	Length	HMAC Key	HMAC Data	Cipher IV	Cipher Key	In FIFO	Out FIFO	Cipher IV Out
	Extent	undefined	undefined	undefined	nil	HMAC In	HMAC Out	undefined
0001_1 ccmp	Length	nil	AES-Ctx	AES Key	In FIFO	In FIFO	Out FIFO	AES-Ctx-Out
	Extent	undefined	undefined	undefined	nil	nil	nil	undefined
0010_1 srtp	Length	HMAC Key	AES-Ctx	AES Key	In FIFO	Out FIFO	HMAC Out	AES-Ctx-Out
	Extent	undefined	undefined	undefined	In FIFO	In FIFO	nil	undefined
1000_1 outbound tls_ssl_ block	Length	MAC Key	Cipher IV	Cipher Key	In FIFO Auth & Cipher	In FIFO Cipher Only	Out FIFO	Cipher IV Out
	Extent	undefined	undefined	undefined	In FIFO Auth only	MAC Out	nil	undefined

**Table 14-10. Descriptor Format by Type (continued)**

Descriptor Type	field type	Pointer Dword1	Pointer Dword2	Pointer Dword 3	Pointer Dword4	Pointer Dword 5	Pointer Dword 6	Pointer Dword 7
1000_1 inbound	Length	MAC Key	Cipher IV	Cipher Key	nil	In FIFO Auth & Cipher	Out FIFO	Cipher IV Out
tls_ssl_block	Extent	undefined	undefined	undefined	In FIFO Auth only	MAC In	MAC Out	undefined
1010_1	Length	nil	nil	nil	In1 (opt)	In2	In3	Out
raid_xor	Extent	nil	nil	nil	undefined	undefined	undefined	undefined
others		Reserved						

## 14.4 Execution Units

Execution unit (EU) is the term used for a functional block that performs the mathematical manipulations required by protocols used in cryptographic processing. The EUs are compatible with IPsec, SSL/TLS, iSCSI, SRTP, and 802.11i processing, and can work together to perform high level cryptographic tasks.

The following execution units are used in the SEC 2.2:

- Data encryption standard execution unit (DEU) for DES and 3DES as specified by FIPS 46-3
- Advanced encryption standard execution unit (AESU) implementing the Rijndael symmetric key cipher per FIPS-197. AESU can perform AES modes CBC, ECB, and CTR modes per NIST SP 800-38A, and CCM mode per NIST SP 800-38C.
- Message digest execution unit (MDEU), implementing MD5 per RFC 1321, and SHA-1, SHA-224, and SHA-256 per FIPS-180-2. In addition, HMAC is implemented per FIPS 198.

In addition, the two symmetric execution units, DEU and AESU, share their input and output FIFOs.

Working together, the EUs can perform high-level cryptographic tasks, such as IPsec Encapsulating Security Protocol (ESP). The remainder of this chapter provides details about the execution units themselves.

### 14.4.1 Data Encryption Standard Execution Unit (DEU)


This section contains details about the data encryption standard execution unit (DEU), including modes of operation, status and control registers, and shared symmetric FIFOs.

Most of the registers described here would not normally be accessed by the host. They are documented here mainly for debug purposes. In typical operation, the DEU is used through channel-controlled access, which means that most reads and writes of DEU registers are directed by the SEC channel. Driver software would perform host-controlled register accesses only on a few registers for initial configuration and error handling.

#### 14.4.1.1 DEU Mode Register (DEUMR)

The DEU mode register (DEUMR) contains three bits that are used to program DEU operation.

The DEUMR is cleared when the DEU is reset or re-initialized. Setting a reserved mode bit will generate a data error. If the mode register is modified during processing, a context error will be generated.



**Figure 14-7. DEU Mode Register (DEUMR)**

Table 14-11 describes DEUMR fields.

**Table 14-11. DEUMR Field Descriptions**

Bits	Name	Description
The following bits are described for information only. They are not under direct user control.		
0–54	—	Reserved
55	—	Reserved
The following bits are controlled through the MODE0 field of the descriptor header.		
56–60	—	Reserved
61	CE	CBC/ECB. If set, the DEU operates in cipher-block-chaining mode. If not set, DEU operates in electronic codebook mode. 0 ECB mode 1 CBC mode
62	TS	Triple/Single DES. If set, the DEU operates the Triple DES algorithm; if not set, DEU operates the single DES algorithm. 0 Single DES 1 Triple DES
63	ED	Encrypt/decrypt. If set, the DEU operates the encryption algorithm; if not set, DEU operates the decryption algorithm. 0 Perform decryption 1 Perform encryption

#### 14.4.1.2 DEU Key Size Register (DEUKSR)

The value of the DEU key size register (DEUKSR), shown in Figure 14-8, indicates the number of bytes of key memory that should be used in encrypting or decrypting. If the DEUMR is set for single DES, any value other than 8 bytes will automatically generate a key size error in the DEU interrupt status register (DEUISR). If the mode bit is set for triple DES, any value other than 16 bytes (112 bits for 2-key triple DES (K1 = K3) or 24 bytes (168 bits for 3-key triple DES) will generate an error. Triple DES always uses K1 to encrypt, K2 to decrypt, K3 to encrypt (any write to K1 duplicates that value into K3 in case 2-key 3DES is desired).




Figure 14-8. DEU Key Size Register (DEUKSR)

Table 14-12 shows the DEUKSR fields.

Table 14-12. DEUKSR Field Descriptions

Bits	Name	Description
0–51	—	Reserved
52–63	Key Size	8 bytes = 0x08 (only legal value if mode is single DES) 16 bytes = 0x10 (for 2 key 3DES, K1 = K3) 24 bytes = 0x18 (for 3 key 3DES)

### 14.4.1.3 DEU Data Size Register (DEUDSR)

The DEU data size register (DEUDSR), shown in Figure 14-9, stores the number of bits in the final message block, which must be 64. All data to be processed by the DEU must be a multiple of the DES algorithm block size of 64 bits; the DEU does not automatically pad messages out to 64-bit blocks. If a data size that is not a multiple of 64 bits is written, a data size error will be generated. Only bits 58–63 are checked to determine if there is a data size error. Since all upper bits are ignored, the entire message length (in bits) can be written to this register.

DEUDSR is cleared when the DEU is reset or re-initialized.




Figure 14-9. DEU Data Size Register (DEUDSR)

### 14.4.1.4 DEU Reset Control Register (DEURCR)

The DEU reset control register (DEURCR), shown in [Figure 14-10](#), allows three levels reset of just DEU, as defined by the three self-clearing bits.

Field	0	60	61	62	63	
	—			RI	MI	SR
Reset	0					
R/W	R/W					
Addr	DEU 0x3_2018					

**Figure 14-10. DEU Reset Control Register (DEURCR)**

[Table 14-13](#) describes DEURCR fields.

**Table 14-13. DEURCR Field Descriptions**

Bits	Names	Description
0–60	—	Reserved
61	RI	Reset interrupt. Writing this bit active high causes DEU interrupts signaling DONE and ERROR to be reset. It further resets the state of the DEU interrupt status register (DEUISR). 0 Do not reset 1 Reset interrupt logic
62	MI	Module initialization is nearly the same as software reset, except that the interrupt control register remains unchanged. this module initialization includes execution of an initialization routine, completion of which is indicated by the RESET_DONE bit in the DEU status register 0 Do not reset 1 Reset most of DEU
63	SR	Software reset is functionally equivalent to hardware reset (the RESET# pin), but only for DEU. All registers and internal state are returned to their defined reset state. Upon negation of SW_RESET, the DEU will enter a routine to perform proper initialization of the parameter memories. The RESET_DONE bit in the DEU status register (DEUSR) will indicate when this initialization routine is complete 0 Do not reset 1 Full DEU reset



### 14.4.1.5 DEU Status Register (DEUSR)

The DEU status register (DEUSR), displayed in [Figure 14-11](#), contains 6 fields that reflect the state of DEU internal signals. The DEUSR is read-only. Writing to this location will result in address error being reflected in the DEU interrupt status register (DEUISR).

Field	0	39	40	47	48	55	56	57	58	59	60	61	62	63	
	—			OFL		IFL		—		HALT	—		IE	ID	RD
Reset	0														
R/W	R														
Addr	DEU 0x3_2028														

**Figure 14-11. DEU Status Register (DEUSR)**

[Table 14-14](#) describes the DEUSR fields.

**Table 14-14. DEUSR Field Descriptions**

Bits	Name	Description
0–39	—	Reserved
40–47	OFL	The number of dwords currently in the output FIFO
48–55	IFL	The number of dwords currently in the input FIFO
56–57	—	Reserved
58	HALT	Halt. Indicates that the DEU has halted due to an error. 0 DEU not halted 1 DEU halted <b>Note:</b> Because the error causing the DEU to stop operating may be masked before reaching the interrupt status register (ISR), the DEU interrupt status register (DEUISR) is used to provide a second source of information regarding errors preventing normal operation.
59–60	—	Reserved
61	IE	Interrupt error. This status bit reflects the state of the ERROR interrupt signal, as sampled by the controller interrupt status register ( <a href="#">Section 14.6.4.3, “Interrupt Status Register (ISR)”</a> ). 0 DEU is not signaling error 1 DEU is signaling error
62	ID	Interrupt done. This status bit reflects the state of the DONE interrupt signal, as sampled by the controller interrupt status register ( <a href="#">Section 14.6.4.3, “Interrupt Status Register (ISR)”</a> ). 0 DEU is not signaling done 1 DEU is signaling done
63	RD	Reset done. This status bit, when high, indicates that DEU has completed its reset sequence, as reflected in the signal sampled by the appropriate channel. 0 Reset in progress 1 Reset done <b>Note:</b> Reset Done resets to 0, but has typically switched to 1 by the time a user checks the register, indicating the EU is ready for operation.

### 14.4.1.6 DEU Interrupt Status Register (DEUI SR)

The DEU interrupt status register (DEUI SR), shown in [Figure 14-12](#), records occurrences of errors. Each bit in the DEUI SR can only be set if the corresponding bit of the DEU interrupt control register (DEUI CR) is zero (see [Section 14.4.1.7, “DEU Interrupt Control Register \(DEUI CR\)”](#)).

If the DEUI SR is non-zero, the DEU halts and the DEU error interrupt signal is asserted to the controller (see [Section 14.6.4.3, “Interrupt Status Register \(ISR\)”](#)). In addition, if the DEU is being operated through channel-controlled access, an interrupt signal is generated to the channel to which this EU is assigned. The EU error then appears in bit 55 of the channel pointer status register (see [Table 14-35](#)) and generates a channel error interrupt to the controller.

	0	49	50	51	52	53	54	55	56	57	58	59	60	61	62	63
Field	—			KPE	IE	ERE	CE	KSE	DSE	ME	AE	OFE	IFE	IFU	IFO	OFO
Reset	0															
R/W	R															
Addr	DEU 0x3_2030															

**Figure 14-12. DEU Interrupt Status Register (DEUI SR)**

[Table 14-15](#) describes DEUI SR fields.

**Table 14-15. DEUI SR Field Descriptions**

Bits	Name	Description
0–49	—	Reserved
50	KPE	Key parity error. Defined parity bits in the keys written to the key registers did not reflect odd parity correctly. (Note that key register 2 and key register 3 are checked for parity only if the appropriate DEU mode register bit indicates triple DES. Also, key register 3 is checked only if key size reg = 24. Key register 2 is checked only if key size reg = 16 or 24.) 0 No error detected 1 Key parity error
51	IE	Internal error. An internal processing error was detected while performing encryption. 0 No error detected 1 Internal error <b>Note:</b> This bit will be asserted any time an enabled error condition occurs and can only be cleared by setting the corresponding bit in the interrupt control register or by resetting the DEU.
52	ERE	Early read error. The DEU IV register was read while the DEU was performing encryption. 0 No error detected 1 Early read error
53	CE	Context error. A DEU key register, the key size register, data size register, mode register, or IV register was modified while DEU was performing encryption. 0 No error detected 1 Context error
54	KSE	Key size error. An inappropriate value (8 being appropriate for single DES, and 16 and 24 being appropriate for triple DES) was written to the DEU key size register 0 No error detected 1 Key size error

Table 14-15. DEUISR Field Descriptions (continued)

Bits	Name	Description
55	DSE	Data size error (DSE): A value was written to the DEU data size register that is not a multiple of 64 bits. 0 No error detected 1 Data size error
56	ME	Mode error. An illegal value was detected in the mode register. Note: writing to reserved bits in mode register is likely source of error. 0 No error detected 1 Mode error
57	AE	Address error. An illegal read or write address was detected within the DEU address space. 0 No error detected 1 Address error
58	OFE	Output FIFO error. The DEU output FIFO was detected non-empty upon write of DEU data size register. 0 No error detected 1 Output FIFO non-empty error
59	IFE	Input FIFO error. The DEU input FIFO was detected non-empty upon generation of DONE interrupt. 0 No error detected 1 Input FIFO non-empty error
60	IFU	Input FIFO underflow. The DEU input FIFO has been read while empty. 0 No error detected 1 Input FIFO has had underflow error
61	IFO	Input FIFO overflow. The DEU input FIFO has been pushed while full. 0 No error detected 1 Input FIFO has overflowed <b>Note:</b> When operated through channel-controlled access, the SEC implements flow control, and FIFO size is not a limit to data input. When operated through host-controlled access, the DEU cannot accept FIFO inputs larger than 256 bytes without overflowing.
62	OFU	Output FIFO underflow. The DEU output FIFO has been read while empty. 0 No error detected 1 Output FIFO has underflow error
63	OFO	Output FIFO overflow. The DEU output FIFO has been pushed while full. 0 No error detected 1 Output FIFO has overflowed

#### 14.4.1.7 DEU Interrupt Control Register (DEUICR)

The DEU interrupt control register (DEUICR), shown in [Figure 14-13](#), controls the result of detected errors. For a given error (as defined in [Section 14.4.1.6, “DEU Interrupt Status Register \(DEUISR\)”](#)), if the corresponding bit in the DEUICR is set, the error is ignored; no bit is set in the DEUISR, and no error interrupt occurs. If the corresponding bit is not set, then upon detection of an error, the DEUISR is updated to reflect the error, causing assertion of the error interrupt signal, and causing the module to halt processing.

Field	0	49	50	51	52	53	54	55	56	57	58	59	60	61	62	63
	—		KPE	IE	ERE	CE	KSE	DSE	ME	AE	OFE	IFE	IFU	IFO	OFU	OFO
Reset	3000															
R/W	R/W															
Addr	DEU 0x3_2038															

Figure 14-13. DEU Interrupt Control Register (DEUICR)

Table 14-16 describes DEUICR fields.

Table 14-16. DEUICR Field Descriptions

Bits	Name	Description
0–49	—	Reserved
50	KPE	Key parity error. The defined parity bits in the keys written to the key registers did not reflect odd parity correctly. (Note that key register 2 and key register 3 are only checked for parity if the appropriate DEU mode register bit indicates triple DES. 0 Key parity error enabled 1 Key parity error disabled
51	IE	Internal error. An internal processing error was detected while performing encryption. 0 Internal error enabled 1 Internal error disabled
52	ERE	Early read error. The DEU IV register was read while the DEU was performing encryption. 0 Early read error enabled 1 Early read error disabled
53	CE	Context error. A DEU key register, the key size register, the data size register, the mode register, or IV register was modified while DEU was performing encryption. 0 Context error enabled 1 Context error disabled
54	KSE	Key size error. An inappropriate value (8 being appropriate for single DES, and 16 and 24 being appropriate for Triple DES) was written to the DEU key size register 0 Key size error enabled 1 Key size error disabled
55	DSE	Data size error (DSE): A value that is not a multiple of 64 bits was written to the DEU data size register. 0 Data size error enabled 1 Data size error disabled
56	ME	Mode error. An illegal value was detected in the mode register. 0 Mode error enabled 1 Mode error disabled
57	AE	Address error. An illegal read or write address was detected within the DEU address space. 0 Address error enabled 1 Address error disabled
58	OFE	Output FIFO error. The shared symmetric output FIFO was detected non-empty upon write of DEU data size register 0 Output FIFO non-empty error enabled 1 Output FIFO non-empty error disabled

Table 14-16. DEUICR Field Descriptions (continued)

Bits	Name	Description
59	IFE	Input FIFO error. The shared symmetric input FIFO was detected non-empty upon generation of done interrupt 0 Input FIFO non-empty error enabled 1 Input FIFO non-empty error disabled
60	IFU	Input FIFO underflow. The shared symmetric input FIFO has been read while empty. 0 Input FIFO underflow error enabled 1 Input FIFO underflow error disabled
61	IFO	Input FIFO overflow. The shared symmetric input FIFO has been pushed while full. 0 Input FIFO overflow error enabled 1 Input FIFO overflow error disabled <b>Note:</b> When operated through channel-controlled access, SEC implements flow control, and FIFO size is not a limit to data input. When operated through host-controlled access, the DEU cannot accept FIFO inputs larger than 256 bytes without overflowing.
62	OFU	Output FIFO underflow. The shared symmetric output FIFO has been read while empty. 0 Output FIFO underflow error enabled 1 Output FIFO underflow error disabled
63	OFO	Output FIFO overflow. The shared symmetric output FIFO has been pushed while full. 0 Output FIFO overflow error enabled 1 Output FIFO overflow error disabled

#### 14.4.1.8 DEU End-of-Message Register (DEUEMR)

The DEU end-of-message register (DEUEMR), shown in Figure 14-14, indicates a DES operation may be completed. After the final message block is written to the input FIFO, the DEUEMR must be written. The value in the data size register will be used to determine how many bits of the final message block (always 64) will be processed. Note that this register has no data size, and during the write operation, the host data bus is not read. Hence, any data value is accepted. Normally, a write operation with a zero data value is performed. Moreover, no read operation from this register is meaningful, but no error is generated, and a zero value is always returned. Writing to the DEUEMR is merely a trigger causing the DEU to process the final block of a message, allowing it to signal DONE.

	0	63
Field	DEU End-of-Message	
Reset	0	
R/W	W	
Addr	DEU 0x3_2050	

Figure 14-14. DEU End-of-Message Register (DEUEMR)

#### 14.4.1.9 DEU IV Register (DEUIV)

For CBC mode, the initialization vector is written to and read from the DEU IV register (DEUIV). The value of this register changes as a result of the encryption process and reflects the context of the DEU. Reading this memory location while the module is processing data generates an error interrupt.

### 14.4.1.10 DEU Key Registers (DEUK $n$ )

The DEU uses three write-only key registers (DEUK1, DEUK2, and DEUK3) to perform encryption and decryption. In single DES mode, only DEUK1 may be written. The value written to DEUK1 is simultaneously written to DEUK3, auto-enabling the DEU for 112-bit triple DES if the key size register indicates 2-key 3DES is to be performed (key size = 16 bytes). To operate in 168-bit triple DES, DEUK1 must be written first, followed by a write to DEUK2 and DEUK3.

Reading any of these memory locations generates an address error interrupt.

### 14.4.1.11 DEU FIFOs

DEU uses the symmetric shared input FIFO/output FIFO pair to hold data before and after the encryption process. These FIFOs are multiply addressable, but those multiple addresses point only to the appropriate end of the appropriate FIFO. A write to anywhere in the DEU FIFO address space causes the 64-bit-word to be pushed onto the shared symmetric input FIFO, and for that FIFO to be configured as reserved for DEU, and a read from anywhere in the DEU FIFO address space causes a 64-bit-word to be popped off of the shared symmetric output FIFO. Overflows and underflows caused by reading or writing the shared symmetric FIFOs are reflected in the DEUIR.

## 14.4.2 Message Digest Execution Unit (MDEU)

This section contains details about the message digest execution unit (MDEU), including modes of operation, status and control registers, and FIFOs.

Most of the registers described here would not normally be accessed by the host. They are documented here mainly for debug purposes. In typical operation, the MDEU is used through channel-controlled access, which means that most reads and writes of MDEU registers are directed by the SEC channel. Driver software would perform host-controlled register accesses only on a few registers for initial configuration and error handling.

### 14.4.2.1 MDEU Mode Register (MDEUMR)

The MDEU mode register (MDEUMR) is used to program the function of the MDEU. Bits 56–63 of the MDEUMR are specified by the user through the MODE0 or MODE1 field of the descriptor header. The remaining bits are supplied by the channel and thus are not under direct user control.

The MDEUMR has two configurations, determined by the value of the NEW bit (see [Figure 14-15](#) and [Figure 14-16](#)). The ‘old’ configuration (NEW = 0) is used by most descriptor types and is backward compatible with previous versions of SEC 2.0. The ‘new’ configuration (NEW = 1) is only used by descriptor type 1000\_1 ‘tls\_ssl\_block’.

The MDEUMR is cleared when the MDEU is reset or re-initialized. Setting a reserved mode bit will generate a data error. If the mode register is modified during processing, a context error is generated.

Field	0	53	54	55	56	57	58	59	60	61	62	63
Reset	0											
R/W	R/W											
Addr	MDEU 0x3_6000											

**Figure 14-15. MDEU Mode Register (MDEUMR) in 'Old' Configuration**


Table 14-17 describes MDEUMR fields in 'old' configuration.

**Table 14-17. MDEUMR in 'Old' Configuration**

Bits	Name	Description
The following bits are described for information only. They are not under direct user control.		
0–53	—	Reserved
54	NEW=0	Determines the configuration of the MDEU mode register (MDEUMR). This table shows the configuration for NEW = 0.
55	—	Reserved, must be set to zero
The following bits are controlled through the MODE0 or MODE1 fields of the descriptor header.		
56	CONT	Continue. Most operations will require this bit to be cleared. It is set only when the data to be hashed is spread across multiple descriptors. The value programmed in PD must be opposite to the value in this bit. 0 Do autopadding and complete the message digest. Used when the entire hash is performed with one descriptor, or on the last of a sequence of descriptors. 1 This hash will be continued in a subsequent descriptor. Do not autopad and do not complete the message digest.
57	CICV	Compare integrity check values 0 Normal operation; no ICV comparison 1 After the message digest (ICV) is computed, compare it to the data in the MDEU's input FIFO. If the ICVs do not match, send an error interrupt to the channel. The number of bytes to be compared is given by the ICV size register. Only applicable to descriptor types that provide for reading an ICV in value.
58	SMAC	Specifies whether to perform an SSL-MAC operation 0 Normal operation 1 Perform an SSL3.0 MAC operation. This requires a key and key length. If this is set then the HMAC bit should be 0.
59	INIT	Initialization bit. Most operations will require this bit to be set. Cleared only for operations that load context from a known intermediate hash value. 0 Do not initialize digest registers. In this case the registers must be loaded from a hash context pointer in the descriptor. When the data to be hashed is spread across multiple descriptors, this bit must be 0 on all but the first descriptor. 1 Do an algorithm-specific initialization of the digest registers.
60	HMAC	Specifies whether to perform an HMAC operation 0 Normal operation 1 Perform an HMAC operation. This requires a key and key length. If this is set then the SMAC bit should be 0.

**Table 14-17. MDEUMR in ‘Old’ Configuration (continued)**

Bits	Name	Description
61	PD	This bit must be programmed opposite to the CONT bit.
62–63	ALG	Message digest algorithm selection 00 SHA-160 algorithm (full name for SHA-1) 01 SHA-256 algorithm 10 MD5 algorithm 11 SHA-224 algorithm



**Figure 14-16. MDEU Mode Register (MDEUMR) in ‘New’ Configuration**

Table 14-18 describes MDEUMR fields in ‘new’ configuration.

**Table 14-18. MDEUMR in ‘New’ Configuration**

Bits	Name	Description
The following bits are described for information only. They are not under direct user control.		
0–52	—	Reserved
53	STIB	SSL/TLS inbound, block cipher 0 Normal operation. 1 Special operation only for SSL/TLS inbound, block cipher. Upon receiving end-of-message, the MDEU performs a calculation involving the last valid byte of data written into its input FIFO (which is Pad Length) to compute a final data size. The MDEU then processes the amount of data specified by this data size, and completes the message digest.
54	NEW=1	Determines the configuration of the MDEU mode register (MDEUMR). This table shows the configuration for NEW = 1.
55	—	Reserved, must be set to zero
The following bits are controlled through the MODE0 or MODE1 fields of the descriptor header.		
56	CONT	Continue. Most operations will require this bit to be cleared. Set only when the data to be hashed is spread across multiple descriptors. 0 Do autopadding and complete the message digest. Used when the entire hash is performed with one descriptor, or on the last of a sequence of descriptors. 1 This hash will be continued in a subsequent descriptor. Do not autopad and do not complete the message digest.
57	CICV	Compare integrity check values 0 Normal operation; no ICV comparison. 1 After the message digest (ICV) is computed, compare it to the data in the MDEU's input FIFO. If the ICVs do not match, send an error interrupt to the channel. The number of bytes to be compared is given by the ICV size register.



**Table 14-18. MDEUMR in ‘New’ Configuration (continued)**

Bits	Name	Description
58	SMAC	Specifies whether to perform an SSL-MAC operation 0 Normal operation 1 Perform an SSL3.0 MAC operation. This requires a key and key length. If this is set then the HMAC bit should be 0.
59	INIT	Initialization bit. Most operations will require this bit to be set. Cleared only for operations that load context from a known intermediate hash value. 0 Do not initialize digest registers. In this case the registers must be loaded from a hash context pointer in the descriptor. When the data to be hashed is spread across multiple descriptors, this bit is set on all but the first descriptor. 1 Do an algorithm-specific initialization of the digest registers.
60	HMAC	Specifies whether to perform an HMAC operation 0 Normal operation 1 Perform an HMAC operation. This requires a key and key length. If this is set then the SMAC bit should be 0.
61	EALG	The EALG (Extended Algorithm bit) and ALG (Algorithm) bits together specify the message digest algorithm, as follows:
62–63	ALG	000 SHA-160 algorithm (full name for SHA-1) 001 SHA-256 algorithm 010 MD5 algorithm 011 SHA-224 algorithm Others: Reserved

#### 14.4.2.2 Recommended Settings for MDEUMR

The most common task likely to be executed through the MDEU is HMAC generation. HMACs are used to provide message integrity within a number of security protocols, including IPsec, and TLS. The SSL 3.0 protocol uses a slightly different ‘SSL-MAC’. If an HMAC or SSL-MAC is to be performed using a single descriptor (with the MDEU acting as sole or secondary EU), the following mode register bit settings should be used:

**Table 14-19. Mode Register—HMAC or SSL-MAC Generated by Single Descriptor**

Bits	Field	Value	
		For HMAC	For SSL-MAC
56	CONT	0 (off)	0 (off)
58	SMAC	0 (on)	1 (on)
59	INIT	1 (on)	1 (on)
60	HMAC	1 (on)	0 (on)

To generate an HMAC for a message that is spread across a sequence of descriptors, the following mode register bit settings should be used:

**Table 14-20. Mode Register—HMAC Generated Across a Sequence of Descriptors**

Bits	Field	Value		
		First Descriptor	Middle Descriptor(s)	Final Descriptor
56	CONT	1 (on)	1 (on)	0 (off)
59	INIT	1 (on)	0 (off)	0 (off)
60	HMAC	1 (on)	0 (off)	1 (on)

All descriptors other than the final descriptor must output the intermediate message digest for the next descriptor to reload as MDEU context.

SSL-MAC operations cannot be spread across a sequence of descriptors.

Additional information on descriptors can be found in [Section 14.3, “Descriptor Overview.”](#)

### 14.4.2.3 MDEU Key Size Register (MDEUKSR)

The MDEU key size register (MDEUKSR), shown in [Figure 14-17](#), indicates the number of bytes of key memory that should be used in HMAC generation. The MDEU supports at most 64 bytes of key. The MDEU will generate a key size error if the value written to the MDEUKSR exceeds 64 bytes.

Field	0	56	57	63
Reset	0			
R/W	R/W			
Addr	MDEU 0x3_6008			
	—			Key Size (bytes)

**Figure 14-17. MDEU Key Size Register (MDEUKSR)**

### 14.4.2.4 MDEU Data Size Register (MDEUDSR)

The MDEU data size register (MDEUDSR), shown in [Figure 14-18](#), indicates the number of bits of data to be processed. The MDEU decrements this number as it processes data. A read of this register provides a snapshot of how much data remains to be processed.

The Data Size field is a 21-bit signed number. Values written to this register are added to the current register value. Multiple writes are allowed. The MDEU processes data when there is a positive value in this register and there is data available in the private MDEU input FIFO. (Negative values can arise in inbound processing, when it is necessary to hold back data from the MDEU until the pad length has been decrypted.)

Since the MDEU does not support bit offsets, bits 61–63 must be written as 0. Furthermore, when the CONT bit of the MDEU mode register (MDEUMR) is high, the data size must be a multiple of the 512-bit block size (that is, bits 55–63 must be written as 0). Violating either of these conditions causes a data size error (DSE in the MDEUISR).

The MDEUDSR is cleared when the MDEU is reset or re-initialized. At the end of processing, its contents has been decremented down to zero (unless there is an error interrupt).

### NOTE

Writing to the MDEUDSR allows the MDEU to enter auto-start mode. Therefore, the required context registers must be written prior to writing the data size.

Field	0	42	43	63
Reset	0			
R/W	R/W			
Addr	MDEU 0x3_6010			

Figure 14-18. MDEU Data Size Register (MDEUDSR)

### 14.4.2.5 MDEU Reset Control Register (MDEURCR)

The MDEU reset control register (MDEURCR), shown in Figure 14-19, allows three levels reset of just the MDEU, as defined by the three self-clearing bits.

Field	0	60	61	62	63
Reset	0				
R/W	R/W				
Addr	MDEU 0x3_6018				

Figure 14-19. MDEU Reset Control Register (MDEURCR)

Table 14-21 describes MDEURCR fields.


Table 14-21. MDEURCR Field Descriptions

Bits	Name	Description
0–60	—	Reserved
61	RI	Reset interrupt. Writing this bit active high causes MDEU interrupts signaling DONE and ERROR to be reset. It further resets the state of the MDEUISR. 0 No reset 1 Reset interrupt logic
62	MI	Module initialization is nearly the same as software reset, except that the MDEUICR remains unchanged. 0 No reset 1 Reset most of MDEU
63	SR	Software reset is functionally equivalent to hardware reset (the RESET# pin), but only for the MDEU. All registers and internal state are returned to their defined reset state. 0 No reset 1 Full MDEU reset

### 14.4.2.6 MDEU Status Register (MDEUSR)

The MDEU status register (MDEUSR), as seen in [Figure 14-20](#), reflects the state of the MDEU internal signals. The majority of these internal signals reflect the state of low-level MDEU functions, such as data padding, key padding, and so on, and are not important to the user, however the user should be aware that reads of this register especially during processing are likely to return non-zero values for many bits between 0–57. The four signals shown are those which are most likely to be of interest to the user.

The MDEUSR is read-only. Writing to this location will result in address error being reflected in the MDEUISR.



**Figure 14-20. MDEU Status Register (MDEUSR)**

[Table 14-22](#) describes the MDEUSR fields.

**Table 14-22. MDEUSR Field Descriptions**

Bits	Name	Description
0–57	—	Reserved
58	HALT	Halt. Indicates that the MDEU has halted due to an error. 0 MDEU not halted 1 MDEU halted <b>Note:</b> Because the error causing the MDEU to stop operating may be masked before reaching the interrupt status register, the MDEUISR is used to provide a second source of information regarding errors preventing normal operation.
59–60	ICR	Integrity check result 00 No integrity check was performed. 01 The integrity check passed. 10 The integrity check failed. 11 Reserved
61	IE	Interrupt error. This status bit reflects the state of the ERROR interrupt signal, as sampled by the controller interrupt status register ( <a href="#">Section 14.6.4.3, “Interrupt Status Register (ISR)”</a> ). 0 MDEU is not signaling error 1 MDEU is signaling error

Table 14-22. MDEUSR Field Descriptions (continued)

Bits	Name	Description
62	ID	Interrupt done. This status bit reflects the state of the DONE interrupt signal, as sampled by the controller interrupt status register (Section 14.6.4.3, “Interrupt Status Register (ISR”). 0 MDEU is not signaling done 1 MDEU is signaling done
63	RD	Reset done. This status bit, when high, indicates that MDEU has completed its reset sequence, as reflected in the signal sampled by the channel. 0 Reset in progress 1 Reset done <b>Note:</b> Reset done resets to 0, but has typically switched to 1 by the time a user checks the register, indicating the EU is ready for operation.

### 14.4.2.7 MDEU Interrupt Status Register (MDEUIR)

The MDEU interrupt status register (MDEUIR), shown in Figure 14-21, tracks the state of possible errors, if those errors are not masked, through the MDEU interrupt control register (MDEUICR).

	0		49	50	51	52	53	54	55	56	57	58	60	61	62	63
Field	—			ICE	—	IE	ERE	CE	KSE	DSE	ME	AE	—	IFO	—	
Reset	64'h1000															
R/W	R															
Addr	MDEU 0x3_6030															

Figure 14-21. MDEU Interrupt Status Register (MDEUIR)

Table 14-23 describes the MDEUIR fields.

Table 14-23. MDEUIR Field Descriptions


Bits	Name	Description
0–48	—	Reserved
49	ICE	Integrity check error 0 No error detected 1 Integrity check error detected. An ICV check was performed and the supplied ICV did not match the one computed by the MDEU.
50	—	Reserved
51	IE	Internal error. Indicates the MDEU has been locked up and requires a reset before use. 0 No internal error detected 1 Internal error detected <b>Note:</b> This bit will be asserted any time an enabled error condition occurs and can only be cleared by resetting the MDEU.
52	ERE	Early read error. The MDEU context was read before the MDEU completed the hashing operation. 0 No error detected 1 Early read error

**Table 14-23. MDEUI SR Field Descriptions (continued)**

Bits	Name	Description
53	CE	Context error. The MDEU key register, key size register, or data size register was modified while MDEU was hashing. 0 No error detected 1 Context error
54	KSE	Key size error. A value greater than 64 bytes was written to the MDEU key size register (MDEUKSR). 0 No error detected 1 Key size error
55	DSE	Data size error. A value not a multiple of 512 bits while the MDEU mode register (MDEUMR) CONT bit is high. 0 No error detected 1 Data size error
56	ME	Mode error. Will be set if any of these error conditions is detected: <ul style="list-style-type: none"> <li>Any reserved bit of the mode register is set</li> <li>The ALG field of the mode register contains an illegal value</li> </ul> 0 No error detected 1 Mode error
57	AE	Address error. An illegal read or write address was detected within the MDEU address space. 0 No error detected 1 Address error
58–60	—	Reserved
61	IFO	Input FIFO overflow. The MDEU input FIFO has been pushed while full. 0 No overflow detected 1 Input FIFO has overflowed <b>Note:</b> When operated through channel-controlled access, the SEC implements flow control, and FIFO size is not a limit to data input size. When operated through host-controlled access, the MDEU cannot accept FIFO inputs larger than 256 bytes without overflowing.
62–63	—	Reserved

### 14.4.2.8 MDEU Interrupt Control Register (MDEUI CR)

The MDEU interrupt control register (MDEUI CR), shown in Figure 14-22, controls the result of detected errors. For a given error (as defined in Section 14.4.2.7, “MDEU Interrupt Status Register (MDEUI SR)”), if the corresponding bit in this register is set, then the error is disabled; no error interrupt occurs and the MDEU interrupt status register (MDEUI SR) is not updated to reflect the error. If the corresponding bit is not set, then upon detection of an error, the interrupt status register is updated to reflect the error, causing assertion of the error interrupt signal, and causing the module to halt processing.



**Figure 14-22. MDEU Interrupt Control Register (MDEUI CR)**

Table 14-23 describes the MDEUICR fields.

**Table 14-24. MDEUICR Field Descriptions**

Bits	Name	Description
0–48	—	Reserved
49	ICE	Integrity check error. The supplied ICV did not match the one computed by the MDEU. 0 Integrity check error enabled 1 Integrity check error disabled
50	—	Reserved
51	IE	Internal error. An internal processing error was detected while performing hashing. 0 Internal error enabled 1 Internal error disabled
52	ERE	Early read error. The MDEU register was read while the MDEU was performing hashing. 0 Early read error enabled 1 Early read error disabled
53	CE	Context error. The MDEU key register, the key size register, the data size register, or the mode register, was modified while the MDEU was performing hashing. 0 Context error enabled 1 Context error disabled
54	KSE	Key size error. A value outside the bounds 64 bytes was written to the MDEU key size register (MDEUKSR). 0 Key size error enabled 1 Key size error disabled
55	DSE	Data size error. An inconsistent value was written to the MDEU data size register (MDEUDSR). 0 Data size error enabled 1 Data size error disabled
56	ME	Mode error. An illegal value was detected in the mode register (MDEUMR). 0 Mode error enabled 1 Mode error disabled
57	AE	Address error. An illegal read or write address was detected within the MDEU address space. 0 Address error enabled 1 Address error disabled
58–60	—	Reserved
61	IFO	Input FIFO overflow. The MDEU input FIFO has been pushed while full. 0 Input FIFO overflow error enabled 1 Input FIFO overflow error disabled
62–63	—	Reserved

#### 14.4.2.9 MDEU ICV Size Register

The MDEU ICV size register, shown in [Figure 14-23](#), stores the number of bytes of the ICV result to be compared if the MDEU performs ICV comparison (see Section 14.4.2.1 "MDEU Mode Register (MDEUMR)").

The MDEU ICV size register is cleared when the MDEU is reset or re-initialized.

Field	0	56	57	63
Reset	0			
R/W	R/W			
Addr	MDEU 0x3_6040			

Figure 14-23. MDEU ICV Size Register

#### 14.4.2.10 MDEU End-of-Message Register (MDEUEMR)

The end-of-message register in the MDEU (MDEUEMR), shown in Figure 14-24, is used to indicate that an authentication operation may be completed. After the final message block is written to the private MDEU input FIFO, the MDEUEMR must be written. The value in the data size register will be used to determine how many bits of the final message block (always 512) will be processed. Note that this register has no data size, and during the write operation, the host data bus is not read. Hence, any data value is accepted. Normally, a write operation with a zero data value is performed. Moreover, no read operation from this register is meaningful, but no error is generated, and a zero value is always returned. Writing to the MDEUEMR is merely a trigger causing the MDEU to process the final block of a message, allowing it to signal DONE.

Field	0	63
Reset	0	
R/W	W	
Addr	MDEU 0x3_6050	

Figure 14-24. MDEU End-of-Message Register (MDEUEMR)

#### 14.4.2.11 MDEU Context Registers

In the MDEU, context consists of the hash plus the message length count. Write access to the MDEU context register block allows continuation of a previous hash. Reading these registers provide the resulting message digest or HMAC, along with an aggregate bit count.

#### NOTE

SHA-1, SHA-224, and SHA-256 are big endian. MD5 is little endian. The MDEU module internally reverses the endianness of the five registers A, B, C, D, and E upon writing to or reading from the MDEU context if the MDEU mode register indicates MD5 is the hash of choice. Most other endian considerations are performed as 8-byte swaps. In this case, 4-byte endianness swapping is performed within the A, B, C, D, and E fields as individual registers. Reading this memory location while the module is not done will generate an error interrupt.



After a power-on reset, all the MDEU context register values are cleared to 0. Figure 14-25 shows how the MDEU context registers are initialized if the INIT bit is set in the MDEU mode register. All registers are initialized, regardless of mode selected, however only the appropriate context register values are used in hash generation per the mode selected. The user typically does not care about the MDEU context register initialization values, however they are documented for completeness in the event the user reads these registers using host-controlled access. MDEU reset through the MDEU reset control register (MDEURCR) (Figure 14-19) or SEC global software reset (Figure 14-47) does not clear these registers.

	0	31	32	63	
<b>Name</b>	<b>A</b>		<b>B</b>		Context offset 0x3_6100
MD-5	0x01234567		0x89ABCDEF		
SHA-1	0x67452301		0xEFCDAB89		
SHA-224	0xC1059ED8		0x367CD507		
SHA-256	0x6A09E667		0xBB67AE85		
<b>Name</b>	<b>C</b>		<b>D</b>		Context offset 0x3_6108
MD-5	0xFEDCBA98		0x76543210		
SHA-1	0x98BADCFE		0x10325476		
SHA-224	0x3070DD17		0xF70E5939		
SHA-256	0x3C6EF372		0xA54FF53A		
<b>Name</b>	<b>E</b>		<b>F</b>		Context offset 0x3_6110
MD-5	0xF0E1D2C3		0x8C68059B		
SHA-1	0xC3D2E1F0		0x9B05688C		
SHA-224	0xFFC00B31		0x68581511		
SHA-256	0x510E527F		0x9B05688C		
<b>Name</b>	<b>G</b>		<b>H</b>		Context offset 0x3_6118
MD-5	0xABD9831F		0x19CDE05B		
SHA-1	0x1F83D9AB		0x5BE0CD19		
SHA-224	0x64F98FA7		0xBEFA4FA4		
SHA-256	0x1F83D9AB		0x5BE0CD19		
<b>Name</b>	<b>Message Length Count</b>				Context offset 0x3_6120
Reset	0				

Figure 14-25. MDEU Context Registers

#### 14.4.2.12 MDEU Key Registers

The MDEU maintains eight 64-bit registers for writing an HMAC key. The IPAD and OPAD operations are performed automatically on the key data when required.

**NOTE**

SHA-1, SHA-224, and SHA-256 are big endian. MD5 is little endian. The MDEU module internally reverses the endianness of the key upon writing to or reading from the MDEU key registers if the MDEU mode register (MDEUMR) indicates that MD5 is the hash of choice.

**14.4.2.13 MDEU FIFOs**

MDEU uses a private input FIFO to hold data to be hashed. The input FIFO is multiply addressable, but those multiple addresses point only to the write (push) end of the FIFO. A write to anywhere in the MDEU FIFO address space causes the 64-bit-words to be pushed onto the MDEU input FIFO, and a read from anywhere in the MDEU FIFO address space returns all zeros.

**NOTE**

SHA-1, SHA-224, and SHA-256 are big endian. MD5 is little endian. The MDEU module internally reverses the endianness of the key upon writing to or reading from the MDEU key registers if the MDEU mode register (MDEUMR) indicates that MD5 is the hash of choice.

**14.4.3 Advanced Encryption Standard Execution Unit (AESU)**

This section contains details about the advanced encryption standard execution unit (AESU), including modes of operation, status and control registers, and FIFOs.

Most of the registers described here would not normally be accessed by the host. They are documented here mainly for debug purposes. In typical operation, the AESU is used through channel-controlled access, which means that most reads and writes of AESU registers are directed by the SEC channel. Driver software would perform host-controlled register accesses only on a few registers for initial configuration and error handling.

**14.4.3.1 AESU Mode Register (AESUMR)**

The AESU mode register (AESUMR), shown in [Figure 14-26](#), contains 7 bits that are used to program the AESU.

AESUMR is cleared when the AESU is reset or re-initialized. Setting a reserved mode bit will generate a data error. If the mode register is modified during processing, a context error will be generated.

	0	50	51	53	54	55	56	57	58	59	60	61	62	63
Field	—		SCM		—		ECM		FM	IM	RDK	CM	ED	
Reset	0													
R/W	R/W													
Addr	AESU 0x3_4000													

**Figure 14-26. AESU Mode Register (AESUMR)**

Table 14-25 describes AESUMR fields.

**Table 14-25. AESUMR Field Descriptions**

Bits	Name	Description
The following bits are described for information only. They are not under direct user control.		
0–50	—	Reserved
51–53	SCM	Sub-cipher-mode. Specifies additional options specific to particular cipher modes. <ul style="list-style-type: none"> <li>• XOR cipher mode: specifies the number of sources to be XORed together. Valid values are 2 and 3.</li> <li>• For all other cipher modes, this field must be 0.</li> </ul>
54–55	—	Reserved, must be set to zero.
The following bits are controlled through the MODE0 field of the descriptor header.		
56–57	ECM	Extend cipher mode. Used in combination with bits 61–62 “Cipher Mode” to define the mode of AES operation. See Table 14-26 for mode bit combinations.
58	FM	Final MAC (FM). Processes final message block and generates final MAC tag at end of message processing (CCM mode only) <ul style="list-style-type: none"> <li>0 Do not generate final MAC tag</li> <li>1 Generate final MAC tag after CCM processing is complete.</li> </ul>
59	IM	Initialize MAC(IM). Initializes AESU for new message (CCM mode only) <ul style="list-style-type: none"> <li>0 Do not initialize (context will be loaded by host)</li> <li>1 Initialize new message with nonce</li> </ul>
60	RDK	Restore decrypt key (RDK). Specifies that key data write will contain pre-expanded key (decrypt mode only). See note below on use of RDK bit. <ul style="list-style-type: none"> <li>0 Expand the user key prior to decrypting the first block</li> <li>1 Do not expand the key. The expanded decryption key will be written following the context switch.</li> </ul>
61–62	CM	Cipher mode. Used in combination with bits 56–57 (Extend Cipher Mode) to define the mode of AES operation. See Table 14-26 for mode bit combinations.
63	ED	Encrypt/Decrypt. If set, AESU operates the encryption algorithm; if not set, AESU operates the decryption algorithm. <ul style="list-style-type: none"> <li>0 Perform decryption</li> <li>1 Perform encryption</li> </ul> Note: This bit is ignored if CM is set to ‘11’ (CTR mode).

**Table 14-26. AES Cipher Modes**

Mode	ECM (56–57)	CM (61–62)
ECB	00	00
CBC	00	01
CTR	00	11
SRT <sup>1</sup>	01	11
CCM (without ICV comparison)	10	00
CCM with ICV comparison	11	00

**Table 14-26. AES Cipher Modes (continued)**

Mode	ECM (56–57)	CM (61–62)
XOR	11	11
Reserved	all others	

<sup>1</sup> SRT is not a new AES mode, it is an AESU method of performing AES-CTR mode with reduced context loading overhead specifically for performing SRTP. It should be used with descriptor type 0010\_0 'srtp'. See [Section 14.4.3.9.3, “Context for SRT Mode,”](#) for more information on how SRT mode reduces context loading overhead.

### NOTE

Note on Restore decrypt key (RDK)—In most networking applications, the decryption of an AES protected packet will be performed as a single operation. However, if circumstances dictate that the decryption of a message should be split across multiple descriptors, the AESU allows the user to save the decrypt key, and the active AES context, to memory for later re-use. This saves the internal AESU processing overhead associated with regenerating the decryption key schedule (~12 AESU clock cycles for the first block of data to be decrypted.)

The use of RDK is completely optional, as the Input time of the preserved decrypt key may exceed the ~12 cycles required to restore the decrypt key for processing the first block.

To use RDK, the following procedure is recommended:

The descriptor type used in decryption of the first portion of the message is '0100\_0- AESU Key Expand Output'. The AESU mode must be 'Decrypt'. See [Table 14-7](#) for more information. The descriptor will cause the SEC to write the contents of the Context registers and the key registers (containing the expanded decrypt key) to memory.

To process the remainder of the message, use a 'common' descriptor type (0001\_0), and set the 'restore decrypt key' mode bit. Load the context registers and the expanded decrypt key with previously saved key and context data from the first message. The key size is written as before (16, 24, or 32 bytes).

### 14.4.3.2 AESU Key Size Register (AESUKSR)

The AESU key size register (AESUKSR), shown in [Figure 14-27](#), stores the number of bytes in the key (16, 24, 32). Any key data beyond the number of bytes in the key size register will be ignored. The AESUKSR is cleared when the AESU is reset or re-initialized. If a key size other than 16, 24, or 32 bytes is specified, an illegal key size error will be generated. If the key size register is modified during processing, a context error will be generated.

Field	0	51	52	63
Reset	0			
R/W	R/W			
Addr	AESU 0x3_4008			

Figure 14-27. AESU Key Size Register (AESUKSR)

### 14.4.3.3 AESU Data Size Register (AESUDSR)

The AESU data size register (AESUDSR), shown in Figure 14-28, stores the number of bits in the final message block. Acceptable sizes vary depending on the AES mode selected. In ECB, CBC, and CTR mode, the message processed by the AESU must be a multiple of 128 bits; the AESU does not automatically pad messages out to 128-bit blocks. In CCM mode, data size must be a multiple of 8 bits. In XOR mode the data size must be a multiple of 256 bits (32 bytes). If an improper data size is written, a data size error is generated. Only the lowest 3, 7, or 8 bits of the data size register are checked to determine if there is a data size error. Since all upper bits are ignored, the entire message length (in bits) can be written to this register.

The AESUDSR is cleared when the AESU is reset or re-initialized.

Writing to the AESUDSR signals the AESU to start processing data from the shared symmetric input FIFO as soon as it is available. If the value of data size is modified during processing, a context error is generated.

Field	0	51	52	63
Reset	0			
R/W	R/W			
Addr	AESU 0x3_4010			

Figure 14-28. AESU Data Size Register (AESUDSR)

### 14.4.3.4 AESU Reset Control Register (AESURCR)

The AESU reset control register (AESURCR), shown in Figure 14-29, allows three levels reset of just AESU, as defined by the three self-clearing bits.

Field	0	60	61	62	63
Reset	0				
R/W	R/W				
Addr	AESU 0x3_4018				

Figure 14-29. AESU Reset Control Register (AESURCR)


Table 14-13 describes AESURCR fields.

**Table 14-27. AESURCR Field Descriptions**

Bits	Names	Description
0–60	—	Reserved
61	RI	Reset interrupt. Writing this bit active high causes AESU interrupts signaling DONE and ERROR to be reset. It further resets the state of the AESU interrupt status register (AESUISR). 0 Do not reset 1 Reset interrupt logic
62	MI	Module initialization is nearly the same as software reset, except that the interrupt control register remains unchanged. This module initialization includes execution of an initialization routine, completion of which is indicated by the RESET_DONE bit in the AESU status register (AESUSR) 0 Do not reset 1 Reset most of AESU
63	SR	Software reset is functionally equivalent to hardware reset (the RESET# pin), but only for AESU. All registers and internal state are returned to their defined reset state. Upon negation of SW_RESET, the AESU will enter a routine to perform proper initialization of the parameter memories. The RESET_DONE bit in the AESU status register will indicate when this initialization routine is complete 0 Do not reset 1 Full AESU reset

### 14.4.3.5 AESU Status Register (AESUSR)

The AESU status register (AESUSR), shown in Figure 14-30, is a read-only register that reflects the state of six status outputs. Writing to this location will result in an address error being reflected in the AESU interrupt status register (AESUISR).



**Figure 14-30. AESU Status Register (AESUSR)**

Table 14-28 describes AESUSR fields.

**Table 14-28. AESUSR Field Descriptions**

Bits	Name	Description
0–39	—	Reserved
40–47	OFL	The number of dwords currently in the output FIFO
48–55	IFL	The number of dwords currently in the input FIFO
56–57	—	Reserved

Table 14-28. AESUSR Field Descriptions (continued)

Bits	Name	Description
58	HALT	Halt. Indicates that the AESU has halted due to an error. 0 AESU not halted 1 AESU halted <b>Note:</b> Because the error causing the AESU to stop operating may be masked before reaching the interrupt status register, the AESU interrupt status register (AESUISR) is used to provide a second source of information regarding errors preventing normal operation.
59–60	ICR	Integrity check result 00 No integrity check was performed. 01 The integrity check passed. 10 The integrity check failed. 11 Reserved
61	IE	Interrupt error. This status bit reflects the state of the ERROR interrupt signal, as sampled by the controller interrupt status register (Section 14.6.4.3, “Interrupt Status Register (ISR)”) 0 AESU is not signaling error 1 AESU is signaling error
62	ID	Interrupt done. This status bit reflects the state of the DONE interrupt signal, as sampled by the controller interrupt status register (Section 14.6.4.3, “Interrupt Status Register (ISR)”) 0 AESU is not signaling done 1 AESU is signaling done
63	RD	Reset done. This status bit, when high, indicates that AESU has completed its reset sequence, as reflected in the signal sampled by the channel. 0 Reset in progress 1 Reset done <b>Note:</b> Reset done resets to 0, but has typically switched to 1 by the time a user checks the register, indicating the EU is ready for operation.

#### 14.4.3.6 AESU Interrupt Status Register (AESUISR)

The AESU interrupt status register (AESUISR), shown in Figure 14-31, tracks the state of possible errors, if those errors are not masked, through the AESU interrupt control register (AESUICR).

	0	48	49	50	51	52	53	54	55	56	57	58	59	60	61	62	63
Field	—	ICE	—	IE	ERE	CE	KSE	DSE	ME	AE	OFE	IFE	IFU	IFO	OFU	—	
Reset	0																
R/W	R																
Addr	AESU 0x3_4030																

Figure 14-31. AESU Interrupt Status Register (AESUISR)

Table 14-29 describes AESUISR fields.

**Table 14-29. AESUISR Field Descriptions**

Bits	Name	Description
0–48	—	Reserved
49	ICE	Integrity check error 0 No error detected 1 Integrity check error detected. An ICV check was performed and the supplied ICV did not match the one computed by the MDEU.
50	—	Reserved
51	IE	Internal error. An internal processing error was detected while the AESU was processing. 0 No error detected 1 Internal error <b>Note:</b> This bit will be asserted any time an enabled error condition occurs and can only be cleared by setting the corresponding bit in the interrupt control register (AESUICR) or by resetting the AESU.
52	ERE	Early read error. The AESU IV register was read while the AESU was processing. 0 No error detected 1 Early read error
53	CE	Context error. An AESU key register, the key size register, data size register, mode register, or IV register was modified while AESU was processing 0 No error detected 1 Context error
54	KSE	Key size error. An inappropriate value (not 16, 24 or 32 bytes) was written to the AESU key size register (AESUKSR). 0 No error detected 1 Key size error
55	DSE	Data size error (DSE): A value was written to the AESU data size register that is not a proper size. See <a href="#">Section 14.4.3.3, “AESU Data Size Register (AESUDSR)”</a> 0 No error detected 1 Data size error
56	ME	Mode error. Indicates that invalid data was written to a register or a reserved mode bit was set. 0 Valid data 1 Reserved or invalid mode selected
57	AE	Address error. An illegal read or write address was detected within the AESU address space. 0 No error detected 1 Address error
58	OFE	Output FIFO error. The shared symmetric output FIFO was detected non-empty upon write of AESU data size register. 0 No error detected 1 Output FIFO non-empty error
59	IFE	Input FIFO error. The shared symmetric input FIFO was detected non-empty upon generation of done interrupt. 0 No error detected 1 Input FIFO non-empty error
60	IFU	Input FIFO underflow. The AESU input FIFO has been read while empty. 0 No error detected 1 Input FIFO has had underflow error



Table 14-29. AESUISR Field Descriptions (continued)

Bits	Name	Description
61	IFO	Input FIFO overflow. The shared symmetric Input FIFO has been pushed while full. 0 No error detected 1 Input FIFO has overflowed <b>Note:</b> When operated through channel-controlled access, the SEC implements flow control, and FIFO size is not a limit to data input. When operated through host-controlled access, the AESU cannot accept FIFO inputs larger than 256 bytes without overflowing.
62	OFU	Output FIFO underflow. The shared symmetric output FIFO has been read while empty. 0 No error detected 1 Output FIFO has underflow error
63	—	Reserved

### 14.4.3.7 AESU Interrupt Control Register (AESUICR)

The AESU interrupt control register (AESUICR), shown in Figure 14-32, controls the result of detected errors. For a given error (as defined in Section 14.4.3.6, “AESU Interrupt Status Register (AESUISR)”), if the corresponding bit in this register is set, then the error is ignored; no error interrupt occurs and the interrupt status register (AESUISR) is not updated to reflect the error. If the corresponding bit is not set, then upon detection of an error, AESUISR is updated to reflect the error, causing assertion of the error interrupt signal, and causing the module to halt processing.

	0	48	49	50	51	52	53	54	55	56	57	58	59	60	61	62	63
Field	—	ICE	—	IE	ERE	CE	KSE	DSE	ME	AE	OFE	IFE	RSV	IFO	OFU	—	
Reset	1000																
R/W	R/W																
Addr	AESU 0x3_4038																

Figure 14-32. AESU Interrupt Control Register (AESUICR)

Table 14-30 describes the AESUICR fields.

Table 14-30. AESUICR Field Descriptions

Bits	Name	Description
0–48	—	Reserved
49	ICE	Integrity check error. The supplied ICV did not match the one computed by the MDEU. 0 Integrity check error enabled 1 Integrity check error disabled
50	—	Reserved
51	IE	Internal error. An internal processing error was detected while the AESU was processing. 0 Internal error enabled 1 Internal error disabled
52	ERE	Early read error. The AESU IV register was read while the AESU was processing. 0 Early read error enabled 1 Early read error disabled

Table 14-30. AESUICR Field Descriptions (continued)

Bits	Name	Description
53	CE	Context error. An AESU key register, the key size register, data size register, mode register, or IV register was modified while the AESU was processing. 0 Context error enabled 1 Context error disabled
54	KSE	Key size error. An inappropriate value (non 16, 24 or 32 bytes) was written to the AESU key size register (AESUKSR) 0 Key size error enabled 1 Key size error disabled
55	DSE	Data size error. Indicates that the number of bits to process is out of range. 0 Data size error enabled 1 Data size error disabled
56	ME	Mode error. Indicates that invalid data was written to a register or a reserved mode bit was set. 0 Mode error enabled 1 Mode error disabled
57	AE	Address error. An illegal read or write address was detected within the AESU address space. 1 Address error disabled 0 Address error enabled
58	OFE	Output FIFO error. The shared symmetric output FIFO was detected non-empty upon write of AESU data size register (AESUDSR) 0 Output FIFO non-empty error enabled 1 Output FIFO non-empty error disabled
59	IFE	Input FIFO error. The shared symmetric input FIFO was detected non-empty upon generation of done interrupt 0 Input FIFO non-empty error enabled 1 Input FIFO non-empty error disabled
60	—	Reserved
61	IFO	Input FIFO overflow. The shared symmetric input FIFO has been pushed while full. 0 Input FIFO overflow error enabled 1 Input FIFO overflow error disabled
62	OFU	Output FIFO underflow. The shared symmetric output FIFO has been read while empty. 0 Output FIFO underflow error enabled 1 Output FIFO underflow error disabled
63	—	Reserved

#### 14.4.3.8 AESU End-of-Message Register (AESUEMR)

The AESU end-of-message register (AESUEMR), shown in [Figure 14-33](#), is used to indicate an AES operation may be completed. After the final message block is written to the shared symmetric input FIFO, the AESUEMR must be written. The value in the data size register (AESUDSR) is used to determine how many bits of the final message block (always 128) will be processed. Writing to the AESUEMR causes the AESU to process the final block of a message, allowing it to signal DONE. A read of the AESUEMR will always return a zero value.

Field	0	63
Reset	AESU End of Message	
R/W	0	
Addr	W	
	AESU 0x3_4050	

Figure 14-33. AESU End-of-Message Register (AESUEMR)

### 14.4.3.9 AESU Context Registers

There are seven 64-bit context data registers that allow the host to read/write the contents of the context used to process the message. The context must be written prior to the key data. If the context registers are written during message processing, a context error will be generated. All context registers are cleared when a hard/soft reset or initialization is performed.

The context registers must be read when changing context and restored to their original values to resume processing an interrupted message (CBC, CTR, and CCM modes). For CTR and CCM mode, all seven 64-bit context registers must be read to retrieve context, and all seven must be written back to restore context. Effectively, the user must read the four empty 'place holder' context registers in addition to the three context registers holding the Counter and Counter Modulus Exponent when in CTR mode. The contents of the 'empty' context registers need not be preserved, but when restoring the CTR mode context, the 'empty' registers must be filled with 32 bytes of zeros before writing the saved Counter and Counter Modulus Exponent.

Context should be loaded with the lower bytes in the lowest 64-bit context register. The context registers are summarized in Figure 14-34.

Cipher Mode	Context Register (64 bits each)						
	1	2	3	4	5	6	7
ECB	—	—	—	—	—	—	—
CBC	IV1 <sup>1</sup>	IV2 <sup>1</sup>	—	—	—	—	—
CTR	—	—	—	—	Counter <sup>1</sup>		Counter Modulus Exponent <sup>1</sup>
SRT	Counter <sup>1</sup>		Counter Modulus Exponent (M) <sup>1</sup>	—	—	—	—
CCM	IV <sup>1</sup> / MAC Tag		Encrypted MAC <sup>2</sup> /Decrypted MAC/Encrypted Counter		Counter <sup>1</sup>		Counter Modulus Exponent <sup>1</sup> /header size/MAC size <sup>3</sup>

<sup>1</sup> Must be written at the start of a new message.

<sup>2</sup> Must be written at start of new CCM decryption.

<sup>3</sup> Header size/MAC size is only used if AES-CCM processing is suspended and resumed.

Figure 14-34. AESU Context Registers

### 14.4.3.9.1 Context for CBC Mode

Within the Context register, for use in CBC mode, are two 64-bit context data registers that allow the host to read/write the contents of the initialization vector (IV):

- IV1 holds the *least* significant bytes of the initialization vector (bytes 1–8).
- IV2 holds the *most* significant bytes of the initialization vector (bytes 9–16).

The IV must be written prior to the message data. If the IV registers are written during message processing, or the CBC mode bit is not set, a context error will be generated.

The IV registers may only be read after processing has completed, as indicated by the assertion of Interrupt Done DONE in the AESU status register as shown in [Section 14.4.3.5, “AESU Status Register \(AESUSR\).”](#) If the IV registers are read prior to assertion of Interrupt Done, an early read error will be generated.

The IV registers must be read when changing context and restored to resume processing an interrupted message (CBC mode only).

### 14.4.3.9.2 Context for Counter Mode

In counter mode, a random 128-bit initial counter value is incremented modulo  $2^M$  with each block processed. The running counter is encrypted and eXclusive-ORed with the plaintext to derive the ciphertext, or with the ciphertext to recover the plaintext. The modulus exponent M can be set between 8 and 128 in multiples of 8. The value of M is specified by writing to context register 3 as described in [Figure 14-34](#).

### 14.4.3.9.3 Context for SRT Mode

As was noted in the AESU mode register, SRT is not a new AES mode, it is an AESU method of performing AES-CTR mode with reduced context loading overhead specifically for performing SRTP. It should be used with descriptor type 0010\_0 ‘srtp’. As with counter mode, a random 128-bit initial counter value is incremented modulo  $2^M$  with each block processed. The running counter is encrypted and eXclusive-ORed with the plaintext to derive the ciphertext, or with the ciphertext to recover the plaintext. The modulus exponent M can be set between 8 and 128 in multiples of 8. The value of M is specified by writing to context register 3 as described in [Figure 14-34](#).

The only difference between SRT mode and CTR mode is in SRT mode, the AES Context is loaded and read through context registers 1–3, with no requirement to access context registers 4–7. In CTR mode, context registers 1–4 must be loaded with zeros, with the Counter and Modulus being loaded into and read from context registers 5–7.

### 14.4.3.9.4 Context for CCM Mode

The SEC AESU is capable of performing single pass encryption and MAC generation. The host is required to order the CCM context is such a way that the context can be fetched as a contiguous string into the context registers, prior to encryption/MAC generation or decryption/MAC validation. The context register contents for CCM mode is summarized in [Figure 14-35](#) and further described below.

		Context Registers						
		1	2	3	4	5	6	7
Encrypt (outbound)	Inputs	IV		0		Initial Counter		Counter Modulus Exponent
	Outputs	MAC	0	MIC	0			
Decrypt (inbound)	Inputs	IV		MIC	0	Initial Counter		Counter Modulus Exponent
	Outputs	Computed MAC	0	Decrypted MAC	0			

**Figure 14-35. AESU CCM Context Registers**

The context for CCM encryption/MAC generation is:

- Reg 1–2, session specific 128-bit initialization vector (from memory)
- Reg 3–4, 128 bits of zero padding
- Reg 5–6, session specific counter (initial counter value) (from memory)
- Reg 7, counter modulus exponent (msb<--lsb). Should be fixed at 0x0000\_0080.

Note that the counter modulus for CCM mode is currently defined as  $2^{128}$  making the exponent 128. This value has been made programmable in the SEC in case the final version of 802.11i uses a different counter modulus. Because this is a programmable field, it must be generated and stored along with other session specific information for loading into the AESU context register prior to CCM encryption.

CCM encryption processing—With the session specific key and context, the AESU will perform the following operations.

1. Initialize the IV, and encrypt with the symmetric key.
2. In CBC fashion, take the output of step 1, hash with the first block of plaintext, and encrypt with the symmetric key.
3. Continue as in step 2 until the final block of plaintext has been processed. The result of the encryption of the final block of plaintext with the symmetric key is the MAC tag. The full 128 bits of MAC data is written to context registers 1–2, for use in the next phase of CCM processing.

Once the MAC tag has been generated (step 3), the MAC tag, along with the plaintext is encrypted with the AESU operating in counter mode.

4. The first item to be encrypted in counter mode is the counter (initial counter value) from context registers 5–6. The counter is encrypted with the symmetric key, and the result is hashed with the MAC tag (retrieved from context registers 1–2) to produce the MIC (encrypted MAC), which is then stored in context registers 3–4. At the completion of CCM encrypt processing, this MIC is output to memory (per the descriptor pointer) for the host to append to the 802.11i frame. Note that the MIC written out to memory by the AESU is the full 128 bits. The host must only append the most-significant 64 bits to the frame as the MIC.

5. The counter value is incremented, and is then encrypted with the symmetric key. The result is then hashed with the first block of plaintext to produce the first block of cipher text. The ciphertext is placed in the shared symmetric output FIFO.
6. The counter continues to be incremented, and encrypted with the symmetric key, with the result hashed with each successive block of plaintext, until all plaintext has been converted to ciphertext. The SEC controller will manage FIFO reads and writes, fetching plaintext and writing ciphertext per the pointers provided in the descriptor. When all ciphertext and the MIC have been output, the CCM encrypt operation is complete.

The context for CCM decryption/MAC generation is:

- Reg 1–2, session specific 128-bit initialization vector (from memory)
- Reg 3–4, MIC (from received frame) + 64 bits of zero padding
- Reg 5–6, session specific counter (initial counter value) (from memory)
- Reg 7, counter modulus exponent (msb<--lsb). Should be fixed at 0x0000\_0080.

Note that the counter modulus for CCM mode is currently defined as  $2^{128}$ , making the exponent 128. This value has been made programmable in the SEC in case the final version of the IEEE 802.11i standard uses a different counter modulus. Because this is a programmable field, it must be generated and stored along with other session specific information for loading into the AESU context register prior to CCM decryption.

CCM decryption processing is the reverse of encryption;

With the session specific key and context, the AESU will perform the following operations.

1. Initialize the IV, and encrypt with the symmetric key. Simultaneously, the counter (initial counter value) from context registers 5–6 is encrypted with the symmetric key. The result is hashed with the encrypted MAC (from context registers 3–4), and the resulting original MAC is written to context registers 3–4, overwriting the encrypted MAC.

Note that the counter is encrypted with the symmetric key, however, the AESU should be set for 'decrypt' to perform the counter and CBC processes in the correct order.

2. The IEEE Std. 802.11 frame header is hashed with the encrypted IV. (The AESU automatically determines the header length.) Simultaneously, the counter is incremented, and is then encrypted with the symmetric key. The result is then hashed with the first block of ciphertext to produce the first block of plaintext. The plaintext is placed in the shared symmetric output FIFO, while simultaneously, in CBC fashion, a copy of the first block of plaintext is hashed with the output of encryption of the IEEE Std. 802.11 frame header. The output is encrypted with the symmetric key.
3. As each ciphertext block is converted to plaintext, the plaintext is CBC encrypted. When the final plaintext block has been processed, the CBC MAC (MAC tag) is written to context registers 1–2. The first 64 bits of the MAC tag are compared to the MAC tag recovered in step 1.

Note that for both encrypt and decrypt operations, if the IEEE Std. 802.11 frame is being processed as a whole (not split across multiple descriptors), the 'Initialize' and 'Final MAC' bits should be set in the AESU mode register.

### 14.4.3.9.5 AESU Key Registers

The AESU key registers hold from 16, 24, or 32 bytes of key data, with the first 8 bytes of key data written to key 1. Any key data written to bytes beyond the value written to the key size register will be ignored. The key data registers are cleared when the AESU is reset or re-initialized. If these registers are modified during message processing, a context error will be generated.

The key data registers may be read when changing context in decrypt mode. To resume processing, the value read must be written back to the key registers and the ‘restore decrypt key’ bit must be set in the mode register. This eliminates the overhead of expanding the key prior to starting decryption when switching context.

### 14.4.3.9.6 AESU FIFOs

The AESU fetches data 128 bits at a time from the shared symmetric input FIFO. During processing, the input data is encrypted or decrypted with the key and initialization vector (CBC mode only) and the results are placed in the shared symmetric output FIFO. The output size is the same as the input size.

Writing to the AESU FIFO address space places 64 bits of message data into the input FIFO and configures the shared symmetric FIFOs to be reserved by AESU. The input FIFO may be written any time the IFW signal is asserted (as indicated in the AESU status register). This will indicate that the number of bytes of available space is at or above the threshold specified in the mode register. There is no limit on the total number of bytes in a message. The number of bits in the final message block must be set in the data size register.

Reading from the AESU FIFO address space will pop 64 bits of message data from the shared symmetric output FIFO. The output FIFO may be read any time the OFR signal is asserted (as indicated in the AESU status register). This will indicate that the number of bytes in the output FIFO is at or above the threshold specified in the mode register.

## 14.5 Channel

The channel in the SEC manages the execution of each cryptographic task, making use of one or more of the SEC’s execution units (EUs). Control information and data pointers for a given task are stored in the form of a descriptor (see [Section 14.3.1, “Descriptor Structure”](#)) in system memory or in the channel itself. A descriptor determines what EUs will be used, how they will be configured, where to fetch needed data, and where to store the results. To invoke cryptographic tasks, the host constructs a descriptor, and writes a pointer to the descriptor into the channel’s fetch FIFO. The fetch FIFO can store up to 24 pointers.

Operations performed by the channel include the following (not necessarily in this order):

- If the channel is idle and its fetch FIFO is non-empty, read the next descriptor pointer from the fetch FIFO, and use this pointer to read the descriptor into the channel’s descriptor buffer.
- Request from the controller the assignment of one or more EUs for the use of the channel. Where necessary, configure the secondary EU to snoop input or output data intended for the primary EU.
- Upon notification of completion of the EU reset sequence, initialize mode registers in the assigned EU.
- Initialize EUs and write to EU registers such as key size and text-data size.

- Transfer data parcels (up to 32 Kbytes) from system memory into assigned EU input registers and FIFOs. This may involve using link tables to gather input data that has been split into multiple segments which are stored in various locations of system memory. For the RAID-XOR descriptor type, the channel rotates among three data sources, fetching 32 bytes from each source.
- Transfer data parcels (up to 32 Kbytes) from assigned EU output registers and FIFOs to system memory space. This may involve using link tables to scatter output data into multiple segments which are stored in various locations of system memory.
- Initialize the end-of-message register (where applicable) in the assigned EU upon completion of last EU write indicated by the descriptor. The channel will wait for an indication from the EU that processing of input text-data is complete before proceeding with further activity after writing end-of-message.
- Reset assigned EU(s).
- Release assigned EU(s).
- When a descriptor has been completely processed, provide feedback to the host, in the form of interrupt and/or descriptor header write-back to system memory.
- When descriptor processing is halted due to an error, provide feedback to the host through an interrupt.

The channel will wait indefinitely for the controller to complete a requested activity before continuing to the next step of descriptor processing.

The channel can generate two types of done notification signals when it completes operation on a descriptor—an interrupt and/or a writeback of the descriptor header. The done interrupt is enabled by the CDIE bit and the done writeback is enabled by the CDWE bit of the channel configuration register (Table 14-31). Table 14-5 shows the DONE field that is written back in if writeback is enabled. In addition, status writeback can also be used to signal processing completion. Any descriptor can have status writeback occur as a result of the AWSE bit of the channel configuration register. Or, by setting IWSE, status writeback will occur when any ICV-checking descriptor completes.

The selected done notification can be performed at the end of processing of every descriptor, or only on selected descriptors. If the NT field is 0 in the channel configuration register, then done notification is performed after every descriptor. If the NT field is 1, done notification is only performed on descriptors in which the DN bit is set in the packet header (Table 14-4).



## 14.5.1 Channel Registers

### 14.5.1.1 Crypto-Channel Configuration Register (CCCR)

The crypto-channel configuration register (CCCR) contains five operational bits permitting configuration of the channel as shown in Figure 14-36. Table 14-31 describes the CCR.




Figure 14-36. Crypto-Channel Configuration Register (CCCR)

Table 14-31. CCCR Field Descriptions

Bits	Names	Description
0–29	—	Reserved, set to zero
30	CON	Continue bit 0 No special action. 1 Causes the same channel reset actions as bit R, except that the fetch FIFO and the lower half of the CCR register are not cleared. After the reset sequence is complete, this bit automatically returns to 0 and the channel resumes normal operation, servicing the next descriptor pointer in the fetch FIFO, if any.
31	R	Reset channel 0 No special action. 1 Causes a software reset of the channel, clearing all its internal state. The details of the software reset actions depend upon what the channel is doing when the bit is set: <ul style="list-style-type: none"> <li>If the R bit is set while the channel is requesting an EU assignment from the controller, the channel cancels its request by asserting the release output signals. The channel then resets all its registers, clears the R bit, and return the channel state machine to the idle state.</li> <li>If the R bit is set after the channel has been assigned an EU, the channel requests a write from the controller to set the software reset bit of the EU. If a secondary EU has been reserved, the channel requests a write to reset that EU as well. The channel next asserts the appropriate release signal to notify the controller that the channel has finished with the reserved EU(s). The channel then resets all the registers, clears the RESET bit and returns the channel state machine to the idle state.</li> </ul>
32–54	—	Reserved, set to zero
55	BS	Burst size—The SEC accesses long text-data parcels in main memory through bursts of programmable size: 0 Burst size is 64 bytes 1 Burst size is 128 bytes

Table 14-31. CCCR Field Descriptions (continued)

Bits	Names	Description
56	IWSE	ICV writeback status enable 0 No special action. 1 If the descriptor calls for ICV comparison, then at the completion of descriptor processing, write back the status of all EUs into the header dword.
57	AWSE	Always writeback status enable 0 No special action. 1 At the completion of processing each descriptor, write back the status of all EUs into the header dword. In this case, IWSE has no effect.
59	CDWE	Channel done writeback enable 0 Channel done writeback disabled. 1 Channel done writeback enabled. Upon completion of descriptor processing, if the NT bit is set for Global, or if the DN (Done Notification) bit is set in the header word of the descriptor, notify the host by writing back the descriptor header with the writeback information shown in <a href="#">Table 14-5</a> . This enables the host to poll the memory location of the original descriptor header to determine if that descriptor has been completed.
60	—	Reserved, set to zero
61	NT	Notification type. This bit controls when the channel will generate channel done notification. Channel done notification can take the form of an interrupt or modified header writeback or both, depending on the state of the CDIE and CDWE control bits. 0 Global notification. The channel will generate channel done notification (if enabled) at the end of each descriptor. 1 Selected notification. The channel will generate channel done notification (if enabled) at the end of every descriptor with the DONE bit set in the descriptor header.
62	CDIE	Channel done interrupt enable 0 Channel done interrupt disabled 1 Channel Done Interrupt enabled. Upon completion of descriptor processing, if the NT bit is set for Global, or if the DN (Done Notification) bit is set in the header word of the descriptor, then notify the host by asserting an interrupt. Refer to <a href="#">Section 14.5.2, “Channel Interrupts,”</a> for complete description of channel interrupt operation.
63	—	Reserved, set to zero

### 14.5.1.2 Crypto-Channel Pointer Status Register (CCPSR)

The crypto-channel pointer status register (CCPSR) contains status fields and counters that provide the user with status information regarding the channel's actual processing of a given descriptor.

	0-2	3-7	8-11	12-15	16-19	20-23	24-31						
Field	—	FF_COUNTER	—	G_STATE	—	S_STATE	CHN_STATE						
Reset	0x0_0000												
R/W	R												
Addr	Channel_1 0x01110												
	32-37	38	39	40	41	42	43	44	45	46	47	48-59	60-63
Field	—	MI	MO	PR	SR	PG	SG	PRD	SRD	PD	SD	Error	PAIR_PTR
Reset	0x0_0007												
R/W	R												
Addr	Channel_1 0x01114												

**Figure 14-37. Crypto-Channel Pointer Status Register (CCPSR)**

Table 14-32 describes the CCPSR fields.

**Table 14-32. CCPSR Field Descriptions**

Bits	Names	Description
0-2	—	Reserved
3-7	FF_COUNTER	Fetch FIFO counter. This 5-bit counter indicates how many fetch pointers are currently stored in the FIFO.
8-11	—	Reserved
12-15	G_STATE	Gather state machine state. This field reflects the state of the channel Gather control state machine. The value of this field indicates which stage the channel is while performing gather function. <a href="#">Table 14-33</a> shows the meaning of all possible values of the G_STATE field. G_State is documented for information only. The User will not typically care about the gather state machine.
16-19	—	Reserved.
20-23	S_STATE	Scatter state machine state. This field reflects the state of the channel Scatter control state machine. The value of this field indicates which stage the channel is while performing scatter function. <a href="#">Table 14-33</a> shows the meaning of all possible values of the S_STATE field. S_State is documented for information only. The User will not typically care about the scatter state machine.
24-31	CHN_STATE	State. State of the channel state machine. This field reflects the state of the channel control state machine. The value of this field indicates exactly which stage the channel is in the sequence of fetching and processing data descriptors. <a href="#">Table 14-34</a> shows the meaning of all possible values of the STATE field. <b>Note:</b> CHN_State is documented for information only. The User will not typically care about the channel state machine.
32-37	—	Reserved, set to zero

Table 14-32. CCPSR Field Descriptions (continued)

Bits	Names	Description
38	MI	Multi_EU_IN. The Multi_EU_IN bit reflects the type of snooping the channel will perform, as programmed by the “Snoop Type” bit in the descriptor header. 0 Data input snooping by secondary EU disabled. 1 Data input snooping by secondary EU enabled.
39	MO	Multi_EU_OUT. The Multi_EU_OUT bit reflects the type of snooping the channel will perform, as programmed by the “Snoop Type” bit in the descriptor header. 0 Data output snooping by secondary EU disabled. 1 Data output snooping by secondary EU enabled.
40	PR	PRI_REQ. Request primary EU assignment. 0 Primary EU Assignment Request is inactive. 1 The channel is requesting assignment of primary EU to the channel. The channel will assert the EU request signal indicated by the op0 field in the Descriptor Header register as long as this bit remains set. The PRI_REQ bit is set when descriptor processing is initiated by the channel and the Op_0 field in the descriptor header contains a valid EU identifier. This bit is cleared when the request is granted, which will be reflected in the status register by the setting the PRI_GRANT bit.
41	SR	SEC_REQ. Request secondary EU assignment. 0 Secondary EU Assignment Request is inactive. 1 The channel is requesting assignment of secondary EU to the channel. The channel will assert the EU request signal indicated by the Op_1 field in the descriptor header register as long as this bit remains set. The SEC_REQ bit is set when descriptor processing is initiated by the channel and the Op_1 field in the descriptor header contains a valid EU identifier. This bit is cleared when the request is granted, which will be reflected in the status register by the setting the SEC_GRANT bit.
42	PG	Primary EU granted. The PRI_GRANT bit reflects the state of the EU grant signal for the requested primary EU from the controller. 0 The primary EU grant signal is inactive. 1 The EU grant signal is active indicating the controller has assigned the requested primary EU to the channel.
43	SG	Secondary EU granted. The SEC_GRANT bit reflects the state of the EU grant signal for the requested secondary EU from the controller. 0 The secondary EU grant signal is inactive. 1 The EU grant signal is active indicating the controller has assigned the requested secondary EU to the channel.
44	PRD	Primary EU reset done. The PRI_RST_DONE bit reflects the state of the reset done signal from the assigned primary EU. 0 The assigned primary EU reset done signal is inactive. 1 The assigned primary EU reset done signal is active indicating its reset sequence has completed and it is ready to accept data.
45	SRD	Secondary EU reset done. The SEC_RST_DONE bit reflects the state of the reset done signal from the assigned secondary EU. 0 The assigned secondary EU reset done signal is inactive. 1 The assigned secondary EU reset done signal is active indicating its reset sequence has completed and it is ready to accept data.

**Table 14-32. CCPSR Field Descriptions (continued)**

Bits	Names	Description
46	PD	Primary EU done. The PRI_DONE bit reflects the state of the done interrupt from the assigned primary EU. 0 The assigned primary EU done interrupt is inactive. 1 The assigned primary EU done interrupt is active indicating the EU has completed processing and is ready to provide output data.
47	SD	Secondary EU done. The SEC_DONE bit reflects the state of the done interrupt from the assigned secondary EU. 0 The assigned secondary EU done interrupt is inactive. 1 The assigned secondary EU done interrupt is active indicating the EU has completed processing and is ready to provide output data.
48–59	Error	Channel error status. This field reflects the error status of the channel. When a channel error interrupt is generated, this field will reflect the source of the error. The bits in the ERROR field are registered at specific stages in the descriptor processing flow. Once registered, an error can only be cleared only by resetting the channel or writing the appropriate registers to initiate the processing of a new descriptor. <a href="#">Table 14-35</a> lists the conditions which can cause a channel error and how they are represented in the ERROR field.
60–63	PAIR_PTR	Descriptor buffer register length/pointer pair. This field indicates which of the length/pointer pairs are currently being processed by the channel. <a href="#">Table 14-36</a> shows the meaning of all possible values of the PAIR_PTR field.

[Table 14-33](#) shows the values for G\_STATE and S\_STATE fields, which are the states for the gather and scatter state machines).

**Table 14-33. G\_STATE and S\_STATE Field Values**

Value	Gather State Machine:
0x0	GS_IDLE
0x1	GS_LOAD_POINTER
0x2	GS_LOAD_POINTER_DONE
0x3	GS_LOAD_NEXT_POINTER
0x4	GS_PROCESS_POINTER
0x5	GS_TRANS_BLOCK
0x6	GS_TRANS_BLOCK_DONE
0x7	GS_TRANS_BYTES
0x8	GS_TRANS_BYTES_DONE
0x9	GS_INC_PAIR_PTR
0xA	GS_UPDATE
0xB	GS_DONE
0xC	GS_ERROR
0xD	GS_RELOAD
0xE	GS_TRANS_INBOUND
0xF	GS_TRANS_INBOUND_DONE

Table 14-34 shows the values of channel states.

**Table 14-34. CHN\_STATE Field Values**

Value	Channel State	Value	Channel State
0x00	IDLE	0x22	EVALUATE_RESET
0x01	PROCESS_HEADER	0x23	RESET_WRITE_RESET_PRI
0x02	FETCH_DESCRIPTOR	0x24	RESET_RELEASE_PRI_CHA
0x03	CHANNEL_DONE	0x25	RESET_WRITE_RESET_SEC
0x04	CHANNEL_DONE_IRQ	0x26	RESET_RELEASE_SEC_CHA
0x05	CHANNEL_DONE_WRITEBACK	0x27	RESET_CHANNEL
0x06	CHANNEL_DONE_NOTIFICATION	0x28	WRITE_DATASIZE_PRI_POST
0x07	CHANNEL_ERROR	0x29	RESET_RELEASE_ALL
0x08	REQUEST_PRI_CHA	0x2A	RESET_RELEASE_ALL_DELAY
0x09	INC_DATA_PAIR_POINTER	0x2B	REQUEST_SEC_CHA
0x0A	DELAY_DATA_PAIR_UPDATE	0x2C	WRITE_DATASIZE_SEC
0x0B	EVALUATE_DATA_PAIRS	0x2D	WRITE_ICV_SIZE
0x0C	WRITE_RESET_PRI	0x2E	WRITE_SEC_CHA_GO_SNOOPOUT
0x0D	RELEASE_PRI_CHA	0x2F	WRITE_PRI_CHA_GO_SNOOPIN
0x0E	WRITE_RESET_SEC	0x30	WRITE_SEC_CHA_GO_SNOOPIN
0x0F	RELEASE_SEC_CHA	0x31	DELAY_1CYCLE
0x10	PROCESS_DATA_PAIRS	0x33	TRANS_EXTENT_READ
0x11	WRITE_MODE_PRI	0x34	TRANS_EXTENT3
0x12	WRITE_MODE_SEC	0x35	TRANS_EXTENT4
0x13	WRITE_DATASIZE_PRI	0x36	XOR_WRITE_READ_REG
0x14	DELAY_RNGA_DONE	0x37	DELAY_SEC_DONE_TLS
0x15	WRITE_DATASIZE_SEC_SNOOPIN	0x38	MAC_TO_CIPHER
0x16	TRANS_REQUEST_WRITE_SNOOPIN	0x39	MAC_TO_CIPHER_DONE
0x17	DELAY_PRI_SEC_DONE	0x3A	READ_PRI_STATUS
0x18	TRANS_REQUEST_WRITE	0x3C	READ_SEC_STATUS
0x19	WRITE_KEY_SIZE	Others	Reserved
0x1B	DELAY_PRI_DONE		
0x1E	WRITE_DATASIZE_SEC_SNOOPOUT		
0x1F	TRANS_REQUEST_READ_SNOOPOUT		
0x20	DELAY_SEC_DONE		
0x21	TRANS_REQUEST_READ		

Table 14-35 shows the bit positions of each potential error. Multiple errors are possible.

**Table 14-35. Crypto-Channel Pointer Status Register Error Field Definitions**

Value	Error
48	DOF. Double fetch FIFO write overflow error. This bit is set when the channel Fetch FIFO is full, SOF is set, and another write has been made to the fetch FIFO. When this bit is set the channel will stop, and an error interrupt will be activated. The channel will not start again until a Continue or Reset is given through the CCR register. This bit can be cleared by writing '1' to this bit in the CPSR register.
49	SOF. Single fetch FIFO write overflow error. This bit is set when the channel Fetch FIFO is full and another write has been made to the Fetch FIFO. The channel will set this bit and activate an error interrupt. The channel continues processing, but the descriptor pointer is lost. The host must clear this bit by writing '1' to this bit in the CPSR register.
50	MDTE. A Master data transfer error was received from the master bus interface. When the SEC, while acting as a bus master, detects an error, the controller passes the error to the channel in use. The channel halts and activates an interrupt. The channel can only be restarted by writing a '1' to the Continue or Reset bit in the channel configuration register, or resetting the whole SEC.
51	Scatter/Gather data length zero error. A zero length Scatter/Gather data pointer was detected.
52	Fetch pointer zero error. An all zero fetch pointer was detected.
53	Illegal descriptor header. Possible causes of an illegal descriptor header are: <ul style="list-style-type: none"> <li>• Invalid primary EU indicated by op0 field in descriptor header.<sup>1</sup></li> <li>• Invalid secondary EU indicated by op1 field in descriptor header.</li> <li>• Descriptor type field in descriptor header indicates secondary EU transaction when not in snoop mode</li> </ul>
54	Invalid EU assignment request. Indicates the channel has been assigned one or more EUs not requested by the descriptor header.
55	EU error detected. An EU assigned to this channel has generated an error interrupt. This error may also be reflected in the controller's interrupt status register.
56	Gather boundary error. Indicates a gather pointer straddles both a primary and secondary EU's data transfer.
57	Gather return/length error. Indicates the total data size covered by a gather link table did not match the total data size from the main descriptor.
58	Scatter boundary error. Indicates a scatter pointer straddles both a primary and secondary EU's data transfer.
59	Scatter return/length error. Indicates the total data size covered by a scatter link table did not match the total data size from the main descriptor.

<sup>1</sup> Invalid opcode includes any opcode valid on other versions of the SEC but not valid on SEC 2.2.

### NOTE

EU error bit (bit 55) can only be cleared by first clearing the error source in the assigned EU which caused it to be set.

Table 14-36 shows the possible values of the PAIR\_PTR field in the CCPSR.

**Table 14-36. Crypto-Channel Pointer Status Register PAIR\_PTR Field Values**

Value	Error
0x00	Processing header or pointer dword 0
0x01	Processing pointer dword 1
0x02	Processing pointer dword 2

**Table 14-36. Crypto-Channel Pointer Status Register PAIR\_PTR Field Values (continued)**

Value	Error
0x03	Processing pointer dword 3
0x04	Processing pointer dword 4
0x05	Processing pointer dword 5
0x06	Processing pointer dword 6
0x07	Complete (or not yet begun) processing of header dword and pointer dwords
0x08–FF	Reserved

### 14.5.1.3 Crypto-Channel Current Descriptor Pointer Register (CDPR)

The crypto-channel current descriptor pointer register (CDPR), shown in [Figure 14-38](#), contains the address of the descriptor which the channel is currently processing.

Field	0 — 31 32 63
Reset	0x0000_0000
R/W	R
Addr	Channel_1 0x3_1140

**Figure 14-38. Crypto-Channel Current Descriptor Pointer Register (CDPR)**

The bits in the CDPR perform the functions described in [Table 14-37](#).

**Table 14-37. CDPR Field Descriptions**

Bits	Names	Description
0–31	—	Reserved, set to zero.
32–63	CUR_DES_PTR_ADRS	Current descriptor pointer address. Pointer to system memory location of the current descriptor. This field reflects the starting location in system memory of the descriptor currently loaded into the DB. This value is updated whenever the channel requests a fetch of a descriptor from the controller. The value from the fetch FIFO is transferred to the current descriptor pointer register immediately after the fetch is completed. This address will be used as the destination for writeback of the modified header dword, if header writeback notification is enabled.

### 14.5.1.4 Fetch FIFO (FF)

The channel contains a fetch FIFO to store a queue of pointers to descriptors that the channel will process.

The fetch FIFO, displayed in [Figure 14-39](#), contains the addresses of the first byte of descriptors to be processed. In typical operation, the host CPU will create a descriptor in memory containing all relevant mode and location information for the SEC, then ‘launch’ the SEC by writing the address of the descriptor to the fetch FIFO.



The fetch FIFO can hold up to 24 descriptor pointers at a time. When the end of the current descriptor is reached, the descriptor pointed to by the next location in the fetch FIFO will be read to launch the next descriptor.

The Fetch Address is written into the FIFO only if the write includes the least significant byte (bits 56–63). If Extended Address mode is used, the Extended Fetch Address must be written before or concurrent with the Fetch Address.

Specifying a FETCH\_ADRS of 0 causes the channel to generate an error and stop.

	0	31	32	63
Field	—			FETCH_ADRS
Reset	0x0000_0000			
R/W	W			
Addr	Channel_1 0x3_01148			

**Figure 14-39. Fetch FIFO Register (FF)**

Table 14-38 describes the fetch FIFO fields.

**Table 14-38. Fetch FIFO Field Descriptions**

Bits	Names	Description
0–31	—	Reserved, set to zero
32–63	FETCH ADRS	Fetch address. Pointer to system memory location of a descriptor the host wants the SEC to fetch.

### 14.5.1.5 Descriptor Buffer (DB)

The descriptor buffer (DB) consists of 8 dword registers (DB0–DB7), and contains the current descriptor being processed by the channel. These registers are read-only, since the descriptor is always fetched from system memory.

For more information about the fields in a descriptor, see [Section 14.3.1, “Descriptor Structure.”](#)

	0	15	16	17	23	24	27	28	31	32	63	
DB0	Header								Reserved			
DB1	Length0	J1	Extent0	—	—	—	—	—	—	—	Pointer0	
DB2	Length1	J2	Extent1	—	—	—	—	—	—	—	Pointer1	
DB3	Length2	J3	Extent2	—	—	—	—	—	—	—	Pointer2	
DB4	Length3	J4	Extent3	—	—	—	—	—	—	—	Pointer3	
DB5	Length4	J5	Extent4	—	—	—	—	—	—	—	Pointer4	
DB6	Length5	J6	Extent5	—	—	—	—	—	—	—	Pointer5	
DB7	Length6	J7	Extent6	—	—	—	—	—	—	—	Pointer6	
Address	Channel_1 0x3_1180–0x3_11BF											

**Figure 14-40. Descriptor Buffer (DB)**

## 14.5.2 Channel Interrupts

The channel can assert both DONE and ERROR interrupts to the controller. When the interrupt generation conditions have been met, the channel will assert the appropriate interrupt. The status of the registered channel interrupts is available in the controller interrupt status register. The channel does not have an internal interrupt mask, but the controller can be programmed to block channel interrupts through its interrupt mask register (see [Section 14.6.4.2, “Interrupt Mask Register \(IMR\)”](#)).

### 14.5.2.1 Channel Done Interrupt

Whether and when a channel DONE interrupt is generated depends on the setting of the crypto-channel configuration register NT and CDIE bits in the CCCR (see [Figure 14-36](#)). Assuming the CDIE (Channel Done Interrupt Enable) is set, the channel will generate an interrupt event after every successfully completed descriptor (Notification Type set to Global), or after each successfully completed descriptor with the DN (Done Notification) bit set in the header word of the descriptor.

Even if multiple Channel Done interrupt events are generated by the channel before the first can be cleared by the host, the interrupt events are not lost. The controller queues channel done interrupts from the channel (see [Section 14.6.3, “Controller Interrupts”](#)).

### 14.5.2.2 Channel Error Interrupt

The channel error interrupt is generated when an error condition occurs during descriptor processing. The channel error interrupt will be asserted as soon as the error condition is detected. The type of error condition is reflected the ERROR field of the channel pointer status register (CPSR). Refer to [Table 14-35](#) for a complete listing of error types.

### 14.5.2.3 Channel Reset

Channel reset is asserted when the host sets the RESET bit in the channel configuration register (CCR). The effect of software reset on the channel varies according to what the channel is doing when the bit is set:

- If the RESET bit is set while the channel is requesting an EU assignment from the controller, the channel will cancel its request by asserting the release output signals. The channel will then reset all the registers, clear the RESET bit and return the control state machine to the idle state.
- If the RESET bit is set after the channel has been dynamically assigned an EU, the channel will request a write from the controller to set the software reset bit of the EU. A write to reset the secondary (MDEU) EU will also be requested if one has been reserved for snooping. The channel will then assert the appropriate release output signal to notify the controller that the channel has finished with the reserved EU(s). The channel will then reset all the registers, clear the RESET bit and return the control state machine to the idle state.

## 14.6 Controller

The controller within the SEC is responsible for overseeing the operations of the execution units (EUs), the interface to the host processor, and the management of the channels. The controller interfaces to the host through the master/slave bus interface and to the channels and EUs through internal buses. All

transfers between the host and the EUs are moderated by the controller. Some of the main functions of the controller are as follows:

- Provide arbitration for bus access and control bus accesses
- Control the internal bus accesses to the EUs
- Assign EUs to the channel
- Monitor interrupts from the channel and EUs and pass to host
- Realign read and write data to the proper byte alignment

### 14.6.1 Assignment of EUs to Channel

Assignment of a EU to the channel is done dynamically.<sup>1</sup> The channel requests an EU, the controller checks to see if the requested EU is available, and if it is, the controller grants the channel assignment of the EU.

When requested, the controller will assert the grant signal pertaining to the request from the channel. The grant signal will remain asserted until the channel is done and releases the EU.

In some cases, the channel may request two EUs. The channel will do this by first requesting the primary EU, then requesting the secondary EU. Once the controller has granted both EUs, the channel is then capable of requesting that the secondary EU snoop the bus. Snooping status is indicated in the MI and MO bits of [Table 14-32](#).

In all cases, the controller assigns the primary EU to the requesting channel as the EUs become available. The controller does not wait until both EUs are available before issuing grants to the channel which is requesting two EUs.

### 14.6.2 Bus Interface

The controller in the SEC (refer to) has the ability to be a bus master or a slave. This means that the controller can issue read and write commands to the bus, and it can also be written to and read from by the host.

The controller is the sole bus master in the SEC. All other modules are slave-only devices. A channel may request access to system resources including the bus. In these cases, the channel must provide the starting address of the transfer for the bus(es) requested. All subsequent addresses are generated by the controller. All addresses will be sequential.

#### 14.6.2.1 Arbitration for Use of the Controller and Buses

The controller attempts to maximize utilization of the system bus by grouping outstanding bus requests from the channel by request type (read or write). The controller will perform all write requests to the system bus, followed by all read requests, then repeat.

<sup>1</sup> The security engine in the MPC8313E is a single-channel implementation, but other variants of the SEC have been implemented with 1, 2, and 4 channels. Although not necessary for a single-channel SEC, dynamic assignment of EUs to the channel is maintained to improve software compatibility with other SEC-enhanced processors.

The SEC does not dynamically adjust its own transaction priorities. System software, however, can adjust SEC transaction priority in realtime, with the change in priority taking effect immediately.

### 14.6.2.2 Master Read

Here is more detail on the sequence of events for an system bus read with the controller as master:

1. Channel asserts its bus read request to the controller.
2. Channel furnishes external read address, internal write address, and transfer length.
3. Controller sends request acknowledge to channel.
4. Controller asserts request to the system bus through the Magenta master interface.
5. Controller waits for system bus read to begin.
6. When bus read begins, controller receives data from the master interface and performs a write to the appropriate internal address supplied by the channel. Data may be realigned byte-wise by the controller if either:
  - The read did not begin on a 32-bit word boundary, or
  - The previous write to an execution unit's input FIFO did not end on a 32-bit word boundary.
7. Transfer continues until the bus read is completed and the controller has written all data to the appropriate internal address. The master interface will continue making bus requests until the full data length has been read.

When the SEC performs a transaction as master, it is possible for the intended slave to terminate the transfer due to an error. SEC transaction requests are posted to the MPC8313E target queue, after which the MPC8313E takes responsibility for completing the transaction or signaling error. An error in an SEC-initiated transaction will also be reported by the SEC through the channel interrupt status register. The host will be able to determine which channel generated the interrupt by checking the ISR for the channel ERROR bit.

### 14.6.2.3 Master Write

Here is more detail on the sequence of events for an system bus write with the controller as master:

1. Channel asserts its bus write request to the controller.
2. Channel furnishes internal read address, external write address, and transfer length.
3. Controller sends request acknowledge to channel.
4. Controller performs a read from the appropriate internal address supplied by the channel, loads the write data into its FIFO, asserts a request to the system bus through the Magenta master interface, and waits for the system bus to become available.
5. When the system bus becomes available, controller writes data from its FIFO to the master interface.

## 14.6.3 Controller Interrupts

All interrupt outputs from other SEC blocks are fed to the controller as interrupt conditions. In addition, the controller itself detects some interrupt conditions. The controller maintains an interrupt status register

(ISR) with bits corresponding to all of these possible interrupt conditions. If an interrupt condition occurs and the corresponding bit of the interrupt mask register (IMR) is set, the associated ISR bit is set, indicating the presence of a pending interrupt. Whenever any bits are set in the interrupt status register, the controller asserts its interrupt output line to the host.

To handle an interrupt, the host must read the interrupt status register to determine the source. It may then need to do further reads of interrupt status registers of other blocks to get more detailed information about the cause. In some cases, the host may need to take action to clear the root cause of the interrupt. After that, the host can clear the desired bit of the interrupt status register by writing a 1 to the corresponding bit of the interrupt clear register (ICR). If the cause of the interrupt condition has not been cleared, or if there is some other interrupt condition from the same source, then the interrupt status register bit will clear for a cycle and go high again, and the interrupt output line to the host remains high. If the ISR bit is successfully cleared and no other interrupt conditions are present, the controller de-asserts its interrupt output. If any interrupts are still pending in the interrupt status register, the interrupt output remains asserted.

Note that EU interrupt conditions may be blocked at two different levels. There is an interrupt control register in each EU which can block particular interrupt conditions before they reach the EU's interrupt status register, and in addition, bits of the controller's interrupt mask register (IMR) must be set to allow interrupt conditions to reach the interrupt status register. Interrupt conditions from the channel and controller can only be blocked through the IMR.

For typical operation it is suggested that the IMR be programmed as follows: unmask channel interrupts while masking EU interrupts. Errors or Done signals coming from the EUs eventually cause the channel to signal an Error or Done interrupt.

The channel can generate frequent interrupts, especially if it is configured to interrupt at the completion of each descriptor. To make sure that the host receives the right number of interrupts, the channel Done interrupt has a special 'queueing' feature. If multiple Channel Done interrupts are generated before the first is cleared, then the additional interrupts are queued by the controller. When the host clears channel interrupt, if there are no other interrupts queued from that channel, then the channel Done interrupt is de-asserted. If other interrupts remain in the queue, the controller will de-assert the interrupt for one cycle and then re-assert it again.

## 14.6.4 Controller Registers

The controller registers are described in detail in the following sections.

### 14.6.4.1 EU Assignment Status Register (EUASR)

The EU assignment status register (EUASR), displayed in [Figure 14-41](#), is used to check the assignment status of a EU to the channel.

A 1-bit field indicates to the channel whether or not the EU is assigned.

	0	3	4	7	8	11	12-14	15	16	19	20-22	23	24	27	28-30	31				
Field	—			—			—			MDEU	—			AESU	—			DEU		
Reset	0xF			0x0			0xF			0x0			0	0xF			0x0			0
R/W	R																			
Addr	0x3_1028																			
	32	35	36	39	40	43	44	47	48	51	52	55	56	59	60	63				
Field	—								—		—		—		—					
Reset	0xFFFF0								0xF		0x0		0xF		0x0					
R/W	R																			
Addr	0x3_102C																			

Figure 14-41. EU Assignment Status Register (EUASR)

### 14.6.4.2 Interrupt Mask Register (IMR)

The SEC controller generates the single interrupt output from all possible interrupt sources. These sources can be individually enabled by the interrupt mask register (IMR). If unmasked, the interrupt source value, when active, is captured into the interrupt status register (ISR). Figure 14-42 shows the bit positions of each potential interrupt source. Each interrupt source is individually unmasked by setting its corresponding bit. At reset, all bits are disabled. The bit fields are described in Table 14-39.

For normal operation, the IMR should be programmed as follows: Unmask the channel interrupts while masking EU interrupts. The channels will generate the appropriate interrupts to the host.

	0													14	15
Field	—													ITO	
Subfield															
Reset	0x0000														
R/W	R/W														
Addr	0x3_1008														
	16	19	20	21	22	23	24					29	30	31	
Field	—			DONE Overflow				—				CHN_1			
Subfield							CH1					Err	Dn		
Reset	0x0000														
R/W	R/W														
Addr	0x3_1008														
	32													47	
Field	—														
Subfield															
Reset	0x0000														
R/W	R/W														
Addr	0x3_100C														
	48				53	54	55	56	57	58	59	60	62	63	63
Field	—				MDEU		—		AESU		—		DEU		
Subfield					Err	Dn			Err	Dn			Err	Dn	
Reset	0x0000														
R/W	R/W														
Addr	0x3_100C														

Figure 14-42. Interrupt Mask Register (IMR)

Table 14-39 describes the register field names in the interrupt mask register (IMR), interrupt status register (ISR), and interrupt clear register (ICR).

**Table 14-39. Field Names in Interrupt Mask, Interrupt Status, and Interrupt Clear Registers**

Bits	Names	Description
15	ITO	Internal time out 0 No internal time out 1 An internal time out was detected <b>Note:</b> Internal time out is an indication that the channel or EU has failed to respond to a slave read or write within 16 cycles, which would only occur in an impending hang condition. Assertion of this interrupt indicates the SEC controller has completed the transaction to avoid a hang, however the 'completed' transaction does not result in a successful read or write, and the interrupt advises the system that the slave transaction was unsuccessful.
20–23	Done Overflow	Done overflow (one bit for each channel) 0 No done overflow 1 Done overflow error. Indicates that more than 15 done interrupts were queued from the channel without an interrupt clear from the host.
30–31	Err and Dn bits for channel	Err 0 No error detected 1 Error detected. Indicates that channel status register must be read to determine exact cause of the error. Dn 0 Not DONE 1 DONE bit indicates that the interrupting channel has completed its operation.
Multiple	Err and Dn bits for execution units (AESU, and so on.)	Err 0 No error detected. 1 Error detected. Indicates that execution unit status register must be read to determine exact cause of the error. Dn 0 Not DONE. 1 DONE bit indicates that the interrupting EU has completed its operation.
0–14, 16–19, 24–29, 32–53, 56–57, 60–61	—	Reserved, set to zero.

### 14.6.4.3 Interrupt Status Register (ISR)

The interrupt status register (ISR) contains fields representing all possible sources of interrupts. The interrupt status register is cleared either by a reset, or by writing the appropriate bits active in the interrupt clear register (ICR). Figure 14-43 shows the bit positions of each potential interrupt source. The bit fields are described in Table 14-39.



Field	—														ITO
Subfield	—														
Reset	0x0000														
R/W	R														
Addr	0x 3_1010														
Field	—			DONE Overflow				—			CHN_1				
Subfield	—						CH1	—			Err	Dn			
Reset	0x0000														
R/W	R														
Addr	0x 3_1010														
Field	—														
Subfield	—														
Reset	0x0000														
R/W	R														
Addr	0x 3_1014														
Field	—				MDEU		—		AESU		—		DEU		
Subfield	—				Err	Dn	—		Err	Dn	—		Err	Dn	
Reset	0x0000														
R/W	R														
Addr	0x 3_1014														

Figure 14-43. Interrupt Status Register (ISR)

#### 14.6.4.4 Interrupt Clear Register (ICR)

The interrupt control register (ICR) provides a means of clearing the ISR. When a bit in the ICR is written with a 1, the corresponding bit in the ISR is cleared, clearing the interrupt output pin  $\overline{IRQ}$  (assuming the cleared bit in the ISR is the only interrupt source). If the input source to the ISR is a steady-state signal that remains active, the appropriate ISR bit, and subsequently  $\overline{IRQ}$ , will be reasserted shortly thereafter.

Figure 14-44 shows the bit positions of each interrupt source that can be cleared by this register. The complete bit definitions for the ICR can be found in Figure 14-44. The bit fields are described in Table 14-39.

When an ICR bit is written, it will automatically clear itself one cycle later. That is, it is not necessary to write a '0' to a bit position which has been written with a '1'.

**NOTE**

Interrupts are registered and sent based upon the conditions which cause them. If the cause of an interrupt is not removed, the interrupt will return a few cycles after it has been cleared using the ICR.

	0													14	15	
Field	—														ITO	
Subfield																
Reset	0x0000															
R/W	W															
Addr	0x 3_1018															
	16	19	20	21	22	23	24					29	30	31		
Field	—			DONE Overflow				—				CHN_1				
Subfield							CH1					Err	Dn			
Reset	0x0000															
R/W	W															
Addr	0x 3_1018															
	32															47
Field	—															
Subfield																
Reset	0x0000															
R/W	W															
Addr	0x 3_101C															
	48				53	54	55	56	57	58	59	60	62	62	63	
Field	—				MDEU		—		AESU		—		DEU			
Subfield					Err	Dn			Err	Dn			Err	Dn		
Reset	0x0000															
R/W	W															
Addr	0x 3_101C															

**Figure 14-44. Interrupt Clear Register (ICR)**

### 14.6.4.5 Identification Register (ID)

The read-only identification register (ID), displayed in [Figure 14-45](#), contains a 64-bit value that uniquely identifies the version of SEC 2.2. The value of ID is always 0x0000\_0000\_0002\_00A0, indicating that this is the first version of the SEC 2.2.

	0	56	57	58	59	60	61	62	63	
Field	—						VERSION			
Reset	0x0000_0000_0002_00A0									
R/W	R									
Addr	0x 3_1020									

**Figure 14-45. ID Register (ID)**

### 14.6.4.6 IP Block Revision Register

The read-only IP block revision register, displayed in [Figure 14-46](#), contains a 64-bit value that uniquely identifies the version of SEC 2.2. The value of this register is always 0x0000\_0000\_0002\_00A0, indicating that this is the first version of SEC2.2.

	0	56	57	58	59	60	61	62	63	
Field	—						VERSION			
Reset	0x0000_0000_0002_00A0									
R/W	R									
Addr	0x 3_1BF8									

**Figure 14-46. IP Block Revision Register**

### 14.6.4.7 Master Control Register (MCR)

The master control register (MCR), shown in [Figure 14-47](#), controls certain functions in the controller and provides a means for software to reset the SEC.

Field	0	21	22	23	24	29	30	31		
Field	—					PRIORITY	—		GIH	SWR
Reset	0x0000_0000									
R/W	R/W									
Addr	0x3_1030									
Field	32	39	40	47	48	55	56	63		
Field										
Reset	0x0000_0000									
R/W	R/W									
Addr	0x 3_1034									

**Figure 14-47. Master Control Register (MCR)**

[Table 14-40](#) describes the MCR fields.

**Table 14-40. MCR Field Descriptions**

Bits	Names	Description
0–21	—	Reserved
22–23	Priority	Priority on master bus. The setting of these bits determines the transaction priority level the SEC asserts to the device internal arbiter. The SEC does not dynamically alter its priority level based on system congestion or SEC utilization, however software may change the SEC priority level in realtime. 00 Lowest priority (default) 01 Next lowest priority 10 Next highest priority 11 Highest priority
24–29	—	Reserved
30	GIH	Global inhibit. Writing 1 to this bit prevents output IPM_SNOOP from being asserted. 0 Permit assertion of IPM_SNOOP, allowing the system cache to snoop bus transactions initiated by the SEC. 1 Prevent assertion of IPM_SNOOP, preventing the system cache from snooping
31	SWR	Software reset. Writing 1 to this bit will cause a global software reset. Upon completion of the reset, this bit will be automatically cleared. 0 Do not reset 1 Global reset

### 14.6.5 Snooping by Caches

SEC transactions can be snooped by the MPC8313E cache if defined as global. This definition is programmed in the master control register MCR[GI]. See [14.6.4.7, “Master Control Register \(MCR\),”](#) for more details. Note that SEC transactions are defined as global by default.

## 14.6.6 Interrupts

The SEC generates a single interrupt to the MPC8313E programmable interrupt controller. The user allows interrupts from the SEC to be reported to the CPU by setting the mask bit in SIMSR\_H[SEC].

The user can control which events cause an interrupt by configuring the SEC interrupt mask register (IMR). These events are:

- Done (of a channel or an execution unit)
- Error (of a channel or an execution unit)

When the user detects an interrupt request from the SEC, it should further read the SEC interrupt status register (ISR) to determine the source of that interrupt. To clear an interrupt, the user should write 1 to the bits in the SEC interrupt clear register (ICR) corresponding to the pending ISR bits.

Events may be further masked per channel by setting or clearing the related fields in the crypto-channel configuration registers. It is suggested that the user leave channel interrupts unmasked, while masking the interrupts from the EUs. Errors or Done signals coming from the EUs eventually cause the channel to signal an error or Done interrupt. Clearing an interrupt before eliminating the condition which caused the interrupt will cause the interrupt to be asserted again a few cycles later.

## 14.7 Power Saving Mode

The SEC can be disabled by clearing SCCR[ENCCM]. See [Section 4.5.2.3, “System Clock Control Register \(SCCR\),”](#) for more information.



# Chapter 15

## Enhanced Three-Speed Ethernet Controllers

### 15.1 Overview

The enhanced three-speed Ethernet controllers (eTSECs) of the device interface to 10 Mbps, 100 Mbps, and 1 Gbps Ethernet/IEEE 802.3™ networks. For Ethernet, an external PHY or SerDes device is required to complete the interface to the media. Each eTSEC supports multiple standard media-independent interfaces. Multiple eTSECs are available, providing flexible options for connectivity and control access at different speeds.

The eTSEC provides the flexibility to accelerate the identification and retrieval of standard and non-standard protocols carried over Ethernet, including both IP versions 4 and 6 and TCP/UDP. CPU-intensive parsing and checksum operations can be optionally off-loaded to an eTSEC to accelerate existing TCP/IP stacks. On transmission, varying fractions of link bandwidth can be allocated to each of multiple transmit queues through a modified weighted round-robin scheduler. On receive, an arbitrary set of queue selection rules can be programmed into each eTSEC to implement flexible quality of service or firewall strategies based on high-level protocol identification. Without enabling these advanced features, each eTSEC emulates a PowerQUICC II Pro TSEC, allowing existing driver software to be re-used with minimal change. Each eTSEC is organized as shown in [Figure 15-1](#).




Figure 15-1. eTSEC Block Diagram

## 15.2 Features

The eTSECs of the device include these distinctive features:

- IEEE 802.3, 802.3u, 802.3x, 802.3z, 802.3ac, 802.3ab compatible
- Support for different Ethernet physical interfaces:
  - 10/100 Mbps IEEE 802.3 MII and RMII
  - 10/100 Mbps RGMII
  - 1000 Mbps full-duplex RGMII and RTBI



- 10/100 Mbps SGMII
- 1000 Mbps full-duplex SGMII
- TCP/IP off-load
  - IP v4 and IP v6 header recognition on receive
  - IP v4 header checksum verification and generation
  - TCP and UDP checksum verification and generation
  - Per-packet configurable off-load
  - Recognition of VLAN, stacked-VLAN, 802.2, PPPoE session, MPLS stacks, ARP, and ESP/AH IP-Security headers
- Quality of service (QoS) support
  - Transmission from up to eight queues
    - Priority-based queue selection
    - Modified weighted round-robin queue selection with fair bandwidth allocation
  - Reception to up to eight physical queues
    - 64 virtual receive queues overlaid on 8 physical buffer descriptor rings
    - Table-oriented queue filing strategy based on 16 header fields or flags
    - Frame rejection support for filtering applications
    - Filing based on Ethernet, IP, and TCP/UDP properties, including VLAN fields, Ether-type, IP protocol type, IP TOS or differentiated services, IP source and destination addresses, TCP/UDP port numbers
- Interrupt coalescing
  - Packet-count-based thresholds for both receive and transmit
  - Timer-based thresholds
- Full- and half-duplex Ethernet support (1000 Mbps supports only full duplex):
  - IEEE 802.3 full-duplex flow control (automatic PAUSE frame generation or software programmed PAUSE frame generation and recognition)
  - Programmable maximum frame length supports jumbo frames (up to 9.6 Kbytes) and IEEE 802.1 virtual local area network (VLAN) tags and priority
  - VLAN insertion and deletion
    - Per-frame VLAN control word or default VLAN for each eTSEC
    - Extracted VLAN control word passed to software separately
    - Programmable VLAN tag to support metropolitan bridging
  - Retransmission following a collision
  - Support for CRC generation and verification of inbound/outbound packets
  - Programmable Ethernet preamble insertion and extraction of up to 7 bytes
- MAC address recognition:
  - Exact match on primary and virtual 48-bit unicast addresses
    - VRRP and HSRP support for seamless router fail-over

- In addition to primary station address, up to fifteen additional exact-match MAC addresses supported
  - Broadcast address (accept/reject)
  - Hash table match on up to 256 unicast/multicast or 512 multicast-only addresses
  - Promiscuous mode
- Remote network monitoring (RMON) statistics support
  - 32-bit byte counters
  - Carry/Overflow of counter interrupts
- Backward compatibility with MPC8349E (PowerQUICC II Pro) TSEC
  - PowerQUICC II Pro buffer descriptor (BD) format and rings supported
  - Common register memory map, with specific exceptions:
    - Out-of-sequence transmit BD not supported
    - Internal DMA BD pointers and data counts not visible
    - MINFLR register not supported
  - Reset state of eTSEC defaults to common PowerQUICC II Pro TSEC subset
  - TSEC\_ID register permits TSEC versus enhanced TSEC differentiation
- Hardware assist for 1588 compliant timestamping
  - Per packet timestamp tag for Receive
  - Programmable timestamp capture for Transmit
  - Recognition of PTP packet
  - Periodic Pulse Generation
  - Self-correcting precision timer with nano-second resolution
  - Phase aligned adjustable (divide by N) clock output
  - Two 64-bit alarm (future time) registers for future time comparison

### 15.3 Modes of Operation

The eTSEC's primary operational modes are the following:

- Full- and half-duplex operation
 

This is determined by the MACCFG2 register's full-duplex bit (MACCFG2[Full Duplex]). Full-duplex mode is intended for use on point-to-point links between switches or end node to switch. Half-duplex mode is used in connections between an end node and a repeater or between repeaters.

If configured in half-duplex mode (10- and 100-Mbps operation; MACCFG2[Full Duplex] is cleared), the MAC complies with the IEEE CSMA/CD access method.

If configured in full-duplex mode (10/100/1000 Mbps operation; MACCFG2[Full Duplex] is set), the MAC supports flow control. If flow control is enabled, it allows the MAC to receive or send PAUSE frames.

- 10- and 100-Mbps MII interface operation
 

The MAC–PHY interface operates in MII mode by setting MACCFG2[I/F Mode] = 01. The MII is the media-independent interface defined by the 802.3 standard for 10/100 Mbps operation. The speed of operation is determined by the TSEC<sub>n</sub>\_TX\_CLK and TSEC<sub>n</sub>\_RX\_CLK signals, which are driven by the transceiver. The transceiver either auto-negotiates the speed, or it may be controlled by software using the serial management interface (MDC/MDIO signals) to the transceiver.

Clause 22.2.4 of the IEEE 802.3 specification describes the MII management interface.
- 10- and 100-Mbps RMII interface operation
 

The RMII is the reduced media-independent interface defined by the RMII Consortium (March 1998) for 10/100 Mbps operation. The speed of operation is determined by the TSEC<sub>n</sub>\_TX\_CLK signal, which is driven by the transceiver.
- MAC address recognition options
 

The options supported are promiscuous, broadcast, exact unicast address match, exact unicast virtual address match to support router redundancy, and multicast hash match. For detailed descriptions refer to [Section 15.6.2.7, “Frame Recognition.”](#)

eTSEC supports automatic LAN-initiated wake-up during power management through the AMD Magic Packet™ protocol, as described in [Section 15.6.2.8, “Magic Packet Mode.”](#)
- Receive frame parsing options
 

Frame parsing options are to disable parsing (no TCP/IP off-load), IP header parsing, and TCP or UDP parsing. Parsing must be enabled to make use of receive queue filing algorithms. The options are detailed in [Section 15.6.3, “TCP/IP Off-Load.”](#)
- Receive queue selection options
 

Received frames are by default sent to a single buffer descriptor ring. If multiple receive queues are enabled, a receive queue filer can be programmed with selection criteria to differentiate received frames and file them to different buffer descriptor rings. See [Section 15.6.4, “Quality of Service \(QoS\) Provision,”](#) for detailed descriptions.
- TCP/IP transmit options
 

Frames for transmission may be sent as-is, with IP header processing, or TCP header processing. The transmit buffer descriptors, described in [Section 15.6.7.2, “Transmit Data Buffer Descriptors \(TxBD\),”](#) enable these options and operate with parameters prepended to frame buffers, as described in [Section 15.6.3, “TCP/IP Off-Load.”](#)
- Transmit queue selection options
 

The options supported are single transmit queue, priority-based queue selection, and modified weighted round-robin queueing. These options are described further in [Section 15.5.3.2.1, “Transmit Control Register \(TCTRL\).”](#)
- RMON support
 

Standard Ethernet interface management information base (MIBs) can be generated through the RMON MIB counters.

- Internal loop back supported for all interfaces except when configured for half-duplex operation. Internal loop back mode is selected through the loop back bit in the MACCFG1 register. See [Section 15.7.1, “Interface Mode Configuration,”](#) for details.

## 15.4 External Signals Description

This section defines the eTSEC interface signals. The buses are described using the bus convention used in IEEE 802.3 because the PHY follows this same convention. (That is, TxD[3:0] means 0 is the lsb.) Note that except for external physical interfaces the buses and registers follow a big-endian format, where 0 denotes the msb.

Each eTSEC network interface supports multiple options:

- The MII option requires 18 I/O signals (including the MDIO and MDC MII management interface) and supports both a data and a management interface to the PHY (transceiver) device. The MII option supports both 10- and 100-Mbps Ethernet rates.
- The RGMII, RTBI, and RMII options are reduced-pin implementations of the GMII, TBI, and MII interfaces, respectively.
- SGMII interfaces are offered via the SerDes interface signals.
- 1588 timer signals

[Table 15-1](#) lists the network interface signals.

**Table 15-1. eTSEC<sub>n</sub> Network Interface Signal Properties**

Signal Name	Function	Reset State
TSEC <sub>n</sub> _COL	MII—collision, input	—
TSEC <sub>n</sub> _CRS	MII—carrier sense, input	—
TSEC <sub>n</sub> _GTX_CLK	RTBI, RGMII—inverted transmit clock feedback, output MII, RMII—transmit clock feedback when transmission is enabled, zero otherwise, output	0
EC_GTX_CLK125	Oscillator source for RGMII, RTBI transmit clock, input, shared by all eTSECs	—
EC_MDC	Management clock, output.	0
EC_MDIO	Management data, bidirectional.	Hi-Z (input)
TSEC <sub>n</sub> _RX_CLK	MII, RGMII—receive clock, input	—
TSEC <sub>n</sub> _RX_DV	MII—receive data valid, input RGMII (RX_CLK rising)—receive data valid, input RGMII (RX_CLK falling)—receive error, input RTBI (RX_CLK rising)—receive code group (RCG) bit 4, input RTBI (RX_CLK falling)—receive code group (RCG) bit 9, input RMII—CRS_DV carrier sense/data valid, input	—
TSEC <sub>n</sub> _RXD[3:0]	MII—Receive data bits 3:0, input RGMII (RX_CLK rising) —Receive data bits 3:0, input RGMII (RX_CLK falling)—Receive data bits 7:4, input RTBI (RX_CLK rising)—RCG bits 3:0, input RTBI (RX_CLK falling)—RCG bits 8:5, input RMII—RXD[1:0] receive data bits, input RMII—RXD[3:2] are unused	—

Table 15-1. eTSEC<sub>n</sub> Network Interface Signal Properties (continued)

Signal Name	Function	Reset State
TSEC <sub>n</sub> _RX_ER	MII, RMII—Receive error, input RGMII, RTBI—Unused	—
TSEC <sub>n</sub> _TX_CLK	MII—transmit clock, input RMII—reference transmit and receive clock, input RGMII, RTBI—unused	—
TSEC <sub>n</sub> _TXD[3:0]	MII—Transmit data bits 3:0, output RGMII (TX_CLK rising)—Transmit data bits 3:0, output RGMII (TX_CLK falling)—Transmit data bits 7:4, output RTBI (TX_CLK rising)—TCG bits 3:0, output RTBI (TX_CLK falling)—TCG bits 8:5, output RMII—TXD[1:0] transmit data bits, output RMII—TXD[3:2] unused, output driven zero	0000
TSEC <sub>n</sub> _TX_ER	MII—transmit error, output RGMII, RTBI, RMII—unused, output driven zero	0
TSEC <sub>n</sub> _TX_EN	MII, RMII—Transmit data valid, output RGMII (TX_CLK rising)—Transmit data enabled, output RGMII (TX_CLK falling)—Transmit error, output RTBI (TX_CLK rising)—TCG bit 4, output RTBI (TX_CLK falling)—TCG bit 9, output	0
TSEC_1588_CLK	1588—Clock input External high precision timer reference clock input (chip external input pin).	—
TSEC_1588_GCLK	1588—Clock output Phase aligned timer clock divider output (chip external output pin).	0
TSEC_1588_TRIG1	1588—Trigger in 1 External timer trigger input 1. This is an asynchronous general purpose input (chip external input pin).	—
TSEC_1588_TRIG2	1588—Trigger in 2 External timer trigger input 2. This is an asynchronous general purpose input (chip external input pin).	—
TSEC_1588_PP1	1588—Pulse out 1 Timer pulse per period 1. It is phase aligned with 1588 timer clock (chip external output pin).	0
TSEC_1588_PP2	1588—Pulse out 2 Timer pulse per period 2. It is phase aligned with 1588 timer clock (chip external output pin).	0
TSEC_1588_PP3	1588—Pulse out 3 Timer pulse per period 3. It is phase aligned with 1588 timer clock (chip external output pin).	0
TSEC_1588_ALARM1	1588—Timer alarm 1 Timer current time is equal to or greater than alarm time comparator register. User reprograms the TSEC_1588_ALARM <sub>n</sub> _H/L register to deactivate this output (chip external output pin).	0
TSEC_1588_ALARM2	1588—Timer alarm 2 Timer current time is equal to or greater than alarm time comparator register. User reprograms the TSEC_1588_ALARM <sub>n</sub> _H/L register to deactivate this output (chip external output pin).	0
<u>SD_REF_CLK</u> , <u>SD_REF_CLK</u>	SerDes PLL reference clock (and complement)	—

Table 15-1. eTSEC<sub>n</sub> Network Interface Signal Properties (continued)

Signal Name	Function	Reset State
TXA/ $\overline{\text{TXA}}$	Serial transmitter, lane A, positive data (and negative data, complement)	—
RXA/ $\overline{\text{RXA}}$	Serial receiver, lane A, positive data (and negative data, complement)	—

### 15.4.1 Detailed Signal Descriptions

Below is a description of the eTSEC interface signals. For RGMII mode details please refer to the Hewlett-Packard reduced gigabit media-independent interface (RGMII) specification version 1.2a, dated 9/22/2000. RMII mode details follow the RMII Consortium Specification, dated 3/20/1998. All other modes follow the IEEE 802.3 standard, 2000 Edition. Input signals not used are internally disabled. Except for TSEC<sub>n</sub>\_GTX\_CLK, output signals not used are driven low.

Table 15-2. eTSEC Signals—Detailed Signal Descriptions

Signal	I/O	Description
TSEC <sub>n</sub> _COL	I	Collision input. The behavior of this signal is not specified while in full-duplex mode.
		<b>State Meaning</b> Asserted/Negated—In MII mode, this signal is asserted upon detection of a collision, and must remain asserted while the collision persists. This signal is not used in the following modes: <ul style="list-style-type: none"> <li>• RMII</li> <li>• RTBI</li> <li>• RGMII</li> </ul>
		<b>Timing</b> Asserted/Negated—This signal is not required to transition synchronously with TSEC <sub>n</sub> _TX_CLK or TSEC <sub>n</sub> _RX_CLK.
TSEC <sub>n</sub> _CRS	I	Carrier sense input. In RTBI mode, this signal is used as SDET (signal detect). In RTBI mode SDET is tied high internally. This signal is not used in the following modes: <ul style="list-style-type: none"> <li>• RMII</li> <li>• RGMII</li> </ul>
		<b>State Meaning</b> Asserted/Negated—In MII mode, TSEC <sub>n</sub> _CRS is asserted while the transmit or receive medium is not idle. In the event of a collision, TSEC <sub>n</sub> _CRS must remain asserted for the duration of the collision.
		<b>Timing</b> Asserted/Negated—This signal is not required to transition synchronously with TSEC <sub>n</sub> _TX_CLK or TSEC <sub>n</sub> _RX_CLK.
TSEC <sub>n</sub> _GTX_CLK	O	Gigabit transmit clock. This signal is an output from the eTSEC into the PHY. TSEC <sub>n</sub> _GTX_CLK is a 125-MHz clock that provides a timing reference for TX_EN, TXD, and TX_ER in the following modes: <ul style="list-style-type: none"> <li>• RTBI</li> </ul> In RGMII mode, TSEC <sub>n</sub> _GTX_CLK becomes the transmit clock and provides timing reference during 1000Base-T (125 MHz), 100Base-T (25 MHz) and 10Base-T (2.5 MHz) transmissions. This signal feeds back the uninverted transmit clock in MII mode, but feeds back an inverted transmit clock in RTBI or RGMII modes. This signal is driven low unless transmission is enabled.

Table 15-2. eTSEC Signals—Detailed Signal Descriptions (continued)

Signal	I/O	Description
EC_GTX_CLK125	I	Gigabit transmit 125-MHz source. This signal must be generated externally with a crystal or oscillator, or is sometimes provided by the PHY. EC_GTX_CLK125 is a 125-MHz input into the eTSEC and is used to generate all 125-MHz related signals and clocks in the following modes: <ul style="list-style-type: none"> <li>• RTBI</li> <li>• RGMII</li> </ul> This input is not used in these modes: <ul style="list-style-type: none"> <li>• RMII</li> <li>• SGMII</li> <li>• MII</li> </ul>
EC_MDC	O	Management data clock. This signal is a clock (typically 2.5 MHz) supplied by the MAC (IEEE set minimum period of 400 ns or a frequency of 2.5 MHz, but the device may be configured up to 12.5 MHz if supported by the PHY at that speed.) The frequency can be modified by writing to MIIMCFG[28:31] of the eTSEC1 controller.
EC_MDIO	I/O	Management data input/output.
		<b>State Meaning</b> Asserted/Negated—EC_MDIO is a bidirectional signal to input PHY-supplied status during management read cycles and output control during MII management write cycles. Addressed using eTSEC1 memory-mapped registers.
		<b>Timing</b> Asserted/Negated—This signal is required to be synchronous with the EC_MDC signal.
TSEC <sub>n</sub> _RX_CLK	I	Receive clock. In MII or RGMII mode, the receive clock TSEC <sub>n</sub> _RX_CLK is a continuous clock (2.5, 25, or 125 MHz) that provides a timing reference for TSEC <sub>n</sub> _RX_DV, TSEC <sub>n</sub> _RXD, and TSEC <sub>n</sub> _RX_ER. In RTBI mode it is a 125-MHz receive clock. In RMII mode this clock is not used for the receive clock, as RMII uses a shared reference clock.
TSEC <sub>n</sub> _RX_DV	I	Receive data valid. In MII mode, if TSEC <sub>n</sub> _RX_DV is asserted, the PHY is indicating that valid data is present on the MII interface. In RGMII mode, TSEC <sub>n</sub> _RX_DV becomes RX_CTL. The RX_DV and RX_ERR are received on this signal on the rising and falling edges of TSEC <sub>n</sub> _RX_CLK. In RTBI mode, TSEC <sub>n</sub> _RX_DV represents receive code group (RCG) bit 4 and 9. On the positive edge of the TSEC <sub>n</sub> _RX_CLK, RCG[4] and RCG[3:0] represent the first half of the 10-bit encoded symbol. On the negative edge of the TSEC <sub>n</sub> _RX_CLK, RCG[9] and RCG[8:5] represent the second half of the 10-bit encoded symbol. In RMII mode the PHY asserts TSEC <sub>n</sub> _RX_DV (CRS_DV) when the receive medium is non-idle. This signal asserts asynchronously with respect to the RMII reference clock, but negates synchronously to indicate loss of carrier.
TSEC <sub>n</sub> _RXD[3:0]	I	Receive data in. In MII mode, TSEC <sub>n</sub> _RXD[3:0] represents a nibble of data to be transferred from the PHY to the MAC when TSEC <sub>n</sub> _RX_DV is asserted. A completely-formed SFD must be passed across the MII. While TSEC <sub>n</sub> _RX_DV is not asserted, TSEC <sub>n</sub> _RXD has no meaning. In RGMII mode, data bits 3:0 are received on the rising edge of TSEC <sub>n</sub> _RX_CLK and data bits 7:4 are received on the falling edge of TSEC <sub>n</sub> _RX_CLK. In RTBI mode, TSEC <sub>n</sub> _RXD[3:0] represents RCG[3:0] on the rising edge of TSEC <sub>n</sub> _RX_CLK and RCG[8:5] are received on the falling edge of TSEC <sub>n</sub> _RX_CLK. In RMII mode, TSEC <sub>n</sub> _RXD[1:0] represents RXD[1:0], which is considered valid when TSEC <sub>n</sub> _RX_DV (CRS_DV) is asserted, or invalid otherwise.

Table 15-2. eTSEC Signals—Detailed Signal Descriptions (continued)

Signal	I/O	Description
TSEC <sub>n</sub> _RX_ER	I	Receive error
		<b>State Meaning</b>
TSEC <sub>n</sub> _TX_CLK	I	Transmit clock in. In MII mode, TSEC <sub>n</sub> _TX_CLK is a continuous clock (2.5 or 25 MHz) that provides a timing reference for the TSEC <sub>n</sub> _TX_EN, TSEC <sub>n</sub> _TXD, and TSEC <sub>n</sub> _TX_ER signals. In RMI mode this signal is the reference clock shared between transmit and receive, and is supplied by the PHY. This signal is not used in the eTSEC RTBI or RGMII modes.
TSEC <sub>n</sub> _TXD[3:0]	O	Transmit data out. DVIn MII mode, TSEC <sub>n</sub> _TXD[3:0] represent a nibble of data to be sent from the MAC to the PHY when TSEC <sub>n</sub> _TX_EN is asserted and have no meaning while TSEC <sub>n</sub> _TX_EN is negated. In RGMII or RTBI mode, data bits 3:0 are transmitted on the rising edge of TSEC <sub>n</sub> _GTX_CLK, and data bits 7:4 are transmitted on the falling edge of TSEC <sub>n</sub> _GTX_CLK. In RMI mode TSEC <sub>n</sub> _TXD[1:0] represents TXD[1:0], which is valid data sent to the PHY when TSEC <sub>n</sub> _TX_EN is asserted, or undefined otherwise. Note that some of these signals are also used during reset to configure the eTSEC interface mode.
TSEC <sub>n</sub> _TX_EN	O	Transmit data valid. In MII, or RMI mode, if TSEC <sub>n</sub> _TX_EN is asserted, the MAC is indicating that valid data is present on the MII's TSEC <sub>n</sub> _TXD signals. In RGMII mode, TSEC <sub>n</sub> _TX_EN becomes TX_CTL. TX_EN and TX_ERR are asserted on this signal on rising and falling edges of the TSEC <sub>n</sub> _GTX_CLK, respectively. In RTBI mode, TSEC <sub>n</sub> _TX_EN represents TCG[4] on the rising edge and TCG[9] on the falling edge of TSEC <sub>n</sub> _GTX_CLK, respectively. Together with TCG[3:0] and TCG[8:5], they represent the 10-bit encoded symbol.
TSEC <sub>n</sub> _TX_ER	O	Transmit error. In MII mode, assertion of TSEC <sub>n</sub> _TX_ER for one or more clock cycles while TSEC <sub>n</sub> _TX_EN is asserted causes the PHY to transmit one or more illegal symbols. Asserting TSEC <sub>n</sub> _TX_ER has no effect while operating at 10 Mbps or while TSEC <sub>n</sub> _TX_EN is negated. This signal transitions synchronously with respect to TSEC <sub>n</sub> _TX_CLK. This signal is not used in the eTSEC RMI, RTBI, or RGMII modes and is driven low.
TSEC_1588_CLK	I	1588 clock in. External high precision timer reference clock input (chip external input pin).
TSEC_1588_GCLK	O	1588 clock out. Phase aligned timer clock divider output (chip external output pin).
TSEC_1588_TRIG1	I	1588 trigger in 1. External timer trigger input 1. This is an asynchronous general purpose input (chip external input pin).
TSEC_1588_TRIG2	i	1588 trigger in 2. External timer trigger input 2. This is an asynchronous general purpose input (chip external input pin).
TSEC_1588_PP1	O	1588 pulse out 1. Timer pulse per period 1. It is phase aligned with 1588 timer clock (chip external output pin)
TSEC_1588_PP2	O	1588 pulse out 2. Timer pulse per period 2. It is phase aligned with 1588 timer clock (chip external output pin)
TSEC_1588_PP3	O	1588 pulse out 3. Timer pulse per period 3. It is phase aligned with 1588 timer clock (chip external output pin)



**Table 15-2. eTSEC Signals—Detailed Signal Descriptions (continued)**

Signal	I/O	Description
TSEC_1588_ALARM1	O	1588 timer alarm 1. Timer current time is equal to or greater than alarm time comparator register. User reprograms the TSEC_1588_ALARMn_H/L register to deactivate this output (chip external output pin)
TSEC_1588_ALARM2	O	1588 timer alarm 2. Timer current time is equal to or greater than alarm time comparator register. User reprograms the 1588_ALARMn_H/L register to deactivate this output (chip external output pin)

## 15.5 Memory Map/Register Definition

The eTSECs use a software model that is a superset of the PowerQUICC II Pro TSEC functionality and is similar to that employed by the Fast Ethernet function supported on the Freescale MPC8260 CPM FCC and in the FEC of the MPC860T.

The eTSEC device is programmed by a combination of control/status registers (CSRs) and buffer descriptors. The CSRs are used for mode control, interrupts, and to extract status information. The descriptors are used to pass data buffers and related buffer status or frame information between the hardware and software.

All accesses to and from the registers must be made as 32-bit accesses. There is no support for accesses of sizes other than 32 bits. Writes to reserved register bits must always store 0, as writing 1 to reserved bits may have unintended side-effects. Reads from unmapped register addresses return zero. Unless otherwise specified, the read value of reserved bits in mapped registers is not defined, and must not be assumed to be 0.

This section of the document defines the memory map and describes the registers in detail. The buffer descriptor is described in [Section 15.6.7, “Buffer Descriptors.”](#)

### 15.5.1 Top-Level Module Memory Map

Each of the eTSECs is allocated 4 Kbytes of memory-mapped space. The space for each eTSEC is divided as indicated in [Table 15-3](#).

**Table 15-3. Module Memory Map Summary**

Address Offset	Function
000–0FF	eTSEC general control/status registers
100–2FF	eTSEC transmit control/status registers
300–4FF	eTSEC receive control/status registers
500–5FF	MAC registers
600–7FF	RMON MIB registers
800–8FF	Hash table registers
900–9FF	—
A00–AFF	FIFO control/status registers
B00–BFF	DMA system registers

**Table 15-3. Module Memory Map Summary (continued)**

C00–C3F	Lossless Flow Control registers
C40–DFF	—
E00–EFF	1588 Hardware Assist

## 15.5.2 Detailed Memory Map

The eTSEC memory mapped registers are accessed by reading and writing to an address comprised of the base address (specified in IMMRBAR as defined in [Chapter 2, “Memory Map.”](#)) plus the block base address, plus the offset of the specific register to be accessed. Note that all memory-mapped registers must only be accessed as 32-bit quantities.

[Table 15-4](#) lists the offset, name, and a cross-reference to the complete description of each register. The offsets to the memory map table are applicable to each eTSEC. Block base addresses are as follows:

- eTSEC1 starts at 0x2\_4000 address offset
- eTSEC2 starts at 0x2\_5000 address offset

In this table and in the register figures and field descriptions, the following access definitions apply:

- Reserved fields are always ignored for the purposes of determining access type.
- R/W, R, and W (read/write, read only, and write only) indicate that all the non-reserved fields in a register have the same access type.
- w1c indicates that all of the non-reserved fields in a register are cleared by writing ones to them.
- Mixed indicates a combination of access types.
- Special is used when no other category applies. In this case the register figure and field description table should be read carefully.

**Table 15-4. Module Memory Map**

eTSEC1 Offset	Name <sup>1</sup>	Access <sup>2</sup>	Reset	Section/Page
<b>eTSEC General Control and Status Registers</b>				
0x2_4000	TSEC_ID*—Controller ID register	R	0x0124_0106	<a href="#">15.5.3.1.1/15-22</a>
0x2_4004	TSEC_ID2*—Controller ID register	R	0x0030_00F0	<a href="#">15.5.3.1.2/15-23</a>
0x2_4008– 0x2_400C	Reserved	—	—	—
0x2_4010	IEVENT—Interrupt event register	w1c	0x0000_0000	<a href="#">15.5.3.1.3/15-24</a>
0x2_4014	IMASK—Interrupt mask register	R/W	0x0000_0000	<a href="#">15.5.3.1.4/15-27</a>
0x2_4018	EDIS—Error disabled register	R/W	0x0000_0000	<a href="#">15.5.3.1.5/15-29</a>
0x2_401C	Reserved	—	—	—
0x2_4020	ECNTRL—Ethernet control register	R/W	0x0000_0000	<a href="#">15.5.3.1.6/15-31</a>
0x2_4024	Reserved	—	—	—
0x2_4028	PTV—Pause time value register	R/W	0x0000_0000	<a href="#">15.5.3.1.7/15-33</a>
0x2_402C	DMACTRL—DMA control register	R/W	0x0000_0000	<a href="#">15.5.3.1.8/15-34</a>

Table 15-4. Module Memory Map (continued)

eTSEC1 Offset	Name <sup>1</sup>	Access <sup>2</sup>	Reset	Section/Page
0x2_4030	Reserved	—	—	—
0x2_4034– 0x2_40FC	Reserved	—	—	—
<b>eTSEC Transmit Control and Status Registers</b>				
0x2_4100	TCTRL—Transmit control register	R/W	0x0000_0000	<a href="#">15.5.3.2.1/15-35</a>
0x2_4104	TSTAT—Transmit status register	w1c	0x0000_0000	<a href="#">15.5.3.2.2/15-37</a>
0x2_4108	DFVLAN*—Default VLAN control word	R/W	0x8100_0000	<a href="#">15.5.3.2.3/15-41</a>
0x2_410C	Reserved	—	—	—
0x2_4110	TXIC—Transmit interrupt coalescing register	R/W	0x0000_0000	<a href="#">15.5.3.2.4/15-42</a>
0x2_4114	TQUEUE*—Transmit queue control register	R/W	0x0000_8000	<a href="#">15.5.3.2.5/15-43</a>
0x2_4118– 0x2_413C	Reserved	—	—	—
0x2_4140	TR03WT*—TxBD Rings 0–3 round-robin weightings	R/W	0x0000_0000	<a href="#">15.5.3.2.6/15-43</a>
0x2_4144	TR47WT*—TxBD Rings 4–7 round-robin weightings	R/W	0x0000_0000	<a href="#">15.5.3.2.7/15-44</a>
0x2_4148– 0x2_4180	Reserved	—	—	—
0x2_4180	TBDBPH*—Tx data buffer pointer high bits	R/W	0x0000_0000	<a href="#">15.5.3.2.8/15-45</a>
0x2_4184	TBPTR0—TxBD pointer for ring 0	R/W	0x0000_0000	<a href="#">15.5.3.2.9/15-45</a>
0x2_4188	Reserved	—	—	—
0x2_418C	TBPTR1*—TxBD pointer for ring 1	R/W	0x0000_0000	<a href="#">15.5.3.2.9/15-45</a>
0x2_4190	Reserved	—	—	—
0x2_4194	TBPTR2*—TxBD pointer for ring 2	R/W	0x0000_0000	<a href="#">15.5.3.2.9/15-45</a>
0x2_4198	Reserved	—	—	—
0x2_419C	TBPTR3*—TxBD pointer for ring 3	R/W	0x0000_0000	<a href="#">15.5.3.2.9/15-45</a>
0x2_41A0	Reserved	—	—	—
0x2_41A4	TBPTR4*—TxBD pointer for ring 4	R/W	0x0000_0000	<a href="#">15.5.3.2.9/15-45</a>
0x2_41A8	Reserved	—	—	—
0x2_41AC	TBPTR5*—TxBD pointer for ring 5	R/W	0x0000_0000	<a href="#">15.5.3.2.9/15-45</a>
0x2_41B0	Reserved	—	—	—
0x2_41B4	TBPTR6*—TxBD pointer for ring 6	R/W	0x0000_0000	<a href="#">15.5.3.2.9/15-45</a>
0x2_41B8	Reserved	—	—	—
0x2_41BC	TBPTR7*—TxBD pointer for ring 7	R/W	0x0000_0000	<a href="#">15.5.3.2.9/15-45</a>
0x2_41C0– 0x2_4200	Reserved	—	—	—
0x2_4204	TBASE0—TxBD base address of ring 0	R/W	0x0000_0000	<a href="#">15.5.3.2.10/15-46</a>

Table 15-4. Module Memory Map (continued)

eTSEC1 Offset	Name <sup>1</sup>	Access <sup>2</sup>	Reset	Section/Page
0x2_4208	Reserved	—	—	—
0x2_420C	TBASE1*—TxBD base address of ring 1	R/W	0x0000_0000	15.5.3.2.10/15-46
0x2_4210	Reserved	—	—	—
0x2_4214	TBASE2*—TxBD base address of ring 2	R/W	0x0000_0000	15.5.3.2.10/15-46
0x2_4218	Reserved	—	—	—
0x2_421C	TBASE3*—TxBD base address of ring 3	R/W	0x0000_0000	15.5.3.2.10/15-46
0x2_4220	Reserved	—	—	—
0x2_4224	TBASE4*—TxBD base address of ring 4	R/W	0x0000_0000	15.5.3.2.10/15-46
0x2_4228	Reserved	—	—	—
0x2_422C	TBASE5*—TxBD base address of ring 5	R/W	0x0000_0000	15.5.3.2.10/15-46
0x2_4230	Reserved	—	—	—
0x2_4234	TBASE6*—TxBD base address of ring 6	R/W	0x0000_0000	15.5.3.2.10/15-46
0x2_4238	Reserved	—	—	—
0x2_423C	TBASE7*—TxBD base address of ring 7	R/W	0x0000_0000	15.5.3.2.10/15-46
0x2_4240– 0x2_427C	Reserved	—	—	—
0x2_4280	TMR_TXTS1_ID* - Tx timestamp identification tag (set 1)	R/W	0x0000_0000	15.5.3.2.11/15-47
0x2_4284	TMR_TXTS2_ID* - Tx timestamp identification tag (set 2)	R/W	0x0000_0000	15.5.3.2.11/15-47
0x2_4288– 0x2_42BC	Reserved	—	—	—
0x2_42C0	TMR_TXTS1_H* - Tx timestamp high (set 1)	R/W	0x0000_0000	15.5.3.2.12/15-47
0x2_42C4	TMR_TXTS1_L* - Tx timestamp high (set 1)	R/W	0x0000_0000	15.5.3.2.12/15-47
0x2_42C8	TMR_TXTS2_H* - Tx timestamp high (set 2)	R/W	0x0000_0000	15.5.3.2.12/15-47
0x2_42CC	TMR_TXTS2_L* - Tx timestamp high (set 2)	R/W	0x0000_0000	15.5.3.2.12/15-47
0x2_42D0– 0x2_42FC	Reserved	—	—	—
<b>eTSEC Receive Control and Status Registers</b>				
0x2_4300	RCTRL—Receive control register	R/W	0x0000_0000	15.5.3.3.1/15-48
0x2_4304	RSTAT—Receive status register	w1c	0x0000_0000	15.5.3.3.2/15-50
0x2_4308– 0x2_430C	Reserved	—	—	—
0x2_4310	RXIC—Receive interrupt coalescing register	R/W	0x0000_0000	15.5.3.3.3/15-52
0x2_4314	RQUEUE*—Receive queue control register.	R/W	0x0080_0080	15.5.3.3.4/15-53
0x2_4318– 0x2_432C	Reserved	—	—	—
0x2_4330	RBIFX*—Receive bit field extract control register	R/W	0x0000_0000	15.5.3.3.5/15-54

Table 15-4. Module Memory Map (continued)

eTSEC1 Offset	Name <sup>1</sup>	Access <sup>2</sup>	Reset	Section/Page
0x2_4334	RQFAR*—Receive queue filing table address register	R/W	0x0000_0000	15.5.3.3.6/15-55
0x2_4338	RQFCR*—Receive queue filing table control register	R/W	0xn <sub>nnnn</sub> _n <sub>nnnn</sub>	15.5.3.3.7/15-56
0x2_433C	RQFPR*—Receive queue filing table property register	R/W	0xn <sub>nnnn</sub> _n <sub>nnnn</sub>	15.5.3.3.8/15-57
0x2_4340	MRBLR—Maximum receive buffer length register	R/W	0x0000_0000	15.5.3.3.9/15-61
0x2_4344– 0x2_4380	Reserved	—	—	—
0x2_4380	RBDBPH*—Rx data buffer pointer high bits	R/W	0x0000_0000	15.5.3.3.10/15-62
0x2_4384	RBPTR0—RxBd pointer for ring 0	R/W	0x0000_0000	15.5.3.3.11/15-62
0x2_4388	Reserved	—	—	—
0x2_438C	RBPTR1*—RxBd pointer for ring 1	R/W	0x0000_0000	15.5.3.3.11/15-62
0x2_4390	Reserved	—	—	—
0x2_4394	RBPTR2*—RxBd pointer for ring 2	R/W	0x0000_0000	15.5.3.3.11/15-62
0x2_4398	Reserved	—	—	—
0x2_439C	RBPTR3*—RxBd pointer for ring 3	R/W	0x0000_0000	15.5.3.3.11/15-62
0x2_43A0	Reserved	—	—	—
0x2_43A4	RBPTR4*—RxBd pointer for ring 4	R/W	0x0000_0000	15.5.3.3.11/15-62
0x2_43A8	Reserved	—	—	—
0x2_43AC	RBPTR5*—RxBd pointer for ring 5	R/W	0x0000_0000	15.5.3.3.11/15-62
0x2_43B0	Reserved	—	—	—
0x2_43B4	RBPTR6*—RxBd pointer for ring 6	R/W	0x0000_0000	15.5.3.3.11/15-62
0x2_43B8	Reserved	—	—	—
0x2_43BC	RBPTR7*—RxBd pointer for ring 7	R/W	0x0000_0000	15.5.3.3.11/15-62
0x2_43C0– 0x2_44400	Reserved	—	—	—
0x2_4404	RBASE0—RxBd base address of ring 0	R/W	0x0000_0000	15.5.3.3.12/15-63
0x2_4408	Reserved	—	—	—
0x2_440C	RBASE1*—RxBd base address of ring 1	R/W	0x0000_0000	15.5.3.3.12/15-63
0x2_4410	Reserved	—	—	—
0x2_4414	RBASE2*—RxBd base address of ring 2	R/W	0x0000_0000	15.5.3.3.12/15-63
0x2_4418	Reserved	—	—	—
0x2_441C	RBASE3*—RxBd base address of ring 3	R/W	0x0000_0000	15.5.3.3.12/15-63
0x2_4420	Reserved	—	—	—
0x2_4424	RBASE4*—RxBd base address of ring 4	R/W	0x0000_0000	15.5.3.3.12/15-63
0x2_4428	Reserved	—	—	—
0x2_442C	RBASE5*—RxBd base address of ring 5	R/W	0x0000_0000	15.5.3.3.12/15-63

Table 15-4. Module Memory Map (continued)

eTSEC1 Offset	Name <sup>1</sup>	Access <sup>2</sup>	Reset	Section/Page
0x2_4430	Reserved	—	—	—
0x2_4434	RBASE6*—RxBD base address of ring 6	R/W	0x0000_0000	15.5.3.3.12/15-63
0x2_4438	Reserved	—	—	—
0x2_443C	RBASE7*—RxBD base address of ring 7	R/W	0x0000_0000	15.5.3.3.12/15-63
0x2_4440– 0x2_44BC	Reserved	—	—	—
0x2_44C0	TMR_RXTS_H* - Rx timer timestamp register high	R/W	0x0000_0000	15.5.3.3.13/15-63
0x2_44C4	TMR_RXTS_L* - Rx timer timestamp register low	R/W	0x0000_0000	15.5.3.3.13/15-63
0x2_44C8– 0x2_44FC	Reserved	—	—	—
<b>eTSEC MAC Registers</b>				
0x2_4500	MACCFG1—MAC configuration register 1	R/W	0x0000_0000	15.5.3.5.1/15-67
0x2_4504	MACCFG2—MAC configuration register 2	R/W	0x0000_7000	15.5.3.5.2/15-68
0x2_4508	IPGIFG—Inter-packet/inter-frame gap register	R/W	0x4060_5060	15.5.3.5.3/15-70
0x2_450C	HAFDUP—Half-duplex control	R/W	0x00A1_F037	15.5.3.5.4/15-71
0x2_4510	MAXFRM—Maximum frame length	R/W	0x0000_0600	15.5.3.5.5/15-72
0x2_4514– 0x2_451C	Reserved	—	—	—
0x2_4520	MIIMCFG—MII management configuration	R/W	0x0000_0007	15.5.3.5.6/15-72
0x2_4524	MIIMCOM—MII management command	R/W	0x0000_0000	15.5.3.5.7/15-73
0x2_4528	MIIMADD—MII management address	R/W	0x0000_0000	15.5.3.5.8/15-74
0x2_452C	MIIMCON—MII management control	WO	0x0000_0000	15.5.3.5.9/15-75
0x2_4530	MIIMSTAT—MII management status	R	0x0000_0000	15.5.3.5.10/15-75
0x2_4534	MIIMIND—MII management indicator	R	0x0000_0000	15.5.3.5.11/15-76
0x2_4538	Reserved	—	—	—
0x2_453C	IFSTAT—Interface status	R	0x0000_0000	15.5.3.5.12/15-76
0x2_4540	MACSTNADDR1—MAC station address register 1	R/W	0x0000_0000	15.5.3.5.13/15-77
0x2_4544	MACSTNADDR2—MAC station address register 2	R/W	0x0000_0000	15.5.3.5.14/15-78

Table 15-4. Module Memory Map (continued)

eTSEC1 Offset	Name <sup>1</sup>	Access <sup>2</sup>	Reset	Section/Page
0x2_4548	MAC01ADDR1*—MAC exact match address 1, part 1	R/W	0x0000_0000	15.5.3.5.15/15-78 15.5.3.5.16/15-79
0x2_454C	MAC01ADDR2*—MAC exact match address 1, part 2	R/W	0x0000_0000	
0x2_4550	MAC02ADDR1*—MAC exact match address 2, part 1	R/W	0x0000_0000	
0x2_4554	MAC02ADDR2*—MAC exact match address 2, part 2	R/W	0x0000_0000	
0x2_4558	MAC03ADDR1*—MAC exact match address 3, part 1	R/W	0x0000_0000	
0x2_455C	MAC03ADDR2*—MAC exact match address 3, part 2	R/W	0x0000_0000	
0x2_4560	MAC04ADDR1*—MAC exact match address 4, part 1	R/W	0x0000_0000	
0x2_4564	MAC04ADDR2*—MAC exact match address 4, part 2	R/W	0x0000_0000	
0x2_4568	MAC05ADDR1*—MAC exact match address 5, part 1	R/W	0x0000_0000	
0x2_456C	MAC05ADDR2*—MAC exact match address 5, part 2	R/W	0x0000_0000	
0x2_4570	MAC06ADDR1*—MAC exact match address 6, part 1	R/W	0x0000_0000	15.5.3.5.15/15-78 15.5.3.5.16/15-79
0x2_4574	MAC06ADDR2*—MAC exact match address 6, part 2	R/W	0x0000_0000	
0x2_4578	MAC07ADDR1*—MAC exact match address 7, part 1	R/W	0x0000_0000	
0x2_457C	MAC07ADDR2*—MAC exact match address 7, part 2	R/W	0x0000_0000	
0x2_4580	MAC08ADDR1*—MAC exact match address 8, part 1	R/W	0x0000_0000	
0x2_4584	MAC08ADDR2*—MAC exact match address 8, part 2	R/W	0x0000_0000	
0x2_4588	MAC09ADDR1*—MAC exact match address 9, part 1	R/W	0x0000_0000	
0x2_458C	MAC09ADDR2*—MAC exact match address 9, part 2	R/W	0x0000_0000	
0x2_4590	MAC10ADDR1*—MAC exact match address 10, part 1	R/W	0x0000_0000	
0x2_4594	MAC10ADDR2*—MAC exact match address 10, part 2	R/W	0x0000_0000	
0x2_4598	MAC11ADDR1*—MAC exact match address 11, part 1	R/W	0x0000_0000	
0x2_459C	MAC11ADDR2*—MAC exact match address 11, part 2	R/W	0x0000_0000	
0x2_45A0	MAC12ADDR1*—MAC exact match address 12, part 1	R/W	0x0000_0000	
0x2_45A4	MAC12ADDR2*—MAC exact match address 12, part 2	R/W	0x0000_0000	
0x2_45A8	MAC13ADDR1*—MAC exact match address 13, part 1	R/W	0x0000_0000	
0x2_45AC	MAC13ADDR2*—MAC exact match address 13, part 2	R/W	0x0000_0000	
0x2_45B0	MAC14ADDR1*—MAC exact match address 14, part 1	R/W	0x0000_0000	
0x2_45B4	MAC14ADDR2*—MAC exact match address 14, part 2	R/W	0x0000_0000	
0x2_45B8	MAC15ADDR1*—MAC exact match address 15, part 1	R/W	0x0000_0000	
0x2_45BC	MAC15ADDR2*—MAC exact match address 15, part 2	R/W	0x0000_0000	
0x2_45C0– 0x2_467C	Reserved	—	—	—
<b>eTSEC Transmit and Receive Counters</b>				
0x2_4680	TR64—Transmit and receive 64-byte frame counter	R/W	0x0000_0000	15.5.3.6.1/15-80
0x2_4684	TR127—Transmit and receive 65- to 127-byte frame counter	R/W	0x0000_0000	15.5.3.6.2/15-80
0x2_4688	TR255—Transmit and receive 128- to 255-byte frame counter	R/W	0x0000_0000	15.5.3.6.3/15-81

Table 15-4. Module Memory Map (continued)

eTSEC1 Offset	Name <sup>1</sup>	Access <sup>2</sup>	Reset	Section/Page
0x2_468C	TR511—Transmit and receive 256- to 511-byte frame counter	R/W	0x0000_0000	15.5.3.6.4/15-81
0x2_4690	TR1K—Transmit and receive 512- to 1023-byte frame counter	R/W	0x0000_0000	15.5.3.6.5/15-82
0x2_4694	TRMAX—Transmit and receive 1024- to 1518-byte frame counter	R/W	0x0000_0000	15.5.3.6.6/15-82
0x2_4698	TRMGV—Transmit and receive 1519- to 1522-byte good VLAN frame count	R/W	0x0000_0000	15.5.3.6.7/15-83
<b>eTSEC Receive Counters</b>				
0x2_469C	RBYT—Receive byte counter	R/W	0x0000_0000	15.5.3.6.8/15-83
0x2_46A0	RPKT—Receive packet counter	R/W	0x0000_0000	15.5.3.6.9/15-83
0x2_46A4	RFCS—Receive FCS error counter	R/W	0x0000_0000	15.5.3.6.10/15-84
0x2_46A8	RMCA—Receive multicast packet counter	R/W	0x0000_0000	15.5.3.6.11/15-84
0x2_46AC	RBCA—Receive broadcast packet counter	R/W	0x0000_0000	15.5.3.6.12/15-85
0x2_46B0	RXCF—Receive control frame packet counter	R/W	0x0000_0000	15.5.3.6.13/15-85
0x2_46B4	RXPF—Receive PAUSE frame packet counter	R/W	0x0000_0000	15.5.3.6.14/15-86
0x2_46B8	RXUO—Receive unknown OP code counter	R/W	0x0000_0000	15.5.3.6.15/15-86
0x2_46BC	RALN—Receive alignment error counter	R/W	0x0000_0000	15.5.3.6.16/15-87
0x2_46C0	RFLR—Receive frame length error counter	R/W	0x0000_0000	15.5.3.6.17/15-87
0x2_46C4	RCDE—Receive code error counter	R/W	0x0000_0000	15.5.3.6.18/15-88
0x2_46C8	RCSE—Receive carrier sense error counter	R/W	0x0000_0000	15.5.3.6.19/15-88
0x2_46CC	RUND—Receive undersize packet counter	R/W	0x0000_0000	15.5.3.6.20/15-89
0x2_46D0	ROVR—Receive oversize packet counter	R/W	0x0000_0000	15.5.3.6.21/15-89
0x2_46D4	RFRG—Receive fragments counter	R/W	0x0000_0000	15.5.3.6.22/15-90
0x2_46D8	RJBR—Receive jabber counter	R/W	0x0000_0000	15.5.3.6.23/15-90
0x2_46DC	RDRP—Receive drop counter	R/W	0x0000_0000	15.5.3.6.24/15-91
<b>eTSEC Transmit Counters</b>				
0x2_46E0	TBYT—Transmit byte counter	R/W	0x0000_0000	15.5.3.6.25/15-91
0x2_46E4	TPKT—Transmit packet counter	R/W	0x0000_0000	15.5.3.6.26/15-92
0x2_46E8	TMCA—Transmit multicast packet counter	R/W	0x0000_0000	15.5.3.6.27/15-92
0x2_46EC	TBCA—Transmit broadcast packet counter	R/W	0x0000_0000	15.5.3.6.28/15-93
0x2_46F0	TXPF—Transmit PAUSE control frame counter	R/W	0x0000_0000	15.5.3.6.29/15-93
0x2_46F4	TDFR—Transmit deferral packet counter	R/W	0x0000_0000	15.5.3.6.30/15-94
0x2_46F8	TEDF—Transmit excessive deferral packet counter	R/W	0x0000_0000	15.5.3.6.31/15-94
0x2_46FC	TSCL—Transmit single collision packet counter	R/W	0x0000_0000	15.5.3.6.32/15-95
0x2_4700	TMCL—Transmit multiple collision packet counter	R/W	0x0000_0000	15.5.3.6.33/15-95
0x2_4704	TLCL—Transmit late collision packet counter	R/W	0x0000_0000	15.5.3.6.34/15-96



Table 15-4. Module Memory Map (continued)

eTSEC1 Offset	Name <sup>1</sup>	Access <sup>2</sup>	Reset	Section/Page
0x2_4708	TXCL—Transmit excessive collision packet counter	R/W	0x0000_0000	15.5.3.6.35/15-96
0x2_470C	TNCL—Transmit total collision counter	R/W	0x0000_0000	15.5.3.6.36/15-97
0x2_4710	Reserved	—	—	—
0x2_4714	TDRP—Transmit drop frame counter	R/W	0x0000_0000	15.5.3.6.37/15-97
0x2_4718	TJBR—Transmit jabber frame counter	R/W	0x0000_0000	15.5.3.6.38/15-98
0x2_471C	TFCS—Transmit FCS error counter	R/W	0x0000_0000	15.5.3.6.39/15-98
0x2_4720	TXCF—Transmit control frame counter	R/W	0x0000_0000	15.5.3.6.40/15-99
0x2_4724	TOVR—Transmit oversize frame counter	R/W	0x0000_0000	15.5.3.6.41/15-99
0x2_4728	TUND—Transmit undersize frame counter	R/W	0x0000_0000	15.5.3.6.42/15-100
0x2_472C	TFRG—Transmit fragments frame counter	R/W	0x0000_0000	15.5.3.6.43/15-100
<b>eTSEC Counter Control and TOE Statistics Registers</b>				
0x2_4730	CAR1—Carry register one register <sup>3</sup>	R	0x0000_0000	15.5.3.6.44/15-101
0x2_4734	CAR2—Carry register two register <sup>3</sup>	R	0x0000_0000	15.5.3.6.45/15-102
0x2_4738	CAM1—Carry register one mask register	R/W	0xFE03_FFFF	15.5.3.6.46/15-103
0x2_473C	CAM2—Carry register two mask register	R/W	0x000F_FFFD	15.5.3.6.47/15-105
0x2_4740	RREJ*—Receive filer rejected packet counter	R/W	0x0000_0000	15.5.3.6.48/15-106
0x2_4744– 0x2_47FC	Reserved	—	—	—
<b>Hash Function Registers</b>				
0x2_4800	IGADDR0—Individual/group address register 0	R/W	0x0000_0000	15.5.3.7.1/15-107
0x2_4804	IGADDR1—Individual/group address register 1	R/W	0x0000_0000	
0x2_4808	IGADDR2—Individual/group address register 2	R/W	0x0000_0000	
0x2_480C	IGADDR3—Individual/group address register 3	R/W	0x0000_0000	
0x2_4810	IGADDR4—Individual/group address register 4	R/W	0x0000_0000	
0x2_4814	IGADDR5—Individual/group address register 5	R/W	0x0000_0000	
0x2_4818	IGADDR6—Individual/group address register 6	R/W	0x0000_0000	
0x2_481C	IGADDR7—Individual/group address register 7	R/W	0x0000_0000	
0x2_4820– 0x2_487C	Reserved	—	—	—

Table 15-4. Module Memory Map (continued)

eTSEC1 Offset	Name <sup>1</sup>	Access <sup>2</sup>	Reset	Section/Page
0x2_4880	GADDR0—Group address register 0	R/W	0x0000_0000	15.5.3.7.2/15-107
0x2_4884	GADDR1—Group address register 1	R/W	0x0000_0000	
0x2_4888	GADDR2—Group address register 2	R/W	0x0000_0000	
0x2_488C	GADDR3—Group address register 3	R/W	0x0000_0000	
0x2_4890	GADDR4—Group address register 4	R/W	0x0000_0000	
0x2_4894	GADDR5—Group address register 5	R/W	0x0000_0000	
0x2_4898	GADDR6—Group address register 6	R/W	0x0000_0000	
0x2_489C	GADDR7—Group address register 7	R/W	0x0000_0000	
0x2_48A0– 0x2_4AFC	Reserved	—	—	—
<b>eTSEC DMA Attribute Registers</b>				
0x2_4B00– 0x2_4BF4	Reserved	—	—	—
0x2_4BF8	ATTR—Attribute register	R/W	0x0000_0000	15.5.3.8.1/15-108
<b>eTSEC Lossless Flow Control Registers</b>				
0x2_4C00	RQPRM0*—Receive Queue Parameters register 0	R/W	0x0000_0000	15.5.3.9.1/15-109
0x2_4C04	RQPRM1*—Receive Queue Parameters register 1	R/W	0x0000_0000	
0x2_4C08	RQPRM2*—Receive Queue Parameters register 2	R/W	0x0000_0000	
0x2_4C0C	RQPRM3*—Receive Queue Parameters register 3	R/W	0x0000_0000	
0x2_4C10	RQPRM4*—Receive Queue Parameters register 4	R/W	0x0000_0000	
0x2_4C14	RQPRM5*—Receive Queue Parameters register 5	R/W	0x0000_0000	
0x2_4C18	RQPRM6*—Receive Queue Parameters register 6	R/W	0x0000_0000	
0x2_4C1C	RQPRM7*—Receive Queue Parameters register 7	R/W	0x0000_0000	
0x2_4C20– 0x2_4C40	Reserved	—	—	—
0x2_4C44	RFBPTR0*—Last Free RxBD pointer for ring 0	R/W	0x0000_0000	15.5.3.9.2/15-109
0x2_4C48	Reserved	—	—	—
0x2_4C4C	RFBPTR1*—Last Free RxBD pointer for ring 1	R/W	0x0000_0000	15.5.3.9.2/15-109
0x2_4C50	Reserved	—	—	—
0x2_4C54	RFBPTR2*—Last Free RxBD pointer for ring 2	R/W	0x0000_0000	15.5.3.9.2/15-109
0x2_4C58	Reserved	—	—	—
0x2_4C5C	RFBPTR3*—Last Free RxBD pointer for ring 3	R/W	0x0000_0000	15.5.3.9.2/15-109
0x2_4C60	Reserved	—	—	—
0x2_4C64	RFBPTR4*—Last Free RxBD pointer for ring 4	R/W	0x0000_0000	15.5.3.9.2/15-109
0x2_4C68	Reserved	—	—	—

Table 15-4. Module Memory Map (continued)

eTSEC1 Offset	Name <sup>1</sup>	Access <sup>2</sup>	Reset	Section/Page
0x2_4C6C	RFBPTR5*—Last Free RxBD pointer for ring 5	R/W	0x0000_0000	<a href="#">15.5.3.9.2/15-109</a>
0x2_4C70	Reserved	—	—	—
0x2_4C74	RFBPTR6*—Last Free RxBD pointer for ring 6	R/W	0x0000_0000	<a href="#">15.5.3.9.2/15-109</a>
0x2_4C78	Reserved	—	—	—
0x2_4C7C	RFBPTR7*—Last Free RxBD pointer for ring 7	R/W	0x0000_0000	<a href="#">15.5.3.9.2/15-109</a>
<b>eTSEC Future Expansion Space</b>				
0x2_4CC0 — 0x2_4D94	Reserved	—	—	—
<b>eTSEC IEEE 1588 Registers</b>				
0x2_4E00	TMR_CTRL*—Timer control register	R/W	0x0001_0001	<a href="#">15.5.3.10.1/15-110</a>
0x2_4E04	TMR_TEVENT*—Timestamp event register	W1C	0x0000_0000	<a href="#">15.5.3.10.2/15-112</a>
0x2_4E08	TMR_TEMASK*—Timer event mask register	R/W	0x0000_0000	<a href="#">15.5.3.10.3/15-114</a>
0x2_4E0C	TMR_PEVENT*—Timestamp event register	R/W	0x0000_0000	<a href="#">15.5.3.10.4/15-115</a>
0x2_4E10	TMR_PEMASK*—Timer event mask register	R/W	0x0000_0000	<a href="#">15.5.3.10.5/15-115</a>
0x2_4E14	TMR_STAT*—Timestamp status register	R/W	0x0000_0000	<a href="#">15.5.3.10.6/15-116</a>
0x2_4E18	TMR_CNT_H*—Timer counter high register	R/W	0x0000_0000	<a href="#">15.5.3.10.7/15-117</a>
0x2_4E1C	TMR_CNT_L*—Timer counter low register	R/W	0x0000_0000	<a href="#">15.5.3.10.7/15-117</a>
0x2_4E20	TMR_ADD*—Timer drift compensation addend register	R/W	0x0000_0000	<a href="#">15.5.3.10.8/15-117</a>
0x2_4E24	TMR_ACC*—Timer accumulator register	R/W	0x0000_0000	<a href="#">15.5.3.10.9/15-118</a>
0x2_4E28	TMR_PRSC*—Timer prescale	R/W	0x0000_0002	<a href="#">15.5.3.10.10/15-118</a>
0x2_4E2C	Reserved	—	—	—
0x2_4E30	TMROFF_H*—Timer offset high	R/W	0x0000_0000	<a href="#">15.5.3.10.11/15-119</a>
0x2_4E34	TMROFF_L*—Timer offset low	R/W	0x0000_0000	<a href="#">15.5.3.10.11/15-119</a>
0x2_4E40	TMR_ALARM1_H*—Timer alarm 1 high register	R/W	0xFFFF_FFFF	<a href="#">15.5.3.10.12/15-119</a>
0x2_4E44	TMR_ALARM1_L*—Timer alarm 1 high register	R/W	0xFFFF_FFFF	
0x2_4E48	TMR_ALARM2_H*—Timer alarm 2 high register	R/W	0xFFFF_FFFF	
0x2_4E4C	TMR_ALARM2_L*—Timer alarm 2 high register	R/W	0xFFFF_FFFF	
0x2_4E50– 0x2_4E7C	Reserved	—	—	—
0x2_4E80	TMR_FIPER1*—Timer fixed period interval	R/W	0xFFFF_FFFF	<a href="#">15.5.3.10.13/15-120</a>
0x2_4E84	TMR_FIPER2*—Timer fixed period interval	R/W	0xFFFF_FFFF	
0x2_4E88	TMR_FIPER*3—Timer fixed period interval	R/W	0xFFFF_FFFF	

**Table 15-4. Module Memory Map (continued)**

eTSEC1 Offset	Name <sup>1</sup>	Access <sup>2</sup>	Reset	Section/Page
0x2_4EA0	TMR_ETTS1_H*—Timestamp of general purpose external trigger	R/W	0x0000_0000	15.5.3.10.14/15-121
0x2_4EA4	TMR_ETTS1_L*—Timestamp of general purpose external trigger	R/W	0x0000_0000	
0x2_4EA8	TMR_ETTS2_H*—Timestamp of general purpose external trigger	R/W	0x0000_0000	
0x2_4EAC	TMR_ETTS2_L*—Timestamp of general purpose external trigger	R/W	0x0000_0000	
0x2_4EB0 — 0x2_4FFF	Reserved	—	—	—
<b>Other eTSECs</b>				
0x2_5000– 0x2_5FFF	eTSEC2 REGISTERS <sup>4</sup>			

<sup>1</sup> Registers denoted \* are new to the enhanced TSEC and not supported by PowerQUICC II Pro TSECs.

<sup>2</sup> Key: R = read only, WO = write only, R/W = read and write, LH = latches high, SC = self-clearing.

<sup>3</sup> Cleared on read.

<sup>4</sup> eTSEC2 has the same memory-mapped registers that are described for eTSEC1 from 0x 2\_4000 to 0x2\_4FFF, except the offsets are from 0x 2\_5000 to 0x2\_5FFF.

### 15.5.3 Memory-Mapped Register Descriptions


This section provides a detailed description of all the eTSEC registers. Because all of the eTSEC registers are 32 bits wide, only 32-bit register accesses are supported.

#### 15.5.3.1 eTSEC General Control and Status Registers

This section describes general control and status registers used for both transmitting and receiving Ethernet frames. All of the registers are 32 bits wide.

##### 15.5.3.1.1 Controller ID Register (TSEC\_ID)

The controller ID register (TSEC\_ID) is a read-only register. The TSEC\_ID register is used to identify the eTSEC block and revision.



**Figure 15-2. TSEC\_ID Register**


Table 15-10 describes the fields of the TSEC\_ID register.

**Table 15-5. TSEC\_ID Field Descriptions**

Bits	Name	Description
0–15	TSEC_ID	Value identifying the eTSEC (10/100/1000 Ethernet MAC). 0124 Unique identifier for eTSEC with 8 Rx and 8 Tx BD rings. 0800 Unique identifier for GMAC1 with 8 Rx and 8 Tx BD rings. 0810 Unique identifier for GMAC2 with 8 Rx and 8 Tx BD rings.
16–23	TSEC_REV_MJ	Value identifies the major revision of the eTSEC. 01 Silicon Rev 2.1
24–31	TSEC_REV_MN	Value identifies the minor revision of the eTSEC. 06 Silicon Rev 2.1

### 15.5.3.1.2 Controller ID Register (TSEC\_ID2)

The controller ID register (TSEC\_ID2) is a read-only register. The TSEC\_ID2 register is used to identify the eTSEC block configuration.



**Figure 15-3. TSEC\_ID2 Register**

Table 15-6 describes the fields of the TSEC\_ID2 register.

**Table 15-6. TSEC\_ID2 Field Descriptions**

Bits	Name	Description
0–9	—	Reserved
10–15	TSEC_INT	Interface mode support. See Table 15-7 for settings.
16–23	—	Reserved
24–31	TSEC_CFG	Value identifies configuration options of the eTSEC. 00 eTSEC multiple ring, Rx TOE, Filer and Tx TOE supports are off F0 eTSEC multiple ring, Rx TOE, Filer and Tx TOE supports are on 30 eTSEC multiple ring support is OFF and Rx TOE, Filer and Tx TOE supports are on 50 eTSEC multiple ring and filer supports are OFF and Rx TOE and Tx TOE supports are on

Table 15-7 describes the field settings for TSEC\_ID2[TSEC\_INT].

**Table 15-7. TSEC\_ID2[TSEC\_INT] Field Settings**

Bit	Mode
10	0 Ethernet mode not supported 1 Ethernet mode supported
11–13	Reserved

**Table 15-7. TSEC\_ID2[TSEC\_INT] Field Settings (continued)**

14	0 Can be configured to run in Ethernet normal/full mode 1 Ethernet normal/full mode off
15	0 Can be configured to run in Ethernet reduced mode 1 Ethernet reduced mode off

### 15.5.3.1.3 Interrupt Event Register (IEVENT)

Interrupt events cause bits in the IEVENT register to be set. Software may poll this register at any time to check for pending interrupts. If an event occurs and its corresponding enable bit is set in the interrupt mask register (IMASK), the event also causes a hardware interrupt at the PIC. A bit in the interrupt event register is cleared by writing a 1 to that bit position. A write of 0 has no effect.

Each eTSEC can issue three kinds of hardware interrupt to the PIC:

1. Transmit data frame interrupts—Issued whenever bits TXB or TXF of IEVENT are set to 1 and either transmit interrupt coalescing is disabled or the interrupt coalescing thresholds have been met for TXF. To negate this hardware interrupt, software must clear both TXB and TXF bits.
2. Receive data frame interrupts—Issued whenever bits RXB or RXF of IEVENT are set to 1 and either receive interrupt coalescing is disabled or the interrupt coalescing thresholds have been met for RXF. To negate this hardware interrupt, software must clear both RXB and RXF bits.
3. Error, diagnostic, and special interrupts—Issued whenever bits MAG, GTSC, GRSC, TXC, RXC, BABR, BAPT, LC, CRL, FGPI, FIR, FIQ, DPE, PERR, EBERR, TXE, XFUN, BSY, MSRO, MMRD, or MMRW of IEVENT are set to 1. Software must clear all of these bits to negate an error/diagnostic/special hardware interrupt.
  - Magic Packet reception event is: MAG
  - Operational diagnostics are events on: GTSC, GRSC, TXC, and RXC
  - Interrupts resulting from errors/problems detected in the network or transceiver are: BABR, BAPT, LC, and CRL
  - Interrupts resulting from internal or combination errors are: FIR, FIQ, DPE, PERR, EBERR, TXE, XFUN, and BSY
  - Special function interrupts are: FGPI, MSRO, MMRD, and MMRW

Some of the error interrupts are independently counted in the MIB block counters. Software may choose to mask off these interrupts because these errors are visible to network management through the MIB counters.

Figure 15-4 describes the definition for the IEVENT register.

Offset eTSEC1:0x2\_4010; eTSEC2: 0x2\_5010

Access: w1c

	0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	
R	BABR	RXC	BSY	EBERR	—	MSRO	GTSC	BABT	TXC	TXE	TXB	TXF	—	LC	CRL	XFUN	
W	w1c	w1c	w1c	w1c	—	w1c	w1c	w1c	w1c	w1c	w1c	w1c	—	w1c	w1c	w1c	
Reset	All zeros																
	16	17	18	19	20	21	22	23	24	25	26	27	28	29	30	31	
R	RXB	—			MAG	MMRD	MMWR	GRSC	RXF	—			FGP I	FIR	FIQ	DPE	PERR
W	w1c	—			w1c	w1c	w1c	w1c	w1c	—			w1c	w1c	w1c	w1c	w1c
Reset	All zeros																

Figure 15-4. IEVENT Register Definition

Table 15-8 describes the fields of the IEVENT register.

Table 15-8. IEVENT Field Descriptions

Bits	Name	Description
0	BABR	Babbling receive error. This bit indicates that a frame was received with length in excess of the MAC's maximum frame length register while MACCFG2[Huge Frame] is set. 0 Excessive frame not received. 1 Excessive frame received.
1	RXC	Receive control interrupt. A control frame was received while MACCFG1[Rx_Flow] is set. As soon as the transmitter finishes sending the current frame, a pause operation is performed. 0 Control frame not received. 1 Control frame received.
2	BSY	Busy condition interrupt. Indicates that a frame was received and discarded due to a lack of buffers. 0 No frame received and discarded. 1 Frame received and discarded.
3	EBERR	Internal bus error. This bit indicates that a system bus error occurred while a DMA transaction was underway. As a result, transferred data is expected to be partially or completely invalid. 0 No system bus error occurred. 1 System bus error occurred.
4	—	Reserved
5	MSRO	MIB counter overflow. This interrupt is asserted if the count for one of the MIB counters has exceeded the size of its register. 0 MIB count not exceeding its register size. 1 MIB count exceeds its register size.
6	GTSC	Graceful transmit stop complete. This interrupt is asserted for one of two reasons. Graceful stop means that the transmitter is put into a pause state after completion of the frame currently being transmitted. <ul style="list-style-type: none"> <li>A graceful stop, which was initiated by setting DMACTRL[GTS], is now complete.</li> <li>A transmission of a flow control PAUSE frame, which was initiated by setting TCTRL[TFC_PAUSE], is now complete.</li> </ul> 0 No graceful stop interrupt. 1 Graceful stop requested.

Table 15-8. IEVENT Field Descriptions (continued)

Bits	Name	Description
7	BABT	Babbling transmit error. This bit indicates that the transmitted frame length has exceeded the value in the MAC's maximum frame length register and MACCFG2[Huge Frame] is cleared. Frame truncation occurs when this condition occurs. 0 Transmitted frame length not exceeding maximum frame length. 1 Transmitted frame length exceeding maximum frame length when MACCFG2[Huge Frame] = 0.
8	TXC	Transmit control interrupt. This bit indicates that a control frame was transmitted. 0 Control frame not transmitted. 1 Control frame transmitted.
9	TXE	Transmit error. This bit indicates that an error occurred on the transmitted channel that has caused TSTAT[THLT] to be set by the eTSEC. This bit is set whenever any transmit error occurs that causes the transmitter to halt (EBERR, LC, CRL, XFUN). 0 No transmit channel error occurred. 1 Transmit channel error occurred.
10	TXB	Transmit buffer. This bit indicates that a transmit buffer descriptor was updated whose I (interrupt) bit was set in its status word and was not the last buffer descriptor of the frame. 0 No transmit buffer descriptor updated. 1 Transmit buffer descriptor updated.
11	TXF	Transmit frame interrupt. This bit indicates that a frame was transmitted and that the last corresponding transmit buffer descriptor (TxBD) was updated. This only occurs if the I (interrupt) bit in the status word of the buffer descriptor is set. The specific transmit queue that was updated has its TXF bit set in TSTAT. 0 No frame transmitted/TxBD not updated. 1 Frame transmitted/TxBD updated.
12	—	Reserved
13	LC	Late collision. This bit indicates that a collision occurred beyond the collision window (slot time) in half-duplex mode. The frame is truncated with a bad CRC and the remainder of the frame is discarded. 0 No late collision occurred. 1 Late collision occurred.
14	CRL	Collision retry limit. This bit indicates that the number of successive transmission collisions has exceeded the MAC's half-duplex register's retransmission maximum count (HAFDUP[Retransmission Maximum]). The frame is discarded without being transmitted and transmission of the next frame commences. This only occurs while in half-duplex mode. 0 Successive transmission collisions do not exceed maximum. 1 Successive transmission collisions exceed maximum.
15	XFUN	Transmit FIFO underrun. This bit indicates that the transmit FIFO became empty before the complete frame was transmitted. 0 Transmit FIFO not underrun. 1 Transmit FIFO underrun.
16	RXB	Receive buffer. This bit indicates that a receive buffer descriptor was updated which had the I (Interrupt) bit set in its status word and was not the last buffer descriptor of the frame. 0 Receive buffer descriptor not updated. 1 Receiver buffer descriptor updated.
17–19	—	Reserved
20	MAG	Magic Packet detected when the eTSEC is in Magic Packet detection mode (MACCFG2[MPEN] = 1). 0 No Magic Packet received, or Magic Packet mode was not enabled. 1 A Magic Packet was received while in Magic Packet mode. MACCFG2[MPEN] is also cleared upon receiving the Magic Packet.



Table 15-8. IEVENT Field Descriptions (continued)

Bits	Name	Description
21	MMRD	MII management read completion 0 MII management read not issued or in process. 1 MII management read completed that was initiated by a user through the MII Scan or Read cycle command.
22	MMWR	MII management write completion 0 MII management write not issued or in process. 1 MII management write completed that was initiated by a user write to the MIIMCON register.
23	GRSC	Graceful receive stop complete. This interrupt is asserted if a graceful receive stop is completed. It allows the user to know if the system has completed the stop and it is safe to write to receive registers (status, control or configuration registers) that are used by the system during normal operation. 0 Graceful stop not completed. 1 Graceful stop completed.
24	RXF	Receive frame interrupt. This bit indicates that a frame was received and the last receive buffer descriptor (RxBd) in that frame was updated. This occurs either if the I (interrupt) bit in the buffer descriptor status word is set, or an overrun error occurs. The specific receive queue that was updated has its RXF bit set in RSTAT. 0 Frame not received. 1 Frame received.
25–26	—	Reserved
27	FGPI	Filer generated general purpose interrupt on a set of filer rule match. This bit will be set upon reception of a frame that matches a GPI rule sequence that is specified in the filer. It is synchronized with the setting of RXF. 0 No filer generated interrupt has occurred. 1 The filer has accepted a frame via a matching rule that the RQFCR[GPI] bit set.
28	FIR	The receive queue filer result is invalid, either because not enough time between frames was available to find a matching rule, or no entry in the filer table could be matched. 0 Receive queue filer reached a definite result; however, bit FIQ may still be set if a frame was filed to a disabled RxBd ring. 1 Receive queue filer was unable to reach a definite result. In this case, bit FIQ is also set if no entry in the filer table could provide a rule match.
29	FIQ	Filed frame to invalid receive queue. This bit indicates that either the receive queue filer chose to DMA a received frame to a disabled RxBd ring, or that no rule in the filer table could be matched. 0 Received frames filed to valid queues or rejected. Note that a frame may be rejected if the filer has insufficient time to reach a conclusive result between frames, in which case bit FIR is set. 1 Received frames filed to RxBd rings that are not enabled. The frame is discarded. If bit FIR is also set this indicates that the filer exhausted all of its table entries without a rule match.
30	DPE	Internal data parity error. This bit indicates that the eTSEC has detected a parity error on its stored data, which is likely to compromise the validity of recently transferred frames. 0 No parity errors detected. 1 Data held in the FIFO or filer arrays is expected to be corrupted due to a parity error.
31	PERR	Receive frame parse error for TCP/IP off-load. This bit indicates that a received frame could not be parsed unambiguously, due to encapsulated header type fields contradicting each other. 0 Received frame parsed successfully. 1 Received frame parse revealed header inconsistencies.

#### 15.5.3.1.4 Interrupt Mask Register (IMASK)

The interrupt mask register provides control over which possible interrupt events in the IEVENT register are permitted to participate in generating hardware interrupts to the PIC. All implemented bits in this

register are R/W and cleared upon a hardware reset. If the corresponding bits in both the IEVENT and IMASK registers are set, the PIC receives an interrupt (for each eTSEC these are grouped into transmit, receive, and error/diagnostic interrupts). The interrupt signal remains asserted until either the IEVENT bit is cleared, by writing a 1 to it, or by writing a 0 to the corresponding IMASK bit.

Figure 15-5 describes the IMASK register.

Offset eTSEC1:0x2\_4014; eTSEC2:0x2\_5014

Access: Read/Write




Figure 15-5. IMASK Register Definition

Table 15-9 describes the fields of the IMASK register.

Table 15-9. IMASK Field Descriptions

Bits	Name	Description
0	BREN	Babbling receiver interrupt enable
1	RXCEN	Receive control interrupt enable
2	BSYEN	Busy interrupt enable
3	EBERREN	Ethernet controller bus error enable
4	—	Reserved
5	MSROEN	MIB counter overflow interrupt enable
6	GTSCEN	Graceful transmit stop complete interrupt enable
7	BTEN	Babbling transmitter interrupt enable
8	TXCEN	Transmit control interrupt enable
9	TXEEN	Transmit error interrupt enable

Table 15-9. IMASK Field Descriptions (continued)

Bits	Name	Description
10	TXBEN	Transmit buffer interrupt enable
11	TXFEN	Transmit frame interrupt enable
12	—	Reserved
13	LCEN	Late collision enable
14	CRLEN	Collision retry limit enable
15	XFUNEN	Transmit FIFO underrun enable
16	RXBEN	Receive buffer interrupt enable
17–19	—	Reserved
20	MAGEN	Magic packet received interrupt enable
21	MMRDEN	MII management read completion interrupt enable
22	MMWREN	MII management write completion interrupt enable
23	GRSCEN	Graceful receive stop complete interrupt enable
24	RXFEN	Receive frame interrupt enable
25–26	—	Reserved
27	FGPIEN	Filer general purpose interrupt enable
28	FIREN	Filer invalid result interrupt enable
29	FIQEN	Filed frame to invalid queue interrupt enable
30	DPEEN	Data parity error interrupt enable
31	PERREN	Receive frame parse error enable

### 15.5.3.1.5 Error Disabled Register (EDIS)

Figure 15-6 describes the definition for the EDIS register. The error disabled register allows the user to disable an error interruption, possibly to avoid spurious error indications external to the eTSECs.

Offset eTSEC1:0x2\_4018; eTSEC2:0x2\_5018

Access: Read/Write




Figure 15-6. EDIS Register Definition

Table 15-10 describes the fields of the EDIS register.

**Table 15-10. EDIS Field Descriptions**

Bits	Name	Description
0–1	—	Reserved
2	BSYDIS	Busy disable. 0 Allow eTSEC to report IEVENT[BSY] status and halt buffer descriptor queue if BSY condition occurs. 1 Do not set IEVENT[BSY] and do not halt buffer descriptor queue if BSY condition occurs.
3	EBERRDIS	Ethernet controller bus error disable. 0 Allow eTSEC to report IEVENT[EBERR] status and halt buffer descriptor queue if EBERR condition occurs. 1 Do not set IEVENT[EBERR] and do not halt buffer descriptor queue if EBERR condition occurs.
4–6	—	Reserved
7	BABTDIS	Babbling transmit error disable. 0 Allow eTSEC to report IEVENT[BABT] status and set the buffer descriptor TR field. 1 Do not set IEVENT[BABT] nor the buffer descriptor TR field.
8	—	Reserved
9	TXEDIS	Transmit error disable. 0 Allow eTSEC to report IEVENT[TXE] status. 1 Do not set IEVENT[TXE] if TXE condition occurs.
10–12	—	Reserved
13	LCDIS	Late collision disable. 0 Allow eTSEC to report IEVENT[LC] status, set the buffer descriptor LC field, and halt buffer descriptor queue if LC condition occurs. 1 Do not set IEVENT[LC] nor the buffer descriptor LC field, and do not halt buffer descriptor queue if LC condition occurs.
14	CRLDIS	Collision retry limit disable. 0 Allow eTSEC to report IEVENT[CRL] status, set the buffer descriptor RL field, and halt buffer descriptor queue if CRL condition occurs. 1 Do not set IEVENT[CRL] nor the buffer descriptor RL field, and do not halt buffer descriptor queue if CRL condition occurs.
15	XFUNDIS	Transmit FIFO underrun disable. 0 Allow eTSEC to report IEVENT[XFUN] status, set the buffer descriptor UN field, and halt buffer descriptor queue if XFUN condition occurs. 1 Do not set IEVENT[XFUN] nor the buffer descriptor UN field, and do not halt buffer descriptor queue if XFUN condition occurs.
16–27	—	Reserved
28	FIRDIS	Filer invalid result error disable. 0 Allow eTSEC to report IEVENT[FIR] status. 1 Do not set IEVENT[FIR] if eTSEC fails to reach a definite filer result when attempting to file a received frame, but discard the frame silently.
29	FIQDIS	Filed frame to invalid queue error disable. 0 Allow eTSEC to report IEVENT[FIQ] status. 1 Do not set IEVENT[FIQ] if eTSEC attempts to file a received frame to an invalid (disabled) RxBD ring, but discard the frame silently.

**Table 15-10. EDIS Field Descriptions (continued)**

Bits	Name	Description
30	DPEDIS	Data parity error disable. 0 Allow eTSEC to report IEVENT[DPE] status. 1 Do not set IEVENT[DPE] if a parity error occurs in eTSEC's FIFO or filer arrays.
31	PERRDIS	Receive frame parse error disable. 0 Allow eTSEC to report IEVENT[PERR] status. 1 Do not set IEVENT[PERR] if a parse error occurs on a received frame.


### 15.5.3.1.6 Ethernet Control Register (ECNTRL)

ECNTRL is a register writable by the user to reset, configure, and initialize the eTSEC. Note that the FIFM, GMIIM, RPM, and RMM fields are read-only, having been set after sampling signals at power-on-reset. (Refer to the TSEC mode in [Section 4.3.2.2, “Reset Configuration Word High Register \(RCWHR\)”](#)).

[Figure 15-7](#) describes the definition for the ECNTRL register.

Offset eTSEC1:0x2\_4020; eTSEC2:0x2\_5020

Access: Mixed

**Figure 15-7. ECNTRL Register Definition**

[Table 15-11](#) describes the fields of the ECNTRL register.

**Table 15-11. ECNTRL Field Descriptions**

Bits	Name	Description
0–15	—	Reserved
16	FIFM	FIFO mode. Not supported.
17	CLRCNT	Clear all statistics counters 0 Allow MIB counters to continue to increment. 1 Reset all MIB counters. This bit is self-resetting.
18	AUTOZ	Automatically zero MIB counter values. 0 The user must write the addressed counter zero after a host read. 1 The addressed counter value is automatically cleared to zero after a host read. This is a steady state signal and must be set prior to enabling the Ethernet controller and must not be changed without proper care.

**Table 15-11. ECNTRL Field Descriptions (continued)**

Bits	Name	Description
19	STEN	MIB counter statistics enabled. 0 Statistics not enabled 1 Enables internal counters to update This is a steady state signal and must be set prior to enabling the Ethernet controller and must not be changed without proper care.
20–24	—	Reserved
25	GMIIM	GMII interface mode. If this bit is set, a PHY with a RGMII interface is expected to be connected. If cleared, a PHY with an MII or RMII interface is expected. The user should then set MACCFG2[I/F Mode] accordingly. The state of this status bit is defined during power-on reset. 0 MII or RMII mode interface expected 1 RGMII mode interface expected
27	RPM	Reduced-pin mode for Gigabit interfaces. If this bit is set, a reduced-pin interface is expected on Ethernet interfaces. RPM and RMM are never set together. This register can be pin-configured at reset to 0 or 1. 0 MII in non-reduced-pin mode configuration 1 RGMII or RTBI reduced-pin mode
28	R100M	RGMII/RMII 100 mode. This bit is ignored unless SGMIIIM, RPM or RMM are set and MACCFG2[I/F Mode] is assigned to 10/100 (01). 0 RGMII is in 10 Mbps mode; RMII is in 10 Mbps mode, and every 10th RMII Reference clock is used to transfer data SGMII is in 10 Mbps mode, and every 100th SGMII Reference clock is used to transfer data 1 RGMII is in 100 Mbps mode; RMII is in 100 Mbps mode, and data is transferred on every Reference clock SGMII is in 100 Mbps mode, and every 10th SGMII Reference clock is used to transfer data This bit must be cleared for 1-Gbps SGMII operation.
29	RMM	Reduced-pin mode for 10/100 interfaces. If this bit is set, an RMII pin interface is expected. RMM must be 0 if RPM = 1. This register can be pin-configured at reset to 0 or 1. 0 Non-RMII interface mode 1 RMII interface mode
30	SGMIIM	Serial GMII mode. If this bit is set, a SGMII pin interface is expected to be connected via an on chip SERDES. This register can be pin-configured at reset to 0 or 1. 0 SGMII mode disabled. eTSEC connected via a parallel interface. 1 SGMII mode enabled.
31	—	Reserved

The different interface configurations indicated by registers ECNTRL and MACCFG2 are summarized in [Table 15-12](#).

**Table 15-12. eTSEC Interface Configurations**

Interface Mode	ECNTRL Field							MACCFG2 Field
	FIFM <sup>1</sup>	GMIIM	TBIM	RPM	R100M	RMM	SGMIIM	I/F Mode
RTBI 1Gbps	0	0	1	1	—	—	0	10
RGMII 1Gbps	0	1	0	1	—	—	0	10

Table 15-12. eTSEC Interface Configurations (continued)

Interface Mode	ECNTRL Field							MACCFG2 Field
	FIFM <sup>1</sup>	GMIIM	TBIM	RPM	R100M	RMM	SGMIIM	I/F Mode
RGMI 100 Mbps	0	1	0	1	1	—	0	01
RGMI 10 Mbps	0	1	0	1	0	—	0	01
MII 10/100 Mbps	0	0	0	0	—	0	0	01
RMII 100 Mbps	0	0	0	0	1	1	0	01
RMII 10 Mbps	0	0	0	0	0	1	0	01
SGMI 1 Gbps	0	0	1	0	—	—	1	10
SGMI 100 Mbps	0	0	1	0	1	—	1	01
SGMI 10 Mbps	0	0	1	0	0	—	1	01

<sup>1</sup> FIFO mode not supported.

### 15.5.3.1.7 Pause Time Value Register (PTV)

PTV is a 32-bit register written by the user to store the pause duration used when the eTSEC initiates an IEEE 802.3 PAUSE control frame through TCTRL[TFC\_PAUSE]. The low-order 16 bits (PT) represent the pause time and the high-order 16 bits (PTE) represent the extended pause control parameter. The pause time is measured in units of *pause\_quanta*, equal to 512 bit times. The pause time can range from 0 to 65,535 *pause\_quanta*, or 0 to 33,553,920 bit times. See Section 15.6.2.9, “Flow Control,” for additional details. Figure 15-8 describes the definition for the PTV register.

Offset eTSEC1:0x2\_4028; eTSEC2:0x2\_5028

Access: Read/Write




Figure 15-8. PTV Register Definition


Table 15-13 describes the fields of the PTV register.

Table 15-13. PTV Field Descriptions

Bits	Name	Description
0–15	PTE	Extended pause control. This field allows software to add a 16-bit additional control parameter into the PAUSE frame to be sent when TCTRL[TFC_PAUSE] is set. Note that current IEEE 802.3 PAUSE frame format requires this parameter to be cleared.
16–31	PT	Pause time value. Represents the 16-bit pause quanta (that is, 512 bit times). This pause value is used as part of the PAUSE frame to be sent when TCTRL[TFC_PAUSE] is set. See Section 15.6.2.9, “Flow Control,” on page 15-154 for more information.

### 15.5.3.1.8 DMA Control Register (DMACTRL)

DMACTRL is writable by the user to configure the DMA block. Figure 15-9 describes the definition for the DMACTRL register.



**Figure 15-9. DMACTRL Register**

Table 15-14 describes the fields of the DMACTRL register.

**Table 15-14. DMACTRL Field Descriptions**

Bits	Name	Description
0–15	—	Reserved
16	LE	Little-endian descriptor mode enable. This bit controls both the reading and writing of descriptors; data buffers are always transferred in network byte order. 0 RxBDs and TxBDs are interpreted with big-endian byte ordering, as shown in Section 15.6.7.1, “Data Buffer Descriptors.” 1 RxBDs and TxBDs are interpreted with little-endian byte ordering. That is, the 16 bits of flags are considered a complete half-word unit, the buffer length is considered another complete half-word unit, and the buffer pointer is considered a complete word unit.
17–23	—	Reserved
24	TDSEN	Tx Data snoop enable. 0 Disables snooping of all transmit frames from memory. 1 Enables snooping of all transmit frames from memory.
25	TBDSSEN	TxBD snoop enable. 0 Disables snooping of all transmit BD memory accesses. 1 Enables snooping of all transmit BD memory accesses.
26	—	Reserved
27	GRS	Graceful receive stop. If this bit is set, the Ethernet controller stops receiving frames following completion of the frame currently being received. (That is, after a valid end of frame was received). The contents of the Rx FIFO are then written to memory, and the IEVENT[GRSC] is set to indicate that all current receive buffers have been closed. Because the receive enable bit of the MAC may still be set, the MAC may continue to receive but the eTSEC ignores the receive data until GRS is cleared. If this bit is cleared, the eTSEC scans the input data stream for the start of a new frame (preamble sequence and start of frame delimiter) and the first valid frame received uses the next RxBd. If GRS is set, the user must monitor the graceful receive stop complete (GRSC) bit in the IEVENT register to insure that the graceful receive stop was completed. The user can then clear IEVENT[GRSC] and can write to receive registers that are accessible to both user and the eTSEC hardware without fear of conflict. 0 eTSEC scans input data stream for valid frame. 1 eTSEC stops receiving frames following completion of current frame.



**Table 15-14. DMACTRL Field Descriptions (continued)**

Bits	Name	Description
28	GTS	Graceful transmit stop. If this bit is set, the Ethernet controller stops transmission after all frames that are currently in the Tx FIFO or scheduled have been transmitted, and the GTSC interrupt in the IEVENT register is asserted. A frame that has started reading buffer descriptors or data from memory is read to completion and transmitted before the GTSC interrupt occurs. However, if no frame has been scheduled for transmission and the Tx FIFO is empty, the GTSC interrupt is asserted immediately. Once transmission has completed, clearing GTS “restart” transmit. 0 Controller continues. 1 Controller stops transmission after completion of current frame.
29	TOD	Transmit on demand for TxBD ring 0. This bit is applicable only to the transmitter, and requires both TCTRL[TXSCHED] = 00 and DMACTRL[WOP] = 0. If 1 is written to this bit, the eTSEC immediately begins fetching the next TxBD from ring 0, avoiding waiting the normal polling time to check the TxBD's R bit. This bit is always read as 0. 0 eTSEC continues waiting for the TxBD ring 0 poll timer to expire. 1 eTSEC immediately fetches a new TxBD from ring 0, and resets the poll timer.
30	WWR	Write with response. This bit gives the user the assurance that a BD was updated in memory before it receives an interrupt concerning a transmit or receive frame. 0 Do not wait for acknowledgement from system for BD writes before setting IEVENT bits. 1 Before setting IEVENT bits TXB, TXF, TXE, XFUN, LC, CRL, RXB, RXF, the eTSEC waits for acknowledgement from system that the transmit or receive BD being updated was stored in memory.
31	WOP	Wait or poll for TxBD ring 0. This bit, which is applicable only to the transmitter and when TCTRL[TXSCHED] = 00, provides the user the option for the eTSEC to periodically poll TxBDs or to wait for software to tell eTSEC to fetch a buffer descriptor. While operating in the “Wait” mode, the eTSEC allows two additional reads of a descriptor which is not ready before entering a halt state. No interrupt is driven. To resume transmission, software must clear TSTAT[THLT]. 0 Poll TxBD on ring 0 every 512 serial clocks. 1 Do not poll, but wait for TSTAT[THLT] to be cleared by the user.

### 15.5.3.2 eTSEC Transmit Control and Status Registers


This section describes the control and status registers that are used specifically for transmitting Ethernet frames. All of the registers are 32 bits wide.

#### 15.5.3.2.1 Transmit Control Register (TCTRL)

This register is writable by the user to configure the transmit block. [Figure 15-10](#) describes the TCTRL register.

Offset eTSEC1:0x2\_4100; eTSEC2:0x2\_5100

Access: Mixed



**Figure 15-10. TCTRL Register Definition**

Table 15-15 describes the fields of the TCTRL register.

**Table 15-15. TCTRL Field Descriptions**

Bits	Name	Description
0–16	—	Reserved
17	IPCSSEN	IP header checksum generation enable. When set, the eTSEC offloads IPv4 header checksum generation. See <a href="#">Section 15.6.3.2, “Transmit Path Off-Load and Tx PTP Packet Parsing,” on page 15-161</a> . 0 IP header checksum generation is disabled even if enabled in a transmit frame control block. 1 IP header checksum generation is performed for IPv4 headers as determined by the settings in the current transmit frame control block.
18	TUCSEN	TCP/UDP header checksum generation enable. When set, the eTSEC offloads TCP or UDP header checksum generation. See <a href="#">Section 15.6.3.2, “Transmit Path Off-Load and Tx PTP Packet Parsing,” on page 15-161</a> . 0 TCP or UDP header checksum generation is disabled even if enabled in a transmit frame control block. 1 TCP or UDP header checksum generation is performed as determined by the settings in the current transmit frame control block.
19	VLINS	VLAN (IEEE Std. 802.1Q) tag insertion enable. Applicable only for transmission through the Ethernet MAC. 0 Do not insert a VLAN tag into the frame. 1 Insert a VLAN tag into the frame. If the frame FCB has a valid VLAN field, use the FCB to source the VLAN control word, otherwise take the default VLAN control word from register DFVLAN.
20	THDF	Transmit half-duplex flow control under software control for 10-/100-Mbps half-duplex media. This bit is not self-resetting. 0 Disable back pressure 1 Back pressure is applied to media by raising carrier
21–26	—	Reserved
27	RFC_PAUSE	Receive flow control pause frame (written by the eTSEC). This read-only status bit is set if a flow control pause frame was received and the transmitter is paused for the duration defined in the received pause frame. This bit automatically clears after the pause duration is complete. 0 Pause duration complete. 1 Flow control pause frame received.
28	TFC_PAUSE	Transmit flow control pause frame. Set this bit to transmit a PAUSE frame. If this bit is set, the MAC stops transmission of data frames after the currently transmitting frame completes. Next, the MAC transmits a pause control frame with the duration value obtained from the PTV register. The TXC event occurs after sending the pause control frame. Finally, the controller clears TFC_PAUSE and resumes transmitting data frames as before. Note that pause control frames can still be transmitted if the Tx controller is stopped due to user assertion of DMACTRL[GTS] or reception of a PAUSE frame. 0 No request for Tx PAUSE frame pending or transmission complete. 1 Software request for Tx PAUSE frame pending.

Table 15-15. TCTRL Field Descriptions (continued)

Bits	Name	Description
29–30	TXSCHED	<p>Transmit ring scheduling algorithm. This field determines which scheme the transmit scheduler uses to arbitrate between the enabled TxBD rings. The scheme chosen also controls how the DMACTRL and TQUEUE bits are interpreted. Ring polling is supported only by mode 00; the other modes require software to restart rings with the TSTAT register. TCP/IP offload can be enabled with any scheduling mode.</p> <p>00 Single polled ring mode. TxBD ring 0 is the only ring serviced, even if other rings are enabled and ready. In this scheduler mode, the DMACTRL[WOP] and DMACTRL[TOD] bits control polling and retry behavior. This mode supports ring polling, and allows fetching of a non-ready TxBD to be retried twice.</p> <p>01 Priority scheduling mode. All enabled TxBD rings are serviced in ascending ring index order. Once a non-ready TxBD has been fetched from the lowest-numbered ring, the eTSEC attempts to fetch TxBDs from the next enabled ring having a higher index, until transmission stops for lack of data. TSTAT records whenever a TxBD ring is exhausted.</p> <p>10 Modified weighted round-robin scheduling mode. Each TxBD ring is polled in sequence for frames that are ready for transmission. If a non-ready TxBD is fetched from a ring, that ring is removed from the scheduling pool until software re-enables it. Ready frames are repeatedly transmitted from a chosen ring until its transmission quota is exhausted. The transmission quota for TxBD ring <math>n</math> is set to <math>WT_n \times 64</math> bytes, where <math>WT_n</math> is a weight from the TR03WT/TR47WT registers. If a ring transmits more data than its quota allows, the excess is deducted from its quota on the next transmission opportunity, thereby preventing large frames from monopolizing the eTSEC bandwidth.</p> <p>11 Reserved</p>
31	—	Reserved

### 15.5.3.2.2 Transmit Status Register (TSTAT)

This register is read/write-one-to-clear and is written by the eTSEC to convey DMA status information for each TxBD ring. The halt bit only has meaning for enabled rings. After processing transmit-related interrupts, software should use TSTAT to restart transmission from rings that may have been affected by the interrupt condition. In particular, an error condition that prevents eTSEC from continuing transmission halts DMA from all rings, including the ring that gave rise to the error. [Figure 15-11](#) describes the TSTAT register.

Offset eTSEC1:0x2\_4104; eTSEC2:0x2\_5104

Access: w1c

	0	1	2	3	4	5	6	7	8	15	
R	THLT0	THLT1	THLT2	THLT3	THLT4	THLT5	THLT6	THLT7	—		
W	w1c	w1c	w1c	w1c	w1c	w1c	w1c	w1c	—		
Reset	All zeros										
	16	17	18	19	20	21	22	23	24	31	
R	TXF0	TXF1	TXF2	TXF3	TXF4	TXF5	TXF6	TXF7	—		
W	w1c	w1c	w1c	w1c	w1c	w1c	w1c	w1c	—		
Reset	All zeros										

Figure 15-11. TSTAT Register Definition

Table 15-16 describes the fields of the TSTAT register.

**Table 15-16. TSTAT Field Descriptions**

Bits	Name	Description
0	THLT0	<p>Transmit halt of ring 0. Set by the eTSEC if is no longer processing transmit frames from this TxBD ring, and DMA from this ring is disabled. To re-start transmission from this TxBD ring, this bit must be cleared by writing 1 to it. This bit is set only on a general error condition (as in IEVENT[TXE]), regardless of TQUEUE[EN0], or if no ready TxBDs can be fetched. DMACTRL[GTS] being set by the user does not cause this bit to be set. Software should examine the halted queue's buffer descriptors for repeatable error conditions before taking it out of the halt state. Failure to do so may cause an effective livelock, in which the error condition recurs and halts all queues again.</p> <p>Repeatable error conditions which cause halt include: Bus error:</p> <ul style="list-style-type: none"> <li>• Invalid BD or data address</li> <li>• Uncorrectable error on BD or data read</li> </ul> <p>TxBD programming errors:</p> <ul style="list-style-type: none"> <li>• Ready=1 and length=0</li> </ul>
1	THLT1	<p>Transmit halt of ring 1. Set by the eTSEC if is no longer processing transmit frames from this TxBD ring, and DMA from this ring is disabled. To re-start transmission from this TxBD ring, this bit must be cleared by writing 1 to it. This bit is set only on a general error condition (as in IEVENT[TXE]), regardless of TQUEUE[EN1], or if no ready TxBDs can be fetched. DMACTRL[GTS] being set by the user does not cause this bit to be set.</p> <p>Software should examine the halted queue's buffer descriptors for repeatable error conditions before taking it out of the halt state. Failure to do so may cause an effective livelock, in which the error condition recurs and halts all queues again.</p> <p>Repeatable error conditions which cause halt include: Bus error:</p> <ul style="list-style-type: none"> <li>• Invalid BD or data address</li> <li>• Uncorrectable error on BD or data read</li> </ul> <p>TxBD programming errors:</p> <ul style="list-style-type: none"> <li>• Ready=1 and length=0</li> </ul>
2	THLT2	<p>Transmit halt of ring 2. Set by the eTSEC if is no longer processing transmit frames from this TxBD ring, and DMA from this ring is disabled. To re-start transmission from this TxBD ring, this bit must be cleared by writing 1 to it. This bit is set only on a general error condition (as in IEVENT[TXE]), regardless of TQUEUE[EN2], or if no ready TxBDs can be fetched. DMACTRL[GTS] being set by the user does not cause this bit to be set.</p> <p>Software should examine the halted queue's buffer descriptors for repeatable error conditions before taking it out of the halt state. Failure to do so may cause an effective livelock, in which the error condition recurs and halts all queues again.</p> <p>Repeatable error conditions which cause halt include: Bus error:</p> <ul style="list-style-type: none"> <li>• Invalid BD or data address</li> <li>• Uncorrectable error on BD or data read</li> </ul> <p>TxBD programming errors:</p> <ul style="list-style-type: none"> <li>• Ready=1 and length=0</li> </ul>

Table 15-16. TSTAT Field Descriptions (continued)

Bits	Name	Description
3	THLT3	<p>Transmit halt of ring 3. Set by the eTSEC if is no longer processing transmit frames from this TxBD ring, and DMA from this ring is disabled. To re-start transmission from this TxBD ring, this bit must be cleared by writing 1 to it. This bit is set only on a general error condition (as in IEVENT[TXE]), regardless of TQUEUE[EN3], or if no ready TxBDs can be fetched. DMACTRL[GTS] being set by the user does not cause this bit to be set.</p> <p>Software should examine the halted queue's buffer descriptors for repeatable error conditions before taking it out of the halt state. Failure to do so may cause an effective livelock, in which the error condition recurs and halts all queues again.</p> <p>Repeatable error conditions which cause halt include: Bus error:</p> <ul style="list-style-type: none"> <li>• Invalid BD or data address</li> <li>• Uncorrectable error on BD or data read</li> </ul> <p>TxBD programming errors:</p> <ul style="list-style-type: none"> <li>• Ready=1 and length=0</li> </ul>
4	THLT4	<p>Transmit halt of ring 4. Set by the eTSEC if is no longer processing transmit frames from this TxBD ring, and DMA from this ring is disabled. To re-start transmission from this TxBD ring, this bit must be cleared by writing 1 to it. This bit is set only on a general error condition (as in IEVENT[TXE]), regardless of TQUEUE[EN4], or if no ready TxBDs can be fetched. DMACTRL[GTS] being set by the user does not cause this bit to be set.</p> <p>Software should examine the halted queue's buffer descriptors for repeatable error conditions before taking it out of the halt state. Failure to do so may cause an effective livelock, in which the error condition recurs and halts all queues again.</p> <p>Repeatable error conditions which cause halt include: Bus error:</p> <ul style="list-style-type: none"> <li>• Invalid BD or data address</li> <li>• Uncorrectable error on BD or data read</li> </ul> <p>TxBD programming errors:</p> <ul style="list-style-type: none"> <li>• Ready=1 and length=0</li> </ul>
5	THLT5	<p>Transmit halt of ring 5. Set by the eTSEC if is no longer processing transmit frames from this TxBD ring, and DMA from this ring is disabled. To re-start transmission from this TxBD ring, this bit must be cleared by writing 1 to it. This bit is set only on a general error condition (as in IEVENT[TXE]), regardless of TQUEUE[EN5], or if no ready TxBDs can be fetched. DMACTRL[GTS] being set by the user does not cause this bit to be set.</p> <p>Software should examine the halted queue's buffer descriptors for repeatable error conditions before taking it out of the halt state. Failure to do so may cause an effective livelock, in which the error condition recurs and halts all queues again.</p> <p>Repeatable error conditions which cause halt include: Bus error:</p> <ul style="list-style-type: none"> <li>• Invalid BD or data address</li> <li>• Uncorrectable error on BD or data read</li> </ul> <p>TxBD programming errors:</p> <ul style="list-style-type: none"> <li>• Ready=1 and length=0</li> </ul>

Table 15-16. TSTAT Field Descriptions (continued)

Bits	Name	Description
6	THLT6	<p>Transmit halt of ring 6. Set by the eTSEC if is no longer processing transmit frames from this TxBD ring, and DMA from this ring is disabled. To re-start transmission from this TxBD ring, this bit must be cleared by writing 1 to it. This bit is set only on a general error condition (as in IEVENT[TXE]), regardless of TQUEUE[EN6], or if no ready TxBDs can be fetched. DMACTRL[GTS] being set by the user does not cause this bit to be set.</p> <p>Software should examine the halted queue's buffer descriptors for repeatable error conditions before taking it out of the halt state. Failure to do so may cause an effective livelock, in which the error condition recurs and halts all queues again.</p> <p>Repeatable error conditions which cause halt include: Bus error:</p> <ul style="list-style-type: none"> <li>• Invalid BD or data address</li> <li>• Uncorrectable error on BD or data read</li> </ul> <p>TxBD programming errors:</p> <ul style="list-style-type: none"> <li>• Ready=1 and length=0</li> </ul>
7	THLT7	<p>Transmit halt of ring 7. Set by the eTSEC if is no longer processing transmit frames from this TxBD ring, and DMA from this ring is disabled. To re-start transmission from this TxBD ring, this bit must be cleared by writing 1 to it. This bit is set only on a general error condition (as in IEVENT[TXE]), regardless of TQUEUE[EN7], or if no ready TxBDs can be fetched. DMACTRL[GTS] being set by the user does not cause this bit to be set.</p> <p>Software should examine the halted queue's buffer descriptors for repeatable error conditions before taking it out of the halt state. Failure to do so may cause an effective livelock, in which the error condition recurs and halts all queues again.</p> <p>Repeatable error conditions which cause halt include: Bus error:</p> <ul style="list-style-type: none"> <li>• Invalid BD or data address</li> <li>• Uncorrectable error on BD or data read</li> </ul> <p>TxBD programming errors:</p> <ul style="list-style-type: none"> <li>• Ready=1 and length=0</li> </ul>
8–15	—	Reserved
16	TXF0	Transmit frame event occurred on ring 0. Set by the eTSEC if IEVENT[TXF] was set in relation to transmitting a frame from this ring.
17	TXF1	Transmit frame event occurred on ring 1. Set by the eTSEC if IEVENT[TXF] was set in relation to transmitting a frame from this ring.
18	TXF2	Transmit frame event occurred on ring 2. Set by the eTSEC if IEVENT[TXF] was set in relation to transmitting a frame from this ring.
19	TXF3	Transmit frame event occurred on ring 3. Set by the eTSEC if IEVENT[TXF] was set in relation to transmitting a frame from this ring.
20	TXF4	Transmit frame event occurred on ring 4. Set by the eTSEC if IEVENT[TXF] was set in relation to transmitting a frame from this ring.
21	TXF5	Transmit frame event occurred on ring 5. Set by the eTSEC if IEVENT[TXF] was set in relation to transmitting a frame from this ring.
22	TXF6	Transmit frame event occurred on ring 6. Set by the eTSEC if IEVENT[TXF] was set in relation to transmitting a frame from this ring.

**Table 15-16. TSTAT Field Descriptions (continued)**

Bits	Name	Description
23	TXF7	Transmit frame event occurred on ring 7. Set by the eTSEC if IEVENT[TXF] was set in relation to transmitting a frame from this ring.
24–31	—	Reserved

### 15.5.3.2.3 Default VLAN Control Word Register (DFVLAN)

This register defines the default value for the VLAN Ethertype and control word when VLAN tags are automatically inserted by the eTSEC, and no per-frame VLAN data is supplied by software. On receive, this register defines a customizable VLAN Ethertype for automatic deletion. Note that an Ethertype of 0x8808 (Control Word) is not permitted as a custom VLAN tag. Frames with an Ethertype of 0x8808 are dropped by the receiver. In the case of frames containing stacked VLAN tags, this register defines the tag associated with the outer or metropolitan area VLAN. Figure 15-12 describes the DFVLAN register.


**Figure 15-12. DFVLAN Register Definition**


Table 15-17 describes the fields of the DFVLAN register.

**Table 15-17. DFVLAN Field Descriptions**

Bits	Name	Description
0–15	TAG	This is the default Ethertype used to tag VLAN frames. On transmit, this tag is inserted ahead of the VLAN control word; TAG should be set to 0x8100 for IEEE 802.1Q VLAN. On receive, an Ethertype matching TAG or an Ethertype of 0x8100 marks a VLAN-tagged frame. Note that if using DFVLAN to set a custom ethertype (that is, using a value other than 0x8100), packets received with a custom tag are not counted by any of the RMON counters. Affected counters include TRMGV, RMCA, RBCA, RXCF, RXPF, RXUO, RALN, RFLR, ROVR, RJBR, TMCA, TBCA, TXPF, TXCF.
16–18	PRI	This is the default value used for the IEEE Std. 802.1p frame priority.
19	CFI	This is the default value used for the IEEE Std. 802.1Q canonical format indicator.
20–31	VID	This is the default value used for the virtual-LAN identifier in VLAN-tagged frames. A value of zero is defined as the null VLAN, however field PRI may be still set independently.

### 15.5.3.2.4 Transmit Interrupt Coalescing Register (TXIC)

The TXIC register enables and configures the operational parameters for interrupt coalescing associated with transmitted frames. Figure 15-13 describes the definition for the TXIC register.



**Figure 15-13. TXIC Register Definition**

Table 15-18 describes the fields of the TXIC register.

**Table 15-18. TXIC Field Descriptions**

Bits	Name	Description
0	ICEN	Interrupt coalescing enable 0 Interrupt coalescing is disabled. Interrupts are raised as they are received. 1 Interrupt coalescing is enabled. If the eTSEC transmit frame interrupt is enabled (IMASK[TXFEN] is set), an interrupt is raised when the threshold number of frames is reached (defined by TXIC[ICFT]) or when the threshold timer expires (determined by TXIC[ICTT]).
1	ICCS	Interrupt coalescing timer clock source. 0 The coalescing timer advances count every 64 eTSEC Tx interface clocks (TSEC <sub>n</sub> _GTX_CLK). 1 The coalescing timer advances count every 64 system clocks. This mode is recommended for FIFO operation.
2	—	Reserved
3–10	ICFT	Interrupt coalescing frame count threshold. While interrupt coalescing is enabled (TXIC[ICEN] is set), this value determines how many frames are transmitted before raising an interrupt. The eTSEC threshold counter is reset to ICFT following an interrupt. The value of ICFT must be greater than zero to avoid unpredictable behavior.
11–15	—	Reserved
16–31	ICTT	Interrupt coalescing timer threshold. While interrupt coalescing is enabled (TXIC[ICEN] is set), this value determines the maximum amount of time after transmitting a frame before raising an interrupt. If frames have been transmitted but the frame count threshold has not been met, an interrupt is raised when the threshold timer reaches zero. The threshold timer is reset to the value in this field and begins counting down upon transmission of the first frame having its TxBD[ <i>i</i> ] bit set. The threshold value is represented in units of 64 clock periods as specified by the timer clock source (TXIC[ICCS]). The value of ICTT must be greater than zero to avoid unpredictable behavior.



### 15.5.3.2.5 Transmit Queue Control Register (TQUEUE)

The TQUEUE register, shown in Figure 15-14, selectively enables each of the TxBD rings 0–7. By default, TxBD ring 0 is enabled.




Figure 15-14. TQUEUE Register Definition

Table 15-19 describes the TQUEUE register.


Table 15-19. TQUEUE Field Descriptions

Bits	Name	Description
0–15	—	Reserved
16	EN0	Transmit queue 0 enable. 0 TxBD ring is not queried for transmission. In effect the transmit queue is disabled. 1 TxBD ring is queried for transmission.
17	EN1	Transmit queue 1 enable. 0 TxBD ring is not queried for transmission. In effect the transmit queue is disabled. 1 TxBD ring is queried for transmission.
18	EN2	Transmit queue 2 enable. 0 TxBD ring is not queried for transmission. In effect the transmit queue is disabled. 1 TxBD ring is queried for transmission.
19	EN3	Transmit queue 3 enable. 0 TxBD ring is not queried for transmission. In effect the transmit queue is disabled. 1 TxBD ring is queried for transmission.
20	EN4	Transmit queue 4 enable. 0 TxBD ring is not queried for transmission. In effect the transmit queue is disabled. 1 TxBD ring is queried for transmission.
21	EN5	Transmit queue 5 enable. 0 TxBD ring is not queried for transmission. In effect the transmit queue is disabled. 1 TxBD ring is queried for transmission.
22	EN6	Transmit queue 6 enable. 0 TxBD ring is not queried for transmission. In effect the transmit queue is disabled. 1 TxBD ring is queried for transmission.
23	EN7	Transmit queue 7 enable. 0 TxBD ring is not queried for transmission. In effect the transmit queue is disabled. 1 TxBD ring is queried for transmission.
24–31	—	Reserved

### 15.5.3.2.6 TxBD Ring 0–3 Weighting Register (TR03WT)

When modified weighted round-robin Tx scheduling is enabled (TCTRL[TXSCHEM] = 10), this register determines the weighting applied to each transmit queue for queues 0 to 3. For priority-based scheduling,

TR03WT has no effect. A description of how queue weights affect eTSEC's round-robin algorithm appears in [Section 15.6.4.3.2, "Modified Weighted Round-Robin Queuing \(MWRR\)."](#) Figure 15-15 describes the TR03WT register.



**Figure 15-15. TR03WT Register Definition**


Table 15-20 describes the fields of the TR03WT register.

**Table 15-20. TR03WT Field Descriptions**

Bits	Name	Description
0–7	WT0	Weighting value for TxBD ring 0 when TCTRL[TXSCHED] = 10. On each round of the Tx scheduler, a minimum of $WT0 \times 64$ bytes of data are scheduled for transmission from TxBD ring 0. Clearing this field prevents transmission.
8–15	WT1	Weighting value for TxBD ring 1 when TCTRL[TXSCHED] = 10. On each round of the Tx scheduler, a minimum of $WT1 \times 64$ bytes of data are scheduled for transmission from TxBD ring 1. Clearing this field prevents transmission.
16–23	WT2	Weighting value for TxBD ring 2 when TCTRL[TXSCHED] = 10. On each round of the Tx scheduler, a minimum of $WT2 \times 64$ bytes of data are scheduled for transmission from TxBD ring 2. Clearing this field prevents transmission.
24–31	WT3	Weighting value for TxBD ring 3 when TCTRL[TXSCHED] = 10. On each round of the Tx scheduler, a minimum of $WT3 \times 64$ bytes of data are scheduled for transmission from TxBD ring 3. Clearing this field prevents transmission.

### 15.5.3.2.7 TxBD Ring 4–7 Weighting Register (TR47WT)

When modified weighted round-robin Tx scheduling is enabled (TCTRL[TXSCHED] = 10), this register determines the weighting applied to each enabled transmit queue for queues 4 to 7. For priority-based scheduling, TR47WT has no effect. A description of how queue weights affect eTSEC's modified weighted round-robin algorithm appears in [Section 15.6.4.3.2, "Modified Weighted Round-Robin Queuing \(MWRR\)."](#) Figure 15-16 describes the definition for the TR47WT register.



**Figure 15-16. TR47WT Register Definition**


Table 15-21 describes the fields of the TR47WT register.

**Table 15-21. TR47WT Field Descriptions**

Bits	Name	Description
0–7	WT4	Weighting value for TxBD ring 4 when TCTRL[TXSCHEDED] = 10. On each round of the Tx scheduler, a minimum of WT4 × 64 bytes of data are scheduled for transmission from TxBD ring 4. Clearing this field prevents transmission.
8–15	WT5	Weighting value for TxBD ring 5 when TCTRL[TXSCHEDED] = 10. On each round of the Tx scheduler, a minimum of WT5 × 64 bytes of data are scheduled for transmission from TxBD ring 5. Clearing this field prevents transmission.
16–23	WT6	Weighting value for TxBD ring 6 when TCTRL[TXSCHEDED] = 10. On each round of the Tx scheduler, a minimum of WT6 × 64 bytes of data are scheduled for transmission from TxBD ring 6. Clearing this field prevents transmission.
24–31	WT7	Weighting value for TxBD ring 7 when TCTRL[TXSCHEDED] = 10. On each round of the Tx scheduler, a minimum of WT7 × 64 bytes of data are scheduled for transmission from TxBD ring 7. Clearing this field prevents transmission.

### 15.5.3.2.8 Transmit Data Buffer Pointer High Register (TBDBPH)

The TBDBPH register is written by the user with the most significant address bits common to all TxBD buffer addresses, TxBD[Data Buffer Pointer]. As a consequence, all Tx buffers must be placed in a 4 gigabyte segment of memory whose base address is prefixed by the bits in TBDBPH. The TxBD ring itself can reside in a different memory region (based at TBASEH). Figure 15-17 describes the definition for the TBDBPH register.



**Figure 15-17. TBDBPH Register Definition**

Table 15-22 describes the fields of the TBDBPH register.


**Table 15-22. TBDBPH Field Descriptions**

Bits	Name	Description
0–27	—	Reserved
28–31	TBDBPH	Most significant bits common to all data buffer addresses contained in TxBDs. The user must initialize TBDBPH before enabling the eTSEC transmit function.

### 15.5.3.2.9 Transmit Buffer Descriptor Pointers 0–7 (TBPTR0–TBPTR7)

TBPTR0–TBPTR7 each contains the low-order 32 bits of the next transmit buffer descriptor address for their respective TxBD ring. Figure 15-18 describes the TBPTR registers. These registers takes on the value of their ring’s associated TBASE when the TBASE register is written by software. Software must not write TBPTR0–TBPTR7 while eTSEC is actively transmitting frames. However, TBPTR0– TBPTR7 can be modified when the transmitter is disabled or when no Tx buffer is in use (after a GRACEFUL STOP

TRANSMIT command is issued and the frame completes its transmission) in order to change the next TxBD eTSEC transmits.



**Figure 15-18. TBPTR0–TBPTR7 Register Definition**


Table 15-23 describes the fields of the TBPTR<sub>*n*</sub> register.

**Table 15-23. TBPTR<sub>*n*</sub> Field Descriptions**

Bits	Name	Description
0–28	TBPTR <sub><i>n</i></sub>	Current TxBD pointer for TxBD ring <i>n</i> . Points to the current BD being processed or to the next BD the transmitter uses when it is idling. When the end of the TxBD ring is reached, eTSEC initializes TBPTR <sub><i>n</i></sub> to the value in the corresponding TBASEn. The TBPTR register is internally written by the eTSEC's DMA controller during transmission. The pointer increments by eight (bytes) each time a descriptor is closed successfully by the eTSEC. Note that the three least significant bits of this register are read-only and zero. After an error condition, the eTSEC returns TBPTR <sub><i>n</i></sub> to point to the first BD of the frame partially transmitted.
29–31	—	Reserved

#### 15.5.3.2.10 Transmit Descriptor Base Address Registers (TBASE0–TBASE7)

The TBASEn registers are written by the user with the base address of each TxBD ring *n*. Each such value must be divisible by eight, since the three least significant bits always write as 000. Figure 15-19 describes the definition for the TBASEn registers.



**Figure 15-19. TBASE Register Definition**

Table 15-24 describes the fields of the TBASEn registers.

**Table 15-24. TBASE0–TBASE7 Field Descriptions**

Bits	Name	Description
0–28	TBASEn	Transmit base for ring <i>n</i> . TBASE defines the starting location in the memory map for the eTSEC TxBDs. This field must be 8-byte aligned. Together with setting the W (wrap) bit in the last BD, the user can select how many BDs to allocate for the transmit packets. The user must initialize TBASE before enabling the eTSEC transmit function on the associated ring.
29–31	—	Reserved

### 15.5.3.2.11 Transmit Timestamp Identification Register (TMR\_TXTS1–2\_ID)

Transmit timestamp identification register (TMR\_TXTS $_n$ \_ID). This register holds the identification number of the transmitted frame corresponding to the timestamp captured in TMR\_TXTS $_n$ \_H/L. Each time the eTSEC is instructed to capture the timestamp of an outgoing frame via TxFCB[PTP] the associated field in TxFCB[PTP\_ID] is stored in this register, overwriting the previous value.

This register is read only in normal operation. Figure 15-20 describes the definition for the TMR\_TXTS $_n$ \_ID register.




Figure 15-20. TMR\_TXTS $_n$ \_ID Register Definition

Table 15-25 describes the fields of the TMR\_TXTS $_n$ \_ID register.

Table 15-25. TMR\_TXTS $_n$ \_ID Register Field Descriptions

Bits	Name	Description
0–15	—	Reserved
16–31	TXTS_ID	Tx timestamp identification field

### 15.5.3.2.12 Transmit Timestamp Register (TMR\_TXTS1–2\_H/L)

Transmit stamp register (TMR\_TXTS $_n$ \_H/L). This register holds the value of the TMR\_CNT\_H/L when a frame tagged for timestamp capture (via Tx FCB[PTP]) is transmitted. Upon transmission of the start of frame symbol of such a frame, the value in TMR\_CNT\_H/L is copied into TMR\_TXTS $_n$ \_H/L.

This register is read only in normal operation. Figure 15-21 depicts TMR\_TXTS $_n$ \_H/L.




Figure 15-21. TMR\_TXTS $_n$ \_H/L Register Definition

Table 15-26 describes the fields of the TMR\_TXTS $_n$ \_H/L register.

Table 15-26. TMR\_TXTS $_n$ \_H/L Register Field Descriptions

Bits	Name	Description
0–63	TXTS_H/L	Timestamp field of the transmitted PTP packet's start of frame detection.

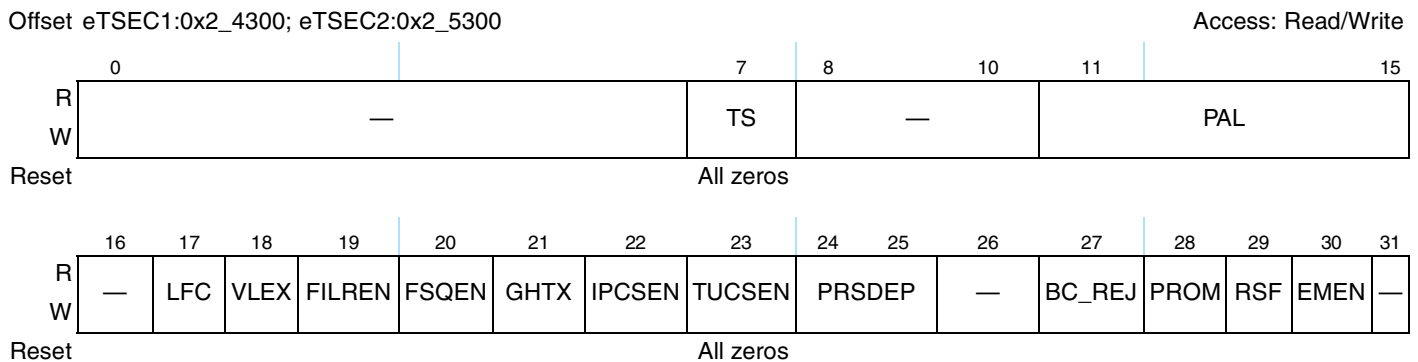
### 15.5.3.3 eTSEC Receive Control and Status Registers

This section describes the control and status registers that are used specifically for receiving Ethernet frames. All of the registers are 32 bits wide.

#### 15.5.3.3.1 Receive Control Register (RCTRL)

The RCTRL register is programmed by the user and controls the operational mode of the receiver. It must be written only after a system reset (at initialization) or after a graceful receive stop has completed.

Figure 15-22 describes the RCTRL register.



**Figure 15-22. RCTRL Register Definition**

Table 15-27 describes the fields of the RCTRL register.

**Table 15-27. RCTRL Field Descriptions**

Bits	Name	Description
0–6	—	Reserved
7	TS	Timestamp incoming packets as padding bytes. PAL field is set to 8 if the PAL field is programmed to less than 8. Must be set to zero if TMR_CTRL[TE]=0.
8–10	—	Reserved
11–15	PAL	Packet alignment padding length. If not zero, PAL (1–31) bytes of zero padding are inserted before the start of each received frame, but following the RxFCB if TOE is enabled. For Ethernet where optional preamble extraction is enabled, the padding appears before the preamble, otherwise the padding precedes the layer 2 header. The value of PAL can be set so that the start of the IP header in the receive data buffer is aligned to a 32-bit boundary. Normally, setting PAL = 2 provides minimal padding to ensure such alignment of the IP header. Note that the minimum zero padding value for this field should be PAL–8 if the TS field is set and 0 when PAL is < 8.
16	—	Reserved

Table 15-27. RCTRL Field Descriptions (continued)

Bits	Name	Description
17	LFC	Lossless flow control. When set, the eTSEC determines the number of free BDs (through RQPARM $n$ [LEN] and RBTPTR $n$ ) in each active ring. Should the free BD count in an active ring drop below its setting for RQPARM $n$ [FBTHR], the eTSEC asserts link layer flow control. For full-duplex ethernet connections, the eTSEC emits a pause frame as if TCTRL[TFC_PAUSE] was set. For FIFO packet interface connections, the RFC signal is asserted. 0 Disabled. This is the default 1 Enabled, calculate the free BDs in each active ring and assert link layer flow control if required.
18	VLEX	Enable automatic VLAN tag extraction and deletion from Ethernet frames. Note that VLEX must be cleared if L2OFF is non-zero. 0 Do not delete VLAN tags from received Ethernet frames. 1 If a VLAN tag is seen after the Ethernet source address, and PRSDEP is non-zero, delete the VLAN tag and return the VLAN control word in the frame control block returned with this frame. Note that if PRSDEP is cleared, VLEX must be cleared as well. (VLAN tag extraction is only supported when the parser is enabled.)
19	FILREN	Filer enable. When set, the receive frame filer is enabled. This file accepted frames to a particular RxBD ring according to rules defined in the filer table. In this case, PRSDEP must not be cleared. 0 Do not search the receive queue filer table for received frames. All received frames are sent to RxBD ring 0 by default. 1 Search the receive queue filer table for received frames, and let the filer determine the index of the RxBD ring for each frame. Note that if PRSDEP is cleared, FILREN must be cleared as well.
20	FSQEN	Enable single-queue mode for the receive frame filer. This bit is ignored unless FILREN is also set. 0 The filer chooses the RxBD ring using the least significant bits of the virtual queue ID as a ring index. 1 The filer always attempts to file received frames to ring 0, regardless of virtual queue ID. This mode is intended for operating the filer as a packet classification engine.
21	GHTX	Group address hash table extend. By default, the group address hash table is 256 entries (as defined by registers GADDR0–GADDR7); registers IGADDR0–IGADDR7 are then used to define the individual address hash table. When this bit is set, the hash table is extended to a total of 512 entries (IGADDR0–IGADDR7 are then the first 256 entries of the extended 512-entry group address hash table). 0 Both the individual and group hash functions are the 8 MSBs of the CRC-32 of the Ethernet destination address. 1 The group hash function is the 9 MSBs of the CRC-32 of the Ethernet destination address. The individual address hash function is unavailable.
22	IPCSEN	IP Checksum verification enable. See <a href="#">Section 15.6.3.3, “Receive Path Off-Load.”</a> 0 IPv4 header checksums are not verified by the eTSEC—even if layer 3 parsing is enabled. 1 Perform IPv4 header checksum verification if PRSDEP > 01.
23	TUCSEN	TCP or UDP Checksum verification enable. See <a href="#">Section 15.6.3.3, “Receive Path Off-Load.”</a> 0 TCP or UDP checksums are not verified by the eTSEC—even if layer 4 parsing is enabled. 1 Perform TCP or UDP checksum verification if PRSDEP = 11.
24–25	PRSDEP	Parser control. The level of parser layer recognition is determined as follows: 00 Parser disabled. Receive frame filer must also be disabled by clearing RCTRL[FILREN]. 01 Only L2 (Ethernet) protocols are recognized. 10 L2 and L3 (IP) protocols are recognized. 11 L2, L3, and L4 (TCP/UDP) protocols are recognized. If this field is non-zero, a TOE frame control block is prepended to the received frame, and the first RxBD points to the FCB. Note that if PRSDEP is cleared, VLEX must be cleared as well. (VLAN tag extraction is only supported when the parser is enabled.) Also, if PRSDEP is cleared, FILREN must also be cleared.

**Table 15-27. RCTRL Field Descriptions (continued)**

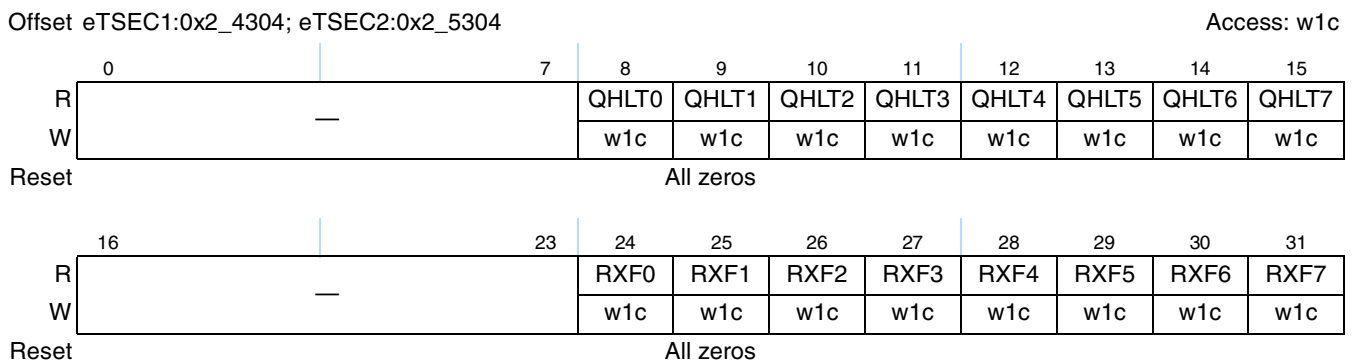
Bits	Name	Description
26	—	Reserved
27	BC_REJ	Broadcast frame reject. If this bit is set, frames with DA (destination address) = FFFF_FFFF_FFFF are rejected unless RCTRL[PROM] is set. If both BC_REJ and RCTRL[PROM] are set, then frames with broadcast DA are accepted and the M (MISS) bit is set in the receive BD.
28	PROM	Promiscuous mode. All Ethernet frames, regardless of destination address, are accepted.
29	RSF	Receive short frame mode. When set, enables the reception of frames shorter than 64 bytes. 0 Ethernet frames less than 64B in length are silently dropped. 1) Frames more than 16B and less than 64B in length are accepted upon a DA match. Note that frames less than or equal to 16B in length are always silently dropped.
30	EMEN	Exact match MAC address enable. If this bit is set, the MAC01ADDR1–MAC15ADDR1 and MAC01ADDR2–MAC15ADDR2 registers are recognized as containing MAC addresses aliasing the MAC’s station address. Setting this bit therefore allows eTSEC to receive Ethernet frames having a destination address matching one of these 15 addresses.
31	—	Reserved

**15.5.3.3.2 Receive Status Register (RSTAT)**

The eTSEC writes to this register under the following conditions:

- A frame interrupt event occurred on one or more RxBD rings
- The receiver runs out of descriptors due to a busy condition on a RxBD ring
- The receiver was halted because an error condition was encountered while receiving a frame

Writing 1 to any bit of this register clears it. Software should clear the QHLT bit to take eTSEC’s receiver function out of halt state for the associated queue. [Figure 15-23](#) describes the definition for the RSTAT register.



**Figure 15-23. RSTAT Register Definition**



Table 15-28 describes the fields of the RSTAT register.

**Table 15-28. RSTAT Field Descriptions**


Bits	Name	Description
0–7	—	Reserved
8	QHLT0	RxBD queue 0 is halted. It is a hardware-initiated stop indication. (DMACTRL[GRS] being set by the user does not cause a QHLT0 to be set.). The current frame and all other frames directed to a halted queue are discarded. A write with a value of 1 re-enables the queue for receiving. 0 This queue is enabled for reception. (That is, it is not halted) 1 All controller receive activity to this queue is halted.
9	QHLT1	RxBD queue 1 is halted. It is a hardware-initiated stop indication. (DMACTRL[GRS] being set by the user does not cause a QHLT1 to be set.). The current frame and all other frames directed to a halted queue are discarded. A write with a value of 1 re-enables the queue for receiving. 0 This queue is enabled for reception. (That is, it is not halted) 1 All controller receive activity to this queue is halted.
10	QHLT2	RxBD queue 2 is halted. It is a hardware-initiated stop indication. (DMACTRL[GRS] being set by the user does not cause a QHLT2 to be set.). The current frame and all other frames directed to a halted queue are discarded. A write with a value of 1 re-enables the queue for receiving. 0 This queue is enabled for reception. (That is, it is not halted) 1 All controller receive activity to this queue is halted.
11	QHLT3	RxBD queue 3 is halted. It is a hardware-initiated stop indication. (DMACTRL[GRS] being set by the user does not cause a QHLT3 to be set.). The current frame and all other frames directed to a halted queue are discarded. A write with a value of 1 re-enables the queue for receiving. 0 This queue is enabled for reception. (That is, it is not halted) 1 All controller receive activity to this queue is halted.
12	QHLT4	RxBD queue 4 is halted. It is a hardware-initiated stop indication. (DMACTRL[GRS] being set by the user does not cause a QHLT4 to be set.). The current frame and all other frames directed to a halted queue are discarded. A write with a value of 1 re-enables the queue for receiving. 0 This queue is enabled for reception. (That is, it is not halted) 1 All controller receive activity to this queue is halted.
13	QHLT5	RxBD queue 5 is halted. It is a hardware-initiated stop indication. (DMACTRL[GRS] being set by the user does not cause a QHLT5 to be set.). The current frame and all other frames directed to a halted queue are discarded. A write with a value of 1 re-enables the queue for receiving. 0 This queue is enabled for reception. (That is, it is not halted) 1 All controller receive activity to this queue is halted.
14	QHLT6	RxBD queue 6 is halted. It is a hardware-initiated stop indication. (DMACTRL[GRS] being set by the user does not cause a QHLT6 to be set.). The current frame and all other frames directed to a halted queue are discarded. A write with a value of 1 re-enables the queue for receiving. 0 This queue is enabled for reception. (That is, it is not halted) 1 All controller receive activity to this queue is halted.
15	QHLT7	RxBD queue 7 is halted. It is a hardware-initiated stop indication. (DMACTRL[GRS] being set by the user does not cause a QHLT7 to be set.). The current frame and all other frames directed to a halted queue are discarded. A write with a value of 1 re-enables the queue for receiving. 0 This queue is enabled for reception. (That is, it is not halted) 1 All controller receive activity to this queue is halted.
16–23	—	Reserved
24	RXF0	Receive frame event occurred on ring 0. Set by the eTSEC if IEVENT[RXF] was set in relation to receiving a frame to this ring.

**Table 15-28. RSTAT Field Descriptions (continued)**

Bits	Name	Description
25	RXF1	Receive frame event occurred on ring 1. Set by the eTSEC if IEVENT[RXF] was set in relation to receiving a frame to this ring.
26	RXF2	Receive frame event occurred on ring 2. Set by the eTSEC if IEVENT[RXF] was set in relation to receiving a frame to this ring.
27	RXF3	Receive frame event occurred on ring 3. Set by the eTSEC if IEVENT[RXF] was set in relation to receiving a frame to this ring.
28	RXF4	Receive frame event occurred on ring 4. Set by the eTSEC if IEVENT[RXF] was set in relation to receiving a frame to this ring.
29	RXF5	Receive frame event occurred on ring 5. Set by the eTSEC if IEVENT[RXF] was set in relation to receiving a frame to this ring.
30	RXF6	Receive frame event occurred on ring 6. Set by the eTSEC if IEVENT[RXF] was set in relation to receiving a frame to this ring.
31	RXF7	Receive frame event occurred on ring 7. Set by the eTSEC if IEVENT[RXF] was set in relation to receiving a frame to this ring.

### 15.5.3.3.3 Receive Interrupt Coalescing Register (RXIC)

The RXIC register enables and configures the operational parameters for interrupt coalescing associated with received frames. [Figure 15-24](#) describes the RXIC register.

**Figure 15-24. RXIC Register Definition**

[Table 15-29](#) describes the fields of the RXIC register.

**Table 15-29. RXIC Field Descriptions**

Bits	Name	Description
0	ICEN	Interrupt coalescing enable 0 Interrupt coalescing is disabled. Interrupts are raised as they are received. 1 Interrupt coalescing is enabled. If the eTSEC receive frame interrupt is enabled (IMASK[RXFEN] is set), an interrupt is raised when the threshold number of frames is reached (defined by RXIC[ICFT]) or when the threshold timer expires (determined by RXIC[ICTT]).
1	ICCS	Interrupt coalescing timer clock source. 0 The coalescing timer advances count every 64 eTSEC Rx interface clocks (TSECn_GTX_CLK). 1 The coalescing timer advances count every 64 system clocks. This mode is recommended for FIFO operation.
2	—	Reserved

Table 15-29. RXIC Field Descriptions (continued)

Bits	Name	Description
3–10	ICFT	Interrupt coalescing frame count threshold. While interrupt coalescing is enabled (RXIC[ICE] is set), this value determines how many frames are received before raising an interrupt. The eTSEC threshold counter is reset to ICFT following an interrupt. The value of ICFT must be greater than zero avoid unpredictable behavior.
11–15	—	Reserved
16–31	ICTT	Interrupt coalescing timer threshold. While interrupt coalescing is enabled (RXIC[ICE] is set), this value determines the maximum amount of time after receiving a frame before raising an interrupt. If frames have been received but the frame count threshold has not been met, an interrupt is raised when the threshold timer reaches zero. The threshold timer is reset to the value in this field and begins counting down upon receiving the first frame having its RxBD[I] bit set. The threshold value is represented in units equal to 64 periods of the clock specified by RXIC[ICCS]. ICTT must be greater than zero to avoid unpredictable behavior.

#### 15.5.3.3.4 Receive Queue Control Register (RQUEUE)

The RQUEUE register enables each of the RxBD rings 0–7. By default, RxBD ring 0 is enabled. Figure 15-25 describes the definition for the RQUEUE register.

Offset eTSEC1:0x2\_4314; eTSEC2:0x2\_5314

Access: Read/Write




Figure 15-25. RQUEUE Register Definition

Table 15-30 describes the RQUEUE register.

Table 15-30. RQUEUE Field Descriptions

Bits	Name	Description
0–23	—	Reserved
24	EN0	Receive queue 0 enable. 0 RxBD ring is not queried for reception. In effect the receive queue is disabled. 1 RxBD ring is queried for reception.
25	EN1	Receive queue 1 enable. 0 RxBD ring is not queried for reception. In effect the receive queue is disabled. 1 RxBD ring is queried for reception.
26	EN2	Receive queue 2 enable. 0 RxBD ring is not queried for reception. In effect the receive queue is disabled. 1 RxBD ring is queried for reception.
27	EN3	Receive queue 3 enable. 0 RxBD ring is not queried for reception. In effect the receive queue is disabled. 1 RxBD ring is queried for reception.
28	EN4	Receive queue 4 enable. 0 RxBD ring is not queried for reception. In effect the receive queue is disabled. 1 RxBD ring is queried for reception.

**Table 15-30. RQUEUE Field Descriptions (continued)**

Bits	Name	Description
29	EN5	Receive queue 5 enable. 0 RxBD ring is not queried for reception. In effect the receive queue is disabled. 1 RxBD ring is queried for reception.
30	EN6	Receive queue 6 enable. 0 RxBD ring is not queried for reception. In effect the receive queue is disabled. 1 RxBD ring is queried for reception.
31	EN7	Receive queue 7 enable. 0 RxBD ring is not queried for reception. In effect the receive queue is disabled. 1 RxBD ring is queried for reception.

### 15.5.3.3.5 Receive Bit Field Extract Control Register (RBIFX)

The RBIFX register provides a set of four 6-bit offsets for locating up to four octets in a received frame and passing them to the receive queue filer as the user-defined ARB property. Through RBIFX a custom ARB filer property can be constructed from arbitrary bytes, which allows frame filing on the basis of bitfields not ordinarily provided to the filer, such as bits from the Ethernet preamble or TCP flags. The value of property ARB is the concatenation of {B0, B1, B2, B3} to 32-bits, where B0–B3 are the bytes as defined by RBIFX.

Figure 15-26 describes the definition for the RBIFX register.

Offset eTSEC1:0x2\_4330; eTSEC2:0x2\_5330

Access: Read/Write


**Figure 15-26. RBIFX Register Definition**

Table 15-31 describes the RBIFX register.

**Table 15-31. RBIFX Field Descriptions**


Bits	Name	Description
0–1	B0CTL	Location of byte 0 of property ARB. 00 Byte 0 is not extracted, and appears as zero in property ARB. 01 Byte 0 is located in the received frame at offset (B0OFFSET – 8) bytes from the first byte of the Ethernet DA. In non-FIFO modes, a negative effective offset points to bytes of the standard Ethernet preamble. Values of B0OFFSET less than 8 are reserved in FIFO modes. 10 Byte 0 is located in the received frame at offset B0OFFSET bytes from the byte after the last byte of the layer 2 header. 11 Byte 0 is located in the received frame at offset B0OFFSET bytes from the byte after the last byte of the layer 3 header.
2–7	B0OFFSET	Offset relative to the header defined by B0CTL that locates byte 0 of property ARB. An effective offset of zero points to the first byte of the specified header.

Table 15-31. RBIFX Field Descriptions (continued)

Bits	Name	Description
8–9	B1CTL	Location of byte 1 of property ARB. 00 Byte 1 is not extracted, and appears as zero in property ARB. 01 Byte 1 is located in the received frame at offset (B1OFFSET – 8) bytes from the first byte of the Ethernet DA. In non-FIFO modes, a negative effective offset points to bytes of the standard Ethernet preamble. Values of B1OFFSET less than 8 are reserved in FIFO modes. 10 Byte 0 is located in the received frame at offset B1OFFSET bytes from the byte after the last byte of the layer 2 header. 11 Byte 0 is located in the received frame at offset B1OFFSET bytes from the byte after the last byte of the layer 3 header.
10–15	B1OFFSET	Offset relative to the header defined by B1CTL that locates byte 1 of property ARB. An effective offset of zero points to the first byte of the specified header.
16–17	B2CTL	Location of byte 2 of property ARB. 00 Byte 2 is not extracted, and appears as zero in property ARB. 01 Byte 2 is located in the received frame at offset (B2OFFSET – 8) bytes from the first byte of the Ethernet DA. In non-FIFO modes, a negative effective offset points to bytes of the standard Ethernet preamble. Values of B2OFFSET less than 8 are reserved in FIFO modes. 10 Byte 0 is located in the received frame at offset B2OFFSET bytes from the byte after the last byte of the layer 2 header. 11 Byte 0 is located in the received frame at offset B2OFFSET bytes from the byte after the last byte of the layer 3 header.
18–23	B2OFFSET	Offset relative to the header defined by B2CTL that locates byte 2 of property ARB. An effective offset of zero points to the first byte of the specified header.
24–25	B3CTL	Location of byte 3 of property ARB. 00 Byte 3 is not extracted, and appears as zero in property ARB. 01 Byte 3 is located in the received frame at offset (B3OFFSET – 8) bytes from the first byte of the Ethernet DA. In non-FIFO modes, a negative effective offset points to bytes of the standard Ethernet preamble. Values of B3OFFSET less than 8 are reserved in FIFO modes. 10 Byte 0 is located in the received frame at offset B3OFFSET bytes from the byte after the last byte of the layer 2 header. 11 Byte 0 is located in the received frame at offset B3OFFSET bytes from the byte after the last byte of the layer 3 header.
26–31	B3OFFSET	Offset relative to the header defined by B3CTL that locates byte 3 of property ARB. An effective offset of zero points to the first byte of the specified header.

### 15.5.3.3.6 Receive Queue Filer Table Address Register (RQFAR)

RQFAR, shown in [Figure 15-27](#), contains the index of the current, indirectly accessible entry of the received queue filer table. Each table entry occupies a pair of 32-bit words, denoted RQCTRL and RQPROP. To access the RQCTRL and RQPROP words of entry  $n$ , write  $n$  to RQFAR. Then read or write the indexed RQCTRL and RQPROP words by reading or writing the RQFCR and RQFPR registers, respectively.



**Figure 15-27. Receive Queue Filer Table Address Register Definition**

Table 15-32 describes the fields of the RQFAR register.


**Table 15-32. RQFAR Field Descriptions**

Bits	Name	Description
0–23	—	Reserved
24–31	RQFAR	Current index of receive queue filer table, which spans a total of 256 entries.

### 15.5.3.3.7 Receive Queue Filer Table Control Register (RQFCR)

RQFCR is accessed to read or write the RQCTRL words in entries of the receive queue filer table. The table entries are described in greater detail in Section 15.6.4.2, “Receive Queue Filer.” The word accessed through RQFCR is defined by the current value of RQFAR.

Figure 15-28 describes the definition for the RQFCR register.



**Figure 15-28. Receive Queue Filer Table Control Register Definition**

Table 15-33 describes the fields of the RQFCR register.

**Table 15-33. RQFCR Field Descriptions**

Bit	Name	Description
0	GPI	General purpose interrupt. When a property matches the value in the RQPROP entry at this index, and REJ = 0 and AND = 0, the filer will instruct the Rx descriptor controller to set IEVENT[FGPI] when the corresponding receive frame is written to memory. If the timer is enabled (TMR_CTRL[TE] = 1), then TMR_PEVENT[RXP] will also be set.
1–15	—	Reserved, should be written with zero.
16–21	Q	Receive queue index, from 0 to 63, inclusive, written into the Rx frame control block associated with the received frame. When a property matches the value in the RQPROP entry at this index, and REJ = 0 and AND = 0, the frame is sent to either RxBD ring 0 (if RCTRL[FSQEN] = 1) or the RxBD ring with index (Q mod 8) and the filer table search is terminated. In the case where RCTRL[FSQEN] = 0, 8 virtual receive queues are overlaid on every RxBD ring, and software needs to consult the RQ field of the Rx frame control block to determine which virtual receive queue was chosen.

Table 15-33. RQFCR Field Descriptions (continued)

Bit	Name	Description
22	CLE	Cluster entry/exit (used in combination with AND bit). This bit brackets clusters, marking the start and end entries of a cluster. Clusters cannot be nested. 0 Regular RQCTRL entry. 1 If entry matches and AND = 1, treat subsequent entries as belonging to a nested cluster and enter the cluster; otherwise skip all entries up to and including the next cluster exit. If AND = 0, exit current cluster.
23	REJ	Reject frame. This bit and its specified action are ignored if AND = 1. 0 If entry matches, accept frame and file it to RxBD ring Q. 1 If entry matches, reject frame and discard it, ignoring Q.
24	AND	Match this entry and the next entry as a pair. 0 Match property[PID] against RQPROP, independent of the next entry. 1 Match property[PID] against RQPROP. If matched and CLE = 0, attempt to match next entry, otherwise, skip all entries up to and including the entry with AND = 0. If matched and CLE = 1, enter cluster of entries, otherwise, skip all entries up to and including the entry with CLE = 1 (cluster exit).
25–26	CMP	Comparison operation to perform on the RQPROP entry at this index when PID > 0. The property value extracted by the frame parser is masked by the 32-bit <i>mask_register</i> prior to comparison against RQPROP. However, the property value is not permanently altered by the value in <i>mask_register</i> . By default, <i>mask_register</i> is initialized to 0xFFFF_FFFF before each frame is processed.  In the case where PID = 0, CMP is interpreted as follows: 00/01 Filer <i>mask_register</i> is set to all 32 bits of RQPROP, and this entry always <i>matches</i> . 10/11 Filer <i>mask_register</i> is set to all 32 bits of RQPROP, and this entry always <i>fails to match</i> .  In the case where PID > 0, CMP is interpreted as follows (& is bit-wise AND operator): 00 <i>property</i> [PID] & <i>mask_register</i> = RQPROP 01 <i>property</i> [PID] & <i>mask_register</i> >= RQPROP 10 <i>property</i> [PID] & <i>mask_register</i> != RQPROP 11 <i>property</i> [PID] & <i>mask_register</i> < RQPROP
27	—	Reserved, should be written with zero.
28–31	PID	Property identifier. The value in the RQPROP entry at this index is interpreted according to PID (see <a href="#">Table 15-34</a> ).

### 15.5.3.3.8 Receive Queue Filer Table Property Register (RQFPR)

RQFPR (see [Figure 15-29](#)) is accessed to read or write the RQPROP words in entries of the receive queue filer table. The table entries are described in greater detail in [Section 15.6.4.2, “Receive Queue Filer.”](#) The word accessed through RQFPR is defined by the current value of RQFAR. [Figure 15-29](#) and [Figure 15-30](#) describe the fields of the RQFPR register according to property ID.




Figure 15-29. Receive Queue Filer Table Property IDs 0, 2–15 Register Definition

Offset eTSEC1:0x2\_433C; eTSEC2:0x2\_533C

Access: Read/Write




Figure 15-30. Receive Queue Filer Table Property ID1 Register Definition

Table 15-34 describes the fields of the RQFPR register.

Table 15-34. RQFPR Field Descriptions

PID <sup>1</sup>	Bit	Name	Description
0000	0–31	MASK	Mask bits to be written to Filer <i>mask_register</i> for masking of property values. The rule match/fail status for this PID is determined by RQCTRL[ <i>CMP</i> ]. Since <i>mask_register</i> is bit-wise ANDed with properties, every bit of MASK that is cleared also results in the corresponding property bit being cleared in comparisons. Therefore setting MASK to 0xFFFF_FFFF ensures that all property bits participate in rule matches.
0001	0–13	—	Reserved
	14	AR	Set if an ARP response packet is seen.
	15	ARQ	Set if an ARP request packet is seen.
	16	EBC	Set if the destination Ethernet address is to the broadcast address.
	17	VLN	Set if a VLAN tag (Ethertype DFVLAN[ <i>TAG</i> ] or 0x8100) was seen in the frame.
	18	CFI	Set to the value of the Canonical Format Indicator in the VLAN control tag if VLAN is set, zero otherwise.
	19	JUM	Set if a jumbo Ethernet frame was parsed.
	20	IPF	Set if a fragmented IPv4 or IPv6 header was encountered. See the descriptions of receive FCB fields IP and PRO in <a href="#">Section 15.6.3.3, “Receive Path Off-Load,”</a> for more information on determining the status of received packets for which IPF is set.
	21	—	Reserved
	22	IP4	Set if an IPv4 header was parsed.
	23	IP6	Set if an IPv6 header was parsed.
	24	ICC	Set if the IPv4 header checksum was checked.
	25	ICV	Set if the IPv4 header checksum was verified correct.
	26	TCP	Set if a TCP header was parsed.
	27	UDP	Set if a UDP header was parsed.
28–29	—	Reserved.	
30	PER	Set on a parse error, such as header inconsistency.	
31	EER	Set on an Ethernet framing error that prevents parsing.	



Table 15-34. RQFPR Field Descriptions (continued)

PID <sup>1</sup>	Bit	Name	Description
0010	0–7	ARB	User-defined arbitrary bit field property: byte 0 extracted. Defaults to 0x00.
	8–15		User-defined arbitrary bit field property: byte 1 extracted. Defaults to 0x00.
	16–23		User-defined arbitrary bit field property: byte 2 extracted. Defaults to 0x00.
	24–31		User-defined arbitrary bit field property: byte 3 extracted. Defaults to 0x00.
0011	0–7	—	Reserved, should be written with zero.
	8–31	DAH	Destination MAC address, most significant 24 bits. Defaults to 0x000000.
0100	0–7	—	Reserved, should be written with zero.
	8–31	DAL	Destination MAC address, least significant 24 bits. Defaults to 0x000000.
0101	0–7	—	Reserved, should be written with zero.
	8–31	SAH	Source MAC address, most significant 24 bits. Defaults to 0x000000.
0110	0–7	—	Reserved, should be written with zero.
	8–31	SAL	Source MAC address, least significant 24 bits. Defaults to 0x000000.

Table 15-34. RQFPR Field Descriptions (continued)

PID <sup>1</sup>	Bit	Name	Description
0111	0–15	—	Reserved, should be written with zero.
	16–31	ETY	<p>Ethertype of next layer protocol, that is, last ethertype if layer 2 headers nest. Defaults to 0xFFFF. Using the filer to match ETY does not work in the case of PPPoE packets, because the PPPoE ethertype in the original packet, 0x8864, is always overwritten with the PPP protocol field. Thus, matches on ETY == 0x8864 always fail.</p> <p>Instead, software should use PID=1 fields IP4 (ETY = 0x0021) and IP6 (ETY = 0x0057) to distinguish PPPoE session packets carrying IPv4 and IPv6 datagrams. Other PPP protocols are encoded in the ETY field, but many of them overlap with real ethertype definitions. Consult IANA and IEEE for possible ambiguities.</p> <p>A value in the length/type field greater than 1500 and less than 1536 is treated as a type encoding by the parser. Since no recognized types exist in this range, the controller will not parse beyond the length/type field of any such frame.</p> <p>Note that the eTSEC filer gets multiple packet attributes as a result of parsing the packet. The behavior of the eTSEC is that it pulls the innermost ethertype found in the packet; this means that in many supported protocols, it is impossible to create a filer rule that matches on the outer ethertype. There are four cases that need to be highlighted.</p> <ol style="list-style-type: none"> <li>1. The jumbo ethertype (0x8870)—In this case, the eTSEC assumes that the following header is LLC/SNAP. LLC/SNAP has an associated Ethertype, and the ETY field is populated with that ethertype. This makes it impossible to file on jumbo frames. In this case, one can use arbitrary extracted bytes to pull the outermost Ethertype.</li> <li>2. The PPPoE ethertype described above.</li> <li>3. The VLAN tag ethertype (0x8100)—In this case, one can use the PID=1 VLN bit to indicate that the packet had a VLAN tag.</li> <li>4. The MPLS tagged packets. In this case, one can use arbitrary extraction bytes to compare to the actual ethertype if a filer rule is intending to file based on an MPLS label existence.</li> </ol> <p>NOTE Users of the eTSEC parser/filer should be aware of a difference in behavior between rev 1 and rev 2 silicon in cases where the Ethernet type/length field contains a value between 1500 and 1536. In rev 2 silicon, values between 1500 and 1536 are interpreted as a type. Since there are currently no valid types in this range publicly defined by IANA, the controller will not parse beyond the length/type field of any such frame.</p> <p>If the same packet is encountered with rev 1 silicon, parser/filer behavior is different. With rev 1 silicon, such packets are treated as payload length. S/W must confirm the parser and filer results by checking the type/length field after the packet has been written to memory to see if it falls in this range.</p>
1000	0–19	—	Reserved, should be written with zero.
	20–31	VID	VLAN network identifier (as per IEEE Std 802.1Q). This value defaults to 0x000 if no VLAN tag was found, or the VLAN tag contained only priority information.
1001	0–28	—	Reserved, should be written with zero.
	29–31	PRI	VLAN user priority (as per IEEE Std 802.1p). This value defaults to 000 (best effort priority) if no VLAN tag was found.
1010	0–23	—	Reserved, should be written with zero.
	24–31	TOS	IPv4 header Type Of Service field or IPv6 Traffic Class field. This value defaults to 0x00 (default RFC 2474 best-effort behavior) if no IP header appeared. Note that for IPv6 the Traffic Class field is extracted using the IP header definition in RFC 2460. IPv6 headers formed using the earlier RFC 1883 have a different format and must be handled with software.

Table 15-34. RQFPR Field Descriptions (continued)

PID <sup>1</sup>	Bit	Name	Description
1011	0–23	—	Reserved, should be written with zero.
	24–31	L4P	Layer 4 protocol identifier as per published IANA specification. This is the last recognized protocol type recognized in the case of IPv6 extension headers. This value defaults to 0xFF to indicate that no layer 4 header was recognized (possibly due to absence of an IP header).
1100	0–31	DIA	Destination IP address. If an IPv4 header was found, this is the entire destination address. If an IPv6 header was found, this is the 32 most significant bits of the 128-bit destination address. This value defaults to 0x0000_0000 if no IP header appeared.
1101	0–31	SIA	Source IP address. If an IPv4 header was found, this is the entire source address. If an IPv6 header was found, this is the 32 most significant bits of the 128-bit source address. This value defaults to 0x0000_0000 if no IP header appeared.
1110	0–15	—	Reserved, should be written with zero.
	16–31	DPT	Destination port number for TCP or UDP headers. This value defaults to 0x0000 if no TCP or UDP headers were recognized.
1111	0–15	—	Reserved, should be written with zero.
	16–31	SPT	Source port number for TCP or UDP headers. This value defaults to 0x0000 if no TCP or UDP headers were recognized.

<sup>1</sup> PID is the property identifier field of the filter table control entry (see RQFCR[PID]) at the same index.

### 15.5.3.3.9 Maximum Receive Buffer Length Register (MRBLR)

The MRBLR register is written by the user. It informs the eTSEC how much space is in the receive buffer pointed to by the RxBD. [Figure 15-31](#) describes the definition for the MRBLR.



Figure 15-31. MRBLR Register Definition

Table 15-35. MRBLR Field Descriptions

Bits	Name	Description
0–15	—	Reserved
16–25	MRBL	Maximum receive buffer length. MRBL is the number of bytes that the eTSEC receiver writes to the receive buffer. The MRBL register is written by the user with a multiple of 64 for all modes. The eTSEC can write fewer bytes to the buffer than the value set in MRBL if a condition such as an error or end-of-frame occurs, but it never exceeds the MRBL value; therefore, user-supplied buffers must be at least as large as the MRBL. MRBL must be set, together with the number of buffer descriptors, to ensure adequate space for received frames. See <a href="#">Section 15.5.3.5.5, “Maximum Frame Length Register (MAXFRM),”</a> for further discussion.
26–31	—	To ensure that MRBL is a multiple of 64, these bits are reserved and should be cleared.

### 15.5.3.3.10 Receive Data Buffer Pointer High Register (RBDBPH)

The RBDBPH register is written by the user with the most significant address bits common to all RxBD buffer addresses, RxBD[Data Buffer Pointer]. As a consequence, Rx buffers must be placed in a 4 Gbyte segment of memory whose base address is prefixed by the bits in RBDBPH. The RxBD ring itself can reside in a different memory region (based at RBASEH). Figure 15-32 describes the definition for the RBDBPH register.




Figure 15-32. RBDBPH Register Definition

Table 15-36 describes the fields of the RBDBPH register.

Table 15-36. RBDBPH Field Descriptions

Bits	Name	Description
0–27	—	Reserved
28–31	RBDBPH	Most significant bits common to all data buffer addresses contained in RxBDs. The user must initialize RBDBPH before enabling the eTSEC receive function.

### 15.5.3.3.11 Receive Buffer Descriptor Pointers 0–7 (RBPTR0–RBPTR7)

RBPTR0–RBPTR7 each contains the low-order 32 bits of the next receive buffer descriptor address for their respective RxBD ring. Figure 15-33 describes the RBPTR registers. These registers take on the value of their ring's associated RBASE when the RBASE register is written by software. Software must not write RBPTR $n$  while eTSEC is actively receiving frames. However, RBPTR $n$  can be modified when the receiver is disabled or when no Rx buffer is in use (after a GRACEFUL STOP RECEIVE command is issued and the frame completes its reception) in order to change the next RxBD eTSEC receives.




Figure 15-33. RBPTR0–RBPTR7 Register Definition


Table 15-23 describes the fields of the  $RBPTR_n$  register.

**Table 15-37.  $RBPTR_n$  Field Descriptions**

Bits	Name	Description
0–28	$RBPTR_n$	Current RxBD pointer for RxBD ring $n$ . Points to the current BD being processed or to the next BD the receiver uses when it is idling. After reset or when the end of the RxBD ring is reached, eTSEC initializes $RBPTR_n$ to the value in the corresponding $RBASE_n$ . The $RBPTR$ register is internally written by the eTSEC's DMA controller during reception. The pointer increments by 8 (bytes) each time a descriptor is closed successfully by the eTSEC. Note that the 3 least-significant bits of this register are read only and zero.
29–31	—	Reserved

### 15.5.3.3.12 Receive Descriptor Base Address Registers ( $RBASE_0$ – $RBASE_7$ )

The  $RBASE_n$  registers are written by the user with the base address of each RxBD ring  $n$ . Each such value must be divisible by eight, since the 3 least-significant bits always write as 000. Figure 15-34 describes the  $RBASE_n$  registers.



**Figure 15-34.  $RBASE$  Register Definition**


Table 15-24 describes the fields of the  $RBASE_n$  registers.

**Table 15-38.  $RBASE_0$ – $RBASE_7$  Field Descriptions**

Bits	Name	Description
0–28	$RBASE_n$	Receive base for ring $n$ . $RBASE$ defines the starting location in the memory map for the eTSEC RxBDs. This field must be 8-byte aligned. Together with setting the W (wrap) bit in the last BD, the user can select how many BDs to allocate for the receive packets. The user must initialize $RBASE$ before enabling the eTSEC receive function on the associated ring.
29–31	—	Reserved

### 15.5.3.3.13 Receive Stamp Register ( $TMR\_RXTS\_H/L$ )

Receive timestamp register ( $RXTS\_H/L$ ). This register holds the value present in  $TMR\_CNT\_H/L$  when the eTSEC detects a new incoming Ethernet frame. This register is only updated when the precision timestamp logic is enable via  $TMR\_CTRL[TE]$ . This register is read only in normal operation. Figure 15-35 describes the definition for the  $RXTS\_H/L$  register.



**Figure 15-35. TMR\_RXTS\_H/L Register Definition**

Table 15-39 describes the fields of the TMR\_RXTS\_H/L register.

**Table 15-39. TMR\_RXTS\_H/L Register Field Descriptions**

Bits	Name	Description
0–63	TMR_RXTS_H/L	Value of the eTSEC precision timer upon detection of a start of frame symbol for the received frame.

### 15.5.3.4 MAC Functionality

This section describes the MAC registers and provides a brief overview of the functionality that can be exercised through the use of these registers, particularly those that provide functionality not explicitly required by the IEEE 802.3 standard. All of the MAC registers are 32 bits wide.

#### 15.5.3.4.1 Configuring the MAC

The MAC configuration registers 1 and 2 provide for configuring the MAC in multiple ways:

- Adjusting the preamble length—The length of the preamble can be adjusted from the nominal seven bytes to some other (non-zero) value. Should custom preamble insertion/extraction be configured, then this register must be left at its default value.
- Varying pad/CRC combinations—Three different pad/CRC combinations are provided to handle a variety of system requirements. Simplest are frames that already have a valid frame check sequence (FCS) field. The other two options include appending a valid CRC or padding and then appending a valid CRC, resulting in a minimum frame of 64 octets. In addition to the programmable register set, the pad/CRC behavior can be dynamically adjusted on a per-packet basis.

#### 15.5.3.4.2 Controlling CSMA/CD

The half-duplex register (HAFDUP) allows control over the carrier-sense multiple access/collision detection (CSMA/CD) logic of the eTSEC. Half-duplex mode is only supported for 10- and 100-Mbps operation. Following the completion of the packet transmission the part begins timing the inter packet gap (IPG) as programmed in the back-to-back IPG configuration register. The system is now free to begin another frame transfer.

In full-duplex mode both the carrier sense (CRS) and collision (COL) indications from the PHY are ignored, but in half-duplex mode the eTSEC defers to CRS, and following a carrier event, times the IPG using the non-back-to-back IPG configuration values that include support for the optional two-thirds/one-third CRS deferral process. This optional IPG mechanism enhances system robustness and ensures fair access to the medium. During the first two-thirds of the IPG, the IPG timer is cleared if CRS

is sensed. During the final one-third of the IPG, CRS is ignored and the transmission begins once IPG is timed. The two-thirds/one-third ratio is the recommended value.

### 15.5.3.4.3 Handling Packet Collisions

While transmitting a packet in half-duplex mode, the eTSEC is sensitive to COL. If a collision occurs, it aborts the packet and outputs the 32-bit jam sequence. The jam sequence is comprised of several bits of the CRC, inverted to guarantee an invalid CRC upon reception. A signal is sent to the system indicating that a collision occurred and that the start of the frame is needed for retransmission. The eTSEC then backs off of the medium for a time determined by the truncated binary exponential back off (BEB) algorithm. Following this back-off time, the packet is retried. The back-off time can be skipped if configured through the half-duplex register. However, this is non-standard behavior and its use must be carefully applied. Should any one packet experience excessive collisions, the packet is aborted. The system should flush the frame and move to the next one in line. If the system requests to send a packet while the eTSEC is deferring to a carrier, the eTSEC simply waits until the end of the carrier event and the timing of IPG before it honors the request.

If packet transmission attempts experience collisions, the eTSEC outputs the jam sequence and waits some amount of time before retrying the packet. This amount of time is determined by a controlled randomization process called truncated binary exponential back-off. The amount of time is an integer number of slot times. The number of slot times to delay before the  $n$ th retransmission attempt is chosen as a uniformly-distributed random integer  $r$  in the range:

$$0 \leq r \leq 2^k, \text{ where } k = \min(n, 10).$$

So after the first collision, the eTSEC backs off either 0 or 1 slot times. After the fifth collision, the eTSEC backs off between 0 and 32 slot times. After the tenth collision, the maximum number of slot times to back off is 1024. This can be adjusted through the half-duplex register. An alternate truncation point, such as 7 for instance, can be programmed. On average, the MAC is more aggressive after seven collisions than other stations on the network.

### 15.5.3.4.4 Controlling Packet Flow

Packet flow can be dealt with in a number of ways within eTSEC. A default retransmit attempt limit of 15 can be reduced using the half-duplex register. The slot time or collision window can be used to gate the retry window and possibly reduce the amount of transmit buffering within the system. The slot time for 10/100 Mbps is 512 bit times. Because the slot time begins at the beginning of the packet (including preamble), the end occurs around the 56th byte of the frame data. Slot time in 1000-Mbps mode is not supported.

Full-duplex flow control is provided for in IEEE 802.3x. Currently the standard does not address flow control in half-duplex environments. Common in the industry, however, is the concept of back pressure. The eTSEC implements the optional back pressure mechanism using the raise carrier method. If the system receive logic wishes to stop the reception of packets in a network-friendly way, transmit half-duplex flow control (THDF) is set (TCTRL[THDF]). If the medium is idle, the eTSEC raises carrier by transmitting preamble. Other stations on the half-duplex network then defer to the carrier.

In the event the preamble transmission happens to cause a collision, the eTSEC ensures the minimum 96-bit presence on the wire, then drops preamble and waits a back-off time depending on the value of the

back-pressure-no-back-off configuration bit HAFDUP[BP No BackOff]. These transmitting-preamble-for-back pressure collisions are not counted. If HAFDUP[BP No BackOff] is set, the eTSEC waits an inter-packet gap before resuming the transmission of preamble following the collision and does not defer. If HAFDUP[BP No BackOff] is cleared, the eTSEC adheres to the truncated BEB algorithm that allows the possibility of packets being received. This also can be detrimental in that packets can now experience excessive collisions, causing them to be dropped in the stations from which they originate. To reduce the likelihood of lost packets and packets leaking through the back pressure mechanism, HAFDUP[BP No BackOff] must be set.

The eTSEC drops carrier (cease transmitting preamble) periodically to avoid excessive defer conditions in other stations on the shared network. If, while applying back pressure, the eTSEC is requested to send a packet, it stops sending preamble, and waits one IPG before sending the packet. HAFDUP[BP No BackOff] applies for any collision that occurs during the sending of this packet. Collisions for packets while half duplex back pressure is asserted are counted. The eTSEC does not defer while attempting to send packets while in back pressure. Again, back pressure is non-standard, yet it can be effective in reducing the flow of receive packets.

#### 15.5.3.4.5 Controlling PHY Links

Control and status to and from the PHY is provided through the two-wire MII management interface described in IEEE 802.3u. The MII management registers (MII management configuration, command, address, control, status, and indicator registers) are used to exercise this interface between a host processor and one or more PHY devices.

The eTSEC MII's registers provide the ability to perform continuous read cycles (called a scan cycle); although, scan cycles are not explicitly defined in the standard. If requested (by setting MIIMCOM[Scan Cycle]), the part performs repetitive read cycles of the PHY status register, for example. In this way, link characteristics may be monitored more efficiently. The different fields in the MII management indicator register (scan, not valid and busy) are used to indicate availability of each read of the scan cycle to the host from MIIMSTAT[PHY scan].

Yet another parameter that can be modified through the MII registers is the length of the MII management interface preamble. After establishing that a PHY supports preamble suppression, the host may so configure the eTSEC. While enabled, the length of MII management frames are reduced from 64 clocks to 32 clocks. This effectively doubles the efficiency of the interface.



### 15.5.3.5 MAC Registers

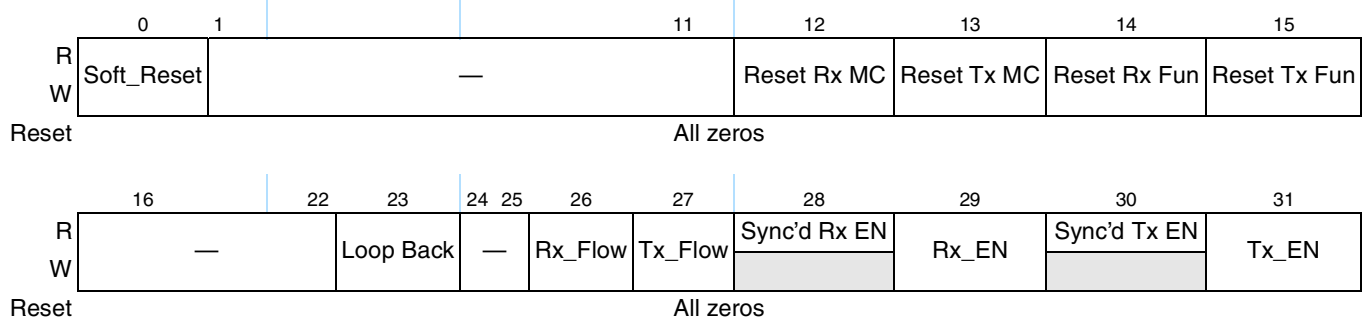
This section describes the MAC registers.

#### 15.5.3.5.1 MAC Configuration 1 Register (MACCFG1)

MACCFG1 is written by the user. [Figure 15-36](#) describes the definition for the MACCFG1 register.

Offset eTSEC1:0x2\_4500; eTSEC2:0x2\_5500

Access: Mixed



**Figure 15-36. MACCFG1 Register Definition**

[Table 15-40](#) describes the fields of the MACCFG1 register.

**Table 15-40. MACCFG1 Field Descriptions**

Bits	Name	Description
0	Soft_Reset	Soft reset. This bit is cleared by default. See <a href="#">Section 15.6.2.2, “Soft Reset and Reconfiguring Procedure,”</a> for more information on setting this bit. 0 Normal operation. 1 Place the entire MAC in reset except for the host interface.
1–11	—	Reserved
12	Reset Rx MC	Reset receive MAC control block. This bit is cleared by default. 0 Normal operation. 1 Place the receive part of the MAC in reset. This block detects control frames and contains the pause timers.
13	Reset Tx MC	Reset transmit MAC control block. This bit is cleared by default. 0 Normal operation. 1 Place the transmit part of the MAC in reset. This block multiplexes data and control frame transfers. It also responds to XOFF PAUSE control frames.
14	Reset Rx Fun	Reset receive function block. This bit is cleared by default. 0 Normal operation. 1 Place the receive function in reset. This block performs the receive frame protocol.
15	Reset Tx Fun	Reset transmit function block. This bit is cleared by default. 0 Normal operation. 1 Place the transmit function in reset. This block performs the frame transmission protocol.
16–22	—	Reserved
23	Loop Back	Loop back. This bit is cleared by default. 0 Normal operation. 1 Loop back the MAC transmit outputs to the MAC receive inputs.

**Table 15-40. MACCFG1 Field Descriptions (continued)**


Bits	Name	Description
24–25	—	Reserved
26	Rx_Flow	Receive flow. This bit is cleared by default. Must be 0 if MACCFG2[Full Duplex] = 0. 0 The receive MAC control ignores PAUSE flow control frames. 1 The receive MAC control detects and acts on PAUSE flow control frames.
27	Tx_Flow	Transmit flow. This bit is cleared by default. Must be 0 if MACCFG2[Full Duplex] = 0. 0 The transmit MAC control may not send PAUSE flow control frames if requested by the system. 1 The transmit MAC control may send PAUSE flow control frames if requested by the system.
28	Sync'd Rx EN	Receive enable synchronized to the receive stream. (Read-only) 0 Frame reception is not enabled. 1 Frame reception is enabled.
29	Rx_EN	Receive enable. This bit is cleared by default. If set, prior to clearing this bit, set DMACTRL[GRS] then confirm subsequent occurrence of the graceful receive stop interrupt (IEVENT[GRSC] is set). 0 The MAC may not receive frames from the PHY. 1 The MAC may receive frames from the PHY.
30	Sync'd Tx EN	Transmit enable synchronized to the transmit stream. (Read-only) 0 Frame transmission is not enabled. 1 Frame transmission is enabled.
31	Tx_EN	Transmit enable. This bit is cleared by default. If set, prior to clearing this bit, set DMACTRL[GTS] then confirm subsequent occurrence of the graceful receive stop interrupt (IEVENT[GTSC] is set). 0 The MAC may not transmit frames from the system. 1 The MAC may transmit frames from the system.

### 15.5.3.5.2 MAC Configuration 2 Register (MACCFG2)

The MACCFG2 register is written by the user. Figure 15-37 describes the definition for the MACCFG2 register.

Offset eTSEC1:0x2\_4504; eTSEC2:0x2\_5504

Access: Read/Write



**Figure 15-37. MACCFG2 Register Definition**

Table 15-41 describes the fields of the MACCFG2 register.

**Table 15-41. MACCFG2 Field Descriptions**


Bits	Name	Description																				
0–15	—	Reserved																				
16–19	Preamble Length	This field determines the length in bytes of the preamble field preceding each Ethernet start-of-frame delimiter byte. Values from 0x3 to 0xF are supported by the controller. The default value of 0x7 should not be altered in order to guarantee reliable operation with IEEE 802.3 compliant hardware.																				
20–21	—	Reserved																				
22–23	I/F Mode	This field determines the type of interface to which the MAC is connected. Its default is 00. 00 Reserved bit mode (not supported) (10 Mbps GENDEC/GPSI) 01 Nibble mode (MII) (10/100 Mbps MII/RMII) 10 Byte mode ( ) (1000 Mbps). Reserved if neither GMII or TBI are supported. 11 Reserved																				
24	PreAM RxEN	User defined preamble enable for received frames. This bit is cleared by default. 0 The MAC skips the Ethernet preamble without returning it. 1 The MAC recovers the received Ethernet preamble and passes it to the driver at the start of each received frame. If the preamble is less than 7 bytes, 0's are prepended to pad it to 7 bytes. Not applicable to or RMII 10/100 modes.																				
25	PreAM TxEN	User defined preamble enable for transmitted frames. This bit is cleared by default. 0 The MAC generates a standard Ethernet preamble. 1 If a user-defined preamble has been passed to the MAC it is transmitted instead of the standard preamble. Otherwise the standard Ethernet preamble is generated. The Preamble Length field should be left at its default setting if a user-defined preamble is transmitted. Not applicable to or RMII 10/100 modes.																				
26	Huge Frame	Huge frame enable. This bit is cleared by default. 0 Limit the length of frames received to less than or equal to the maximum frame length value (MAXFRM[Maximum Frame]) and limit the length of frames transmitted to less than the maximum frame length. See <a href="#">Section 15.6.7, "Buffer Descriptors,"</a> for further details of buffer descriptor bit updating. <table border="1" data-bbox="456 1220 1443 1497"> <thead> <tr> <th>Frame type</th> <th>Frame length</th> <th>Packet truncation</th> <th>Buffer descriptor updated</th> </tr> </thead> <tbody> <tr> <td>Receive or transmit</td> <td>&gt; maximum frame length</td> <td>yes</td> <td>yes</td> </tr> <tr> <td>Receive</td> <td>= maximum frame length</td> <td>no</td> <td>yes</td> </tr> <tr> <td>Transmit</td> <td>= maximum frame length</td> <td>no</td> <td>no</td> </tr> <tr> <td>Receive or transmit</td> <td>&lt; maximum frame length</td> <td>no</td> <td>no</td> </tr> </tbody> </table> 1 Frames are transmitted and received regardless of their relationship to the maximum frame length. Note that if Huge Frame is cleared, the user must ensure that adequate buffer space is allocated for received frames. See <a href="#">Section 15.5.3.5.5, "Maximum Frame Length Register (MAXFRM),"</a> for further information.	Frame type	Frame length	Packet truncation	Buffer descriptor updated	Receive or transmit	> maximum frame length	yes	yes	Receive	= maximum frame length	no	yes	Transmit	= maximum frame length	no	no	Receive or transmit	< maximum frame length	no	no
Frame type	Frame length	Packet truncation	Buffer descriptor updated																			
Receive or transmit	> maximum frame length	yes	yes																			
Receive	= maximum frame length	no	yes																			
Transmit	= maximum frame length	no	no																			
Receive or transmit	< maximum frame length	no	no																			
27	Length check	Length check. This bit is cleared by default. 0 No length field checking is performed. 1 The MAC checks the frame's length field on receive to ensure it matches the actual data field length. Transmitted frames are not checked.																				

**Table 15-41. MACCFG2 Field Descriptions (continued)**

Bits	Name	Description
28	MPEN	Magic packet enable for Ethernet modes. This bit is cleared by default. MPEN should be enabled only after GRACEFUL RECEIVE STOP and GRACEFUL TRANSMIT STOP are completed successfully (in other words, transmission and reception have stopped). 0 Normal receive behavior on receive, or Magic Packet mode has exited with reception of a valid Magic Packet. 1 Commence Magic Packet detection by the MAC provided that frame reception is enabled in MACCFG1. In this mode the MAC ignores all received frames until the specific Magic Packet frame is received, at which point this bit is cleared by the eTSEC, and a maskable interrupt through IEVENT[MAG] occurs.
29	PAD/CRC	Pad and append CRC. This bit is cleared by default. This bit must be set when in half-duplex mode (MACCFG2[Full Duplex] is cleared). 0 Frames presented to the MAC have a valid length and contain a CRC. 1 The MAC pads all transmitted short frames and appends a CRC to every frame regardless of padding requirement.
30	CRC EN	CRC enable. If the configuration bit PAD/CRC ENABLE or the per-packet PAD/CRC ENABLE is set, CRC ENABLE is ignored. This bit is cleared by default. 0 Frames presented to the MAC have a valid length and contain a valid CRC. 1 The MAC appends a CRC on all frames. Clear this bit if frames presented to the MAC have a valid length and contain a valid CRC.
31	Full Duplex	Full duplex configure. This bit is cleared by default. 0 The MAC operates in half-duplex mode only. 1 The MAC operates in full-duplex mode.

**15.5.3.5.3 Inter-Packet Gap/Inter-Frame Gap Register (IPGIFG)**

The IPGIFG register is written by the user. [Figure 15-38](#) describes the definition for IPGIFG.



**Figure 15-38. IPGIFG Register Definition**

[Table 15-42](#) describes the fields of the IPGIFG register.

**Table 15-42. IPGIFG Field Descriptions**

Bits	Name	Description
0	—	Reserved
1–7	Non-Back-to-Back Inter-Packet-Gap, Part 1	This is a programmable field representing the optional carrier sense window referenced in IEEE 802.3/4.2.3.2.1 ‘carrier deference’. If carrier is detected during the timing of IPGR1, the MAC defers to carrier. If, however, carrier becomes active after IPGR1, the MAC continues timing IPGR2 and transmits, knowingly causing a collision, thus ensuring fair access to medium. Its range of values is 0x00 to IPGR2. Its default is 0x40 (64d) which follows the two-thirds/one-third guideline.
8	—	Reserved

Table 15-42. IPGIFG Field Descriptions (continued)

Bits	Name	Description
9–15	Non-Back-to-Back Inter-Packet-Gap, Part 2	This is a programmable field representing the non-back-to-back inter-packet-gap in bits. Its default is 0x60 (96d), which represents the minimum IPG of 96 bits.
16–23	Minimum IFG Enforcement	This is a programmable field representing the minimum number of bits of IFG to enforce between frames. A frame is dropped whose IFG is less than that programmed. The default setting of 0x50 (80d) represents half of the nominal minimum IFG which is 160 bits.
24	—	Reserved
25–31	Back-to-Back Inter-Packet-Gap	This is a programmable field representing the IPG between back-to-back packets. This is the IPG parameter used exclusively in full-duplex mode and in half-duplex mode if two transmit packets are sent back-to-back. Set this field to the number of bits of IPG desired. The default setting of 0x60 (96d) represents the minimum IPG of 96 bits.

### 15.5.3.5.4 Half-Duplex Register (HAFDUP)

The HAFDUP register is written by the user. Figure 15-39 describes the HAFDUP register.

Offset eTSEC1:0x2\_450C; eTSEC2:0x2\_550C

Access: Read/Write




Figure 15-39. Half-Duplex Register Definition

Table 15-43 describes the fields of the HAFDUP register.

Table 15-43. HAFDUP Field Descriptions


Bits	Name	Description
0–7	—	Reserved
8–11	Alternate BEB Truncation	This field is used while ALTERNATE BINARY EXPONENTIAL BACKOFF ENABLE is set. The value programmed is substituted for the Ethernet standard value of ten. Its default is 0xA.
12	Alt BEB	Alternate binary exponential backoff. This bit is cleared by default. 0 The Tx MAC follows the standard binary exponential back off rule. 1 The Tx MAC uses the ALTERNATE BINARY EXPONENTIAL BACKOFF TRUNCATION setting instead of the 802.3 standard tenth collision. The standard specifies that any collision after the tenth uses one less than 210 as the maximum backoff time.
13	BP No BackOff	Back pressure no backoff. This bit is cleared by default. 0 The Tx MAC follows the binary exponential back off rule. 1 The Tx MAC immediately re-transmits, following a collision, during back pressure operation.

**Table 15-43. HAFDUP Field Descriptions (continued)**

Bits	Name	Description
14	No BackOff	No backoff. This bit is cleared by default. 0 The Tx MAC follows the binary exponential back off rule. 1 The Tx MAC immediately re-transmits following a collision.
15	Excess Defer	Excessively deferred. This bit is set by default. 0 The Tx MAC aborts the transmission of a packet that is excessively deferred. 1 The Tx MAC allows the transmission of a packet that is excessively deferred.
16–19	Retransmission Maximum	This is a programmable field specifying the number of retransmission attempts following a collision before aborting the packet due to excessive collisions. The standard specifies the attempt limit to be 0xF (15d). Its default value is 0xF.
20–25	—	Reserved
26–31	Collision Window	This is a programmable field representing the slot time or collision window during which collisions occur in properly configured networks. Because the collision window starts at the beginning of transmission, the preamble and SFD are included. Its default of 0x37 (55d) corresponds to the count of frame bytes at the end of the window.

### 15.5.3.5.5 Maximum Frame Length Register (MAXFRM)

The MAXFRM register is written by the user. [Figure 15-40](#) shows the MAXFRM register.

**Figure 15-40. Maximum Frame Length Register Definition**

[Table 15-44](#) describes the fields of the MAXFRM register.

**Table 15-44. MAXFRM Descriptions**

Bits	Name	Description
0–15	—	Reserved
16–31	Maximum Frame	By default this field is set to 0x0600 (1536 bytes). It sets the maximum Ethernet frame size in both the transmit and receive directions. (Refer to MACCFG2[Huge Frame].)  Note that if MACCFG2[Huge Frame] = 0, the value of this field must be less than or equal to MRBLR[MRBL] × (minimum number of RxBDs per ring). See <a href="#">Section 15.5.3.5.2, “MAC Configuration 2 Register (MACCFG2)”</a> , <a href="#">Section 15.5.3.3.9, “Maximum Receive Buffer Length Register (MRBLR)”</a> , and <a href="#">Section 15.6.7.3, “Receive Buffer Descriptors (RxBd)”</a> .

### 15.5.3.5.6 MII Management Configuration Register (MIIMCFG)

The MIIMCFG register is written by the user to configure all MII management operations. Note that MII management hardware is shared by all eTSECs. Thus, only through the MIIM registers of eTSEC1 can external PHYs be accessed and configured. Note: when an eTSEC is configured to use RTBI, configuration of the RTBI (described in [Section 15.5.4, “Ten-Bit Interface \(TBI\)”](#)) is done through the

MIIM registers for that eTSEC. For example, if a RTBI interface is required on eTSEC2, then the MIIM registers starting at offset 0x2\_5520 are used to configure it.

Figure 15-41 describes the definition for the MIIMCFG register.




Figure 15-41. MII Management Configuration Register Definition

Table 15-45 describes the fields of the MIIMCFG register.

Table 15-45. MIIMCFG Field Descriptions

Bits	Name	Description
0	Reset Mgmt	Reset management. This bit is cleared by default. 0 Allow the MII MGMT to perform mgmt read/write cycles if requested through the host interface. 1 Reset the MII MGMT.
1–26	—	Reserved
27	No Pre	Preamble suppress. This bit is cleared by default. 0 The MII MGMT performs Mgmt read/write cycles with 32 clocks of preamble. 1 The MII MGMT suppresses preamble generation and reduces the Mgmt cycle from 64 clocks to 32 clocks. This is in accordance with IEEE 802.3/22.2.4.4.2.
28	—	Reserved
29–31	MgmtClk	This field determines the clock frequency of the MII management clock (EC_MDC). Its default value is 111. <b>Note:</b> The eTSEC system clock is derived from (CCB Clock)/2. 000 1/4 of the eTSEC system clock divided by 8 001 1/4 of the eTSEC system clock divided by 8 010 1/6 of the eTSEC system clock divided by 8 011 1/8 of the eTSEC system clock divided by 8 100 1/10 of the eTSEC system clock divided by 8 101 1/14 of the eTSEC system clock divided by 8 110 1/20 of the eTSEC system clock divided by 8 111 1/28 of the eTSEC system clock divided by 8

### 15.5.3.5.7 MII Management Command Register (MIIMCOM)

The MIIMCOM register is written by the user. Figure 15-42 describes the definition for MIIMCOM.




Figure 15-42. MIIMCOM Register Definition


Table 15-46 describes the fields of the MIIMCOM register.

**Table 15-46. MIIMCOM Descriptions**

Bits	Name	Description
0–29	—	Reserved
30	Scan Cycle	Scan cycle. This bit is cleared by default. 0 Normal operation. 1 The MII management continuously performs read cycles. This is useful for monitoring link fail, for example.
31	Read Cycle	Read cycle. This bit is cleared by default but is not self-clearing once set. 0 Normal operation. 1 The MII management performs a single read cycle upon the transition of this bit from 0 to 1 using the PHY address (at MIIMADD[PHY Address]) and the register address (at MIIMADD[Register Address]). The 0-to-1 transition of this bit also causes the MIIMIND[Busy] bit to be set. The read is complete when the MIIMIND[Busy] bit clears. Data is returned in register MIIMSTAT[PHY Status].

### 15.5.3.5.8 MII Management Address Register (MIIMADD)

The MIIMADD register is written by the user. Figure 15-43 shows the MIIMADD register.



**Figure 15-43. MIIMADD Register Definition**

Table 15-47 describes the fields of the MIIMADD register.


**Table 15-47. MIIMADD Field Descriptions**

Bits	Name	Description
0–18	—	Reserved
19–23	PHY Address	This field represents the 5-bit PHY address field of Mgmt cycles. Up to 31 PHYs can be addressed (0 is reserved). Its default value is 0x00.
24–26	—	Reserved
27–31	Register Address	This field represents the 5-bit register address field of Mgmt cycles. Up to 32 registers can be accessed. Its default value is 0x00.



### 15.5.3.5.9 MII Management Control Register (MIIMCON)

MIIMCON, shown in [Figure 15-44](#), is written by the user.



**Figure 15-44. MII Mgmt Control Register Definition**


[Table 15-48](#) describes the fields of the MIIMCON register.

**Table 15-48. MIIMCON Field Descriptions**

Bits	Name	Description
0–15	—	Reserved
16–31	PHY Control	If written, an MII Mgmt write cycle is performed using this 16-bit data, the pre-configured PHY address (at MIIMADD[PHY Address]) and the register address (at MIIMADD[Register Address]). Its default value is 0x0000.

### 15.5.3.5.10 MII Management Status Register (MIIMSTAT)

The MIIMSTAT register is read only by the user. [Figure 15-45](#) describes the definition for the MIIMSTAT register.



**Figure 15-45. MIIMSTAT Register Definition**


[Table 15-49](#) describes the fields of the MIIMSTAT register.

**Table 15-49. MIIMSTAT Field Descriptions**

Bits	Name	Description
0–15	—	Reserved
16–31	PHY Status	Following an MII Mgmt read cycle, the 16-bit data can be read from this location. Its default value is 0x0000.

### 15.5.3.5.11 MII Management Indicator Register (MIIMIND)

The MIIMIND register is read-only by the user. Figure 15-46 describes the definition for the MIIMIND register.




**Figure 15-46. MII Mgmt Indicator Register Definition**

**Table 15-50. MIIMIND Field Descriptions**

Bits	Name	Description
0–28	—	Reserved
29	Not Valid	Not valid. 0 MII Mgmt read cycle has completed and the read data is valid. 1 MII Mgmt read cycle has not completed and the read data is not yet valid.
30	Scan	Scan in progress. 0 A scan operation (continuous MII Mgmt read cycles) is not in progress. 1 A scan operation (continuous MII Mgmt read cycles) is in progress.
31	Busy	Busy. 0 MII Mgmt block is not currently performing an MII Mgmt read or write cycle. 1 MII Mgmt block is currently performing an MII Mgmt read or write cycle.

### 15.5.3.5.12 Interface Status Register (IFSTAT)

Figure 15-47 shows the IFSTAT register.



**Figure 15-47. Interface Status Register Definition**

Table 15-51 describes the fields of the FSTAT register.

**Table 15-51. IFSTAT Field Descriptions**

Bits	Name	Description
0–21	—	Reserved

Table 15-51. IFSTAT Field Descriptions (continued)

Bits	Name	Description
22	Excess Defer	Excessive transmission defer. This bit latches high and is cleared when read. This bit is cleared by default. 0 Normal operation. 1 The MAC excessively defers a transmission.
23–31	—	Reserved

### 15.5.3.5.13 MAC Station Address Part 1 Register (MACSTNADDR1)

The MACSTNADDR1 register is written by the user. The value of the station address written into MACSTNADDR1 and MACSTNADDR2 is byte reversed from how it would appear in the DA field of a frame in memory. For example, for a station address of 0x12345678ABCD, MACSTNADDR1 is set to 0xCDAB7856 and MACSTNADDR2 is set to 0x34120000.

Figure 15-48 shows the MACSTNADDR1 register.




Figure 15-48. MAC Station Address Part 1 Register Definition

Table 15-52 describes the fields of the MACSTNADDR1 register.

Table 15-52. MACSTNADDR1 Field Descriptions

Bit	Name	Description
0–7	Station Address, 6th Octet	This field holds the sixth octet of the station address. The sixth octet (station address bits 40–47) defaults to a value of 0x0.
8–15	Station Address, 5th Octet	This field holds the fifth octet of the station address. The fifth octet (station address bits 32–39) defaults to a value of 0x0.
16–23	Station Address, 4th Octet	This field holds the fourth octet of the station address. The fourth octet (station address bits 24–31) defaults to a value of 0x0.
24–31	Station Address, 3rd Octet	This field holds the third octet of the station address. The third octet (station address bits 16–23) defaults to a value of 0x0.

### 15.5.3.5.14 MAC Station Address Part 2 Register (MACSTNADDR2)

The MACSTNADDR2 register is written by the user. Figure 15-49 describes the definition for the MACSTNADDR2 register.




Figure 15-49. MAC Station Address Part 2 Register Definition

Table 15-53 describes the fields of the MACSTNADDR2 register.

Table 15-53. MACSTNADDR2 Field Descriptions

Bit	Name	Description
0–7	Station Address, 2nd Octet	This field holds the second octet of the station address. The second octet (station address bits 8–15) defaults to a value of 0x0.
8–15	Station Address, 1st Octet	This field holds the first octet of the station address. The first octet (station address bits 0–7) defaults to a value of 0x0.
16–31	—	Reserved

### 15.5.3.5.15 MAC Exact Match Address 1–15 Part 1 Registers (MAC01ADDR1–MAC15ADDR1)

The MAC01ADDR1–MAC15ADDR1 registers are written by the user with the unicast or multicast addresses aliasing the MAC. Figure 15-50 describes the definition for all of the fifteen MAC<sub>n</sub>ADDR1 registers. The value of the address written into MAC<sub>x</sub>ADDR1 and MAC<sub>n</sub>ADDR2 is byte reversed from how it would appear in the DA field of a frame in memory. For example, for a MAC address of 0x12345678ABCD, MAC<sub>n</sub>ADDR1 is set to 0xCDAB7856 and MAC<sub>n</sub>ADDR2 is set to 0x34120000. For any valid, non-zero MAC address received, exact match registers can be excluded individually by clearing them to all zero bytes.




Figure 15-50. MAC Exact Match Address *n* Part 1 Register Definition

Table 15-52 describes the fields of a MAC<sub>n</sub>ADDR1 register.

**Table 15-54. MAC<sub>n</sub>ADDR1 Field Descriptions**


Bit	Name	Description
0–7	Exact Match Address, 6th Octet	Holds the sixth octet of the exact match address. The sixth octet (destination address bits 40–47) defaults to a value of 0x0.
8–15	Exact Match Address, 5th Octet	Holds the fifth octet of the exact match address. The fifth octet (destination address bits 32–39) defaults to a value of 0x0.
16–23	Exact Match Address, 4th Octet	Holds the fourth octet of the exact match address. The fourth octet (destination address bits 24–31) defaults to a value of 0x0.
24–31	Exact Match Address, 3rd Octet	Holds the third octet of the exact match address. The third octet (destination address bits 16–23) defaults to a value of 0x0.

### 15.5.3.5.16 MAC Exact Match Address 1–15 Part 2 Registers (MAC01ADDR2–MAC15ADDR2)

The MAC01ADDR2–MAC15ADDR2 registers are written by the user with the unicast or multicast addresses aliasing the MAC. Figure 15-51 describes the definition for all of the fifteen MAC<sub>x</sub>ADDR2 registers.

Offset eTSEC1:0x2\_454C+8×*n*; eTSEC2:0x2\_554C+8×*n*

Access: Read/Write



**Figure 15-51. MAC Exact Match Address x Part 2 Register Definition**

Table 15-53 describes the fields of a MAC<sub>x</sub>ADDR2 register.

**Table 15-55. MAC01ADDR2–MAC15ADDR2 Field Descriptions**

Bit	Name	Description
0–7	Exact Match Address, 2nd Octet	This field holds the second octet of the exact match address. The second octet (destination address bits 8–15) defaults to a value of 0x0.
8–15	Exact Match Address, 1st Octet	This field holds the first octet of the exact match address. The first octet (destination address bits 0–7) defaults to a value of 0x0.
16–31	—	Reserved

### 15.5.3.6 MIB Registers

This section describes the MIB registers. The eTSEC RMON module has 37 separate statistics counters, which simply count or accumulate statistical events that occur as packets transmitted and received. These counters support RMON MIB group 1, RMON MIB group 2 if table counters, RMON MIB group 3, RMON MIB group 9, RMON MIB 2, and the IEEE 802.3 Ethernet MIB.

An interrupt can be generated upon any one counter's rollover condition through a carry interrupt output from the RMON. Each counter's rollover condition can be discretely masked from causing an interrupt by internal masking registers. In addition, each individual counter value may be reset on read access, or all counters may be simultaneously reset by setting ECNTRL[CLRCNT].

The majority of MIB counters are Ethernet-specific.

### NOTE

RMON counters do not comprehend custom VLAN tagged frames. Affected counters include TRMGV, RMCA, RBCA, RXCF, RXPF, RXUO, RALN, RFLR, ROVR, RJBR, TMCA, TBCA, TXPF, TXCF. Specifically, custom VLAN tagged frames are not afforded the ability to be greater than 1518, as compared to the IEEE standard tagged frames.

#### 15.5.3.6.1 Transmit and Receive 64-Byte Frame Counter (TR64)

Figure 15-52 describes the definition for the TR64 register.




Figure 15-52. Transmit and Receive 64-Byte Frame Register Definition

Table 15-56 describes the fields of the TR64 register.

Table 15-56. TR64 Field Descriptions

Bits	Name	Description
0–9	—	Reserved
10–31	TR64	Transmit and receive 64-byte frame counter—Increment for each good or bad frame transmitted and received which is 64 bytes in length, inclusive (excluding preamble and SFD but including FCS bytes).

#### 15.5.3.6.2 Transmit and Receive 65- to 127-Byte Frame Counter (TR127)

Figure 15-53 describes the definition for the TR127 register.




Figure 15-53. Transmit and Receive 65- to 127-Byte Frame Register Definition


Table 15-57 describes the fields of the TR127 register.

**Table 15-57. TR127 Field Descriptions**

Bits	Name	Description
0–9	—	Reserved
10–31	TR127	Transmit and receive 65- to 127-byte frame counter—Increments for each good or bad frame transmitted and received which is 65–127 bytes in length, inclusive (excluding preamble and SFD but including FCS bytes).

### 15.5.3.6.3 Transmit and Receive 128- to 255-Byte Frame Counter (TR255)

Figure 15-54 describes the definition for the TR255 register.



**Figure 15-54. Transmit and Received 128- to 255-Byte Frame Register Definition**


Table 15-58 describes the fields of the TR255 register.

**Table 15-58. TR255 Field Descriptions**

Bits	Name	Description
0–9	—	Reserved
10–31	TR255	Transmit and receive 128- to 255-byte frame counter—Increments for each good or bad frame transmitted and received which is 128–255 bytes in length, inclusive (excluding preamble and SFD but including FCS bytes).

### 15.5.3.6.4 Transmit and Receive 256- to 511-Byte Frame Counter (TR511)

Figure 15-55 describes the definition for the TR511 register.



**Figure 15-55. Transmit and Received 256- to 511-Byte Frame Register Definition**


Table 15-59 describes the fields of the TR511 register.

**Table 15-59. TR511 Field Descriptions**

Bits	Name	Description
0–9	—	Reserved
10–31	TR511	Increments for each good or bad frame transmitted and received which is 256–511 bytes in length, inclusive (excluding preamble and SFD but including FCS bytes).

### 15.5.3.6.5 Transmit and Receive 512- to 1023-Byte Frame Counter (TR1K)

Figure 15-56 shows the TR1K register.



**Figure 15-56. Transmit and Received 512- to 1023-Byte Frame Register Definition**


Table 15-60 describes the fields of the TR1K register.

**Table 15-60. TR1K Field Descriptions**

Bits	Name	Description
0–9	—	Reserved
10–31	TR1K	Increments for each good or bad frame transmitted and received which is 512–1023 bytes in length, inclusive (excluding preamble and SFD but including FCS bytes).

### 15.5.3.6.6 Transmit and Receive 1024- to 1518-Byte Frame Counter (TRMAX)

Figure 15-57 describes the definition for the TRMAX register.



**Figure 15-57. Transmit and Received 1024- to 1518-Byte Frame Register Definition**

Table 15-61 describes the fields of the TRMAX register.

**Table 15-61. TRMAX Field Descriptions**

Bits	Name	Description
0–9	—	Reserved
10–31	TRMAX	Increments for each good or bad frame transmitted and received which is 1024–1518 bytes in length, inclusive (excluding preamble and SFD but including FCS bytes).



### 15.5.3.6.7 Transmit and Receive 1519- to 1522-Byte VLAN Frame Counter (TRMGV)

Figure 15-58 describes the definition for the TRMGV register.




Figure 15-58. Transmit and Received 1519- to 1522-Byte VLAN Frame Register Definition

Table 15-62 describes the fields of the TRMGV register.

Table 15-62. TRMGV Field Descriptions

Bits	Name	Description
0–9	—	Reserved
10–31	TRMGV	Increments for each good or bad frame transmitted and received which is 1519–1522 bytes in length, inclusive (excluding preamble and SFD but including FCS bytes).

### 15.5.3.6.8 Receive Byte Counter (RBYT)

Figure 15-59 shows the RBYT register.




Figure 15-59. Receive Byte Counter Register Definition

Table 15-63 describes the fields of the RBYT register.

Table 15-63. RBYT Field Descriptions

Bits	Name	Description
0–31	RBYT	Receive byte counter. The statistic counter register increments by the byte count of frames received, including those in bad packets, excluding preamble and SFD but including FCS bytes.

### 15.5.3.6.9 Receive Packet Counter (RPKT)

Figure 15-60 describes the definition for the RPKT register.




Figure 15-60. Receive Packet Counter Register Definition


Table 15-64 describes the fields of the RPKT register.

**Table 15-64. RPKT Field Descriptions**

Bits	Name	Description
0–9	—	Reserved
10–31	RPKT	Receive packet counter. Increments for each frame received packet (including bad packets, all unicast, broadcast, and multicast packets).

### 15.5.3.6.10 Receive FCS Error Counter (RFCS)

Figure 15-61 describes the definition for the RFCS register.



**Figure 15-61. Receive FCS Error Counter Register Definition**


Table 15-65 describes the fields of the RFCS register.

**Table 15-65. RFCS Field Descriptions**

Bits	Name	Description
0–15	—	Reserved
16–31	RFCS	Receive FCS error counter. In Ethernet mode, increments for each frame received that has an integral 64–1518 length and contains a frame check sequence error.

### 15.5.3.6.11 Receive Multicast Packet Counter (RMCA)

Figure 15-62 describes the definition for the RMCA register.



**Figure 15-62. Receive Multicast Packet Counter Register Definition**


Table 15-66 describes the fields of the RMCA register.

**Table 15-66. RMCA Field Descriptions**

Bits	Name	Description
0–9	—	Reserved
10–31	RMCA	Receive multicast packet counter. Increments for each multicast frame with valid CRC and of lengths 64 to 1518 (non VLAN) or 1522 (VLAN), excluding broadcast frames. This count does not include range/length errors.

### 15.5.3.6.12 Receive Broadcast Packet Counter (RBCA)

Figure 15-63 describes the definition for the RBCA register.



**Figure 15-63. Receive Broadcast Packet Counter Register Definition**


Table 15-67 describes the fields of the RBCA register.

**Table 15-67. RBCA Field Descriptions**

Bits	Name	Description
0–9	—	Reserved
10–31	RBCA	Receive broadcast packet counter. Increments for each broadcast frame with valid CRC and of lengths 64 to 1518 (non VLAN) or 1522 (VLAN), excluding multicast frames. Does not include range/length errors.

### 15.5.3.6.13 Receive Control Frame Packet Counter (RXCF)

Figure 15-64 describes the definition for the RXCF register.



**Figure 15-64. Receive Control Frame Packet Counter Register Definition**


Table 15-68 describes the fields of the RXCF register.

**Table 15-68. RXCF Field Descriptions**

Bits	Name	Description
0–15	—	Reserved
16–31	RXCF	Receive control frame packet counter. Increments for each MAC control frame received (PAUSE and unsupported) with valid CRC and of lengths 64 to 1518 (non VLAN) or 1522 (VLAN).

### 15.5.3.6.14 Receive Pause Frame Packet Counter (RXPF)

Figure 15-65 describes the definition for the RXPF register.



**Figure 15-65. Receive Pause Frame Packet Counter Register Definition**


Table 15-69 describes the fields of the RXPF register.

**Table 15-69. RXPF Field Descriptions**

Bits	Name	Description
0–15	—	Reserved
16–31	RXPF	Receive PAUSE frame packet counter. Increments each time a PAUSE MAC control frame is received with valid CRC and of lengths 64 to 1518 (non VLAN) or 1522 (VLAN).

### 15.5.3.6.15 Receive Unknown Opcode Packet Counter (RXUO)

Figure 15-66 describes the definition for the RXUO register.



**Figure 15-66. Receive Unknown Opcode Packet Counter Register Definition**

Table 15-70 describes the fields of the RXUO register.

**Table 15-70. RXUO Field Descriptions**

Bits	Name	Description
0–15	—	Reserved
16–31	RXUO	Receive unknown opcode counter. Increments each time a MAC control frame is received which contains an opcode other than PAUSE, but the frame has valid CRC and length 64 to 1518 (non VLAN) or 1522 (VLAN).

### 15.5.3.6.16 Receive Alignment Error Counter (RALN)

Figure 15-67 describes the definition for the RALN register.



**Figure 15-67. Receive Alignment Error Counter Register Definition**


Table 15-71 describes the fields of the RALN register.

**Table 15-71. RALN Field Descriptions**

Bits	Name	Description
0–15	—	Reserved
16–31	RALN	Receive alignment error counter. Increments for each received frame from 64 to 1518 (non VLAN) or 1522 (VLAN) which contains an invalid FCS and is not an integral number of bytes.

### 15.5.3.6.17 Receive Frame Length Error Counter (RFLR)

Figure 15-68 describes the definition for the RFLR register.



**Figure 15-68. Receive Frame Length Error Counter Register Definition**

Table 15-72 describes the fields of the RFLR register.

**Table 15-72. RFLR Field Descriptions**

Bits	Name	Description
0–15	—	Reserved
16–31	RFLR	Receive frame length error counter. Increments for each frame received in which the 802.3 length field did not match the number of data bytes actually received (46–1500 bytes). The counter does not increment if the length field is not a valid 802.3 length, such as an Ethertype value.

### 15.5.3.6.18 Receive Code Error Counter (RCDE)

Figure 15-69 describes the definition for the RCDE register.




Figure 15-69. Receive Code Error Counter Register Definition

Table 15-73 describes the fields of the RCDE register.

Table 15-73. RCDE Field Descriptions

Bits	Name	Description
0–15	—	Reserved
16–31	RCDE	Receive code error counter. Increments each time a valid carrier is present and at least one invalid data symbol is detected.

### 15.5.3.6.19 Receive Carrier Sense Error Counter (RCSE)

Figure 15-70 describes the definition for the RCSE register.




Figure 15-70. Receive Carrier Sense Error Counter Register Definition

Table 15-74 describes the fields of the RCSE register.

Table 15-74. RCSE Field Descriptions

Bits	Name	Description
0–15	—	Reserved
16–31	RCSE	Receive false carrier counter. Counts the number of times that the carrier sense condition was lost or never asserted when attempting to transmit a frame on a particular interface. The count represented by an instance of this object is incremented at most once per transmission attempt, even if the carrier sense condition fluctuates during a transmission attempt. The event is reported along with the statistics generated on the next received frame, as defined by a 1 on TSEC <sub>n</sub> _RX_ER and an 0xE on TSEC <sub>n</sub> _RXD. Only one false carrier condition can be detected and logged between frames.

### 15.5.3.6.20 Receive Undersize Packet Counter (RUND)

Figure 15-71 describes the definition for the RUND register.



**Figure 15-71. Receive Undersize Packet Counter Register Definition**


Table 15-75 describes the fields of the RUND register.

**Table 15-75. RUND Field Descriptions**

Bits	Name	Description
0–15	—	Reserved
16–31	RUND	Receive undersize packet counter. Increments each time a frame is received which is less than 64 bytes in length and contains a valid FCS and were otherwise well formed. This count does not include range length errors.

### 15.5.3.6.21 Receive Oversize Packet Counter (ROVR)

Figure 15-72 describes the definition for the ROVR register.



**Figure 15-72. Receive Oversize Packet Counter Register Definition**


Table 15-76 describes the fields of the ROVR register.

**Table 15-76. ROVR Field Descriptions**

Bits	Name	Description
0–15	—	Reserved
16–31	ROVR	Receive oversize packet counter. Increments each time a frame is received which exceeded 1518 (non VLAN) or 1522 (VLAN) and contains a valid FCS and was otherwise well formed.

### 15.5.3.6.22 Receive Fragments Counter (RFRG)

Figure 15-73 describes the definition for the RFRG register.



**Figure 15-73. Receive Fragments Counter Register Definition**


Table 15-77 describes the fields of the RFRG register.

**Table 15-77. RFRG Field Descriptions**

Bits	Name	Description
0–15	—	Reserved
16–31	RFRG	Receive fragments counter. Increments for each frame received which is less than 64 bytes in length and contains an invalid FCS. This includes integral and non-integral lengths.

### 15.5.3.6.23 Receive Jabber Counter (RJBR)

Figure 15-74 describes the definition for the RJBR register.



**Figure 15-74. Receive Jabber Counter Register Definition**

Table 15-78 describes the fields of the RJBR register.

**Table 15-78. RJBR Field Descriptions**

Bits	Name	Description
0–15	—	Reserved
16–31	RJBR	Receive jabber counter. Increments for frames received which exceed 1518 (non VLAN) or 1522 (VLAN) bytes and contain an invalid FCS. This includes alignment errors.



### 15.5.3.6.24 Receive Dropped Packet Counter (RDRP)

Figure 15-75 describes the definition for the RDRP register.



Figure 15-75. Receive Dropped Packet Counter Register Definition

Table 15-79 describes the fields of the RDRP register.

Table 15-79. RDRP Field Descriptions

Bits	Name	Description
0–15	—	Reserved
16–31	RDRP	Receive dropped packets counter. Increments for frames received which are streamed to system but are later dropped due to lack of system resources.

### 15.5.3.6.25 Transmit Byte Counter (TBYT)

Figure 15-76 depicts the TBYT register.




Figure 15-76. Transmit Byte Counter Register Definition


Table 15-80 describes the fields of the TBYT register.

Table 15-80. TBYT Field Descriptions

Bits	Name	Description
0–31	TBYT	Transmit byte counter. Increments by the number of bytes that were put on the wire including fragments of frames that were involved with collisions. This count does not include preamble/SFD or jam bytes, except for half-duplex flow control (back-pressure triggered by TCTRL[THDF] = 1). For THDF, the sum total of 'phantom' preamble bytes transmitted for flow control purposes is included in the TBYT increment value of the next frame to be transmitted, up to 65,535 bytes of frame and phantom preamble. Note that the value of TBYT may be greater than the actual number of bytes transmitted if the frame is truncated because it exceeds MAXFRM.

### 15.5.3.6.26 Transmit Packet Counter (TPKT)

Figure 15-77 describes the definition for the TPKT register.



**Figure 15-77. Transmit Packet Counter Register Definition**


Table 15-81 describes the fields of the TPKT register.

**Table 15-81. TPKT Field Descriptions**

Bits	Name	Description
0–9	—	Reserved
10–31	TPKT	Transmit packet counter. Increments for each transmitted packet (including bad packets, excessive deferred packets, excessive collision packets, late collision packets, all unicast, broadcast, and multicast packets).

### 15.5.3.6.27 Transmit Multicast Packet Counter (TMCA)

Figure 15-78 describes the definition for the TMCA register.



**Figure 15-78. Transmit Multicast Packet Counter Register Definition**


Table 15-82 describes the fields of the TMCA register.

**Table 15-82. TMCA Field Descriptions**

Bits	Name	Description
0–9	—	Reserved
10–31	TMCA	Transmit multicast packet counter. Increments for each multicast valid frame transmitted (excluding broadcast frames) with valid CRC and of lengths 64 to 1518 (non VLAN) or 1522 (VLAN).

### 15.5.3.6.28 Transmit Broadcast Packet Counter (TBCA)

Figure 15-79 describes the definition for the TBCA register.



**Figure 15-79. Transmit Broadcast Packet Counter Register Definition**


Table 15-83 describes the fields of the TBCA register.

**Table 15-83. TBCA Field Descriptions**

Bits	Name	Description
0–9	—	Reserved
10–31	TBCA	Transmit broadcast packet counter. Increments for each broadcast frame transmitted (excluding multicast frames) with valid CRC and of lengths 64 to 1518 (non VLAN) or 1522 (VLAN).

### 15.5.3.6.29 Transmit Pause Control Frame Counter (TXPF)

Figure 15-80 describes the definition for the TXPF register.



**Figure 15-80. Transmit Pause Control Frame Counter Register Definition**


Table 15-84 describes the fields of the TXPF register.

**Table 15-84. TXPF Field Descriptions**

Bits	Name	Description
0–15	—	Reserved
16–31	TXPF	Transmit PAUSE frame packet counter. Increments each time a valid PAUSE MAC control frame is transmitted with valid CRC and of lengths 64 to 1518 (non VLAN) or 1522 (VLAN).

### 15.5.3.6.30 Transmit Deferral Packet Counter (TDFR)

Figure 15-81 describes the definition for the TDFR register.



**Figure 15-81. Transmit Deferral Packet Counter Register Definition**


Table 15-85 describes the fields of the TDFR register.

**Table 15-85. TDFR Field Descriptions**

Bits	Name	Description
0–19	—	Reserved
20–31	TDFR	Transmit deferral packet counter. Increments for each frame, which was deferred on its first transmission attempt. This count does not include frames involved in collisions.

### 15.5.3.6.31 Transmit Excessive Deferral Packet Counter (TEDF)

Figure 15-82 describes the definition for the TEDF register.



**Figure 15-82. Transmit Excessive Deferral Packet Counter Register Definition**

Table 15-86 describes the fields of the TEDF register.

**Table 15-86. TEDF Field Descriptions**

Bits	Name	Description
0–19	—	Reserved
20–31	TEDF	Transmit excessive deferral packet counter. Increments for frames aborted which were deferred for an excessive period of time (3036 byte times).

### 15.5.3.6.32 Transmit Single Collision Packet Counter (TSCL)

Figure 15-83 describes the definition for the TSCL register.




Figure 15-83. Transmit Single Collision Packet Counter Register Definition

Table 15-87 describes the fields of the TSCL register.

Table 15-87. TSCL Field Descriptions

Bits	Name	Description
0–19	—	Reserved
20–31	TSCL	Transmit single collision packet counter. Increments for each frame transmitted which experienced exactly one collision during transmission.

### 15.5.3.6.33 Transmit Multiple Collision Packet Counter (TMCL)

Figure 15-84 describes the definition for the TMCL register.



Figure 15-84. Transmit Multiple Collision Packet Counter Register Definition


Table 15-88 describes the fields of the TMCL register.

Table 15-88. TMCL Field Descriptions

Bits	Name	Description
0–19	—	Reserved
20–31	TMCL	Transmit multiple collision packet counter. Increments for each frame transmitted which experienced 2–15 collisions (including any late collisions) during transmission as defined using the Half_Duplex[RETRANSMISSION MAXIMUM] field.

### 15.5.3.6.34 Transmit Late Collision Packet Counter (TLCL)

Figure 15-85 describes the definition for the TLCL register.



**Figure 15-85. Transmit Late Collision Packet Counter Register Definition**


Table 15-89 describes the fields of the TLCL register.

**Table 15-89. TLCL Field Descriptions**

Bits	Name	Description
0–19	—	Reserved
20–31	TLCL	Transmit late collision packet counter. Increments for each frame transmitted which experienced a late collision during a transmission attempt. Late collisions are defined using the collision window field of the half-duplex [26:31] register.

### 15.5.3.6.35 Transmit Excessive Collision Packet Counter (TXCL)

Figure 15-86 describes the definition for the TXCL register.



**Figure 15-86. Transmit Excessive Collision Packet Counter Register Definition**

Table 15-90 describes the fields of the TXCL register.

**Table 15-90. TXCL Field Descriptions**

Bits	Name	Description
0–19	—	Reserved
20–31	TXCL	Transmit excessive collision packet counter. Increments for each frame that experienced 16 collisions during transmission and was aborted.

### 15.5.3.6.36 Transmit Total Collision Counter (TNCL)

Figure 15-87 describes the definition for the TNCL register.




Figure 15-87. Transmit Total Collision Counter Register Definition

Table 15-91 describes the fields of the TNCL register.

Table 15-91. TNCL Field Descriptions

Bits	Name	Description
0–19	—	Reserved
20–31	TNCL	Transmit total collision counter. Increments by the number of collisions experienced during the transmission of a frame as defined as the simultaneous presence of signals on the DO and RD circuits (That is, transmitting and receiving at the same time). <b>Note:</b> This count does not include collisions that result in an excessive collision condition.

### 15.5.3.6.37 Transmit Drop Frame Counter (TDRP)

Figure 15-88 describes the definition for the TDRP register.



Figure 15-88. Transmit Drop Frame Counter Register Definition

Table 15-92 describes the fields of the TDRP register.

Table 15-92. TDRP Field Descriptions

Bits	Name	Description
0–15	—	Reserved
16–31	TDRP	Transmit drop frame counter. Increments each time a memory error or an underrun has occurred.

### 15.5.3.6.38 Transmit Jabber Frame Counter (TJBR)

Figure 15-89 describes the definition for the TJBR register.




Figure 15-89. Transmit Jabber Frame Counter Register Definition

Table 15-93 describes the fields of the TJBR register.

Table 15-93. TJBR Field Descriptions

Bits	Name	Description
0–19	—	Reserved
20–31	TJBR	Transmit jabber frame counter. Increments for each oversized transmitted frame with an incorrect FCS value.

### 15.5.3.6.39 Transmit FCS Error Counter (TFCS)

Figure 15-90 describes the definition for the TFCS register.




Figure 15-90. Transmit FCS Error Counter Register Definition

Table 15-94 describes the fields of the TFCS register.


Table 15-94. TFCS Field Descriptions

Bits	Name	Description
0–19	—	Reserved
20–31	TFCS	Transmit FCS error counter. Increments for every valid sized packet with an incorrect FCS value.



### 15.5.3.6.40 Transmit Control Frame Counter (TXCF)

Figure 15-91 describes the definition for the TXCF register.



**Figure 15-91. Transmit Control Frame Counter Register Definition**

Table 15-95 describes the fields of the TXCF register.

**Table 15-95. TXCF Field Descriptions**

Bits	Name	Description
0–19	—	Reserved
20–31	TXCF	Transmit control frame counter. Increments for every control frame with valid CRC and of lengths 64 to 1518 (non VLAN) or 1522 (VLAN).

### 15.5.3.6.41 Transmit Oversize Frame Counter (TOVR)

Figure 15-92 describes the definition for the TOVR register.



**Figure 15-92. Transmit Oversized Frame Counter Register Definition**


Table 15-96 describes the fields of the TOVR register.

**Table 15-96. TOVR Field Descriptions**

Bits	Name	Description
0–19	—	Reserved
20–31	TOVR	Transmit oversize frame counter. Increments for each oversized transmitted frame with a correct FCS value.

### 15.5.3.6.42 Transmit Undersize Frame Counter (TUND)

Figure 15-93 describes the definition for the TUND register.



**Figure 15-93. Transmit Undersize Frame Counter Register Definition**


Table 15-97 describes the fields of the TUND register.

**Table 15-97. TUND Field Descriptions**

Bits	Name	Description
0–19	—	Reserved
20–31	TUND	Transmit undersize frame counter. Increments for every frame less than 64 bytes, with a correct FCS value.

### 15.5.3.6.43 Transmit Fragment Counter (TFRG)

Figure 15-94 describes the definition for the TFRG register.



**Figure 15-94. Transmit Fragment Counter Register Definition**

Table 15-98 describes the fields of the TFRG register.

**Table 15-98. TFRG Field Descriptions**

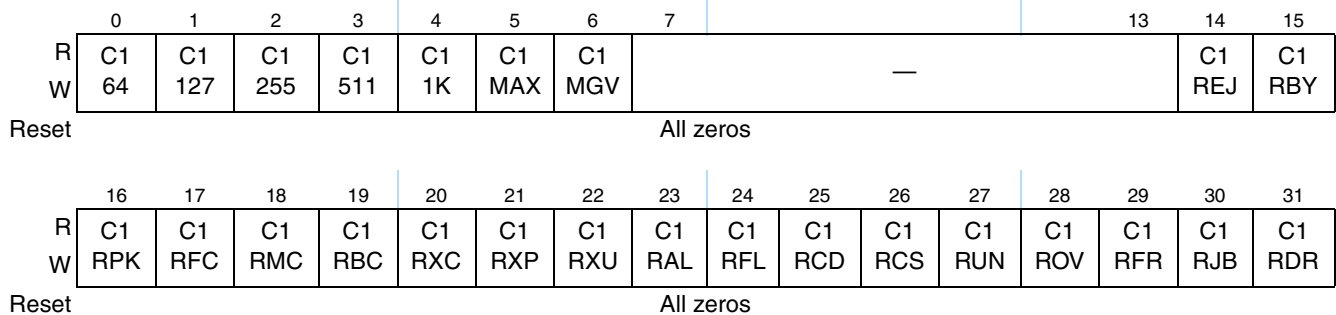
Bits	Name	Description
0–19	—	Reserved
20–31	TFRG	Transmit fragment counter. Increments for every frame less than 64 bytes, with an incorrect FCS value.

### 15.5.3.6.44 Carry Register 1 (CAR1)

Carry register bits are cleared on carry register writes when the respective bits are set. [Figure 15-95](#) describes the definition for the CAR1 register.

Offset eTSEC1:0x2\_4730; eTSEC2:0x2\_5730

Access: Read/Write



**Figure 15-95. Carry Register 1 (CAR1) Register Definition**

[Table 15-99](#) describes the fields of the CAR1 register.

**Table 15-99. CAR1 Field Descriptions**

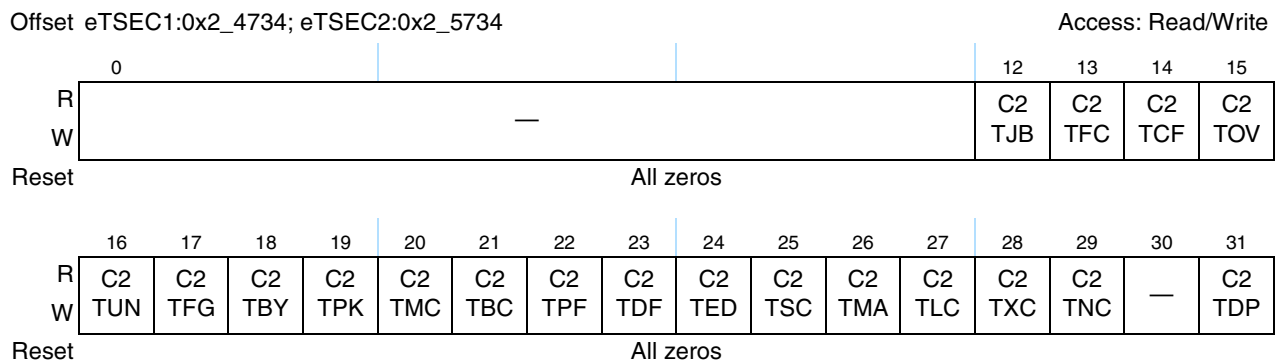
Bits	Name	Description
0	C164	Carry register 1 TR64 counter carry bit
1	C1127	Carry register 1 TR127 counter carry bit
2	C1255	Carry register 1 TR255 counter carry bit
3	C1511	Carry register 1 TR511 counter carry bit
4	C11K	Carry register 1 TR1K counter carry bit
5	C1MAX	Carry register 1 TRMAX counter carry bit
6	C1MGV	Carry register 1 TRMGV counter carry bit
7–13	—	Reserved
14	C1REJ	Carry register 1 RREJ counter carry bit
15	C1RBY	Carry register 1 RBYT counter carry bit
16	C1RPK	Carry register 1 RPKT counter carry bit
17	C1RFC	Carry register 1 RFCS counter carry bit
18	C1RMC	Carry register 1 RMCA counter carry bit
19	C1RBC	Carry register 1 RBCA counter carry bit
20	C1RXC	Carry register 1 RXCF counter carry bit
21	C1RXP	Carry register 1 RXPf counter carry bit
22	C1RXU	Carry register 1 RXUO counter carry bit
23	C1RAL	Carry register 1 RALN counter carry bit
24	C1RFL	Carry register 1 RFLR counter carry bit

**Table 15-99. CAR1 Field Descriptions (continued)**

Bits	Name	Description
25	C1RCD	Carry register 1 RCDE counter carry bit
26	C1RCS	Carry register 1 RCSE counter carry bit
27	C1RUN	Carry register 1 RUND counter carry bit
28	C1ROV	Carry register 1 ROVR counter carry bit
29	C1RFR	Carry register 1 RFRG counter carry bit
30	C1RJB	Carry register 1 RJBR counter carry bit
31	C1RDR	Carry register 1 RDRP counter carry bit

### 15.5.3.6.45 Carry Register 2 (CAR2)

Figure 15-96 describes the definition for the CAR2 register.

**Figure 15-96. Carry Register 2 (CAR2) Register Definition**

Carry register bits are cleared on carry register write when the respective bits are set. Table 15-100 describes the fields of the CAR2 register.

**Table 15-100. CAR2 Field Descriptions**

Bits	Name	Description
0–11	—	Reserved
12	C2TJB	Carry register 2 TJBR counter carry bit
13	C2TFC	Carry register 2 TFCS counter carry bit
14	C2TCF	Carry register 2 TXCF counter carry bit
15	C2TOV	Carry register 2 TOVR counter carry bit
16	C2TUN	Carry register 2 TUND counter carry bit
17	C2TFG	Carry register 2 TFRG counter carry bit
18	C2TBY	Carry register 2 TBYT counter carry bit
19	C2TPK	Carry register 2 TPKT counter carry bit

**Table 15-100. CAR2 Field Descriptions (continued)**

Bits	Name	Description
20	C2TMC	Carry register 2 TMCA counter carry bit
21	C2TBC	Carry register 2 TBCA counter carry bit
22	C2TPF	Carry register 2 TXPF counter carry bit
23	C2TDF	Carry register 2 TDFR counter carry bit
24	C2TED	Carry register 2 TEDF counter carry bit
25	C2TSC	Carry register 2 TSCL counter carry bit
26	C2TMA	Carry register 2 TMCL counter carry bit
27	C2TLC	Carry register 2 TLCL counter carry bit
28	C2TXC	Carry register 2 TXCL counter carry bit
29	C2TNC	Carry register 2 TNCL counter carry bit
30	—	Reserved, should be cleared
31	C2TDP	Carry register 2 TDRP counter carry bit

### 15.5.3.6.46 Carry Mask Register 1 (CAM1)

While one of the below mask bits are cleared, the corresponding carry bit in CAR1 is allowed to cause interrupt indications in register IEVENT[MSR0]. These bits all default to a set state. [Figure 15-97](#) describes the definition for the CAM1 register.

Offset eTSEC1:0x2\_4738; eTSEC2:0x2\_5738 Access: Read/Write

	0	1	2	3	4	5	6	7						13	14	15
R	M1	M1	M1	M1	M1	M1	M1								M1	M1
W	64	127	255	511	1K	MAX	MGV								REJ	RBY
Reset	1	1	1	1	1	1	1	0	0	0	0	0	0	0	1	1
	16	17	18	19	20	21	22	23	24	25	26	27	28	29	30	31
R	M1	M1	M1	M1	M1	M1	M1	M1	M1	M1	M1	M1	M1	M1	M1	M1
W	RPK	RFC	RMC	RBC	RXC	RXP	RXU	RAL	RFL	RCD	RCS	RUN	ROV	RFR	RJB	RDR
Reset	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1

**Figure 15-97. Carry Mask Register 1 (CAM1) Register Definition**

[Table 15-101](#) describes the fields of the CAM1 register.

**Table 15-101. CAM1 Field Descriptions**


Bits	Name	Description
0	M164	Mask register 1 TR64 counter carry bit mask
1	M1127	Mask register 1 TR127 counter carry bit mask
2	M1255	Mask register 1 TR255 counter carry bit mask
3	M1511	Mask register 1 TR511 counter carry bit mask

**Table 15-101. CAM1 Field Descriptions (continued)**

Bits	Name	Description
4	M11k	Mask register 1 TR1K counter carry bit mask
5	M1MAX	Mask register 1 TRMAX counter carry bit mask
6	M1MGV	Mask register 1 TRMGV counter carry bit mask
7–13	—	Reserved
14	M1REJ	Mask register 1 RREJ counter carry bit mask
15	M1RBY	Mask register 1 RBYT counter carry bit mask
16	M1RPK	Mask register 1 RPKT counter carry bit mask
17	M1RFC	Mask register 1 RFCS counter carry bit mask
18	M1RMC	Mask register 1 RMCA counter carry bit mask
19	M1RBC	Mask register 1 RBCA counter carry bit mask
20	M1RXC	Mask register 1 RXCF counter carry bit mask
21	M1RXP	Mask register 1 RXPF counter carry bit mask
22	M1RXU	Mask register 1 RXUO counter carry bit mask
23	M1RAL	Mask register 1 RALN counter carry bit mask
24	M1RFL	Mask register 1 RFLR counter carry bit mask
25	M1RCD	Mask register 1 RCDE counter carry bit mask
26	M1RCS	Mask register 1 RCSE counter carry bit mask
27	M1RUN	Mask register 1 RUND counter carry bit mask
28	M1ROV	Mask register 1 ROVR counter carry bit mask
29	M1RFR	Mask register 1 RFRG counter carry bit mask
30	M1RJB	Mask register 1 RJBR counter carry bit mask
31	M1RDR	Mask register 1 RDRP counter carry bit mask

### 15.5.3.6.47 Carry Mask Register 2 (CAM2)

While one of the below mask bits are cleared, the corresponding carry bit in CAR2 is allowed to cause interrupt indications in register IEVENT[MSR0]. These bits default to a set state. Figure 15-98 describes the definition for the CAM2 register.



**Figure 15-98. Carry Mask Register 2 (CAM2) Register Definition**

Table 15-102 describes the fields of the CAM2 register.

**Table 15-102. CAM2 Field Descriptions**

Bits	Name	Description
0–11	—	Reserved
12	M2TJB	Mask register 2 TJBR counter carry bit mask
13	M2TFC	Mask register 2 TFCS counter carry bit mask
14	M2TCF	Mask register 2 TXCF counter carry bit mask
15	M2TOV	Mask register 2 TOVR counter carry bit mask
16	M2TUN	Mask register 2 TUND counter carry bit mask
17	M2TFG	Mask register 2 TFRG counter carry bit mask
18	M2TBY	Mask register 2 TBYT counter carry bit mask
19	M2TPK	Mask register 2 TPKT counter carry bit mask
20	M2TMC	Mask register 2 TMCA counter carry bit mask
21	M2TBC	Mask register 2 TBCA counter carry bit mask
22	M2TPF	Mask register 2 TXPF counter carry bit mask
23	M2TDF	Mask register 2 TDFR counter carry bit mask
24	M2TED	Mask register 2 TEDF counter carry bit mask
25	M2TSC	Mask register 2 TSCL counter carry bit mask
26	M2TMA	Mask register 2 TMCL counter carry bit mask
27	M2TLC	Mask register 2 TLCL counter carry bit mask
28	M2TXC	Mask register 2 TXCL counter carry bit mask
29	M2TNC	Mask register 2 TNCL counter carry bit mask

**Table 15-102. CAM2 Field Descriptions (continued)**

Bits	Name	Description
30	—	Reserved
31	M2TDP	Mask register 2 TDRP counter carry bit mask

### 15.5.3.6.48 Receive Filer Rejected Packet Counter (RREJ)

Figure 15-99 describes the definition for the RREJ register.


**Figure 15-99. Receive Filer Rejected Packet Counter Register Definition**

Table 15-67 describes the fields of the RREJ register.

**Table 15-103. RREJ Field Descriptions**

Bits	Name	Description
0–9	—	Reserved
10–31	RREJ	Receive filer rejected packet counter. Increments for each frame with valid CRC received, but rejected by the receive queue filer—either due to a matching rule that asserted the REJ flag or due to filing to a RxBD ring that was not enabled (see IEVENT[FIQ] error).

### 15.5.3.7 Hash Function Registers

This section provides detailed descriptions of the registers used for hash functions. All of the registers are 32 bits wide. The DA field of every received frame is processed through a 32-bit CRC generator (CRC-32 polynomial), and the 8 or 9 most significant bits of the CRC are mapped to a hash table entry. The user can enable a hash entry by setting its bit. A hash entry usually represents a set of addresses. A hash table hit occurs if the DA CRC result points to an enabled hash entry. Software may need to further filter the address in order to eliminate false-positive hits in the hash table.

If RCTRL[GHTX] = 0, the 8 most significant bits of the CRC are used as the hash table index. In this case, registers IGADDR0–IGADDR7 comprise a 256-entry hash table exclusively for individual (unicast) address matching, while registers GADDR0–GADDR7 comprise a 256-entry hash table for group (multicast) address matching. If RCTRL[GHTX] = 1, the group hash table is extended to all 512 entries, and the 9 most significant bits of the CRC are used as the hash table index. In this case, registers IGADDR0–IGADDR7 hold hash table entries 0–255 for group addresses, while registers GADDR0–GADDR7 hold entries 256–511 of the extended group hash table.

See Section 15.6.2.7.2, “Hash Table Algorithm,” for more information on the hash algorithm.



### 15.5.3.7.1 Individual/Group Address Registers 0–7 (IGADDR $n$ )

The IGADDR $n$  registers are written by the user. Together these registers represent, depending on RCTRL[GHTX], either the 256 entries of the individual address hash table, or the first 256 entries of the extended group address hash table used in the address recognition process. The user can enable a hash entry by setting the appropriate bit. A hash table hit occurs if the DA CRC-32 result points to an enabled hash entry.

Figure 15-100 describes the definition for the IGADDR $n$  register.




Figure 15-100. IGADDR $n$  Register Definition

Table 15-105 describes the fields of the IGADDR $n$  register.

Table 15-104. IGADDR $n$  Field Descriptions

Bits	Name	Description
0–31	IGADDR $n$	Represents the 32-bit value associated with the corresponding register. When RCTRL[GHTX] = 0, IGADDR0 contains entries 0–31 of the 256-entry individual hash table and IGADDR7 represents entries 224–255. When RCTRL[GHTX] = 1, IGADDR0 contains entries 0–31 of the 512-entry extended group hash table and IGADDR7 represents entries 224–255.

### 15.5.3.7.2 Group Address Registers 0–7 (GADDR $n$ )

The GADDR $n$  registers are written by the user. Together these registers represent, depending on RCTRL[GHTX], either the 256 entries of the group address hash table, or the last 256 entries of the extended group address hash table used in the address recognition process. The user can enable a hash entry by setting the appropriate bit. A hash table hit occurs if the DA CRC result points to an enabled hash entry. Figure 15-101 describes the definition for the GADDR $n$  register.




Figure 15-101. GADDR $n$  Register Definition

Table 15-105 describes the fields of the GADDR $n$  register.

**Table 15-105. GADDR $n$  Field Descriptions**

Bits	Name	Description
0–31	GADDR $n$	Represents the 32-bit value associated with the corresponding register. When RCTRL[GHTX] = 0, GADDR0 contains entries 0–31 of the 256-entry group hash table and GADDR7 represents entries 224–255. When RCTRL[GHTX] = 1, GADDR0 contains entries 256–287 of the 512-entry extended group hash table and GADDR7 represents entries 480–511.


### 15.5.3.8 DMA Attribute Registers

This section describes the two eTSEC DMA attribute registers.

#### 15.5.3.8.1 Attribute Register (ATTR)

The attribute register defines memory access attributes and transaction types used to access buffer descriptors, to write receive data, and to read transmit data. Snoop enable attributes may be set for reading buffer descriptors and for reading transmit data.

Figure 15-102 describes the definition for the ATTR register.



**Figure 15-102. ATTR Register Definition**

Table 15-106 describes the fields of the ATTR register.

**Table 15-106. ATTR Field Descriptions**

Bits	Name	Description
0–23	—	Reserved
24	RDSSEN	Rx data snoop enable. 0 Disables snooping of all receive frames data to memory. 1 Enables snooping of all receive frames data to memory.
25	RBDSSEN	RxBD snoop enable. 0 Disables snooping of all receive BD memory accesses. 1 Enables snooping of all receive BD memory accesses.
26–31	—	Reserved

### 15.5.3.9 Lossless Flow Control Configuration Registers

When enabled through RCTRL[LFC], the eTSEC tracks location of the last free BD in each Rx BD ring through the value of RFBPTR $n$ . Using this pointer and the ring length stored in RQPRM $n$ [LEN], the eTSEC continuously calculates the number of free BDs in the ring. Whenever the calculated number of free BDs in the ring drops below the pause threshold specified in RQPRM $n$ [FBTHR], the eTSEC issues link layer flow control. It continues to assert flow control until the free BD count for each active ring reaches or exceeds RQPRM $n$ [FBTHR]. See section 15.6.5.1/15-175 for the theory of operation of these registers.

#### 15.5.3.9.1 Receive Queue Parameters 0–7 (RQPRM0–PQPRM7)

The RQPRM $n$  registers specify the minimum number of BDs required to prevent flow control being asserted and the total number of Rx BDs in their respective ring. Whenever the free BD count calculated by the eTSEC for any active ring drops below the value of RQPRM $n$ [FBTHR] for that ring, link level flow control is asserted. Software must not write to RQPRM $n$  while LFC is enabled and the eTSEC is actively receiving frames. However, software may modify these registers after disabling LFC by clearing RCTRL[LFC]. Note that packets may be lost due to lack of RxBDs while RCTRL[LFC] is clear. Software can prevent packet loss by manually generating pause frames (through TCTRL[TFC\_PAUSE]) to cover the time when RCTRL[LFC] is clear. Figure 15-103 describes the definition for the RQPRM $n$  register.




Figure 15-103. RQPRM Register Definition

Table 15-107 describes the fields of the RQPRM register.

Table 15-107. RQPRM Field Descriptions


Bits	Name	Description
0–7	FBTHR	Free BD threshold. Minimum number of BDs required for normal operation. If the eTSEC calculated number of free BDs drops below this threshold, link layer flow control is asserted.
8–31	LEN	Ring length. Total number of Rx BDs in this ring.

#### 15.5.3.9.2 Receive Free Buffer Descriptor Pointer Registers 0–7 (RFBPTR0–RFBPTR7)

The RFBPTR $n$  registers specify the location of the last free buffer descriptor in their respective ring. These registers live in the same 32b address space – and must share the same 4 most significant bits – as RBPTR $n$ . That is, RFBPTR $n$  and its associated RBPTR $n$  must remain in the same 256MB page. Like RBPTR $n$ , whenever RBASE $n$  is updated, RFBPTR $n$  is initialized to the value of RBASE $n$ . This indicates that the ring is completely empty. As buffers are freed and their respective BDs are returned (by setting the EMPTY bit) to the ring, software is expected to update this register. The eTSEC then performs modulo arithmetic involving RBASE $n$ , RBPTR $n$  and RFBPTR $n$  to determine the number of free BDs remaining in the ring.

If, at any stage, the value written to  $RFBPTR_n$  matches that of the respective  $RBPTR_n$  the eTSEC free BD calculation assumes that the ring is now completely empty. For more information on the recommended use of these registers, see [Section 15.6.5.1, “Back Pressure Determination through Free Buffers.”](#)

[Figure 15-104](#) describes the definition for the  $RFBPTR_n$  register.



**Figure 15-104. RFBPTR0–RFBPTR7 Register Definition**

[Table 15-108](#) describes the fields of the  $RFBPTR_n$  registers.

**Table 15-108. RFBPTR0–RFBPTR7 Field Descriptions**

Bits	Name	Description
0–28	RFBPTR	Pointer to the last free BD in RxBD Ring <i>n</i> . When $RBASE_n$ is updated, eTSEC initializes RFBPTR $_n$ to the value in the corresponding $RBASE_n$ . Software may update this register at any time to inform the eTSEC the location of the last free BD in the ring. Note that the 3 least-significant bits of this register are read only and zero.
29–31	—	Reserved.

### 15.5.3.10 Hardware Assist for IEEE1588 Compliant Timestamping

IEEE 1588 compliant timestamping on this device is accomplished using the per-port transmit timestamping registers within each Ethernet controller memory space (See [Section 15.5.3.2.11, “Transmit Timestamp Identification Register \(TMR\\_TXTS1–2\\_ID\),”](#) and [Section 15.5.3.2.12, “Transmit Timestamp Register \(TMR\\_TXTS1–2\\_H/L\).”](#)) in conjunction with the following common registers, which are located within the memory space for eTSEC1. Because the common 1588 timestamping registers exist within the eTSEC1 memory space, the eTSEC1 controller must remain enabled in order to use 1588 timestamping for any Ethernet port.

#### 15.5.3.10.1 Timer Control Register (TMR\_CTRL)

This register is used to reset, configure, and initialize the eTSEC precision timer clock. The control of all timer function is performed via programming eTSEC1. The register in eTSEC1 is shared for all eTSECs. [Figure 15-7](#) describes the definition for the TMR\_CTRL register.

Register fields not described below are reserved.

Offset eTSEC1:0x2\_4E00

Access: Mixed

	0	1	2	3	4	5	6											15
R	ALM1P	ALM2P	—	FS	PP1L	PP2L	TCLK_PERIOD											
W																		
Reset	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	1	
	16	17	18	19	20	21	22	23	24	25	26	27	28	29	30	31		
R	RTPE	FRD			ESFDP	ESFDE	ETEP2	ETEP1	COPH	CIPH	TMSR		BYP	TE	CKSEL			
W																		
Reset	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	1		

Figure 15-105. TMR\_CTRL Register Definition

Table 15-109 describes the fields of the TMR\_CTRL register. Register fields not described below are reserved.

Table 15-109. TMR\_CTRL Register Field Descriptions

Bits	Name	Description
0	ALM1P	Alarm1 output polarity 0 active high output 1 active low output
1	ALM2P	Alarm2 output polarity 0 active high output 1 active low output
3	FS	FIPER start indication 0 Fiper is enabled through timer enable 1 Fiper is enabled through timer enable and alarm indication.
4	PP1L	Fiper1 pulse loopback mode enabled. 0 Trigger1 input is based upon normal external trigger input. 1 Fiper1 pulse is looped back into Trigger1 input.
5	PP2L	Fiper2 pulse loopback mode enabled. 0 Trigger2 input is based upon normal external trigger input. 1 Fiper2 pulse is looped back into Trigger2 input.
6–15	TCLK_PERIOD	1588 timer reference clock period. The timer clock counter will increment by TCLK_PERIOD every time the accumulator register overflows. This clock period must be larger than the clock period of the timer reference clock. For applications where user does not want the clock period to be added, they can program this field to 1 to count the clock ticks. This field defaulted to 1 to count overflow ticks. For nanosecond granularity on 1588 timer counter rate, the TCLK_PERIOD should be calculated using the following equation: $TCLK\_PERIOD = 10^9 / \text{Nominal\_Frequency}$
16	RTPE	Record Tx Timestamp to PAL Enable. When set, and FCB[PTP] is set, the 8-byte timestamp for the packet is written to the PAL located in external memory location at an offset of 16 bytes from the start of the Data Buffer Pointer of the first TxBD. For guidelines on using the RTPE bit, refer to <a href="#">Section 15.6.6.5, “Timestamp Insertion on Transmit Packets.”</a>
17	FRD	FIPER Realignment Disable 0 Fiper Realignment is enabled. 1 Fiper Realignment is disabled.


**Table 15-109. TMR\_CTRL Register Field Descriptions (continued)**

Bits	Name	Description
20	ESFDP	External Tx/Rx SFD Polarity. 0 Timestamp on rising edge of external SFD indication. 1 Timestamp on falling edge of external SFD indication.
21	ESFDE	External Tx/Rx SFD Enable. 0 Timestamp PTP TX frame based on MAC's SFD indication. 1 Timestamp PTP TX frame based on external SFD indication from PHY.
22	ETEP2	External trigger 2 edge polarity 0 Timestamp on the rising edge of the external trigger 1 Timestamp on the falling edge of the external trigger
23	ETEP1	External trigger 1 edge polarity 0 Timestamp on the rising edge of the external trigger 1 Timestamp on the falling edge of the external trigger
24	COPH	Generated clock (TSEC_1588_GCLK) output phase. 0 non-inverted divided clock is output 1 inverted divided clock is output
25	CIPH	External oscillator input clock phase. 0 non-inverted frequency tuned timer input clock 1 inverted frequency tuned timer input clock (NOTE: this setting is reserved if CKSEL=01.)
26	TMSR	Timer soft reset. When enabled, it resets all the timer registers and state machines. 0 normal operation 1 place entire timer in reset except control and config registers NOTE: Prior to initiating timer reset (setting TMSR), must gracefully stop receiver (See MACCFG1[RX_EN] description). User programmable registers are not reset by the soft reset e.g. TMR_CTRL, TMR_TEMASK, TMR_PEMASK, TMR_ADD, TMR_PRSC, TMROFF_H/L, TMR_ALARMn, and TMR_FIPERn.
28	BYP	Bypass drift compensated clock 0 64-bit clock counter is incremented on the accumulator overflow 1 64-bit clock counter is directly driven from the external oscillator ignoring accumulator overflow
29	TE	1588 timer enable. If not enabled, all the timer registers and state machines are disabled. 0 timer not enabled 1 timer enabled and resume normal operation
30–31	CKSEL	1588 Timer reference clock source select. 00 External high precision timer reference clock (TSEC_TMR_CLK) 01 eTSEC system clock 10 Reserved 11 RTC clock input. Note that the 1588 reference clock must be no slower than 1/7 the Rx_clk frequency. The default clock select is eTSEC system clock, which is always active when eTSEC is enabled. The user must ensure the corresponding clock source is active before changing the 1588 refclk selection to external reference, RTC, or TX clock. Selecting an inactive 1588 reference clock may cause boundedly undefined behavior in the ethernet controller and on accesses to the 1588 registers.

### 15.5.3.10.2 Timer Event Register (TMR\_TEVENT)

The eTSEC precision timer implementation can generate additional interrupts that are independent of the frame based events that controlled via IEVENT. The timer interrupts are not affected by any interrupt

coalescing that may be specified in TXIC/RXIC. Software may poll this register at any time to check for pending interrupts. If an event occurs and its corresponding enable bit is set in the event mask register (TEMASK), the event also causes a hardware interrupt at the PIC. A bit in the timer event register is cleared by writing a 1 to that bit position. Figure 15-4 describes the definition for the TMR\_TEVENT register.



**Figure 15-106. TMR\_TEVENT Register Definition**

Table 15-110 describes the fields of the TMR\_TEVENT register fields for the timer.

**Table 15-110. TMR\_TEVENT Register Field Descriptions**

Bits	Name	Description
0–6	—	Reserved
6	ETS2	External trigger 2 timestamp sampled 0 external trigger timestamp not sampled 1 external trigger timestamp sampled
7	ETS1	External trigger 1 timestamp sampled 0 external trigger timestamp not sampled 1 external trigger timestamp sampled
8–13	—	Reserved
14	ALM2	Current time equaled alarm time register 2 0 alarm time has not be reached yet 1 alarm time has been reached
15	ALM1	Current time equaled alarm time register 1 0 alarm time has not be reached yet 1 alarm time has been reached
16–23	—	Reserved
24	PP1	Indicates that a periodic pulse has been generated based on FIPER1 register. 0 periodic pulse not generated 1 periodic pulse generated
25	PP2	Indicates that a periodic pulse has been generated based on FIPER2 register. 0 periodic pulse not generated 1 periodic pulse generated

**Table 15-110. TMR\_TEVENT Register Field Descriptions (continued)**

Bits	Name	Description
26	PP3	Indicates that a periodic pulse has been generated based on FIPER3 register. 0 periodic pulse not generated 1 periodic pulse generated
27–31	—	Reserved

**15.5.3.10.3 Timer Event Mask Register (TMR\_TEMASK)**

Timer event mask register. The event mask register provides control over which possible interrupt events in the TMR\_TEVENT register are permitted to participate in generating hardware interrupts to the PIC. All implemented bits in this register are R/W and cleared upon a hardware reset. Figure 15-111 describes the definition for the TMR\_TEMASK register.

Offset eTSEC1:0x2\_4E08

Access: Read/Write


**Table 15-111. TMR\_TEMASK Register Definition**

Table 15-112 describes the fields of the TMR\_TEMASK register fields for the timer.


**Table 15-112. TMR\_TEMASK Register Field Descriptions**

Bits	Name	Description
0–5	—	Reserved
6	ETS2EN	External trigger 2 timestamp sample event enable
7	ETS1EN	External trigger 1 timestamp sample event enable
8–13	—	Reserved
14	ALM2EN	Timer ALM1 event enable
15	ALM1EN	Timer ALM2 event enable
16–23	—	Reserved
24	PP1EN	Periodic pulse event 1 enable
25	PP2EN	Periodic pulse event 2 enable
26–31	—	Reserved



### 15.5.3.10.4 Timer PTP Packet Event Register (TMR\_PEVENT)

The eTSEC precision timer logic can generate interrupts upon the capture of a timestamp due to either transmission or reception of a frame. If an event occurs and its corresponding enable bit is set in the event mask register (PEMASK), the event also causes a hardware interrupt at the PIC. A bit in the timer event register is cleared by writing a 1 to that bit position. Figure 15-107 describes the definition for the TMR\_PEVENT register.



**Figure 15-107. TMR\_PEVENT Register Definition**

Table 15-113 describes the fields of the TMR\_PEVENT register fields for the timer.

**Table 15-113. TMR\_PEVENT Register Field Descriptions**

Bits	Name	Description
0–21	—	Reserved
22	TXP2	Indicates that a PTP frame has been transmitted and its timestamp is stored in TXTS2 register. 0 PTP packet not transmitted 1 PTP packet has been transmitted
23	TXP1	Indicates that a PTP frame has been transmitted and its timestamp is stored in TXTS1 register. 0 PTP packet not transmitted 1 PTP packet has been transmitted
24–30	—	Reserved
31	RXP	Indicates that a PTP frame has been received 0 PTP packet not received 1 PTP packet has been received

### 15.5.3.10.5 Timer Event Mask Register (TMR\_PEMASK)

Timer event mask register. The event mask register provides control over which possible interrupt events in the TMR\_PEVENT register are permitted to participate in generating hardware interrupts to the PIC. All implemented bits in this register are R/W and cleared upon a hardware reset. Figure 15-108 describes the definition for the TMR\_PEMASK register.

Offset eTSEC1:0x2\_4E10

Access: Read/Write


**Figure 15-108. TMR\_PEMASK Register Definition**

Table 15-114 describes the fields of the TMR\_PEMASK register fields for the timer.

**Table 15-114. TMR\_PEMASK Register Field Descriptions**

Bits	Name	Description
0–21	—	Reserved
22	TXP2EN	Transmit PTP packet event 2 enable
23	TXP1EN	Transmit PTP packet event 1 enable
24–30	—	Reserved
31	RXPEN	Receive PTP packet event enable

### 15.5.3.10.6 Timer Status Register (TMR\_STAT)

This register requires the eTSEC filer to be enabled (via RCTRL[FILREN]). When eTSEC generates an interrupt based on the timestamp event for a received packet, the queue ID which the incoming packet will be sent to is captured in this register. This register update is synchronized with the RXF interrupt of the corresponding received packet. Writing 1 to any bit of this register clears it. Figure 15-115 describes the definition for the TMR\_STAT register.


**Table 15-115. TMR\_STAT Register Definition**


Table 15-116 describes the fields of the TMR\_STAT register.

**Table 15-116. TMR\_STAT Register Field Descriptions**

Bits	Name	Description
0–25	—	Reserved
26–31	STAT_VEC	Timer general purpose status vector. It will store the 6-bit queue number generated by the filer. User to decode this status vector. For example, user can encode received PTP packet message types (Sync, Delay_req, Follow_up, Delay_resp, Management) in the filer virtual queue field.

### 15.5.3.10.7 Timer Counter Register (TMR\_CNT\_H/L)

The timer register (TMR\_CNT\_H/L) represents accurate time in terms clock ticks or in nano-seconds. Writes to these registers will override the previous time. The register in eTSEC1 is shared for all eTSECs. This is a read/write register. Figure 15-109 describes the definition for the TMR\_CNT\_H/L register.



**Figure 15-109. TMR\_CNT\_H Register Definition**

Table 15-117 describes the fields of the TMR\_CNT\_H/L register.


**Table 15-117. TMR\_CNT\_H/L Register Field Descriptions**

Bits	Name	Description
0–63	TMR_CNT_H/L	Value of the current time counter. Current time is calculated by adding TMROFF_H/L with the TMR_CNT_H/L counter. This register can be written through the register writes. Writes to the TMR_CNT_L register copies the written value into the shadow TMR_CNT_L register. Writes to the TMR_CNT_H register copies the values written into the shadow TMR_CNT_H register. Contents of the shadow registers are copied into the TMR_CNT_L and TMR_CNT_H registers following a write into the TMR_CNT_H register. Writes to these registers have precedence over the timer increment. The user must write to TMR_CNT_L register first.  Reads from the TMR_CNT_L register copies the entire 64-bit clock time of the read enable into the TMR_CNT_H/L shadow registers. Read instruction from the TMR_CNT_H register reads the value stored in the TMR_CNT_H shadow register. The user must read the TMR_CNT_L register first to get correct 64-bit TMR_CNT_H/L counter values.

### 15.5.3.10.8 Timer Drift Compensation Addend Register (TMR\_ADD)

Timer drift compensation addend register (TMR\_ADD) is used to hold timer frequency compensation value (FreqCompensationValue). The nominal frequency of the clock counter is determined by the FreqDivRatio and the clock frequency (FreqClock). This register is programmed with  $2^{32}/\text{FreqDivRatio}$ . Frequency division ratio (FreqDivRatio) is the ratio between the frequency of the oscillator (TimerOsc) and the desired clock frequency (NominalFreq). FreqDivRatio is a design constant chosen to be greater than 1.0001. The ADDEND value is added to the 32-bit accumulator register at every rising edge of the oscillator clock (TimerOsc). The clock counter is incremented at every carry pulse of the accumulator.

Only one of this register is required for the entire group of eTSECs. Figure 15-110 describes the definition of the TMR\_ADD register.



**Figure 15-110. TMR\_ADD Register Definition**


Table 15-118 describes the fields of the TMR\_ADD register fields for the timer.

**Table 15-118. TMR\_ADD Register Field Descriptions**

Bits	Name	Description
0–31	ADDEND	Timer drift compensation addend register value. It is programmed with a value of $2^{32}/\text{FreqDivRatio}$ . For example, TimerOsc = 50 MHz NominalFreq = 40 MHz FreqDivRatio = 1.25 ADDEND = $\text{ceil}(2^{32}/1.25) = 0xCCCC\_CCCD$

### 15.5.3.10.9 Timer Accumulator Register (TMR\_ACC)

Timer accumulator register accumulates the value of the addend register into it. An overflow pulse of the accumulator is used to increment the timer clock by TMR\_CTRL[TCLK\_PERIOD]. This register is read only in normal operation. The register in eTSEC1 is shared for all eTSECs. Figure 15-111 describes the definition of the TMR\_ACC register.



**Figure 15-111. TMR\_ACC Register Definition**

Table 15-119 describes the fields of the TMR\_ACC register.

**Table 15-119. TMR\_ACC Register Field Descriptions**

Bits	Name	Description
0–31	TMR_ACC	32-bit timer accumulator register

### 15.5.3.10.10 Timer Prescale Register (TMR\_PRSC)

Timer generated output clock prescale register. It is used to adjust output clock frequency that is put onto the 1588 clock output signal. The register in eTSEC1 is shared for all eTSECs. Figure 15-112 describes the definition for the TMR\_PRSC register.




Figure 15-112. TMR\_PRSC Register Definition

Table 15-120 describes the fields of the TMR\_PRSC register.

Table 15-120. TMR\_PRSC Register Field Descriptions

Bits	Name	Description
0–15	—	Reserved
16–31	PRSC_OCK	Output clock division/prescale factor. Output clock is generated by dividing the timer input clock by this number. Programmed value in this field must be greater than 1. Any value less than 1 is treated as 2.

### 15.5.3.10.11 Timer Offset Register (TMROFF\_H/L)

The timer offset register is used to provide current time by adding its value to the clock counter. Figure 15-113 describes the definition of the TMROFF\_H/L register.

#### NOTE

All TMROFF\_H registers in a device should be set to the same value, and all TMROFF\_L registers in a device should be set to the same value. Otherwise, the precision time protocol may not work.




Figure 15-113. TMROFF\_H/L Register Definition

Table 15-121 describes the fields of the TMROFF\_H/L register.

Table 15-121. TMROFF\_H/L Register Field Descriptions

Bits	Name	Description
0–63	TMROFF_H/L	Offset value of the clock counter. Current time in is calculated by adding TMROFF_H/L with the timer's counter TMR_CNT_H/L register.

### 15.5.3.10.12 Alarm Time Comparator Register (TMR\_ALARM1–2\_H/L)

Alarm time comparator register (TMR\_ALARM<sub>n</sub>\_H/L). This register holds alarm time for comparison with the current time counter. There are two of these registers for eTSEC1 which are shared amongst all eTSECs. Figure 15-114 describes the definition for the TMR\_ALARM<sub>n</sub>\_H/L register.

Offset eTSEC1:0x2\_4E40+8×n

Access: Read/Write




Figure 15-114. TMR\_ALARM1-2\_H/L Register Definition

Table 15-122 describes the fields of the TMR\_ALARM $n$ \_H/L register.

Table 15-122. TMR\_ALARM $n$ \_H/L Register Field Descriptions

Bits	Name	Description
0–63	ALARM_H/L	Alarm time comparator register. The corresponding alarm event in TMR_TEVENT is set when the current time counter becomes equal to or greater than the alarm time compare value in TMR_ALARM $n$ _L/H. Writing the TMR_ALARM $n$ _L register deactivates the alarm event after it has fired. Writing the TMR_ALARM $n$ _L followed by the TMR_ALARM $n$ _H register rearms the alarm function with the new compare value. The value programmed in this register must be an integer multiple of TMR_CTRL[TCLK_PERIOD] in order to get correct result. This register is reset to all ones to avoid false alarm after reset. In FS mode the alarm trigger is used as an indication to the fiper start down counting. Only alarm 1 supports this mode. In FS mode, alarm polarity bit should be configured to 0 (rising edge).

### 15.5.3.10.13 Timer Fixed Interval Period Register (TMR\_FIPER1–3)

Timer fixed interval period pulse generator register. It is used to generate periodic pulses. This register is reset with 0xFFFF\_FFFF to prevent any false pulse upon initialization. The down count register loads the value programmed in the fixed period interval (FIPER). FIPER register must be programmed before the timer is enabled. At every tick of the timer accumulator overflow, the counter decrements by the value of TMR\_CTRL[TCLK\_PERIOD]. It generates a pulse when the down counter value reaches zero. It reloads the down counter in the cycle following a pulse.

Should a user wish to use the TMR\_FIPER1 register to generate a 1 PPS event, the following setup should be used:

- Program TMR\_FIPER1 to a value that will generate a pulse every second,
- Program TMR\_ALARM1 to the correct time for the first PPS event
- Enable the timer

The eTSEC will then wait for TMR\_ALARM1 to expire before enabling the count down of TMR\_FIPER1. The end result will be that TMR\_FIPER1 will pulse every second after the original timer ALARM1 expired.

Note:

In the case where the PPS signals are required to be phased aligned to the prescale output clock, the alarm value should be configured to **1 clock period less** than the wanted value.

In order to keep tracking the prescale output clock, each time before enabling the FIPER, the user must reset the FIPER by writing a new value to the register. The ratio between the prescale register value and the FIPER value should be devisable by the clk period.

$$\text{FIPER\_VALUE} = (\text{prescale\_value} \times \text{tclk\_per} \times N) - \text{tclk\_per}$$

For example:

prescale = 9

clock period = 10

The FIPER can get the following values: 80, 170, 260 .....

The three registers in eTSEC1 are shared for all eTSECs. Figure 15-115 describes the definition for the TMR\_FIPER register.




Figure 15-115. TMR\_FIPER $n$  Register Definition

Table 15-123 describes the fields of the TMR\_FIPER register.

Table 15-123. TMR\_FIPER Register Field Descriptions

Bits	Name	Description
0–31	FIPER	Fixed interval pulse period register. This field must be programmed to an integer multiple of TMR_CTRL[TCLK_PERIOD] value to ensure a period pulse being generated correctly.

#### 15.5.3.10.14 External Trigger Stamp Register (TMR\_ETTS1–2\_H/L)

General purpose external trigger -stamp register (TMR\_ETTS $n$ \_H/L). This register holds time at the programmable edge of the external trigger. The registers in eTSEC1 are shared for all eTSECs. This register is read only in normal operation. Figure 15-116 describes the definition for the TMR\_ETTS $n$ \_H/L register.




Figure 15-116. TMR\_ETTS1-2\_H/L Register Definition

Table 15-124 describes the fields of the TMR\_ETTS $n$ \_H/L register.

Table 15-124. TMR\_ETTS1-2\_H Register Field Descriptions

Bits	Name	Description
0–63	ETTS_H/L	Time stamp field at the programmable edge of the external trigger.

## 15.5.4 Ten-Bit Interface (TBI)

This section describes the ten-bit interface (TBI) and the TBI MII set of registers.

### 15.5.4.1 TBI Transmit Process

The eTSEC's TBI implements the transmit portion of the physical coding sublayer as found in Clause 36 of IEEE 802.3z. In SerDes mode, packets conveyed across the GMII are encapsulated and encoded into 10-bit symbols and output to the SerDes. In GMII mode, the GMII signals are passed through to the attached GMII PHY.

#### 15.5.4.1.1 Packet Encapsulation

If TX\_EN is de-asserted the eTSEC outputs an idle stream. If TX\_EN is asserted, a Start\_of\_Packet symbol is output. This symbol replaces the first byte of the preamble field. All other bytes of the packet pass through an 8B10B encoding module. After the last byte of the FCS field is signaled through the GMII, the MAC de-asserts TX\_EN. The eTSEC then outputs an End\_of\_Packet symbol. Then, depending on the position of the End\_of\_Packet symbols (being in either an odd or even position) the eTSEC outputs one or two Carrier\_Extend symbols. Following the last Carrier\_Extend symbol, the eTSEC resumes sending idle codes. If, during a packet, the eTSEC wishes to mark a byte invalid, TX\_ER is asserted. The eTSEC, upon detection of TX\_ER, substitutes the data symbol for an Error\_Propagation symbol.

#### 15.5.4.1.2 8B10B Encoding

Every eight-bit data octet has two (not necessarily different) ten-bit symbols associated with it. Depending on the running disparity (the cumulative difference of ones and zeroes) the eTSEC module chooses the appropriate symbol.

Special encapsulation symbols are called ordered\_sets. Ordered\_sets are comprised of one to four ten-bit symbols. Ordered\_sets can be found in Clause 36 of the IEEE 802.3z specification.

#### 15.5.4.1.3 Preamble Shortening

Because the idle ordered\_set comprises two symbols and begins on an even symbols boundary, packets can only begin on an even boundary. However, the GMII has no such restriction and may signal TX\_EN on an odd boundary. If this happens, the eTSEC delays the Start\_of\_Packet symbol, effectively ignoring the first byte of preamble; thus, a seven octet preamble becomes six octets on the Ten-Bit Interface.

### 15.5.4.2 TBI Receive Process

The eTSEC's TBI Implements the receive portion of the physical coding sublayer as found in Clause 36 of IEEE 802.3z. The Receive portion includes the Synchronization state machine. In SerDes mode, the eTSEC first attempts to acquire synchronization on the link by examining received symbols. Once synchronization is acquired, received packets are decoded and sent across the Receive GMII interface. In GMII mode, the GMII signals are passed through to the MAC.



### 15.5.4.2.1 Synchronization

The eTSEC examines received symbols looking for the seven bit ‘comma’ string embedded in some special symbols. Both the idle ordered\_set and the Configuration ordered\_set contain a symbol which has the comma. Once a certain number of codes with comma are detected, the eTSEC is considered to have acquired synchronization.

### 15.5.4.2.2 Auto-Negotiation for 1000BASE-X

Once synchronization is acquired, ordered\_sets are decoded. If Configuration ordered\_sets are received, the eTSEC decodes the two octet data field and the sixteen-bit Configuration data is stored and used to Auto-Negotiate with the link partner. In the Receive Configuration Register (RXCR[15:0]) an internal register used to receive all the link partners informations and used to compare to local ability during negotiation. Not visible to user. If, during Auto-Negotiation an invalid symbol is detected, Auto-Negotiation re-starts. After Auto-Negotiation is completed the TBI MII Status Register SR[AN done] in set. In this mode, packets may be received from the link partner.

### 15.5.4.3 TBI MII Set Register Descriptions

This section describes the TBI MII registers. All of the TBI registers are 16 bits wide. The TBI registers are accessed at the offset of the TBI physical address. The eTSEC’s TBI physical address is stored in the TBIPA register. Writing to the TBI registers is performed in a way similar to writing to an external PHY, using the MII management interface. Using TBIPA in place of the PHY address, in the MIIMADD[PHY Address] field, and setting the MIIMADD[Register Address] to the appropriate address offset that corresponds to the register that one wants to read or write (see [Table 15-125](#)), the user can read (set MIIMCOM[read cycle]) or write (writing to MIIMCON[PHY control]) to the TBI block. Refer to the TBI physical address register in [Section 15.5.3.1, “eTSEC General Control and Status Registers,”](#) and the TBI MII register set in [Table 15-125](#). Notice that jitter diagnostics and TBI control are not IEEE 802.3 required registers and are only used for test and control of the eTSEC TBI block. The TBI’s TBI control register (TBI) is for configuring the eTSEC ten-bit interface block. However, because this TBI block has an MII management interface (just like any other PHY), it has an IEEE 802.3 register called the control register (CR).

**Table 15-125. TBI MII Register Set**

Offset Address	Name	Access	Size	Section/page
<b>TEN-BIT INTERFACE (TBI) REGISTERS</b>				<a href="#">15.5.4/15-122</a>
0x00	Control (CR)	R/W <sup>1</sup>	16 bits	<a href="#">15.5.4.3.1/15-124</a>
0x01	Status (SR)	R, LH, LL	16 bits	<a href="#">15.5.4.3.2/15-125</a>
0x02–0x03	Reserved	R	2 bytes	—
0x04	AN advertisement (ANA)	RW, R	16 bits	<a href="#">15.5.4.3.2/15-125</a>
0x05	AN link partner base page ability (ANLPBPA)	R	16 bits	<a href="#">15.5.4.3.4/15-128</a>
0x06	AN expansion (ANEX)	R, LH	16 bits	<a href="#">15.5.4.3.5/15-129</a>
0x07	AN next page transmit (ANNPT)	R/W, R	16 bits	<a href="#">15.5.4.3.6/15-130</a>

Table 15-125. TBI MII Register Set (continued)

Offset Address	Name	Access	Size	Section/page
0x08	AN link partner ability next page (ANLPANP)	R	16 bits	<a href="#">15.5.4.3.7/15-130</a>
0x0F	Extended status (EXST)	R	16 bits	<a href="#">15.5.4.3.8/15-131</a>
0x10	Jitter diagnostics (JD)	R/W	16 bits	<a href="#">15.5.4.3.9/15-132</a>
0x11	TBI control (TBICON)	R/W	16 bits	<a href="#">15.5.4.3.10/15-133</a>

<sup>1</sup> R = Read-only, WO = Write Only, R/W = Read and Write, LH = Latches High, LL = Latches Low, SC = Self-clearing,

### 15.5.4.3.1 Control Register (CR)

Figure 15-117 describes the definition for the CR register.




Figure 15-117. Control Register Definition

Table 15-126 describes the fields of the CR register.

Table 15-126. CR Field Descriptions

Bits	Name	Description
0	PHY Reset	PHY reset. This bit is cleared by default. This bit is self-clearing. 0 Normal operation. 1 The internal state of the TBI is reset. This in turn may change the state of the TBI link partner.
1	—	Reserved
2	Speed[0]	Speed selection. This bit defaults to a cleared state and should always be cleared, which corresponds to 1000 Mbps speed. Setting this field controls the speed at which the TBI operates. The table for Speed[1] provides the appropriate encoding. Its default is bit[2] = '0'; bit[9] = '1'.
3	AN Enable	Auto-negotiation enable. This bit is set by default. 0 The values programmed in bits 2, 7 and 9 determine the operating condition of the link. 1 Auto-negotiation process enabled.
4–5	—	Reserved
6	Reset AN	Reset auto-negotiation. This bit is cleared by default and is self-clearing. 0 Normal operation. 1 The auto-negotiation process restarts. This action is only available if auto-negotiation is enabled.
7	Full Duplex	Duplex mode. This bit is set by default. 0 Reserved. 1 Full-duplex operation.
8	—	Reserved, should be cleared.

Table 15-126. CR Field Descriptions (continued)

Bits	Name	Description															
9	Speed[1]	Speed selection. This bit defaults to a set state and should always be set, which corresponds to 1000 Mbps speed. Setting this field controls the speed at which the TBI operates. The following table provides the appropriate encoding. Its default is bit[2] = '0'; bit[9] = '1'. <table border="1" style="margin-left: auto; margin-right: auto;"> <thead> <tr> <th>Maximum Operating Speed</th> <th>Bit 2</th> <th>Bit 9</th> </tr> </thead> <tbody> <tr> <td>Reserved</td> <td>0</td> <td>0</td> </tr> <tr> <td>Reserved</td> <td>1</td> <td>0</td> </tr> <tr> <td>1000 Mbps</td> <td>0</td> <td>1</td> </tr> <tr> <td>Reserved</td> <td>1</td> <td>1</td> </tr> </tbody> </table>	Maximum Operating Speed	Bit 2	Bit 9	Reserved	0	0	Reserved	1	0	1000 Mbps	0	1	Reserved	1	1
Maximum Operating Speed	Bit 2	Bit 9															
Reserved	0	0															
Reserved	1	0															
1000 Mbps	0	1															
Reserved	1	1															
10–15	—	Reserved															

### 15.5.4.3.2 Status Register (SR)

Figure 15-118 describes the definition for the SR register.

Offset 0x01

Access: Read only




Figure 15-118. Status Register Definition

Table 15-127 describes the fields of the SR register.

Table 15-127. SR Descriptions

Bits	Name	Description
0–6	—	Reserved, should be cleared.
7	Extend Status	This bit indicates that PHY status information is also contained in the Register 15, Extended Status Register. Returns 1 on read. This bit is read-only.
8	—	Reserved, should be cleared.
9	No Pre	MF preamble suppression enable. This bit indicates whether or not the PHY is capable of handling MII management frames without the 32-bit preamble field. Returns 1, indicating support for suppressed preamble MII management frames. This bit is read-only.
10	AN Done	Auto-negotiation complete. This bit is read-only and is cleared by default. 0 Either the auto-negotiation process is underway or the auto-negotiation function is disabled. 1 The auto-negotiation process has completed.
11	Remote Fault	Remote fault. This bit is read-only and is cleared by default. Each read of the status register clears this bit. 0 Normal operation. 1 A remote fault condition was detected. This bit latches high in order for software to detect the condition.

Table 15-127. SR Descriptions (continued)

Bits	Name	Description
12	AN Ability	Auto-negotiation ability. While read as set, this bit indicates that the PHY has the ability to perform auto-negotiation. While read as cleared, this bit indicates the PHY lacks the ability to perform auto-negotiation. Returns 1 on read. This bit is read-only.
13	Link Status	Link status. This bit is read-only and is cleared by default. 0 A valid link is not established. This bit latches low allowing for software polling to detect a failure condition. 1 A valid link is established.
14	—	Reserved, should be cleared.
15	Extend Ability	Extended capability. This bit indicates that the PHY contains the extended set of registers (those beyond control and status). Returns 1 on read. This bit is read-only.

### 15.5.4.3.3 AN Advertisement Register (ANA)

Figure 15-119 describes the definition for the ANA register.




Figure 15-119. AN Advertisement Register Definition

Table 15-128 describes the fields of the ANA register.

Table 15-128. ANA Field Descriptions

Bits	Name	Description															
0	Next Page	Next page configuration. The local device sets this bit to either request next page transmission or advertise next page exchange capability. 0 The local device wishes not to engage in next page exchange. 1 The local device has no next pages but wishes to allow reception of next pages. If the local device has no next pages and the link partner wishes to send next pages, the local device shall send null message codes and have the message page set to 0b000_0000_0001, as defined in annex 28C.															
1	—	Reserved. (Ignore on read)															
2–3	Remote Fault	The local device's remote fault condition is encoded in bits 2 and 3 of the base page. Values are shown in the following table. The default value is 00. Indicate a fault by setting a non-zero remote fault encoding and re-negotiating. <table border="1" style="margin: 10px auto;"> <thead> <tr> <th>RF1 bit[3]</th> <th>RF2 bit[2]</th> <th>Description</th> </tr> </thead> <tbody> <tr> <td>0</td> <td>0</td> <td>No error, link OK</td> </tr> <tr> <td>0</td> <td>1</td> <td>Offline</td> </tr> <tr> <td>1</td> <td>0</td> <td>Link_Failure</td> </tr> <tr> <td>1</td> <td>1</td> <td>Auto-Negotiation_Error</td> </tr> </tbody> </table>	RF1 bit[3]	RF2 bit[2]	Description	0	0	No error, link OK	0	1	Offline	1	0	Link_Failure	1	1	Auto-Negotiation_Error
RF1 bit[3]	RF2 bit[2]	Description															
0	0	No error, link OK															
0	1	Offline															
1	0	Link_Failure															
1	1	Auto-Negotiation_Error															

Table 15-128. ANA Field Descriptions (continued)

Bits	Name	Description															
4–6	—	Reserved, should be cleared.															
7–8	Pause	<p>The local device's PAUSE capability is encoded in bits 7 and 8, and the decodes are shown in the following table. For priority resolution information consult <a href="#">Table 15-129</a>.</p> <table border="1"> <thead> <tr> <th>PAUSE bit[8]</th> <th>ASM_DIR bit[7]</th> <th>Capability</th> </tr> </thead> <tbody> <tr> <td>0</td> <td>0</td> <td>No PAUSE</td> </tr> <tr> <td>0</td> <td>1</td> <td>Asymmetric PAUSE toward link partner</td> </tr> <tr> <td>1</td> <td>0</td> <td>Symmetric PAUSE</td> </tr> <tr> <td>1</td> <td>1</td> <td>Both symmetric PAUSE and Asymmetric PAUSE toward local device</td> </tr> </tbody> </table>	PAUSE bit[8]	ASM_DIR bit[7]	Capability	0	0	No PAUSE	0	1	Asymmetric PAUSE toward link partner	1	0	Symmetric PAUSE	1	1	Both symmetric PAUSE and Asymmetric PAUSE toward local device
PAUSE bit[8]	ASM_DIR bit[7]	Capability															
0	0	No PAUSE															
0	1	Asymmetric PAUSE toward link partner															
1	0	Symmetric PAUSE															
1	1	Both symmetric PAUSE and Asymmetric PAUSE toward local device															
9	Half Duplex	<p>Half-duplex capability.</p> <p>0 Designates local device as not capable of half-duplex operation.</p> <p>1 Designates local device as capable of half-duplex operation.</p>															
10	Full Duplex	<p>Full-duplex capability.</p> <p>0 Designates the local device as not capable of full-duplex operation.</p> <p>1 Designates the local device as capable of full-duplex operation.</p>															
11–15	—	Reserved, should be cleared.															

[Table 15-129](#) describes the resolution of pause priority.

Table 15-129. PAUSE Priority Resolution


Local Device		Link Partner		Local Resolution	Link Partner Resolution
PAUSE	ASM_DIR	PAUSE	ASM_DIR		
0	0	x	x	Disable PAUSE transmit Disable PAUSE receive	Disable PAUSE transmit Disable PAUSE receive
0	1	0	x	Disable PAUSE transmit Disable PAUSE receive	Disable PAUSE transmit Disable PAUSE receive
0	1	1	0	Disable PAUSE transmit Disable PAUSE receive	Disable PAUSE transmit Disable PAUSE receive
0	1	1	1	Enable PAUSE transmit Disable PAUSE receive	Disable PAUSE transmit Enable PAUSE receive
1	0	0	x	Disable PAUSE transmit Disable PAUSE receive	Disable PAUSE transmit Disable PAUSE receive
1	0	1	x	Enable PAUSE transmit Enable PAUSE receive	Enable PAUSE transmit Enable PAUSE receive
1	1	0	0	Disable PAUSE transmit Disable PAUSE receive	Disable PAUSE transmit Disable PAUSE receive

**Table 15-129. PAUSE Priority Resolution (continued)**

Local Device		Link Partner		Local Resolution	Link Partner Resolution
PAUSE	ASM_DIR	PAUSE	ASM_DIR		
1	1	0	1	Disable PAUSE transmit Enable PAUSE receive	Enable PAUSE transmit Disable PAUSE receive
1	1	1	x	Enable PAUSE transmit Enable PAUSE receive	Enable PAUSE transmit Enable PAUSE receive

**15.5.4.3.4 AN Link Partner Base Page Ability Register (ANLPBPA)**

Figure 15-120 describes the definition for the ANLPBPA register.



**Figure 15-120. AN Link Partner Base Page Ability Register Definition**

Table 15-130 describes the fields of the ANLPBPA register.

**Table 15-130. ANLPBPA Field Descriptions**

Bits	Name	Description															
0	Next Page	Next page. This bit is read-only. The link partner sets or clears this bit. 0 Link partner has no subsequent next pages or is not capable of receiving next pages. 1 Link partner either requesting next page transmission or indicating the capability to receive next pages.															
1	—	Reserved. (Ignore on read)															
2–3	Remote Fault	The link partner's remote fault condition is encoded in bits 2 and 3 of the base page. Values are shown in the remote fault encoding field table below. This bit is read-only. <table border="1" style="margin: 10px auto; width: 80%;"> <thead> <tr> <th>RF1 bit[3]</th> <th>RF2 bit[2]</th> <th>Description</th> </tr> </thead> <tbody> <tr> <td>0</td> <td>0</td> <td>No error, link OK</td> </tr> <tr> <td>0</td> <td>1</td> <td>Offline</td> </tr> <tr> <td>1</td> <td>0</td> <td>Link_Failure</td> </tr> <tr> <td>1</td> <td>1</td> <td>Auto-Negotiation_Error</td> </tr> </tbody> </table>	RF1 bit[3]	RF2 bit[2]	Description	0	0	No error, link OK	0	1	Offline	1	0	Link_Failure	1	1	Auto-Negotiation_Error
RF1 bit[3]	RF2 bit[2]	Description															
0	0	No error, link OK															
0	1	Offline															
1	0	Link_Failure															
1	1	Auto-Negotiation_Error															
4–6	—	Reserved, should be cleared.															

Table 15-130. ANLPBPA Field Descriptions (continued)

Bits	Name	Description															
7–8	Pause	Encoding of the link partner's PAUSE capability is shown in the PAUSE encoding table below. For priority resolution information consult. This bit is read-only <table border="1" data-bbox="365 373 1360 653"> <thead> <tr> <th>PAUSE bit[8]</th> <th>ASM_DIR bit[7]</th> <th>Capability</th> </tr> </thead> <tbody> <tr> <td>0</td> <td>0</td> <td>No PAUSE</td> </tr> <tr> <td>0</td> <td>1</td> <td>Asymmetric PAUSE toward link partner</td> </tr> <tr> <td>1</td> <td>0</td> <td>Symmetric PAUSE</td> </tr> <tr> <td>1</td> <td>1</td> <td>Both symmetric PAUSE and Asymmetric PAUSE toward local device</td> </tr> </tbody> </table>	PAUSE bit[8]	ASM_DIR bit[7]	Capability	0	0	No PAUSE	0	1	Asymmetric PAUSE toward link partner	1	0	Symmetric PAUSE	1	1	Both symmetric PAUSE and Asymmetric PAUSE toward local device
PAUSE bit[8]	ASM_DIR bit[7]	Capability															
0	0	No PAUSE															
0	1	Asymmetric PAUSE toward link partner															
1	0	Symmetric PAUSE															
1	1	Both symmetric PAUSE and Asymmetric PAUSE toward local device															
9	Half Duplex	Half-duplex capability. This bit is read-only. 0 Link partner is not capable of half-duplex mode. 1 Link partner is capable of half-duplex mode.															
10	Full Duplex	Full-duplex capability. This bit is read-only. 0 Link partner is not capable of full-duplex mode. 1 Link partner is capable of full-duplex mode.															
11–15	—	Reserved, should be cleared.															

### 15.5.4.3.5 AN Expansion Register (ANEX)

Figure 15-121 describes the definition for the ANEX register.




Figure 15-121. AN Expansion Register Definition

Table 15-131 describes the fields of the ANEX register.

Table 15-131. ANEX Field Descriptions

Bits	Name	Description
0–12	—	Reserved, should be cleared.
13	NP Able	Next page able. This bit is read-only and returns 1 on read. While read as set, indicates local device supports next page function.
14	Page Rx'd	Page received. This bit is read-only. The bit clears on a read to the register. 0 Normal operation. 1 A new page was received and stored in the applicable AN link partner ability or AN next page register. This bit latches high in order for software to detect while polling.
15	—	Reserved, should be cleared.

### 15.5.4.3.6 AN Next Page Transmit Register (ANNPT)

Figure 15-122 describes the definition for the ANNPT register.




Figure 15-122. AN Next Page Transmit Register Definition

Table 15-132 describes the fields of the ANNPT register.

Table 15-132. ANNPT Field Descriptions

Bits	Name	Description
0	Next Page	Next page indication. [Reference MII bit 7.15 in IEEE 802.3, 2000 Edition Clause 28.2.4] 0 Last page. 1 Additional next pages to follow.
1	—	Reserved. (Ignore on read)
2	Msg Page	Message page. [Reference MII bit 7.13] 0 Unformatted page. 1 Message page.
3	Ack2	Acknowledge 2. Used by the next page function to indicate that the device has the ability to comply with the message. [Reference MII bit 7.12] 0 The local device cannot comply with message. 1 The local device complies with message.
4	Toggle	Toggle. Used to ensure synchronization with the link partner during next page exchange. This bit always takes the opposite value of the toggle bit of the previously-exchanged link code word. The initial value in the first next page transmitted is the inverse of bit 11 in the base link code word. [Reference MII bit 7.11] This bit is read-only. 0 Toggle bit of the previously-exchanged link code word was 1. 1 Toggle bit of the previously-exchanged link code word was 0.
5–15	Message/Un-formatted Code Field	Message pages are formatted pages that carry a pre-defined message code, which is enumerated in IEEE 802.3u/Annex 28C. Unformatted code fields take on an arbitrary value. [Reference MII field 7.10:0]

### 15.5.4.3.7 AN Link Partner Ability Next Page Register (ANLPANP)

Figure 15-123 describes the definition for the ANLPANP register.




Figure 15-123. AN Link Partner Ability Next Page Register Definition





Table 15-134 describes the fields of the EXST register.


**Table 15-134. EXST Field Descriptions**

Bits	Name	Description
0	1000X Full	1000X full-duplex capability. Returns 1 on read. This bit is read-only. 0 PHY cannot operation in 1000BASE-X full-duplex mode. 1 PHY can operate in 1000BASE-X full-duplex mode.
1	1000X Half	1000X half-duplex capability. Returns 0 on read. This bit is read-only. 0 PHY cannot operation in 1000BASE-X half-duplex mode. 1 PHY can operate in 1000BASE-X half-duplex mode.
2	1000T Full	1000T full-duplex capability. Returns 1 on read. This bit is read-only. 0 PHY cannot operation in 1000BASE-T full-duplex mode. 1 PHY can operate in 1000BASE-T full-duplex mode.
3	1000T Half	1000T half-duplex capability. Returns 0 on read. This bit is read-only. 0 PHY cannot operation in 1000BASE-T half-duplex mode. 1 PHY can operate in 1000BASE-T half-duplex mode.
4–15	—	Reserved

### 15.5.4.3.9 Jitter Diagnostics Register (JD)

Annex 36A in IEEE 802.3z describes several jitter test patterns. These can be configured to be sent by writing the jitter diagnostics register. See the register description for more information. In may be wise to auto-negotiate and advertise a remote fault signaling of offline prior to beginning the test patterns.

Figure 15-125 describes the definition for the JD register.



**Figure 15-125. Jitter Diagnostics Register Definition**


Table 15-135 describes the fields of the JD register.

**Table 15-135. JD Field Descriptions**

Bits	Name	Description																																				
0	Jitter Enable	Jitter enable. This bit is cleared by default. 0 Normal transmit operation. 1 Enable the TBI to transmit the jitter test patterns defined in IEEE 802.3z 36A.																																				
1–3	Jitter Select	Selects the jitter pattern to be transmitted in diagnostics mode. Encoding of this field is shown in the following table. Default is 00. <table border="1" style="width: 100%; margin-top: 10px;"> <thead> <tr> <th style="text-align: center;">Jitter Pattern Select</th> <th style="text-align: center;">bit[1]</th> <th style="text-align: center;">bit[2]</th> <th style="text-align: center;">bit[3]</th> </tr> </thead> <tbody> <tr> <td>User defined uses custom jitter pattern, bits 6–15</td> <td style="text-align: center;">0</td> <td style="text-align: center;">0</td> <td style="text-align: center;">0</td> </tr> <tr> <td>High frequency (+/- D21.5) 101010101010101010101010101010101010...</td> <td style="text-align: center;">0</td> <td style="text-align: center;">0</td> <td style="text-align: center;">1</td> </tr> <tr> <td>Mixed frequency (+/- K28.5) 1111101011000001010011111010110000010100...</td> <td style="text-align: center;">0</td> <td style="text-align: center;">1</td> <td style="text-align: center;">0</td> </tr> <tr> <td>Low frequency 1111100000111110000011111000001111100000...</td> <td style="text-align: center;">0</td> <td style="text-align: center;">1</td> <td style="text-align: center;">1</td> </tr> <tr> <td>Complex pattern (10'h17c,10'h0c9,10'h0e5,10'h2a3, 10'h17c,...)</td> <td style="text-align: center;">1</td> <td style="text-align: center;">0</td> <td style="text-align: center;">0</td> </tr> <tr> <td>Square Wave (- K28.7) 0011111000001111100000111110000011111000...</td> <td style="text-align: center;">1</td> <td style="text-align: center;">0</td> <td style="text-align: center;">1</td> </tr> <tr> <td>Reserved</td> <td style="text-align: center;">1</td> <td style="text-align: center;">1</td> <td style="text-align: center;">0</td> </tr> <tr> <td>Reserved</td> <td style="text-align: center;">1</td> <td style="text-align: center;">1</td> <td style="text-align: center;">1</td> </tr> </tbody> </table>	Jitter Pattern Select	bit[1]	bit[2]	bit[3]	User defined uses custom jitter pattern, bits 6–15	0	0	0	High frequency (+/- D21.5) 101010101010101010101010101010101010...	0	0	1	Mixed frequency (+/- K28.5) 1111101011000001010011111010110000010100...	0	1	0	Low frequency 1111100000111110000011111000001111100000...	0	1	1	Complex pattern (10'h17c,10'h0c9,10'h0e5,10'h2a3, 10'h17c,...)	1	0	0	Square Wave (- K28.7) 0011111000001111100000111110000011111000...	1	0	1	Reserved	1	1	0	Reserved	1	1	1
Jitter Pattern Select	bit[1]	bit[2]	bit[3]																																			
User defined uses custom jitter pattern, bits 6–15	0	0	0																																			
High frequency (+/- D21.5) 101010101010101010101010101010101010...	0	0	1																																			
Mixed frequency (+/- K28.5) 1111101011000001010011111010110000010100...	0	1	0																																			
Low frequency 1111100000111110000011111000001111100000...	0	1	1																																			
Complex pattern (10'h17c,10'h0c9,10'h0e5,10'h2a3, 10'h17c,...)	1	0	0																																			
Square Wave (- K28.7) 0011111000001111100000111110000011111000...	1	0	1																																			
Reserved	1	1	0																																			
Reserved	1	1	1																																			
4–5	—	Reserved																																				
6–15	Custom Jitter Pattern	Used in conjunction with jitter (pattern) select and jitter (diagnostic) enable; set this field to the desired custom pattern which is continuously transmitted. Its default is 0x000.																																				

#### 15.5.4.3.10 TBI Control Register (TBICON)

Figure 15-126 describes the definition for the TBICON register.



**Figure 15-126. TBI Control Register Definition**

Table 15-136 describes the fields of the TBICON register.

**Table 15-136. TBICON Field Descriptions**

Bits	Name	Description
0	Soft_Reset	Soft reset. This bit is cleared by default. 0 Normal operation. 1 Resets the functional modules in the TBI.
1	—	Reserved. (Ignore on read)
2	Disable Rx Dis	Disable receive disparity. This bit is cleared by default. 0 Normal operation. 1 Disables the running disparity calculation and checking in the receive direction.
3	Disable Tx Dis	Disable transmit disparity. This bit is cleared by default. 0 Normal operation. 1 Disables the running disparity calculation and checking in the transmit direction.
4–6	—	Reserved
7	AN Sense	Auto-negotiation sense enable. This bit is cleared by default. 0 IEEE 802.3z Clause 37 behavior is desired, which results in the link not completing. 1 Allow the auto-negotiation function to sense either a Gigabit MAC in auto-negotiation bypass mode or an older Gigabit MAC without auto-negotiation capability. If sensed, auto-negotiation complete becomes true; however, the page received is low, indicating no page was exchanged. Management can then act accordingly.
8–9	—	Reserved
10	Clock Select	Clock select. This bit is cleared by default. 0 Allow the TBI to accept dual split-phase 62.5 MHz receive clocks. 1 Configure the TBI to accept a 125 MHz receive clock from the SerDes/PHY. The 125 MHz clock must be physically connected to 'PMA receive clock 0' if using a parallel (non-SGMII) Ethernet protocol.
11	MI Mode	This bit describes the configuration mode of the TBI. The user reads a 1 while the TBI is configured in GMII/MII mode (connected to a GMII/MII PHY) and a 0 while configured in TBI mode (connected to a 1000BASE-X SerDes). Its value is the inverse of ECNTRL[TBIM]. 0 TBI mode. 1 GMII mode.
12–15	—	Reserved

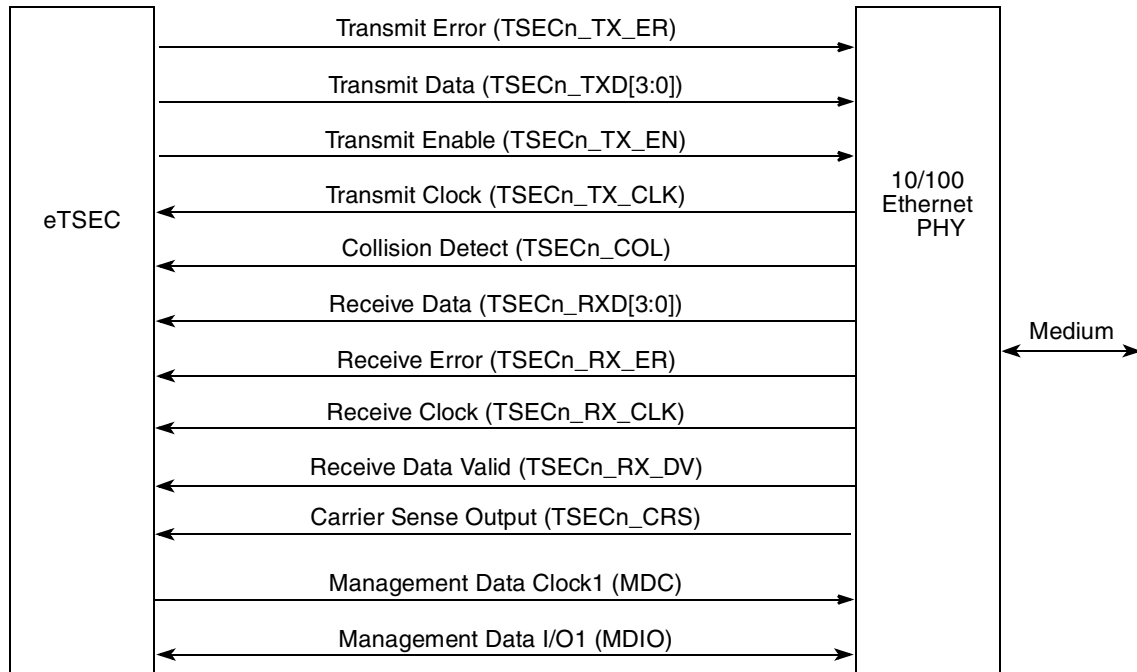
## 15.6 Functional Description

### 15.6.1 Connecting to Physical Interfaces on Ethernet

This section describes how to connect the eTSEC to various interfaces: MII, RMII, RGMII, and RTBI. To avoid confusion, all of the buses follow the bus conventions used in the IEEE 802.3 specification because the PHYs follow the same conventions. (For instance, in the bus TSEC<sub>n</sub>\_TXD[3:0], bit 3 is the msb and bit 0 is the lsb). If a mode does not use all input signals available to a particular eTSEC, those inputs that are not used must be pulled low on the board.

### 15.6.1.1 Media-Independent Interface (MII)

This section describes the media-independent interface (MII) intended to be used between the PHYs and the eTSEC. Figure 15-127 depicts the basic components of the MII including the signals required to establish eTSEC module connection with a PHY.




<sup>1</sup> The management signals (MDC and MDIO) are common to all of the Ethernet controllers' connections in the system, assuming that each PHY has a different management address.

**Figure 15-127. eTSEC-MII Connection**

An MII interface has 18 signals (including the MDC and MDIO signals), as defined by the IEEE 802.3u standard, for connecting to an Ethernet PHY.

### 15.6.1.2 Reduced Media-Independent Interface (RMII)

This section describes the reduced media-independent interface (RMII) intended to be used between the PHYs and the GMII MAC. The RMII is a reduced-pin alternative to the IEEE 802.3u MII. The RMII reduces the number of signals required to interconnect the MAC and the PHY from a maximum of 18 signals (MII) to 10 signals. To accomplish this objective, the data paths are halved in width and clocked at twice the MII clock frequency, while clocks, carrier sense and error signals have been partly combined. For 100 Mbps operation, the reference clock operates at 50 MHz, whereas for 10 Mbps operation, the clock remains at 50MHz, but only every 10th cycle is used. Figure 15-128 depicts the basic components of the reduced media-independent interface and the signals required to establish an eTSEC's connection with a PHY. The RMII is implemented as defined by the RMII Specification of the RMII Consortium, as of March 20, 1998.

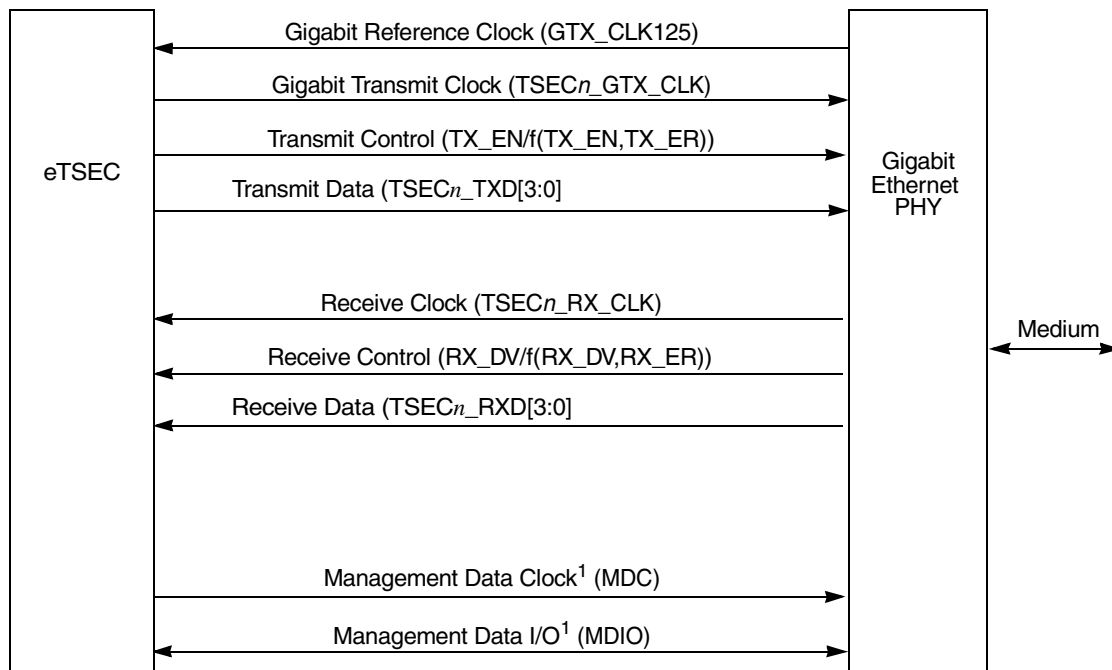


<sup>2</sup> The management signals (MDC and MDIO) are common to all of the Ethernet controllers module connections in the system, assuming that each PHY has a different management address.

**Figure 15-128. eTSEC-RMII Connection**

### 15.6.1.3 Reduced Gigabit Media-Independent Interface (RGMII)

This section describes the reduced gigabit media-independent interface (RGMII) intended to be used between the PHYs and the GMII MAC. The RGMII is an alternative to the IEEE802.3u MII, the IEEE802.3z GMII. The RGMII reduces the number of signals required to interconnect the MAC and the PHY from a maximum of 28 signals (GMII) to 15 signals (GTX\_CLK125 included) in a cost effective and technology independent manner. To accomplish this objective, the data paths and all associated control signals are multiplexed using both edges of the clock. For gigabit operation, the clocks operate at 125MHz, and for 10/100 operation, the clocks operate at 2.5 MHz or 25 MHz, respectively. Note that the GTX\_CLK125 input must be provided at 125 MHz for an RGMII interface, regardless of operation speed (1 Gbps, 100 Mbps, or 10 Mbps). [Figure 15-129](#) depicts the basic components of the gigabit reduced media-independent interface and the signals required to establish the gigabit Ethernet controllers' module connection with a PHY. The RGMII is implemented as defined by the RGMII specification Version 1.2a 9/22/00.

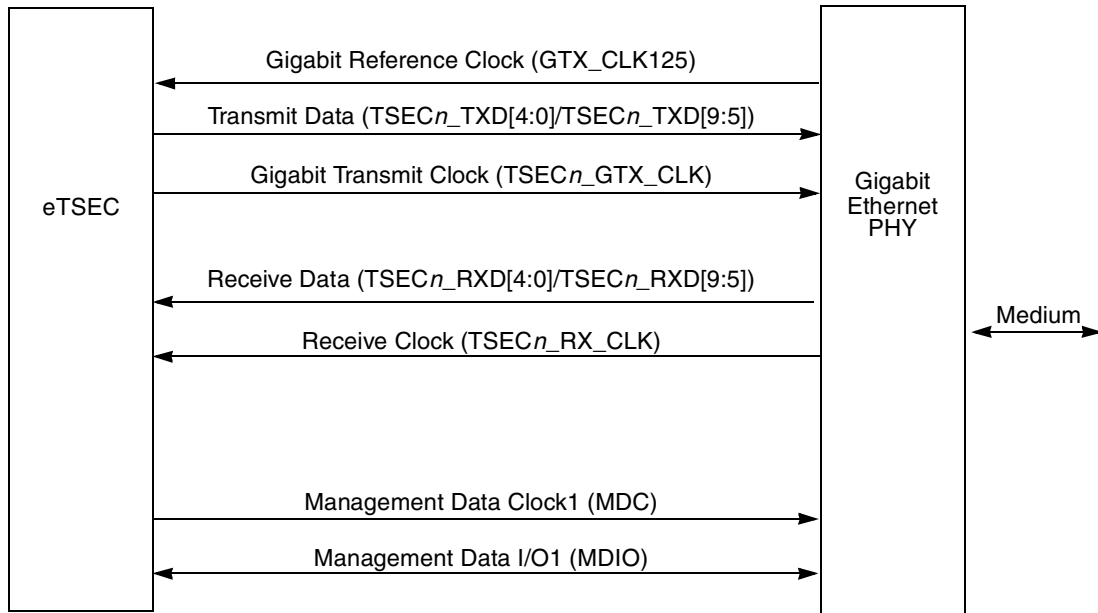


<sup>1</sup> The management signals (MDC and MDIO) are common to all of the gigabit Ethernet controllers module connections in the system, assuming that each PHY has a different management address.

**Figure 15-129. eTSEC-RGMII Connection**

#### 15.6.1.4 Reduced Ten-Bit Interface (RTBI)

This section describes the reduced ten-bit interface (RTBI) intended to be used between the PHYs and the eTSEC to implement a reduced-pin count version of a SerDes interface for optical-fiber devices in 1000BASE-SX/LX applications. [Figure 15-130](#) depicts the basic components of the RTBI including the signals required to establish eTSEC module connection with a PHY. Note that in RTBI the eTSEC immediately begins auto-negotiation with the SerDes.



<sup>1</sup> The management signals (MDC and MDIO) are common to all of the Ethernet controllers' connections in the system, assuming that each PHY has a different management address.

**Figure 15-130. eTSEC-RTBI Connection**

A RTBI interface has 15 signals (GE\_GTX\_CLK125 included), as defined by the RGMII specification Version 1.2a 9/22/00, and is intended to be an alternative to the IEEE 802.3u MII, the IEEE 802.3z GMII and the TBI standard for connecting to an Ethernet PHY.



### 15.6.1.5 Ethernet Physical Interfaces Signal Summary

Table 15-137 describes the signal multiplexing for the following interfaces: GMII, TBI, and RMII.

**Table 15-137. GMII, MII, and RMII Signals Multiplexing**

eTSEC Signals			GMII Interface			MII Interface			RMII Interface		
Frequency [MHz] 125			Frequency [MHz] 125			Frequency [MHz] 25			Frequency [MHz] 50		
Voltage[V] 3.3/2.5			Voltage[V] 3.3			Voltage[V] 3.3			Voltage[V] 3.3		
Signals (TSEC <sub>n</sub> )	I/O	No. of Signals	Signals (TSEC <sub>n</sub> )	I/O	No. of Signals	Signals (TSEC <sub>n</sub> )	I/O	No. of Signals	Signals (TSEC <sub>n</sub> )	I/O	No. of Signals
GTX_CLK	O	1	GTX_CLK	O	1						
TX_CLK	I	1	TX_CLK	I	1	TX_CLK	I	1	REF_CLK	I	1
TxD[0]	O	1	TxD[0]	O	1	TxD[0]	O	1	TxD[0]	O	1
TxD[1]	O	1	TxD[1]	O	1	TxD[1]	O	1	TxD[1]	O	1
TxD[2]	O	1	TxD[2]	O	1	TxD[2]	O	1			
TxD[3]	O	1	TxD[3]	O	1	TxD[3]	O	1			
TX_EN	O	1	TX_EN	O	1	TX_EN	O	1	TX_EN	O	1
TX_ER	O	1	TX_ER	O	1	TX_ER	O	1			
RX_CLK	I	1	RX_CLK	I	1	RX_CLK	I	1			
RxD[0]	I	1	RxD[0]	I	1	RxD[0]	I	1	RxD[0]	I	1
RxD[1]	I	1	RxD[1]	I	1	RxD[1]	I	1	RxD[1]	I	1
RxD[2]	I	1	RxD[2]	I	1	RxD[2]	I	1			
RxD[3]	I	1	RxD[3]	I	1	RxD[3]	I	1			
RX_DV	I	1	RX_DV	I	1	RX_DV	I	1	CRS_DV	I	1
RX_ER	I	1	RX_ER	I	1	RX_ER	I	1	RX_ER	I	1
COL	I	1				COL	I	1			
CRS	I	1				CRS	I	1			
<b>Sum</b>		25	<b>Sum</b>		23	<b>Sum</b>		16	<b>Sum</b>		8

Table 15-138 describes the signal multiplexing for MII and RMII interfaces.

**Table 15-138. MII and RMII Signals Multiplexing**

eTSEC Signals			MII Interface			RMII Interface		
Frequency [MHz] 125			Frequency [MHz] 25			Frequency [MHz] 50		
Voltage[V] 3.3/2.5			Voltage[V] 3.3			Voltage[V] 3.3		
Signals (TSEC <sub>n</sub> )	I/O	No. of Signals	Signals (TSEC <sub>n</sub> )	I/O	No. of Signals	Signals (TSEC <sub>n</sub> )	I/O	No. of Signals
GTX_CLK	O	1						
TX_CLK	I	1	TX_CLK	I	1	REF_CLK	I	1

Table 15-138. MII and RMII Signals Multiplexing (continued)

eTSEC Signals			MII Interface			RMII Interface		
Frequency [MHz] 125			Frequency [MHz] 25			Frequency [MHz] 50		
Voltage[V] 3.3/2.5			Voltage[V] 3.3			Voltage[V] 3.3		
TxD[0]	O	1	TxD[0]	O	1	TxD[0]	O	1
TxD[1]	O	1	TxD[1]	O	1	TxD[1]	O	1
TxD[2]	O	1	TxD[2]	O	1			
TxD[3]	O	1	TxD[3]	O	1			
TX_EN	O	1	TX_EN	O	1	TX_EN	O	1
TX_ER	O	1	TX_ER	O	1			
RX_CLK	I	1	RX_CLK	I	1			
RxD[0]	I	1	RxD[0]	I	1	RxD[0]	I	1
RxD[1]	I	1	RxD[1]	I	1	RxD[1]	I	1
RxD[2]	I	1	RxD[2]	I	1			
RxD[3]	I	1	RxD[3]	I	1			
RX_DV	I	1	RX_DV	I	1	CRS_DV	I	1
RX_ER	I	1	RX_ER	I	1	RX_ER	I	1
COL	I	1	COL	I	1			
CRS	I	1	CRS	I	1			
<b>Sum</b>		25	<b>Sum</b>		16	<b>Sum</b>		8

Table 15-139 describes the signal multiplexing for RGMII, TBI, and RTBI interfaces.

Table 15-139. RGMII, TBI, and RTBI Signals Multiplexing

eTSEC Signals			RGMII Interface			TBI Interface			RTBI Interface		
Frequency [MHz] 125			Frequency [MHz] 125			Frequency [MHz] 62.5			Frequency [MHz] 62.5		
Voltage[V] 3.3/2.5			Voltage[V] 2.5			Voltage[V] 3.3			Voltage[V] 2.5		
Signals (TSEC <sub>n</sub> )	I/O	No. of Signals	Signals (TSEC <sub>n</sub> )	I/O	No. of Signals	Signals (TSEC <sub>n</sub> )	I/O	No. of Signals	Signals (TSEC <sub>n</sub> )	I/O	No. of Signals
GTX_CLK	O	1	GTX_CLK	O	1	GTX_CLK	O	1	GTX_CLK	O	1
TX_CLK	I	1				RX_CLK1	I	1			
TxD[0]	O	1	TxD[0]	O	1	TCG[0]	O	1	TCG[0]/TCG[5]	O	1
TxD[1]	O	1	TxD[1]	O	1	TCG[1]	O	1	TCG[1]/TCG[6]	O	1
TxD[2]	O	1	TxD[2]	O	1	TCG[2]	O	1	TCG[2]/TCG[7]	O	1
TxD[3]	O	1	TxD[3]/	O	1	TCG[3]	O	1	TCG[3]/TCG[8]	O	1
TX_EN	O	1	TX_CTL (TX_EN/ TX_ERR)	O	1	TCG[8]	O	1	TCG[4]/TCG[9]	O	1
TX_ER	O	1				TCG[9]	O	1			

Table 15-139. RGMII, TBI, and RTBI Signals Multiplexing (continued)

eTSEC Signals			RGMII Interface			TBI Interface			RTBI Interface		
Frequency [MHz] 125			Frequency [MHz] 125			Frequency [MHz] 62.5			Frequency [MHz] 62.5		
Voltage[V] 3.3/2.5			Voltage[V] 2.5			Voltage[V] 3.3			Voltage[V] 2.5		
Signals (TSEC <sub>n</sub> )	I/O	No. of Signals	Signals (TSEC <sub>n</sub> )	I/O	No. of Signals	Signals (TSEC <sub>n</sub> )	I/O	No. of Signals	Signals (TSEC <sub>n</sub> )	I/O	No. of Signals
RX_CLK	I	1	RX_CLK	I	1	RX_CLK0	I	1	RX_CLK	I	1
RxD[0]	I	1	RxD[0]/RxD[4]	I	1	RCG[0]	I	1	RCG[0]/RCG[5]	I	1
RxD[1]	I	1	RxD[1]/RxD[5]	I	1	RCG[1]	I	1	RCG[1]/RCG[6]	I	1
RxD[2]	I	1	RxD[2]/RxD[6]	I	1	RCG[2]	I	1	RCG[2]/RCG[7]	I	1
RxD[3]	I	1	RxD[3]/RxD[7]	I	1	RCG[3]	I	1	RCG[3]/RCG[8]	I	1
RX_DV	I	1	RX_CTL (RX_DV/ RX_ERR)	I	1	RCG[8]	I	1	RCG[4]/RCG[9]	I	1
RX_ER	I	1				RCG[9]	I	1			
COL	I	1									
CRS	I	1				SDET	I	1		I	
<b>Sum</b>		25	<b>Sum</b>		12	<b>Sum</b>		24	<b>Sum</b>		12

Table 15-140 describes the signal multiplexing for RGMII and RTBI interfaces.

Table 15-140. RGMII and RTBI Signals Multiplexing

eTSEC Signals			RGMII Interface			RTBI Interface		
Frequency [MHz] 125			Frequency [MHz] 125			Frequency [MHz] 62.5		
Voltage[V] 3.3/2.5			Voltage[V] 2.5			Voltage[V] 2.5		
Signals (TSEC <sub>n</sub> )	I/O	No. of Signals	Signals (TSEC <sub>n</sub> )	I/O	No. of Signals	Signals (TSEC <sub>n</sub> )	I/O	No. of Signals
GTX_CLK	O	1	GTX_CLK	O	1	GTX_CLK	O	1
TX_CLK	I	1						
TxD[0]	O	1	TxD[0]	O	1	TCG[0]/TCG[5]	O	1
TxD[1]	O	1	TxD[1]	O	1	TCG[1]/TCG[6]	O	1
TxD[2]	O	1	TxD[2]	O	1	TCG[2]/TCG[7]	O	1
TxD[3]	O	1	TxD[3]	O	1	TCG[3]/TCG[8]	O	1
TX_EN	O	1	TX_CTL (TX_EN/ TX_ERR)	O	1	TCG[4]/TCG[9]	O	1
TX_ER	O	1						
RX_CLK	I	1	RX_CLK	I	1	RX_CLK	I	1
RxD[0]	I	1	RxD[0]/RxD[4]	I	1	RCG[0]/RCG[5]	I	1
RxD[1]	I	1	RxD[1]/RxD[5]	I	1	RCG[1]/RCG[6]	I	1

Table 15-140. RGMII and RTBI Signals Multiplexing (continued)

eTSEC Signals			RGMII Interface			RTBI Interface		
Frequency [MHz] 125			Frequency [MHz] 125			Frequency [MHz] 62.5		
Voltage[V] 3.3/2.5			Voltage[V] 2.5			Voltage[V] 2.5		
Signals (TSECn_)	I/O	No. of Signals	Signals (TSECn_)	I/O	No. of Signals	Signals (TSECn_)	I/O	No. of Signals
RxD[2]	I	1	RxD[2]/RxD[6]	I	1	RCG[2]/RCG[7]	I	1
RxD[3]	I	1	RxD[3]/RxD[7]	I	1	RCG[3]/RCG[8]	I	1
RX_DV	I	1	RX_CTL (RX_DV/ RX_ERR)	I	1	RCG[4]/RCG[9]	I	1
RX_ER	I	1						
COL	I	1						
CRS	I	1					I	
<b>Sum</b>		25	<b>Sum</b>		12	<b>Sum</b>		12

Table 15-141 describes the signal multiplexing for the RGMII interface.

Table 15-141. RGMII Signals Multiplexing

eTSEC Signals			RGMII Interface		
Frequency [MHz] 125			Frequency [MHz] 125		
Voltage[V] 3.3/2.5			Voltage[V] 2.5		
Signals (TSECn_)	I/O	No. of Signals	Signals (TSECn_)	I/O	No. of Signals
GTX_CLK	O	1	GTX_CLK	O	1
TX_CLK	I	1			
TxD[0]	O	1	TxD[0]	O	1
TxD[1]	O	1	TxD[1]	O	1
TxD[2]	O	1	TxD[2]	O	1
TxD[3]	O	1	TxD[3]	O	1
TX_EN	O	1	TX_CTL (TX_EN/TX_ERR)	O	1
TX_ER	O	1			
RX_CLK	I	1	RX_CLK	I	1
RxD[0]	I	1	RxD[0]/RxD[4]	I	1
RxD[1]	I	1	RxD[1]/RxD[5]	I	1
RxD[2]	I	1	RxD[2]/RxD[6]	I	1
RxD[3]	I	1	RxD[3]/RxD[7]	I	1

**Table 15-141. RGMII Signals Multiplexing (continued)**

eTSEC Signals			RGMII Interface		
Frequency [MHz] 125			Frequency [MHz] 125		
Voltage[V] 3.3/2.5			Voltage[V] 2.5		
Signals (TSECn_)	I/O	No. of Signals	Signals (TSECn_)	I/O	No. of Signals
RX_DV	I	1	RX_CTL (RX_DV/RX_ERR)	I	1
RX_ER	I	1			
COL	I	1			
CRS	I	1			
<b>Sum</b>		25	<b>Sum</b>		12

Table 15-142 describes the signals shared by all interfaces.

**Table 15-142. Shared Signals**

Signals	I/O	No. of Signals	Function
MDIO	I/O	1	Management interface I/O
MDC	O	1	Management interface clock
GTX_CLK125	I	1	Reference clock
<b>Sum</b>			—

### 15.6.1.6 SGMII Interface

SGMII communication using the eTSEC is accomplished through the SerDes interface. See [Table 15-1 on page 15-6](#) for specific signal assignments.

## 15.6.2 Gigabit Ethernet Controller Channel Operation

This section describes the operation of the eTSEC. First, the software initialization sequence is described. Next, the software (Ethernet driver) interface for transmitting and receiving frames is reviewed. Frame filtering and receive filing algorithm features are also discussed. The section concludes with interrupt handling, inter-packet gap time, and loop back descriptions.

### 15.6.2.1 Initialization Sequence

This sections describes which registers are reset due to a hard or software reset and what registers the user must initialize prior to enabling the eTSEC.

### 15.6.2.1.1 Hardware Controlled Initialization

A hard reset occurs when the system powers up. All eTSEC's registers and control logic are reset to their default states after a hard reset has occurred. In this state, each eTSEC behaves like a PowerQUICC II Pro device, except for the absence of out-of-sequence TxBD features. That is, initially TCP/IP off-load is disabled and only single RxBD and TxBD rings are accessible.

### 15.6.2.1.2 User Initialization

After the system has undergone a hard reset, software must initialize certain basic eTSEC registers. Other registers can also be initialized during this time, but they are optional and must be determined based on the requirements of the system. See [Table 15-3](#) for the register list. [Table 15-143](#) describes the minimum steps for register initialization.

**Table 15-143. Steps for Minimum Register Initialization**

Description
1. Set and clear MACCFG1 [Soft_Reset]
2. Initialize MACCFG2
3. Initialize MAC station address
4. Set up the PHY using the MII Mgmt Interface
5. Configure the GMII
6. Clear IEVENT
7. Initialize IMASK
8. Initialize RCTRL
9. Initialize DMACTRL

After the initialization of registers is performed, the user must execute the following steps in the order described below to bring the eTSEC into a functional state (out of reset):

1. Write to the MACCFG1 register and set the appropriate bits. These need to include RX\_EN and TX\_EN. To enable flow control, Rx\_Flow and Tx\_Flow should also be set.
2. For the transmission of Ethernet frames, TxBDs must first be built in memory, linked together as a ring, and pointed to by the TBASE $n$  registers. A minimum of two buffer descriptors per ring is required, unless the ring is disabled. Setting the ring to a size of one causes the same frame to be transmitted twice. If TCP/IP off-load is to be enabled, the TxBD[TOE] bit must be set for each frame.
3. Likewise, for the reception of Ethernet frames, the receive queue (or queues) must be ready, with its RxBD pointed to by the RBASE $n$  registers. If TCP/IP off-load is to be enabled, RCTRL[PRSDEP] must be set to the required off-load level. Both transmit and receive can be gracefully stopped after transmission and reception begins.

4. Clearing DMACTRL[GTS] triggers the transmission of frame data if the transmitter had been previously stopped. The DMACTRL[GRS] must be cleared if the receiver had been previously stopped. Refer to the DMACTRL register section, and [Section 15.6.7.1, “Data Buffer Descriptors,”](#) for more information.

### 15.6.2.2 Soft Reset and Reconfiguring Procedure

Before issuing a soft-reset to and/or reconfiguring the MAC with new parameters, the user must properly shutdown the DMA and make sure it is in an idle state for the entire duration. User must gracefully stop the DMA by setting both GRS and GTS bits in the DMACTRL register, then wait for both GRSC and GTSC bits to be set in the IEVENT register before resetting the MAC or changing parameters. Both GRS and GTS bits must be cleared before re-enabling the MAC to resume the DMA.

During the MAC configuration, if a new set of Tx buffer descriptors are used, the user must load the pointers into the TBASE registers. Likewise if a new set of Rx buffer descriptors are used, the RBASE registers must be written with new pointers.

Following is a procedure to gracefully reset and reconfigure the MAC:

1. Set GRS/GTS bits in DMACTRL register
2. Poll GRSC/GTSC bits in IEVENT register until both are set
3. Set SOFT\_RESET bit in MACCFG1 register (Note that SOFT\_RESET must remain set for at least 3 TX clocks before proceeding.)
4. Clear SOFT\_RESET bit in MACCFG1 register
5. Load TBASE0–TBASE7 with new Tx BD pointers
6. Load RBASE0–RBASE7 with new Rx BD pointers
7. Setup other MAC registers (MACCFG2, MAXFRM, and so on)
8. Setup group address hash table (GADDR0–GADDR15) if address filtering is required
9. Setup receive frame filter table (through RQFAR, RQFCR, and RQFPR) if filtering to multiple RxBD rings is required
10. Setup WWR, WOP, TOD bits in DMACTRL register
11. Enable transmit queues in TQUEUE, and ensure that the transmit scheduling mode is correctly set in TCTRL.
12. Enable receive queues in RQUEUE, and optionally set TOE functionality in RCTRL.
13. Clear THLT and TXF bits in TSTAT register by writing 1 to them
14. Clear QHLT and RXF bits in RSTAT register by writing 1 to them.
15. Clear GRS/GTS bits in DMACTRL (do not change other bits)
16. Enable Tx\_EN/Rx\_EN in MACCFG1 register

### 15.6.2.3 Gigabit Ethernet Frame Transmission

The Ethernet transmitter requires little core intervention. After the software driver initializes the system, the eTSEC begins to poll the first transmit buffer descriptor (TxBD) in TxBD ring 0 every 512 transmit clocks. If TxBD[R] is set, and the TxBD ring is scheduled for transmission, the eTSEC begins copying the

associated transmit buffer from memory to its Tx FIFO. The transmitter takes data from the Tx FIFO and transmits data to the MAC. The MAC transmits the data through the GMII interface to the physical media. The transmitter, once initialized, runs until the end-of-frame (EOF) condition is detected unless a collision within the collision window occurs (half-duplex mode) or an abort condition is encountered.

If the user has a frame ready to transmit, setting the DMACTRL[TOD] eliminates waiting for the next poll and a DMA transfer of the transmit data buffers can begin immediately. The transmission begins once all data for the frame is loaded into the Tx FIFO or sufficient transmit data (determined by the Tx FIFO threshold register) is in the Tx FIFO. If the line is not busy, the MAC transmit logic asserts TX\_EN and sends the 7-octet preamble sequence, 1-octet start of frame delimiter, and frame information in that order. If the line is busy, the controller waits for the carrier sense signal, CRS, to remain inactive for 60 bit times (60 clocks) and transmission begins after an additional 36 bit times (96 bit times after CRS became active). In full-duplex mode, because collisions are ignored, frame transmission maintains only the interframe gap (96 bit times) regardless of CRS.

In half-duplex mode (MACCFG2[Full Duplex] is cleared) the MAC defers transmission if the line is busy (CRS asserted). Before transmitting, the MAC waits for carrier sense to become inactive, at which point it then determines if CRS remains negated for 60 clocks. If so, transmission begins after an additional 36 bit times (96 bit times after CRS originally became negated). If CRS continues to be asserted, the MAC follows a specified back-off procedure and tries to retransmit the frame until the retry limit is reached. Data stored in the Tx FIFO is re-transmitted in case of a collision. This avoids unnecessary memory traffic.

The transmitter also monitors for an abort condition and terminates the current frame if an abort condition is encountered. In full-duplex mode the protocol is independent of network activity, and only the transmit inter-frame gap must be enforced.

The transmitter implements full-duplex flow control. If a flow control frame is received, the MAC does not service the transmitter's request to send data until the pause duration is over. If the MAC is currently sending data after a pause frame has been received and processed, the MAC finishes sending the current frame, then suspends subsequent frames (except a pause frame) until the pause duration is over. In addition, the transmitter supports transmission of flow control frames through TCTRL[TFC\_PAUSE]. The transmit pause frame is generated internally based on the PAUSE register that defines the pause value to be sent. Note that it is possible to send a pause frame while the pause timer has not expired.

The MAC automatically appends FCS (32-bit CRC) bytes to the frame if any of the following values are set:

- TxBD[PAD/CRC] is set in first TxBD
- TxBD[TC] is set in first TxBD
- MACCFG2[PAD/CRC] is set
- MACCFG2[CRC] is set

The TX\_EN is negated after the FCS is sent. This notifies the PHY of the need to generate the illegal Manchester encoding that signifies the end of an Ethernet frame. Following the transmission of the FCS, the Ethernet controller writes the frame status bits into the BD and clears TxBD[R]. If the end of the current buffer is reached and TxBD[L] is cleared (a frame is comprised of multiple buffer descriptors), only TxBD[R] is cleared.



For both half- and full-duplex modes, an interrupt can be issued depending on TxBD[I]. The Ethernet controller then proceeds to the next TxBD in the table. In this way, the core can be interrupted after each frame, after each buffer, or after a specific buffer is sent. If TxBD[PAD/CRC] is set, the Ethernet controller pads any frame shorter than 64 bytes with zero bytes to make up the minimum length.

To pause transmission, or rearrange the transmit queue, set DMACTRL[GTS]. This can be useful for transmitting expedited data ahead of previously-linked buffers or for error situations. If this bit is set, the eTSEC transmitter performs a graceful transmit stop. The Ethernet controller stops immediately if no transmission is in progress or continues transmission until all queued frames in the Tx FIFO have been disposed of. The IEVENT[GTSC] interrupt occurs once the graceful transmit stop operation is completed. After the DMACTRL[GTS] is cleared, the eTSEC resumes transmission with the next frame.

While the eTSEC is in 10/100Mbps mode it sends bytes least-significant nibble first and each nibble is sent lsb first. While it is in 1000Mbps mode it sends bytes LSB first.

#### 15.6.2.4 Gigabit Ethernet Frame Reception

The eTSEC Ethernet receiver is designed to work with little core intervention and can perform data extraction, address recognition, CRC checking, short frame checking, and maximum frame-length checking.

After a hardware reset, the software driver clears the RSTAT register and sets MACCFG1[RX\_EN]. The Ethernet receiver is enabled and immediately starts processing receive frames. The MAC checks for when TSEC<sub>n</sub>\_RX\_DV is asserted and as long as TSEC<sub>n</sub>\_COL remains negated (full-duplex mode ignores TSEC<sub>n</sub>\_COL), the MAC looks for the start of a frame by searching for a valid preamble/SFD (start of frame delimiter) header, which is stripped (unless MACCFG2[PreAM RxEN] is set) and the frame begins to be processed. If a valid header is not found, the frame is ignored.

If the receiver detects the first bytes of a frame, the eTSEC controller begins to perform the frame recognition function through destination address (DA) recognition (see [Section 15.6.2.7, “Frame Recognition”](#)). Based on this match the frame can be accepted or rejected. The receiver can filter frames based on individual (unicast), group (multicast), and broadcast addresses. Because Ethernet receive frame data is not written to memory until the internal frame recognition algorithm is complete, system bus usage is not wasted on frames unwanted by this station.

If a frame is accepted, the Ethernet controller fetches the receive buffer descriptor (RxBd) from either queue 0 or the queue determined by the filer. If the RxBd is not being used by software (RxBd[E] is set), the eTSEC starts transferring the incoming frame. RxBd[F] is set for the first RxBd used for any particular receive frame. If the current RxBd is not available for the received frame, a receive busy error condition is raised in IEVENT[BSY].

After the buffer is filled, the eTSEC clears RxBd[E] and, if RxBd[I] is set, generates an interrupt. If the incoming frame is larger than the buffer, the Ethernet controller fetches the next RxBd in the table. If it is empty, the controller continues receiving the rest of the frame. In half-duplex mode, if a collision is detected during the frame, no RxBds are used; thus, no collision frames are presented to the user except late collisions, which indicate LAN problems.

The RxBd length is determined by the MRBL field in the maximum receive buffer length register (MRBLR). The smallest valid value is 64 bytes, with larger values being some integral multiple of 64

bytes. During reception, the Ethernet controller checks for frames that are too short or too long. After the frame ends (CRS is negated), the receive CRC field is checked and written to the data buffer. The data length written to the last RxBD in the Ethernet frame is the length of the entire frame, which enables the software to recognize an oversized frame condition.

Receive frames are not truncated when they exceed maximum frame bytes in the MAC's maximum frame register if MACCFG2[Huge Frame] is set, yet the babbling receiver error interrupt occurs (IEVENT[BABR] is set) and RxBD[LG] is set.

After the receive frame is complete, the Ethernet controller sets RxBD[L], updates the frame status bits in the RxBD, and clears RxBD[E]. If RxBD[I] is set, the Ethernet controller next generates an interrupt (that can be masked) indicating that a frame was received and is in memory. The Ethernet controller then waits for a new frame.

To interrupt reception or rearrange the receive queue, DMACTRL[GRS] must be set. If this bit is set, the eTSEC receiver performs a graceful receive stop. The Ethernet controller stops immediately if no frames are being received or continues receiving until the current frame either finishes or an error condition occurs. The IEVENT[GRSC] interrupt event is signaled after the graceful receive stop operation is completed. While in this mode the user can write to registers that are accessible to both the user and the eTSEC hardware without fear of conflict, and finally clear IEVENT[GRSC]. After DMACTRL[GRS] is cleared, the eTSEC scans the input data stream for the start of a new frame (preamble sequence and start of frame delimiter), it resumes receiving, and the first valid frame received is placed in the next available RxBD.

### 15.6.2.5 Ethernet Preamble Customization

By default eTSEC generates a standard Ethernet preamble sequence prior to transmitting frames. However, the user can substitute a custom preamble sequence for the purpose of controlling switching equipment at the receiver, particularly at 100/1000Mbps speeds.; in any RMII mode only the standard preamble can be transmitted

eTSEC normally searches for and discards the standard Ethernet preamble sequence upon receiving frames. Part of the received preamble sequence can be optionally recovered and returned as part of the frame data, making it visible to user software. Note however, that , and preamble cannot be recovered in any RMII mode. Note that it is also possible for the first two bytes of custom preamble (PreOct0 and PreOct1) to be lost in during conversion to ten-bit code groups in the PCS sub-layer. Thus is it recommended that any custom preamble start at PreOct2.

#### 15.6.2.5.1 User-Defined Preamble Transmission

To substitute a custom preamble, the user must ensure that:

- MACCFG2[PreAm TxEN] bit is set
- The first TxBD of every frame containing a custom preamble has its PRE bit set
- An 8-byte custom preamble sequence appears before the Ethernet DA field in the first transmit data buffer

The definition of the 8-byte custom preamble sequence is shown in [Figure 15-131](#).

Byte Offsets	0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15
0–1	PreOct0							PreOct1								
2–3	PreOct2							PreOct3								
4–5	PreOct4							PreOct5								
6–7	PreOct6															

**Figure 15-131. Definition of Custom Preamble Sequence**

The fields of the custom preamble sequence are described in [Table 15-144](#). It should be noted that use of preamble octets matching the standard start of frame delimiter (0xD5) can be expected to trigger premature frame reception by the receiving station.

**Table 15-144. Custom Preamble Field Descriptions**

Bytes	Bits	Name	Description
0–1	0–7	PreOct0	Octet #0 of custom transmit preamble. This is the first octet of preamble sent.
	8–15	PreOct1	Octet #1 of custom transmit preamble. This is the second octet of preamble sent.
2–3	0–7	PreOct2	Octet #2 of custom transmit preamble. This is the third octet of preamble sent.
	8–15	PreOct3	Octet #3 of custom transmit preamble. This is the fourth octet of preamble sent.
4–5	0–7	PreOct4	Octet #4 of custom transmit preamble. This is the fifth octet of preamble sent.
	8–15	PreOct5	Octet #5 of custom transmit preamble. This is the sixth octet of preamble sent.
6–7	0–7	PreOct6	Octet #6 of custom transmit preamble. This is the seventh octet of preamble sent. The last octet (the start of frame delimiter) is generated by the MAC automatically.
	8–15	—	Reserved; should be cleared.

### 15.6.2.5.2 User-Visible Preamble Reception

To return the received preamble, the user must ensure that:

- MACCFG2[PreAm RxEN] bit is set
- Space for an 8-byte preamble sequence is allowed before the Ethernet DA field in the first receive data buffer of each frame

The definition of the 8-byte received preamble sequence is shown in [Figure 15-132](#).

Byte Offsets	0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15
0–1	PreOct0							PreOct1								
2–3	PreOct2							PreOct3								
4–5	PreOct4							PreOct5								
6–7	PreOct6															

**Figure 15-132. Definition of Received Preamble Sequence**

The fields of the received preamble sequence are described in [Table 15-145](#). Should the received preamble be shorter than the 7-octet sequence defined by IEEE Std. 802.3, initial bytes of the received preamble

sequence hold undefined values. The standard start of frame delimiter (0xD5) is always omitted. Note that preamble extraction is not possible in RMII mode.

**Table 15-145. Received Preamble Field Descriptions**

Bytes	Bits	Name	Description
0–1	0–7	PreOct0	Octet #0 of received preamble. This is the first octet of preamble received.
	8–15	PreOct1	Octet #1 of received preamble. This is the second octet of preamble received.
2–3	0–7	PreOct2	Octet #2 of received preamble. This is the third octet of preamble received.
	8–15	PreOct3	Octet #3 of received preamble. This is the fourth octet of preamble received.
4–5	0–7	PreOct4	Octet #4 of received preamble. This is the fifth octet of preamble received.
	8–15	PreOct5	Octet #5 of received preamble. This is the sixth octet of preamble received.
6–7	0–7	PreOct6	Octet #6 of received preamble. This is the seventh octet of preamble received. The last octet (the start of frame delimiter) is discarded.
	8–15	—	Reserved

### 15.6.2.6 RMON Support

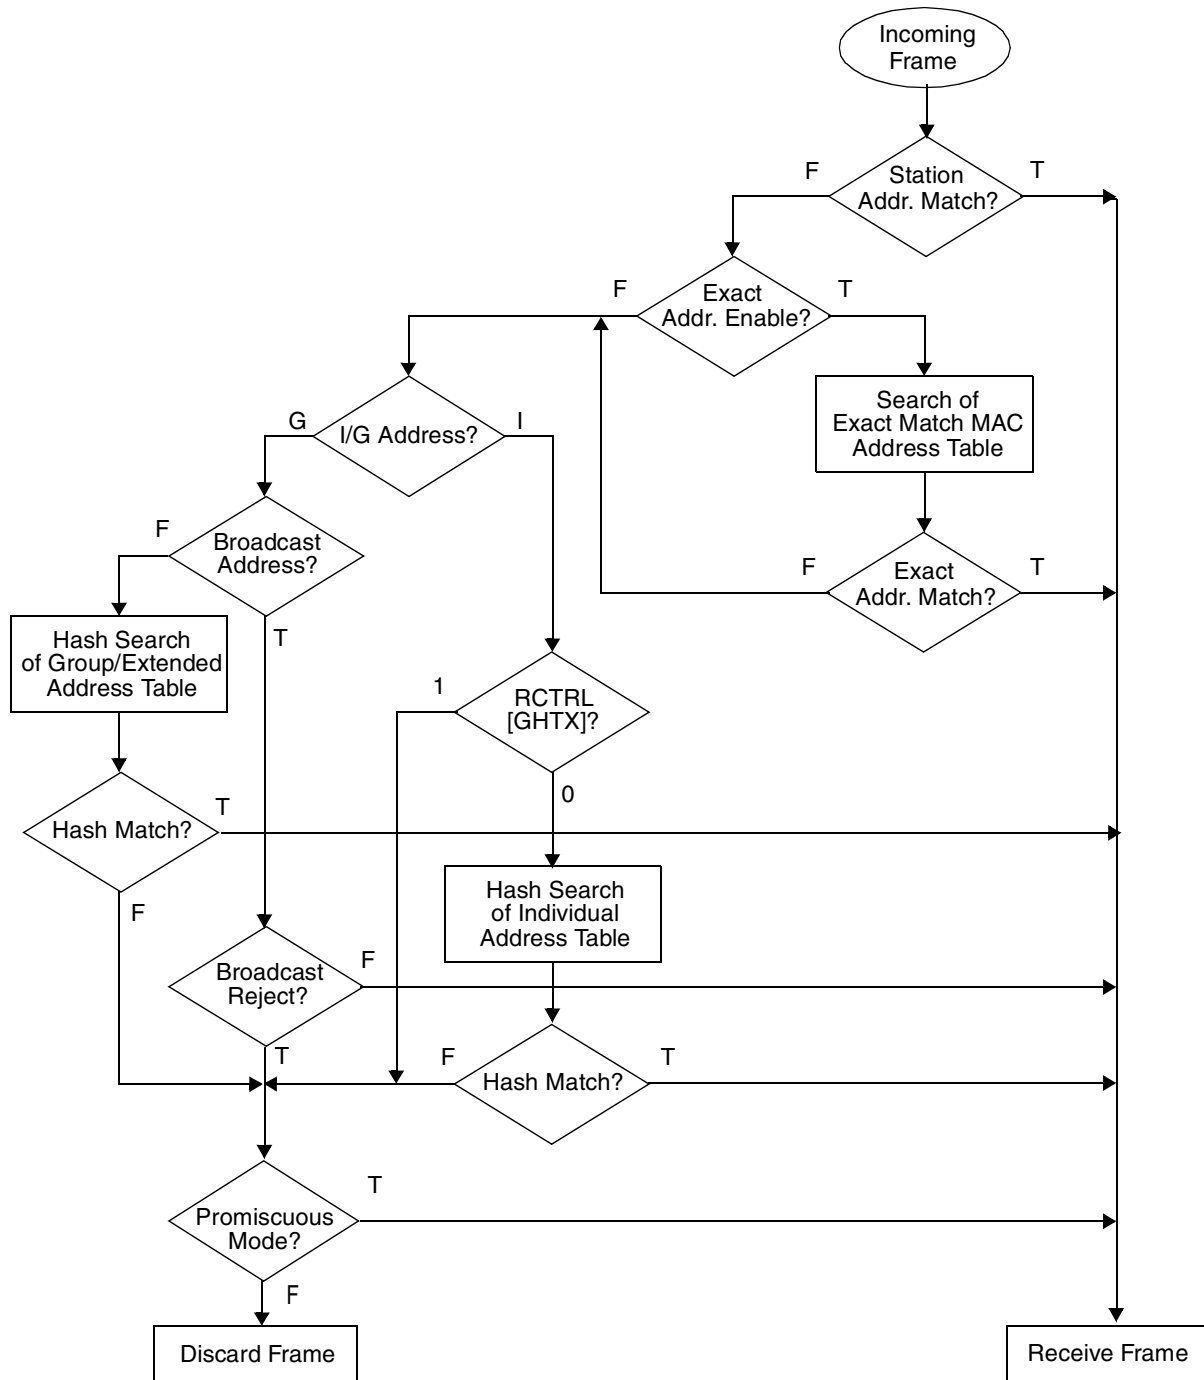
Using promiscuous mode, the eTSEC can automatically gather network statistics required for remote network interface monitoring. The RMON MIB group 1, RMON MIB group 2, RMON MIB group 3, RMON MIB group 9, RMON MIB2, and the IEEE 802.3 Ethernet MIB are supported. For RMON statistics and their corresponding counters, see the memory map.

### 15.6.2.7 Frame Recognition

The Ethernet controller performs frame recognition using destination address (DA) recognition. A frame can be rejected or accepted based on the outcome.

### 15.6.2.7.1 Destination Address Recognition and Frame Filtering

The eTSEC can perform layer 2 frame filtering on the basis of destination Ethernet address (DA), as illustrated by the flowchart in [Figure 15-133](#).



**Figure 15-133. Ethernet Address Recognition Flowchart**

In promiscuous mode, the eTSEC accepts all received frames regardless of DA. Note, however, that Ethernet frame filtering simply restricts the traffic seen by the receive queue filter. Therefore even in

promiscuous mode it remains possible to program the filter to reject frames based on their higher-layer header contents.

In the case of an individual address, the DA field of the received frame is compared with the physical address that the user programs in the station address registers (MACSTNADDR1 and MACSTNADDR2). If the DA does not match the station address, and exact MAC address matching is enabled through RCTRL[EMEN], the controller performs address recognition on the multiple MAC addresses written to the MACxADDR1 and MACxADDR2 registers. These virtual addresses give a particular eTSEC the ability to mirror other MACs on the network, which caters for router redundancy protocols, such as HSRP and VRRP.

If exact MAC address matching is not enabled, the eTSEC determines whether DA is a group or individual address. If DA is the standard broadcast address, and broadcast addresses are not rejected, the frame is accepted. If any other group address is received, the eTSEC looks-up the DA by means of the group hash table. The group hash table may be extended to 512 entries if RCTRL[GHTX] = 1. Otherwise, an individual address is hashed into the 256-entry individual hash table when RCTRL[GHTX] = 0.

### 15.6.2.7.2 Hash Table Algorithm

The hash table process used in the group hash filtering operates as follows. By default, the Ethernet controller maps any 48-bit destination address into one of 256 bins, represented by the 256 bits in IGADDR0–IGADDR7 for individual addresses, and the 256 bits in GADDR0–GADDR7 for group addresses. But in the case where RCTRL[GHTX] is set, both sets of registers are combined into an extended group-only hash table of 512 bits, where IGADDR0–IGADDR7 contain the first 256 bits and GADDR0–GADDR7 contain the last 256 bits. No individual-address table exists in extended mode.

The 48-bit destination address received by the MAC is passed through the Ethernet CRC-32 algorithm to produce a hash value. The CRC polynomial used is:

$$x^{32} + x^{26} + x^{23} + x^{22} + x^{16} + x^{12} + x^{11} + x^{10} + x^8 + x^7 + x^5 + x^4 + x^2 + x + 1$$

The MAC initializes its CRC register to 0xFFFFFFFF before computing a CRC on the 6 bit-reversed octets of the DA. A non-optimized sample of C code for computing the DA hash is listed in [Figure 15-134](#). The 9 most significant bits of the raw, uninverted CRC are used as the hash table index, H[8:0]. If RCTRL[GHTX] = 0, bits H[8:6] select one of the 8 IGADDR or GADDR registers, while bits H[5:1] select a bit within the 32-bit register. If RCTRL[GHTX] = 1, bits H[8:5] select one of the 16 registers in the {IGADDR, GADDR} set, while bits H[4:0] select a bit within the 32-bit register. For example, if H[8:5] = 7, IGADDR7 is selected, whereas H[8:5] = 9 selects GADDR1.

```

/* Wrapper macros for 256-bucket and 512-bucket hash tables:
   Pass 6-byte Ethernet MAC address as parameter. */
#define TSEC_HASH256(macaddr) ((crc32(macaddr) >> 24) & 0xff)
#define TSEC_HASH512(macaddr) ((crc32(macaddr) >> 23) & 0x1ff)

/* CRC constants. Note: CRC-32 polynomial is bit-reversed. */
#define CRC_POLYNOMIAL 0xedb88320
#define CRC_INITIAL    0xffffffff
#define MAC_ADDRLEN    6
#define BITS_PER_BYTE  8

/* crc32() Takes the array of bytes, macaddr[], representing an
   Ethernet MAC address and returns the CRC-32 result over these bytes,
   where each byte is used in bit-reversed form (Ethernet bit order).
   Index 0 of macaddr[] is the first byte of the address on the wire.
   Test case: the result of crc32 on {0x00, 0x01, 0x02, 0x03, 0x04, 0x05}
   should be 0xad0c28f3.
   */
unsigned long crc32(unsigned char macaddr[MAC_ADDRLEN])
{
    unsigned long crc, result;
    int byte, i;

    /* CRC-32 algorithm starts by inverting first 4 bytes */
    crc = CRC_INITIAL;
    /* add each byte to running CRC accumulator */
    for (byte = 0; byte < MAC_ADDRLEN; ++byte) {
        crc ^= macaddr[byte];
        /* shift CRC right to perform but reversal on byte of address */
        for (i = 0; i < BITS_PER_BYTE; ++i)
            if (crc & 1)
                crc = (crc >> 1) ^ CRC_POLYNOMIAL;
            else
                crc >>= 1;
    }
    /* finally, reverse bits of result to get CRC in normal bit order */
    for (result = 0, i = 4*BITS_PER_BYTE-1; i >= 0; crc >>= 1, --i)
        result |= (crc & 1) << i;
    return result;
}

```

**Figure 15-134. Sample C Code for Computing eTSEC Hash Table Indices**

If the CRC hash table index selects a bit that is set in the hash table, the frame is accepted. If 32 group addresses are stored in the hash table and random group addresses are received, the extended hash table prevents roughly 480/512 (93.8%) of the group address frames from reaching memory. Software must further filter those that reach memory to determine if they contain the correct addresses. Alternatively, small multicast groups can be held in the exact match MAC address registers, which guarantees that only correct frames are admitted.

The effectiveness of the hash table declines as the number of addresses increases. For instance, as the number of addresses stored in the 512-bin hash table increases, the vast majority of the hash table bits are set, preventing only a small fraction of frames from reaching memory.

**NOTE**

The hash table cannot be used to reject frames that match a set of selected addresses because unintended addresses can map to the same bit in the hash table. The receive queue filter may be used to reject frames with unintended address hits in the hash table.

**15.6.2.8 Magic Packet Mode**

eTSEC implements the AMD Magic Packet™ specification for LAN-initiated power management. This mode is normally entered with the rest of the system in a low-power sleep mode. Software must enable normal receive function in the Ethernet MAC, and then finally set the MACCFG2[MPEN] bit to enable Magic Packet detection before the system enters a reduced mode. While the rest of the system is operating in low-power mode, the enabled eTSEC continues to receive Ethernet frames, but discards them immediately. Upon receipt of any frame whose contents contain the valid Magic Packet sequence, the eTSEC exits out of Magic Packet mode, thus clearing MACCFG2[MPEN], and raises an error/diagnostic interrupt through IEVENT[MAG], which causes the surrounding system to wake-up. Frames received after Magic Packet mode has exited are received into software buffers as usual. Software can abort Magic Packet mode by writing 0 to MACCFG2[MPEN] at any time.

AMD specify a Magic Packet™ to be any Ethernet frame containing a valid Ethernet header (Destination and Source Addresses) and valid FCS (CRC-32), and whose payload includes the specific Magic Packet byte sequence at any offset from the start of frame. The specific byte sequence comprises an unbroken stream of 102 bytes, the first 6 bytes of which are 0xFFFFFFFF\_FFFFFFFF, followed by 16 copies of the MAC's unique IEEE station address in the normal byte order for Ethernet addresses. For example, if the station address were 0x112233\_445566, then the MAC would have to receive 0xFFFFFFFF\_FFFFFFFF, 0x112233\_445566, ..., 0x112233\_445566 in any payload to detect a Magic Packet. Only frames addressed specifically to the MAC's station address or a valid multicast or broadcast address can be examined for the Magic Packet sequence.

**15.6.2.9 Flow Control**

Because collisions cannot occur in full-duplex mode, gigabit Ethernet can operate at the maximum rate. If the rate becomes too fast for a station's receiver, the station's transmitter can send flow-control frames to reduce the rate. Flow-control instructions are transferred by special frames of minimum frame size. The length/type fields of these frames have a special value.

Table 15-146 lists the flow-control frame structure.

**Table 15-146. Flow Control Frame Structure**

Size [Octets]	Description	Value	Comment
7	Preamble		—
1	SFD		Start frame delimiter
6	Destination address	01-80-C2-00-00-01	Multicast address reserved for use in MAC frames (or MAC station address)
6	Source address		—
2	Length/type	88-08	Control frame type



Table 15-146. Flow Control Frame Structure (continued)

Size [Octets]	Description	Value	Comment
2	MAC opcode	00-01	Pause command
2	MAC parameter		Pause time as defined by the PTV[PT] field. The pause period is measured in pause_quanta, a speed independent constant of 512 bit-times (unlike slot time). The most-significant octet is transmitted first.
2	Extended MAC parameter		Pause time extended as defined by the PTV[PTE] field. The most significant octet is transmitted first.
40	Reserved	—	—
4	FCS		Frame check sequence (CRC)

If flow-control mode is enabled (MACCFG1[Rx\_Flow] is set) and the receiver identifies a pause-flow control frame, transmission stops for the time specified in the control frame. Since the pause timer commences counting immediately upon receipt of a PAUSE frame, regardless of whether transmission is currently in progress, a sufficiently large pause time must be received to stop transmission past a frame of MTU size. During a pause, only a control frame can be sent (TCTRL[TFC\_PAUSE] is set). Normal transmission resumes after the pause timer stops counting, or resumes immediately if a pause frame with a zero time-out is received. If another pause-control frame is received during the pause, the period changes to the new value received.

### 15.6.2.10 Interrupt Handling

The following describes what usually occurs within a eTSEC interrupt handler:

- If an interrupt occurs, read IEVENT to determine interrupt sources. IEVENT bits to be handled in this interrupt handler are normally cleared at this time. There are three kinds of interrupts:
  - Receive data frame interrupts, when bits RXB or RXF in IEVENT are set
  - Transmit data frame interrupts, when bits TXB or TXF in IEVENT are set
  - Error, diagnostic, and special interrupts (all bits in IEVENT other than RXB, RXF, TXB, or TXF)
- Process the TxBDs to reuse them if the IEVENT[TXB, TXF or TXE] were set. Consult register bits TSTAT[TXF0–TXF7] to determine which TxBD rings gave rise to the transmit interrupt in the case of TXF. If the transmit speed is fast or the interrupt delay is long, more than one transmit buffer may have been sent by the eTSEC; thus, it is important to check more than just one TxBD during the interrupt handler. One common practice is to process all TxBDs in the interrupt handler until one is found with R set.
- Obtain data from RxBD rings if IEVENT[RXC, RXB or RXF] is set. Consult register bits RSTAT[RXF0–RXF7] to determine which RxBD rings gave rise to the receive interrupt in the case of RXF. If the receive speed is fast or the interrupt delay is long, the eTSEC may have received more than one RxBD; thus, it is important to check more than just one RxBD during interrupt handling. Typically, all RxBDs in the interrupt handler are processed until one is found with E set. Because the eTSEC pre-fetches BDs, the BD table must be big enough so that there is always another empty BD to pre-fetch, otherwise a BSY error occurs.

- Clear any set halt or frame interrupt bits in TSTAT and RSTAT registers, or DMACTRL[GTS] and DMACTRL[GRS] by writing 1s to these bits.
- Continue normal execution.

**Table 15-147. Non-Error Transmit Interrupts**

Interrupt	Description	Action Taken by the eTSEC
GTSC	Graceful transmit stop complete: transmitter is put into a pause state after completion of the frame currently being transmitted.	None
TXC	Transmit control: Instead of the next transmit frame, a control frame was sent.	None
TXB	Transmit buffer: A transmit buffer descriptor, that is not the last one in the frame, was updated in one of the enabled TxBD rings.	Programmable 'write with response' TxBD to memory before setting IEVENT[TXB].
TXF	Transmit frame: A frame from an enabled TxBD ring was transmitted and the last transmit buffer descriptor (TxBD) of that frame was updated.	Programmable 'write with response' to memory on the last TxBD before setting IEVENT[TXF].

**Table 15-148. Non-Error Receive Interrupts**

Interrupt	Description	Action Taken by the eTSEC
GRSC	Graceful receive stop complete: Receiver is put into a pause state after completion of the frame currently being received.	None
RXC	Receive control: A control frame was received. As soon as the transmitter finishes sending the current frame, a pause operation is performed.	None
RXB	Receive buffer: A receive buffer descriptor, that is not the last one of the frame, was updated in one of the enabled RxBD rings.	Programmable 'write with response' RxBD to memory before setting IEVENT[RXB].
RXF	Receive frame: A frame was received to an enabled RxBD ring and the last receive buffer descriptor (RxBD) of that frame was updated.	Programmable 'write with response' to memory on the last RxBD before setting IEVENT[RXF].

### 15.6.2.10.1 Interrupt Coalescing

Interrupt coalescing offers the user the ability to contour the behavior of the eTSEC with regard to frame interrupts. Separate but identical mechanisms exist for both transmitted frames and received frames. In either case, frame interrupts require that software set the I-bit in RxBDs or TxBDs, and disable buffer interrupts (IEVENT[RXB] or IEVENT[TXB]). Particular rings can remain free of interrupts by ensuring that the I-bit is consistently cleared in all BDs. While interrupt coalescing is enabled, a transmit or receive frame interrupt is raised either when a counter threshold-defined number of frames is received/transmitted or the timer threshold-defined period of time has elapsed, whichever occurs first. Disabling and then re-enabling interrupt coalescing forces reset of the coalescing timers and counters to reflect changes made to the threshold registers.

### 15.6.2.10.2 Interrupt Coalescing By Frame Count Threshold

To avoid interrupt bandwidth congestion due to frequent, consecutive interrupts, the user may enable and configure interrupt coalescing to deliberately group frame interrupts, reducing the total number of

interrupts raised. The number of frames received or transmitted prior to an interrupt being raised is determined by the frame threshold field (ICFT) in the appropriate interrupt coalescing configuration register (RXIC or TXIC). The frame threshold field may be assigned a value between 1 and 255. The internal transmit or receive frame counter decrements from this initial value each time a frame is transmitted or received. Upon reaching zero, an interrupt is raised, the appropriate threshold counter is reset to the value in the ICFT field, and then eTSEC continues counting frames while the interrupt is active. The appropriate threshold counter is also reset to the value in the ICFT field if an interrupt is raised subject to the corresponding threshold timer.

### 15.6.2.10.3 Interrupt Coalescing By Timer Threshold

To avoid stale frame interrupts, the user may also assign a timer threshold, beyond which any frame interrupts not yet raised are forced. The timer threshold fields of the receive and transmit interrupt coalescing configuration registers (RXIC[ICTT] and TXIC[ICTT]) are defined in units equivalent to 64 interface clocks or system clocks, depending on the setting of the ICCS field in RXIC and TXIC.

After transmitting a frame, the transmit interrupt coalescing threshold time begins counting down from the value in TXIC[ICTT]. An interrupt is raised when the counter reaches zero. In the event of graceful transmit stop completion before the coalescing timer expires, the eTSEC issues two interrupts, the first for GTS, the second for TXF (due to timer expiration of a pending event). To prevent the second interrupt from affecting servicing of the GTS event, it is recommended that the user mask out the TXF event during execution of the service routine. After receiving a frame, the receive interrupt coalescing threshold time begins counting down from the value in RXIC[ICTT]. An interrupt is raised when the counter reaches zero. In the event of graceful receive stop completion before the coalescing timer expires, the eTSEC issues two interrupts, the first for GRS, the second for RXF (due to timer expiration of a pending event). To prevent the second interrupt from affecting servicing of the GRS event, it is recommended that the user mask out the RXF event during execution of the service routine.

The interrupt coalescing timer thresholds (transmit and receive, operating independently) may be values ranging from 0x0001 to 0xFFFF. Table 15-149 specifies the range of possible timing thresholds subject to timer clock source, the interface or system frequency, and the value of the RXIC[ICTT] or TXIC[ICTT] field.

**Table 15-149. Interrupt Coalescing Timing Threshold Ranges**

ICCS (Clock Source)	eTSEC Interface Format and Frequency or eTSEC System Frequency	Interrupt Coalescing Threshold Time	
		Minimum (ICTT = 0x0001)	Maximum (ICTT = 0xFFFF)
0 (I/F clock)	10Base-T at 2.5 MHz	25.6 $\mu$ s	1.68 s
0 (I/F clock)	100Base-T at 25 MHz	2.56 $\mu$ s	168 ms
0 (I/F clock)	1000Base-T at 125 MHz	0.51 $\mu$ s	33.6 ms
1 (sys. clock)	eTSEC operating at 133 MHz	0.48 $\mu$ s	31.4 ms
1 (sys. clock)	eTSEC operating at 166 MHz	0.38 $\mu$ s	25.2 ms

The transmit timer threshold counter is reset to the value in TXIC[ICTT] and begins counting down on transmission of the frame following an interrupt.

The receive timer threshold counter is reset to the value in RXIC[ICTT] and begins counting down on receiving the frame following an interrupt.

### 15.6.2.11 Inter-Frame Gap Time

If a station must transmit, it waits until the LAN becomes silent for a specified period (inter-frame gap, or IFG). The minimum inter-packet gap (IPG) time for back-to-back transmission is set by IPGIFG[Back-to-Back Inter-Packet-Gap]. The receiver receives back-to-back frames with the minimum interframe gap (IFG) as set in IPGIFG[Minimum IFG Enforcement]. If multiple frames are ready to transmit, the Ethernet controller follows the minimum IPG as long as the following restrictions are met:

- The first TxBD pointer, TBPTR $n$ , of any given frame is located at a 16-byte aligned address.
- Each TxBD[Data Length] is greater-than or equal to 64 bytes.

If the first TxBD alignment restriction is not met, the back-to-back IPG may be as many as 32 cycles. If the TxBD size restriction is not met, the back-to-back IPG may be significantly longer.

In half-duplex mode, after a station begins sending, it continually checks for collisions on the LAN. If a collision is detected, the station forces a jam signal (all ones) on its frame and stops transmitting. Collisions usually occur close to the beginning of a packet. The station then waits a random time period (back-off) before attempting to send again. After the back-off completes, the station waits for silence on the LAN (carrier sense negated) and then begins retransmission (retry) on the LAN. Retransmission begins 36 bit times after carrier sense is negated for at least 60 bit times. If the frame is not successfully sent within a specified number of retries, an error is indicated (collision retry limit exceeded).

### 15.6.2.12 Internal and External Loop Back

Setting MACCFG1[Loop Back] causes the MAC transmit outputs to be looped back to the MAC receive inputs. Clearing this bit results in normal operation. This bit is cleared by default. Clearing this bit results in normal operation.

### 15.6.2.13 Error-Handling Procedure

The eTSEC reports frame reception and transmission error conditions using the channel BDs, the error counters, and the IEVENT register.

Transmission errors are described in [Table 15-150](#).

**Table 15-150. Transmission Errors**

Error	Response
Transmitter underrun	Transmitter underrun can occur either after frame transmission has commenced, or in response to an incomplete sequence of TxBDs. In the former case, the controller sends 32 bits that ensure a CRC error, and terminates buffer transmission. In the latter case, the relevant transmit queue is halted. In all cases, the eTSEC closes the buffer, sets TxBD[UN], IEVENT[XFUN], and IEVENT[TXE]. The controller resumes transmission after TSTAT[THLT] is cleared (and DMACTRL[GTS] is cleared).
Retransmission attempts limit expired	The controller terminates buffer transmission, sets TxBD[RL], closes the buffer, IEVENT[CRL], and IEVENT[TXE]. Transmission resumes after TSTAT[THLT] is cleared (and DMACTRL[GTS] is cleared).

**Table 15-150. Transmission Errors (continued)**

Error	Response
Late collision	The controller terminates buffer transmission, sets TxBD[LC], closes the buffer, IEVENT[LC], and IEVENT[TXE]. The controller resumes transmission after TSTAT[THLT] is cleared (and DMACTRL[GTS] is cleared).
Memory read error	A system bus error occurred during a DMA transaction. The controller sets IEVENT[EBERR], DMA stops sending data to the FIFO which causes an underrun error, and therefore TxBD[UN] is set, but IEVENT[XFUN] is not set. The TSTAT[THLT] is set. Transmits are continued once TSTAT[THLT] is cleared.
Data parity error	Data in the transmit FIFO was potentially corrupted. The controller sets IEVENT[DPE], but otherwise continues transmission until halted explicitly.
Babbling transmit error	A frame is transmitted which exceeds the MAC's Maximum Frame Length and MACCFG2[Huge Frame] is a 0. The controller sets IEVENT[BABT] and continues without interruption.

Reception errors are described in [Table 15-151](#).

**Table 15-151. Reception Errors**

Error	Description
Overrun error	The Ethernet controller maintains an internal FIFO buffer for receiving data. If a receiver FIFO buffer overrun occurs, the controller sets RxBDOV, sets RxBDL, closes the buffer, increments the discarded frame counter (RDRP), and sets IEVENT[RXF]. The receiver then enters hunt mode (seeking start of a new frame).
Busy error	A frame is received and discarded due to a lack of buffers. The controller sets IEVENT[BSY] and increments the discarded frame counter (RDRP). In addition, the RSTAT[QHLT $n$ ] bit is set. RDRP increments for each frame that is received while the receiver is halted due to a busy condition. The halted queue resumes reception once the RSTAT[QHLT $n$ ] bit is cleared.
Filed frame to invalid queue error	A frame is received and discarded as a result of the filer directing it to an RxBDRing that is currently not enabled. The controller sets IEVENT[FIQ] and increments the discarded frame counter (RDRP).
Parser error	If the receive frame parser is enabled, a parse error can be flagged as a result of inconsistencies discovered between fields of the embedded packet headers. For example, the L2 header may indicate an IPv4 header, but the IP version number fails to match. In the event of a parse error, parsing is terminated at the inconsistent header, and the RxFCB[PERR] field indicates at which layer of the protocol stack the error was discovered. Receiver function continues regardless of parse errors, but IEVENT[PERR] is set. The receive queue filer may operate with reduced or default information in some cases; therefore, filer rule sets should be constructed so as to be tolerant of malformed frames. <b>Note:</b> Any values in the length/type field between 1500 and 1536 is treated as a length, however, only illegal packets exist with this length/type since these are not valid lengths and not valid types. These are treated by the MAC logic as out of range. Software must confirm the parser and filer results by checking the type/length field after the packet has been written to memory to see if it falls in this range.
Non-octet error (dribbling bits)	The Ethernet controller handles a nibble of dribbling bits if the receive frame terminates as non-octet aligned and it checks the CRC of the frame on the last octet boundary. If there is a CRC error, the frame non-octet aligned (RxBDOV) error is reported, IEVENT[RXF] is set, and the alignment error counter increments. The eTSEC relies on the statistics collector block to increment the receive alignment error counter (RALN). If there is no CRC error, no error is reported.

**Table 15-151. Reception Errors (continued)**

Error	Description
CRC error	If a CRC error occurs, the controller sets RxB[CR], closes the buffer, and sets IEVENT[RXF]. This eTSEC relies on the statistics collector block to record the event. After receiving a frame with a CRC error, the receiver then enters hunt mode.
Memory read error	A system bus error occurred during a DMA transaction. The controller sets IEVENT[EBERR] and discards the frame and increments the discarded frame counter (RDRP). In addition the RSTAT[QHLT $n$ ] bit is set. The halted queue resumes reception once the RSTAT[QHLT $n$ ] bit is cleared.
Data parity error	Data in the receive FIFO or filter table was potentially corrupted. The controller sets IEVENT[DPE], but otherwise continues reception until halted explicitly.
Babbling receive error	A frame is received that exceeds the MAC's maximum frame length. The controller sets IEVENT[BABR] and continues.

### 15.6.3 TCP/IP Off-Load

Each eTSEC provides hardware support for accelerating the basic functions of TCP/IP packet transmission and reception. By default, these features are disabled and must be explicitly enabled through RCTRL and TCTRL. In this configuration, the eTSEC processes frames as vanilla Ethernet frames and none of the multi-ring QoS/CoS receive services or per-frame VLAN insertion and deletion are available. Operate eTSEC in this default configuration when using existing TCP/IP stack software that has not been modified to take advantage of TOE.

TOE can be enabled independently for Rx and Tx and at various levels. Receive TOE functions are controlled by RCTRL and transmit functions through a combination of TCTRL[TUCSEN] and the Tx frame control block.

On receive, according to RCTRL[PRSDEP], eTSEC can parse frames at layer 2 of the stack only (Ethernet headers and switching headers), layers 2 to 3 (including IPv4 or IPv6), or layers 2 to 4 (including TCP and UDP). TOE provides protocol header recognition, header verification (IPv4 header checksum verification), and TCP/UDP payload checksum verification including verification of associated pseudo-header checksums. For large frames off-load of checksum verification saves a significant fraction of the CPU cycles that would otherwise be spent by the TCP/IP stack. IP packet fragmentation and re-assembly, and TCP stream establishment and tear-down are not performed in hardware. The frame parser sets RQFPR[IPF] status flag encountering a fragmented frame. The frame parser in eTSEC searches a maximum of 512 bytes from the start of a received frame when attempting to locate headers; headers deeper than 512 bytes are assumed not to exist, and any associated receive status flags in the frame control block remain cleared.

On transmit, TOE provides IPv4 and TCP/UDP header checksum generation. Like receive TOE, checksum generation reduces CPU load significantly for TCP/IP stacks modified to exploit eTSEC TOE functions. The eTSEC does not checksum transmitted packets with IPv6 routing headers or calculate TCP/UDP checksums from IP fragments. If a transmitted TCP segment requires checksum generation but IPv6 extension headers would prevent eTSEC from calculating the pseudo-header checksum, software can calculate just the pseudo-header checksum in advance and supply it to the eTSEC as part of per-frame TOE configuration.

### 15.6.3.1 Frame Control Blocks

Frame control blocks (FCBs) are 8-byte blocks of TOE control and/or status data that are passed between software (driver and TCP/IP stack) and each eTSEC. A FCB always precedes the frame it applies to, and is present only when TOE functions are being used. As Figure 15-135 shows, the first BD of each frame points to the initial data buffer and the FCB. The initial data buffer must be at least 8 bytes long to contain the FCB without breaking it. Custom or received Ethernet preamble sequences also follow the FCB if preambles are visible.

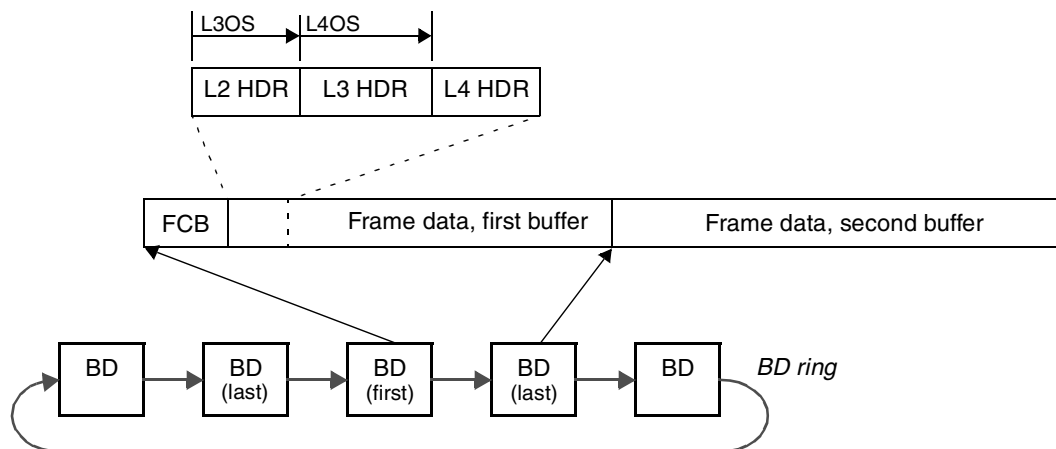


Figure 15-135. Location of Frame Control Blocks for TOE Parameters

For TxBD rings, FCBs are assumed present when the TxBD[TOE/UN] bit is set by user software. The eTSEC ignores the TxBD[TOE/UN] bit in all BDs other than those pointing to initial data buffers, therefore FCBs must not be inserted in second and subsequent data buffers. Since TxBD[TOE/UN] can be set under software discretion, TOE acceleration for transmit may be applied on a frame-by-frame basis.

In the case of RxBD rings, FCBs are inserted by the eTSEC whenever RCTRL[PRSDEP] is set to a non-zero value. Only one FCB is inserted per frame, in the buffer pointed to by the RxBD with bit F set. TOE acceleration for receive is enabled for all frames in this case.

### 15.6.3.2 Transmit Path Off-Load and Tx PTP Packet Parsing

TOE functions for transmit are defined by the contents of the Tx FCB. Figure 15-136 describes the definition for the Tx FCB.

	0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15
Offset + 0	VLN	IP	IP6	TUP	UDP	CIP	CTU	NPH								PTP
Offset + 2	L4OS							L3OS								
Offset + 4	PHCS															
Offset + 6	VLCTL															

Figure 15-136. Transmit Frame Control Block

The user instructs the Tx packet to be timestamped via setting bit 15 in the TxFCB to mark a PTP packet. TxFCB[VLCTL] can be translated as the Tx PTP packet identification number. BD[TOE] has to be set to enable transmit PTP packet timestamping. TxFCB[PTP] bit takes precedence over TxFCB[VLN] bit. It disables per packet VLAN tag insertion. On a PTP packet, VLAN tag can be inserted from the DFVLAN register. A proposed TxFCB update for the PTP packet is shown in [Figure 15-143](#).

The contents of the Tx FCB are defined in [Table 15-152](#).

**Table 15-152. Tx Frame Control Block Description**

Bytes	Bits	Name	Description
0–1	0	VLN	VLAN control word valid. This bit is ignored when the PTP bit is set. VLAN tag is read from the DFVLAN register if PTP=1. 0 Ignore VLCTL field. 1 If VLAN tag insertion is enabled for eTSEC, use the VLCTL field as the VLAN control word.
	1	IP	Layer 3 header is an IP header. 0 Ignore layer 3 and higher headers. 1 Assume that the layer 3 header is an IPv4 or IPv6 header, and take L3OS field as valid.
	2	IP6	IP header is IP version 6. Valid only if IP = 1. 0 IP header version is 4. 1 IP header version is 6.
	3	TUP	Layer 4 header is a TCP or UDP header. 0 Do not process any layer 4 header. 1 Assume that the layer 4 header is either TCP or UDP (see UDP bit), and offload checksumming on the basis that the IP header has no extension headers.
	4	UDP	UDP protocol at layer 4. 0 Layer 4 protocol is either TCP (if TUP = 1) or undefined. 1 Layer 4 protocol is UDP if TUP = 1.
0–1	5	CIP	Checksum IP header enable. 0 Do not generate an IP header checksum. 1 Generate an IPv4 header checksum.
	6	CTU	Checksum TCP or UDP header enable. 0 Do not generate a TCP or UDP header checksum. RFC 768 advises that UDP packets not requiring checksum validation should have their checksum field set to zero. 1 Generate a TCP header checksum if IP = 1 and TUP = 1 and UDP = 0.
	7	NPH	Disable calculation of TCP or UDP pseudo-header checksum. This bit should be set if IP options need to be consulted in forming the pseudo-header checksum, as eTSEC does not examine IP options or extension headers for TCP/IP offload on transmit. 0 Calculate TCP or UDP pseudo-header checksum as normal, assuming that the IP header has no options. 1 Do not calculate a TCP or UDP pseudo-header checksum, but instead use the value in field PHCS when determining the overall TCP or UDP checksum.
	8–14	—	Reserved
	15	PTP	Indication to the transmitter that this is a PTP packet. Enabling PTP disables per packet VLAN tag insertion. Instead, VLAN tag will be read from the DFVLAN when the PTP field is true. 0 Do not attempt to capture transmission event time 1 Valid PTP_ID field. When this packet is transmitted, capture the time of transmission. Must be clear if TMR_CTRL[TE] is clear.



Table 15-152. Tx Frame Control Block Description (continued)

Bytes	Bits	Name	Description
2–3	0–7	L4OS	Layer 4 header offset from start of layer 3 header. The layer 4 header starts L4OS octets after the layer 3 header if it is present. The maximum layer 3 header length supported is thus 255 bytes, which may prevent TCP/IP offload on particularly large IPv6 headers.
	8–15	L3OS	Layer 3 header offset from start of frame not including the 8 bytes for this FCB. The layer 3 header starts L3OS octets from the start of the frame including any custom preamble header that may be present. The maximum layer 2 header length supported is thus 255 bytes.
4–5	0–15	PHCS	Pseudo-header checksum (16-bit one's complement sum with carry wraparound, but without result inversion) for TCP or UDP packets, calculated by software. Valid only if NPH = 1.
6–7	0–15	VLCTL/ PTP_ID	VLAN control word for insertion in the transmitted VLAN tag. Valid only if VLN = 1. Tx PTP packet identification number. This number will be copied into the Tx PTP packet timestamp identification field. PTP field takes precedence over VLN field.

### 15.6.3.3 Receive Path Off-Load

Upon receive, the Rx FCB returns the status of frame parse and TOE functions applied to the accompanying frame. [Figure 15-137](#) describes the definition for the Rx FCB.

	0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15
Offset + 0	VLN	IP	IP6	TUP	CIP	CTU	EIP	ETU					PERR			
Offset + 2				RQ				PRO								
Offset + 4																
Offset + 6	VLCTL															

Figure 15-137. Receive Frame Control Block

The contents of the Rx FCB are defined in [Table 15-153](#).

**Table 15-153. Rx Frame Control Block Descriptions**

Bytes	Bits	Name	Description
0–1	0	VLN	VLAN tag recognized. This bit is set only if RCTRL[VLEX] is set. 0 No VLAN tag recognized. 1 IEEE Std. 802.1Q VLAN tag found; VLAN control word in VLCTL is valid.
	1	IP	IP header found at layer 3. RCTRL[PRSDEP] must be set to 10 or 11 in order to enable IP discovery. See also IP6 bit of FCB. 0 No layer 3 header recognized. 1 An IP header was recognized at layer 3; the IANA protocol identifier for the next header can be found in PRO; see PRO for more information.  If S/W is relying on the RxFCB for the parse results, any RxFCB[IP] bits set with the corresponding RxFCB[PRO] = 0xFF indicates a fragmented packet (or that this packet had a back-to-back IPv6 routing extension header). Additionally, RQFPR[IPF] (see <a href="#">Section 15.5.3.3.8, “Receive Queue Filer Table Property Register (RQFPR)”</a> ) indicates that the packet was fragmented.
	2	IP6	IP version 6 header found at layer 3. 0 No IPv6 header was found. 1 The layer 3 header was an IPv6 header provided IP = 1.
	3	TUP	TCP or UDP header found at layer 4. RCTRL[PRSDEP] must be set to 10 or 11 in order to enable TCP/UDP discovery. 0 No layer 4 header recognized. 1 The layer 4 header was recognized as either TCP (PRO = 0x06) or UDP (PRO = 0x11).
	4	CIP	IPv4 header checksum checked. RCTRL[PRSDEP] must be set to 10 or 11 in order to enable IPv4 checksum verification. 0 IPv4 header checksum not verified, either because verification was disabled or a valid IPv4 header could not be located. 1 IPv4 header checksum was verified by the eTSEC, and bit EIP indicates result.
	5	CTU	TCP or UDP header checksum checked. RCTRL[PRSDEP] must be set to 11 in order to enable layer 4 checksum verification. 0 TCP or UDP header checksum not verified, either because verification was disabled or a valid TCP or UDP header could not be located. If a UDP header with zero checksum was located, this bit is cleared in accordance with RFC 768. 1 TCP or UDP header checksum was verified by the eTSEC, and ETU indicates result.
	6	EIP	IPv4 header checksum verification error. Not valid unless CIP = 1. 0 No checksum error in IPv4 header. 1 Error in header checksum only if IP = 1 and IP6 = 0.
0–1	7	ETU	TCP or UDP header checksum verification error. Not valid unless CTU = 1. 0 No checksum error in TCP or UDP header. 1 Error in header checksum only if PRO = 0x06 or PRO = 0x11.
	8–11	—	Reserved
	12–13	PER	Parse error. 00 No error in L2 to L4 parse 01 Reserved 10 Inconsistent or unsupported L3 header sequence 11 Reserved
	14	—	Reserved
	15	GFPF	General-purpose filer event packet. This packet was filed based on matching a GPI rule sequence.

Table 15-153. Rx Frame Control Block Descriptions (continued)

Bytes	Bits	Name	Description
2–3	0–1	—	Reserved
	2–7	RQ	Receive queue index. This index was selected by the eTSEC Rx Filer (from a matching Filer rule's RQCTRL[Q] field) when it accepted the associated frame. If filing is not enabled, RQ is zero. Note that the 3 least significant bits of RQ correspond with the RxBD ring index whenever RCTRL[FSQEN] = 0.
	8–15	PRO	<p>If IP = 1, PRO is set as follows:</p> <ul style="list-style-type: none"> <li>• PRO=0xFF for a fragment header or a back to back route header</li> <li>• PRO=0xnn for an unrecognized header, where nn is the next protocol field</li> <li>• PRO=(TCP/UDP header), as defined in the IANA specification, if TCP or UDP header is found</li> </ul> <p>If IP = 0, PRO is undefined.</p> <p>Note that the eTSEC parser logic stops further parsing when encountering an IP datagram that has indicated that it has fragmented the upper layer protocol. This in general means that there is likely no layer 4 header following the IP header and extension headers. eTSEC leaves the RxFCB[PRO] and RQFPR[L4P] fields 0xFF in this case, which usually means that there was no IP header seen. In this case RxFCB[IP] and optionally RxFCB[IP6] is set. IP header checksumming operates and performs as intended. Most of the time, the eTSEC updates the RxFCB[PRO] field and RQFPR[L4P] fields with whatever value was found in the protocol field of the IP header. See <a href="#">Section 15.5.3.3.8, “Receive Queue Filer Table Property Register (RQFPR)”</a>, for a description of RQFPR.</p>
4–5	0–15	—	Reserved
6–7	0–15	VLCTL	VLAN control word as per IEEE Std. 802.1Q. The lower 12 bits comprise the VLAN identifier. Valid only if VLN = 1.

## 15.6.4 Quality of Service (QoS) Provision

This section describes the quality of service support features of this device. It includes a parser which extracts vital packet properties and passes them to the filer which essentially acts as a frame classifier.

### 15.6.4.1 Receive Parser

The receive parser parses the incoming frame data and generates filer properties and frame control block (FCB). The receive parser composes of the Ethernet header parser and L3/L4 parser.

The Ethernet header parser parses only L2 (ethertype) headers. It is enabled by RCTRL[PRSDEP] != 0. It has the following key features:

- Extraction of 48-bit MAC destination and source addresses
- Extraction and recognition of the first 2-byte ethertype field
- Extraction and recognition of the final 2-byte ethertype field
- Extraction of 2-byte VLAN control field
- Walk through MPLS stack and find layer 3 protocol
- Walk through VLAN stack and find layer 3 protocol
- Recognition of the following ethertypes for inner layer parsing
  - LLC and SNAP header

- JUMBO and SNAP header
- IPV4
- IPV6
- VLAN
- MPLSU/MPLSM
- PPOES
- ARP

For stack L2 (that is, more than one ethertypes) header, the Ethernet parser traverses through the header until it finds the last valid ethertype or the ethertype is unsupported. Below is a description of what the Ethernet header parser recognizes for stack L2 header.

**Table 15-154. Supported Stack L2 Ethernet Headers**

Column—Current L2 Ethertype Row—Next Supported L2 Ethertype	LLC/ SNAP	JUMBO/ SNAP	IPV4	IPV6	VLAN	MPLSU	MPLSM	PPOES	ARP
LLC/SNAP	N	N	Y	Y	Y	Y	Y	Y	Y
JUMBO/SNAP	N	N	Y	Y	Y	Y	Y	Y	Y
IPV4	N	N	N	N	N	N	N	N	N
IPV6	N	N	N	N	N	N	N	N	N
VLAN	Y	Y	Y	Y	Y	Y	Y	Y	Y
MPLSU	N	N	Y*	Y*	N	y	Y	N	N
MPLSM	N	N	Y*	Y*	N	Y	Y	N	N
PPOES	N	N	Y	Y	N	Y	Y	N	N
ARP	N	N	N	N	N	N	N	N	N

**Note:** \* means that it is the next protocol


The L3 parser is enabled by RCTRL[PRSDEP] = 10 or 11. It begins when the Ethernet parser ends and a valid IPv4/v6 ethertype is found. The L4 header is enabled by RCTRL[PRSDEP] = 11. It begins when the L3 parser ends and a valid TCP/UDP next protocol is found and no fragment frame is found. The primary functionalities of L3(IPv4/6) and L4(TCP/UDP) parsers are as follows:

- IP recognition (v4/v6, ARP, encapsulated protocol)
- IP header checksum verification
- IPv4/6 over IPv4/6 (tunneling)—parse headers and find layer 4 protocol
- IP layer 4 protocol/next header extraction

- Stop parsing on unrecognized next header/protocol
- IPv4 support
  - IPv4 source and destination addresses
  - 8-bit IPv4 type of service
  - IP layer 4 protocol / next header support
    - IPV4
    - IPV4 Fragment. Parser stops after a fragment is found
    - TCP/UDP
- IPv6 support
  - The first 4 bytes of the IPv6 source address extraction
  - The first 4 bytes of the IPv6 destination address extraction
  - IPv6 source address hash for pseudo header calculation
  - IPv6 destination address hash for pseudo header calculation
  - 8-bit IPv6 traffic class field extraction
  - Payload length field extraction
  - IP layer 4 protocol/next header support
    - IPV6
    - IPV6 fragment. Parser stops after a fragment is found
    - IPV6 route
    - IPV6 hop/destination
    - TCP/UDP
- L4 (TCP/UDP) support
  - Extraction of 16-bit source port number extraction
  - Extraction of 16-bit destination port number extraction
  - TCP checksum calculation (including pseudo header)
  - UDP checksum calculation if the checksum field is not zero (including pseudo header)

#### 15.6.4.2 Receive Queue Filer

The receive queue filer receives protocol header properties extracted from the incoming frame by the eTSEC frame parse engine. A property is defined to be a field extracted from a packet header, such as a TCP port number or VLAN identifier. As soon as the last identifiable header has been recognized, the filer commences searching the receive queue filer table, comparing properties in the table against properties extracted from the frame. This table is illustrated in [Figure 15-138](#). Software populates the table with property values, stored to the RQPROP field, and indicates how to match and interpret the properties by setting flags in the RQCTRL field. The eTSEC memory map provides access to these fields by way of an address register (RQFAR) and two porthole registers (RQFCR and RQFPR).



**Figure 15-138. Structure of the Receive Queue Filer Table**

### 15.6.4.2.1 Filing Rules

Unless the filer is disabled, every received frame from the Ethernet MAC initiates a search of the receive queue filer table, starting at entry 0. The table search is terminated as soon as an entry is found whose contents match a property of the frame. Accordingly, software must guarantee that at least one entry results in a match—even if only to set a default receive queue index.

Since eTSEC searches the table at a rate of two entries every system clock cycle, all 256 entries can be searched in the time taken to receive a 64-byte Ethernet frame.

Each entry of the receive queue filer table specifies a simple match rule for determining how to process the received frame. The elements of a filing rule, expressed in the RQCTRL and RQPROP fields, are summarized as follows:

- The PID field in RQCTRL identifies what property is being matched against RQPROP. The eTSEC supports 16 properties, some of which are different portions of the same header field. Reserved or unused bits in RQPROP are read as zero. See [Section 15.5.3.3.8, “Receive Queue Filer Table Property Register \(RQFPR\),” on page 15-57](#) for a list of all properties and their associated PID values.
- The Q field in RQCTRL identifies which one of 64 virtual receive queues the frame should be filed to (sent through DMA) in the event of a filing rule match that accepts the frame. The physical RxBD ring this queue maps to is controlled by the RCTRL[FSQEN] bit. If RCTRL[FSQEN] = 0, the three least significant bits of the Q field indicate which physical RxBD ring hosts the queue. If RCTRL[FSQEN] = 1, RxBD ring 0 hosts all receive queues, but the RxFCB[RQ] field allows software to distinguish queues by ID. In all cases if Q maps to a RxBD ring that is not currently enabled, the frame is discarded with an IEVENT[FIQ] error.
- The REJ field in RQCTRL controls whether the frame is to be rejected (REJ = 1) or filed (REJ = 0) upon a filing rule match. Rejected frames occupy Rx FIFO space, but do not consume memory bus cycles.
- The CMP field in RQCTRL determines how property PID is compared against RQPROP. Equality, inequality, greater-or-equal, and less-than compares are available.

- The AND field in RQCTRL allows more than one comparison in a sequence to be chained together as a Boolean AND condition. Setting AND = 1 defers evaluation of the rule until the next entry has been matched, which may, in turn, have AND set. If any comparison involving AND = 1 fails, the entire chained sequence fails. A typical use for AND is to combine a pair of comparisons in a range match; the first such entry has AND = 1, the second has AND = 0 and its values of Q and REJ take effect.
- The CLE field in RQCTRL offers a way to bracket a set of consecutive—perhaps related—rules into a rule cluster. A cluster must be preceded by a guard rule, which simply determines whether the cluster rules can be evaluated. If the guard rule succeeds and its last entry has both CLE = 1 and AND = 1, the cluster rules that follow are enabled. The cluster ends at the first entry where CLE = 1 and AND = 0, which may also belong to a rule that files or rejects a frame. If the guard rule fails, all rules in the cluster are skipped, including mask\_register assignments. Clusters must not be nested.
- The GPI field offers the user the ability to interrupt the core upon matching a rule that causes a frame to be filed to memory. Once the last RxBD corresponding to that frame is written to memory, the IEVENT[FGPI] event will be asserted. This bit will be set regardless of any interrupt coalescing that may be set.

#### 15.6.4.2.2 Comparing Properties with Bit Masks

By default, extracted properties are compared arithmetically according to the CMP field in each RQCTRL word. This permits point value matches in each table entry, and range checks across a pair of table entries combined with the AND attribute in RQCTRL. However, inspection of the parse flags, Ethernet preamble, and IP addresses typically requires “don’t care” bit fields in the properties to be cleared as part of the comparison. The eTSEC provides a dedicated 32-bit register, known as the mask\_register, for performing such masking operations. At the start of each table search by the filer, mask\_register is reset to 0xFFFF\_FFFF, which ensures that no masking occurs.

Filer rules may be configured to assign specific bit patterns to mask\_register. Such rules can be configured to either match always (useful for implementing a default rule and specifying an associated receive queue), or fail always (which prevents termination of the filer table search). Once mask\_register has been assigned, it retains its value until it is reassigned or the table search terminates. All properties are non-destructively bit-wise ANDed with mask\_register prior to comparison in subsequent rules, which allows an entire cluster of rules to make use of a common mask. Individual masks for specific rules can also be created simply by combining a mask\_register assignment (match always form) with a regular rule using the AND attribute.

To create a mask\_register assignment rule, it is necessary to select PID = 0 in RQCTRL, and choose CMP such that the rule either matches (CMP = 01) or fails (CMP = 11). In this entry, RQPROP is then considered to be the assigned bit vector.

### 15.6.4.2.3 Special-Case Rules

It is frequently useful to create rules that are guaranteed to succeed or fail, specifically to enforce a default filing decision or act as null entries. Suggested constructions for such rules are shown in [Table 15-155](#).

**Table 15-155. Special Filer Rules**

Rule Description	RQCTRL Fields						RQPROP Word	RQCTRL Word <sup>1</sup>
	CLE	REJ	AND	Q	CMP	PID		
Default file—Always file frame to ring Q	0	0	0	Q	01	0000	0x0000_0000	0x0000_0q20
Default reject—Always discard frame	0	1	0	000_000	01	0000	0x0000_0000	0x0000_0120
Empty rule in AND—Always matches	0/1 <sup>2</sup>	0	1	000_000	01	0000	0xFFFF_FFFF	0x0000_00A0
Empty rule in rule set—Always fails	0/1 <sup>3</sup>	0	0	000_000	11	0000	0xFFFF_FFFF	0x0000_0060

<sup>1</sup> Hexadecimal digits *qq* denotes field Q shifted left 2 bits.

<sup>2</sup> Set CLE = 1 if the empty rule guards a cluster.

<sup>3</sup> Set CLE = 1 if the empty rule occurs at the end of a cluster.

### 15.6.4.2.4 Filer Interrupt Events

The filer can produce three interrupt events in IEVENT. Event FIR indicates an error condition where the filer was unable to provide a definite result, either because no rule in the table succeeded, or because frames arrived too rapidly to complete searching of the table. Event FIQ indicates that the filer accepted a frame to a RxB ring that was not enabled in RQCTRL (this can also occur if the filer is disabled, but RxB ring 0—default queue or FSQEN mode queue—is not enabled). FIQ is also asserted in the case where no rule in the entire table succeeded. The various combinations of these interrupt events and their interpretation appear in [Table 15-156](#).

**Table 15-156. Receive Queue Filer Interrupt Events**

IEVENT[FIR]	IEVENT[FIQ]	Description
0	0	No error. The filer successfully rejected or filed a frame.
0	1	Illegal queue error. The filer accepted a frame to a RxB ring that is disabled (including ring 0 if filing is disabled).
1	0	Partial search error. The filer did not have sufficient time to complete its search of the filer table.
1	1	No matching rule error. The filer searched all 256 entries of the filer table without finding a rule that succeeds.

A functional interrupt is provided via use of the general purpose interrupt (GPI) bit in the filer table. When a property matches the value in the RQPROP entry at this index, and REJ = 0 and AND = 0, the filer will set IEVENT[FGPI] when the corresponding receive frame is written to memory. This allows the user to set up a filer rule where the core will be interrupted upon the reception of ‘special’ frames.

If the timer is enabled (TMR\_CTRL[TE] = 1), then the interrupt dedicated for timer events (in addition to the usual receive, transmit and error interrupts) will be asserted.



### 15.6.4.2.5 Setting Up the Receive Queue Filer Table

The eTSEC frame parser always provides values for all properties, even where the relevant headers are not available. In the latter case, the filer is given default properties that can be used to avoid conflict with normal, defined property values. Accordingly, the rules in the filer table can be partitioned into rule sets such that if all rules in a given set fail (due to headers being unavailable), lower priority rule sets can be subsequently searched until either a rule set provides a match or a single default—catch-all—rule specifies a definite receive queue. For example, an IEEE 802.1p priority rule set may be followed by an IP TOS rule set, followed by a default rule; thus, if no VLAN tag appears in the received frame, the TOS rules are checked, or the default is activated should no IP header be present.

The rule cluster feature is used to conditionalize evaluation of rule sets. Typically, this avoids evaluating rules based on properties that may not be valid or relevant to the filing or filtering decision. For example, TCP-related rules might be clustered behind a guard rule that checks that a TCP header has appeared and the IP address matches our home address. Property 1—the parse flags property—is provided specifically to check the characteristics of the received frame and the parser error status. The mask\_register is typically assigned beforehand to extract specific flags, in which case care should be taken that mask\_register be reassigned an appropriate mask vector for following comparisons.

In many cases it is possible to write the entire filer table before using eTSEC, as the rule set is static. However, dynamic rule updates can be supported by pre-allocating partially instantiated rule sets, which software rewrites as necessary. Rules that are not instantiated should be composed of empty entries, as indicated in [Table 15-155](#). In many cases empty entries can be overwritten by software without stopping eTSEC's receive function.

### 15.6.4.2.6 Filer Example—802.1p Priority Filing

This example, shown in [Table 15-157](#), illustrates how to file frames according to layer 2 802.1p priority. This matches against property 1001, comparing each specific priority level in order to associate them with a RxBD ring index. Note that if a VLAN tag does not appear in the frame, the parser passes priority 0 to the filer, which always matches the rule at entry 7 and terminate the table search.

**Table 15-157. Filer Table Example—802.1p Priority Filing**

Table Entry	RQCTRL Fields						RQPROP	Comment	RQCTRL Word
	CLE	REJ	AND	Q	CMP	PID			
0	0	0	0	000_000	00	1001	0x0000_0007	File priority 7 to ring 0	0x0000_0009
1	0	0	0	000_001	00	1001	0x0000_0006	File priority 6 to ring 1	0x0000_0409
2	0	0	0	000_010	00	1001	0x0000_0005	File priority 5 to ring 2	0x0000_0809
3	0	0	0	000_011	00	1001	0x0000_0004	File priority 4 to ring 3	0x0000_0C09
4	0	0	0	000_100	00	1001	0x0000_0003	File priority 3 to ring 4	0x0000_1009
5	0	0	0	000_101	00	1001	0x0000_0002	File priority 2 to ring 5	0x0000_1409

**Table 15-157. Filer Table Example—802.1p Priority Filing (continued)**

Table Entry	RQCTRL Fields						RQPROP	Comment	RQCTRL Word
	CLE	REJ	AND	Q	CMP	PID			
6	0	0	0	000_110	00	1001	0x0000_0001	File priority 1 to ring 6	0x0000_1809
7	0	0	0	000_111	00	1001	0x0000_0000	File undefined 802.1p or priority 0 to ring 7—Default always matches	0x0000_1C09

#### 15.6.4.2.7 Filer Example—IP Diff-Serv Code Points Filing

This example demonstrates use of rule priority for determining class selector codepoints (RFC 2474) from the IP TOS property. An example filer table is shown in [Table 15-158](#). The example relies on the fact that the first rule matched terminates the search, hence successively lower Diff-Serv codepoint ranges can be compared in each step until the default (zero or greater) range is reached. By default, property 1010 (IP TOS) takes the value 0x00 if no IP headers were recognized, therefore the table search always terminates.

**Table 15-158. Filer Table Example—IP Diff-Serv Code Points Filing**

Table Entry	RQCTRL Fields						RQPROP	Comment	RQCTRL Word
	CLE	REJ	AND	Q	CMP	PID			
0	0	0	0	001_000	01	1010	0x0000_00E0	File class 7 to queue 8 (TOS >= 0xE0)	0x0000_202A
1	0	0	0	001_001	01	1010	0x0000_00C0	File class 6 to queue 9 (TOS >= 0xC0)	0x0000_242A
2	0	0	0	001_010	01	1010	0x0000_00A0	File class 5 to queue 10 (TOS >= 0xA0)	0x0000_282A
3	0	0	0	001_011	01	1010	0x0000_0080	File class 4 to queue 11 (TOS >= 0x80)	0x0000_2C2A
4	0	0	0	000_100	01	1010	0x0000_0060	File class 3 to queue 4 (TOS >= 0x60)	0x0000_102A
5	0	0	0	001_100	01	1010	0x0000_0040	File class 2 to queue 12 (TOS >= 0x40)	0x0000_302A
6	0	0	0	010_100	01	1010	0x0000_0020	File class 1 to queue 20 (TOS >= 0x20)	0x0000_502A
7	0	0	0	011_100	01	1010	0x0000_0000	File class 0 to queue 28 (TOS >= 0x00) or file to ring 4 by default	0x0000_702A

#### 15.6.4.2.8 Filer Example—TCP and UDP Port Filing

This example demonstrates rule clusters and AND-combined entries for filing packets based on transport protocol and well-known port numbers in a termination application. An example filer table is shown in [Table 15-159](#). The example contains two clusters; the first is entered only for TCP packets, the second is entered only for UDP packets. A default filing rule catches the case where neither TCP nor UDP headers are found. Each cluster compares source port number (property 1111) against a list of server ports, and files the packets accordingly. Note that entries 1 and 2 form an AND rule for checking that the port number >= 20 and port number < 22. Entries 4 and 5 are initially set up to always fail (zero port number), and thus comprise empty entries that can be used at a later time.

Table 15-159. Filer Table Example—TCP and UDP Port Filing

Table Entry	RQCTRL Fields						RQPROP	Comment	RQCTRL Word
	CLE	REJ	AND	Q	CMP	PID			
0	1	0	1	000_000	00	1011	0x0000_0006	Enter cluster if layer 4 is TCP	0x0000_028B
1	0	0	1	000_000	01	1111	0x0000_0014	AND rule—FTP from TCP ports 20 and 21: file to ring 2	0x0000_00AF
2	0	0	0	000_010	11	1111	0x0000_0016		0x0000_086F
3	0	0	0	000_011	00	1111	0x0000_0017	telnet from TCP port 23: file to ring 3	0x0000_0C0F
4	0	0	0	000_000	00	1111	0x0000_0000	<i>empty entry reserved for future use</i>	0x0000_000F
5	0	0	0	000_000	00	1111	0x0000_0000	<i>empty entry reserved for future use</i>	0x0000_000F
6	1	0	0	000_001	01	0000	0x0000_0000	end cluster; default TCP: file to ring 1	0x0000_0620
7	1	0	1	000_000	00	1011	0x0000_0011	Enter cluster if layer 4 is UDP	0x0000_028B
8	0	0	0	000_101	00	1111	0x0000_0801	NFS from UDP port 2049	0x0000_140F
9	0	0	0	000_111	00	1111	0x0000_0208	Route from UDP port 520	0x0000_000F
10	0	0	0	000_110	00	1111	0x0000_0045	TFTP from UDP port 69	0x0000_180F
11	1	0	0	000_100	01	0000	0x0000_0000	End cluster; default UDP: file to ring 4	0x0000_1220
12	0	0	0	000_000	01	0000	0x0000_0000	By default, file to ring 0	0x0000_0020

### 15.6.4.3 Transmission Scheduling

Each eTSEC can maintain multiple TxBD rings (or transmission queues) to satisfy QoS requirements. The ability to choose from a number of transmission streams dynamically is especially important during periods of network congestion. Certain application such as voice and video streaming are delay sensitive, but loss insensitive. For instance, VoIP applications require little bandwidth, but are highly sensitive to latency. Conversely, FTP or SMTP protocols are delay insensitive, but loss sensitive.

eTSEC has a transmission scheduler that implements a programmable QoS regime. The scheduler is responsible for choosing which of the prefetched TxBDs shall be processed next, and accordingly issuing DMA requests to service the data stream described by the chosen BD(s). The scheduler cycle is one of:

1. decide on a TxBD queue,
2. transmit exactly one frame from that queue, and
3. return to deciding on another queue, in step 1.

If TCTRL[TXSCHEDED] is set to 00, no transmission scheduling occurs, and only TxBD ring 0 is polled for new data to transmit, with DMACTRL controlling waiting or polling. TCTRL[TXSCHEDED], if not zero, can be programmed to invoke one of two scheduling algorithms, namely priority-based queuing (PBQ), and modified weighted round-robin queuing (MWRR). In all cases where TCTRL[TXSCHEDED] is not zero, the scheduler can choose from among 1 to 8 TxBD rings per eTSEC, with individual rings being enabled by the setting of TQUEUE[EN0–EN7] bits. For example, TxBD rings 3, 4, and 7 may be enabled for scheduling by setting EN3, EN4, and EN7, and clearing all other EN bits.

### 15.6.4.3.1 Priority-Based Queuing (PBQ)

PBQ is the simplest scheduler decision policy. The enabled TxBD rings are assigned a priority value based on their index. Rings with a lower index have precedence over rings with higher indices. For example, TxBD ring 0 has higher priority than TxBD ring 1, and TxBD ring 1 has higher priority than TxBD ring 2, and so on.

The scheduling decision is then achieved as follows:

```

loop
  priority_ring = null;
  ring = 0;
  while priority_ring == null and ring <= 7 loop
    if enabled(ring) and not ring_empty(ring) then
      priority_ring = ring;
    endif
    ring = ring + 1;
  endloop
  if priority_ring >= 0 then
    while not ring_empty(priority_ring) loop
      transmit_frame(priority_ring);
    endloop
  endif
endloop

```

In practice a protocol stack or device driver can abuse PBQ by attempting to queue too much traffic onto high priority rings. It is recommended that the highest priority ring should normally not be used at all except for frames requiring the utmost urgent transmission. This allows emergency traffic to overtake backlogged queues out of sequence.

### 15.6.4.3.2 Modified Weighted Round-Robin Queuing (MWRR)

eTSEC implements a modified weighted round-robin scheduling algorithm across all enabled TxBD rings when TCTRL[TXSCHEM] = 10. In MWRR, the weights in the TR03WT and TR47WT registers determine the ideal size of each transmit slot, as measured in multiples of 64 bytes. Thus, to set a transmit slot of 512 bytes, a weight of 512/64 or 8 needs to be set for the ring. In this mode TxBD rings 1–7 are selected in round-robin fashion, whereas TxBD ring 0, if enabled with ready data for transmission, is always selected in between other rings so as to expedite transmission from ring 0.

The scheduling decision is then achieved as follows:

```

for ring = 1..7 and enabled(ring) loop
  credit[ring] = 0;
endloop
for ring = 1..7 and enabled(ring) loop
  if not ring_empty(0) then
    credit[0] = credit[0] + weight[0];
    while credit[0] > 0 loop
      transmit_frame(0);
      credit[0] = credit[0] - frame_size;
      if ring_empty(0) then
        credit[0] = 0;
      endif
    endloop
  endif
endloop
endif

```

```

if not ring_empty(ring) then
    credit[ring] = credit[ring] + weight[ring];
endif
while credit[ring] > 0 loop
    transmit_frame(ring);
    credit[ring] = credit[ring] - frame_size;
    if ring_empty(ring) then
        credit[ring] = 0;
    endif
endloop
endloop

```

The algorithm checks registers TQUEUE[EN0–EN7] for `enabled()`, TSTAT[THLT0–THLT7] for `ring_empty()`, and TRxWT for `weight()`. For TxBD ring  $k$ , having a weight  $WT_k$ , the long term average throughput for that ring is:

$$\text{rate of queue}[k] \text{ (} K = 1 \text{ to } 7) = (\text{available bandwidth}) * WT_k / (\text{sum}(WT_i) + 6WT_0)$$

$$\text{rate of queue}(0) = (\text{available bandwidth}) * 7 * WT_0 / (\text{sum}(WT_i) + 6WT_0)$$

where  $i = 0$  to 7

## 15.6.5 Lossless Flow Control

The eTSEC DMA subsystem is designed to be able to support simultaneous receive and transmit traffic at gigabit line rates. If the host memory has sufficient bandwidth to support such line rates, then the principle cause of overflow on receive traffic is due to a lack of Rx BDs. Thus, the long term receive throughput is determined by the rate at which software can process receive traffic. If a user desires to prevent dropped packets, they can inform the far-end link to stop transmission while the software processing catches up with the backlog.

To avoid overflow in the latter case, back pressure must be applied to the far-end transmitter before the Rx descriptor controller encounters a non-empty BD and halts with a BSY error. As there is lag between application of back-pressure and response of the far-end, the pause request must be issued while there are still BDs free in the ring. In the traditional eTSEC descriptor ring programming model, there is no way for hardware to know how many free BDs are available, so software must initiate any pause requests required during operation. If software is backlogged, the request may not be issued in time to prevent BSY errors. To allow the eTSEC to generate the pause request automatically, additional information (a pointer to the last free BD and ring length) is required.

### 15.6.5.1 Back Pressure Determination through Free Buffers

Ultimately, the rate of data reception is determined by how quickly software can release buffers back into the receive ring(s). Each time a buffer is freed, the associated BD has its empty bit set and hardware is free to consume both. Thus the number of free BDs in a given Rx ring indicates how close hardware is to the end of that ring. To prevent data loss, back pressure should be applied when the number of free BDs drops below some critical level. The number of BDs that can be consumed by an incoming packet stream while back-pressure takes effect is determined by several factors, such as: receive traffic profile, transmit traffic profile, Rx buffer size, physical transmission time between eTSEC and far-end device and intra-device latency. Theoretically, the worst case is as follows:

$$\text{FreeBDsRequired} = \frac{\text{MaxFrameSize}}{\text{MinFrameSize} + \text{IFG}} + \frac{\text{MaxFrameSize}}{\text{RxBufferSize}} + \text{LinkDelay}$$

This case comes about when:

- The eTSEC has just started transmitting a large frame and thus cannot send out a pause frame
- Upon reception of the pause request the far-end has just started transmission of a large frame
- The eTSEC receives a burst of short frames with minimum inter-frame-gap (96bit times for ethernet)

Once the user has determined the worst case scenario for their application, they program the required free BD threshold into the eTSEC (through RQPRM[PBTHR]). Since different BD rings may have different sizes and expected packet arrival rates, a separate threshold is provided for each active ring. It is recommended that a threshold of at least fourBDs is the practical minimum for gigabit ethernet links.

For the Rx descriptor controller to determine the number of free BDs remaining in the ring, it needs to know the following:

1. The location of the current BD being used by hardware
2. The location of the last BD that was released (freed) by software
3. The length of the Rx BD ring.

For each active ring, the current BD pointer (RBPTR<sub>n</sub>) is maintained by the eTSEC. Software knows both the size of the Rx ring and the location of the last freed BD. By providing the eTSEC with those values (through RQPRM[LEN] and RFBPTR respectively) the eTSEC always know how many receive buffers are available to be consumed by incoming data.

The number of guaranteed free BDs in the ring is then determined by:

When RFBPTR<sub>n</sub> < RBPTR<sub>n</sub>

$$\text{FreeBDs} = \text{RQPRM}_n[\text{LEN}] - \text{RBPTR}_n + \text{RFBPTR}_n$$

When RFBPTR<sub>n</sub> > RBPTR<sub>n</sub>

$$\text{FreeBDs} = \text{RFBPTR}_n - \text{RBPTR}_n$$

When RBPTR<sub>n</sub> = RFBPTR<sub>n</sub> the number of free BDs in the ring is either one (since RFBPTR<sub>n</sub> points to a free BD) or equal to the ring length. Since the BD pointed to by RBPTR<sub>n</sub> may be either in use or about to be used, it is not considered in the free BD count. To resolve the case where the two pointers collide, the following logic applies:

If RBASE<sub>n</sub> was updated and thus initializes both RBPTR<sub>n</sub> and RFBPTR<sub>n</sub>, the ring is deemed empty.

If RFBPTR<sub>n</sub> is updated by a software write and matches RBPTR<sub>n</sub>, the ring is deemed empty.

If HW updates RBPTR<sub>n</sub> and the result matches RFBPTR<sub>n</sub>, the ring is deemed to have one BD remaining. Upon writing this BD back to memory (indicating the buffer is occupied) the ring is deemed to be full.

**Important.** There is a possibility that if software is severely backlogged in updating  $RFBPTR_n$ , the hardware could wrap around the ring entirely, consume exactly the remaining number of BDs and not halt with a BSY error. If software then increments  $RFBPTR_n$  to the next address (thereby equalling  $RBPTR_n$ ), the hardware assumes the ring is now empty (when in fact there is only a single BD freed up). This results in the hardware failing to maintain back pressure on the far end. Upon software incrementing  $RFBPTR_n$  a subsequent time, the wrap condition is successfully detected and hardware recognizes a nearly full ring (rather than a nearly empty one). Since software can increment  $RFBPTR_n$  by any amount, it is not possible for hardware to determine in this case whether the user has cleared the entire ring or just one BD. Users can eliminate the possibility of this condition occurring by ensuring that  $RFBPTR_n$  is incremented by at least two BDs each time (that is, clear at least two buffers whenever the RxBD unload routine is called).

Once the eTSEC determines that this threshold has been reached, back pressure is applied accordingly. The type of back pressure that is applied varies according to the physical interface that is used.

- **Half duplex Ethernet:** No support in this mode.
- **Full duplex Ethernet:** An IEEE 802.3 PAUSE frame (see sect. 15.6.2.9/15-154) is issued as if the TCTRL[TFC\_PAUSE] bit was set. An internal counter tracks the time the far end controller is expected to remain in pause (based on the setting of PTV[PT]). When that counter reaches half the value of PTV[PT], the eTSEC reissues a pause frame if the free BD calculation for any ring is below the threshold for that ring. For example, if PTV[PT] is set to 10 quanta, a pause frame is re-issued when five quanta have elapsed if the free BD threshold is still not met. A practical minimum for PTV[PT] of 4 quanta is recommended.
- **FIFO packet interface:** Link layer flow control is asserted through use of the RFC signal (CRS pin). Flow control is asserted for the entire time that free BD threshold is not met. The same mechanism is used for both GMII-style and encoded packet modes.

## 15.6.5.2 Software Use of Hardware-Initiated Back Pressure

### 15.6.5.2.1 Initialization

Software configures  $RBASE_n$  and  $RQPRM_n[LEN]$  according to the parameters for that ring. Then the number of free BDs that are required to prevent the eTSEC from automatically asserting flow control are programmed in  $RQPRM_n[FBTHR]$ . The receiver is then enabled.

Note: the act of programming  $RBASE_n$  initializes  $RFBPTR_n$  to the start of the of the ring. When the ring is in this initial empty state, there is no concept of a last freed BD. In this case, the calculated number of free BDs is the size of the ring. Since the BD that the hardware is currently pointing to is to be considered in-use, the free BD count is actually one higher than the total available. As soon as the hardware consumes a BD (by writing it back to memory),  $RBPTR_n$  advances and the free BD count reflects the correct number of available free BDs.

### 15.6.5.2.2 Operation

As software frees BDs from the ring, it writes the physical address of the BD just freed to  $RFBPTR_n$ . The eTSEC asserts flow control if the distance (using modulo arithmetic) between  $RBPTR_n$  and  $RFBPTR_n$  is  $< RQPRM_n[FBTHR]$ . In multi-ring operation, if the free BD count of **any** active ring drops below the

threshold for that ring, flow control is asserted. Once enough BDs are freed for **all** active rings to meet their respective free BD thresholds, application of back pressure cases.

Note: The eTSEC does not issue an exit pause frame (that is, pause frame with PTV of 0x0000) once all active rings have sufficient BDs. Instead, it waits for the far-end pause timer to expire and start re-transmission.

## 15.6.6 Hardware Assist for IEEE Std. 1588-Compatible Timestamping

There is a push in industrial control applications to use Ethernet as the principal link layer for communications. This requires Ethernet to be used for both data transfer and real-time control. For real-time systems, each node is required to be synchronized to a master clock. The precision of this clock is dictated by the application, but generally needs to be of the order of <1uSec for high-speed machinery (for example, printing presses).

IEEE 1588 [1588] specifies a mechanism for synchronizing multiple nodes to a master clock. Support for 1588 can be done entirely in software running on a host CPU, but applications that require sub 10uSec accuracy will need hardware support for accurate timestamping of incoming packets.

The eTSEC includes a new timer clock module to support the IEEE Std. 1588 timer standard. The following sections describe the features, programming model, and implementation information.

### 15.6.6.1 Features

- 64-bit free running timer running from an external oscillator or internal clock
- Programmable timer oscillator clock selection
- Self-correcting precision timer with nano-second resolution
- Time stamp all incoming packets inline
  - Maskable interrupts on received PTP packet's filter rule match
- Time stamp transmit packets when instructed in the TxFCB
  - Maskable interrupts on transmit timestamp capture
- Two Tx timestamp registers per eTSEC with 16-bit tag for each of them to support burst mode.
- Time stamp capture on two general-purpose external triggers
  - Maskable interrupts on GPIO timestamp trigger
  - Programmable polarity of external trigger (GPIO) edge
- Two 64-bit alarm (future time) registers for future time comparison
  - Maskable interrupts on alarm
- Three programmable timer output pulse period phase aligned with 1588 timer clock
  - Maskable interrupts associated with each pulse
- Separate maskable timer interrupt event register



- Recognition of incoming PTP packet through filter rule match
- Phase aligned adjustable (divide by N) clock output
- Supports all Ethernet modes supported by the eTSEC, including full- and half-duplex modes
- Supports both master and slave modes
- Supports timestamp of nano-second resolution

### 15.6.6.2 Timer Logic Overview

The 1588 timer module can be partitioned into four different sub-modules as shown in [Figure 15-139](#).




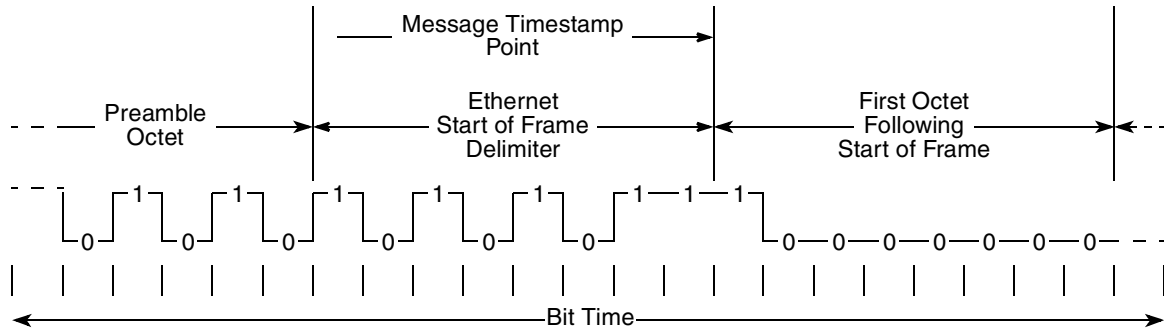
Figure 15-139. 1588 Timer Design Partition

### 15.6.6.3 Timestamp Insertion on the Received Packets

Every incoming packet's 8-byte timestamp is inserted into the packet data buffer as padding alignment bytes. Timestamp insertion into the data buffer requires `RCTRL[PAL]` to be set to a value greater than or equal to 8 and the control bit `RCTRL[TS]` bit to be set.

#### 15.6.6.3.1 Timestamp Point

The required timestamp point, as specified in the IEEE 1588 Specification Sep-2004 (IEC 61588 First Edition), is shown in [Figure 15-140](#). From this, it is clear that the end of the SFD is the critical point in the MII data stream.



**Figure 15-140. Ethernet Sampling Points for 1588**

The sample point coincides with the cycle after the SFD (Start of Frame Delimiter) detection by the MAC. For received frames, this will be at least 4 bit times (MII) or 8 bit times (GMII) after the message timestamp point specified in [1588]. For transmission, the eTSEC sample point precedes the sample point specified in [1588] by at least 4-bit times (MII) or 8-bit times (GMII). For a particular mode, the eTSEC sample point is a consistent number of bit times relative to the SFD detection. Thus, the offset from the [1558] specified sample point can be accounted for in the PTP software implementation.

#### 15.6.6.4 PTP Packet Parsing

PTP packets are typically embedded within a UDP payload with special IP source and destination address and special source and destination ports numbers. Special fields of interest of a PTP packet are listed in [Table 15-160](#).

**Table 15-160. PTP Payload Special Fields**

Layer	Octet (Offset from the SFD)	Field	Value	eTSEC filter PID	Comments
Ethernet	12-13	Length/Packet	0x0800	ETY-RQPFR[P ID=0111]	IPv4
IP header	22	Time to live	0x00	RBIFX-choose an arbitrary extraction byte	Must be 0
IP header	23	IP Protocol	0x11	L4P-RQPFR[P ID=1011]	UDP
IP header	26-29	Source IP Address IANA defines 4 multicast address for the PTP packet		SIA-RQPFR[PI D=1101]	

Table 15-160. PTP Payload Special Fields (continued)

Layer	Octet (Offset from the SFD)	Field	Value	eTSEC filer PID	Comments
IP header	30-33	Destination IP Address IANA defines 4 multicast address for the PTP packet	224.0.1.129 224.0.1.130 224.0.1.131 224.0.1.132	DIA-RQPFR[PID=1100]	DefaultPTPdomain AlternatePTPdomain1 AlternatePTPdomain2 AlternatePTPdomain3
UDP header	34-35	Source port number		SPT-RQPFR[PID=1011]	
UDP header	36-37	Destination port number	319 320	DPT-RQPFR[PID=1011]	EventPort GeneralPort
UDP data	74	Control	0x0 0x1 0x2 0x3 0x4	RBIFX-choose an arbitrary extraction byte	Sync Delay_req Follow_up Delay_resp Management

A representation of the PTP packet is shown in Figure 15-141.




Figure 15-141. PTP Packet Format

#### 15.6.6.4.1 General Purpose Filer Rule

The eTSEC receive filer has been enhanced with the addition of a general-purpose event bit. This event bit can be used in conjunction with filing table rules to identify 1588 packets and indicate these packets by setting special timer status register bits (TMR\_STAT). Additionally, 1588 packets can be easily identified by upper-layer software by using the filer to queue all PTP packets to one or more predefined virtual queues. See Section 15.6.4.2.1, “Filing Rules,” for further information.

#### 15.6.6.5 Timestamp Insertion on Transmit Packets

Software has the option to write the timestamp of the transmitted frame to memory in the padding alignment bytes (PAL) located between the TxFCB and the frame data. It is required that a minimum of two TxBDs are used. The first points to the start of the 8-byte TxFCB. The second points to the start of frame data. In memory, the TxFCB, and at least the first 16 bytes of the TxPAL must be adjacent, i.e., located in contiguous memory locations, as depicted in Figure 15-142.

The first TxBD[TOE] bit is set. When the TMR\_CTRL[Record Timestamp In PAL Enable] and TxFCB[PTP] bits are set, the timestamp is written to memory location TxBD[Data Buffer Pointer]+16.

The second TxBD's Data Length must either contain the full frame length, or a value greater than the TxThreshold setting. Refer to [Table 15-161](#). When timestamps are inserted into the TxPAL, the TMR\_TXTSn\_H/L and TMR\_TXTSn\_ID registers still function normally.

### 15.6.6.5.1 Interrupts

The TxPAL is updated with a timestamp before closing the second TxBD. The TxBD[I] bit can be set for the second TxBD frame to cause an interrupt (via IEVENT[TXF]) after the timestamp has been written to the TxPAL.

When timestamps are inserted into the TxPAL, the TMR\_TXTSn\_H/L and TMR\_TXTSn\_ID registers still function normally. Therefore, the 1588 interrupt can be triggered by using the TMR\_PEVENT register bits TXP1, and TXP2.

**Table 15-161. Timestamp Insertion Programming Requirements**

Requirement	Behavior if requirement is not met
TMR_CTRL[RTPE]=1	If TMR_CTRL[RTPE]=0, then no timestamp is written to a TxPAL.
TxBD[TOE]=1	If TxBD[TOE]=0, then no timestamp is written to a TxPAL.
First TxBD[Data Buffer Pointer] is 8-byte aligned	The timestamp will be written to address First TxBD[Data Buffer Pointer] + 0x10 rounded down to the nearest 8-byte aligned address, except at 4K page boundaries, in which case the timestamp may be invalid, and the Second TxBD close status will be lost.
First TxBD[Data Length]=8, 8 bytes for TxFCB	If L2 or frame data is included in the Length, the buffer immediately following the FCB is transmitted on the line and the frame data stored in memory will be overwritten with a timestamp value after the frame is transmitted.
TxFCB[PTP]=1	If TxBD[PTP]=0, then no timestamp is written to a TxPAL.
The TxFCB is followed immediately by a minimum of 16 bytes for the TxPAL	The timestamp will be written to address First TxBD[Data Buffer Pointer] + 0x10.
Second TxBD[Data Buffer Pointers] points to start of L2 or frame data	If there is only one TxBD used to transfer a PTP frame, then no timestamp is written to a TxPAL.
Second TxBD[Data Length] >= FIFO_TX_THR or includes the entire frame	If this condition is not true, the timestamp in TxPAL is invalid.

Figure 15-142 depicts the buffer format requirements for timestamp insertion on transmit packets.

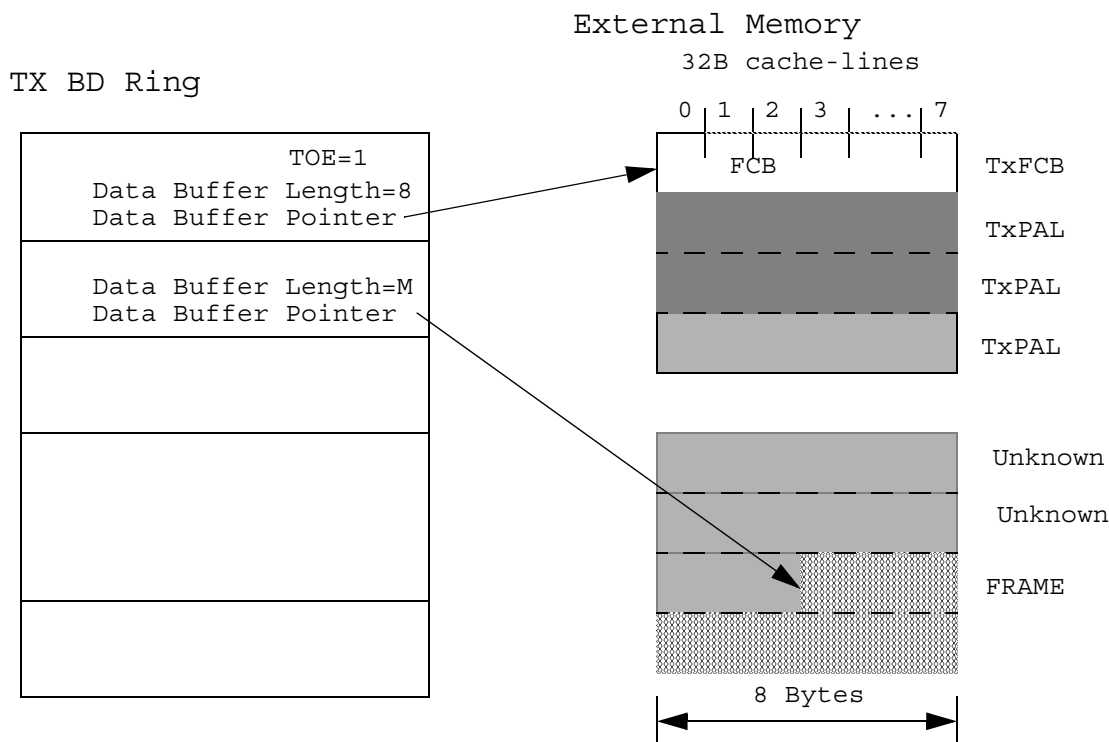


Figure 15-142. Buffer Format for Transmit timestamp Insertion

### 15.6.6.5.2 Error Condition

When an error is encountered after a PTP packet has begun to be processed, the timestamp written to the TxPAL is zero. Subsequent frames may be flushed by eTSEC. There will be no timestamp update to TxPAL for the subsequent flushed frames.

### 15.6.6.6 Tx PTP Packet Parsing

Software instructs the Tx packet to be timestamped via setting bit 15 in the TxFCB to mark a PTP packet. TxFCB[VLCTL] can be translated as the Tx PTP packet identification number. BD[TOE] must be set to enable transmit PTP packet timestamping. TxFCB[PTP] bit takes precedence over TxFCB[VLN] bit. It disables per packet VLAN tag insertion. On a PTP packet, a VLAN tag can be inserted from the DFVLAN register. The TxFCB for the PTP packet is shown in Figure 15-143.

	0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15
Offset + 0	VLN	IP	IP6	TUP	UDP	CIP	CTU	NPH								PTP
Offset + 2	L4OS							L3OS								
Offset + 4	PHCS															
Offset + 6	VLCTL/PTP_ID															

Figure 15-143. Transmit Frame Control Block

The contents of the Tx FCB are defined in [Table 15-162](#).

**Table 15-162. Tx Frame Control Block Description**

Bytes	Bits	Name	Description
0–1	0	VLN	VLAN control word valid. This bit is ignored when the PTP bit is set. VLAN tag is read from the DFVLAN register if PTP=1. 0 Ignore VLCTL field. 1 If VLAN tag insertion is enabled for eTSEC, use the VLCTL field as the VLAN control word.
	1	IP	Layer 3 header is an IP header. 0 Ignore layer 3 and higher headers. 1 Assume that the layer 3 header is an IPv4 or IPv6 header, and take L3OS field as valid.
	2	IP6	IP header is IP version 6. Valid only if IP = 1. 0 IP header version is 4. 1 IP header version is 6.
	3	TUP	Layer 4 header is a TCP or UDP header. 0 Do not process any layer 4 header. 1 Assume that the layer 4 header is either TCP or UDP (see UDP bit), and offload checksumming on the basis that the IP header has no extension headers.
	4	UDP	UDP protocol at layer 4. 0 Layer 4 protocol is either TCP (if TUP = 1) or undefined. 1 Layer 4 protocol is UDP if TUP = 1.
0–1	5	CIP	Checksum IP header enable. 0 Do not generate an IP header checksum. 1 Generate an IPv4 header checksum.
	6	CTU	Checksum TCP or UDP header enable. 0 Do not generate a TCP or UDP header checksum. RFC 768 advises that UDP packets not requiring checksum validation should have their checksum field set to zero. 1 Generate a TCP header checksum if IP = 1 and TUP = 1 and UDP = 0.
	7	NPH	Disable calculation of TCP or UDP pseudo-header checksum. This bit should be set if IP options need to be consulted in forming the pseudo-header checksum, as eTSEC does not examine IP options or extension headers for TCP/IP offload on transmit. 0 Calculate TCP or UDP pseudo-header checksum as normal, assuming that the IP header has no options. 1 Do not calculate a TCP or UDP pseudo-header checksum, but instead use the value in field PHCS when determining the overall TCP or UDP checksum.
	8–14	—	Reserved
	15	PTP	Indication to the transmitter that this is a PTP packet. Enabling PTP disables per packet VLAN tag insertion. Instead, VLAN tag will be read from the DFVLAN when the PTP field is true. 0 Do not attempt to capture transmission event time 1 Valid PTP_ID field. When this packet is transmitted, capture the time of transmission. Must be clear if TMR_CTRL[TE] is clear.
2–3	0–7	L4OS	Layer 4 header offset from start of layer 3 header. The layer 4 header starts L4OS octets after the layer 3 header if it is present. The maximum layer 3 header length supported is thus 255 bytes, which may prevent TCP/IP offload on particularly large IPv6 headers.
	8–15	L3OS	Layer 3 header offset from start of frame not including the 8 bytes for this FCB. The layer 3 header starts L3OS octets from the start of the frame including any custom preamble header that may be present. The maximum layer 2 header length supported is thus 255 bytes.

Table 15-162. Tx Frame Control Block Description (continued)

Bytes	Bits	Name	Description
4–5	0–15	PHCS	Pseudo-header checksum (16-bit one's complement sum with carry wraparound, but without result inversion) for TCP or UDP packets, calculated by software. Valid only if NPH = 1.
6–7	0–15	VLCTL/ PTP_ID	VLAN control word for insertion in the transmitted VLAN tag. Valid only if VLN = 1. Tx PTP packet identification number. This number will be copied into the Tx PTP packet timestamp identification field. PTP field takes precedence over VLN field.

## 15.6.7 Buffer Descriptors

The eTSEC buffer descriptor (BD) is modeled after the MPC8260 Fast Ethernet controller BD for ease of reuse across the PowerQUICC network processor family. Drawing from the MPC8260 FEC BD programming model, the eTSEC descriptor base registers point to the beginning of BD rings. The eTSEC BD also expands upon the MPC8260 BD model to accommodate the eTSEC's unique features. However, the 8-byte data BD format is designed to be compatible with the existing MPC8260 BD model.

### 15.6.7.1 Data Buffer Descriptors

Data buffers are used in the transmission and reception of Ethernet frames (see [Figure 15-144](#)). Data BDs encapsulate all information necessary for the eTSEC to transmit or receive an Ethernet frame. Within each data BD there is a status field, a data length field, and a data pointer. The BD completely describes an Ethernet packet by centralizing status information for the data packet in the status field of the BD and by containing a data BD pointer to the location of the data buffer. Software is responsible for setting up the BDs in memory. Because of pre-fetching, a minimum of four buffer descriptors per ring are required. This applies to both the transmit and the receive descriptor rings. Transmit rings are limited to a maximum size of 65536 BDs due to BD and frame data prefetching. Software also must have the data pointer pointing to a legal memory location. Within the status field, there exists an ownership bit which defines the current state of the buffer (pointed to by the data pointer). Other bits in the status field of the buffer descriptor are used to communicate status/control information between the eTSEC and the software driver.

Because there is no next BD pointer in the transmit/receive BD (see [Figure 15-145](#)), all BDs must reside sequentially in memory. The eTSEC increments the current BD location appropriately to the next BD location to be processed. There is a wrap bit in the last BD that informs the eTSEC to loop back to the beginning of the BD chain. Software must initialize the TBASE and RBASE registers that point to the beginning transmit and receive BDs for eTSEC.

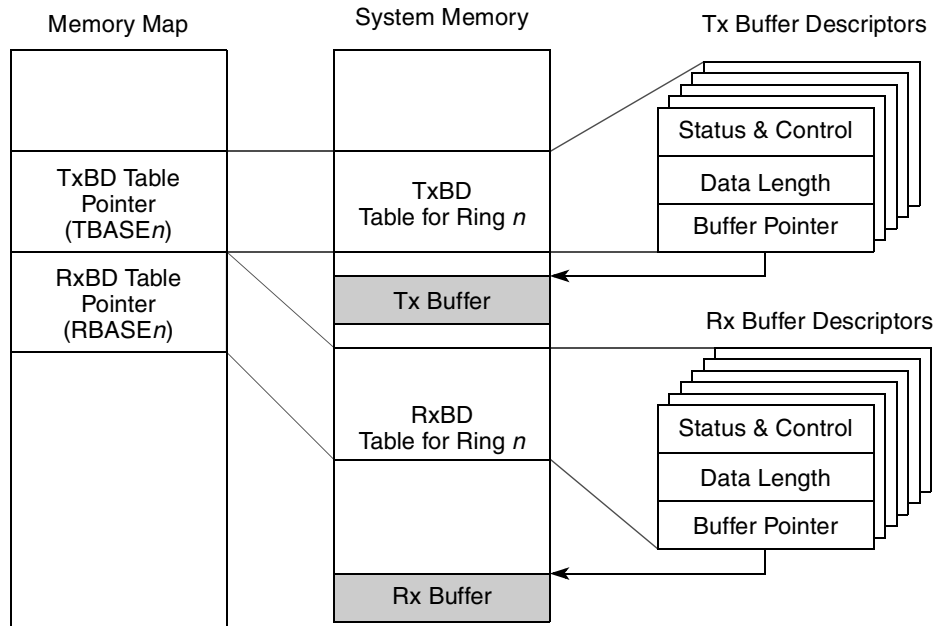


Figure 15-144. Example of eTSEC Memory Structure for BDs




Figure 15-145. Buffer Descriptor Ring

### 15.6.7.2 Transmit Data Buffer Descriptors (TxBD)

Data is presented to the eTSEC for transmission by arranging it in memory buffers referenced by the TxBDs. In the TxBD the user initializes the R, PAD, W, I, L, TC, PRE, HFE, CF, and TOE bits and the length (in bytes) in the first word, and the buffer pointer in the second word. Unused fields or fields written by the eTSEC must be initialized to zero.



The eTSEC clears the R bit in the first word of the BD after it finishes using the data buffer. The transfer status bits are then updated. Additional transmit frame status can be found in statistic counters in the MIB block.

Software must expect eTSEC to prefetch multiple TxBDs, and for TCP/IP checksumming an entire frame must be read from memory before a checksum can be computed. Accordingly, the R bit of the first TxBD in a frame must not be set until at least one entire frame can be fetched from this TxBD onwards. If eTSEC prefetches TxBDs and fails to reach a last TxBD (with bit L set), it halts further transmission from the current TxBD ring and report an underrun error as IEVENT[XFUN]; this indicates that an incomplete frame was fetched, but remained unprocessed. The relevant TBPTR register points to the next unread TxBD following the error.

Figure 15-146 defines the TxBD.

	0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15
Offset + 0	R	PAD/CRC	W	I	L	TC	PRE/DEF	0	HFE/LC	CF/RL	RC			TOE/UN	TR	
Offset + 2	DATA LENGTH															
Offset + 4	TX DATA BUFFER POINTER															
Offset + 6																

**Figure 15-146. Transmit Buffer Descriptor**

The TxBD definition is interpreted by eTSEC hardware as if TxBDs mapped to C data structures in the manner illustrated by Figure 15-147.

```

typedef unsigned short uint_16; /* choose 16-bit native type */
typedef unsigned int uint_32; /* choose 32-bit native type */
typedef struct txbd_struct {
    uint_16 flags;
    uint_16 length;
    uint_32 bufptr;
} txbd;

```

**Figure 15-147. Mapping of TxBDs to a C Data Structure**

The TxBD fields are detailed in [Table 15-163](#).

**Table 15-163. Transmit Data Buffer Descriptor (TxBD) Field Descriptions**

Offset	Bits	Name	Description
0–1	0	R	Ready, written by eTSEC and user. 0 The data buffer associated with this BD is not ready for transmission. The user is free to manipulate this BD or its associated data buffer. The eTSEC clears this bit after the buffer is transmitted or after an error condition is encountered. 1 The data buffer, which is prepared for transmission by the user, was not transmitted or is currently being transmitted. No fields of this BD may be written by the user once this bit is set.
	1	PAD/CRC	Padding for frames. (Valid only while it is set in the first BD and MACCFG2[PAD enable] is cleared). If MACCFG2[PAD enable] is set, this bit is ignored. 0 Do not add padding to short frames. 1 Add PAD to frames. PAD bytes are inserted until the length of the transmitted frame equals 64 bytes. Unlike the MPC8260 which PADs up to MINFLR value, the eTSEC PADs always up to the IEEE minimum frame length of 64 bytes. CRC is always appended to frames.
	2	W	Wrap. Written by user. 0 The next buffer descriptor is found in the consecutive location. 1 The next buffer descriptor is found at the location defined in TBASE.
	3	I	Interrupt. Written by user. 0 No interrupt is generated after this buffer is serviced. 1 IEVENT[TXB] or IEVENT[TXF] are set after this buffer is serviced. These bits can cause an interrupt if they are enabled (That is, IEVENT[TXBEN] or IEVENT[TXFEN] are set).
	4	L	Last in frame. Written by user. 0 The buffer is not the last in the transmit frame. 1 The buffer is the last in the transmit frame.

Table 15-163. Transmit Data Buffer Descriptor (TxBD) Field Descriptions (continued)

Offset	Bits	Name	Description
0–1	5	TC	Tx CRC. Written by user. (Valid only while it is set in first BD and TxBD[PAD/CRC] is cleared and MACCFG2[PAD/CRC enable] is cleared and MACCFG2[CRC enable] is cleared.) If MACCFG2[PAD/CRC enable] is set or MACCFG2[CRC enable] is set, this bit is ignored in ethernet modes. 0 End transmission immediately after the last data byte with no hardware generated CRC appended, unless TxBD[PAD/CRC] is set. 1 Transmit the CRC sequence after the last data byte.
	6	PRE	Transmit user-defined Ethernet preamble. Written by user. Valid only if set in the first BD of a frame, and MACCFG2[PreAm TxEN] is set. 0 This frame does not contain Ethernet preamble bytes for transmission. 1 This frame includes a user-defined Ethernet preamble sequence prior to the destination address in the data buffer.
		DEF	Defer indication. The eTSEC updates this bit after transmitting a frame (TxBD[L] is set) 0 This frame was not deferred. 1 This frame did not have a collision before it was sent but it was sent late because of deferring
	7	—	Reserved
	8	HFE	Huge frame enable. Written by user. Valid only if set in the first BD of a frame and MACCFG2[Huge Frame] is cleared. If MACCFG2[Huge Frame] is set, this bit is ignored. 0 Truncate transmit frame if its length is greater than the MAC's maximum frame length. 1 Allow large frames to be transmitted without truncation.
		LC	Late collision. Written by the eTSEC. 0 No late collision. 1 A collision occurred after 64 bytes are sent. The eTSEC terminates the transmission and updates LC.
	9	CF	Control Frame. Written by user. Valid only if set in the first BD of a frame. 0 Regular frame; transmission is deferred when eTSEC is in PAUSE. 1 Control frame; transmission starts even if eTSEC is in PAUSE.
		RL	Retransmission Limit. Written by the eTSEC. 0 Transmission before maximum retry limit is hit. 1 The transmitter failed (max. retry limit + 1) attempts to successfully send a message due to repeated collisions. The eTSEC terminates the transmission and updates RL.
	10–13	RC	Retry Count. Written by the eTSEC. 0 The frame is sent correctly the first time. x One or more attempts where needed to send the transmit frame. If this field is 15, then 15 or more retries were needed. The Ethernet controller updates RC after sending the buffer.

**Table 15-163. Transmit Data Buffer Descriptor (TxBD) Field Descriptions (continued)**

Offset	Bits	Name	Description
0–1	14	UN	Underrun. Written by the eTSEC. 0 No underrun encountered (data was retrieved from external memory in time to send a complete frame). 1 The Ethernet controller encountered a transmitter underrun condition while sending the associated buffer. This could also have occurred in relation to a bus error causing IEVENT[EBERR]. The eTSEC terminates the transmission and updates UN.
		TOE	TCP/IP off-load enable. Written by user. Valid only if set in the first BD of a frame. 0 No TCP/IP off-load acceleration is applied to the frame prior to transmission. 1 eTSEC looks for a TOE Frame Control Block preceding the frame, and applies TCP/IP off-load acceleration as controlled by the FCB.
	15	TR	Truncation. Written by the eTSEC. Set in the last TxBD (TxBD[L] is set) when IEVENT[BABT] occurs for a frame (a frame length greater than or equal to the value set in the maximum frame length register is encountered, the HFE bit in the BD is cleared, and MACCFG2[Huge Frame] is cleared). The frame is sent truncated.
2–3	0–15	Data Length	Data length is the number of octets the eTSEC should transmit from this BD's data buffer. It is never modified by the eTSEC. This field must be greater than zero, as zero indicates a BD not ready.
4–7	0–31	TX Data Buffer Pointer	The transmit buffer pointer contains the address of the associated data buffer. The data buffer pointer for the first BD of a TxPAL-enabled frame must be aligned on an 8-byte boundary. There are no alignment restrictions for the data buffer pointers of the second or subsequent BDs of a TxPAL-enabled frame, or for non-TxPAL frames.

### 15.6.7.3 Receive Buffer Descriptors (RxB D)

In the RxB D the user initializes the E, I, and W bits in the first word and the pointer in second word. If the data buffer is used, the eTSEC modifies the E, L, F, M, BC, MC, LG, NO, CR, OV, and TR bits and writes the length of the used portion of the buffer in the first word. The M, BC, MC, LG, NO, CR, OV, and TR bits in the first word of the buffer descriptor are only modified by the eTSEC if the L (last BD in frame) bit is set. The first word of the RxB D contains control and status bits. Its formats are detailed below.

The number of buffer descriptors in a ring is set using the W bit to indicate that the next buffer wraps back to the beginning of the ring. See [Section 15.5.3.5.5, “Maximum Frame Length Register \(MAXFRM\),”](#) for information on setting the size of the buffer ring.

Figure 15-148 defines the RxB D.

	0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15
Offset + 0	E	RO1	W	I	L	F	0	M	BC	MC	LG	NO	SH	CR	OV	TR
Offset + 2	DATA LENGTH															
Offset + 4	RX DATA BUFFER POINTER															
Offset + 6																

**Figure 15-148. Receive Buffer Descriptor**

The RxB D definition is interpreted by eTSEC hardware as if RxB Ds mapped to C data structures in the manner illustrated by [Figure 15-149](#).

```

typedef unsigned short uint_16; /* choose 16-bit native type */
typedef unsigned int uint_32; /* choose 32-bit native type */
typedef struct rxbd_struct {
    uint_16 flags;
    uint_16 length;
    uint_32 bufptr;
} rxbd;

```

**Figure 15-149. Mapping of RxBDs to a C Data Structure**

Table 15-164 describes the fields of the RxBD.

**Table 15-164. Receive Buffer Descriptor Field Descriptions**

Offset	Bits	Name	Description
0-1	0	E	Empty, written by the eTSEC (when cleared) and by the user (when set). 0 The data buffer associated with this BD is filled with received data, or data reception is aborted due to an error condition. The status and length fields have been updated as required. 1 The data buffer associated with this BD is empty, or reception is currently in progress.
	1	RO1	Receive software ownership bit. This field is reserved for use by software. This read/write bit is not modified by hardware, nor does its value affect hardware.
	2	W	Wrap, written by user. 0 The next buffer descriptor is found in the consecutive location. 1 The next buffer descriptor is found at the location defined in RBASE.
	3	I	Interrupt, written by user. 0 No interrupt is generated after this buffer is serviced. 1 IEVENT[RXB] or IEVENT[RXF] are set after this buffer is serviced. This bit can cause an interrupt if enabled (IMASK[RXBEN] or IMASK[RXFEN]). If the user wants to be interrupted only if RXF occurs, then the user must disable RXB (IMASK[RXBEN] is cleared) and enable RXF (IMASK[RXFEN] is set).
	4	L	Last in frame, written by the eTSEC. 0 The buffer is not the last in a frame. 1 The buffer is the last in a frame.
	5	F	First in frame, written by the eTSEC. 0 The buffer is not the first in a frame. 1 The buffer is the first in a frame.
	6	—	Reserved
	7	M	Miss, written by the eTSEC. (This bit is valid only if the L-bit is set and eTSEC is in promiscuous mode.) This bit is set by the eTSEC for frames that were accepted in promiscuous mode, but were flagged as a “miss” by the internal address recognition; thus, while in promiscuous mode, the user can use the M-bit to quickly determine whether the frame was destined to this station. 0 The frame was received because of an address recognition hit. 1 The frame was received because of promiscuous mode.

Table 15-164. Receive Buffer Descriptor Field Descriptions (continued)

Offset	Bits	Name	Description
0–1	8	BC	Broadcast. Written by the eTSEC. (Only valid if L is set.) Is set if the DA is broadcast (FF-FF-FF-FF-FF-FF).
	9	MC	Multicast. Written by the eTSEC. (Only valid if L is set.) Is set if the DA is multicast and not BC.
	10	LG	Rx frame length violation, written by the eTSEC (only valid if L is set). A frame length greater than or equal to the maximum frame length was recognized; in this case LG is set regardless of the setting of MACCFG2[Huge Frame]. If MACCFG2[Huge Frame] is cleared, the frame is truncated to the value programmed in the maximum frame length register. This bit is valid only if the L bit is set.
	11	NO	Rx non-octet aligned frame, written by the eTSEC (only valid if L is set). A frame that contained a number of bits not divisible by eight was received.
	12	SH	Short frame, written by the eTSEC (only valid if L is set). A frame length less than the minimum 64B that is defined for ethernet. was recognized, provided RCTRL[RSF] is set.
	13	CR	Rx CRC error, written by the eTSEC (only valid if L is set). This frame contains a CRC error and is an integral number of octets in length. This bit is also set if a receive code group error is detected.
	14	OV	Overflow, written by the eTSEC (only valid if L is set). A receive FIFO overflow occurred during frame reception. If this bit is set, the other status bits, M, LG, NO, CR and TR lose their normal meaning and are zero.
	15	TR	Truncation, written by the eTSEC (only valid if L is set). Set if the receive frame is truncated. This can happen if a frame length greater than the maximum frame length is received and MACCFG2[Huge Frame] is cleared. If this bit is set, the frame must be discarded and the other error bits must be ignored as they may be incorrect.
2–3	0–15	Data Length	Data length, written by the eTSEC. Data length is the number of octets written by the eTSEC into this BD's data buffer if L is cleared (the value is equal to MRBLR), or, if L is set, the length of the frame including CRC, FCB (if RCTRL[PRSDEP > 00]), preamble (if MACCFG2[PreAmRxEn]=1), timestamp (if RCTRL[TS]=1) and any padding (RCTRL[PAL]).
4–7	0–31	RX Data Buffer Pointer	Receive buffer pointer, written by the user. The receive buffer pointer, which always points to the first location of the associated data buffer, must be 8-byte aligned. For best performance, use 64-byte aligned receive buffer pointer addresses. The buffer must reside in memory external to the eTSEC.

## 15.7 Initialization/Application Information

### 15.7.1 Interface Mode Configuration

This section describes how to configure the eTSEC in different supported interface modes. These include the following:

- MII
- RMII

- RGMII
- SGMII
- RTBI

The pinout, the data registers that must be initialized, as well as speed selection options are described.

### 15.7.1.1 MII Interface Mode

Table 15-165 describes the signal configurations required for MII interface mode.

**Table 15-165. MII Interface Mode Signal Configuration**

eTSEC Signals			MII Interface		
			Frequency [MHz] 25		
			Voltage [V] 3.3		
Signals	I/O	No. of Signals	Signals	I/O	No. of Signals
GTX_CLK	O	1	leave unconnected		
TX_CLK	I	1	TX_CLK	I	1
TxD[0]	O	1	TxD[0]	O	1
TxD[1]	O	1	TxD[1]	O	1
TxD[2]	O	1	TxD[2]	O	1
TxD[3]	O	1	TxD[3]	O	1
TX_EN	O	1	TX_EN	O	1
TX_ER	O	1	TX_ER	O	1
RX_CLK	I	1	RX_CLK	I	1
RxD[0]	I	1	RxD[0]	I	1
RxD[1]	I	1	RxD[1]	I	1
RxD[2]	I	1	RxD[2]	I	1
RxD[3]	I	1	RxD[3]	I	1
RX_DV	I	1	RX_DV	I	1
RX_ER	I	1	RX_ER	I	1
COL	I	1	COL	I	1
CRS	I	1	CRS	I	1
<b>Sum</b>		17	<b>Sum</b>		16

Table 15-166 describes the shared signals of the MII interface.

**Table 15-166. Shared MII Signals**

eTSEC Signals	I/O	No. of Signals	MII Signals	I/O	No. of Signals	Function
MDIO	I/O	1	MDIO	I/O	1	Management interface I/O
MDC	O	1	MDC	O	1	Management interface clock
ECGTX_CLK125	I	1	not used	I	0	Reference clock
<b>Sum</b>			<b>Sum</b>			

Table 15-167 describes the register initializations required to configure the eTSEC in MII mode.

**Table 15-167. MII Mode Register Initialization Steps**

Set Soft_Reset, MACCFG1[1000_0000_0000_0000_0000_0000_0000_0000]
Clear Soft_Reset, MACCFG1[0000_0000_0000_0000_0000_0000_0000_0000]
Initialize MACCFG2, for MII, half duplex operation. Set I/F Mode bit, MACCFG2[0000_0000_0000_0000_0111_0001_0000_0100] (This example has Full Duplex = 0, Preamble count = 7, PAD/CRC append = 1)
Initialize ECNTRL, ECNTRL[0000_0000_0000_0000_0001_0000_0000_0000] (This example has Statistics Enable = 1)
Initialize MAC Station Address, MACSTNADDR2[0110_0000_0000_0010_0000_0000_0000_0000] Set station address to 02_60_8C_87_65_43, for example.
Initialize MAC Station Address, MACSTNADDR1[0100_0011_0110_0101_1000_0111_1000_1100] Set station address to 02_60_8C_87_65_43, for example.
Reset the management interface. MIIMCFG[1000_0000_0000_0000_0000_0000_0000_0111]
Setup the MII Mgmt clock speed, MIIMCFG[0000_0000_0000_0000_0000_0000_0000_0101] set source clock divide by 14 for example to insure that MDC clock speed is not greater than 2.5 MHz
Read MII Mgmt Indicator register and check for Busy = 0, MIIMIND ---> [0000_0000_0000_0000_0000_0000_0000_0000] This indicates that the eTSEC MII Mgmt bus is idle.
Set up the MII Mgmt for a write cycle to the external PHY Auxiliary Control and Status Register to configure the PHY through the Management interface (overrides configuration signals of the PHY). MIIMADD[0000_0000_0000_0000_0000_0000_0001_1100]
Perform an MII Mgmt write cycle to the external PHY Writing to MII Mgmt Control with 16-bit data intended for the external PHY register, MIIMCON[0000_0000_0000_0000_0000_0000_0000_0100]



**Table 15-167. MII Mode Register Initialization Steps (continued)**

<p>Check to see if MII Mgmt write is complete  Read MII Mgmt Indicator register and check for Busy = 0,  MIIMIND ---&gt; [0000_0000_0000_0000_0000_0000_0000_0000]  This indicates that the write cycle was completed.</p>
<p>Set up the MII Mgmt for a write cycle to the external PHY Extended PHY control register #1 to set up the interface mode selection.  MIIMADD[0000_0000_0000_0000_0000_0000_0001_0111]</p>
<p>Perform an MII Mgmt write cycle to the external PHY.  Write to MII Mgmt Control with 16-bit data intended for the external PHY register,  MIIMCON[0000_0000_0000_0000_0000_0000_0000_0000]</p>
<p>Check to see if MII Mgmt write is complete.  Read MII Mgmt Indicator register and check for Busy = 0,  MIIMIND ---&gt; [0000_0000_0000_0000_0000_0000_0000_0000]  This indicates that the write cycle was completed.</p>
<p>Set up the MII Mgmt for a write cycle to the external PHY Mode control register to set up the interface mode selection.  MIIMADD[0000_0000_0000_0000_0000_0000_0000_0000]</p>
<p>Perform an MII Mgmt write cycle to the external PHY.  Write to MII Mgmt Control with 16-bit data intended for the external PHY register,  MIIMCON[0000_0000_0000_0000_00uu_00uu_0u00_0000]  where u is user defined based on desired configuration.</p>
<p>Check to see if MII Mgmt write is complete  Read MII Mgmt Indicator register and check for Busy = 0,  MIIMIND ---&gt; [0000_0000_0000_0000_0000_0000_0000_0000]  This indicates that the write cycle was completed.</p>
<p>If auto-negotiation was enabled in the PHY, check to see if PHY has completed Auto-Negotiation.  Set up the MII Mgmt for a read cycle to PHY MII Mgmt register (write the PHY address and Register address),  MIIMADD[0000_0000_0000_0000_0000_0000_0000_0001]  The PHY Status register is at address 0x1 and in this case the PHY Address is 0x00.</p>
<p>Perform an MII Mgmt read cycle of Status Register.  Clear MIIMCOM[Read Cycle].  Set MIIMCOM[Read Cycle].  (Uses the PHY address (0) and Register address (1) placed in MIIMADD register),  When MIIMIND[BUSY]=0,  read the MIIMSTAT register and check bit 10 (AN Done and Link is up)  MIIMSTAT ---&gt; [0000_0000_0000_0000_0000_0000_0010_0100]  Other information about the link is also returned.(Extend Status, No pre, Remote Fault, An Ability, Link status, extend Ability)</p>
<p>Check auto-negotiation attributes in the PHY as necessary.</p>
<p>Clear IEVENT register,  IEVENT[0000_0000_0000_0000_0000_0000_0000_0000]</p>
<p>Initialize IMASK (Optional)  IMASK[0000_0000_0000_0000_0000_0000_0000_0000]</p>
<p>Initialize MACnADDR1/2 (Optional)  MACnADDR1/2[0000_0000_0000_0000_0000_0000_0000_0000]</p>
<p>Initialize GADDR<sub>n</sub> (Optional)  GADDR<sub>n</sub>[0000_0000_0000_0000_0000_0000_0000_0000]</p>
<p>Initialize RCTRL (Optional)  RCTRL[0000_0000_0000_0000_0000_0000_0000_0000]</p>

**Table 15-167. MII Mode Register Initialization Steps (continued)**

Initialize DMACTRL (Optional) DMACTRL[0000_0000_0000_0000_0000_0000_0000_0000]
Initialize (Empty) Transmit Descriptor ring and fill buffers with Data Initialize TBASE0–TBASE7, TBASE0–TBASE7[LLLL_LLLL_LLLL_LLLL_LLLL_LLLL_LLLL_L000]
Initialize (Empty) Receive Descriptor ring and fill with empty buffers Initialize RBASE0–RBASE7, RBASE0–RBASE7[LLLL_LLLL_LLLL_LLLL_LLLL_LLLL_LLLL_L000]
Enable Transmit Queues Initialize TQUEUE
Enable Receive Queues Initialize RQUEUE
Enable Rx and Tx, MACCFG1[0000_0000_0000_0000_0000_0000_0000_0101]

### 15.7.1.2 RGMII Interface Mode

Table 15-168 shows the signals configurations required for RGMII interface mode.

**Table 15-168. RGMII Interface Mode Signal Configuration**

eTSEC Signals			RGMII Interface		
			Frequency [MHz] 125		
			Voltage [V] 2.5		
Signals	I/O	No. of Signals	Signals	I/O	No. of Signals
GTX_CLK	O	1	GTX_CLK	O	1
TX_CLK	I	1	not used		
TxD[0]	O	1	TxD[0]/TxD[4]	O	1
TxD[1]	O	1	TxD[1]/TxD[5]	O	1
TxD[2]	O	1	TxD[2]/TxD[6]	O	1
TxD[3]	O	1	TxD[3]/TxD[7]	O	1
TX_EN	O	1	TX_CTL (TX_EN/TX_ERR)	O	1
TX_ER	O	1	leave unconnected		
RX_CLK	I	1	RX_CLK	I	1
RxD[0]	I	1	RxD[0]/RxD[4]	I	1
RxD[1]	I	1	RxD[1]/RxD[5]	I	1
RxD[2]	I	1	RxD[2]/RxD[6]	I	1
RxD[3]	I	1	RxD[3]/RxD[7]	I	1
RX_DV	I	1	RX_CTL (RX_DV/RX_ERR)	I	1

**Table 15-168. RGMII Interface Mode Signal Configuration (continued)**

eTSEC Signals			RGMII Interface		
			Frequency [MHz] 125		
			Voltage [V] 2.5		
Signals	I/O	No. of Signals	Signals	I/O	No. of Signals
RX_ER	I	1	not used		
COL	I	1	not used		
CRS	I	1	not used		
<b>Sum</b>		17	<b>Sum</b>		12

Table 15-169 describes the shared signals for the RGMII interface.

**Table 15-169. Shared RGMII Signals**

eTSEC Signals	I/O	No. of Signals	GMII Signals	I/O	No. of Signals	Function
MDIO	I/O	1	MDIO	I/O	1	Management interface I/O
MDC	O	1	MDC	O	1	Management interface clock
GTX_CLK125	I	1	GTX_CLK125	I	1	Reference clock
<b>Sum</b>			<b>Sum</b>			

Table 15-170 describes the register initializations required to configure the eTSEC in RGMII mode.

**Table 15-170. RGMII Mode Register Initialization Steps**

Set Soft_Reset, MACCFG1[1000_0000_0000_0000_0000_0000_0000]
Clear Soft_Reset, MACCFG1[0000_0000_0000_0000_0000_0000_0000]
Initialize MACCFG2, MACCFG2[0000_0000_0000_0000_0111_0010_0000_0101] (I/F Mode = 2, Full Duplex = 1)
Initialize ECNTRL, ECNTRL[0000_0000_0000_0000_0001_0000_0000_0000] (This example has RGMII 10Mbps mode, Statistics Enable = 1)
Initialize MAC Station Address, MACSTNADDR2[0110_0000_0000_0010_0000_0000_0000_0000] to 02608C:876543, for example.
Initialize MAC Station Address, MACSTNADDR1[0100_0011_0110_0101_1000_0111_1000_1100] to 02608C:876543, for example.

**Table 15-170. RGMII Mode Register Initialization Steps (continued)**

<p>Setup the MII Mgmt clock speed,  MIIMCFG[0000_0000_0000_0000_0000_0000_0101]  Set source clock divide by 14, for example, to insure that TSEC_MDC clock speed is not greater than 2.5 MHz.</p>
<p>Read MII Mgmt Indicator register and check for Busy = 0,  MIIMIND ---&gt; [0000_0000_0000_0000_0000_0000_0000]  This indicates that the eTSEC MII Mgmt bus is idle.</p>
<p>Set up the MII Mgmt for a write cycle to external the PHY AN Advertisement register  (write the PHY address and Register address),  MIIMADD[0000_0000_0000_0000_0001_0001_0000_0100]  The AN Advertisement register is at offset address 0x04 from the external PHY address. (in this case 0x11)</p>
<p>Perform an MII Mgmt write cycle to the external PHY.  Write to MII Mgmt Control with 16-bit data intended for the external PHY AN Advertisement register,  MIIMCON[0000_0000_0000_0000_u0uu_uuuu_uuuu_uuuu]  Where u must be selected by the user for proper system configuration.</p>
<p>Check to see if MII Mgmt write is complete.  Read MII Mgmt Indicator register and check for Busy = 0,  MIIMIND ---&gt; [0000_0000_0000_0000_0000_0000_0000]  This indicates that the write cycle was completed.</p>
<p>Set up the MII Mgmt for a write cycle to the external PHY Control register (write the PHY address and Register address),  MIIMADD[0000_0000_0000_0000_0001_0001_0000_0000]  The control register (CR) is at offset address 0x00 from the external PHY address. (in this case 0x11)</p>
<p>Perform an MII Mgmt write cycle to the external PHY.  Write to MII Mgmt Control with 16-bit data intended for the external PHY Control register,  MIIMCON[0000_0000_0000_0000_0001_0010_0000_0000]  This enables the external PHY to restart Auto-Negotiations using the configuration set in the AN Advertisement register.</p>
<p>Check to see if MII Mgmt write is complete.  Read MII Mgmt Indicator register and check for Busy = 0,  MIIMIND ---&gt; [0000_0000_0000_0000_0000_0000_0000]  This indicates that the write cycle was completed.</p>
<p>Check to see if PHY has completed Auto-Negotiation.  Set up the MII Mgmt for a read cycle to the PHY MII Mgmt register (write the PHY address and Register address),  MIIMADD[0000_0000_0000_0000_0000_0010_0000_0001]  The PHY Status register is at address 0x1 and in this case the PHY Address is 0x2.</p>
<p>Perform an MII Mgmt read cycle of Status Register.  Clear MIIMCOM[Read Cycle]  Set MIIMCOM[Read Cycle]  (Uses the PHY address (2) and Register address (2) placed in MIIMADD register)  When MIIMIND[BUSY]=0,  read the MIIMSTAT register and check bit 10. (AN Done)  MIIMSTAT ---&gt; [0000_0000_0000_0000_0000_0000_0010_0000]  Other information about the link is also returned. (Extend Status, No pre, Remote Fault, An Ability, Link status, extend Ability)</p>
<p>Perform an MII Mgmt read cycle of AN Expansion Register.  Setup MIIMADD[0000_0000_0000_0000_0001_0001_0000_0110]  Clear MIIMCOM[Read Cycle]  Set MIIMCOM[Read Cycle]  (Uses the PHY address (0x11) and Register address (6) placed in MIIMADD register)  When MIIMIND[BUSY]=0,  read the MII Mgmt AN Expansion register and check bits 13 and 14. (NP Able and Page Rx'd)  MII Mgmt AN Expansion ---&gt; [0000_0000_0000_0000_0000_0000_0000_0110]</p>

**Table 15-170. RGII Mode Register Initialization Steps (continued)**

<p>Perform an MII Mgmt read cycle of AN Link Partner Base Page Ability Register. (Optional)  Setup MIIMADD[0000_0000_0000_0000_0001_0001_0000_0101]  Clear MIIMCOM[Read Cycle]  Set MIIMCOM[Read Cycle]  (Uses the PHY address (0x11) and Register address (5) placed in MIIMADD register)  When MIIMIND[BUSY]=0,  read the MII Mgmt AN Link Partner Base Page Ability register and check bits 9 and 10. (Half and Full Duplex)  MII Mgmt AN Link Partner Base Page Ability ---&gt; [0000_0000_0000_0000_0000_000x_1x10_0000]</p>
<p>Clear IEVENT register,  IEVENT[0000_0000_0000_0000_0000_0000_0000_0000]</p>
<p>Initialize IMASK (Optional)  IMASK[0000_0000_0000_0000_0000_0000_0000_0000]</p>
<p>Initialize MACnADDR1/2 (Optional)  MACnADDR1/2[0000_0000_0000_0000_0000_0000_0000_0000]</p>
<p>Initialize GADDR<sub>n</sub> (Optional)  GADDR<sub>n</sub>[0000_0000_0000_0000_0000_0000_0000_0000]</p>
<p>Initialize RCTRL (Optional)  RCTRL[0000_0000_0000_0000_0000_0000_0000_0000]</p>
<p>Initialize DMACTRL (Optional)  DMACTRL[0000_0000_0000_0000_0000_0000_0000_0000]</p>
<p>Initialize (Empty) Transmit Descriptor ring and fill buffers with Data  Initialize TBASE0–TBASE7,  TBASE0–TBASE7[LLLL_LLLL_LLLL_LLLL_LLLL_LLLL_LLLL_L000]</p>
<p>Initialize (Empty) Receive Descriptor ring and fill with empty buffers  Initialize RBASE0–RBASE7,  RBASE0–RBASE7[LLLL_LLLL_LLLL_LLLL_LLLL_LLLL_LLLL_L000]</p>
<p>Enable Transmit Queues  Initialize TQUEUE</p>
<p>Enable Receive Queues  Initialize RQUEUE</p>
<p>Enable Rx and Tx,  MACCFG1[0000_0000_0000_0000_0000_0000_0000_0101]</p>

### 15.7.1.3 RMII Interface Mode

Table 15-171 shows the signals configurations required for RMII interface mode.

**Table 15-171. RMII Interface Mode Signal Configuration**

eTSEC Signals			RMII Interface		
			Frequency [MHz] 50		
			Voltage [V] 3.3		
Signals	I/O	No. of Signals	Signals	I/O	No. of Signals
GTX_CLK	O	1	leave unconnected		
TX_CLK	I	1	REF_CLK	I	1
TxD[0]	O	1	TxD[0]	O	1
TxD[1]	O	1	TxD[1]	O	1
TxD[2]	O	1	leave unconnected		
TxD[3]	O	1	leave unconnected		
TX_EN	O	1	TX_EN	O	1
TX_ER	O	1	leave unconnected		
RX_CLK	I	1	leave unconnected		
RxD[0]	I	1	RxD[0]	I	1
RxD[1]	I	1	RxD[1]	I	1
RxD[2]	I	1	not used		
RxD[3]	I	1	not used		
RX_DV	I	1	CRS_DV	I	1
RX_ER	I	1	RX_ER	I	1
COL	I	1	not used		
CRS	I	1	not used		
<b>Sum</b>		17	<b>Sum</b>		8

Table 15-172 describes the shared signals for the RMII interface.

**Table 15-172. Shared RMII Signals**

eTSEC Signals	I/O	No. of Signals	GMII Signals	I/O	No. of Signals	Function
MDIO	I/O	1	MDIO	I/O	1	Management interface I/O
MDC	O	1	MDC	O	1	Management interface clock
TX_CLK	I	1	REF_CLK	I	1	Reference clock
<b>Sum</b>		3	<b>Sum</b>		3	

Table 15-173 describes the register initializations required to configure the eTSEC in RMII mode.

**Table 15-173. RMII Mode Register Initialization Steps**

Set Soft_Reset, MACCFG1[1000_0000_0000_0000_0000_0000_0000_0000]
Clear Soft_Reset, MACCFG1[0000_0000_0000_0000_0000_0000_0000_0000]
Initialize MACCFG2, MACCFG2[0000_0000_0000_0000_0111_0010_0000_0101] (I/F Mode = 2, Full Duplex = 1)
Initialize ECNTRL, ECNTRL[0000_0000_0000_0000_0001_0000_0001_0000] (Used to setup Reduced-Pin mode = 1, and TBIM = 0, statistics enable = 1)
Initialize MAC Station Address MACSTNADDR2[0110_0000_0000_0010_0000_0000_0000_0000] to 02608C:876543 for example
Initialize MAC Station Address MACSTNADDR1[0100_0011_0110_0101_1000_0111_1000_1100] to 02608C:876543 for example
Setup the MII Mgmt clock speed, MIIMCFG[0000_0000_0000_0000_0000_0000_0000_1101] set system clock divide by 14 for example to insure that MDC clock speed = 2.5 MHz
Read MII Mgmt Indicator register and check for Busy = 0, MIIMIND ---> [0000_0000_0000_0000_0000_0000_0000_0000] This indicates that the eTSEC MII Mgmt bus is idle.
Set up the MII Mgmt for a write cycle to external the PHY AN Advertisement register (write the PHY address and Register address), MIIMADD[0000_0000_0000_0000_0001_0001_0000_0100] The AN Advertisement register is at offset address 0x04 from the external PHY address. (in this case 0x11)
Perform an MII Mgmt write cycle to the external PHY. Write to MII Mgmt Control with 16-bit data intended for the external PHY AN Advertisement register, MIIMCON[0000_0000_0000_0000_u0uu_uuuu_uuuu_uuuu] Where u must be selected by the user for proper system configuration.
Check to see if MII Mgmt write is complete. Read MII Mgmt Indicator register and check for Busy = 0, MIIMIND ---> [0000_0000_0000_0000_0000_0000_0000_0000] This indicates that the write cycle was completed.
Set up the MII Mgmt for a write cycle to the external PHY Control register (write the PHY address and Register address), MIIMADD[0000_0000_0000_0000_0001_0001_0000_0000] The control register is at offset address 0x00 from the external PHY address. (in this case 0x11)
Perform an MII Mgmt write cycle to the external PHY. Write to MII Mgmt Control with 16-bit data intended for the external PHY Control register, MIIMCON[0000_0000_0000_0000_0001_0010_0000_0000] This enables the external PHY to restart Auto-Negotiations using the configuration set in the AN Advertisement register.
Check to see if MII Mgmt write is complete. Read MII Mgmt Indicator register and check for Busy = 0, MIIMIND ---> [0000_0000_0000_0000_0000_0000_0000_0000] This indicates that the write cycle was completed.

**Table 15-173. RMII Mode Register Initialization Steps (continued)**

<p>Check to see if PHY has completed Auto-Negotiation.  Set up the MII Mgmt for a read cycle to the PHY MII Mgmt register (write the PHY address and Register address),  MIIMADD[0000_0000_0000_0000_0000_0010_0000_0001]  The PHY Status register is at address 0x1 and in this case the PHY Address is 0x2.</p>
<p>Perform an MII Mgmt read cycle of Status Register.  Clear MIIMCOM[Read Cycle]  Set MIIMCOM[Read Cycle]  (Uses the PHY address (2) and Register address (1) placed in MIIMADD register)  When MIIMIND[BUSY]=0,  read the MIIMSTAT register and check bit 10. (AN Done)  MIIMSTAT ---&gt; [0000_0000_0000_0000_0000_0010_0000]  Other information about the link is also returned. (Extend Status, No pre, Remote Fault, An Ability, Link status, extend Ability)</p>
<p>Perform an MII Mgmt read cycle of AN Expansion Register.  Setup MIIMADD[0000_0000_0000_0000_0001_0001_0000_0110]  Clear MIIMCOM[Read Cycle]  Set MIIMCOM[Read Cycle]  (Uses the PHY address (0x11) and Register address (6) placed in MIIMADD register)  When MIIMIND[BUSY]=0,  read the MII Mgmt AN Expansion register and check bits 13 and 14. (NP Able and Page Rx'd)  MII Mgmt AN Expansion ---&gt; [0000_0000_0000_0000_0000_0000_0000_0110]</p>
<p>Perform an MII Mgmt read cycle of AN Link Partner Base Page Ability Register. (Optional)  Setup MIIMADD[0000_0000_0000_0000_0001_0001_0000_0101]  Clear MIIMCOM[Read Cycle]  Set MIIMCOM[Read Cycle]  (Uses the PHY address (0x11) and Register address (5) placed in MIIMADD register)  When MIIMIND[BUSY]=0,  read the MII Mgmt AN Link Partner Base Page Ability register and check bits 9 and 10. (Half and Full Duplex)  MII Mgmt AN Link Partner Base Page Ability ---&gt; [0000_0000_0000_0000_0000_000x_x110_0000]</p>
<p>Setting up the MII Mgmt for a write cycle to TBI MII Mgmt register (write the TBI's address and Register address),  MIIMADD[0000_0000_0000_0000_0001_0000_0000_1011]  the TBI control register is at offset address 0x11 from TBIPA</p>
<p>Perform an MII Mgmt write cycle  Writing to MII Mgmt Control with 16-bit data intended for TBI's MII Mgmt control register (TBI control),  MIIMCON[0000_0000_0000_0000_0000_0010_0001_0000]  This configures the TBI control to GMII mode and AN sense</p>
<p>Check to see if MII Mgmt write is complete  Read MII Mgmt Indicator register and check for Busy = 0,  MIIMIND ---&gt; [0000_0000_0000_0000_0000_0000_0000_0000]  This indicate that the write cycle was completed</p>
<p>Perform an MII Mgmt read cycle (Optional)  Set MIIMCOM[Read Cycle]  (Uses the TBI address and Register address placed in MIIMADD register),  read the MIIMSTAT register and verify that  MIIMSTAT ---&gt; [0000_0000_0000_0000_0000_0010_0001_0000]</p>
<p>Check to see if PHY has completed Auto-Negotiation  Setting up the MII Mgmt for a read cycle to PHY's MII Mgmt register (write the PHY's address and Register address),  MIIMADD[0000_0000_0000_0000_0000_0010_0000_0010]  the PHY Status control register is at address 0x2 and lets say the PHY Address is 0x2</p>



**Table 15-173. RMII Mode Register Initialization Steps (continued)**

<p>Perform an MII Mgmt read cycle of Status Register Set MIIMCOM[Read Cycle] (Uses the PHY address (2) and Register address (2) placed in MIIMADD register), read the MIIMSTAT register and check bit 10 (AN Done) MIIMSTAT ---&gt; [0000_0000_0000_0000_0000_0000_0010_0000] other information about the link is also returned (Extend Status, No pre, Remote Fault, An Ability, Link status, extend Ability)</p>
<p>Perform an MII Mgmt read cycle of AN Expansion Register MIIMADD[0000_0000_0000_0000_0000_0010_0000_0110] Set MIIMCOM[Read Cycle] (Uses the PHY address (2) and Register address (6) placed in MIIMADD register), read the MII Mgmt AN Expansion register and check bits 13 and 14 (NP Able and Page Rx'd) MII Mgmt AN Expansion ---&gt; [0000_0000_0000_0000_0000_0000_0000_0110]</p>
<p>Perform an MII Mgmt read cycle of AN Link Partner Base Page Ability Register (Optional) MIIMADD[0000_0000_0000_0000_0000_0010_0000_0101] Set MIIMCOM[Read Cycle] (Uses the PHY address (2) and Register address (5) placed in MIIMADD register), read the MII Mgmt AN Link Partner Base Page Ability register and check bits 9 and 10 (Half and Full Duplex) MII Mgmt AN Link Partner Base Page Ability ---&gt; [0000_0000_0000_0000_0000_0000_00X_1110_0000]</p>
<p>Clear IEVENT register, IEVENT[0000_0000_0000_0000_0000_0000_0000_0000]</p>
<p>Initialize IMASK (Optional) IMASK[0000_0000_0000_0000_0000_0000_0000_0000]</p>
<p>Initialize GADDR<sub>n</sub> (Optional) GADDR<sub>n</sub>[0000_0000_0000_0000_0000_0000_0000_0000]</p>
<p>Initialize RCTRL (Optional) RCTRL[0000_0000_0000_0000_0000_0000_0000_0000]</p>
<p>Initialize DMACTRL (Optional) DMACTRL[0000_0000_0000_0000_0000_0000_0000_0000]</p>
<p>Initialize (Empty) Transmit Descriptor ring and fill buffers with Data Initialize TBASE0–TBASE7, TBASE0–TBASE7[LLLL_LLLL_LLLL_LLLL_LLLL_LLLL_LLLL_L000]</p>
<p>Initialize (Empty) Receive Descriptor ring and fill with empty buffers Initialize RBASE0–RBASE7, RBASE0–RBASE7[LLLL_LLLL_LLLL_LLLL_LLLL_LLLL_LLLL_L000]</p>
<p>Enable Transmit Queues Initialize TQUEUE</p>
<p>Enable Receive Queues Initialize RQUEUE</p>
<p>Enable Rx and Tx, MACCFG1[0000_0000_0000_0000_0000_0000_0000_0101]</p>

### 15.7.1.4 RTBI Interface Mode

Table 15-174 describes the signal configurations required for RTBI interface mode.

**Table 15-174. RTBI Interface Mode Signal Configuration**

eTSEC Signal s			RTBI Interface		
			Frequency [MHz] 125		
			Voltage [V] 2.5		
Signals	I/O	No. of Signals	Signals	I/O	No. of Signals
GTX_CLK	O	1	GTX_CLK	O	1
TX_CLK	I	1	not used		
TxD[0]	O	1	TCG[0]/TCG[5]	O	1
TxD[1]	O	1	TCG[1]/TCG[6]	O	1
TxD[2]	O	1	TCG[2]/TCG[7]	O	1
TxD[3]	O	1	TCG[3]/TCG[8]	O	1
TX_EN	O	1	TCG[4]/TCG[9]	O	1
TX_ER	O	1	leave unconnected		
RX_CLK	I	1	RX_CLK	I	1
RxD[0]	I	1	RCG[0]/RCG[5]	I	1
RxD[1]	I	1	RCG[1]/RCG[6]	I	1
RxD[2]	I	1	RCG[2]/RCG[7]	I	1
RxD[3]	I	1	RCG[3]/RCG[8]	I	1
RX_DV	I	1	RCG[4]/RCG[9]	I	1
RX_ER	I	1	not used		
COL	I	1	not used		
CRS	I	1	not used		
<b>Sum</b>		17	<b>sum</b>		12

Table 15-175 describes the shared signals for the RTBI interface.

**Table 15-175. Shared RTBI Signals**

eTSEC Signals	I/O	No. of Signals	GMII Signals	I/O	No. of Signals	Function
MDIO	I/O	1	MDIO	I/O	1	Management interface I/O
MDC	O	1	MDC	O	1	Management interface clock
ECGTX_CLK125	I	1	GTX_CLK125	I	1	Reference clock
Sum		3	Sum		3	

Table 15-176 describes the register initializations required to configure the eTSEC in RTBI mode.

**Table 15-176. RTBI Mode Register Initialization Steps**

<p>Set Soft_Reset, MACCFG1[1000_0000_0000_0000_0000_0000_0000_0000]</p>
<p>Clear Soft_Reset, MACCFG1[0000_0000_0000_0000_0000_0000_0000_0000]</p>
<p>Initialize MACCFG2, MACCFG2[0000_0000_0000_0000_0111_0010_0000_0101] (I/F Mode = 2, Full Duplex = 1)</p>
<p>Initialize ECNTRL, ECNTRL[0000_0000_0000_0000_0001_0000_0000_0000] (This example has Statistics Enable = 1)</p>
<p>Initialize MAC Station Address, MACSTNADDR2[0110_0000_0000_0010_0000_0000_0000_0000] to 02608C:876543, for example.</p>
<p>Initialize MAC Station Address, MACSTNADDR1[0100_0011_0110_0101_1000_0111_1000_1100] to 02608C:876543, for example.</p>
<p>Assign a Physical address to the TBI, TBIPA[0000_0000_0000_0000_0000_0000_0001_0000] set to 16, for example.</p>
<p>Setup the MII Mgmt clock speed, MIIMCFG[0000_0000_0000_0000_0000_0000_0000_0101] Set source clock divide by 14, for example, to insure that TSEC_MDC clock speed is not greater than 2.5 MHz.</p>
<p>Read MII Mgmt Indicator register and check for Busy = 0, MIIMIND ---&gt; [0000_0000_0000_0000_0000_0000_0000_0000] This indicates that the eTSEC MII Mgmt bus is idle.</p>
<p>Set up the MII Mgmt for a read cycle to TBI's Control register (write the TBI's address and Register address), MIIMADD[0000_0000_0000_0000_0001_0000_0000_0000] The control register (CR) is at offset address 0x0 from TBIPA.</p>
<p>Perform an MII Mgmt read cycle to verify state of TBI Control Register(Optional) Clear MIIMCOM[Read Cycle] Set MIIMCOM[Read Cycle] (Uses the TBI address and Register address placed in MIIMADD register), When MIIMIND[BUSY]=0, read the MIIMSTAT and look for AN Enable and other bit information.</p>
<p>Set up the MII Mgmt for a write cycle to TBI's AN Advertisement register (write the PHY address and Register address), MIIMADD[0000_0000_0000_0000_0001_0000_0000_0100] The AN Advertisement register is at offset address 0x04 from the TBI's address. (in this case 0x10)</p>
<p>Perform an MII Mgmt write cycle to TBI. Write to MII Mgmt Control with 16-bit data intended for TBI's AN Advertisement register, MIIMCON[0000_0000_0000_0000_0000_0001_1010_0000] This advertises to the Link Partner that the TBI supports PAUSE and Full Duplex mode and does not support Half Duplex mode</p>
<p>Check to see if MII Mgmt write is complete. Read MII Mgmt Indicator register and check for Busy = 0, MIIMIND ---&gt; [0000_0000_0000_0000_0000_0000_0000_0000] This indicates that the write cycle was completed.</p>

**Table 15-176. RTBI Mode Register Initialization Steps (continued)**

<p>Set up the MII Mgmt for a write cycle to TBI's Control register (write the PHY address and Register address), MIIMADD[0000_0000_0000_0000_0001_0000_0000_0000] The control register (CR) is at offset address 0x00 from the TBI's address. (in this case 0x10)</p>
<p>Perform an MII Mgmt write cycle to TBI. Writing to MII Mgmt Control with 16-bit data intended for TBI's Control register, MIIMCON[0000_0000_0000_0000_0001_0010_0000_0000] This enables the TBI to restart Auto-Negotiations using the configuration set in the AN Advertisement register.</p>
<p>Check to see if MII Mgmt write is complete. Read MII Mgmt Indicator register and check for Busy = 0, MIIMIND ---&gt; [0000_0000_0000_0000_0000_0000_0000_0000] This indicates that the write cycle was completed.</p>
<p>Check to see if PHY has completed Auto-Negotiation. Set up the MII Mgmt for a read cycle to the PHY MII Mgmt register (write the PHY address and Register address), MIIMADD[0000_0000_0000_0000_0001_0000_0000_0001] The PHY Status control register is at address 0x1 and in this case the PHY Address is 0x10.</p>
<p>Perform an MII Mgmt read cycle of Status Register. Clear MIIMCOM[Read Cycle] Set MIIMCOM[Read Cycle] (Uses the PHY address (2) and Register address (2) placed in MIIMADD register), When MIIMIND[BUSY]=0, read the MIIMSTAT register and check bit 10 (AN Done) MIIMSTAT ---&gt; [0000_0000_0000_0000_0000_0000_0010_0000] Other information about the link is also returned. (Extend Status, No pre, Remote Fault, An Ability, Link status, extend Ability)</p>
<p>Perform an MII Mgmt read cycle of AN Expansion Register. Setup MIIMADD[0000_0000_0000_0000_0001_0000_0000_0110] Clear MIIMCOM[Read Cycle] Set MIIMCOM[Read Cycle] (Uses the PHY address (0x10) and Register address (6) placed in MIIMADD register), When MIIMIND[BUSY]=0, read the MII Mgmt AN Expansion register and check bits 13 and 14. (NP Able and Page Rx'd) MII Mgmt AN Expansion ---&gt; [0000_0000_0000_0000_0000_0000_0000_0110]</p>
<p>Perform an MII Mgmt read cycle of AN Link Partner Base Page Ability Register. (Optional) Setup MIIMADD[0000_0000_0000_0000_0001_0000_0000_0101] Clear MIIMCOM[Read Cycle] Set MIIMCOM[Read Cycle] (Uses the PHY address (0x10) and Register address (5) placed in MIIMADD register), When MIIMIND[BUSY]=0, read the MII Mgmt AN Link Partner Base Page Ability register and check bits 9 and 10. (Half and Full Duplex) MII Mgmt AN Link Partner Base Page Ability ---&gt; [0000_0000_0000_0000_0000_0000_00x_x110_0000]</p>
<p>Clear IEVENT register, IEVENT[0000_0000_0000_0000_0000_0000_0000_0000]</p>
<p>Initialize IMASK (Optional) IMASK[0000_0000_0000_0000_0000_0000_0000_0000]</p>
<p>Initialize MACnADDR1/2 (Optional) MACnADDR1/2[0000_0000_0000_0000_0000_0000_0000_0000]</p>
<p>Initialize GADDR<sub>n</sub> (Optional) GADDR<sub>n</sub>[0000_0000_0000_0000_0000_0000_0000_0000]</p>

**Table 15-176. RTBI Mode Register Initialization Steps (continued)**

Initialize RCTRL (Optional) RCTRL[0000_0000_0000_0000_0000_0000_0000_0000]
Initialize DMACTRL (Optional) DMACTRL[0000_0000_0000_0000_0000_0000_0000_0000]
Initialize (Empty) Transmit Descriptor ring and fill buffers with Data Initialize TBASE0–TBASE7, TBASE0–TBASE7[LLLL_LLLL_LLLL_LLLL_LLLL_LLLL_LLLL_L000]
Initialize (Empty) Receive Descriptor ring and fill with empty buffers Initialize RBASE0–RBASE7, RBASE0–RBASE7[LLLL_LLLL_LLLL_LLLL_LLLL_LLLL_LLLL_L000]
Enable Transmit Queues Initialize TQUEUE
Enable Receive Queues Initialize RQUEUE
Enable Rx and Tx, MACCFG1[0000_0000_0000_0000_0000_0000_0000_0101]

### 15.7.1.5 SGMII Interface Support

**Table 15-177. SGMII Interface Signal Configuration (4-Wire)**

SerDes Signals			SGMII Interface		
Frequency [MHz] 1250			Frequency [MHz] 1250		
Voltage [V] LVDS			Voltage [V] LVDS		
Signals	I/O	No. of Signals	Signals	I/O	No. of Signals
$TX_n/\overline{TX}_n$	O	2	TXD	O	2
$RX_n/\overline{RX}_n$	I	2	RXD	I	2
<b>Sum</b>		4	<b>Sum</b>		4

SGMII mode initialization sequence is very similar to TBI mode initialization. Additional initialization is required for the SerDes. An example of SGMII mode initialization sequence is shown in [Table 15-178](#).

#### NOTE

SGMII mode utilizes the internal TBI PHY. The internal TBI PHY only auto-negotiates at 1 Gbps. However, 10 Mbps and 100 Mbps speeds are supported in SGMII mode. It is recommended that the external PHY inform the MAC if the desired link speed is not 1 Gbps. Software can perform MII management cycles to determine the external PHY link speed and program ECNTRL and MACCFG2 accordingly.

**Table 15-178. SGMII Mode Register Initialization Steps**

<i>Initialize SerDes to select SGMII. The initialization sequence should be prepended with SerDes initialization.</i>
Set Soft_Reset, MACCFG1[1000_0000_0000_0000_0000_0000_0000_0000]
Clear Soft_Reset, MACCFG1[0000_0000_0000_0000_0000_0000_0000_0000]
Initialize MACCFG2, MACCFG2[0000_0000_0000_0000_0111_0010_0000_0101] (I/F Mode = 2, Full Duplex = 1) (Set I/F mode = 1 in SGMII 10/100 Mbps speed)
Initialize ECNTRL, ECNTRL[0000_0000_0000_0000_0001_0000_0010_0010] (This example has Statistics Enable = 1, TBIM = 1, SGMIIIM = 1) (Set R100M = 1 in SGMII 100 Mbps speed)
Initialize MAC Station Address MACSTNADDR2[0110_0000_0000_0010_0000_0000_0000_0000] to 02608C:876543, for example.
Initialize MAC Station Address MACSTNADDR1[0100_0011_0110_0101_1000_0111_1000_1100] to 02608C:876543, for example.
Assign a Physical address to the TBI, TBIPA[0000_0000_0000_0000_0000_0000_0001_0000] set to 16, for example.
Setup the MII Mgmt clock speed, MIIMCFG[0000_0000_0000_0000_0000_0000_0000_0101] set source clock divide by 14 for example to insure that MDC clock speed is not greater than 2.5 MHz
Read MII Mgmt Indicator register and check for Busy = 0, MIIMIND ---> [0000_0000_0000_0000_0000_0000_0000_0000] This indicates that the eTSEC MII Mgmt bus is idle.
Set up the MII Mgmt for a read cycle to TBI's Control register (write the PHY address and Register address), MIIMADD[0000_0000_0000_0000_0001_0000_0000_0000] the control register (CR) is at offset address 0x00 from the TBI's address.
Perform an MII Mgmt read cycle to verify state of TBI Control Register (optional) Clear MIIMCOM[Read Cycle] Set MIIMCOM[Read Cycle] (Uses the TBI address and Register address placed in MIIMADD register), When MIIMIND[BUSY] = 0, read the MIIMSTAT and look for AN Enable and other bit information.
Set up the MII Mgmt for a write cycle to TBICON register (write the PHY address and Register address), MIIMADD[0000_0000_0000_0000_0001_0000_0001_0001] The TBICON register is at offset address 0x11 from the TBI's address.
Perform an MII Mgmt write cycle to TBI. Writing to MII Mgmt Control with 16-bit data intended for TBICON register, MIIMCON[0000_0000_0000_0000_0000_0000_0010_0000] This sets TBI in single clock mode and MII Mode off to enable communication with SerDes.

**Table 15-178. SGMII Mode Register Initialization Steps (continued)**

<p>Check to see if MII Mgmt write is complete.  Read MII Mgmt Indicator register and check for Busy = 0,  MIIMIND ---&gt; [0000_0000_0000_0000_0000_0000_0000_0000]  This indicates that the write cycle was completed.</p>
<p>Set up the MII Mgmt for a write cycle to TBI's AN Advertisement register (write the PHY address and Register address),  MIIMADD[0000_0000_0000_0000_0001_0000_0000_0100]  The AN Advertisement register is at offset address 0x04 from the TBI's address.</p>
<p>Perform an MII Mgmt write cycle to TBI.  Writing to MII Mgmt Control with 16-bit data intended for TBI's AN Advertisement register,  MIIMCON[0000_0000_0000_0000_0000_0001_1010_0000]  This advertises to the Link Partner that the TBI supports PAUSE and Full Duplex mode and does not support Half Duplex mode.</p>
<p>Check to see if MII Mgmt write is complete.  Read MII Mgmt Indicator register and check for Busy = 0,  MIIMIND ---&gt; [0000_0000_0000_0000_0000_0000_0000_0000]  This indicates that the write cycle was completed.</p>
<p><i>Additional SerDes setup as required</i></p>
<p>Set up the MII Mgmt for a write cycle to TBI's Control register (write the PHY address and Register address),  MIIMADD[0000_0000_0000_0000_0001_0000_0000_0000]  the control register (CR) is at offset address 0x00 from the TBI's address.</p>
<p>Perform an MII Mgmt write cycle to TBI.  Writing to MII Mgmt Control with 16-bit data intended for TBI's Control register,  MIIMCON[0000_0000_0000_0000_0001_0011_0100_0000]  This enables the TBI to restart Auto-Negotiations using the configuration set in the AN Advertisement register.</p>
<p>Check to see if MII Mgmt write is complete.  Read MII Mgmt Indicator register and check for Busy = 0,  MIIMIND ---&gt; [0000_0000_0000_0000_0000_0000_0000_0000]  This indicates that the write cycle was completed.</p>
<p>Check to see if PHY has completed Auto-Negotiation.  Set up the MII Mgmt for a read cycle to PHY MII Mgmt register (write the PHY address and Register address),  MIIMADD[0000_0000_0000_0000_0001_0000_0000_0001]  The PHY Status control register is at address 0x1 and in this case the PHY Address is 0x10.</p>
<p>Perform an MII Mgmt read cycle of Status Register.  Clear MIIMCOM[Read Cycle]  Set MIIMCOM[Read Cycle]  (Uses the PHY address (2) and Register address (2) placed in MIIMADD register),  When MIIMIND[BUSY] = 0,  read the MIIMSTAT register and check bit 10 (AN Done)  MII Mgmt AN Expansion ---&gt; [0000_0000_0000_0000_0000_0000_0000_0110]  Other information about the link is also returned. (Extend Status, No pre, Remote Fault, An Ability, Link status, extend Ability)</p>
<p>Perform an MII Mgmt read cycle of AN Expansion Register.  Setup MIIMADD[0000_0000_0000_0000_0001_0000_0000_0110]  Clear MIIMCOM[Read Cycle]  Set MIIMCOM[Read Cycle]  (Uses the PHY address (0x10) and Register address (6) placed in MIIMADD register),  When MIIMIND[BUSY] = 0,  read the MII Mgmt AN Expansion register and check bits 13 and 14 (NP Able and Page Rx'd)  MII Mgmt AN Expansion ---&gt; [0000_0000_0000_0000_0000_0000_0000_0110]</p>

**Table 15-178. SGMII Mode Register Initialization Steps (continued)**

<p>Perform an MII Mgmt read cycle of AN Link Partner Base Page Ability Register. (Optional)  Setup MIIMADD[0000_0000_0000_0000_0001_0000_0000_0101]  Clear MIIMCOM[Read Cycle]  Set MIIMCOM[Read Cycle]  (Uses the PHY address (0x10) and Register address (5) placed in MIIMADD register),  When MIIMIND[BUSY] = 0,  read the MII Mgmt AN Link Partner Base Page Ability register and check bits 9 and 10. (Half and Full Duplex)  MII Mgmt AN Link Partner Base Page Ability ---&gt; [0000_0000_0000_0000_0000_000x_1110_0000]</p>
<p>Clear IEVENT register,  IEVENT[0000_0000_0000_0000_0000_0000_0000_0000]</p>
<p>Initialize IMASK (Optional)  IMASK[0000_0000_0000_0000_0000_0000_0000_0000]</p>
<p>Initialize MACnADDR1/2 (Optional)  MACnADDR1/2[0000_0000_0000_0000_0000_0000_0000_0000]</p>
<p>Initialize GADDR<sub>n</sub> (Optional)  GADDR<sub>n</sub>[0000_0000_0000_0000_0000_0000_0000_0000]</p>
<p>Initialize RCTRL (Optional)  RCTRL[0000_0000_0000_0000_0000_0000_0000_0000]</p>
<p>Initialize DMACTRL (Optional)  DMACTRL[0000_0000_0000_0000_0000_0000_0000_0000]</p>
<p>Initialize (Empty) Transmit Descriptor ring and fill buffers with Data  Initialize TBASE0–TBASE7,  TBASE0–TBASE7[LLLL_LLLL_LLLL_LLLL_LLLL_LLLL_LLLL_L000]</p>
<p>Initialize (Empty) Receive Descriptor ring and fill with empty buffers  Initialize RBASE0–RBASE7,  RBASE0–RBASE7[LLLL_LLLL_LLLL_LLLL_LLLL_LLLL_LLLL_L000]</p>
<p>Enable Transmit Queues  Initialize TQUEUE</p>
<p>Enable Receive Queues  Initialize RQUEUE</p>
<p>Enable Rx and Tx,  MACCFG1[0000_0000_0000_0000_0000_0000_0000_0101]</p>



## Chapter 16

# Universal Serial Bus Interface

This chapter describes the universal serial bus (USB) interface of the device. The USB interface implements many industry standards. However, it is beyond the scope of this document to document the intricacies of these standards. Instead, it is left to the reader to refer to the governing specifications.

The following documents are available from the USB Implementers Forum web page at <http://www.usb.org/developers/docs/>.

- *Universal Serial Bus Revision 2.0 Specification*
- *On-The-Go Supplement to the USB 2.0 Specification, Revision 1.0a*

The following documents are available from the Intel USB Specifications web page at <http://www.intel.com/technology/usb/spec.htm>.

- *Enhanced Host Controller Interface (EHCI) Specification for Universal Serial Bus, Revision 1.0*
- *USB 2.0 Transceiver Macrocell Interface (UTMI) Specification, Version 1.05*

The following documents are available from the ULPI web page at <http://www.ulpi.org/>.

- *UTMI+ Specification, Revision 1.0*
- *UTMI Low Pin-Count Interface (ULPI) Specification, Revision 1.0*

## 16.1 Introduction

The device implements a dual-role (DR) USB module. This module may be connected to an external port. Collectively the module and external port are called the USB interface. The USB interface is shown in Figure 16-1.




Figure 16-1. USB Interface Block Diagram

### 16.1.1 Overview

The USB DR module is a USB 2.0-compliant serial interface engine for implementing a USB interface. The registers and data structures for the module are based on the *Enhanced Host Controller Interface Specification for Universal Serial Bus* (EHCI) from Intel Corporation. The DR module can act as a device or host controller. Interfaces to negotiate the host or device role on the bus in compliance with the On-The-Go (OTG) supplement to the USB specification are also provided.

The DR module supports the required signaling for USB transceiver macrocell interface (UTMI) and UTMI low pin count interface (ULPI) transceivers (PHYs). The PHY interfacing to the ULPI is an external PHY.

The module contains a chaining DMA (direct memory access) engine that reduces the interrupt load on the application processor and reduces the total system bus bandwidth that must be dedicated to servicing the USB interface requirements.

### 16.1.2 Features

The USB DR module includes the following features:

- Complies with USB specification rev 2.0
- Supports operation as a standalone USB host controller

- Supports enhanced host controller interface (EHCI)
- Supports high-speed (480 Mbps), full-speed (12 Mbps), and low-speed (1.5 Mbps) operation. Low speed is only supported in host mode.
- Supports external PHY with ULPI (UTMI + low-pin interface)
- Supports operation as a standalone USB device
  - Supports one upstream facing port
  - Supports three programmable, bidirectional USB endpoints
- Host and device support
- OTG (on-the-go) support, which includes both device and host capability, with external PHY (ULPI)

### 16.1.3 Modes of Operation

The USB DR module has three basic operating modes: host, device, and OTG. The module can be configured to use one of two different PHY interfaces: ULPI or UTMI.

#### NOTE

Only high-speed and full-speed operations are supported in device mode.

## 16.2 External Signals

This section contains detailed descriptions of all the USB dual-role controller signals. Many of the signals for the PHY interfaces are muxed onto the same pins in order to reduce pin count. [Table 16-1](#) describes the signals, indicating which interface supports each signal.

**Table 16-1. USB External Signals**

Signal	I/O	Description
USBDR_D0_ENABLEN	I/O	ULPI—Use as USBDR_D0USB_DP
USBDR_D1_SER_TXD	I/O	ULPI—Use as USBDR_D1USB_DM
USBDR_D2_VMO_SE0	I/O	ULPI—Use as USBDR_D2USB_VBUS
USBDR_D3_SPEED	I/O	ULPI—Use as USBDR_D3USB_PHY_GND
USBDR_D4_DP	I/O	ULPI—Use as USBDR_D4USB_TPA
USBDR_D5_DM	I/O	ULPI—Use as USBDR_D5USB_RBIAS
USBDR_D6_SER_RCV	I/O	ULPI—Use as USBDR_D6USB_PLL_PWR3
USBDR_D7_DRVVBUS	I/O	ULPI—Use as USBDR_D7USB_PLL_GND
USBDR_DIR_DPPULLUP	I/O	ULPI—Use as USBDR_DIRUSB_PLL_PWR1
USBDR_STP_SUSPEND	O	ULPI—Use as USBDR_STPUSB_VSSA_BIAS
USBDR_PWRFAULT	I	ULPI—Use as USBDR_PWRFAULTUSB_VDDA_BIAS
USBDR_PCTL0	O	ULPI—Use as USBDR_PCTL0USB_VSSA

**Table 16-1. USB External Signals (continued)**

Signal	I/O	Description
USBDR_PCTL1	O	ULPI—Use as USBDR_PCTL1USB_VDDA
USBDR_CLK	I	ULPI—Use as USBDR_CLK

### 16.2.1 UTMI Interface

UTMI was developed to specify a standard interface between USB 2.0 controllers and USB 2.0 PHY's. This interface is made available to support applications that may use a UTMI-compliant PHY. UTMI+ extensions are not supported by the USB DR module. Functionality added by UTMI+ is available in the ULPI interface. Only the integrated PHY uses the UTMI interface; therefore, there are no external UTMI signals to be documented in this manual. The integrated USB PHY has four dedicated external signals, which are only used when the MPC8313E is a host: USBDR\_DRIVE\_VBUS, USBDR\_PWRFAULT, USBDR\_PCTL0, and USBDR\_PCTL1. These signals are also part of ULPI Interface as described below.

### 16.2.2 ULPI Interface

The ULPI (UTMI low pin count interface) is a reduced pin-count (12 signals) extension of the UTMI+ specification. Pin count is reduced by converting relatively static signals to register bits, and providing a bidirectional, generic data bus that carries USB and register data. This interface minimizes pin count requirements for external PHYs. [Table 16-2](#) describes the signals for the ULPI interface.

**Table 16-2. ULPI Signal Descriptions**

Signal	I/O	Description	
USBDR_DIR	I	Direction. USBDR_DIR controls the direction of the data bus. When the PHY has data to transfer to USB port, it drives USBDR_DIR high to take ownership of the bus. When the PHY has no data to transfer it drives USBDR_DIR low and monitors the bus for link activity. The PHY pulls USBDR_DIR high whenever the interface cannot accept data from the link.	
		<b>State Meaning</b>	Asserted—PHY has data to transfer to the link. Negated—PHY has no data to transfer.
		<b>Timing</b>	Synchronous to PHY_CLK.
USBDR_NXT	I	Next data. The PHY asserts USBDR_NXT to throttle the data. When USB port is sending data to the PHY, USBDR_NXT indicates when the current byte has been accepted by the PHY. The USB port places the next byte on the data bus in the following clock cycle. When the PHY is sending data to USB port, USBDR_NXT indicates when a new byte is available for USB port to consume.	
		<b>State Meaning</b>	Asserted—PHY is ready to transfer byte. Negated—PHY is not ready.
		<b>Timing</b>	Synchronous to PHY_CLK.

Table 16-2. ULPI Signal Descriptions (continued)

Signal	I/O	Description
USBDR_STP	O	Stop. USBDR_STP indicates the end of a transfer on the bus.
		<b>State Meaning</b> Asserted—USB asserts this signal for 1 clock cycle to stop the data stream currently on the bus. If USB port is sending data to the PHY, USBDR_STP indicates the last byte of data was previously on the bus. If the PHY is sending data to USB port, USBDR_STP forces the PHY to end its transfer, negate USBDR_DIR and relinquish control of the data bus to the USB port. Negated—Indicates normal operation.
		<b>Timing</b> Synchronous to PHY_CLK.
USBDR_PWRFAULT	I	Power fault. USBDR_PWRFAULT indicates whether a power fault occurred on the USB port Vbus.
		<b>State Meaning</b> Asserted—Indicates that a Vbus fault occurred. Applications that support power switching must shut down Vbus power. Negated—Indicates normal operation.
		<b>Timing</b> Synchronous to PHY_CLK.
USBDR_PCTL0	O	Port control 0. USBDR_PCTL0 controls the port status indicator LED 0 when in host mode.
		<b>State Meaning</b> Asserted—LED on. Negated—LED off.
		<b>Timing</b> Synchronous to PHY_CLK.
USBDR_PCTL1	O	Port control 1. USBDR_PCTL1 controls the port status indicator LED 1 when in host mode.
		<b>State Meaning</b> Asserted—LED on. Negated—LED off.
		<b>Timing</b> Synchronous to PHY_CLK.
USBDR_TXRXD[7:0]	I/O	Data bit $n$ . USBDR_TXRXD $n$ is bit $n$ of the 8-bit (USBDR_TXRXD7–USBDR_TXRXD0), uni-directional data bus used to carry USB, register, and interrupt data between the PHY and the USB controller.
		<b>State Meaning</b> Asserted—Data bit $n$ is 1. Negated—Data bit $n$ is 0.
		<b>Timing</b> Synchronous to PHY_CLK.

### 16.2.3 PHY Clocks

The USBDR\_CLK input provides the clocking signal for the ULPI PHY interface. The clock is 60 MHz. Detailed clock specifications are given in the appropriate hardware specifications document.

## 16.3 Memory Map/Register Definitions

This section provides the memory map and detailed descriptions of all USB interface registers. The memory map of the USB interface is shown in [Table 16-3](#).

Table 16-3. USB Interface Memory Map

Offset	Register	Access	Reset	Section/Page
<b>USB DR Controller Registers</b>				
0x2_3000–0x2_30FF	Reserved, should be cleared	—	—	—
0x2_3100	CAPLENGTH—Capability register length	R	0x40	16.3.1.1/16-8
0x2_3102	HCVERSION—Host interface version number <sup>1</sup>	R	0x0100	16.3.1.2/16-8
0x2_3104	HCSPARAMS—Host ctrl. structural parameters <sup>1</sup>	R	0x0001_0011	16.3.1.3/16-9
0x2_3108	HCCPARAMS—Host ctrl. capability parameters <sup>1</sup>	R	0x0000_0006	16.3.1.4/16-10
0x2_3120	DCVERSION—Device interface version number	R	0x0001	16.3.1.5/16-10
0x2_3124	DCCPARAMS—Device controller parameters	R	0x0000_0183	16.3.1.6/16-11
0x2_3140	USBCMD—USB command	Mixed	0x0008_0000	16.3.2.1/16-12
0x2_3144	USBSTS—USB status	Mixed	0x0000_0000	16.3.2.2/16-14
0x2_3148	USBINTR—USB interrupt enable	R/W	0x0000_0000	16.3.2.3/16-17
0x2_314C	FRINDEX—USB frame index	R/W	0x0000_0000	16.3.2.4/16-18
0x2_3154	PERIODICLISTBASE—Frame list base address <sup>1</sup>	R/W	0x0000_0000	16.3.2.6/16-19
	DEVICEADDR—USB device address	R/W	0x0000_0000	16.3.2.7/16-20
0x2_3158	ASYNCLISTADDR—Next asynchronous list addr (host mode) <sup>1</sup>	R/W	0x0000_0000	16.3.2.8/16-20
	ENDPOINTLISTADDR—Address at endpoint list (device mode)	R/W	0x0000_0000	16.3.2.9/16-21
0x2_3160	BURSTSIZE—Programmable burst size	R/W	0x0000_1010	16.3.2.10/16-22
0x2_3164	TXFILLTUNING—Host TT transmit pre-buffer packet tuning	R/W	0x0000_0000	16.3.2.11/16-22
0x2_3170	ULPI VIEWPORT—ULPI Register Access	Mixed	0x0000_0000	16.3.2.12/16-24
0x2_3180	CONFIGFLAG—Configured flag register	R	0x0000_0001	16.3.2.13/16-26
0x2_3184	PORTSC—Port status/control	Mixed	0x1000_0000	16.3.2.14/16-26
0x2_31A4	OTGSC—On-The-Go status and control <sup>1</sup>	Mixed	0x0000_0C20	16.3.2.15/16-31
0x2_31A8	USBMODE—USB device mode	R/W	0x0000_0000	16.3.2.16/16-34
0x2_31AC	ENDPTSETUPSTAT—Endpoint setup status	R/W	0x0000_0000	16.3.2.17/16-35
0x2_31B0	ENDPOINTPRIME—Endpoint initialization	R/W	0x0000_0000	16.3.2.18/16-35
0x2_31B4	ENDPTFLUSH—Endpoint de-initialize	R/W	0x0000_0000	16.3.2.19/16-36
0x2_31B8	ENDPTSTATUS—Endpoint status	R	0x0000_0000	16.3.2.20/16-36
0x2_31BC	ENDPTCOMPLETE—Endpoint complete	w1c	0x0000_0000	16.3.2.21/16-37
0x2_31C0	ENDPTCTRL0—Endpoint control 0	Mixed	0x0080_0080	16.3.2.22/16-38
0x2_31C4	ENDPTCTRL1—Endpoint control 1	R/W	0x0000_0000	16.3.2.23/16-39
0x2_31C8	ENDPTCTRL2—Endpoint control 2	R/W	0x0000_0000	16.3.2.23/16-39

Table 16-3. USB Interface Memory Map (continued)

Offset	Register	Access	Reset	Section/Page
0x2_31CA– 0x2_31D4	Reserved	—	—	—
0x2_3400	SNOOP1—Snoop 1	R/W	0x0000_0000	16.3.2.24/16-40
0x2_3404	SNOOP2—Snoop 2	R/W	0x0000_0000	16.3.2.24/16-40
0x2_3408	AGE_CNT_THRESH—Age count threshold	R/W	0x0000_0000	16.3.2.25/16-41
0x2_340C	PRI_CTRL—Priority control	R/W	0x0000_0000	16.3.2.26/16-42
0x2_3410	SI_CTRL—System interface control	R/W	0x0000_0000	16.3.2.27/16-43
0x2_3500	CONTROL—Control	Mixed	0x0000_0000	16.3.2.28/16-43
0x2_3504– 0x2_3FFF	Reserved, should be cleared	—	—	—

<sup>1</sup> This register has separate functions for the host and device operation; the host function is listed first in the table.

The following sections provide details about the registers in the USB memory map.

#### NOTE

Memory may be viewed from either a big-endian or little-endian byte ordering perspective depending on the processor configuration. In big-endian mode, the most-significant byte of word 0 is located at address 0 and the least-significant byte of word 0 is located at address 3. In little-endian mode, the least-significant byte of word 0 is located at address 0 and the most-significant byte of word 0 is located at address 3. Within registers, bits are numbered within a word starting with bit 31 as the most-significant bit. By convention USB registers use little-endian byte ordering. In the USB DR module, these are the registers from offsets 0x00 to 0x1FF. The registers associated with the internal system interface (0x400 and above) use big-endian byte ordering.

### 16.3.1 Capability Registers

The capability registers specify the software limits, restrictions, and capabilities of the host/device controller implementation. Most of these registers are defined by the EHCI specification. Registers that are not defined by the EHCI specification are noted in their descriptions.

### 16.3.1.1 Capability Registers Length (CAPLENGTH)

CAPLENGTH is used as an offset to add to the register base address to find the beginning of the operational register space, that is, the location of the USBCMD register. Figure 16-2 shows CAPLENGTH.




Figure 16-2. Capability Registers Length (CAPLENGTH)

Table 16-4 provides bit descriptions for the CAPLENGTH register.

Table 16-4. CAPLENGTH Register Field Descriptions

Bits	Name	Description
7-0	CAPLENGTH	Capability registers length. Value is 0x40.

### 16.3.1.2 Host Controller Interface Version (HCIVERSION)

HCIVERSION contains a BCD encoding of the EHCI revision number supported by this host controller. The most-significant byte of the register represents a major revision and the least-significant byte is the minor revision. Figure 16-3 shows the HCIVERSION register.




Figure 16-3. Host Controller Interface Version (HCIVERSION)

Table 16-5 provides bit descriptions for the HCIVERSION register.

Table 16-5. HCIVERSION Register Field Descriptions

Bits	Name	Description
15-0	—	EHCI revision number. Value is 0x0100 indicating version 1.0.



### 16.3.1.3 Host Controller Structural Parameters (HCSPARAMS)

HCSPARAMS contains structural parameters such as the number of downstream ports. Figure 16-4 shows the HCSPARAMS register.




Figure 16-4. Host Controller Structural Parameters (HCSPARAMS)

Table 16-6 provides bit descriptions for the HCSPARAMS register.

Table 16-6. HCSPARAMS Register Field Descriptions

Bits	Name	Description
31–28	—	Reserved, should be cleared.
27–24	N_TT	Number of transaction translators. This is a non-EHCI field. This field indicates the number of embedded transaction translators associated the module. The reset value of this field is always 1 after the USBDR controller is configured as a host by writing 0x3 to USBMODE; else, the reset value is always 0. See Section 16.9.1, “Embedded Transaction Translator Function.”
23–20	N_PTT	Ports per transaction translator. This is a non-EHCI field. The number of ports assigned to each transaction translator. This is equal to N_PORTS.
19–17	—	Reserved, should be cleared.
16	PI	Port indicators. Indicates whether the ports support port indicator control. The reset value of this field is always 0 after the USBDR controller is configured as a host by writing 0x3 to USBMODE; else, the reset value is always 1. 1 The port status and control registers include a R/W field for controlling the state of the port indicator.
15–12	N_CC	Number of companion controllers associated with the DR controller. Always 0.
11–8	N_PCC	Number ports per CC. This field indicates the number of ports supported per internal companion controller. Always 0.
7–5	—	Reserved, should be cleared.
4	PPC	Power port control. Indicates whether the host controller supports port power control. The reset value of this field is always 0 after the USBDR controller is configured as a host by writing 0x3 to USBMODE; else, the reset value is always 1. 1 Ports have power port switches.
3–0	N_PORTS	Number of ports. Number of physical downstream ports implemented for host applications. The value of this field determines how many port registers are addressable in the operational register. The reset value of this field is always 0 after the USBDR controller is configured as a host by writing 0x3 to USBMODE; else, the reset value is always 1.

### 16.3.1.4 Host Controller Capability Parameters (HCCPARAMS)

HCCPARAMS identifies multiple mode control (time-base bit functionality) addressing capability. Figure 16-5 shows the HCCPARAMS register.




Figure 16-5. Host Control Capability Parameters (HCCPARAMS)

Table 16-7 provides bit descriptions for the HCCPARAMS register.

Table 16-7. HCCPARAMS Register Field Descriptions

Bits	Name	Description
31–16	—	Reserved, should be cleared.
15–8	EECP	EHCI extended capabilities pointer. Indicates the existence of a capabilities list. A value of 0x00 indicates no extended capabilities are implemented. A non-zero value in this register indicates the offset in PCI configuration space of the first EHCI extended capability. The pointer value must be 0x40 or greater if implemented to maintain the consistency of the PCI header defined for this class of device. This field is always 0.
7–4	IST	Isochronous scheduling threshold. Indicates, relative to the current position of the executing host controller, where software can reliably update the isochronous schedule. When bit 7 is zero, the value of the least significant 3 bits indicates the number of microframes a host controller can hold a set of isochronous data structures (one or more) before flushing the state. When bit 7 is a one, then host software assumes the host controller may cache an isochronous data structure for an entire frame. This field is always 0.
3	—	Reserved, should be cleared.
2	ASP	Asynchronous schedule park capability. Indicates whether the USB DR module supports the park feature for high-speed queue heads in the asynchronous schedule. The feature can be disabled or enabled and set to a specific level by using the asynchronous schedule park mode enable and asynchronous schedule park mode count fields in the USBCMD register. This field is always 1 (park feature supported).
1	PFL	Programmable frame list flag. Indicates whether system software can specify and use a frame list length less than 1024 elements. Frame list size is configured via the USBCMD register frame list size field. The frame list must always be aligned on a 4K page boundary. This requirement ensures that the frame list is always physically contiguous. This field is always 1.
0	ADC	64-bit addressing capability. Always 0; 64-bit addressing is not supported. 0 Data structures use 32-bit address memory pointers

### 16.3.1.5 Device Controller Interface Version (DCIVERSION)—Non-EHCI

This register is not defined in the EHCI specification. DCIVERSION is a two-byte register containing a BCD encoding of the device controller interface. The most-significant byte of the register represents a major revision and the least-significant byte is the minor revision. Figure 16-6 shows the DCIVERSION register.

Offset 0x2\_3120

Access: Read-only



**Figure 16-6. Device Interface Version (DCIVERSION)**

Table 16-8 provides bit descriptions for the DCIVERSION register.

**Table 16-8. DCIVERSION Register Field Descriptions**

Bits	Name	Description
15–0	DCIVERSION	Device interface revision number.

### 16.3.1.6 Device Controller Capability Parameters (DCCPARAMS)—Non-EHCI

This register is not defined in the EHCI specification. This register describes the overall host/device capability of the DR module. Figure 16-7 shows the DCCPARAMS register.

Offset 0x2\_3124

Access: Read-only



**Figure 16-7. Device Control Capability Parameters (DCCPARAMS)**

Table 16-9 provides bit descriptions for the DCCPARAMS register.

**Table 16-9. DCCPARAMS Register Field Descriptions**

Bits	Name	Description
31–9	—	Reserved, should be cleared.
8	HC	Host capable. Always 1, indicating the USB DR controller can operate as an EHCI compatible USB 2.0 host
7	DC	Device capable. Always 1, indicating the USB DR controller can operate as an USB 2.0 device. 1 Device capability. 0 No device capability (host only).
6–5	—	Reserved, should be cleared.
4–0	DEN	Device endpoint number. Indicates the number of endpoints built into the device controller. Always 0x3.

## 16.3.2 Operational Registers

The operational registers are comprised of dynamic control or status registers that may be read-only, read/write, or read/write-1-to-clear. The following sections define the operational registers.

### 16.3.2.1 USB Command Register (USBCMD)

The module executes the command indicated in this register.

Offset 0x2\_3140

Access: Mixed




Figure 16-8. USB Command Register (USBCMD)

Table 16-10. USBCMD Register Field Descriptions

Bits	Name	Description
31–24	—	Reserved, should be cleared.
23–16	ITC	Interrupt threshold control. The system software uses this field to set the maximum rate at which the USB DR module will issue interrupts. ITC contains the maximum interrupt interval measured in microframes. Valid values are shown below. 0x00 Immediate (no threshold) 0x01 1 microframe 0x02 2 microframes 0x04 4 microframes 0x08 8 microframes 0x10 16 microframes 0x20 32 microframes 0x40 40 microframes
15	FS2	See bits 3–2 below. This is a non-EHCI bit.
14	ATDTW	Add dTD TripWire. This is a non-EHCI bit. Used as a semaphore when a dTD is added to an active (primed) endpoint. This bit is set and cleared by software. This bit shall also be cleared by hardware when is state machine is hazard region where adding a dTD to a primed endpoint may go unrecognized. More information on the use of this bit is described in <a href="#">Section 16.9.2, “Device Operation.”</a>
13	SUTW	Setup tripwire. This is a non-EHCI bit. Used as a semaphore when the 8 bytes of setup data read extracted from a QH by the DCD. If the setup lockout mode is off (See USBMODE) then there exists a hazard when new setup data arrives and the DCD is copying setup from the QH for a previous setup packet. This bit is set and cleared by software and will be cleared by hardware when a hazard exists. More information on the use of this bit is described in <a href="#">Section 16.9.2, “Device Operation.”</a>
12	—	Reserved, should be cleared.
11	ASPE	Asynchronous schedule park mode enable. Software uses this bit to enable or disable park mode. The reset value of this field is always 1 after the USBDR controller is configured as a host by writing 0x3 to USBMODE; else, the reset value is always 0. 0 Disabled 1 Enabled
10	—	Reserved, should be cleared.

Table 16-10. USBCMD Register Field Descriptions (continued)

Bits	Name	Description
9–8	ASP	Asynchronous schedule park mode count. The reset value of this field is always 0x3 after the USBDR controller is configured as a host by writing 0x3 to USBMODE; else, the reset value is always 0x0. It contains a count of the number of successive transactions the host controller is allowed to execute from a high-speed queue head on the Asynchronous schedule before continuing traversal of the Asynchronous schedule. Valid values are 0x1H to 0x3H. Software must not write a zero to this field when ASPE is set as this will result in undefined behavior.
7	LR	Light host/device controller reset (OPTIONAL). Not implemented. Always 0.
6	IAA	Interrupt on async advance doorbell. Used as a doorbell by software to tell the USB DR controller to issue an interrupt the next time it advances asynchronous schedule. Software must write a 1 to this bit to ring the doorbell. When the controller has evicted all appropriate cached schedule states, it sets USBSTS[AAI]. If USBINTR[AAE] is set, the host controller will assert an interrupt at the next interrupt threshold. The controller clears this bit after it has set USBSTS[AAI]. Software should not set this bit when the asynchronous schedule is inactive. Doing so will yield undefined results. This bit is only used in host mode. Setting this bit when the USB DR module is in device mode is selected will result in undefined results.
5	ASE	Asynchronous schedule enable. Controls whether the controller skips processing the asynchronous schedule. Only used in host mode. 0 Do not process the asynchronous schedule 1 Use the ASYNCLISTADDR register to access the asynchronous schedule.
4	PSE	Periodic schedule enable. Controls whether the controller skips processing the periodic schedule. Only used in host mode. 0 Do not process the periodic schedule. 1 Use the PERIODICLISTBASE register to access the periodic schedule.
3–2	FS	Frame list size. Together with bit 15 these bits make the FS[2:0] field. This field is read/write only if programmable frame list flag in the HCCPARAMS registers is set to 1. This field specifies the size of the frame list that controls which bits in FRINDEX should be used for the frame list current index. Only used in host mode. Note that values below 256 elements are not defined in the EHCI specification. 000 1024 elements (4096 bytes) 001 512 elements (2048 bytes) 010 256 elements (1024 bytes) 011 128 elements (512 bytes) 100 64 elements (256 bytes) 101 32 elements (128 bytes) 110 16 elements (64 bytes) 111 8 elements (32 bytes)

**Table 16-10. USBCMD Register Field Descriptions (continued)**

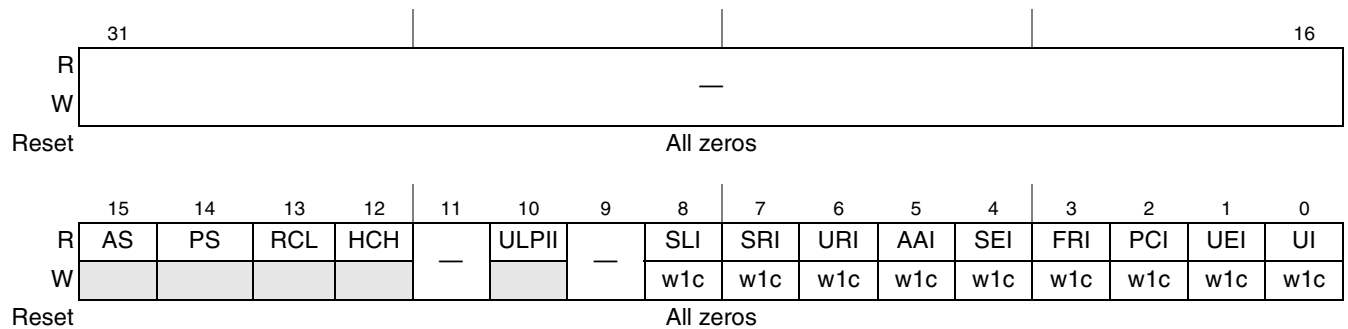
Bits	Name	Description
1	RST	<p>Controller reset. Software uses this bit to reset the controller. This bit is cleared by the controller when the reset process is complete. Software cannot terminate the reset process early by writing a zero to this register.</p> <p>Host mode:</p> <ul style="list-style-type: none"> <li>When software sets this bit, the host controller resets its internal pipelines, timers, counters, state machines etc. to their initial value. Any transaction currently in progress on USB is immediately terminated. A USB reset is not driven on downstream ports. Software should not set this bit when USBSTS[HCH] is a zero. Attempting to reset an actively running host controller will result in undefined behavior.</li> </ul> <p>Device mode:</p> <ul style="list-style-type: none"> <li>When software sets this bit, the USB DR controller resets its internal pipelines, timers, counters, state machines etc. to their initial value. Any transaction currently in progress on USB is immediately terminated. Writing a one to this bit in device mode is not recommended.</li> </ul>
0	RS	<p>Run/Stop.</p> <p>Host mode:</p> <ul style="list-style-type: none"> <li>When this bit is set, the controller proceeds with the execution of the schedule. The controller continues execution as long as this bit is set. When this bit is set to 0, the host controller completes the current transaction on the USB and then halts. The USBSTS[HCH] bit indicates when the USB DR controller has finished the transaction and has entered the stopped state. Software should not write a one to this field unless the controller is in the halted state (that is, USBSTS[HCH] is a one).</li> </ul> <p>Device mode:</p> <ul style="list-style-type: none"> <li>Setting this bit will cause the USB DR controller to enable a pull-up on D+ and initiate an attach event. This control bit is not directly connected to the pull-up enable, as the pull-up will become disabled upon transitioning into high-speed mode. Software should use this bit to prevent an attach event before the controller has been properly initialized. Clearing this bit will cause a detach event.</li> </ul> <p>0 Stop 1 Run</p>

### 16.3.2.2 USB Status Register (USBSTS)

This register indicates various states of the USB DR module and any pending interrupts. This register does not indicate status resulting from a transaction on the serial bus. Software clears certain bits in this register by writing a 1 to them (indicated by a w1c in the bit's W cell in [Figure 16-9](#)).

Offset 0x2\_3144

Access: Mixed



**Figure 16-9. USB Status Register (USBSTS)**

Table 16-11. USBSTS Register Field Descriptions

Bits	Name	Description
31–16	—	Reserved, should be cleared.
15	AS	Asynchronous schedule status. Reports the current real status of the asynchronous schedule. The USB DR controller is not required to immediately disable or enable the asynchronous schedule when software transitions USBCMD[ASE]. When this bit and USBCMD[ASE] have the same value, the asynchronous schedule is either enabled (1) or disabled (0). Only used in host mode. 0 Disabled 1 Enabled
14	PS	Periodic schedule status. Reports the current real status of the periodic schedule. The USB DR controller is not required to immediately disable or enable the periodic schedule when software transitions USBCMD[PSE]. When this bit and USBCMD[PSE] have the same value, the periodic schedule is either enabled (1) or disabled (0). Only used in host mode. 0 Disabled 1 Enabled
13	RCL	Reclamation. Used to detect an empty asynchronous schedule. Only used by the host mode. 0 Non-empty asynchronous schedule 1 Empty asynchronous schedule
12	HCH	HC halted. This bit is a zero whenever USBCMD[RS] is a one. The USB DR controller sets this bit to one after it has stopped executing because of USBCMD[RS] being cleared, either by software or by the host controller hardware (for example, internal error). Only used in host mode. 0 Running 1 Halted
11	—	Reserved, should be cleared.
10	ULPII	ULPI interrupt. An event completion to the viewport register sets this bit. If the ULPI enables the USBINTR[ULPIE] to be set, the USB interrupt (UI) will occur.
9	—	Reserved, should be cleared.
8	SLI	DCSuspend. This is a non-EHCI bit. When a device controller enters a suspend state from an active state, this bit is set. The device controller clears the bit upon exiting from a suspend state. Only used by the device controller. 0 Active 1 Suspended
7	SRI	Host mode: <ul style="list-style-type: none"> <li>This is a non-EHCI status bit. In host mode, this bit will be set every 125 us, provided the PHY clock is present and running (for example, the port is NOT suspended), and can be used by the host controller driver as a time base.</li> </ul> Device mode: <ul style="list-style-type: none"> <li>SOF received. When the USB DR controller detects a Start Of (Micro)Frame, this bit will be set. When a SOF is extremely late, the DR controller will automatically set this bit to indicate that an SOF was expected. Therefore, this bit will be set roughly every 1 msec in device FS mode and every 125 msec in HS mode and will be synchronized to the actual SOF that is received. Because the controller is initialized to FS before connect, this bit will be set at an interval of 1 msec during the prelude to the connect and chirp.</li> </ul> Software writes a 1 to this bit to clear it.

**Table 16-11. USBSTS Register Field Descriptions (continued)**

Bits	Name	Description
6	URI	USB reset received. This is a non-EHCI bit. When the USB DR controller detects a USB reset and enters the default state, this bit will be set. Software can write a 1 to this bit to clear the USB reset received status bit. Only used by the device mode. 0 No reset received 1 Reset received
5	AAI	Interrupt on async advance. System software can force the controller to issue an interrupt the next time the USB DR controller advances the asynchronous schedule by writing a one to USBCMD[IAA]. This status bit indicates the assertion of that interrupt source. Only used by the host mode. 0 No async advance interrupt 1 Async advance interrupt
4	SEI	System error. This bit is set whenever an error is detected on the system bus. If USBINTR[SEE] is set, an interrupt will be generated. The interrupt and status bits will remain asserted until cleared by writing a 1 to this bit. Additionally, when in host mode, USBCMD[RS] is cleared, effectively disabling the USB DR controller. For the USB DR controller in device mode, an interrupt is generated, but no other action is taken. 0 Normal operation 1 Error
3	FRI	Frame list rollover. The controller sets this bit to a one when the frame list index rolls over from its maximum value to zero. The exact value at which the rollover occurs depends on the frame list size. For example, if the frame list size (as programmed in USBCMD[FS]) is 1024, FRINDEX rolls over every time FRINDEX [1 3] toggles. Similarly, if the size is 512, the USB DR controller sets this bit to a one every time FHINDEX [12] toggles. Only used by the host mode.
2	PCI	Host mode: <ul style="list-style-type: none"> <li>Port change detect. The controller sets this bit when a connect status occurs on any port, a port enable/disable change occurs, an over current change occurs, or PORTSC[FPR] is set as the result of a J-K transition on the suspended port.</li> </ul> Device mode: <ul style="list-style-type: none"> <li>The USB DR controller sets this bit when it enters the full or high-speed operational state. When the it exits the full or high-speed operation states due to reset or suspend events, the notification mechanisms are USBSTS[URI] and USBSTS[SLI], respectively.</li> </ul> This bit is not EHCI compatible.
1	UEI (USBERRINT)	USB error interrupt (USBERRINT). When completion of a USB transaction results in an error condition, this bit is set by the controller. This bit is set along with the UI, if the TD on which the error interrupt occurred also had its interrupt on complete (IOC) bit set. See Section 4.15.1 in EHCI for a complete list of host error interrupt conditions. Also see <a href="#">Table 16-91</a> in this chapter for more information on device error matrix. For the USB DR controller in device mode, only resume signaling is detected, all others are ignored. 0 No error 1 Error detected
0	UI (USBINT)	USB interrupt (USBINT). This bit is set by the controller when the cause of an interrupt is a completion of a USB transaction where the transfer descriptor (TD) has an interrupt on complete (IOC) bit set. This bit is also set by the controller when a short packet is detected. A short packet is when the actual number of bytes received was less than the expected number of bytes.

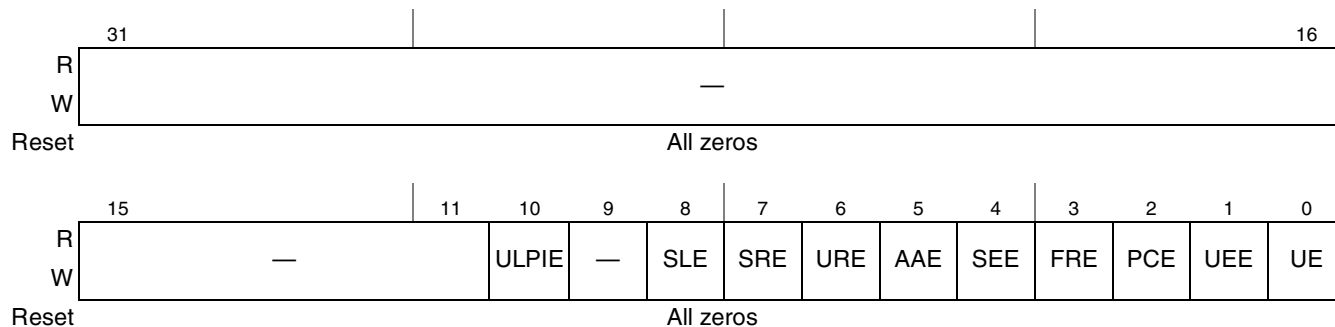


### 16.3.2.3 USB Interrupt Enable Register (USBINTR)

The interrupts to software are enabled with this register. An interrupt is generated when a bit is set and the corresponding interrupt is active. The USB status register (USBSTS) still shows interrupt sources even if they are disabled by the USBINTR register, allowing polling of interrupt events by the software.

Offset 0x2\_3148

Access: Read/Write



**Figure 16-10. USB Interrupt Enable (USBINTR)**

**Table 16-12. USBINTR Register Field Descriptions**

Bits	Name	Description
31–11	—	Reserved, should be cleared.
10	ULPIE	ULPI interrupt enable. An event completion to the viewport register sets the USBSTS[ULPII]. If the ULPI enables ULPIE bit to be set, then the USBINT (USBSTS[UI]) will occur. 0 Disable 1 Enable
9	—	Reserved, should be cleared.
8	SLE	Sleep enable. This is a non-EHCI bit. When this bit is a one, and USBSTS[SLI] transitions, the USB DR controller will issue an interrupt. The interrupt is acknowledged by software writing a one to USBSTS[SLI]. Only used in device mode. 0 Disable 1 Enable
7	SRE	SOF received enable. This is a non-EHCI bit. When this bit is a one, and USBSTS[SRI] is a one, the controller will issue an interrupt. The interrupt is acknowledged by software clearing USBSTS[SRI]. 0 Disable 1 Enable
6	URE	USB reset enable. This is a non-EHCI bit. When this bit is a one, USBSTS[URI] is a one, the device controller will issue an interrupt. The interrupt is acknowledged by software clearing USBSTS[URI] bit. Only used in device mode. 0 Disable 1 Enable
5	AAE	Interrupt on async advance enable. When this bit is a one, and USBSTS[AAI] is a one, the controller will issue an interrupt at the next interrupt threshold. The interrupt is acknowledged by software clearing USBSTS[AAI]. Only used in host mode. 0 Disable 1 Enable

Table 16-12. USBINTR Register Field Descriptions (continued)

Bits	Name	Description
4	SEE	System error enable. When this bit is a one, and USBSTS[SEI] is a one, the controller will issue an interrupt. The interrupt is acknowledged by software clearing USBSTS[SEI]. 0 Disable 1 Enable
3	FRE	Frame list rollover enable. When this bit is a one, and USBSTS[FRI] is a one, the controller will issue an interrupt. The interrupt is acknowledged by software clearing USBSTS[FRI]. Only used by the host mode. 0 Disable 1 Enable
2	PCE	Port change detect enable. When this bit is a one, and USBSTS[PCI] is a one, the controller will issue an interrupt. The interrupt is acknowledged by software clearing USBSTS[PCI]. 0 Disable 1 Enable
1	UEE	USB error interrupt enable. When this bit is a one, and USBSTS[UEI] is a one, the controller will issue an interrupt at the next interrupt threshold. The interrupt is acknowledged by software clearing USBSTS[UEI]. 0 Disable 1 Enable
0	UE	USB interrupt enable. When this bit is a one, and USBSTS[UI] is a one, the DR controller will issue an interrupt at the next interrupt threshold. The interrupt is acknowledged by software clearing USBSTS[UI]. 0 Disable 1 Enable

#### 16.3.2.4 Frame Index Register (FRINDEX)

In host mode, this register is used by the controller to index the periodic frame list. The register updates every 125 microseconds (once each microframe). Bits N-3 are used to select a particular entry in the periodic frame list during periodic schedule execution. The number of bits used for the index depends on the size of the frame list as set by system software in USBCMD[FS].

This register must be written as a DWord. Byte writes produce undefined results. This register cannot be written unless the USB DR controller is in the Halted state as indicated by the USBSTS[HCH]. A write to this register while USBCMD[RS] is set produces undefined results. Writes to this register also affect the SOF value.

In device mode, this register is read-only and, the USB DR controller updates the FRINDEX[13-3] register from the frame number indicated by the SOF marker. Whenever a SOF is received by the USB bus, FRINDEX[13-3] is checked against the SOF marker. If FRINDEX[13-3] is different from the SOF marker, FRINDEX[13-3] is set to the SOF value and FRINDEX[2-0] is cleared (that is, SOF for 1 msec frame). If FRINDEX[13-3] is equal to the SOF value, FRINDEX[2-0] is incremented (that is, SOF for 125- $\mu$ sec microframe.)

Offset 0x2\_314C

Access: Read/Write




Figure 16-11. USB Frame Index (FRINDEX)

Table 16-13. FRINDEX Register Field Descriptions

Bits	Name	Description
31–14	—	Reserved, should be cleared.
13–0	FRINDEX	Frame index. The value in this register increments at the end of each time frame (for example, microframe). Bits N–3 are used for the Frame List current index. This means that each location of the frame list is accessed 8 times (frames or microframes) before moving to the next index. In device mode, the value is the current frame number of the last frame transmitted. It is not used as an index. In either mode, bits 2–0 indicate the current microframe.

Table 16-14 illustrates values of N based on the value of the Frame List Size in the USBCMD register, when used in host mode.

Table 16-14. FRINDEX N Values

USBCMD[FS]	Frame List Size	FRINDEX N value
000	1024 elements (4096 bytes)	12
001	512 elements (2048 bytes)	11
010	256 elements (1024 bytes)	10
011	128 elements (512 bytes)	9
100	64 elements (256 bytes)	8
101	32 elements (128 bytes)	7
110	16 elements (64 bytes)	6
111	8 elements (32 bytes)	5

### 16.3.2.5 Control Data Structure Segment Register (CTRLDSSEGMENT)

The CTRLDSSEGMENT register is not implemented on the MPC8313E.

### 16.3.2.6 Periodic Frame List Base Address Register (PERIODICLISTBASE)

This register contains the beginning address of the Periodic Frame List in the system memory. The host controller driver loads this register prior to starting the schedule execution by the controller. The memory structure referenced by this physical memory pointer is assumed to be 4-Kbyte aligned. The contents of this register are combined with the frame index register (FRINDEX) to enable the controller to step through the Periodic Frame List in sequence.

Note that this register is shared between the host and device mode functions. In host mode, it is the PERIODICLISTBASE register; in device mode, it is the DEVICEADDR register. See Section 16.3.2.7, “Device Address Register (DEVICEADDR)—Non-EHCI,” for more information.

Offset 0x2\_3154

Access: Read/Write




Figure 16-12. Periodic Frame List Base Address (PERIODICLISTBASE)

Table 16-15. PERIODICLISTBASE Register Field Descriptions

Bits	Name	Description
31–12	PERBASE	Base address. Correspond to memory address signal [31:12]. Only used in the host mode.
11–0	—	Reserved, should be cleared.

### 16.3.2.7 Device Address Register (DEVICEADDR)—Non-EHCI

This register is not defined in the EHCI specification. In device mode, the upper seven bits of this register represent the device address. After any controller reset or a USB reset, the device address is set to the default address (0). The default address will match all incoming addresses. Software shall reprogram the address after receiving a SET\_ADDRESS descriptor.

Note that this register is shared between the host and device mode functions. In device mode, it is the DEVICEADDR register; in host mode, it is the PERIODICLISTBASE register. See Section 16.3.2.6, “Periodic Frame List Base Address Register (PERIODICLISTBASE),” for more information.

Offset 0x2\_3154

Access: Read/Write




Figure 16-13. Device Address (DEVICEADDR)

Table 16-16. DEVICEADDR Register Field Descriptions

Bits	Name	Description
31–25	USBADR	Device address. This field corresponds to the USB device address.
24–0	—	Reserved, should be cleared.


### 16.3.2.8 Current Asynchronous List Address Register (ASYNCLISTADDR)

This 32-bit register contains the address of the next asynchronous queue head to be executed by the host. Bits 4–0 of this register cannot be modified by the system software and always return zeros when read.

Note that this register is shared between the host and device mode functions. In host mode, it is the ASYNCLISTADDR register; in device mode, it is the ENDPOINTLISTADDR register. See [Section 16.3.2.9, “Endpoint List Address Register \(ENDPOINTLISTADDR\)—Non-EHCI,”](#) for more information.

Offset 0x2\_3158

Access: Read/Write



**Figure 16-14. Current Asynchronous List Address (ASYNCLISTADDR)**

**Table 16-17. ASYNCLISTADDR Register Field Descriptions**

Bits	Name	Description
31–5	ASYBASE	Link pointer low (LPL). These bits correspond to memory address signals [31:5]. This field may only reference a queue head (QH). Only used by the host controller.
4–0	—	Reserved, should be cleared.

### 16.3.2.9 Endpoint List Address Register (ENDPOINTLISTADDR)—Non-EHCI

This register is not defined in the EHCI specification. In device mode, this register contains the address of the top of the endpoint list in system memory. Bits 10–0 of this register cannot be modified by the system software and always return zeros when read. The memory structure referenced by this physical memory pointer is assumed to be 64-bytes. The queue head is actually a 48-byte structure, but must be aligned on 64-byte boundary. However, the ENDPOINTLISTADDR[EPBASE] has a granularity of 2 Kbytes, so in practice the queue head should be 2-Kbyte aligned.

Note that this register is shared between the host and device mode functions. In device mode, it is the ENDPOINTLISTADDR register; in host mode, it is the ASYNCLISTADDR register. See [Section 16.3.2.8, “Current Asynchronous List Address Register \(ASYNCLISTADDR\),”](#) for more information.

Offset 0x2\_3158

Access: Read/Write



**Figure 16-15. Endpoint List Address (ENDPOINTLISTADDR)**

**Table 16-18. ENDPOINTLISTADDR Register Field Descriptions**


Bits	Name	Description
31–11	EPBASE	Endpoint list address. Address of the top of the endpoint list.
10–0	—	Reserved, should be cleared.

### 16.3.2.10 Master Interface Data Burst Size Register (BURSTSIZE)—Non-EHCI

This register is not defined in the EHCI specification. This register is used to control and dynamically change the burst size used during data movement on the initiator (master) interface.

Offset 0x2\_3160

Access: Read/Write



**Figure 16-16. Master Interface Data Burst Size (BURSTSIZE)**

**Table 16-19. BURSTSIZE Register Field Descriptions**

Bits	Name	Description
31–16	—	Reserved, should be cleared.
15–8	TXPBURST	Programable TX burst length. This register represents the maximum length of a burst in 32-bit words while moving data from system memory to the USB bus. Must not be set to greater than 16.
7–0	RXPBURST	Programable RX burst length. This register represents the maximum length of a burst in 32-bit words while moving data from the USB bus to system memory. Must not be set to greater than 16.

### 16.3.2.11 Transmit FIFO Tuning Controls Register (TXFILLTUNING)—Non-EHCI

This register is not defined in the EHCI specification. This register is used to control and dynamically change the burst size used during data movement on device DMA transfers. It is only used in host mode.

The fields in this register control performance tuning associated with how the USB DR module posts data to the TX latency FIFO before moving the data onto the USB bus. The specific areas of performance include the how much data to post into the FIFO and an estimate for how long that operation should take in the target system.

Definitions:

$T_0$  = Standard packet overhead

$T_I$  = Time to send data payload

$T_s$  = Total Packet Flight Time (send-only) packet ( $T_s = T_0 + T_I$ )

$T_{ff}$  = Time to fetch packet into TX FIFO up to specified level.

$T_p$  = Total Packet Time (fetch and send) packet ( $T_p = T_{ff} + T_s$ )

Upon discovery of a transmit (OUT/SETUP) packet in the data structures, host controller checks to ensure  $T_p$  remains before the end of the [micro]frame. If so it proceeds to pre-fill the TX FIFO. If at any time during the pre-fill operation the time remaining the [micro]frame is  $< T_s$  then the packet attempt ceases and the packet is tried at a later time. Although this is not an error condition and the module eventually recovers, a mark is made in the scheduler health counter to note the occurrence of a back-off event. When a back-off event is detected, the partial packet fetched may need to be discarded from the latency buffer to make room for periodic traffic that will begin after the next SOF. Too many back-off events can waste bandwidth and power on the system bus and thus should be minimized (not necessarily eliminated). Back-offs can be minimized with use of the TSCHEALTH ( $T_{ff}$ ) parameter described below.

Offset 0x2\_3164

Access: Read/Write




Figure 16-17. Transmit FIFO Tuning Controls (TXFILLTUNING)

Table 16-20. TXFILLTUNING Register Field Descriptions

Bits	Name	Description
31–22	—	Reserved, should be cleared.
21–16	TXFIFOTHRES	FIFO burst threshold. Control the number of data bursts that are posted to the TX latency FIFO in host mode before the packet begins on to the bus. The minimum value is 2 and this value should be a low as possible to maximize USB performance. A higher value can be used in systems with unpredictable latency and/or insufficient bandwidth where the FIFO may underrun because the data transferred from the latency FIFO to USB occurs before it can be replenished from system memory. This value is ignored if USBMODE[SDIS] (stream disable bit) is set. When USBMODE[SDIS] is set, the host controller behaves as if TXFIFOTHRES is set to the maximum value.
15–13	—	Reserved, should be cleared.
12–8	TXSCHHEALTH	Scheduler health counter. Increment when the host controller fails to fill the TX latency FIFO to the level programmed by TXFIFOTHRES before running out of time to send the packet before the next Start-Of-Frame. This health counter measures the number of times this occurs to provide feedback to selecting a proper TXSCHOH. Writing to this register clears the counter and this counter stops counting after reaching the maximum of 31.

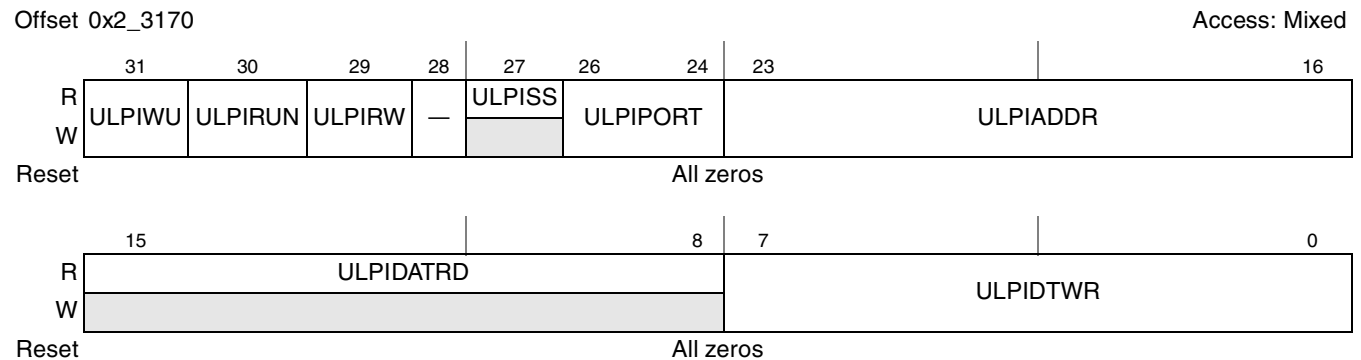
**Table 16-20. TXFILLTUNING Register Field Descriptions (continued)**

Bits	Name	Description
7–0	TXSCHOH	<p>Scheduler overhead. These bits add an additional fixed offset to the schedule time estimator described above as <math>T_{ff}</math>. As an approximation, the value chosen for this register should limit the number of back-off events captured in the TXSCHHEALTH to less than 10 per second in a highly utilized bus. Choosing a value that is too high for this register is not desired as it can needlessly reduce USB utilization.</p> <p>The time unit represented in this register is 1.267<math>\mu</math>s when a device is connected in high-speed mode.</p> <p>The time unit represented in this register is 6.333<math>\mu</math>s when a device is connected in low-/full-speed mode.</p> <p>For most applications, TXSCHOH can be set to 4 or less. A good value to begin with is: <math>TXFIFOTHRES \times (BURSTSIZE \times 4 \text{ bytes-per-word}) \div (40 \times TimeUnit)</math>, always rounded to the next higher integer. <i>TimeUnit</i> is either 1.267 or 6.333 as noted earlier in this description. For example, if TXFIFOTHRES is 5 and BURSTSIZE is 8, then set TXSCHOH to <math>5 \times (8 \times 4) \div (40 \times 1.267) = 4</math> for a high-speed link. If this value of TXSCHOH results in a TXSCHHEALTH count of 0 per second, try lowering the value by 1 if optimizing performance is desired. If TXSCHHEALTH exceeds 10 per second, try raising the value by 1.</p> <p>If streaming mode is disabled via the USBMODE register, treat TXFIFOTHRES as the maximum value for purposes of the TXSCHOH calculation.</p>

### 16.3.2.12 ULPI Register Access (ULPI VIEWPORT)

The register provides indirect access to the ULPI PHY register set. Although the controller modules perform access to the ULPI PHY register set, there may be extraordinary circumstances where software may need direct access. Be advised that writes to the ULPI through the ULPI viewport can substantially harm standard USB operations. Currently no usage model has been defined where software should need to execute writes directly to the ULPI. Note that executing read operations through the ULPI viewport should have no harmful side effects to standard USB operations. Also note that if the ULPI interface is not enabled, this register will always read zeros.

ULPI VIEWPORT is shown in [Figure 16-18](#).



**Figure 16-18. ULPI Register Access (ULPI VIEWPORT)**



Table 16-21. ULPI VIEWPORT Field Descriptions

Bits	Name	Description
31	ULPIWU	ULPI Wake Up. Writing 1 to this bit begins the wakeup operation. This bit automatically transitions to 0 after the wakeup is complete. Once this bit is set, it can not be cleared by software. <b>Note:</b> The driver must never execute a wakeup and a read/write operation at the same time.
30	ULPIRUN	ULPI Run. Writing 1 to this bit begins a read/write operation. This bit automatically transitions to 0 after the read/write is complete. Once this bit is set, it can not be cleared by software. <b>Note:</b> The driver must never execute a wakeup and a read/write operation at the same time.
29	ULPIRW	This bit selects between running a read or write operation to the ULPI. 0 Read 1 Write
28	—	Reserved, should be cleared.
27	ULPISS	This bit represents the state of the ULPI interface. Before reading this bit, the ULPIPORT field should be set accordingly if used with the multi-port host. Otherwise, this field should always remain 0. 0 Any other state (that is, carkit, serial, low power). 1 Normal Sync State.
26–24	ULPIPORT	For wakeup or read/write operations this value selects the port number to which the ULPI PHY is attached. Valid values are 0 and 1.
23–16	ULPIADDR	When a read or write operation is commanded, the address of the operation is written to this field.
15–8	ULPIDATRD	After a read operation completes, the result is placed in this field.
7–0	ULPIDTWR	When a write operation is commanded, the data to be sent is written to this field.

There are two operations that can be performed with the ULPI viewport, wakeup and read /write operations. The wakeup operation is used to put the ULPI interface into normal operation mode and re-enable the clock if necessary. A wakeup operation is required before accessing the registers when the ULPI interface is operating in low power mode, serial mode, or carkit mode. The ULPI state can be determined by reading the sync state bit (ULPISS). If this bit is set, then the ULPI interface is running in normal operation mode and can accept read/write operations. If the ULPISS is cleared, then read/write operations will not be able execute. Undefined behavior results if a read or write operation is performed when ULPISS is cleared. To execute a wakeup operation, write all 32-bits of the ULPI Viewport where ULPIPORT is constructed appropriately and the ULPIWU bit is set and the ULPIRUN bit is cleared. Poll the ULPI Viewport until ULPIWU is cleared for the operation to complete.

To execute a read or write operation, write all 32-bits of the ULPI Viewport where ULPIDATWR, ULPIADDR, ULPIPORT, ULPIRW are constructed appropriately and the ULPIRUN bit is set. Poll the ULPI Viewport until ULPIRUN is cleared for the operation to complete. For read operations, ULPIDATRD is valid once ULPIRUN is cleared.

The polling method above can be replaced with interrupts using the ULPI interrupt defined in the USBSTS and USBINTR registers. When a wakeup or read/write operation completes, the ULPI interrupt is set.

### 16.3.2.13 Configure Flag Register (CONFIGFLAG)

This EHCI register is not used in this implementation. A read from this register returns a constant of a 0x0000\_0001 to indicate that all port routings default to this host controller.




Figure 16-19. Configure Flag Register (CONFIGFLAG)

Table 16-22. CONFIGFLAG Register Field Descriptions

Bits	Name	Description
31–1	—	Reserved.
0	CF	Configure flag. Always 1 indicating all port routings default to this host.

### 16.3.2.14 Port Status and Control Register (PORTSC)

The port status and control (PORTSC) register is only reset when power is initially applied or in response to a controller reset. The initial conditions of a port are:

- No device connected
- Port disabled

If the port has port power control, this state remains until software applies power to the port by setting port power to one.

In device mode, the USB DR controller does not support power control. Port control in device mode is only used for status port reset, suspend, and current connect status. It is also used to initiate test mode or force signaling and allows software to put the PHY into low power suspend mode and disable the PHY clock.




Figure 16-20. Port Status and Control (PORTSC)

Table 16-23. PORTSC Register Field Descriptions

Bits	Name	Description
31–30	PTS	Port transceiver select. This register bit is used to control which parallel transceiver interface is selected. 00 UTMI parallel interface 01 Reserved, should be cleared 10 ULPI parallel interface 11 Reserved This bit is not defined in the EHCI specification.
29	—	Reserved, should be cleared
28	—	Reserved
27–26	PSPD	Port speed. This read-only register field indicates the speed at which the port is operating. This bit is not defined in the EHCI specification. 00 Full-speed 01 Low-speed 10 High-speed 11 Undefined
25	—	Reserved, should be cleared
24	PFSC	Port force full-speed connect. Used to disable the chirp sequence that allows the port to identify itself as a HS port. This is useful for testing FS configurations with a HS host, hub or device. 0 Allow the port to identify itself as high speed. 1 Force the port to only connect at full speed. This bit is not defined in the EHCI specification. This bit is for debugging purposes.
23	PHCD	PHY low power suspend. This bit is not defined in the EHCI specification. Host mode: <ul style="list-style-type: none"> <li>The PHY can be put into low power suspend—when the downstream device has been put into suspend mode or when no downstream device is connected. Low power suspend is completely under the control of software.</li> </ul> Device mode: <ul style="list-style-type: none"> <li>The PHY can be put into low power suspend—when the device is not running (USBCMD[RS] = 0b) or suspend signaling is detected on the USB. Low power suspend will be cleared automatically when the resume signaling has been detected or when forcing port resume.</li> </ul> 0 Normal PHY operation. 1 Signal the PHY to enter low power suspend mode Reading this bit indicates the status of the PHY. <b>Note:</b> If there is no clock connected to the USBDR_CLK signals, PHCD must be set and the following registers should not be written: DEVICE_ADDR/PERIODICLISTBASE, PORTSC, ENDPTCTRL0, ENDPTCTRL1, ENDPTCTRL2.
22	WKOC	Wake on over-current enable. Writing this bit to a one enables the port to be sensitive to over-current conditions as wake-up events. This field is zero if Port Power (PP) is zero. This bit is (OTG/host mode only) for use by an external power control circuit.
21	WKDS	Wake on disconnect enable. Writing this bit to a one enables the port to be sensitive to device disconnects as wake-up events. This field is zero if Port Power(PP) is zero or in device mode. This bit is (OTG/host mode only) for use by an external power control circuit.

**Table 16-23. PORTSC Register Field Descriptions (continued)**

Bits	Name	Description
20	WLCN	Wake on connect enable. Writing this bit to a one enables the port to be sensitive to device connects as wake-up events. This field is zero if Port Power(PP) is zero or in device mode. This bit is (OTG/host mode only) for use by an external power control circuit.
19–16	PTC	Port test control. Any other value than zero indicates that the port is operating in test mode. 0000 Not Enabled 0001 J_STATE 0010 K_STATE 0011 SEQ_NAK 0100 Packet 0101 FORCE_ENABLE 0110–1111 Reserved, should be cleared Refer to Chapter 7 of the USB Specification Revision 2.0 [3] for details on each test mode.
15–14	PIC	Port indicator control. Control the link indicator signals. These signals are valid for host mode only. 00 Off 01 Amber 10 Green 11 Undefined Refer to the USB Specification Revision 2.0 [3] for a description on how these bits are to be used. This field is output from the module on the USB port control signals for use by an external LED driving circuit.
13	PO	Port owner. Unconditionally goes to a 0 when the configured bit in the CONFIGFLAG register makes a 0 to 1 transition. This bit unconditionally goes to 1 whenever the Configured bit is zero. System software uses this field to release ownership of the port to a selected the module (in the event that the attached device is not a high-speed device). Software writes a one to this bit when the attached device is not a high-speed device. A one in this bit means that an internal companion controller owns and controls the port. Port owner hand-off is not implemented in this design, therefore this bit is always 0.
12	PP	Port power. Represents the current setting of the switch (0=off, 1=on). When power is not available on a port (that is, PP equals a 0), the port is non-functional and will not report attaches, detaches, etc. When an over-current condition is detected on a powered port, the PP bit in each affected port is transitioned by the host controller driver from a one to a zero (removing power from the port). This feature is implemented in the host/OTG controller (PPC = 1). In a device-only implementation port power control is not necessary, thus PPC and PP = 0.
11–10	LS	Line status. Reflect the current logical levels of the USB D+ (bit 11) and D– (bit 10) signal lines. The use of line status by the host controller driver is not necessary (unlike EHCI), because the connection of FS and LS is managed by hardware. 00 SE0 10 J-state 01 K-state 11 Undefined
9	—	Reserved, should be cleared

Table 16-23. PORTSC Register Field Descriptions (continued)

Bits	Name	Description
8	PR	<p>Port reset.</p> <p>Host mode:</p> <ul style="list-style-type: none"> <li>When software writes a one to this bit the bus-reset sequence as defined in the USB Specification Revision 2.0 is started. This bit will automatically change to zero after the reset sequence is complete. This behavior is different from EHCI where the host controller driver is required to set this bit to a zero after the reset duration is timed in the driver.</li> </ul> <p>Device mode:</p> <ul style="list-style-type: none"> <li>This bit is a read only status bit. Device reset from the USB bus is also indicated in the USBSTS register.</li> </ul> <p>1 Port is in reset. 0 Port is not in reset. This field is zero if Port Power(PP) is zero.</p>
7	SUSP	<p>Suspend.</p> <p>Host mode:</p> <ul style="list-style-type: none"> <li>The port enabled bit (PE) and suspend (SUSP) bit define the port states as follows:</li> </ul> <p>0x Disable 10 Enable 11 Suspend</p> <ul style="list-style-type: none"> <li>When in suspend state, downstream propagation of data is blocked on this port, except for port reset. The blocking occurs at the end of the current transaction if a transaction was in progress when this bit was written to 1. In the suspend state, the port is sensitive to resume detection. Note that the bit status does not change until the port is suspended and that there may be a delay in suspending a port if there is a transaction currently in progress on the USB.</li> <li>The module unconditionally sets this bit to zero when software clears the FPR bit. A write of zero to this bit is ignored by the host controller. If host software sets this bit to a one when the port is not enabled (that is, port enabled bit is a zero) the results are undefined.</li> <li>This field is zero if Port Power (PP) is zero in host mode.</li> </ul> <p>Device mode:</p> <p>1 Port in suspend state. 0 Port not in suspend state. Default. In device mode this bit is a read-only status bit.</p>
6	FPR	<p>Force port resume. This bit is not-EHCI compatible.</p> <p>1 Resume detected/driven on port. 0 No resume (K-state) detected/driven on port.</p> <p>Host mode:</p> <ul style="list-style-type: none"> <li>Software sets this bit to one to drive resume signaling. The controller sets this bit to one if a J-to-K transition is detected while the port is in the Suspend state. When this bit transitions to a one a J-to-K transition is detected, USBSTS[PCI] (port change detect) is also set. This bit will automatically change to zero after the resume sequence is complete. This behavior is different from EHCI where the host controller driver is required to set this bit to a zero after the resume duration is timed in the driver.</li> <li>Note that when the controller owns the port, the resume sequence follows the defined sequence documented in the USB Specification Revision 2.0. The resume signaling (Full-speed 'K') is driven on the port as long as this bit remains a one. This bit will remain a one until the port has switched to the high-speed idle. Writing a zero has no affect because the port controller will time the resume operation clear the bit the port control state switches to HS or FS idle.</li> <li>This field is zero if Port Power (PP) is zero in host mode.</li> </ul> <p>Device mode:</p> <ul style="list-style-type: none"> <li>After the device has been in Suspend State for 5 msec or more, software must set this bit to one to drive resume signaling before clearing. The USB DR controller will set this bit to one if a J-to-K transition is detected while the port is in the Suspend state. The bit will be cleared when the device returns to normal operation. Also, when this bit transitions to a one because a J-to-K transition detected, USBSTS[PCI] is also set.</li> </ul>

**Table 16-23. PORTSC Register Field Descriptions (continued)**

Bits	Name	Description
5	OCC	Over-current change. This bit gets set when there is a change to over-current active. Software clears this bit by writing a one to this bit position. Host/OTG mode: <ul style="list-style-type: none"> <li>The user can provide over-current detection to the USB<sub>n</sub>_PWRFAULT signal for this condition.</li> </ul> Device mode: <ul style="list-style-type: none"> <li>This bit must always be 0.</li> </ul> 1 Over current detect. 0 No over current.
4	OCA	Over-current active. This bit will automatically transition from one to zero when the over current condition is removed. Host/OTG mode: <ul style="list-style-type: none"> <li>The user can provide over-current detection to the USB<sub>n</sub>_PWRFAULT signal for this condition.</li> </ul> Device mode: <ul style="list-style-type: none"> <li>This bit must always be 0.</li> </ul> 1 Port currently in over-current condition. 0 Port not in over-current condition.
3	PEC	Port enable/disable change. For the root hub, this bit gets set only when a port is disabled due to disconnect on the port or due to the appropriate conditions existing at the EOF2 point (See Chapter 11 of the USB Specification). Software clears this by writing a one to it.[ In device mode: <ul style="list-style-type: none"> <li>The device port is always enabled. (This bit will be zero.)</li> </ul> 1 Port disabled. 0 No change. This field is zero if Port Power(PP) is zero.
2	PE	Port enabled/disabled. Host mode: <ul style="list-style-type: none"> <li>Ports can only be enabled by the controller as a part of the reset and enable. Software cannot enable a port by writing a one to this field. Ports can be disabled by either a fault condition (disconnect event or other fault condition) or by the host software. Note that the bit status does not change until the port state actually changes. There may be a delay in disabling or enabling a port due to other host and bus events.</li> <li>When the port is disabled, (0) downstream propagation of data is blocked except for reset.</li> <li>This field is zero if Port Power(PP) is zero in host mode.</li> </ul> Device mode: <ul style="list-style-type: none"> <li>The device port is always enabled. (This bit will be one.)</li> </ul>

**Table 16-23. PORTSC Register Field Descriptions (continued)**

Bits	Name	Description
1	CSC	<p>Connect change status.</p> <p>Host mode:</p> <ul style="list-style-type: none"> <li>This bit indicates a change has occurred in the port's Current Connect Status. the controller sets this bit for all changes to the port device connect status, even if system software has not cleared an existing connect status change. For example, the insertion status changes twice before system software has cleared the changed condition, hub hardware will be 'setting' an already-set bit (i.e., the bit will remain set). Software clears this bit by writing a one to it.</li> </ul> <p>1 Connect Status has changed. 0 No change.</p> <ul style="list-style-type: none"> <li>This field is zero if Port Power(PP) is zero.</li> </ul> <p>Device mode:</p> <ul style="list-style-type: none"> <li>This bit is undefined.</li> </ul>
0	CCS	<p>Current connect status.</p> <p>Host mode:</p> <p>1 Device is present 0 No device present.</p> <p>This field is zero if Port Power(PP) is zero in host mode.</p> <p>In device mode:</p> <p>1 Attached 0 Not attached.</p> <p>A one indicates that the device successfully attached and is operating in either high-speed or full-speed as indicated by the High Speed Port bit in this register. A zero indicates that the device did not attach successfully or was forcibly disconnected by the software writing a zero to USBCMD[RS] (run bit). It does not state the device being disconnected or suspended.</p>

### 16.3.2.15 On-The-Go Status and Control (OTGSC)—Non-EHCI

This register is not defined in the EHCI specification. The USB DR module implements one On-The-Go (OTG) status and control register corresponding to Port 0.

The OTGSC register has four sections:

- OTG interrupt enables (Read/Write)
- OTG interrupt status (Read/Write to Clear)
- OTG status inputs (Read Only)
- OTG controls (Read/Write)

The status inputs are de-bounced using a 1-msec time constant. Values on the status inputs that do not persist for more than 1 msec will not cause an update of the status inputs, or cause an OTG interrupt.

Offset 0x2\_31A4

Access: Mixed

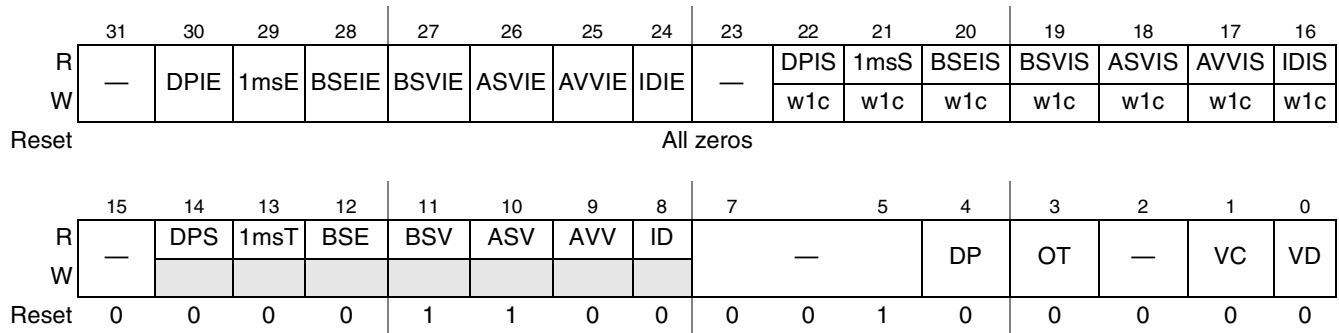


Figure 16-21. OTG Status Control (OTGSC)

Table 16-24. OTGSC Register Field Descriptions

Bits	Name	Description
31	—	Reserved, should be cleared.
30	DPIE	Data pulse interrupt enable 1 Enable 0 Disable
29	1msE	1-millisecond timer Interrupt enable 1 Enable 0 Disable
28	BSEIE	B session end interrupt enable 1 Enable 0 Disable
27	BSVIE	B session valid interrupt enable 1 Enable 0 Disable
26	ASVIE	A session valid interrupt enable 1 Enable 0 Disable
25	AVVIE	A VBus valid interrupt enable 1 Enable 0 Disable
24	IDIE	USB ID interrupt enable. 1 Enable 0 Disable
23	—	Reserved, should be cleared.
22	DPIS	Data pulse interrupt status. Set when data bus pulsing occurs on DP or DM. Data bus pulsing is only detected when USBMODE[CM] = Host (11) and PORTSC[PP] (port power) = Off (0). Software must write a one to clear this bit.
21	1msS	1-millisecond timer interrupt status. Set once every millisecond. Software must write a one to clear this bit.
20	BSEIS	B session end interrupt status. Set when VBus has fallen below the B session end threshold. Software must write a one to clear this bit.



Table 16-24. OTGSC Register Field Descriptions (continued)

Bits	Name	Description
19	BSVIS	B session valid interrupt status. Set when VBus has either risen above or fallen below the B session valid threshold (0.8 VDC). Software must write a one to clear this bit.
18	ASVIS	A session valid interrupt status. Set when VBus has either risen above or fallen below the A session valid threshold (0.8 VDC). Software must write a one to clear this bit.
17	AVVIS	A VBus valid interrupt status. Set when VBus has either risen above or fallen below the VBus valid threshold (4.4 VDC) on an A device. Software must write a one to clear this bit.
16	IDIS	USB ID interrupt status. Set when a change on the ID input has been detected. Software must write a one to clear this bit.
15	—	Reserved, should be cleared.
14	DPS	Data bus pulsing status 1 Pulsing detected on port 0 No pulsing on port
13	1msT	1 millisecond timer toggle. This bit toggles once per millisecond.
12	BSE	B session end 1 VBus is below the B session end threshold. 0 VBus is above the B session end threshold.
11	BSV	B session valid 1 VBus is above the B session valid threshold. 0 VBus is below the B session valid threshold.
10	ASV	A session valid 1 VBus is above the A session valid threshold. 0 VBus is below the A session valid threshold.
9	AVV	A VBus valid 1 VBus is above the A VBus valid threshold. 0 VBus is below the A VBus valid threshold.
8	ID	USB ID 1 B device 0 A device
7–5	—	Reserved, writes should preserve reset value.
4	DP	Data pulsing 1 The pullup on DP is asserted for data pulsing during SRP. 0 The pullup on DP is not asserted.
3	OT	OTG termination. This bit must be set when the OTG device is in device mode. 1 Enable pulldown on DM 0 Disable pulldown on DM
2	—	Reserved, should be cleared.
1	VC	VBUS charge. Setting this bit causes the VBus line to be charged. This is used for VBus pulsing during SRP.
0	VD	VBUS discharge. Setting this bit causes VBus to discharge through a resistor.

### 16.3.2.16 USB Mode Register (USBMODE)—Non-EHCI

This register is not defined in the EHCI specification. This register controls the operating mode of the module.





Figure 16-22. USB Mode (USBMODE)

Table 16-25. USBMODE Register Field Descriptions

Bits	Name	Description
31–5	—	Reserved, should be cleared.
4	SDIS	Stream disable In host mode, setting this bit ensures that overruns/underruns of the latency FIFO are eliminated for low bandwidth systems where the RX and TX buffers are sufficient to contain the entire packet. Enabling stream disable also has the effect of ensuring the TX latency is filled to capacity before the packet is launched onto the USB. Note that time duration to pre-fill the FIFO becomes significant when stream disable is active. See TXFILLTUNING to characterize the adjustments needed for the scheduler when using this feature. Also note that in systems with high system bus utilization, setting this bit will ensure no overruns or underruns during operation, at the expense of link utilization. For those who desire optimal link performance, SDIS can be left clear, and the rules used under the description of the TXFILLTUNING register to limit underruns/overruns. 1 Active. 0 Inactive. In device mode, setting this bit disables double priming on both RX and TX for low bandwidth systems. This mode ensures that when the RX and TX buffers are sufficient to contain an entire packet that the standard double buffering scheme is disabled to prevent overruns/underruns in bandwidth limited systems. Note that in high-speed mode, all packets received will be responded to with a NYET handshake when stream disable is active.
3	SLOM	Setup lockout mode. In device mode, this bit controls behavior of the setup lock mechanism. See <a href="#">Section 16.8.3.5, “Control Endpoint Operation Model.”</a> 1 Setup lockouts off. DCD requires use of setup data buffer tripwire in USBCMD (SUTW). 0 Setup lockouts on
2	—	Reserved, should be cleared.
1–0	CM	Controller mode This register can only be written once after reset. If it is necessary to switch modes, software must reset the controller by writing to USBCMD[RST] before reprogramming this register. 00 Idle (default for combination host/device). 01 Reserved, should be cleared. 10 Device controller (default for device only controller). 11 Host controller (default for host only controller). Defaults to the idle state and needs to be initialized to the desired operating mode after reset.

### 16.3.2.17 Endpoint Setup Status Register (ENDPTSETUPSTAT)—Non-EHCI

This register is not defined in the EHCI specification. This register contains the endpoint setup status. It is only used in device mode.




**Figure 16-23. Endpoint Setup Status (ENDPTSETUPSTAT)**

**Table 16-26. ENDPTSETUPSTAT Register Field Descriptions**

Bits	Name	Description
31–3	—	Reserved, should be cleared.
2–0	ENDPTSETUP STAT	Setup endpoint status. For every setup transaction that is received, a corresponding bit in this register is set. Software must clear or acknowledge the setup transfer by writing a one to a respective bit after it has read the setup data from queue head. The response to a setup packet as in the order of operations and total response time is crucial to limit bus time outs while the setup lockout mechanism is engaged. This register is only used in device mode.

### 16.3.2.18 Endpoint Initialization Register (ENDPTPRIME)—Non-EHCI

This register is not defined in the EHCI specification. This register is used to initialize endpoints. It is only used in device mode.



**Figure 16-24. Endpoint Initialization (ENDPTPRIME)**

**Table 16-27. ENDPTPRIME Register Field Descriptions**


Bits	Name	Description
31–19	—	Reserved, should be cleared.
18–16	PETB	Prime endpoint transmit buffer. For each endpoint a corresponding bit is used to request that a buffer prepared for a transmit operation in order to respond to a USB IN/INTERRUPT transaction. Software should write a one to the corresponding bit when posting a new transfer descriptor to an endpoint. Hardware will automatically use this bit to begin parsing for a new transfer descriptor from the queue head and prepare a transmit buffer. Hardware will clear this bit when the associated endpoint(s) is (are) successfully primed. PETB[2] (bit 18 of the register) corresponds to endpoint 2. Note that these bits will be momentarily set by hardware during hardware re-priming operations when a dTD is retired, and the dQH is updated.

**Table 16-27. ENDPTPRIME Register Field Descriptions (continued)**

Bits	Name	Description
15–3	—	Reserved, should be cleared.
2–0	PERB	Prime endpoint receive buffer. For each endpoint, a corresponding bit is used to request a buffer prepare for a receive operation in order to respond to a USB OUT transaction. Software should write a one to the corresponding bit whenever posting a new transfer descriptor to an endpoint. Hardware will automatically use this bit to begin parsing for a new transfer descriptor from the queue head and prepare a receive buffer. Hardware will clear this bit when the associated endpoint(s) is (are) successfully primed. PERB[2] corresponds to endpoint 2. Note that these bits will be momentarily set by hardware during hardware re-priming operations when a dTD is retired, and the dQH is updated.

### 16.3.2.19 Endpoint Flush Register (ENDPTFLUSH)—Non-EHCI

This register is not defined in the EHCI specification. This register is only used in device mode.




**Figure 16-25. Endpoint Flush (ENDPTFLUSH)**

**Table 16-28. ENDPTFLUSH Register Field Descriptions**

Bits	Name	Description
31–19	—	Reserved, should be cleared.
18–16	FETB	Flush endpoint transmit buffer. Writing a one to a bit(s) in this register will cause the associated endpoint(s) to clear any primed buffers. If a packet is in progress for one of the associated endpoints, then that transfer will continue until completion. Hardware will clear this register after the endpoint flush operation is successful. FETB[2] (bit 18 of the register) corresponds to endpoint 2.
15–3	—	Reserved, should be cleared.
2–0	FERB	Flush endpoint receive buffer. Writing a one to a bit(s) will cause the associated endpoint(s) to clear any primed buffers. If a packet is in progress for one of the associated endpoints, then that transfer will continue until completion. Hardware will clear this register after the endpoint flush operation is successful. FERB[2] corresponds to endpoint 2.

### 16.3.2.20 Endpoint Status Register (ENDPTSTATUS)—Non-EHCI

This register is not defined in the EHCI specification. This register is only used in device mode.



**Figure 16-26. Endpoint Status (ENDPTSTATUS)**

**Table 16-29. ENDPTSTATUS Register Field Descriptions**


Bits	Name	Description
31–19	—	Reserved, should be cleared
18–16	ETBR	Endpoint transmit buffer ready. One bit for each endpoint indicates status of the respective endpoint buffer. This bit is set by the hardware as a response to receiving a command from a corresponding bit in the ENDPTPRIME register. There will always be a delay between setting a bit in the ENDPTPRIME register and endpoint indicating ready. This delay time varies based upon the current USB traffic and the number of bits set in the ENDPTPRIME register. Buffer ready is cleared by USB reset, by the USB DMA system, or through the ENDPTFLUSH register. ETBR[2] (bit 18 of the register) corresponds to endpoint 2. Note that these bits will be momentarily cleared by hardware during hardware endpoint re-priming operations when a dTD is retired, and the dQH is updated.
15–3	—	Reserved, should be cleared
2–0	ERBR	Endpoint receive buffer ready. One bit for each endpoint indicates status of the respective endpoint buffer. This bit is set by the hardware as a response to receiving a command from a corresponding bit in the ENDPTPRIME register. There will always be a delay between setting a bit in the ENDPTPRIME register and endpoint indicating ready. This delay time varies based upon the current USB traffic and the number of bits set in the ENDPTPRIME register. Buffer ready is cleared by USB reset, by the USB DMA system, or through the ENDPTFLUSH register. ERBR[2] corresponds to endpoint 2. Note that these bits will be momentarily cleared by hardware during hardware endpoint re-priming operations when a dTD is retired, and the dQH is updated.

### 16.3.2.21 Endpoint Complete Register (ENDPTCOMPLETE)—Non-EHCI

This register is not defined in the EHCI specification. This register is only used in device mode.

Offset 0x2\_31BC

Access: w1c



**Figure 16-27. Endpoint Complete (ENDPTCOMPLETE)**

**Table 16-30. ENDPTCOMPLETE Register Field Descriptions**

Bits	Name	Description
31–19	—	Reserved, should be cleared
18–16	ETCE	Endpoint transmit complete event. Each bit indicates a transmit event (IN/INTERRUPT) occurred and software should read the corresponding endpoint queue to determine the endpoint status. If the corresponding IOC bit is set in the Transfer Descriptor, then this bit will be set simultaneously with the USBINT. Writing a one will clear the corresponding bit in this register. ETCE[2] (bit 18 of the register) corresponds to endpoint 2.
15–3	—	Reserved, should be cleared
2–0	ERCE	Endpoint receive complete event. Each bit indicates a received event (OUT/SETUP) occurred and software should read the corresponding endpoint queue to determine the transfer status. If the corresponding IOC bit is set in the Transfer Descriptor, then this bit will be set simultaneously with the USBINT. Writing a one will clear the corresponding bit in this register. ERCE[2] corresponds to endpoint 2.

### 16.3.2.22 Endpoint Control Register 0 (ENDPTCTRL0)—Non-EHCI

This register is not defined in the EHCI specification. Every device will implement endpoint 0 as a control endpoint.

Offset 0x2\_31C0

Access: Mixed





Figure 16-28. Endpoint Control 0 (ENDPTCTRL0)

Table 16-31. ENDPTCTRL0 Register Field Descriptions

Bits	Name	Description
31–24	—	Reserved, should be cleared.
23	TXE	TX endpoint enable. Endpoint zero is always enabled. 0 Disable 1 Enable
22–20	—	Reserved, should be cleared.
19–18	TXT	TX endpoint type. Endpoint zero is always a control endpoint (00).
17	—	Reserved, should be cleared.
16	TXS	TX endpoint stall. Software can write a one to this bit to force the endpoint to return a STALL handshake to the Host. It will continue returning STALL until the bit is cleared by software or it will automatically be cleared upon receipt of a new SETUP request. 1 Endpoint stalled 0 Endpoint OK
15–8	—	Reserved, should be cleared.
7	RXE	RX endpoint enable. Endpoint zero is always enabled. 0 Disabled 1 Enabled
6–4	—	Reserved, should be cleared.
3–2	RXT	RX endpoint type. Endpoint zero is always a control endpoint (00).
1	—	Reserved, should be cleared.
0	RXS	RX endpoint stall Software can write a one to this bit to force the endpoint to return a STALL handshake to the host. It will continue returning STALL until the bit is cleared by software or it will automatically be cleared upon receipt of a new SETUP request. 1 Endpoint stalled 0 Endpoint OK

### 16.3.2.23 Endpoint Control Register $n$ (ENDPTCTRL $n$ )—Non-EHCI

These registers are not defined in the EHCI specification. There is an ENDPTCTRL $n$  register of each endpoint in a device.



**Figure 16-29. Endpoint Control 1 to 5 (ENDPTCTRL $n$ )**

**Table 16-32. ENDPTCTRL $n$  Register Field Descriptions**


Bits	Name	Description
31–24	—	Reserved, should be cleared
23	TXE	TX endpoint enable 0 Disabled 1 Enabled
22	TXR	TX data toggle reset. Whenever a configuration event is received for this endpoint, software must write a one to this bit in order to synchronize the data PID's between the Host and device.
21	TXI	TX data toggle inhibit. Used only for test and should always be written as zero. Writing a one to this bit will cause this endpoint to ignore the data toggle sequence and always transmit DATA0 for a data packet. 0 PID sequencing enabled 1 PID sequencing disabled
20	—	Reserved, should be cleared
19–18	TXT	TX endpoint type 00 Control 01 Isochronous 10 Bulk 11 Interrupt <b>Note:</b> When only one endpoint (RX or TX, but not both) of an endpoint pair is used, the unused endpoint should be configured as a bulk type endpoint.
17	TXD	TX endpoint data source. This bit should always be written as 0, which selects the dual port memory/DMA engine as the source.
16	TXS	TX endpoint stall. This bit will be set automatically upon receipt of a SETUP request if this endpoint is not configured as a control endpoint. It will be cleared automatically upon receipt of a SETUP request if this endpoint is configured as a control endpoint. Software can write a one to this bit to force the endpoint to return a STALL handshake to the host. It will continue to returning STALL until this bit is either cleared by software or automatically cleared as above. 0 Endpoint OK 1 Endpoint stalled
15–8	—	Reserved, should be cleared
7	RXE	RX endpoint enable 0 Disabled 1 Enabled

**Table 16-32. ENDPCTRL $n$  Register Field Descriptions (continued)**

Bits	Name	Description
6	RXR	RX data toggle reset. Whenever a configuration event is received for this endpoint, software must write a one to this bit in order to synchronize the data PID's between the Host and device.
5	RXI	RX data toggle inhibit. This bit is only used for test and should always be written as zero. Writing a one to this bit will cause this endpoint to ignore the data toggle sequence and always accept data packets regardless of their data PID. 1 PID sequencing enabled 0 PID sequencing disabled
4	—	Reserved, should be cleared
3–2	RXT	RX endpoint type 00 Control 01 Isochronous 10 Bulk 11 Interrupt <b>Note:</b> When only one endpoint (RX or TX, but not both) of an endpoint pair is used, the unused endpoint should be configured as a bulk type endpoint.
1	RXD	RX endpoint data sink. This bit should always be written as 0, which selects the dual port memory/DMA engine as the sink.
0	RXS	RX endpoint stall. This bit will be set automatically upon receipt of a SETUP request if this endpoint is not configured as a control endpoint. It will be cleared automatically upon receipt a SETUP request if this endpoint is configured as a control endpoint, Software can write a one to this bit to force the endpoint to return a STALL handshake to the host. It will continue to returning STALL until this bit is either cleared by software or automatically cleared as above, 1 Endpoint stalled 0 Endpoint OK

### 16.3.2.24 SNOOP1 and SNOOP2—Non-EHCI

Note that these registers use big-endian byte ordering and are not defined in the EHCI specification. The SNOOP1 and SNOOP2 registers provide snooping control and address range selection function. Transactions that hit a snooping window will generate cache coherent transactions on the internal CSB bus. When the five lower bits (SNOOP $n$ [27–31]) are equal to 00000, snooping is always disabled on the CSB for all DMA transfers. When SNOOP $n$ [27–31] is 01011 through 11110, the twenty upper bits (SNOOP $n$ [0–19]) provide the starting base address for which transactions are snooped. These twenty bits are compared to the twenty upper bits of the address provided by the DMA block of the USB controller. When a match occurs, the five lower bits are decoded as shown below. This provides a snooping region of 4 Kbytes to 2 Gbytes within each starting base address that is programmed by the core. The SNOOP $n$ [20–26] are not used.



**Figure 16-30. Snoop 1 and Snoop 2 (SNOOP $n$ )**



Table 16-33. SNOOP<sub>n</sub> Register Field Descriptions

Bits	Name	Description
0–19	Snoop address	The starting base address for which transactions are snooped.
20–26	—	Reserved, should be cleared
27–31	Snoop Enables	0x00 Snooping disabled 0x0B 4-Kbyte snoop range starting at the value defined by SNOOP <sub>n</sub> [0–19] 0x0C 8-Kbyte snoop range starting at the value defined by SNOOP <sub>n</sub> [0–18] 0x0D 16-Kbyte snoop range starting at the value defined by SNOOP <sub>n</sub> [0–17] 0x0E 32-Kbyte snoop range starting at the value defined by SNOOP <sub>n</sub> [0–16] 0x0F 64-Kbyte snoop range starting at the value defined by SNOOP <sub>n</sub> [0–15] 0x10 128-Kbyte snoop range starting at the value defined by SNOOP <sub>n</sub> [0–14] 0x11 256-Kbyte snoop range starting at the value defined by SNOOP <sub>n</sub> [0–13] 0x12 512-Kbyte snoop range starting at the value defined by SNOOP <sub>n</sub> [0–12] 0x13 1-Mbyte snoop range starting at the value defined by SNOOP <sub>n</sub> [0–11] 0x14 2-Mbyte snoop range starting at the value defined by SNOOP <sub>n</sub> [0–10] 0x15 4-Mbyte snoop range starting at the value defined by SNOOP <sub>n</sub> [0–9] 0x16 8-Mbyte snoop range starting at the value defined by SNOOP <sub>n</sub> [0–8] 0x17 16-Mbyte snoop range starting at the value defined by SNOOP <sub>n</sub> [0–7] 0x18 32-Mbyte snoop range starting at the value defined by SNOOP <sub>n</sub> [0–6] 0x19 64-M byte snoop range starting at the value defined by SNOOP <sub>n</sub> [0–5] 0x1A 31-Mbyte snoop range starting at the value defined by SNOOP <sub>n</sub> [0–4] 0x1B 256-Mbyte snoop range starting at the value defined by SNOOP <sub>n</sub> [0–3] 0x1C 512-Mbyte snoop range starting at the value defined by SNOOP <sub>n</sub> [0–2] 0x1D 1-Gbyte snoop range starting at the value defined by SNOOP <sub>n</sub> [0–1] 0x1E 2-Gbyte snoop range starting at the value defined by SNOOP <sub>n</sub> [0]

### 16.3.2.25 Age Count Threshold Register (AGE\_CNT\_THRESH)—Non-EHCI

Note that this register uses big-endian byte ordering and is not defined in the EHCI specification. The age count threshold (AGE\_CNT\_THRESH) register provides the aging counter threshold value used to determine the priority state of the USB DR controller's internal system interface. This is used to increase the priority state of the module's system interface from zero to one. The actual priority level on the system bus for each state is defined by the PRI\_CTRL register. See [Section 6.3.1.1, “Address Bus Arbitration with PRIORITY\[0:1\],”](#) for more details on bus priority. The threshold value is in units of *csb\_clk* cycles. This register should be written during system initialization or during normal system operation when the system bus interface is idle. It can be read at any time.

If the aging counter is less than the AGE\_CNT\_THRESH value, priority state zero is chosen. If the aging counter is greater than or equal to the AGE\_CNT\_THRESH value, priority state one is chosen.

The aging counter begins to count from zero when a bus access is requested. It increments every bus cycle until the bus transaction completes. At the completion of a bus transaction, the counter is synchronously reset to zero. If there are any outstanding bus requests, the aging counter will then begin counting immediately.

The AGE\_CNT\_THRESH is compared against the value of the aging counter during each clock cycle of the current transaction. If AGE\_CNT\_THRESH is equal to zero, priority state one is always chosen. If the aging counter is less than the AGE\_CNT\_THRESH value, priority state zero is selected. If the aging counter is greater than or equal to the AGE\_CNT\_THRESH value, priority state one is selected.

The two priority states of the aging counter function each have corresponding register bits which are programmed by the CPU. Thus, when the aging counter function is at priority state zero, PRI\_CTRL[30–31] are selected and used to drive bus priority levels. When the aging counter function is at priority state one, PRI\_CTRL[28–29] are selected and used to drive the priority.




Figure 16-31. Age Count Threshold (AGE\_CNT\_THRESH)

Table 16-34. AGE\_CNT\_THRESH Register Field Descriptions

Bits	Name	Description
0–17	—	Reserved, should be cleared
18–31	Threshold	Aging counter threshold value.

The setting of AGE\_CNT\_THRESH is highly dependent on both the mix of other controllers operating on the system bus as well as the kind of traffic moving through the USB controller. A recommended approach is first to try leaving the aging mechanism disabled and see if the USB meets performance requirements. If USB performance does not meet application requirements, try the following settings:

- Set PRI\_CTRL[pri\_lv10] to 0.
- Set tPRI\_CTRL[pri\_lv11] to 3.
- Set AGE\_CNT\_THRESH to 40.

This combination works for a wide variety of applications. If this combination still does not meet application requirements, try lowering AGE\_CNT\_THRESH by 5. On the contrary, the setting 40 may be too conservative for some applications. If USB performance is acceptable at 40, try raising the value in increments of 5. Raising AGE\_CNT\_THRESH benefits the other controllers on the system bus by reducing the frequency that this USB controller raises its priority to the arbiter.

### 16.3.2.26 Priority Control Register (PRI\_CTRL)—Non-EHCI

Note that this register uses big-endian byte ordering and is not defined in the EHCI specification. The priority control (PRI\_CTRL) register sets the priority level for each of two priority states. The priority state is determined by the value programmed in the AGE\_CNT\_THRESH register and the number of *csb\_clk* cycles that a particular transaction takes to complete.




Figure 16-32. Priority Control (PRI\_CTRL)

Table 16-35. PRI\_CTRL Register Field Descriptions

Bits	Name	Description
0–27	—	Reserved, should be cleared
28–29	pri_lv1	Priority level for priority state 1.
30–31	pri_lv0	Priority level for priority state 0.

### 16.3.2.27 System Interface Control Register (SI\_CTRL)—Non-EHCI

Note that this register uses big-endian byte ordering and is not defined in the EHCI specification. The system interface control register (SI\_CTRL) controls various functions pertaining to the internal system interface.

Offset 0x2\_3410

Access: Read/Write




Figure 16-33. System Interface Control Register (SI\_CTRL)

Table 16-36. SI\_CTRL Register Field Descriptions

Bits	Name	Description
0–26	—	Reserved, should be cleared
27	err_disable	When this bit is set, it causes the controller to ignore system bus errors. If cleared the controller responds according to the values set in USBSTS[SEI] and USBINT[SEE]. 0 enable 1 disable
28–30	—	Reserved, should be cleared
31	rd_prefetch_val	Selects whether 32 bytes or 64 bytes are fetched during burst read transactions at the system interface. When this input is LOW 64 bytes are fetched and when it is HIGH 32 bytes are fetched. The setting of rd_prefetch_val must match the setting of the larger of TXPBURST and RXPBURST fields in the BURSTSIZE register. If either of these fields is 64 bytes, then rd_prefetch_val must be left cleared. Otherwise, this value should be set. 0 64-byte fetch 1 32-byte fetch

### 16.3.2.28 USB General Purpose Register (CONTROL)—Non-EHCI

Note that this register uses big-endian byte ordering and is not defined in the EHCI specification. The USB general purpose (CONTROL) register contains the general-purpose IP control register outputs and is shown in [Figure 16-34](#).

Offset 0x2\_3500

Access: Mixed




Figure 16-34. USB General-Purpose Register (CONTROL)

Table 16-37. CONTROL Field Descriptions

Bits	Name	Description
0–13	—	Reserved, must be cleared.
14	PHY_CLK_VALID	Indicates whether the PHY clock is valid (read only). When in UTMI mode, this bit reflects the value of the UTMI PHY ClkValid signal. In ULPI mode, this bit reflects the inverted ULPI DIR. In ULPI mode, this bit is not valid if the USB I/O have not been configured and after the USB_EN signal is asserted. 0 USB PHY clock is not valid 1 USB PHY clock is valid
15	WU_INT	Reflects the state of the wake up interrupt. The wake up interrupt signal is asserted when a wake-up event occurs while in a low-power suspend state. If WU_INT_EN is set, this WU_INT signal generates an interrupt to the system to indicate wake up servicing is required. WU_INT will remain set until the USB controller is exited from the low power by clearing the PORTSC[PHCD] bit. 0 Normal operation or Low Power mode waiting for wakeup event 1 Low power wakeup event has occurred
16–18	—	Reserved, must be cleared.
19–20	CLKIN_SEL[1:0]	Select the clock source for the UTMI PHY PLL reference clock. The reference clock can be sourced from the USB_CLK or divide by 1 or 2 of the SYS_CLK. These bits are not relevant when in ULPI mode. 00 Reference clock is USB_CLK 01 Reference clock is USB_CLK 10 Reference clock is SYS_CLK 11 Reference clock is SYS_CLK divided by 2
21	PHY_CLK_SEL	Select the source of the USB link controller transceiver clock. When cleared the UTMI PHY is the source of the clock. When set, the clock is sourced from the external ULPI PHY. 0 UTMI is clock source 1 ULPI is clock source
22	UTMI_PHY_EN	Enable the UTMI PHY. The UTMI PHY is reset when placed in the disable mode. 0 UTMI PHY disabled 1 UTMI PHY enabled
23	PLL_RESET	Reset the UTMI PHY PLL. This bit is not self clearing and must be cleared to complete the reset sequence. 0 UTMI PHY in normal operating state 1 Put UTMI PHY in reset state

Table 16-37. CONTROL Field Descriptions (continued)

Bits	Name	Description
24–25	REFSEL[1:0]	The REFSEL signals are used to set the frequency value for the UTMI PLL reference clock. These bit fields are not relevant if not in UTMI mode. 00 Reserved 01 Reference clock frequency is 24 MHz. This is the default frequency. 10 Reference clock frequency is 48 MHz 11 Reserved
26	OTG_PORT	Enables the OTG comparators. 0 OTG comparators disabled 1 OTG comparators enabled
27	KEEP_OTG_ON	Keeps the OTG comparators on during low power suspend. 0 OTG comparators disabled during suspend 1 OTG comparators enabled during suspend
28	LSF_EN	This bit is used to enable the UTMI line state filter. When enabled the UTMI linestate[1:0] output of the UTMI PHY are filtered to account for any skew between the USB differential data lines (DP/DM). 0 Line state filter disabled 1 Line state filter enabled
29	USB_EN	UTMI mode: This bit is used to enable the USB interface. It must be set before setting RS bit in USB CMD register. 1 Enable 0 Disable ULPI mode: In safe mode, all USB interface signals are put into input mode or driven inactive, except for SUSPEND_STP which is driven high. Also, the input signal DIR is forced to appear high to the controller. This prevents any start-up problems that otherwise could occur if the PHY and the controller take significantly different times to complete power-on reset. 1 Normal operation 0 Safe mode
30	WU_INT_EN	This bit is used to mask/unmask the system wakeup interrupt signal 0 System wakeup interrupt disabled 1 System wakeup interrupt enabled <b>Note:</b> PORTSC[PHCD] bit must be set for the system wakeup interrupt generation.
31	ULPI_INT_EN	Used to enable the ULPI low power wakeup interrupt from the PHY when the PHY is in low power mode only. 0 ULPI low power wakeup interrupt disabled 1 ULPI low power wakeup interrupt enabled <b>Note:</b> PORTSC[PHCD] bit must be set

## 16.4 Functional Description

The USB DR module can be broken down into functional sub-blocks, which are described below.

### 16.4.1 System Interface

The system interface block contains all the control and status registers that allow a processor to interface to the USB DR module. These registers allow the processor to control the configuration of the module,

ascertain the capabilities of the module, and control the module’s operation. It also has registers to control snoopability and priority of the DMA interface.

### 16.4.2 DMA Engine

The module contains a local DMA engine. The DMA engine interfaces internally to the CSB. It is responsible for moving all of the data to be transferred over the USB between the module and buffers in system memory. Like the system interface block, the DMA engine block uses a simple synchronous bus signaling protocol that eases connections to a number of different standard buses.

The DMA controller must access both control information and packet data from system memory. The control information is contained in link list–based queue structures. The DMA controller has state machines that are able to parse data structures defined in the EHCI specification. In host mode, the data structures are EHCI compliant and represent queues of transfers to be performed by the host controller, including the split-transaction requests that allow an EHCI controller to direct packets to FS and LS devices. In device mode, the data structures are designed to be similar to those in the EHCI specification and are used to allow device responses to be queued for each of the active pipes in the device.

### 16.4.3 FIFO RAM Controller

The FIFO RAM controller is used for context information and to control FIFOs between the protocol engine and the DMA controller. These FIFOs decouple the system processor/memory bus requests from the extremely tight timing required by USB.

The use of the FIFO buffers differs between host and device mode operation. In host mode, a single data channel is maintained in each direction through the buffer memory. In device mode, multiple FIFO channels are maintained for each of the active endpoints in the system.

In host mode, the USB DR module uses a 512-byte Tx buffer and a 512-byte Rx buffer. Device operation uses a single 512-byte Rx buffer and a 512-byte Tx buffer for each endpoint. The 512-byte buffers allow the module to buffer a complete HS bulk packet.

### 16.4.4 PHY Interface

The USB DR module interfaces to any UTMI- or ULPI-compatible PHY. The primary function of the port controller block is to isolate the rest of the module from the transceiver, and to move all of the transceiver signaling into the primary clock domain of the module. This allows the module to run synchronously with the system processor and its associated resources.

Due to pincount limitations the module only supports certain combinations of PHY interfaces and USB functionality. Refer to [Table 16-38](#) for more information.

**Table 16-38. Supported PHY Interfaces**

PHY	Function
UTMI	Host/Device
ULPI	Host/Device/OTG

## 16.5 Host Data Structures

This section defines the interface data structures used to communicate control, status, and data between HCD (software) and the Enhanced Host Controller (hardware). The data structure definitions in this section support a 32-bit memory buffer address space. The interface consists of a periodic schedule, periodic frame list, asynchronous schedule, isochronous transaction descriptors, split-transaction isochronous transfer descriptors, queue heads, and queue element transfer descriptors.

The periodic frame list is the root of all periodic (isochronous and interrupt transfer type) support for the host controller interface. The asynchronous list is the root for all the bulk and control transfer type support. Isochronous data streams are managed using isochronous transaction descriptors. Isochronous split-transaction data streams are managed with split-transaction isochronous transfer descriptors. All interrupt, control, and bulk data streams are managed with queue heads and queue element transfer descriptors. These data structures are optimized to reduce the total memory footprint of the schedule and to reduce (on average) the number of memory accesses needed to execute a USB transaction.

Note that software must ensure that no interface data structure reachable by the EHCI host controller spans a 4K-page boundary.

The data structures defined in this section are (from the host controller's perspective) a mix of read-only and read/writable fields. The host controller must preserve the read-only fields on all data structure writes.

### 16.5.1 Periodic Frame List

[Figure 16-35](#) shows the organization of the periodic schedule. This schedule is for all periodic transfers (isochronous and interrupt). The periodic schedule is referenced from the operational registers space using the PERIODICLISTBASE address register and the FRINDEX register. The periodic schedule is based on an array of pointers called the periodic frame list. The PERIODICLISTBASE address register is combined with the FRINDEX register to produce a memory pointer into the frame list. The periodic frame list implements a sliding window of work over time.

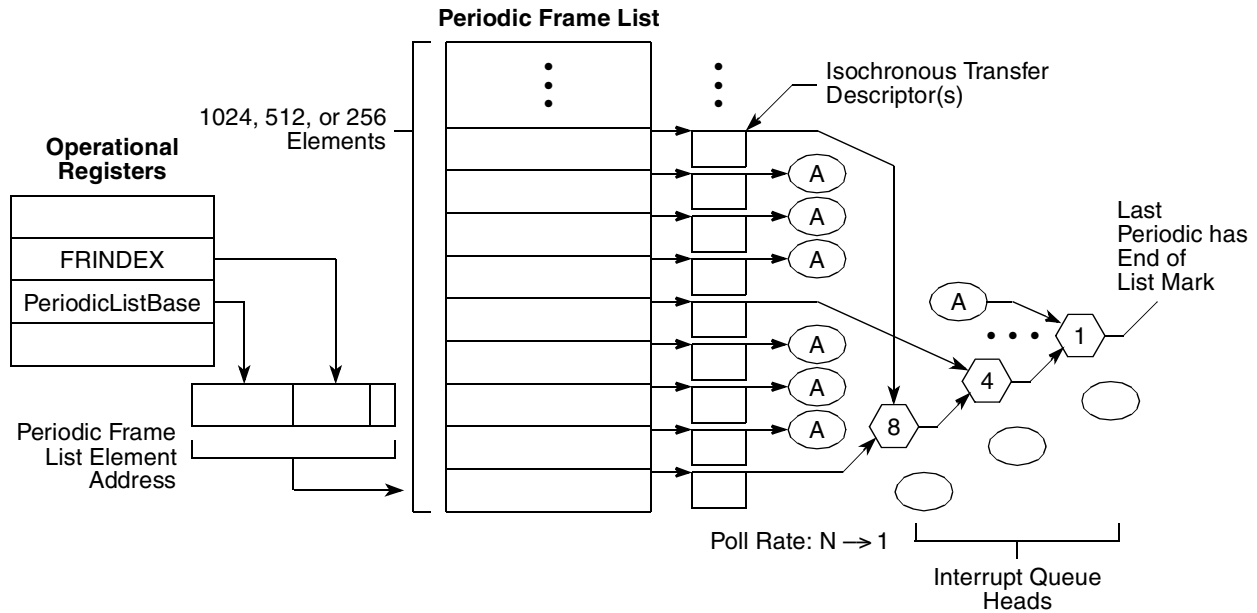


Figure 16-35. Periodic Schedule Organization

Split transaction interrupt, bulk and control are also managed using queue heads and queue element transfer descriptors.

The periodic frame list is a 4K-page aligned array of frame list link pointers. The length of the frame list may be programmable. The programmability of the periodic frame list is exported to system software through the HCCPARAMS register. If non-programmable, the length is 1024 elements. If programmable, the length can be selected by system software as one of 256, 512, or 1024 elements. An implementation must support all three sizes. Programming the size (that is, the number of elements) is accomplished by system software writing the appropriate value into frame list size field in the USBCMD register.

Frame list link pointers direct the host controller to the first work item in the frame’s periodic schedule for the current micro-frame. The link pointers are aligned on DWord boundaries within the frame list.

Figure 16-36 shows the format for the frame list link pointer.




Figure 16-36. Frame List Link Pointer Format

Frame list link pointers always reference memory objects that are 32-byte aligned. The referenced object may be an isochronous transfer descriptor for high-speed devices, a split-transaction isochronous transfer descriptor (for full-speed isochronous endpoints), or a queue head (used to support high-, full- and low-speed interrupt). System software should not place non-periodic schedule items into the periodic schedule. The least-significant bits in a frame list pointer are used to key the host controller in as to the type of object the pointer is referencing.

The least-significant bit is the T bit (bit 0). When this bit is set, the host controller never uses the value of the frame list pointer as a physical memory pointer. The Typ field indicates the exact type of data structure being referenced by this pointer. The value encodings for the Typ field are given in Table 16-39.




**Table 16-39. Typ Field Encodings**

Typ	Description
00	Isochronous transfer descriptor
01	Queue head
10	Split transaction isochronous transfer descriptor
11	Frame span traversal node

### 16.5.2 Asynchronous List Queue Head Pointer

The asynchronous transfer list (based at the ASYNCLISTADDR register) is where all the control and bulk transfers are managed. Host controllers use this list only when it reaches the end of the periodic list, the periodic list is disabled, or the periodic list is empty.



**Figure 16-37. Asynchronous Schedule Organization**

The asynchronous list is a simple circular list of queue heads. The ASYNCLISTADDR register is simply a pointer to the next queue head. This implements a pure round-robin service for all queue heads linked into the asynchronous list.

### 16.5.3 Isochronous (High-Speed) Transfer Descriptor (iTd)

Figure 16-38 illustrates the format of an isochronous transfer descriptor. This structure is used only for high-speed isochronous endpoints. All other transfer types should use queue structures. Isochronous TDs must be aligned on a 32-byte boundary.

31	30	29	28	27	26	25	24	23	22	21	20	19	18	17	16	15	14	13	12	11	10	9	8	7	6	5	4	3	2	1	0	offset
Next Link Pointer																											00	Typ	T	0x00		
Status <sup>1</sup>		Transaction 0 Length <sup>1</sup>										ioc	PG <sup>2</sup>	Transaction 0 Offset <sup>2</sup>										0x04								
Status <sup>1</sup>		Transaction 1 Length <sup>1</sup>										ioc	PG <sup>2</sup>	Transaction 1 Offset <sup>2</sup>										0x08								
Status <sup>1</sup>		Transaction 2 Length <sup>1</sup>										ioc	PG <sup>2</sup>	Transaction 2 Offset <sup>2</sup>										0x0C								
Status <sup>1</sup>		Transaction 3 Length <sup>1</sup>										ioc	PG <sup>2</sup>	Transaction 3 Offset <sup>2</sup>										0x10								
Status <sup>1</sup>		Transaction 4 Length <sup>1</sup>										ioc	PG <sup>2</sup>	Transaction 4 Offset <sup>2</sup>										0x14								
Status <sup>1</sup>		Transaction 5 Length <sup>1</sup>										ioc	PG <sup>2</sup>	Transaction 5 Offset <sup>2</sup>										0x18								
Status <sup>1</sup>		Transaction 6 Length <sup>1</sup>										ioc	PG <sup>2</sup>	Transaction 6 Offset <sup>2</sup>										0x1C								
Status <sup>1</sup>		Transaction 7 Length <sup>1</sup>										ioc	PG <sup>2</sup>	Transaction 7 Offset <sup>2</sup>										0x20								
Buffer Pointer (Page 0)											EndPt	R	Device Address										0x24									
Buffer Pointer (Page 1)											I/O	Maximum Packet Size										0x28										
Buffer Pointer (Page 2)											Reserved										Mult	0x2C										
Buffer Pointer (Page 3)											Reserved										0x30											
Buffer Pointer (Page 4)											Reserved										0x34											
Buffer Pointer (Page 5)											Reserved										0x38											
Buffer Pointer (Page 6)											Reserved										0x3C											

Figure 16-38. Isochronous Transaction Descriptor (iTD)

<sup>1</sup> Host controller read/write; all others read-only.

<sup>2</sup> These fields may be modified by the host controller if the I/O field indicates an OUT.

### 16.5.3.1 Next Link Pointer

The first DWord of an iTD is a pointer to the next schedule data structure.

Table 16-40. Next Schedule Element Pointer

Bits	Name	Description
31–5	Link Pointer	Correspond to memory address signals [31:5], respectively. This field points to another isochronous transaction descriptor (iTD/siTD) or queue head (QH).
4–3	—	Reserved, should be cleared. These bits are reserved and their value has no effect on operation. Software should initialize this field to zero.
2–1	Typ	Indicates to the host controller whether the item referenced is an iTD, siTD or a QH. This allows the host controller to perform the proper type of processing on the item after it is fetched. Value encodings are: 00 iTD (isochronous transfer descriptor) 01 QH (queue head) 10 siTD (split transaction isochronous transfer descriptor) 11 FSTN (frame span traversal node)
0	T	Terminate 1 Link Pointer field is not valid. 0 Link Pointer field is valid.

### 16.5.3.2 iTD Transaction Status and Control List

DWords 1–8 constitute eight slots of transaction control and status. Each transaction description includes:

- Status results field
- Transaction length (bytes to send for OUT transactions and bytes received for IN transactions).
- Buffer offset. The PG and Transaction *n* Offset fields are used with the buffer pointer list to construct the starting buffer address for the transaction.

The host controller uses the information in each transaction description, plus the endpoint information contained in the first three DWords of the buffer page pointer list, to execute a transaction on the USB.

**Table 16-41. iTD Transaction Status and Control**

Bits	Name	Description
31–28	Status	Records the status of the transaction executed by the host controller for this slot. This field is a bit vector with the following encoding: 31 Active. Set by software to enable the execution of an isochronous transaction by the host controller. When the transaction associated with this descriptor is completed, the host controller clears this bit indicating that a transaction for this element should not be executed when it is next encountered in the schedule. 30 Data buffer error. Set by the host controller during status update to indicate that the host controller is unable to keep up with the reception of incoming data (overrun) or is unable to supply data fast enough during transmission (underrun). If an overrun condition occurs, no action is necessary. 29 Babble detected. Set by the host controller during status update when "babble" is detected during the transaction generated by this descriptor. 28 Transaction error (XactErr). Set by the host controller during status update in the case where the host did not receive a valid response from the device (Time-out, CRC, Bad PID, etc.). This bit may only be set for isochronous IN transactions.
27–16	Transaction <i>n</i> Length	For an OUT, this field is the number of data bytes the host controller will send during the transaction. The host controller is not required to update this field to reflect the actual number of bytes transferred during the transfer. For an IN, the initial value of the endpoint to deliver. During the status update, the host controller writes back the field is the number of bytes the host expects the number of bytes successfully received. The value in this register is the actual byte count (for example, 0 zero length data, 1 one byte, 2 two bytes, etc.). The maximum value this field may contain is 0xC00 (3072).
15	ioc	Interrupt on complete. If this bit is set, it specifies that when this transaction completes, the host controller should issue an interrupt at the next interrupt threshold.
14–12	PG	These bits are set by software to indicate which of the buffer page pointers the offset field in this slot should be concatenated to produce the starting memory address for this transaction. The valid range of values for this field is 0 to 6.
11–0	Transaction <i>n</i> Offset	This field is a value that is an offset, expressed in bytes, from the beginning of a buffer. This field is concatenated onto the buffer page pointer indicated in the adjacent PG field to produce the starting buffer address for this transaction.

### 16.5.3.3 iTD Buffer Page Pointer List (Plus)

DWords 9–15 of an isochronous transaction descriptor are nominally page pointers (4K aligned) to the data buffer for this transfer descriptor. This data structure requires the associated data buffer to be contiguous (relative to virtual memory), but allows the physical memory pages to be non-contiguous. Seven page pointers are provided to support the expression of eight isochronous transfers. The seven pointers allow

for 3 (transactions) × 1024 (maximum packet size) × 8 (transaction records) = 24576 bytes to be moved with this data structure, regardless of the alignment offset of the first page.

Since each pointer is a 4K-aligned page pointer, the least-significant 12 bits in several of the page pointers are used for other purposes.

**Table 16-42. Buffer Pointer Page 0 (Plus)**

Bits	Name	Description
31–12	Buffer Pointer (Page 0)	A 4K-aligned pointer to physical memory. Corresponds to memory address bits 31–12.
11–8	EndPt	Selects the particular endpoint number on the device serving as the data source or sink.
7	—	Reserved, should be cleared. Reserved for future use and should be initialized by software to zero.
6–0	Device Address	This field selects the specific device serving as the data source or sink.

**Table 16-43. iTD Buffer Pointer Page 1 (Plus)**

Bits	Name	Description
31–12	Buffer Pointer (Page 1)	This is a 4K aligned pointer to physical memory. Corresponds to memory address bits 31–12.
11	I/O	Direction (I/O). This field encodes whether the high-speed transaction should use an IN or OUT PID. 0 OUT 1 IN
10–0	Maximum Packet Size	This directly corresponds to the maximum packet size of the associated endpoint ( <i>wMaxPacketSize</i> ). This field is used for high-bandwidth endpoints where more than one transaction is issued per transaction description (for example, per micro-frame). This field is used with the <i>Multi</i> field to support high-bandwidth pipes. This field is also used for all IN transfers to detect packet babble. Software should not set a value larger than 1024 (0x400). Any value larger yields undefined results.

**Table 16-44. Buffer Pointer Page 2 (Plus)**

Bits	Name	Description
31–12	Buffer Pointer (Page 2)	This is a 4K-aligned pointer to physical memory. Corresponds to memory address bits 31–12.
11–2	—	Reserved, should be cleared. This bit reserved for future use and should be cleared.
1–0	Mult	Indicates to the host controller the number of transactions that should be executed per transaction description (for example, per micro-frame). 00 Reserved, should be cleared. A zero in this field yields undefined results. 01 One transaction to be issued for this endpoint per micro-frame 10 Two transactions to be issued for this endpoint per micro-frame 11 Three transactions to be issued for this endpoint per micro-frame

**Table 16-45. Buffer Pointer Page 3–6**

Bits	Name	Description
31–12	Buffer Pointer	This is a 4K aligned pointer to physical memory. Corresponds to memory address bits 31–12.
11–2	—	Reserved, should be cleared. These bits reserved for future use and should be cleared.

## 16.5.4 Split Transaction Isochronous Transfer Descriptor (siTD)

All full-speed isochronous transfers through the internal transaction translator are managed using the siTD data structure. This data structure satisfies the operational requirements for managing the split transaction protocol.

		31	30	29	28	27	26	25	24	23	22	21	20	19	18	17	16	15	14	13	12	11	10	9	8	7	6	5	4	3	2	1	0	offset
Next Link Pointer																00	Typ	T	0x00															
I/O	Port Number	0	Hub Address				0000	EndPt				0	Device Address						0x04															
0000_0000_0000_00000										μFrame C-mask						μFrame S-mask						0x08												
ioc	P <sup>1</sup>	0000	Total Bytes to Transfer <sup>1</sup>				μFrame C-prog-mask <sup>1</sup>						Status <sup>1</sup>						0x0C															
Buffer Pointer (Page 0)												Current Offset <sup>1</sup>												0x10										
Buffer Pointer (Page 1)												000_0000				TP <sup>1</sup>	T-count <sup>1</sup>				0x14													
Back Pointer																0000		T	0x18															

**Figure 16-39. Split-Transaction Isochronous Transaction Descriptor (siTD)**

<sup>1</sup> Host controller read/write; all others read-only.

### 16.5.4.1 Next Link Pointer

DWord0 of a siTD is a pointer to the next schedule data structure.

**Table 16-46. Next Link Pointer**

Bits	Name	Description
31–5	Next Link Pointer	This field contains the address of the next data object to be processed in the periodic list and corresponds to memory address signals [31:5], respectively.
4–3	—	Reserved, should be cleared. These bits must be written as zeros.
2–1	Typ	Indicates to the host controller whether the item referenced is an iTD/siTD or a QH. This allows the host controller to perform the proper type of processing on the item after it is fetched. Value encodings are: 00 iTD (isochronous transfer descriptor) 01 QH (queue head) 10 siTD (split transaction isochronous transfer descriptor) 11 FSTN (frame span traversal node)
0	T	Terminate. 0 Link pointer is valid. 1 Link pointer field is not valid.

### 16.5.4.2 siTD Endpoint Capabilities/Characteristics

DWords 1 and 2 specify static information about the full-speed endpoint, the addressing of the parent Companion Controller, and micro-frame scheduling control.

**Table 16-47. Endpoint and Transaction Translator Characteristics**

Bits	Name	Description
31	I/O	Direction (I/O). This field encodes whether the full-speed transaction should be an IN or OUT. 0 OUT 1 IN
30–24	Port Number	This field is the port number of the recipient transaction translator.
23	—	Reserved, should be cleared. Bit reserved and should be cleared.
22–16	Hub Address	This field holds the device address of the companion controllers' hub.
15–12	—	Reserved, should be cleared. Field reserved and should be cleared.
11–8	EndPt	Endpoint Number. Selects the particular endpoint number on the device serving as the data source or sink.
7	—	Reserved, should be cleared. Bit is reserved for future use. It should be cleared.
6–0	Device Address	Selects the specific device serving as the data source or sink.

**Table 16-48. Micro-Frame Schedule Control**

Bits	Name	Description
31–16	—	Reserved, should be cleared. This field reserved for future use. It should be cleared.
15–8	μFrame C-mask	Split completion mask. This field (along with the Active and SplitX- state fields in the status byte) is used to determine during which micro-frames the host controller should execute complete-split transactions. When the criteria for using this field is met, an all-zeros value has undefined behavior. The host controller uses the value of the three low-order bits of the FRINDEX register to index into this bit field. If the FRINDEX register value indexes to a position where the μFrame C-Mask field is a one, this siTD is a candidate for transaction execution. There may be more than one bit in this mask set.
7–0	μFrame S-mask	Split start mask. This field (along with the Active and SplitX-state fields in the Status byte) is used to determine during which micro-frames the host controller should execute start-split transactions. The host controller uses the value of the three low-order bits of the FRINDEX register to index into this bit field. If the FRINDEX register value indexes to a position where the μFrame S-mask field is a one, then this siTD is a candidate for transaction execution. An all zeros value in this field, in combination with existing periodic frame list has undefined results.

### 16.5.4.3 siTD Transfer State

DWords 3–6 manage the state of the transfer.

**Table 16-49. siTD Transfer Status and Control**

Bits	Name	Description
31	ioc	Interrupt on complete 0 Do not interrupt when transaction is complete. 1 Do interrupt when transaction is complete. When the host controller determines that the split transaction has completed it will assert a hardware interrupt at the next interrupt threshold.
30	P	Page select. Indicates which data page pointer should be concatenated with the CurrentOffset field to construct a data buffer pointer 0 Selects Page 0 pointer 1 Selects Page 1 pointer The host controller is not required to write this field back when the siTD is retired (Active bit transitioned from a one to a zero).

Table 16-49. siTD Transfer Status and Control (continued)

Bits	Name	Description	
29–26	—	Reserved, should be cleared. This field reserved for future use and should be cleared.	
25–16	Total Bytes to Transfer	This field is initialized by software to the total number of bytes expected in this transfer. Maximum value is 1023 (3FFh)	
15–8	μFrame C-prog-mask	Split complete progress mask. This field is used by the host controller to record which split-completes have been executed.	
7–0	Status	This field records the status of the transaction executed by the host controller for this slot. This field is a bit vector with the following encoding:	
		<b>Status Bits</b>	<b>Definition</b>
		7	Active. Set by software to enable the execution of an isochronous split transaction by the host controller.
		6	ERR. Set by the host controller when an ERR response is received from the companion controller.
		5	Data buffer error. Set by the host controller during status update to indicate that the host controller is unable to keep up with the reception of incoming data (overrun) or is unable to supply data fast enough during transmission (under run). In the case of an under run, the host controller will transmit an incorrect CRC (thus invalidating the data at the endpoint). If an overrun condition occurs, no action is necessary.
		4	Babble detected. Set by the host controller during status update when "babble" is detected during the transaction generated by this descriptor.
		3	Transaction error (XactErr). Set by the host controller during status update in the case where the host did not receive a valid response from the device (Time-out, CRC, Bad PID, etc.). This bit will only be set for IN transactions.
		2	Missed micro-frame. The host controller detected that a host-induced hold-off caused the host controller to miss a required complete-split transaction.
		1	Split transaction state (SplitXstate). The bit encodings are: 0 Do start split. This value directs the host controller to issue a Start split transaction to the endpoint when a match is encountered in the S-mask. 1 Do complete split. This value directs the host controller to issue a Complete split transaction to the endpoint when a match is encountered in the C-mask.
0	Reserved, should be cleared. Bit reserved for future use and should be cleared.		

#### 16.5.4.4 siTD Buffer Pointer List (Plus)

DWords 4 and 5 are the data buffer page pointers for the transfer. This structure supports one physical page cross. The most-significant 20 bits of each DWord in this section are the 4K (page) aligned buffer pointers. The least-significant 12 bits of each DWord are used as additional transfer state.

Table 16-50. siTD Buffer Pointer Page 0 (Plus)

Bits	Name	Description
31–12	Buffer Pointer (Page 0)	Bits 31–12 are 4K page-aligned, physical memory addresses. These bits correspond to physical address bits 31–12 respectively. The field P specifies the current active pointer
11–0	Current Offset	The 12 least-significant bits of the Page 0 pointer is the current byte offset for the current page pointer (as selected with the page indicator bit (P field)). The host controller is not required to write this field back when the siTD is retired (Active bit transitioned from a one to a zero).

**Table 16-51. siTD Buffer Pointer Page 1 (Plus)**

Bits	Name	Description
31–12	Buffer Pointer (Page 1)	Bits 31–12 are 4K page-aligned, physical memory addresses. These bits correspond to physical address bits 31–12 respectively. The field P specifies the current active pointer
11–5	—	Reserved, should be cleared.
4–3	TP	Transaction position. This field is used with T-count to determine whether to send all, first, middle, or last with each outbound transaction payload. System software must initialize this field with the appropriate starting value. The host controller must correctly manage this state during the lifetime of the transfer. The bit encodings are: 00 All. The entire full-speed transaction data payload is in this transaction (that is, less than or equal to 188 bytes). 01 Begin. This is the first data payload for a full-speed transaction that is greater than 188 bytes. 10 Mid. This is the middle payload for a full-speed OUT transaction that is larger than 188 bytes. 11 End. This is the last payload for a full-speed OUT transaction that was larger than 188 bytes.
2–0	T-Count	Transaction count. Software initializes this field with the number of OUT start-splits this transfer requires. Any value larger than 6 is undefined.

#### 16.5.4.5 siTD Back Link Pointer

DWord 6 of a siTD is simply another schedule link pointer. This pointer is always zero, or references a siTD. This pointer cannot reference any other schedule data structure.

**Table 16-52. siTD Back Link Pointer**

Bits	Name	Description
31–5	Back Pointer	A physical memory pointer to an siTD
4–1	—	Reserved, should be cleared. This field is reserved for future use. It should be cleared.
0	T	Terminate 0 siTD Back Pointer field is valid 1 siTD Back Pointer field is not valid

#### 16.5.5 Queue Element Transfer Descriptor (qTD)

This data structure is only used with a queue head. This data structure is used for one or more USB transactions. This data structure is used to transfer up to 20480 ( $5 \times 4096$ ) bytes. The structure contains two structure pointers used for queue advancement, a DWord of transfer state, and a five-element array of data buffer pointers. This structure is 32 bytes (or one 32-byte cache line). This data structure must be physically contiguous.

The buffer associated with this transfer must be virtually contiguous. The buffer may start on any byte boundary. A separate buffer pointer list element must be used for each physical page in the buffer, regardless of whether the buffer is physically contiguous.

Host controller updates (host controller writes) to stand-alone qTDs only occur during transfer retirement. References in the following bit field definitions of updates to the qTD are to the qTD portion of a queue head.



31	30	29	28	27	26	25	24	23	22	21	20	19	18	17	16	15	14	13	12	11	10	9	8	7	6	5	4	3	2	1	0	offset
Next qTD Pointer																												0000	T	0x00		
Alternate Next qTD Pointer																												0000	T	0x04		
dt <sup>1</sup>				Total Bytes to Transfer <sup>1</sup>												ioc		C_Page <sup>1</sup>		Cerr <sup>1</sup>		PID Code		Status <sup>1</sup>						0x08		
Buffer Pointer (Page 0)														Current Offset <sup>1</sup>														0x0C				
Buffer Pointer (Page 1)														0000_0000_0000														0x10				
Buffer Pointer (Page 2)														0000_0000_0000														0x14				
Buffer Pointer (Page 3)														0000_0000_0000														0x18				
Buffer Pointer (Page 4)														0000_0000_0000														0x1C				

**Figure 16-40. Queue Element Transfer Descriptor (qTD)**

<sup>1</sup> Host controller read/write; all others read-only.

Queue element transfer descriptors must be aligned on 32-byte boundaries.

### 16.5.5.1 Next qTD Pointer

The first DWord of an element transfer descriptor is a pointer to another transfer element descriptor.

**Table 16-53. qTD Next Element Transfer Pointer (DWord 0)**

Bits	Name	Description
31–5	Next qTD Pointer	This field contains the physical memory address of the next qTD to be processed and corresponds to memory address signals [31:5], respectively.
4–1	—	Reserved, should be cleared. These bits are reserved and their value has no effect on operation.
0	T	Terminate. Indicates to the host controller that there are no more valid entries in the queue. 0 Pointer is valid (points to a valid transfer element descriptor) 1 Pointer is invalid

### 16.5.5.2 Alternate Next qTD Pointer

The second DWord of a queue element transfer descriptor is used to support hardware-only advance of the data stream to the next client buffer on short packet. To be more explicit the host controller will always use this pointer when the current qTD is retired due to short packet.

**Table 16-54. qTD Alternate Next Element Transfer Pointer (DWord 1)**

Bits	Name	Description
31–5	Alternate Next qTD Pointer	This field contains the physical memory address of the next qTD to be processed in the event that the current qTD execution encounters a short packet (for an IN transaction). The field corresponds to memory address signals [31:5], respectively.
4–1	—	Reserved, should be cleared. These bits are reserved and their value has no effect on operation.
0	T	Terminate. Indicates to the host controller that there are no more valid entries in the queue. 0 Pointer is valid (points to a valid transfer element descriptor) 1 Pointer is invalid

### 16.5.5.3 qTD Token

The third DWord of a queue element transfer descriptor contains most of the information the host controller requires to execute a USB transaction (the remaining endpoint-addressing information is specified in the queue head). Note that some of the field descriptions in [Table 16-55](#) reference fields are defined in the queue head. See [Section 16.5.6, “Queue Head,”](#) for more information on these fields.

**Table 16-55. qTD Token (DWord 2)**

Bits	Name	Description
31	dt	Data toggle. This is the data toggle sequence bit. The use of this bit depends on the setting of the Data Toggle Control bit in the queue head.
30–16	Total Bytes to Transfer	Total bytes to transfer. This field specifies the total number of bytes to be moved with this transfer descriptor. This field is decremented by the number of bytes actually moved during the transaction, only on the successful completion of the transaction. The maximum value software may store in this field is $5 \times 4K$ (0x5000). This is the maximum number of bytes 5 page pointers can access. If the value of this field is zero when the host controller fetches this transfer descriptor (and the active bit is set), the host controller executes a zero-length transaction and retires the transfer descriptor. It is not a requirement for OUT transfers that total bytes to transfer be an even multiple of QH[Maximum Packet Length]. If software builds such a transfer descriptor for an OUT transfer, the last transaction will always be less than QH[Maximum Packet Length]. Although it is possible to create a transfer up to 20K this assumes the page is 0. When the offset cannot be predetermined, crossing past the 5th page can be guaranteed by limiting the total bytes to 16K. Therefore, the maximum recommended transfer is 16K (0x4000).
15	ioc	Interrupt on complete. If this bit is set, the host controller should issue an interrupt at the next interrupt threshold when this qTD is completed.
14–12	C_Page	Current rage. This field is used as an index into the qTD buffer pointer list. Valid values are in the range 0x0 to 0x4. The host controller is not required to write this field back when the qTD is retired.

Table 16-55. qTD Token (DWord 2) (continued)

Bits	Name	Description	
11–10	Cerr	Error counter. 2-bit down counter that keeps track of the number of consecutive errors detected while executing this qTD. If this field is programmed with a non-zero value during setup, the host controller decrements the count and writes it back to the qTD if the transaction fails. If the counter counts from one to zero, the host controller marks the qTD inactive, sets the Halted bit to a one, and error status bit for the error that caused Cerr to decrement to zero. An interrupt will be generated if USBINTR[UEE] is set. If the host controller driver (HCD) software programs this field to zero during setup, the host controller will not count errors for this qTD and there will be no limit on the retries of this qTD. Note that write-backs of intermediate execution state are to the queue head overlay area, not the qTD.	
		<b>Error</b>	<b>Decrement Counter</b>
		Transaction Error	Yes
		Data Buffer Error	No. Data buffer errors are host problems. They don't count against the device's retries. Note that software must not program Cerr to a value of zero when the EPS field is programmed with a value indicating a full- or low-speed device. This combination could result in undefined behavior.
		Stalled	No. Detection of babble or stall automatically halts the queue head. Thus, count is not decremented
		Babble Detected	No. Detection of babble or stall automatically halts the queue head. Thus, count is not decremented
		No Error	No. If the EPS field indicates a HS device or the queue head is in the asynchronous schedule (and PIDCode indicates an IN or OUT) and a bus transaction completes and the host controller does not detect a transaction error, then the host controller should reset Cerr to extend the total number of errors for this transaction. For example, Cerr should be reset with maximum value (0b11) on each successful completion of a transaction. The host controller must never reset this field if the value at the start of the transaction is 0b00.
9–8	PID Code	This field is an encoding of the token, which should be used for transactions associated with this transfer descriptor. Encodings are: 00 OUT Token generates token (E1H) 01 IN Token generates token (69H) 10 SETUP Token generates token (2DH) (undefined if endpoint is an Interrupt transfer type, for example. $\mu$ Frame S-mask field in the queue head is non-zero.) 11 Reserved, should be cleared	

**Table 16-55. qTD Token (DWord 2) (continued)**

Bits	Name	Description														
7–0	Status	This field is used by the host controller to communicate individual command execution states back to the host controller driver (HCD) software. This field contains the status of the last transaction performed on this qTD. The bit encodings are:														
		<table border="1"> <thead> <tr> <th data-bbox="428 388 659 443">Bits</th> <th data-bbox="659 388 1482 443">Status Field Description</th> </tr> </thead> <tbody> <tr> <td data-bbox="428 443 659 516">7</td> <td data-bbox="659 443 1482 516">Active. Set by software to enable the execution of transactions by the host controller.</td> </tr> <tr> <td data-bbox="428 516 659 711">6</td> <td data-bbox="659 516 1482 711">Halted. Set by the host controller during status updates to indicate that a serious error has occurred at the device/endpoint addressed by this qTD. This can be caused by babble, the error counter counting down to zero, or reception of the STALL handshake from the device during a transaction. Any time that a transaction results in the Halted bit being set, the Active bit is also cleared.</td> </tr> <tr> <td data-bbox="428 711 659 932">5</td> <td data-bbox="659 711 1482 932">Data buffer error. Set by the host controller during status update to indicate that the host controller is unable to keep up with the reception of incoming data (overrun) or is unable to supply data fast enough during transmission (under run). If an overrun condition occurs, the host controller will force a time-out condition on the USB, invalidating the transaction at the source. If the host controller sets this bit to a one, then it remains a one for the duration of the transfer.</td> </tr> <tr> <td data-bbox="428 932 659 1098">4</td> <td data-bbox="659 932 1482 1098">Babble detected. Set by the host controller during status update when babble is detected during the transaction. In addition to setting this bit, the host controller also sets the Halted bit to a one. Since babble is considered a fatal error for the transfer, setting the Halted bit to a one insures that no more transactions occur because of this descriptor.</td> </tr> <tr> <td data-bbox="428 1098 659 1234">3</td> <td data-bbox="659 1098 1482 1234">Transaction error (XactErr). Set by the host controller during status update in the case where the host did not receive a valid response from the device (time-out, CRC, bad PID). If the host controller sets this bit to a one, then it remains a one for the duration of the transfer.</td> </tr> <tr> <td data-bbox="428 1234 659 1425">2</td> <td data-bbox="659 1234 1482 1425">Missed micro-frame. This bit is ignored unless the QH[EPS] field indicates a full- or low-speed endpoint and the queue head is in the periodic list. This bit is set when the host controller detected that a host-induced hold-off caused the host controller to miss a required complete-split transaction. If the host controller sets this bit to a one, then it remains a one for the duration of the transfer.</td> </tr> </tbody> </table>	Bits	Status Field Description	7	Active. Set by software to enable the execution of transactions by the host controller.	6	Halted. Set by the host controller during status updates to indicate that a serious error has occurred at the device/endpoint addressed by this qTD. This can be caused by babble, the error counter counting down to zero, or reception of the STALL handshake from the device during a transaction. Any time that a transaction results in the Halted bit being set, the Active bit is also cleared.	5	Data buffer error. Set by the host controller during status update to indicate that the host controller is unable to keep up with the reception of incoming data (overrun) or is unable to supply data fast enough during transmission (under run). If an overrun condition occurs, the host controller will force a time-out condition on the USB, invalidating the transaction at the source. If the host controller sets this bit to a one, then it remains a one for the duration of the transfer.	4	Babble detected. Set by the host controller during status update when babble is detected during the transaction. In addition to setting this bit, the host controller also sets the Halted bit to a one. Since babble is considered a fatal error for the transfer, setting the Halted bit to a one insures that no more transactions occur because of this descriptor.	3	Transaction error (XactErr). Set by the host controller during status update in the case where the host did not receive a valid response from the device (time-out, CRC, bad PID). If the host controller sets this bit to a one, then it remains a one for the duration of the transfer.	2	Missed micro-frame. This bit is ignored unless the QH[EPS] field indicates a full- or low-speed endpoint and the queue head is in the periodic list. This bit is set when the host controller detected that a host-induced hold-off caused the host controller to miss a required complete-split transaction. If the host controller sets this bit to a one, then it remains a one for the duration of the transfer.
	Bits	Status Field Description														
	7	Active. Set by software to enable the execution of transactions by the host controller.														
	6	Halted. Set by the host controller during status updates to indicate that a serious error has occurred at the device/endpoint addressed by this qTD. This can be caused by babble, the error counter counting down to zero, or reception of the STALL handshake from the device during a transaction. Any time that a transaction results in the Halted bit being set, the Active bit is also cleared.														
	5	Data buffer error. Set by the host controller during status update to indicate that the host controller is unable to keep up with the reception of incoming data (overrun) or is unable to supply data fast enough during transmission (under run). If an overrun condition occurs, the host controller will force a time-out condition on the USB, invalidating the transaction at the source. If the host controller sets this bit to a one, then it remains a one for the duration of the transfer.														
	4	Babble detected. Set by the host controller during status update when babble is detected during the transaction. In addition to setting this bit, the host controller also sets the Halted bit to a one. Since babble is considered a fatal error for the transfer, setting the Halted bit to a one insures that no more transactions occur because of this descriptor.														
	3	Transaction error (XactErr). Set by the host controller during status update in the case where the host did not receive a valid response from the device (time-out, CRC, bad PID). If the host controller sets this bit to a one, then it remains a one for the duration of the transfer.														
2	Missed micro-frame. This bit is ignored unless the QH[EPS] field indicates a full- or low-speed endpoint and the queue head is in the periodic list. This bit is set when the host controller detected that a host-induced hold-off caused the host controller to miss a required complete-split transaction. If the host controller sets this bit to a one, then it remains a one for the duration of the transfer.															

Table 16-55. qTD Token (DWord 2) (continued)

Bits	Name	Description
1		Split transaction state (SplitXstate). This bit is ignored by the host controller unless the QH[EPS] field indicates a full- or low-speed endpoint. When a full- or low-speed device, the host controller uses this bit to track the state of the split- transaction. The functional requirements of the host controller for managing this state bit and the split transaction protocol depends on whether the endpoint is in the periodic or asynchronous schedule. The bit encodings are: 0 Do start split. This value directs the host controller to issue a start split transaction to the endpoint. 1 Do complete split. This value directs the host controller to issue a Complete split transaction to the endpoint.
0		Ping state (P)/ERR. If the QH[EPS] field indicates a high-speed device and the PID Code indicates an OUT endpoint, then this is the state bit for the Ping protocol. The bit encodings are: 0 Do OUT. This value directs the host controller to issue an OUT PID to the endpoint. 1 Do Ping. This value directs the host controller to issue a PING PID to the endpoint. If the QH[EPS] field does not indicate a high-speed device, then this field is used as an error indicator bit. It is set by the host controller whenever a periodic split-transaction receives an ERR handshake.

#### 16.5.5.4 qTD Buffer Page Pointer List

The last five DWords of a queue element transfer descriptor make up an array of physical memory address pointers. These pointers reference the individual pages of a data buffer.

System software initializes the Current Offset field to the starting offset into the current page, where current page is selected with the value in the C\_Page field.

Table 16-56. qTD Buffer Pointer

Bits	Name	Description
31–12	Buffer Pointer (page <i>n</i> )	Each element in the list is a 4K page aligned physical memory address. The lower 12 bits in each pointer are reserved (except for the first one), as each memory pointer must reference the start of a 4K page. The field C_Page specifies the current active pointer. When the transfer element descriptor is fetched, the starting buffer address is selected using C_Page (similar to an array index to select an array element). If a transaction spans a 4K buffer boundary, the host controller must detect the page-span boundary in the data stream, increment C_Page and advance to the next buffer pointer in the list, and conclude the transaction via the new buffer pointer.
11–0	Current Offset (Page 0)/ — (Pages 1–4)	Reserved in all pointers except the first one (that is, Page 0). The host controller should ignore all reserved bits. For the page 0 current offset interpretation, this field is the byte offset into the current page (as selected by C_Page). The host controller is not required to write this field back when the qTD is retired. Software should ensure the reserved fields are initialized to zeros.

## 16.5.6 Queue Head

Figure 16-41 shows the queue head structure.

31	30	29	28	27	26	25	24	23	22	21	20	19	18	17	16	15	14	13	12	11	10	9	8	7	6	5	4	3	2	1	0	offset
Queue Head Horizontal Link Pointer																											00	Typ	T	0x00		
RL			C	Maximum Packet Length				H	dtc	EPS	EndPt		I	Device Address								0x04 <sup>1</sup>										
Mult		Port Number			Hub Addr			µFrame C-mask				µFrame S-mask								0x08 <sup>1</sup>												
Current qTD Pointer <sup>2</sup>														00000								0x0C										
Next qTD Pointer <sup>2</sup>														0000				T <sup>2</sup>				0x10 <sup>3</sup>										
Alternate Next qTD Pointer <sup>2</sup>														NakCnt <sup>2</sup>				T <sup>2</sup>				0x14 <sup>3,4</sup>										
dt <sup>1</sup>	Total Bytes to Transfer <sup>2</sup>						ioc <sup>2</sup>	C_Page <sup>2</sup>	Cerr <sup>2</sup>	PID Code <sup>2</sup>	Status <sup>2</sup>										0x18 <sup>3,4</sup>											
Buffer Pointer (Page 0) <sup>2</sup>									Current Offset <sup>2</sup>													0x1C <sup>3,4</sup>										
Buffer Pointer (Page 1) <sup>2</sup>									0000				C-prog-mask <sup>2</sup>										0x20 <sup>3,4</sup>									
Buffer Pointer (Page 2) <sup>2</sup>									S-bytes <sup>2</sup>						FrameTag <sup>2</sup>								0x24 <sup>3,4</sup>									
Buffer Pointer (Page 3) <sup>2</sup>									0000_0000_0000													0x28 <sup>3</sup>										
Buffer Pointer (Page 4) <sup>2</sup>									0000_0000_0000													0x2C <sup>3</sup>										

Figure 16-41. Queue Head Layout

- <sup>1</sup> Offsets 0x04 through 0x0B contain the static endpoint state.
- <sup>2</sup> Host controller read/write; all others read-only.
- <sup>3</sup> Offsets 0x10 through 0x2F contain the transfer overlay.
- <sup>4</sup> Offsets 0x14 through 0x27 contain the transfer results.

### 16.5.6.1 Queue Head Horizontal Link Pointer

The first DWord of a queue head contains a link pointer to the next data object to be processed after any required processing in this queue has been completed, as well as the control bits defined below.

This pointer may reference a queue head or one of the isochronous transfer descriptors. It must not reference a queue element transfer descriptor.

Table 16-57. Queue Head DWord 0

Bits	Name	Description
31–5	QHLP	Queue head horizontal link pointer. This field contains the address of the next data object to be processed in the horizontal list and corresponds to memory address signals [31:5], respectively.
4–3	—	Reserved, should be cleared. These bits must be written as zeros.

Table 16-57. Queue Head DWord 0 (continued)

Bits	Name	Description
2–1	Typ	Indicates to the hardware whether the item referenced by the link pointer is an iTD, siTD or a QH. This allows the host controller to perform the proper type of processing on the item after it is fetched. 00 iTD (isochronous transfer descriptor) 01 QH (queue head) 10 siTD (split transaction isochronous transfer descriptor) 11 FSTN (frame span traversal node)
0	T	Terminate. 1 Last QH (pointer is invalid). 0 Pointer is valid. If the queue head is in the context of the periodic list, a one bit in this field indicates to the host controller that this is the end of the periodic list. This bit is ignored by the host controller when the queue head is in the asynchronous schedule. Software must ensure that queue heads reachable by the host controller always have valid horizontal link pointers.

### 16.5.6.2 Endpoint Capabilities/Characteristics

The second and third DWords of a Queue Head specify static information about the endpoint. This information does not change over the lifetime of the endpoint. There are three types of information in this region:

- Endpoint characteristics. These are the USB endpoint characteristics, which include addressing, maximum packet size, and endpoint speed.
- Endpoint capabilities. These are adjustable parameters of the endpoint. They affect how the endpoint data stream is managed by the host controller.
- Split transaction characteristics. This data structure manages full- and low-speed data streams for bulk, control, and interrupt with split transactions to USB 2.0 Hub transaction translator. Additional fields exist for addressing the hub and scheduling the protocol transactions (for periodic).

The host controller must not modify the bits in this region.

Table 16-58. Endpoint Characteristics: Queue Head DWord 1

Bits	Name	Description
31–28	RL	Nak count reload. This field contains a value, which is used by the host controller to reload the Nak Counter field.
27	C	Control endpoint flag. If the QH[EPS] field indicates the endpoint is not a high-speed device, and the endpoint is a control endpoint, then software must set this bit to a one. Otherwise, it should always set this bit to a zero.
26–16	Maximum Packet Length	This directly corresponds to the maximum packet size of the associated endpoint (wMaxPacketSize). The maximum value this field may contain is 0x400 (1024).
15	H	Head of reclamation list flag. This bit is set by system software to mark a queue head as being the head of the reclamation list.

**Table 16-58. Endpoint Characteristics: Queue Head DWord 1 (continued)**

Bits	Name	Description
14	dtc	Data toggle control (DTC). Specifies where the host controller should get the initial data toggle on an overlay transition. 0 Ignore DT bit from incoming qTD. Host controller preserves DT bit in the queue head. 1 Initial data toggle comes from incoming qTD DT bit. Host controller replaces DT bit in the queue head from the DT bit in the qTD.
13–12	EPS	Endpoint speed. This is the speed of the associated endpoint. 00 Full-speed (12 Mbps) 01 Low-speed (1.5 Mbps) 10 High-speed (480 Mbps) 11 Reserved, should be cleared This field must not be modified by the host controller.
11–8	EndPt	Endpoint number. Selects the particular endpoint number on the device serving as the data source or sink.
7	I	Inactivate on next transaction. This bit is used by system software to request that the host controller set the Active bit to zero. This field is only valid when the queue head is in the periodic schedule and the EPS field indicates a full- or low-speed endpoint. Setting this bit when the queue head is in the asynchronous schedule or the EPS field indicates a high-speed device yields undefined results.
6–0	Device Address	Selects the specific device serving as the data source or sink.

**Table 16-59. Endpoint Capabilities: Queue Head DWord 2**

Bits	Name	Description
31–30	Mult	High-bandwidth pipe multiplier. This field is a multiplier used to key the host controller as the number of successive packets the host controller may submit to the endpoint in the current execution. The host controller makes the simplifying assumption that software properly initializes this field (regardless of location of queue head in the schedules or other run time parameters). 00 Reserved, should be cleared. A zero in this field yields undefined results. 01 One transaction to be issued for this endpoint per micro-frame 10 Two transactions to be issued for this endpoint per micro-frame 11 Three transactions to be issued for this endpoint per micro-frame
29–23	Port Number	This field is ignored by the host controller unless the EPS field indicates a full- or low-speed device. The value is the port number identifier on the USB 2.0 hub (for hub at device address Hub Addr below), below which the full- or low-speed device associated with this endpoint is attached. This information is used in the split-transaction protocol.
22–16	Hub Addr	This field is ignored by the host controller unless the EPS field indicates a full- or low-speed device. The value is the USB device address of the USB 2.0 hub below which the full- or low-speed device associated with this endpoint is attached. This field is used in the split-transaction protocol.



**Table 16-59. Endpoint Capabilities: Queue Head DWord 2 (continued)**

Bits	Name	Description
15–8	μFrame C-mask	This field is ignored by the host controller unless the EPS field indicates this device is a low- or full-speed device and this queue head is in the periodic list. This field (along with the Active and SplitX-state fields) is used to determine during which micro-frames the host controller should execute a complete-split transaction. When the criteria for using this field are met, a zero value in this field has undefined behavior. This field is used by the host controller to match against the three low-order bits of the FRINDEX register. If the FRINDEX register bits decode to a position where the μFrame C- mask field is a one, then this queue head is a candidate for transaction execution. There may be more than one bit in this mask set.
7–0	μFrame S-mask	Interrupt schedule mask. This field is used for all endpoint speeds. Software should set this field to a zero when the queue head is on the asynchronous schedule. A non-zero value in this field indicates an interrupt endpoint. The host controller uses the value of the three low-order bits of the FRINDEX register as an index into a bit position in this bit vector. If the μFrame S-mask field has a one at the indexed bit position then this queue head is a candidate for transaction execution. If the EPS field indicates the endpoint is a high-speed endpoint, then the transaction executed is determined by the PID_Code field contained in the execution area. This field is also used to support split transaction types: Interrupt (IN/OUT). This condition is true when this field is non-zero and the EPS field indicates this is either a full- or low-speed device. A zero value in this field, in combination with existing in the periodic frame list has undefined results.

### 16.5.6.3 Transfer Overlay

The nine DWords in this area represent a transaction working space for the host controller. The general operational model is that the host controller can detect whether the overlay area contains a description of an active transfer. If it does not contain an active transfer, then it follows the queue head horizontal link pointer to the next queue head. The host controller will never follow the next transfer queue element or alternate queue element pointers unless it is actively attempting to advance the queue. For the duration of the transfer, the host controller keeps the incremental status of the transfer in the overlay area. When the transfer is complete, the results are written back to the original queue element.

The DWord3 of a queue head contains a pointer to the source qTD currently associated with the overlay. The host controller uses this pointer to write back the overlay area into the source qTD after the transfer is complete.

**Table 16-60. Current qTD Link Pointer**

Bits	Name	Description
31–5	Current qTD Pointer	Current element transaction descriptor link pointer. This field contains the address Of the current transaction being processed in this queue and corresponds to memory address signals [31:5], respectively.
4–0	—	Reserved, should be cleared. These bits are ignored by the host controller when using the value as an address to write data. The actual value may vary depending on the usage.

The DWords 4–11 of a queue head are the transaction overlay area. This area has the same base structure as a queue element transfer descriptor. The queue head utilizes the reserved fields of the page pointers to implement tracking the state of split transactions.


This area is characterized as an overlay because when the queue is advanced to the next queue element, the source queue element is merged onto this area. This area serves an execution cache for the transfer.

**Table 16-61. Host-Controller Rules for Bits in Overlay (DWords 5, 6, 8, and 9)**

DWord	QH Offset	Bits	Name	Description
5	0x14	4–1	NakCnt	Nak counter—RW. This field is a counter the host controller decrements whenever a transaction for the endpoint associated with this queue head results in a Nak or Nyet response. This counter is reloaded from RL before a transaction is executed during the first pass of the reclamation list (relative to an Asynchronous List Restart condition). It is also loaded from RL during an overlay.
6	0x18	31	dt	Data toggle. The Data toggle control controls whether the host controller preserves this bit when an overlay operation is performed.
6	0x18	15	ioc	Interrupt on complete. The ioc control bit is always inherited from the source qTD when the overlay operation is performed.
6	0x18	11–10	Cerr	Error counter. Copied from the qTD during the overlay and written back during queue advancement.
6	0x18	0	Status[0]	Ping state (P)/ERR. If the EPS field indicates a high-speed endpoint, then this field should be preserved during the overlay operation.
8	0x20	7–0	C-prog-mask	Split-transaction complete-split progress. Initialized to zero during any overlay. This field is used to track the progress of an interrupt split-transaction.
9	0x24	11–5	S-bytes	Software must ensure that the S-bytes field in a qTD is zero before activating the qTD. Keeps track of the number of bytes sent or received during an IN or OUT split transaction.
9	0x24	4–0	FrameTag	Split-transaction frame tag. Initialized to zero during any overlay. This field is used to track the progress of an interrupt split-transaction.

### 16.5.7 Periodic Frame Span Traversal Node (FSTN)

The periodic frame span traversal node (FSTN) data structure is to be used only for managing full- and low-speed transactions that span a host-frame boundary. Software must not use an FSTN in the asynchronous schedule. An FSTN in the asynchronous schedule results in undefined behavior. Software must not use the FSTN feature with a host controller whose HCIVERSION register indicates a revision implementation under 0x0096. Note that FSTNs were not defined for EHCI implementations before Revision 0.96 of the EHCI Specification and their use may yield undefined results.



**Figure 16-42. Frame Span Traversal Node Structure**

### 16.5.7.1 FTSN Normal Path Pointer

The first DWord of an FSTN contains a link pointer to the next schedule object. This object can be of any valid periodic schedule data type.

**Table 16-62. FTSN Normal Path Pointer**

Bits	Name	Description
31–5	NPLP	Normal path link pointer. Contains the address of the next data object to be processed in the periodic list and corresponds to memory address signals [31:5], respectively.
4–3	—	Reserved, should be cleared. These bits must be written as 0s.
2–1	Typ	Indicates to the host controller whether the item referenced is a iTD/siTD, QH, or FSTN. This allows the host controller to perform the proper type of processing on the item after it is fetched. 00 iTD (isochronous transfer descriptor) 01 QH (queue head) 10 siTD (split transaction isochronous transfer descriptor) 11 FSTN (frame span traversal node)
0	T	Terminate. 0 Link pointer is valid. 1 Link pointer field is not valid.

### 16.5.7.2 FSTN Back Path Link Pointer

The second DWord of an FTSN node contains a link pointer to a queue head. If the T-bit in this pointer is a zero, then this FSTN is a Save-Place indicator. Its Typ field must be set by software to indicate the target data structure is a queue head. If the T-bit in this pointer is set, then this FSTN is the Restore indicator. When the T-bit is a one, the host controller ignores the Typ field.

**Table 16-63. FSTN Back Path Link Pointer**

Bits	Name	Description
31–5	BPLP	Back path link pointer. Contains the address of a queue head. This field corresponds to memory address signals [31:5], respectively.
4–3	—	Reserved, should be cleared. These bits must be written as 0s.
2–1	Typ	Software must ensure this field is set to indicate the target data structure is a Queue Head (01). Any other value in this field yields undefined results.
0	T	Terminate. 0 Link pointer is valid (that is, the host controller may use bits 31–5 (in combination with the CTRLDSSEGMENT register if applicable) as a valid memory address). This value also indicates that this FSTN is a Save-Place indicator. 1 Link pointer field is not valid (that is, the host controller must not use bits 31–5 (in combination with the CTRLDSSEGMENT register if applicable) as a valid memory address). This value also indicates that this FSTN is a Restore indicator.

## 16.6 Host Operations

The general operational model for the USB DR module in host mode is defined by the Enhanced Host Controller Interface (EHCI) Specification. The EHCI specification describes the register-level interface for a host controller for the USB Revision 2.0. It includes a description of the hardware/software interface

between system software and host controller hardware. Information concerning the initialization of the USB module is included in the following section; however, the full details of the EHCI specification are beyond the scope of this document.

### 16.6.1 Host Controller Initialization

After initial power-on or host controller reset (hardware or through USBCMD[RST]), all of the operational registers will be at their default values. After a hardware reset, only the operational registers will be at their default values.

The MPC8313E USB PHY and clock must be configured prior to initialization of the USB controller. Initialization of the MPC8313E USB PHY interface is performed through software control following a power-on reset.

In order to initialize the USB DR module, software should perform the following steps

1. Set the controller mode to host mode. Optionally set USBMODE[SDIS] (streaming disable)

#### NOTE

Transitioning from device mode to host mode requires a host controller reset before modifying USBMODE.

2. Optionally modify the BURSTSIZE register.
3. Program the PTS field of the PORTSC register if using a non-ULPI PHY.
4. Set CONTROL[USB\_EN].
5. Write the appropriate value to the USBINTR register to enable the appropriate interrupts.
6. Write the base address of the periodic frame list to the PERIODICLIST BASE register. If there are no work items in the periodic schedule, all elements of the periodic frame list should have their T-Bits set.
7. Write the USBCMD register to set the desired interrupt threshold, frame list size (if applicable) and turn the controller by setting the RS bit.

At this point, the USB DR module is up and running and the port registers begin reporting device connects. System software can enumerate a port through the reset process (where the port is in the enabled state). At this point, the port is active with SOFs occurring down the enabled port enabled high-speed ports, but the schedules have not yet been enabled. The EHCI host controller will not transmit SOFs to enabled Full- or Low-speed ports.

In order to communicate with devices via the asynchronous schedule, system software must write the ASYNDLISTADDR register with the address of a control or bulk queue head. Software must then enable the asynchronous schedule by writing a one to USBCMD[ASE]. In order to communicate with devices via the periodic schedule, system software must enable the periodic schedule by writing a one to USBCMD[PSE]. Note that the schedules can be turned on before the first port is reset (and enabled).

Any time the USBCMD register is written, system software must ensure the appropriate bits are preserved, depending on the intended operation.

## 16.6.2 Power Port

The HCSPARAMS[PPC] bit indicates whether the USB 2.0 host controller has port power control. When the PPC bit is set, the host controller supports port power switches. Each available switch has an output enable. PPE is controlled based on the state of the combination bits PPC bit, EHCI Configured (CF)-bit and individual Port Power (PP) bits.

## 16.6.3 Reporting Over-Current

Host ports by definition are power providers on USB. Whether the ports are considered high- or low-powered is a platform implementation issue. The EHCI PORTSC register has an over-current status and over-current change bit. The functionality of these bits is specified in the USB Specification Revision 2.0.

The over current detection and limiting logic resides outside the DR logic. The over-current condition effects the following bits in the PORTSC register on the EHCI port:

- Over-current active bit (OCA) is set. When the over-current condition goes away, the OCA will transition from a one to a zero.
- Over-current change bit (OCC) is set. On every transition of OCA, the controller will set OCC to a one. Software sets OCC to a zero by writing a one to this bit.
- Port enabled/disabled bit (PE) is cleared. When this change bit gets set, USBSTS[PCI] (the port change detect bit) is set.
- Port power (PP) bit may optionally be cleared. There is no requirement in USB that a power provider shut off power in an over current condition. It is sufficient to limit the current and leave power applied. When OCC transitions from a zero to a one, the controller also sets USBSTS[PCI] to a one. In addition, if the Port Change Interrupt Enable bit, USBINTR[PCE], is a one, the controller issues an interrupt to the system. Refer to [Table 16-64](#) for summary of behavior for over-current detection when the controller is halted (suspended from a device component point of view).

## 16.6.4 Suspend/Resume

The host controller provides an equivalent suspend and resume model as that defined for individual ports in a USB 2.0 hub. Control mechanisms are provided to allow system software to suspend and resume individual ports. The mechanisms allow the individual ports to be resumed completely through software initiation. Other control mechanisms are provided to parameterize the host controller's response (or sensitivity) to external resume events. In this discussion, host-initiated, or software-initiated resumes are called Resume Events/Actions; bus-initiated resume events are called wake-up events. The classes of wakeup events are:

- Remote-wakeup enabled device asserts resume signaling. In similar kind to USB 2.0 hubs, when in host mode the host controller responds to explicit device resume signaling and wake up the system (if necessary).
- Port connect and disconnect and over-current events. Sensitivity to these events can be turned on or off by using the port control bits in the PORTSC register.

Selective suspend is a feature supported by the PORTSC register. It is used to place specific ports into a suspend mode. This feature is used as a functional component for implementing the appropriate power management policy implemented in a particular operating system. When system software intends to suspend the bus, it should suspend the enabled port, then shut off the controller by setting the USB\_CMD[RS] to a zero.

When a wake event occurs the system will resume operation and system software must set the RS bit to a one and resume the suspended port.

#### 16.6.4.1 Port Suspend/Resume

System software places the USB into suspend mode by writing a one into the appropriate PORTSC Suspend bit. Software must only set the Suspend bit when the port is in the enabled state (Port Enabled bit is a one).

The host controller may evaluate the Suspend bit immediately or wait until a micro-frame or frame boundary occurs. If evaluated immediately, the port is not suspended until the current transaction (if one is executing) completes. Therefore, there may be several micro-frames of activity on the port until the host controller evaluates the Suspend bit. The host controller must evaluate the Suspend bit at least every frame boundary.

System software can initiate a resume on the suspended port by writing a one to PORTSC[FPR]. Software should not attempt to resume a port unless the port reports that it is in the suspended state. If system software sets PORTSC[FPR] when the port is not in the suspended state, the resulting behavior is undefined. In order to assure proper USB device operation, software must wait for at least 10 milliseconds after a port indicates that it is suspended (Suspend bit is a one) before initiating a port resume through PORTSC[FPR]. When PORTSC[FPR] is set, the host controller sends resume signaling down the port. System software times the duration of the resume (nominally 20 milliseconds) then clears PORTSC[FPR]. When the host controller receives the write to transition PORTSC[FPR] to zero, it completes the resume sequence as defined in the USB specification, and clears both PORTSC[FPR] and PORTSC[SUSP]. Software-initiated port resumes do not affect the port change detect bit (USBSTS[PCI]) nor do they cause an interrupt if USBINTR[PCE] (port change interrupt enable) is a one. When a wake event occurs on a suspended port, the resume signaling is detected by the port and the resume is reflected downstream within 100  $\mu$ sec. The port's PORTSC[FPR] bit is set and USBSTS[PCI] is set. If USBINTR[PCE] is a one, the host controller issues a hardware interrupt.

System software observes the resume event on the port, delays a port resume time (nominally 20 milliseconds), then terminates the resume sequence by clearing PORTSC[FPR] in the port. The host controller receives the write of zero to PORTSC[FPR], terminates the resume sequence and clears PORTSC[FPR] and PORTSC[SUSP]. Software can determine that the port is enabled (not suspended) by sampling the PORTSC register and observing that the SUSP and FPR bits are zero. Software must ensure that the host controller is running (that is, USBSTS[HCH] is a zero), before terminating a resume by clearing the port's PORTSC[FPR] bit. If HCH is a one when PORTSC[FPR] is cleared, then SOFs will not occur down the enabled port and the device will return to suspend mode in a maximum of 10 milliseconds.

[Table 16-64](#) summarizes the wake-up events. Whenever a resume event is detected, USBSTS[PCI] is set. If USBINTR[PCE] (port change interrupt enable) is a one, the host controller also generates an interrupt on the resume event. Software acknowledges the resume event interrupt by clearing the USBSTS[PCI].

Table 16-64. Behavior During Wake-Up Events

Port Status and Signaling Type	Signaled Port Response	Device State	
		D0	not D0
Port disabled, resume K-State received	No effect	N/A	N/A
Port suspended, Resume K-State received	Resume reflected downstream on signaled port. PORTSC[FPR] is set. USBSTS[PCI] is set.	[1], [2]	[2]
Port is enabled, disabled or suspended, and the port's WKDSCNNT_E bit, PORTSC[WKDS], is set. A disconnect is detected.	Depending on the initial port state, the PORTSC Connect (CCS) and Enable (PE) status bits are cleared, and the Connect Change status bit (CSC) is set. USBSTS[PCI] is set.	[1], [2]	[2]
Port is enabled, disabled or suspended, and the port's WKDSCNNT_E bit, PORTSC[WKDS], is cleared. A disconnect is detected.	Depending on the initial port state, the PORTSC Connect (CCS) and Enable (PE) status bits are cleared, and the Connect Change status bit (CSC) is set. USBSTS[PCI] is set.	[1], [3]	[3]
Port is not connected and the port's WKCNTNT_E bit is a one. A connect is detected.	PORTSC Connect Status (CCS) and Connect Status Change (CSC) bits are set. USBSTS[PCI] is set.	[1], [2]	[2]
Port is not connected and the port's WKCNTNT_E bit is a zero. A connect is detected.	PORTSC Connect Status (CCS) and Connect Status Change (CSC) bits are set. USBSTS[PCI] is set.	[1], [3]	[3]
Port is connected and the port's WKOC_E bit is a one. An over-current condition occurs.	PORTSC Over-current Active (OCA), Over-current Change (OCC) bits are set. If Port Enable/Disable bit (PE) is a one, it is cleared. USBSTS[PCI] is set	[1], [2]	[2]
Port is connected and the port's WKOC_E bit is a zero. An over-current condition occurs.	PORTSC Over-current Active (OCA), Over-current Change (OCC) bits are set. If Port Enable/Disable bit (PE) is a one, it is cleared. USBSTS[PCI] is set.	[1], [3]	[3]

<sup>1</sup> Hardware interrupt issued if USBINTR[PCE] (port change interrupt enable) is set.

<sup>2</sup> PME# asserted if enabled (Note: PME Status must always be set).

<sup>3</sup> PME# not asserted.

## 16.6.5 Schedule Traversal Rules

The host controller executes transactions for devices using a simple, shared-memory schedule. The schedule is comprised of a few data structures, organized into two distinct lists. The data structures are designed to provide the maximum flexibility required by USB, minimize memory traffic and hardware/software complexity.

System software maintains two schedules for the host controller: a periodic schedule and an asynchronous schedule. The root of the periodic schedule is the PERIODICLISTBASE register. See [Section 16.3.2.6, “Periodic Frame List Base Address Register \(PERIODICLISTBASE\),”](#) for more information. The PERIODICLISTBASE register is the physical memory base address of the periodic frame list. The periodic frame list is an array of physical memory pointers. The objects referenced from the frame list must be valid schedule data structures as defined in [Section 16.5, “Host Data Structures.”](#) In each micro-frame, if the periodic schedule is enabled (see) then the host controller must execute from the periodic schedule before executing from the asynchronous schedule. It will only execute from the asynchronous schedule

after it encounters the end of the periodic schedule. The host controller traverses the periodic schedule by constructing an array offset reference from the PERIODICLISTBASE and the FRINDEX registers (see Figure 16-43). It fetches the element and begins traversing the graph of linked schedule data structures.

The end of the periodic schedule is identified by a next link pointer of a schedule data structure having its T-bit set. When the host controller encounters a T-Bit set during a horizontal traversal of the periodic list, it interprets this as an End-Of-Periodic-List mark. This causes the host controller to cease working on the periodic schedule and transitions immediately to traversing the asynchronous schedule. Once this transition is made, the host controller executes from the asynchronous schedule until the end of the micro-frame.

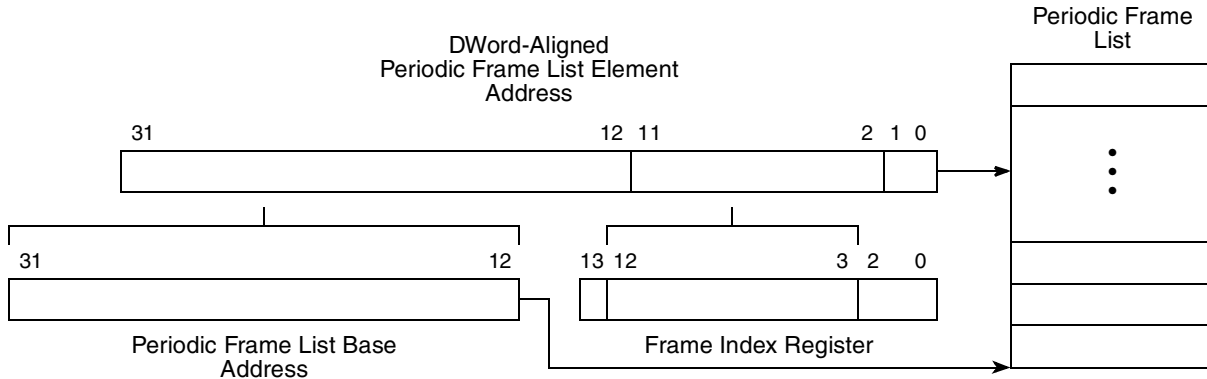


Figure 16-43. Derivation of Pointer into Frame List Array

When the host controller determines that it is time to execute from the asynchronous list, it uses the operational register ASYNCLISTADDR to access the asynchronous schedule, as shown in Figure 16-44.




Figure 16-44. General Format of Asynchronous Schedule List


The ASYNCLISTADDR register contains a physical memory pointer to the next queue head. When the host controller makes a transition to executing the asynchronous schedule, it begins by reading the queue head referenced by the ASYNCLISTADDR register. Software must set queue head horizontal pointer T-bits to a zero for queue heads in the asynchronous schedule.

### 16.6.6 Periodic Schedule Frame Boundaries vs. Bus Frame Boundaries

The USB Specification Revision 2.0 requires that the frame boundaries (SOF frame number changes) of the high-speed bus and the full- and low-speed bus(es) below USB 2.0 hubs be strictly aligned.



Super-imposed on this requirement is that USB 2.0 hubs manage full- and low-speed transactions via a micro-frame pipeline (see start- (SS) and complete- (CS) splits illustrated in Figure 16-45). A simple, direct projection of the frame boundary model into the host controller interface schedule architecture creates tension (complexity for both hardware and software) between the frame boundaries and the scheduling mechanisms required to service the full- and low-speed transaction translator periodic pipelines.

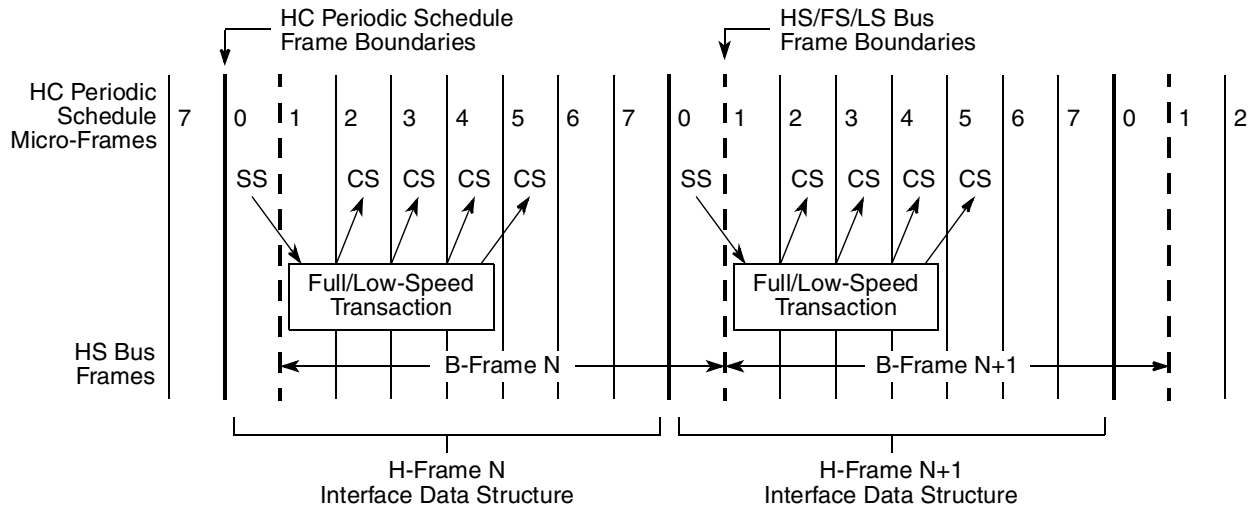


**Figure 16-45. Frame Boundary Relationship Between HS Bus and FS/LS Bus**

The simple projection, as Figure 16-45 illustrates, introduces frame-boundary wrap conditions for scheduling on both the beginning and end of a frame. In order to reduce the complexity for hardware and software, the host controller is required to implement a one micro-frame phase shift for its view of frame boundaries. The phase shift eliminates the beginning of frame and frame-wrap scheduling boundary conditions.

The implementation of this phase shift requires that the host controller use one register value for accessing the periodic frame list and another value for the frame number value included in the SOF token. These two values are separate, but tightly coupled. The periodic frame list is accessed via the Frame List Index Register (FRINDEX). Bits FRINDEX[2–0], represent the micro-frame number. The SOF value is coupled to the value of FRINDEX[13–3]. Both FRINDEX[13–3] and the SOF value are incremented based on FRINDEX[2–0]. It is required that the SOF value be delayed from the FRINDEX value by one micro-frame. The one micro-frame delay yields a host controller periodic schedule and bus frame boundary relationship as illustrated in Figure 16-46. This adjustment allows software to trivially schedule the periodic start and complete-split transactions for full-and low-speed periodic endpoints, using the natural alignment of the periodic schedule interface.

Figure 16-46 illustrates how periodic schedule data structures relate to schedule frame boundaries and bus frame boundaries. To aid the presentation, two terms are defined. The host controller's view of the 1-millisecond boundaries is called H-Frames. The high-speed bus's view of the 1-millisecond boundaries is called B-Frames.



**Figure 16-46. Relationship of Periodic Schedule Frame Boundaries to Bus Frame Boundaries**

H-Frame boundaries for the host controller correspond to increments of FRINDEX[13–3]. Micro-frame numbers for the H-Frame are tracked by FRINDEX[2–0]. B-Frame boundaries are visible on the high-speed bus via changes in the SOF token's frame number. Micro-frame numbers on the high-speed bus are only derived from the SOF token's frame number (that is, the high-speed bus will see eight SOFs with the same frame number value). H-Frames and B-Frames have the fixed relationship (that is, B-Frames lag H-Frames by one micro-frame time) illustrated in Figure 16-46. The host controller's periodic schedule is naturally aligned to H-Frames. Software schedules transactions for full- and low-speed periodic endpoints relative the H-Frames. The result is these transactions execute on the high-speed bus at exactly the right time for the USB 2.0 hub periodic pipeline. As described in Section 16.3.2.4, “Frame Index Register (FRINDEX),” the SOF Value can be implemented as a shadow register (in this example, called SOFV), which lags the FRINDEX register bits [13–3] by one micro-frame count. Table 16-65 illustrates the required relationship between the value of FRINDEX and the value of SOFV. This lag behavior can be accomplished by incrementing FRINDEX[13–3] based on carry-out on the 7 to 0 increment of FRINDEX[2–0] and incrementing SOFV based on the transition of 0 to 1 of FRINDEX[2–0].

Software is allowed to write to FRINDEX. Section 16.3.2.4, “Frame Index Register (FRINDEX),” provides the requirements that software should adhere when writing a new value in FRINDEX.

Table 16-65. Operation of FRINDEX and SOFV (SOF Value Register)

Current			Next		
FRINDEX[13–3]	SOFV	FRINDEX[2–0]	FRINDEX[13–3]	SOFV	FRINDEX[2–0]
N	N	111	N+1	N	000
N+1	N	000	N+1	N+1	001
N+1	N+1	001	N+1	N+1	010
N+1	N+1	010	N+1	N+1	011
N+1	N+1	011	N+1	N+1	100
N+1	N+1	100	N+1	N+1	101
N+1	N+1	101	N+1	N+1	110
N+1	N+1	110	N+1	N+1	111

### 16.6.7 Periodic Schedule

The periodic schedule traversal is enabled or disabled through USBCMD[PSE] (periodic schedule enable). If USBCMD[PSE] is cleared, then the host controller simply does not try to access the periodic frame list via the PERIODICLISTBASE register. Likewise, when USBCMD[PSE] is a one, then the host controller does use the PERIODICLISTBASE register to traverse the periodic schedule. The host controller will not react to modifications to USBCMD[PSE] immediately. In order to eliminate conflicts with split transactions, the host controller evaluates USBCMD[PSE] only when FRINDEX[2–0] is zero. System software must not disable the periodic schedule if the schedule contains an active split transaction work item that spans the 0b000 micro-frame. These work items must be removed from the schedule before USBCMD[PSE] is cleared. USBSTS[PS] (periodic schedule status) indicates status of the periodic schedule. System software enables (or disables) the periodic schedule by setting (or clearing) USBCMD[PSE]. Software then can poll USBSTS[PS] to determine when the periodic schedule has made the desired transition. Software must not modify USBCMD[PSE] unless the value of USBCMD[PSE] equals that of USBSTS[PS].

The periodic schedule is used to manage all isochronous and interrupt transfer streams. The base of the periodic schedule is the periodic frame list. Software links schedule data structures to the periodic frame list to produce a graph of scheduled data structures. The graph represents an appropriate sequence of transactions on the USB. Figure 16-47 illustrates isochronous transfers (using iTDs and siTDs) with a period of one are linked directly to the periodic frame list. Interrupt transfers (are managed with queue heads) and isochronous streams with periods other than one are linked following the period-one iTD/siTDs. Interrupt queue heads are linked into the frame list ordered by poll rate. Longer poll rates are linked first (for example, closest to the periodic frame list), followed by shorter poll rates, with queue heads with a poll rate of one, on the very end.

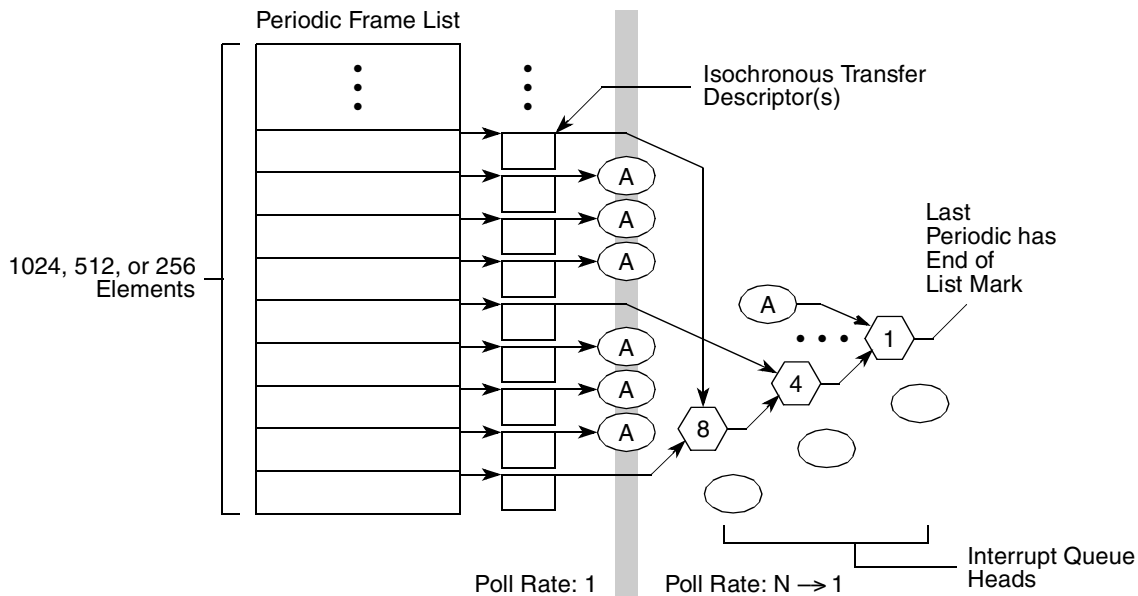


Figure 16-47. Example Periodic Schedule

## 16.6.8 Managing Isochronous Transfers Using iTDs

The structure of an iTD is presented in Isochronous (High-Speed) Transfer Descriptor (iTID). There are four distinct sections to an iTD:

- The first field is the Next Link Pointer. This field is for schedule linkage purposes only.
- Transaction description array. This area is an eight-element array. Each element represents control and status information for one micro-frame's worth of transactions for a single high-speed isochronous endpoint.
- The buffer page pointer array is a 7-element array of physical memory pointers to data buffers. These are 4K aligned pointers to physical memory.
- Endpoint capabilities. This area utilizes the unused low-order 12 bits of the buffer page pointer array. The fields in this area are used across all transactions executed for this iTD, including endpoint addressing, transfer direction, maximum packet size and high-bandwidth multiplier.

### 16.6.8.1 Host Controller Operational Model for iTDs

The host controller uses FRINDEX register bits 12–3 to index into the periodic frame list. This means that the host controller visits each frame list element eight consecutive times before incrementing to the next periodic frame list element. Each iTD contains eight transaction descriptions, which map directly to FRINDEX register bits 2–0. Each iTD can span 8 micro-frames worth of transactions. When the host controller fetches an iTD, it uses FRINDEX register bits 2–0 to index into the transaction description array. When the first iTD in the periodic list is traversed after periodic schedule is enabled, the value of FRINDEX[2:0] may be other than 0, so the first transaction issued by the controller may be any of the eight available active transactions. If the active bit in the Status field of the indexed transaction description is cleared, the host controller ignores the iTD and follows the Next pointer to the next schedule data structure.

When the indexed active bit is a one the host controller continues to parse the iTD. It stores the indexed transaction description and the general endpoint information (device address, endpoint number, maximum packet size, etc.). It also uses the Page Select (PG) field to index the buffer pointer array, storing the selected buffer pointer and the next sequential buffer pointer. For example, if PG field is a 0, then the host controller will store Page 0 and Page 1.

The host controller constructs a physical data buffer address by concatenating the current buffer pointer (as selected using the current transaction description's PG field) and the transaction description's Transaction Offset field. The host controller uses the endpoint addressing information and I/O-bit to execute a transaction to the appropriate endpoint. When the transaction is complete, the host controller clears the active bit and writes back any additional status information to the Status field in the currently selected transaction description.

The data buffer associated with the iTD must be virtually contiguous memory. Seven page pointers are provided to support eight high-bandwidth transactions regardless of the starting packet's offset alignment into the first page. A starting buffer pointer (physical memory address) is constructed by concatenating the page pointer (example: page 0 pointer) selected by the active transaction descriptions' PG (example value: 0b00) field with the transaction offset field. As the transaction moves data, the host controller must detect when an increment of the current buffer pointer will cross a page boundary. When this occurs the host controller simply replaces the current buffer pointer's page portion with the next page pointer (example: page 1 pointer) and continues to move data. The size of each bus transaction is determined by the value in the Maximum Packet Size field. An iTD supports high-bandwidth pipes via the Mult (multiplier) field. When the Mult field is 1, 2, or 3, the host controller executes the specified number of Maximum Packet sized bus transactions for the endpoint in the current micro-frame. In other words, the Mult field represents a transaction count for the endpoint in the current micro-frame. If the Mult field is zero, the operation of the host controller is undefined. The transfer description is used to service all transactions indicated by the Mult field.

For OUT transfers, the value of the Transaction  $n$  Length field represents the total bytes to be sent during the micro-frame. The Mult field must be set by software to be consistent with Transaction  $n$  Length and Maximum Packet Size. The host controller will send the bytes in Maximum Packet Sized portions. After each transaction, the host controller decrements it's local copy of Transaction  $n$  Length by Maximum Packet Size. The number of bytes the host controller sends is always Maximum Packet Size or Transaction  $n$  Length, whichever is less. The host controller advances the transfer state in the transfer description, updates the appropriate record in the iTD and moves to the next schedule data structure. The maximum sized transaction supported is  $3 \times 1024$  bytes.

For IN transfers, the host controller issues Mult transactions. It is assumed that software has properly initialized the iTD to accommodate all of the possible data. During each IN transaction, the host controller must use Maximum Packet Size to detect packet babble errors. The host controller keeps the sum of bytes received in the Transaction  $n$  Length field. After all transactions for the endpoint have completed for the micro-frame, Transaction  $n$  Length contains the total bytes received. If the final value of Transaction  $n$  Length is less than the value of Maximum Packet Size, then less data than was allowed for was received from the associated endpoint. This short packet condition does not set USBSTS[UI] (USB interrupt). The host controller will not detect this condition. If the device sends more than Transaction  $n$  Length or Maximum Packet Size bytes (whichever is less), then the host controller will set the Babble Detected bit and clear the Active bit. Note, that the host controller is not required to update the iTD field Transaction  $n$


Length in this error scenario. If the Mult field is greater than one, then the host controller will automatically execute the value of Mult transactions. The host controller will not execute all Mult transactions if:

- The endpoint is an OUT and Transaction  $n$  Length goes to zero before all the Mult transactions have executed (ran out of data), or
- The endpoint is an IN and the endpoint delivers a short packet, or an error occurs on a transaction before Mult transactions have been executed. The end of micro-frame may occur before all of the transaction opportunities have been executed. When this happens, the transfer state of the transfer description is advanced to reflect the progress that was made, the result written back to the iTD and the host controller proceeds to processing the next micro-frame.

### 16.6.8.2 Software Operational Model for iTDs

A client buffer request to an isochronous endpoint may span 1 to N micro-frames. When N is larger than one, system software may have to use multiple iTDs to read or write data with the buffer (if N is larger than eight, it must use more than one iTD).

Figure 16-48 illustrates the simple model of how a client buffer is mapped by system software to the periodic schedule (that is, the periodic frame list and a set of iTDs). On the right is the client description of its request. The description includes a buffer base address plus additional annotations to identify which portions of the buffer should be used with each bus transaction. In the middle is the iTD data structures used by the system software to service the client request. Each iTD can be initialized to service up to 24 transactions, organized into eight groups of up to three transactions each. Each group maps to one micro-frame's worth of transactions. The EHCI controller does not provide per-transaction results within a micro-frame. It treats the per-micro-frame transactions as a single logical transfer. On the left is the host controller's frame list. System software establishes references from the appropriate locations in the frame list to each of the appropriate iTDs. If the buffer is large, then system software can use a small set of iTDs to service the entire buffer. System software can activate the transaction description records (contained in each iTD) in any pattern required for the particular data stream.



**Figure 16-48. Example Association of iTDs to Client Request Buffer**

As noted above, the client request includes a pointer to the base of the buffer and offsets into the buffer to annotate which buffer sections are to be used on each bus transaction that occurs on this endpoint. System software must initialize each transaction description in an iTD to ensure it uses the correct portion of the client buffer. For example, for each transaction description, the PG field is set to index the correct physical buffer page pointer and the Transaction Offset field is set relative to the correct buffer pointer page (for example, the same one referenced by the PG field). When the host controller executes a transaction it selects a transaction description record based on FRINDEX[2–0]. It then uses the current Page Buffer Pointer (as selected by the PG field) and concatenates to the transaction offset field. The result is a starting buffer address for the transaction. As the host controller moves data for the transaction, it must watch for a page wrap condition and properly advance to the next available Page Buffer Pointer. System software must not use the Page 6 buffer pointer in a transaction description where the length of the transfer will wrap a page boundary. Doing so yields undefined behavior. The host controller hardware is not required to alias the page selector to page zero. USB 2.0 isochronous endpoints can specify a period greater than one. Software can achieve the appropriate scheduling by linking iTDs into the appropriate frames (relative to the frame list) and by setting appropriate transaction description elements active bits to a one.

### 16.6.8.2.1 Periodic Scheduling Threshold

The Isochronous Scheduling Threshold field in the HCCPARAMS capability register is an indicator to system software as to how the host controller pre-fetches and effectively caches schedule data structures. It is used by system software when adding isochronous work items to the periodic schedule. The value of this field indicates to system software the minimum distance it can update isochronous data (relative to the current location of the host controller execution in the periodic list) and still have the host controller process them.

The iTD and siTD data structures each describe 8 micro-frames worth of transactions. The host controller is allowed to cache one (or more) of these data structures in order to reduce memory traffic. There are three basic caching models that account for the fact the isochronous data structures span 8 micro-frames. The three caching models are: no caching, micro-frame caching and frame caching.

When software is adding new isochronous transactions to the schedule, it always performs a read of the FRINDEX register to determine the current frame and micro-frame the host controller is currently executing. Of course, there is no information about where in the micro-frame the host controller is, so a constant uncertainty factor of one micro-frame has to be assumed. Combining the knowledge of where the host controller is executing with the knowledge of the caching model allows the definition of simple algorithms for how closely software can reliably work to the executing host controller.

No caching is indicated with a value of zero in the Isochronous Scheduling Threshold field. The host controller may pre-fetch data structures during a periodic schedule traversal (per micro-frame) but will always dump any accumulated schedule state at the end of the micro-frame. At the appropriate time relative to the beginning of every micro-frame, the host controller always begins schedule traversal from the frame list. Software can use the value of the FRINDEX register (plus the constant 1 uncertainty-factor) to determine the approximate position of the executing host controller. When no caching is selected, software can add an isochronous transaction as near as 2 micro-frames in front of the current executing position of the host controller.

Frame caching is indicated with a non-zero value in bit [7] of the Isochronous Scheduling Threshold field. In the frame-caching model, system software assumes that the host controller caches one (or more) isochronous data structures for an entire frame (8 micro-frames). Software uses the value of the FRINDEX register (plus the constant 1 uncertainty) to determine the current micro-frame/frame (assume modulo 8 arithmetic in adding the constant 1 to the micro-frame number). For any current frame N, if the current micro-frame is 0 to 6, then software can safely add isochronous transactions to Frame N + 1. If the current micro-frame is 7, then software can add isochronous transactions to Frame N + 2.

Micro-frame caching is indicated with a non-zero value in the least-significant 3 bits of the Isochronous Scheduling Threshold field. System software assumes the host controller caches one or more periodic data structures for the number of micro-frames indicated in the Isochronous Scheduling Threshold field. For example, if the count value were 2, then the host controller keeps a window of 2 micro-frames worth of state (current micro-frame, plus the next) on-chip. On each micro-frame boundary, the host controller releases the current micro-frame state and begins accumulating the next micro-frame state.

### 16.6.9 Asynchronous Schedule

The asynchronous schedule traversal is enabled or disabled through USBCMD[ASE] (asynchronous schedule enable). If USBCMD[ASE] is cleared, then the host controller simply does not try to access the asynchronous schedule via the ASYNCLISTADDR register. Likewise, if USBCMD[ASE] is set, the host controller does use the ASYNCLISTADDR register to traverse the asynchronous schedule. Modifications to USBCMD[ASE] are not necessarily immediate. Rather the new value of the bit will only be taken into consideration the next time the host controller needs to use the value of the ASYNCLISTADDR register to get the next queue head.

USBSTS[AS] indicates status of the asynchronous schedule. System software enables (or disables) the asynchronous schedule by writing a one (or zero) to USBCMD[ASE]. Software then can poll



USBSTS[AS] to determine when the asynchronous schedule has made the desired transition. Software must not modify USBCMD[ASE] unless the value of USBCMD[ASE] equals that of the USBSTS[AS] (asynchronous schedule status).

The asynchronous schedule is used to manage all Control and Bulk transfers. Control and Bulk transfers are managed using queue head data structures. The asynchronous schedule is based at the ASYNCLISTADDR register. The default value of the ASYNCLISTADDR register after reset is undefined and the schedule is disabled when USBCMD[ASE] is cleared.

Software may only write this register with defined results when the schedule is disabled, for example, USBCMD[ASE] and the USBSTS[AS] are cleared. System software enables execution from the asynchronous schedule by writing a valid memory address (of a queue head) into this register. Then software enables the asynchronous schedule by setting USBCMD[ASE]. The asynchronous schedule is actually enabled when USBSTS[AS] is set.

When the host controller begins servicing the asynchronous schedule, it begins by using the value of the ASYNCLISTADDR register. It reads the first referenced data structure and begins executing transactions and traversing the linked list as appropriate. When the host controller completes processing the asynchronous schedule, it retains the value of the last accessed queue head's horizontal pointer in the ASYNCLISTADDR register. Next time the asynchronous schedule is accessed, this is the first data structure that is serviced. This provides round-robin fairness for processing the asynchronous schedule.

A host controller completes processing the asynchronous schedule when one of the following events occur:

- The end of a micro-frame occurs.
- The host controller detects an empty list condition
- The schedule has been disabled through USBCMD[ASE].

The queue heads in the asynchronous list are linked into a simple circular list as shown in [Figure 16-44](#). Queue head data structures are the only valid data structures that may be linked into the asynchronous schedule. An isochronous transfer descriptor (iT<sub>D</sub> or siT<sub>D</sub>) in the asynchronous schedule yields undefined results.

The maximum packet size field in a queue head is sized to accommodate the use of this data structure for all non-isochronous transfer types. The USB Specification, Revision 2.0 specifies the maximum packet sizes for all transfer types and transfer speeds. System software should always parameterize the queue head data structures according to the core specification requirements.

### 16.6.9.1 Adding Queue Heads to Asynchronous Schedule

This is a software requirement section. There are two independent events for adding queue heads to the asynchronous schedule. The first is the initial activation of the asynchronous list. The second is inserting a new queue head into an activated asynchronous list.

Activation of the list is simple. System software writes the physical memory address of a queue head into the ASYNCLISTADDR register, then enables the list by setting USBCMD[ASE] to a one.

When inserting a queue head into an active list, software must ensure that the schedule is always coherent from the host controllers' point of view. This means that the system software must ensure that all queue

head pointer fields are valid. For example qTD pointers have T-Bits set or reference valid qTDs and the Horizontal Pointer references a valid queue head data structure. The following algorithm represents the functional requirements:

```

InsertQueueHead (pQHeadCurrent, pQueueHeadNew)
--
-- Requirement: all inputs must be properly initialized.
--
-- pQHeadCurrent is a pointer to a queue head that is
-- already in the active list
-- pQHeadNew is a pointer to the queue head to be added
--
-- This algorithm links a new queue head into a existing
-- list
--
pQueueHeadNew.HorizontalPointer = pQueueHeadCurrent.HorizontalPointer
pQueueHeadCurrent.HorizontalPointer = physicalAddressOf (pQueueHeadNew)
End InsertQueueHead

```

### 16.6.9.2 Removing Queue Heads from Asynchronous Schedule

This is a software requirement section. There are two independent events for removing queue heads from the asynchronous schedule. The first is shutting down (deactivating) the asynchronous list. The second is extracting a single queue head from an activated list. Software deactivates the asynchronous schedule by setting USBCMD[ASE] to a zero. Software can determine when the list is idle when USBSTS[AS] is cleared. The normal mode of operation is that software removes queue heads from the asynchronous schedule without shutting it down. Software must not remove an active queue head from the schedule. Software should first deactivate all active qTDs, wait for the queue head to go inactive, then remove the queue head from the asynchronous list. Software removes a queue head from the asynchronous list using the following algorithm. Software merely must ensure all of the link pointers reachable by the host controller are kept consistent.

```

UnlinkQueueHead (pQHeadPrevious, pQueueHeadToUnlink, pQHeadNext)
--
-- Requirement: all inputs must be properly initialized.
--
-- pQHeadPrevious is a pointer to a queue head that
-- references the queue head to remove
-- pQHeadToUnlink is a pointer to the queue head to be
-- removed
-- pQHeadNext is a pointer to a queue head still in the
-- schedule. Software provides this pointer with the
-- following strict rules:
-- if the host software is one queue head, then
-- pQHeadNext must be the same as
-- QueueheadToUnlink.HorizontalPointer. If the host
-- software is unlinking a consecutive series of
-- queue heads, QHeadNext must be set by software to
-- the queue head remaining in the schedule.
--
-- This algorithm unlinks a queue head from a circular list
--
pQueueHeadPrevious.HorizontalPointer = pQueueHeadToUnlink.HorizontalPointer
pQueueHeadToUnlink.HorizontalPointer = pQHeadNext
End UnlinkQueueHead

```

If software removes the queue head with the H-bit set, it must select another queue head still linked into the schedule and set its H-bit. This should be completed before removing the queue head. The requirement is that software keep one queue head in the asynchronous schedule, with its H-bit set. At the point software has removed one or more queue heads from the asynchronous schedule, it is unknown whether the host controller has a cached pointer to them. Similarly, it is unknown how long the host controller might retain the cached information, as it is implementation dependent and may be affected by the actual dynamics of the schedule load. Therefore, once software has removed a queue head from the asynchronous list, it must retain the coherency of the queue head (link pointers). It cannot disturb the removed queue heads until it knows that the host controller does not have a local copy of a pointer to any of the removed data structures.

The method software uses to determine when it is safe to modify a removed queue head is to handshake with the host controller. The handshake mechanism allows software to remove items from the asynchronous schedule, then execute a simple, lightweight handshake that is used by software as a key that it can free (or reuse) the memory associated the data structures it has removed from the asynchronous schedule.

The handshake is implemented with three bits in the host controller. The first bit is a command bit (USBCMD[IAA]—interrupt on async advance doorbell) that allows software to inform the host controller that something has been removed from its asynchronous schedule. The second bit is a status bit (USBSTS[AAI]—interrupt on async advance) that the host controller sets after it has released all on-chip state that may potentially reference one of the data structures just removed. When the host controller sets this status bit, it also clears the command bit. The third bit is an interrupt enable (USBINTR[AAE]—interrupt on async advance enable) that is matched with the status bit. If the status bit is set and the interrupt enable bit is set, then the host controller asserts a hardware interrupt.

**Figure 16-49** illustrates a general example where consecutive queue heads (B and C) are unlinked from the schedule using the algorithm above. Before the unlink operation, the host controller has a copy of queue head A.

The unlink algorithm requires that as software unlinks each queue head, the unlinked queue head is loaded with the address of a queue head that will remain in the asynchronous schedule.

When the host controller observes that doorbell bit being set, it makes a note of the local reachable schedule information. In this example, the local reachable schedule information includes both queue heads (A & B). It is sufficient that the host controller can set the status bit (and clear the doorbell bit) as soon as it has traversed beyond current reachable schedule information (that is, traversed beyond queue head (B) in this example).

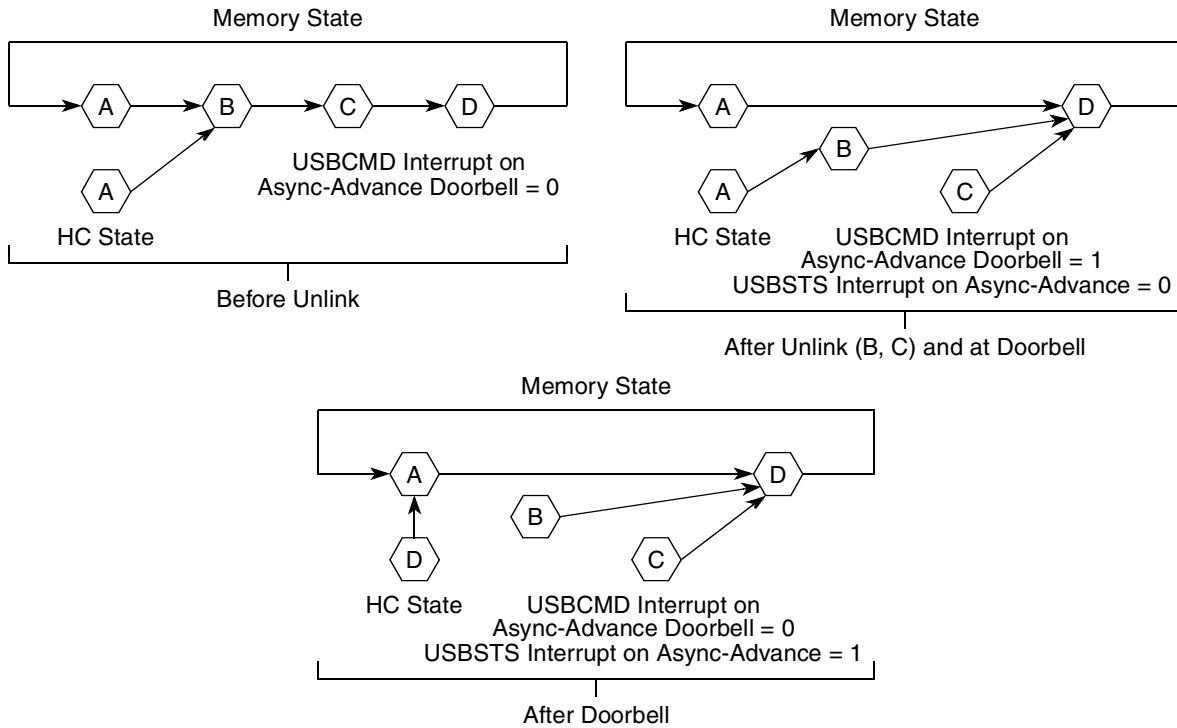


Figure 16-49. Generic Queue Head Unlink Scenario

Alternatively, a host controller implementation is allowed to traverse the entire asynchronous schedule list (for example, observed the head of the queue (twice)) before setting USBSTS[AAI].

Software may re-use the memory associated with the removed queue heads after it observes USBSTS[AAI] is set, following assertion of the doorbell. Software should acknowledge the interrupt on async advance status as indicated in the USBSTS register, before using the doorbell handshake again

### 16.6.9.3 Empty Asynchronous Schedule Detection

EHCI uses two bits to detect when the asynchronous schedule is empty. The queue head data structure (see Figure 16-41) defines an H-bit in the queue head, which allows software to mark a queue head as being the head of the reclaim list. host controller also keeps a 1-bit flag in the USBSTS register (Reclamation) that is cleared when the host controller observes a queue head with the H-bit set. The reclamation flag in the status register is set when any USB transaction from the asynchronous schedule is executed (or whenever the asynchronous schedule starts, see Section 16.6.9.4, “Asynchronous Schedule Traversal: Start Event.”

If the controller ever encounters an H-bit of one and a Reclamation bit of zero, the controller simply stops traversal of the asynchronous schedule.

An example illustrating the H-bit in a schedule is shown in Figure 16-50

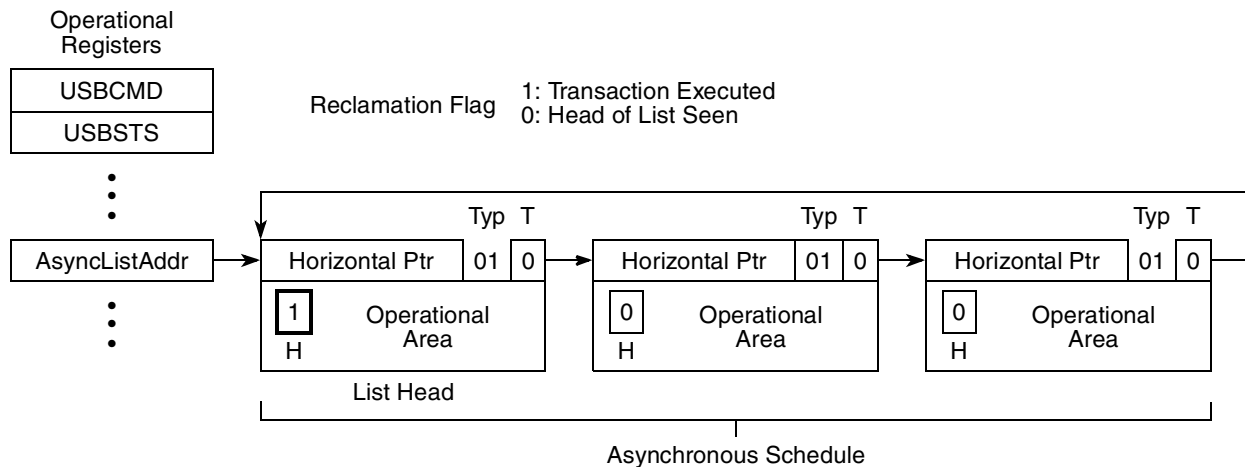


Figure 16-50. Asynchronous Schedule List with Annotation to Mark Head of List

#### 16.6.9.4 Asynchronous Schedule Traversal: Start Event

Once the host controller has idled itself using the empty schedule detection, it naturally activates and begins processing from the Periodic Schedule at the beginning of each micro-frame. In addition, it may have idled itself early in a micro-frame. When this occurs (idles early in the micro-frame) the host controller must occasionally reactivate during the micro-frame and traverse the asynchronous schedule to determine whether any progress can be made. Asynchronous schedule Start Events are defined to be:

- Whenever the host controller transitions from the periodic schedule to the asynchronous schedule. If the periodic schedule is disabled and the asynchronous schedule is enabled, then the beginning of the micro-frame is equivalent to the transition from the periodic schedule, or
- The asynchronous schedule traversal restarts from a sleeping state.

#### 16.6.9.5 Reclamation Status Bit (USBSTS Register)

The operation of the empty asynchronous schedule detection feature depends on the proper management of the Reclamation bit (RCL) in the USBSTS register. The host controller tests for an empty schedule just after it fetches a new queue head while traversing the asynchronous schedule. The host controller sets USBSTS[RCL] whenever an asynchronous schedule traversal Start Event occurs. USBSTS[RCL] is also set whenever the host controller executes a transaction while traversing the asynchronous schedule. The host controller clears USBSTS[RCL] whenever it finds a queue head with its H-bit set. Software should only set a queue head's H-bit if the queue head is in the asynchronous schedule. If software sets the H-bit in an interrupt queue head, the resulting behavior is undefined. The host controller may clear USBSTS[RCL] when executing from the periodic schedule.

#### 16.6.10 Managing Control/Bulk/Interrupt Transfers via Queue Heads

This section presents an overview of how the host controller interacts with queuing data structures.

Queue heads use the Queue Element Transfer Descriptor (qTD) structure defined in [Section 16.5.5, “Queue Element Transfer Descriptor \(qTD\).”](#)

One queue head is used to manage the data stream for one endpoint. The queue head structure contains static endpoint characteristics and capabilities. It also contains a working area from where individual bus transactions for an endpoint are executed. Each qTD represents one or more bus transactions, which is defined in the context of the EHCI specification as a transfer.

The general processing model for the host controller's use of a queue head is simple:

- Read a queue head,
- Execute a transaction from the overlay area,
- Write back the results of the transaction to the overlay area
- Move to the next queue head.


If the host controller encounters errors during a transaction, the host controller will set one of the error reporting bits in the queue head's Status field. The Status field accumulates all errors encountered during the execution of a qTD (that is, the error bits in the queue head Status field are sticky until the transfer (qTD) has completed). This state is always written back to the source qTD when the transfer is complete. On transfer (for example, buffer or halt conditions) boundaries, the host controller must auto-advance (without software intervention) to the next qTD. Additionally, the hardware must be able to halt the queue so no additional bus transactions will occur for the endpoint and the host controller will not advance the queue.

#### 16.6.10.1 Buffer Pointer List Use for Data Streaming with qTDs

A qTD has an array of buffer pointers, which is used to reference the data buffer for a transfer. The EHCI specification requires that the buffer associated with the transfer be virtually contiguous. This means that if the buffer spans more than one physical page, it must obey the following rules:

- The first portion of the buffer must begin at some offset in a page and extend through the end of the page.
- The remaining buffer cannot be allocated in small chunks scattered around memory. For each 4K chunk beyond the first page, each buffer portion matches to a full 4K page. The final portion, which may only be large enough to occupy a portion of a page, must start at the top of the page and be contiguous within that page.

Figure 16-51 illustrates these requirements.



**Figure 16-51. Example Mapping of qTD Buffer Pointers to Buffer Pages**

The buffer pointer list in the qTD is long enough to support a maximum transfer size of 20K bytes. This case occurs when all five buffer pointers are used and the first offset is zero. A qTD handles a 16Kbyte buffer with any starting buffer alignment.

The host controller uses the C\_Page field as an index value to determine which buffer pointer in the list should be used to start the current transaction. The host controller uses a different buffer pointer for each physical page of the buffer. This is always true, even if the buffer is physically contiguous.

The host controller must detect when the current transaction spans a page boundary and automatically move to the next available buffer pointer in the page pointer list. The next available pointer is reached by incrementing C\_Page and pulling the next page pointer from the list. Software must ensure there are sufficient buffer pointers to move the amount of data specified in the Bytes to Transfer field.

Figure 16-51 illustrates a nominal example of how System software would initialize the buffer pointers list and the C\_Page field for a transfer size of 16383 bytes. C\_Page is cleared. The upper 20-bits of Page 0 references the start of the physical page. Current Offset (the lower 12-bits of queue head Dword 7) holds the offset in the page for example, 2049 (for example, 4096-2047). The remaining page pointers are set to reference the beginning of each subsequent 4K page.

For the first transaction on the qTD (assuming a 512-byte transaction), the host controller uses the first buffer pointer (page 0 because C\_Page is cleared) and concatenates the Current Offset field. The 512 bytes are moved during the transaction, the Current Offset and Total Bytes to Transfer are adjusted by 512 and written back to the queue head working area.

During the 4th transaction, the host controller needs 511 bytes in page 0 and one byte in page 1. The host controller will increment C\_Page (to 1) and use the page 1 pointer to move the final byte of the transaction. After the 4th transaction, the active page pointer is the page 1 pointer and Current Offset has rolled to one, and both are written back to the overlay area. The transactions continue for the rest of the buffer, with the

host controller automatically moving to the next page pointer (that is, C\_Page) when necessary. There are three conditions for how the host controller handles C\_Page.

- The current transaction does not span a page boundary. The value of C\_Page is not adjusted by the host controller.
- The current transaction does span a page boundary. The host controller must detect the page cross condition and advance to the next buffer while streaming data to/from the USB.
- The current transaction completes on a page boundary (that is, the last byte moved for the current transaction is the last byte in the page for the current page pointer). The host controller must increment C\_Page before writing back status for the transaction.

Note that the only valid adjustment the host controller may make to C\_Page is to increment by one.

### 16.6.10.2 Adding Interrupt Queue Heads to the Periodic Schedule

The link path(s) from the periodic frame list to a queue head establishes in which frames a transaction can be executed for the queue head. Queue heads are linked into the periodic schedule so they are polled at the appropriate rate. System software sets a bit in a queue head's S-Mask to indicate which micro-frame within a 1 millisecond period a transaction should be executed for the queue head. Software must ensure that all queue heads in the periodic schedule have S-Mask set to a non-zero value. An S-mask with a zero value in the context of the periodic schedule yields undefined results.

If the desired poll rate is greater than one frame, system software can use a combination of queue head linking and S-Mask values to spread interrupts of equal poll rates through the schedule so that the periodic bandwidth is allocated and managed in the most efficient manner possible. Some examples are illustrated in [Table 16-66](#).

**Table 16-66. Example Periodic Reference Patterns for Interrupt Transfers**

Frame # Reference Sequence	Description
0, 2, 4, 6, 8, ... S-Mask = 0x01	A queue head for the bInterval of 2 milliseconds (16 micro-frames) is linked into the periodic schedule so that it is reachable from the periodic frame list locations indicated in the previous column. In addition, the S-Mask field in the queue head is set to 0x01, indicating that the transaction for the endpoint should be executed on the bus during micro-frame 0 of the frame.
0, 2, 4, 6, 8, ... S-Mask = 0x02	Another example of a queue head with a bInterval of 2 milliseconds is linked into the periodic frame list at exactly the same interval as the previous example. However, the S-Mask is set to 0x02 indicating that the transaction for the endpoint should be executed on the bus during micro-frame 1 of the frame.

### 16.6.10.3 Managing Transfer Complete Interrupts from Queue Heads

The host controller sets an interrupt to be signaled at the next interrupt threshold when the completed transfer (qTD) has an Interrupt on Complete (IOC) bit set, or whenever a transfer (qTD) completes with a short packet. If system software needs multiple qTDs to complete a client request (that is, like a control transfer) the intermediate qTDs do not require interrupts. System software may only need a single interrupt to notify it that the complete buffer has been transferred. System software may set IOC's to occur more frequently. A motivation for this may be that it wants early notification so that interface data structures can be re-used in a timely manner.



## 16.6.11 Ping Control

USB 2.0 defines an addition to the protocol for high-speed devices called Ping. Ping is required for all USB 2.0 High-speed bulk and control endpoints. Ping is not allowed for a split-transaction stream. This extension to the protocol eliminates the bad side-effects of Naking OUT endpoints. The Status field has a Ping State bit, which the host controller uses to determine the next actual PID it will use in the next transaction to the endpoint (see [Table 16-55](#)). The Ping State bit is only managed by the host controller for queue heads that meet all of the following criteria:

- The queue head is not an interrupt
- The EPS field equals High-Speed
- The PIDCode field equals OUT

[Table 16-67](#) illustrates the state transition table for the host controller's responsibility for maintaining the PING protocol. Refer to Chapter 8 in the *USB Specification, Revision 2.0* for detailed description on the Ping protocol.

**Table 16-67. Ping Control State Transition Table**

Current	Event		Next
	Host	Device	
Do Ping	PING	Nak	Do Ping
Do Ping	PING	Ack	Do OUT
Do Ping	PING	XactErr <sup>1</sup>	Do Ping
Do Ping	PING	Stall	N/C <sup>2</sup>
Do OUT	OUT	Nak	Do Ping
Do OUT	OUT	Nyet	Do Ping <sup>3</sup>
Do OUT	OUT	Ack	Do OUT
Do OUT	OUT	XactErr <sup>1</sup>	Do Ping
Do OUT	OUT	Stall	N/C <sup>2</sup>

<sup>1</sup> Transaction Error (XactErr) is any time the host misses the handshake.

<sup>2</sup> No transition change required for the Ping State bit. The Stall handshake results in the endpoint being halted (for example, Active cleared and Halt set). Software intervention is required to restart queue.

<sup>3</sup> A Nyet response to an OUT means that the device has accepted the data, but cannot receive any more at this time. Host must advance the transfer state and additionally, transition the Ping State bit to Do Ping.

The Ping State bit is described in [Table 16-55](#). The defined ping protocol allows the host to be imprecise on the initialization of the ping protocol (that is, start in Do OUT when we don't know whether there is space on the device or not). The host controller manages the Ping State bit. System software sets the initial value in the queue head when it initializes a queue head. The host controller preserves the Ping State bit across all queue advancements. This means that when a new qTD is written into the queue head overlay area, the previous value of the Ping State bit is preserved.

## 16.6.12 Split Transactions

USB 2.0 defines extensions to the bus protocol for managing USB 1.x data streams through USB 2.0 hubs. This section describes how the host controller uses the interface data structures to manage data streams with full- and low-speed devices, connected below a USB 2.0 hub, utilizing the split transaction protocol. Refer to the USB 2.0 Specification for the complete definition of the split transaction protocol. Full- and low-speed devices are enumerated identically as high-speed devices, but the transactions to the full- and low-speed endpoints use the split-transaction protocol on the high-speed bus. The split transaction protocol is an encapsulation of (or wrapper around) the full- or low-speed transaction. The high-speed wrapper portion of the protocol is addressed to the USB 2.0 hub and transaction translator below which the full- or low-speed device is attached.

EHCI uses dedicated data structures for managing full-speed isochronous data streams. Control, Bulk and Interrupt are managed using the queuing data structures. The interface data structures need to be programmed with the device address and the transaction translator number of the USB 2.0 hub operating as the low-/full-speed host controller for this link. The following sections describe the details of how the host controller processes and manages the split transaction protocol.

### 16.6.12.1 Split Transactions for Asynchronous Transfers

A queue head in the asynchronous schedule with an EPS field indicating a full-or low-speed device indicates to the host controller that it must use split transactions to stream data for this queue head. All full-speed bulk and full-, low-speed control are managed via queue heads in the asynchronous schedule.

Software must initialize the queue head with the appropriate device address and port number for the transaction translator that is serving as the full-/low-speed host controller for the links connecting the endpoint. Software must also initialize the split transaction state bit (SplitXState) to Do-Start-Split. Finally, if the endpoint is a control endpoint, then system software must set the Control Transfer Type (C) bit in the queue head to a one. If this is not a control transfer type endpoint, the C bit must be initialized by software to be a zero. This information is used by the host controller to properly set the Endpoint Type (ET) field in the split transaction bus token. When the C bit is a zero, the split transaction token's ET field is set to indicate a bulk endpoint. When the C bit is a one, the split transaction token's ET field is set to indicate a control endpoint. Refer to Chapter 8 of *USB Specification, Revision 2.0* for details.




Figure 16-52. Host Controller Asynchronous Schedule Split-Transaction State Machine

### 16.6.12.1.1 Asynchronous—Do-Start-Split

Do-Start-Split is the state which software must initialize a full- or low-speed asynchronous queue head. This state is entered from the Do-Complete-Split state only after a complete-split transaction receives a valid response from the transaction translator that is not a Nyet handshake.

For queue heads in this state, the host controller executes a start-split transaction to the transaction translator. If the bus transaction completes without an error and PID Code indicates an IN or OUT transaction, then the host controller reloads the error counter (Cerr). If it is a successful bus transaction and the PID Code indicates a SETUP, the host controller will not reload the error counter. If the transaction translator responds with a Nak, the queue head is left in this state, and the host controller proceeds to the next queue head in the asynchronous schedule.

If the host controller times out the transaction (no response, or bad response) the host controller decrements Cerr and proceeds to the next queue head in the asynchronous schedule.

### 16.6.12.1.2 Asynchronous—Do-Complete-Split

This state is entered from the Do-Start-Split state only after a start-split transaction receives an Ack handshake from the transaction translator.

For queue heads in this state, the host controller executes a complete-split transaction to the transaction translator. If the transaction translator responds with a Nyet handshake, the queue head is left in this state, the error counter is reset and the host controller proceeds to the next queue head in the asynchronous schedule. When a Nyet handshake is received for a bus transaction where the queue head's PID Code indicates an IN or OUT, the host controller reloads the error counter (Cerr). When a Nyet handshake is received for a complete-split bus transaction where the queue head's PID Code indicates a SETUP, the host controller must not adjust the value of Cerr.

Independent of PID Code, the following responses have the indicated effects:

- Transaction Error (XactErr). Timeout/data CRC failure. The error counter (Cerr) is decremented by one and the complete split transaction is immediately retried (if possible). If there is not enough time in the micro-frame to execute the retry, the host controller ensures that the next time the host controller begins executing from the Asynchronous schedule, it must begin executing from this queue head. If another start-split (for some other endpoint) is sent to the transaction translator before the complete-split is really completed, the transaction translator could dump the results (which were never delivered to the host). This is why the core specification states the retries must be immediate. When the host controller returns to the asynchronous schedule in the next micro-frame, the first transaction from the schedule will be the retry for this endpoint. If Cerr went to zero, the host controller halts the queue.
- NAK. The target endpoint Nak'd the full- or low-speed transaction. The state of the transfer is not advanced and the state is exited. If the PID Code is a SETUP, then the Nak response is a protocol error. The XactErr status bit is set and the Cerr field is decremented.
- STALL. The target endpoint responded with a STALL handshake. The host controller sets the halt bit in the status byte, retires the qTD but does not attempt to advance the queue.

If the PID Code indicates an IN, then any of following responses are expected:

- **DATA0/1.** On reception of data, the host controller ensures the PID matches the expected data toggle and checks CRC. If the packet is good, the host controller advances the state of the transfer (for example, moves the data pointer by the number of bytes received, decrements the BytesToTransfer field by the number of bytes received, and toggles the dt bit). The host controller then exits this state. The response and advancement of transfer may trigger other processing events, such as retirement of the qTD and advancement of the queue.

If the data sequence PID does not match the expected, the data is ignored, the transfer state is not advanced and this state is exited.

If the PID Code indicates an OUT/SETUP, then any of following responses are expected:

- **ACK.** The target endpoint accepted the data, so the host controller must advance the state of the transfer. The Current Offset field is incremented by Maximum Packet Length or Bytes to Transfer, whichever is less. The Bytes To Transfer field is decremented by the same amount and the data toggle bit (dt) is toggled. The host controller then exits this state.

Advancing the transfer state may cause other processing events such as retirement of the qTD and advancement of the queue.


## 16.6.12.2 Split Transaction Interrupt

Split-transaction Interrupt-IN/OUT endpoints are managed using the same data structures used for high-speed interrupt endpoints. They both co-exist in the periodic schedule. Queue heads/qTDs offer the set of features required for reliable data delivery, which is characteristic to interrupt transfer types. The split-transaction protocol is managed completely within this defined functional transfer framework. For example, for a high-speed endpoint, the host controller will visit a queue head, execute a high-speed transaction (if criteria are met) and advance the transfer state (or not) depending on the results of the entire transaction. For low- and full-speed endpoints, the details of the execution phase are different (that is, takes more than one bus transaction to complete), but the remainder of the operational framework is intact.

### 16.6.12.2.1 Split Transaction Scheduling Mechanisms for Interrupt

Full- and low-speed Interrupt queue heads have an EPS field indicating full- or low-speed and have a non-zero S-mask field. The host controller can detect this combination of parameters and assume the endpoint is a periodic endpoint. Low- and full-speed interrupt queue heads require the use of the split transaction protocol. The host controller sets the Endpoint Type (ET) field in the split token to indicate the transaction is an interrupt. These transactions are managed through a transaction translator's periodic pipeline. Software should not set these fields to indicate the queue head is an interrupt unless the queue head is used in the periodic schedule.

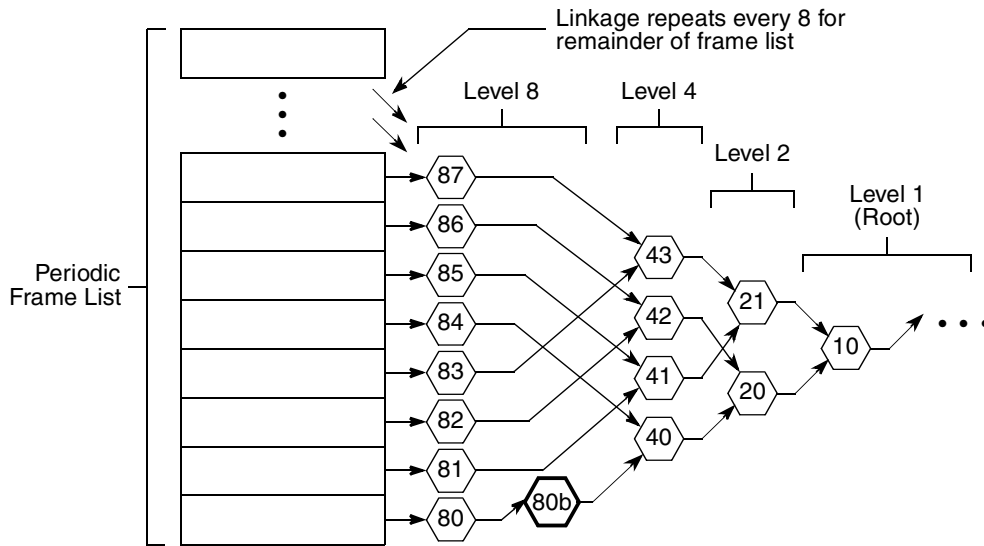
System software manages the per/transaction translator periodic pipeline by budgeting and scheduling exactly during which micro-frames the start-splits and complete-splits for each endpoint will occur. The characteristics of the transaction translator are such that the high-speed transaction protocol must execute during explicit micro-frames, or the data or response information in the pipeline is lost. [Figure 16-53](#) illustrates the general scheduling boundary conditions that are supported by the EHCI periodic schedule and queue head data structure. The S and  $C_n$  labels indicate micro-frames where software can schedule start-splits and complete splits (respectively).



**Figure 16-53. Split Transaction, Interrupt Scheduling Boundary Conditions**

The scheduling cases are:

- Case 1: The normal scheduling case is where the entire split transaction is completely bounded by a frame (H-Frame in this case).
- Case 2a through Case 2c: The USB 2.0 hub pipeline rules states clearly, when and how many complete-splits must be scheduled to account for earliest to latest execution on the full/low-speed link. The complete-splits may span the H-Frame boundary when the start-split is in micro-frame 4 or later. When this occurs, the H-Frame to B-Frame alignment requires that the queue head be reachable from consecutive periodic frame list locations. System software cannot build an efficient schedule that satisfies this requirement unless it uses FSTNs. [Figure 16-54](#) illustrates the general layout of the periodic schedule.



**Figure 16-54. General Structure of EHCI Periodic Schedule Utilizing Interrupt Spreading**

The periodic frame list is effectively the leaf level a binary tree, which is always traversed leaf to root. Each level in the tree corresponds to a  $2^N$  poll rate. Software can efficiently manage periodic bandwidth on the USB by spreading interrupt queue heads that have the same poll rate requirement across all the available paths from the frame list. For example, system software can schedule eight poll rate 8 queue heads and account for them once in the high-speed bus bandwidth allocation.

When an endpoint is allocated an execution footprint that spans a frame boundary, the queue head for the endpoint must be reachable from consecutive locations in the frame list. An example would be if  $8_{0b}$  were such an endpoint. Without additional support on the interface, to get  $8_{0b}$  reachable at the correct time, software would have to link  $8_1$  to  $8_{0b}$ . It would then have to move  $4_1$  and everything linked after into the same path as  $4_0$ . This upsets the integrity of the binary tree and disallows the use of the spreading technique.

FSTN data structures are used to preserve the integrity of the binary-tree structure and enable the use of the spreading technique. [Section 16.5.7, “Periodic Frame Span Traversal Node \(FSTN\),”](#) defines the hardware and software operational model requirements for using FSTNs.

The following queue head fields are initialized by system software to instruct the host controller when to execute portions of the split-transaction protocol.

- **SplitXState.** This is a single bit residing in the Status field of a queue head ([Table 16-55](#)). This bit is used to track the current state of the split transaction.
- **Frame S-mask.** This is a bit-field where-in system software sets a bit corresponding to the micro-frame (within an H-Frame) that the host controller should execute a start-split transaction. This is always qualified by the value of the SplitXState bit in the Status field of the queue head. For example, referring to [Figure 16-53](#), case one, the S-mask would have a value of `0b0000_0001` indicating that if the queue head is traversed by the host controller, and the SplitXState indicates Do\_Start, and the current micro-frame as indicated by `FRINDEX[2-0]` is 0, then execute a start-split transaction.

- **Frame C-mask.** This is a bit-field where system software sets one or more bits corresponding to the micro-frames (within an H-Frame) that the host controller should execute complete-split transactions. The interpretation of this field is always qualified by the value of the SplitXState bit in the Status field of the queue head. For example, referring to [Figure 16-53](#), case one, the C-mask would have a value of 0b0001\_1100 indicating that if the queue head is traversed by the host controller, and the SplitXState indicates Do\_Complete, and the current micro-frame as indicated by FRINDEX[2–0] is 2, 3, or 4, then execute a complete-split transaction. It is software's responsibility to ensure that the translation between H-Frames and B-Frames is correctly performed when setting bits in S-mask and C-mask.

### 16.6.12.2.2 Host Controller Operational Model for FSTNs

The FSTN data structure is used to manage Low/Full-speed interrupt queue heads that need to be reached from consecutive frame list locations (that is, boundary cases 2a through 2c). An FSTN is essentially a back pointer, similar in intent to the back pointer field in the siTD data structure.

This feature provides software a simple primitive to save a schedule position, redirect the host controller to traverse the necessary queue heads in the previous frame, then restore the original schedule position and complete normal traversal.

There are four components to the use of FSTNs:

- FSTN data structure, defined in [Section 16.5.7](#), “[Periodic Frame Span Traversal Node \(FSTN\)](#).”
- A Save Place indicator; this is always an FSTN with its Back Path Link Pointer[T] bit cleared.
- A Restore indicator; this is always an FSTN with its Back Path Link Pointer[T] bit set.
- Host controller FSTN traversal rules.

When the host controller encounters an FSTN during micro-frames 2 through 7 it simply follows the node's Normal Path Link Pointer to access the next schedule data structure. Note that the FSTN's Normal Path Link Pointer[T] bit may set, which the host controller must interpret as the end of periodic list mark.

When the host controller encounters a Save-Place FSTN in micro-frames 0 or 1, it saves the value of the Normal Path Link Pointer and sets an internal flag indicating that it is executing in Recovery Path mode. Recovery Path mode modifies the host controller's rules for how it traverses the schedule and limits which data structures are considered for execution of bus transactions. The host controller continues executing in Recovery Path mode until it encounters a Restore FSTN or it determines that it has reached the end of the micro-frame.

The rules for schedule traversal and limited execution while in Recovery Path mode are:


- Always follow the Normal Path Link Pointer when it encounters an FSTN that is a Save-Place indicator. The host controller must not recursively follow Save-Place FSTNs. Therefore, while executing in Recovery Path mode, it must never follow an FSTN's Back Path Link Pointer.
- Do not process an siTD or iTD data structure; simply follow its Next Link Pointer.
- Do not process a QH (Queue Head) whose EPS field indicates a high-speed device; simply follow its Horizontal Link Pointer.
- When a QH's EPS field indicates a Full/Low-speed device, the host controller only considers it for execution if its SplitXState is DoComplete (note: this applies whether the PID Code indicates an

IN or an OUT). Refer to the *EHCI Specification* for a complete list of additional conditions that must be met in general for the host controller to issue a bus transaction. Note that the host controller must not execute a Start-split transaction while executing in Recovery Path mode. Refer to the *EHCI Specification* for special handling when in Recovery Path mode.

- Stop traversing the recovery path when it encounters an FSTN that is a Restore indicator. The host controller unconditionally uses the saved value of the Save-Place FSTN's Normal Path Link Pointer when returning to the normal path traversal. The host controller must clear the context of executing a Recovery Path when it restores schedule traversal to the Save-Place FSTN's Normal Path Link Pointer.

If the host controller determines that there is not enough time left in the micro-frame to complete processing of the periodic schedule, it abandons traversal of the recovery path, and clears the context of executing a recovery path. The result is that at the start of the next consecutive micro-frame, the host controller starts traversal at the frame list.

An example traversal of a periodic schedule that includes FSTNs is illustrated in Figure 16-55.



**Figure 16-55. Example Host Controller Traversal of Recovery Path via FSTNs**

In frame N (micro-frames 0-7), for this example, the host controller traverses all of the schedule data structures utilizing the Normal Path Link Pointers in any FSTNs it encounters. This is because the host controller has not yet encountered a Save-Place FSTN so it is not executing in Recovery Path mode. When it encounters the Restore FSTN, (Restore-N), during micro-frames 0 and 1, it uses Restore-N. Normal Path Link Pointer to traverse to the next data structure (that is, normal schedule traversal). This is because the host controller must use a Restore FSTN's Normal Path Link Pointer when not executing in a Recovery-Path mode. The nodes traversed during frame N include: {82.0, 82.1, 82.2, 82.3, 42, 20, Restore-N, 10 ... }.

In frame N+1 (micro-frames 0 and 1), when the host controller encounters Save-Place FSTN (Save-N), it observes that Save-N.Back Path Link Pointer.T-bit is zero (definition of a Save-Place indicator). The host controller saves the value of Save-N. Normal Path Link Pointer and follows Save-N.Back Path Link



Pointer. At the same time, it sets an internal flag indicating that it is now in Recovery Path mode (the recovery path is annotated in [Figure 16-55](#) with a large dashed line). The host controller continues traversing data structures on the recovery path and executing only those bus transactions as noted above, on the recovery path until it reaches Restore FSTN (Restore-N). Restore-N.Back Path Link Pointer.T-bit is set (definition of a Restore indicator), so the host controller exits Recovery Path mode by clearing the internal Recovery Path mode flag and commences (restores) schedule traversal using the saved value of the Save-Place FSTN's Normal Path Link Pointer (for example, Save-N.Normal Path Link Pointer). The nodes traversed during these micro-frames include: {8<sub>3,0</sub>, 8<sub>3,1</sub>, 8<sub>3,2</sub>, Save-A, 8<sub>2,2</sub>, 8<sub>2,3</sub>, 4<sub>2</sub>, 2<sub>0</sub>, Restore-N, 4<sub>3</sub>, 2<sub>1</sub>, Restore-N, 10 ...}.

In frame N+1 (micro-frames 2-7), when the host controller encounters Save-Path FSTN Save-N, it unconditionally follows Save-N.Normal Path Link Pointer. The nodes traversed during these micro-frames include: {8<sub>3,0</sub>, 8<sub>3,1</sub>, 8<sub>3,2</sub>, Save-A, 4<sub>3</sub>, 2<sub>1</sub>, Restore-N, 1<sub>0</sub> ...}.

### 16.6.12.2.3 Software Operational Model for FSTNs

Software must create a consistent, coherent schedule for the host controller to traverse. When using FSTNs, system software must adhere to the following rules:

- Each Save-Place indicator requires a matching Restore indicator.  
The Save-Place indicator is an FSTN with a valid Back Path Link Pointer and T-bit equal to zero. Note that Back Path Link Pointer[Typ] field must be set to indicate the referenced data structure is a queue head. The Restore indicator is an FSTN with its Back Path Link Pointer[T] bit set.  
A Restore FSTN may be matched to one or more Save-Place FSTNs. For example, if the schedule includes a poll-rate 1 level, then system software only needs to place a Restore FSTN at the beginning of this list in order to match all possible Save-Place FSTNs.
- If the schedule does not have elements linked at a poll-rate level of one, and one or more Save-Place FSTNs are used, then System Software must ensure the Restore FSTN's Normal Path Link Pointer's T-bit is set, as this will be use to mark the end of the periodic list.
- When the schedule does have elements linked at a poll rate level of one, a Restore FSTN must be the first data structure on the poll rate one list. All traversal paths from the frame list converge on the poll-rate one list. System software must ensure that Recovery Path mode is exited before the host controller is allowed to traverse the poll rate level one list.
- A Save-Place FSTN's Back Path Link Pointer must reference a queue head data structure. The referenced queue head must be reachable from the previous frame list location. In other words, if the Save-Place FSTN is reachable from frame list offset N, then the FSTN's Back Path Link Pointer must reference a queue head that is reachable from frame list offset N-1.

Software should make the schedule as efficient as possible. What this means in this context is that software should have no more than one Save-Place FSTN reachable in any single frame. Note there will be times when two (or more, depending on the implementation) could exist as full-/low-speed footprints change with bandwidth adjustments. This could occur, for example when a bandwidth rebalance causes system software to move the Save-Place FSTN from one poll rate level to another. During the transition, software must preserve the integrity of the previous schedule until the new schedule is in place.

#### 16.6.12.2.4 Tracking Split Transaction Progress for Interrupt Transfers

To correctly maintain the data stream, the host controller must be able to detect and report errors where data is lost. For interrupt-IN transfers, data is lost when it makes it into the USB 2.0 hub, but the USB 2.0 host system is unable to get it from the USB 2.0 hub and into the system before it expires from the transaction translator pipeline. When a lost data condition is detected, the queue is halted, thus signaling system software to recover from the error. A data-loss condition exists whenever a start-split is issued, accepted and successfully executed by the USB 2.0 hub, but the complete-splits get unrecoverable errors on the high-speed link, or the complete-splits do not occur at the correct times. One reason complete-splits might not occur at the right time would be due to host-induced system hold-offs that cause the host controller to miss bus transactions because it cannot get timely access to the schedule in system memory.

The same condition can occur for an interrupt-OUT, but the result is not an endpoint halt condition, but rather effects only the progress of the transfer. The queue head has the following fields to track the progress of each split transaction. These fields are used to keep incremental state about which (and when) portions have been executed.

- **C-prog-mask.** This is an eight-bit bit-vector where the host controller keeps track of which complete-splits have been executed. Due to the nature of the transaction translator periodic pipeline, the complete-splits need to be executed in-order. The host controller needs to detect when the complete-splits have not been executed in order. This can only occur due to system hold-offs where the host controller cannot get to the memory-based schedule. C-prog-mask is a simple bit-vector that the host controller sets one of the C-prog-mask bits for each complete-split executed. The bit position is determined by the micro-frame number in which the complete-split was executed. The host controller always checks C-prog-mask before executing a complete-split transaction. If the previous complete-splits have not been executed then it means one (or more) have been skipped and data has potentially been lost.
- **FrameTag.** This field is used by the host controller during the complete-split portion of the split transaction to tag the queue head with the frame number (H-Frame number) when the next complete split must be executed.
- **S-bytes.** This field can be used to store the number of data payload bytes sent during the start-split (if the transaction was an OUT). The S-bytes field must be used to accumulate the data payload bytes received during the complete-splits (for an IN).

#### 16.6.12.2.5 Split Transaction Execution State Machine for Interrupt

In the following section, all references to micro-frame are in the context of a micro-frame within an H-Frame.

As with asynchronous Full- and Low-speed endpoints, a split-transaction state machine is used to manage the split transaction sequence. Aside from the fields defined in the queue head for scheduling and tracking the split transaction, the host controller calculates one internal mechanism that is also used to manage the split transaction. The internal calculated mechanism is:

- **cMicroFrameBit.** This is a single-bit encoding of the current micro-frame number. It is an eight-bit value calculated by the host controller at the beginning of every micro-frame. It is calculated from the three least significant bits of the FRINDEX register (that is,  $cMicroFrameBit = (1 \text{ shifted-left}(FRINDEX[2-0]))$ ). The cMicroFrameBit has at most one bit asserted, which always

corresponds to the current micro-frame number. For example, if the current micro-frame is 0, then cMicroFrameBit will equal 0b0000\_0001.

The variable cMicroFrameBit is used to compare against the S-mask and C-mask fields to determine whether the queue head is marked for a start- or complete-split transaction for the current micro-frame.

Figure 16-56 illustrates how a complete interrupt split transaction is managed. There are two phases to each split transaction. The first is a single start-split transaction, which occurs when the SplitXState is at Do\_Start and the single bit in cMicroFrameBit has a corresponding bit active in QH[S-mask]. The transaction translator does not acknowledge the receipt of the periodic start-split, so the host controller unconditionally transitions the state to Do\_Complete. Due to the available jitter in the transaction translator pipeline, there will be more than one complete-split transaction scheduled by software for the Do\_Complete state. This translates simply to the fact that there are multiple bits set in the QH[C-mask] field.

The host controller keeps the queue head in the Do\_Complete state until the split transaction is complete (see definition below), or an error condition triggers the three-strikes-rule (for example, after the host tries the same transaction three times, and each encounters an error, the host controller stops retrying the bus transaction and halts the endpoint, thus requiring system software to detect the condition and perform system-dependent recovery).




Figure 16-56. Split Transaction State Machine for Interrupt

### 16.6.12.2.6 Periodic Interrupt—Do-Start-Split

This is the state software must initialize a full- or low-speed interrupt queue head StartXState bit. This state is entered from the Do\_Complete Split state only after the split transaction is complete. This occurs when

one of the following events occur: The transaction translator responds to a complete-split transaction with one of the following:

- **NAK.** A NAK response is a propagation of the full- or low-speed endpoint's NAK response.
- **ACK.** An ACK response is a propagation of the full- or low-speed endpoint's ACK response. Only occurs on an OUT endpoint.
- **DATA 0/1.** Only occurs for INs. Indicates that this is the last of the data from the endpoint for this split transaction.
- **ERR.** The transaction on the low-/full-speed link below the transaction translator had a failure (for example, timeout, bad CRC, etc.).
- **NYET (and Last).** The host controller issued the last complete-split and the transaction translator responded with a NYET handshake. This means that the start-split was not correctly received by the transaction translator, so it never executed a transaction to the full- or low-speed endpoint, see Section Periodic Interrupt - Do Complete Split for the definition of 'Last'.

Each time the host controller visits a queue head in this state (once within the Execute Transaction state), bit-wise ANDs QH[S-mask] with cMicroFrameBit to determine whether to execute a start-split. If the result is non-zero, then the host controller issues a start-split transaction. If the PID Code field indicates an IN transaction, the host controller must zero-out the QH[S-bytes] field. After the split-transaction has been executed, the host controller sets up state in the queue head to track the progress of the complete-split phase of the split transaction. Specifically, it records the expected frame number into QH[FrameTag] field, sets C-prog-mask to zero (0x00), and exits this state. Note that the host controller must not adjust the value of Cerr as a result of completion of a start-split transaction.

#### 16.6.12.2.7 Periodic Interrupt—Do-Complete-Split

This state is entered unconditionally from the Do Start Split state after a start-split transaction is executed on the bus. Each time the host controller visits a queue head in this state (once within the Execute Transaction state), it checks to determine whether a complete-split transaction should be executed now.

There are four tests to determine whether a complete-split transaction should be executed.

- **Test A.** cMicroFrameBit is bit-wise ANDed with QH[C-mask] field. A non-zero result indicates that software scheduled a complete-split for this endpoint, during this micro-frame.
- **Test B.** QH[FrameTag] is compared with the current contents of FRINDEX[7–3]. An equal indicates a match.
- **Test C.** The complete-split progress bit vector is checked to determine whether the previous bit is set, indicating that the previous complete-split was appropriately executed. An example algorithm for this test is provided below:

```
Algorithm Boolean CheckPreviousBit(QH.C-prog-mask, QH.C-mask, cMicroFrameBit)
Begin
-- Return values:
-- TRUE - no error
-- FALSE - error
--
Boolean rvalue = TRUE;
previousBit = cMicroframeBit logical-rotate-right(1)
-- Bit-wise anding previousBit with C-mask indicates
-- whether there was an intent
```

```

-- to send a complete split in the previous micro-frame. So,
-- if the
-- 'previous bit' is set in C-mask, check C-prog-mask to
-- make sure it
-- happened.
If (previousBit bitAND QH.C-mask) then
    If not(previousBit bitAND QH.C-prog-mask) then
        rvalue = FALSE;
    End if
End If
-- If the C-prog-mask already has a one in this bit position,
-- then an aliasing
-- error has occurred. It will probably get caught by the
-- FrameTag Test, but
-- at any rate it is an error condition that as detectable here
-- should not allow
-- a transaction to be executed.
If (cMicroFrameBit bitAND QH.C-prog-mask) then
    rvalue = FALSE;
End if
return (rvalue)
End Algorithm

```

- Test D. Check to see if a start-split should be executed in this micro-frame. Note this is the same test performed in the Do Start Split state. Whenever it evaluates to TRUE and the controller is NOT processing in the context of a Recovery Path mode, it means a start-split should occur in this micro-frame. Test D and Test A evaluating to TRUE at the same time is a system software error. Behavior is undefined.

If (A .and. B .and. C .and. not(D)) then the host controller will execute a complete-split transaction. When the host controller commits to executing the complete-split transaction, it updates QH[C-prog-mask] by bit-ORing with cMicroFrameBit. On completion of the complete-split transaction, the host controller records the result of the transaction in the queue head and sets QH[FrameTag] to the expected H-Frame number. The effect to the state of the queue head and thus the state of the transfer depends on the response by the transaction translator to the complete-split transaction. The following responses have the effects (note that any responses that result in decrementing of the Cerr will result in the queue head being halted by the host controller if the result of the decrement is zero):

- NYET (and Last). On each NYET response, the host controller checks to determine whether this is the last complete-split for this split transaction. Last is defined in this context as the condition where all of the scheduled complete-splits have been executed. If it is the last complete-split (with a NYET response), then the transfer state of the queue head is not advanced (never received any data) and this state exited. The transaction translator must have responded to all the complete-splits with NYETs, meaning that the start-split issued by the host controller was not received. The start-split should be retried at the next poll period.
- The test for whether this is the Last complete split can be performed by XOR QH[C-mask] with QH[C-prog-mask]. If the result is all zeros then all complete-splits have been executed. When this condition occurs, the XactErr status bit is set and the Cerr field is decremented.
- NYET (and not Last). See above description for testing for Last. The complete-split transaction received a NYET response from the transaction translator. Do not update any transfer state (except for C-prog-mask and FrameTag) and stay in this state. The host controller must not adjust Cerr on this response.

- Transaction Error (XactErr). Timeout, data CRC failure, etc. The Cerr field is decremented and the XactErr bit in the Status field is set. The complete split transaction is immediately retried (if Cerr is non-zero). If there is not enough time in the micro-frame to complete the retry and the endpoint is an IN, or Cerr is decremented to a zero from a one, the queue is halted. If there is not enough time in the micro-frame to complete the retry and the endpoint is an OUT and Cerr is not zero, then this state is exited (that is, return to Do Start Split). This results in a retry of the entire OUT split transaction, at the next poll period. Refer to Chapter 11 Hubs (specifically the section on full- and low-speed interrupts) in the *USB Specification Revision 2.0* for detailed requirements on why these errors must be immediately retried.
- ACK. This can only occur if the target endpoint is an OUT. The target endpoint ACK'd the data and this response is a propagation of the endpoint ACK up to the host controller. The host controller must advance the state of the transfer. The Current Offset field is incremented by Maximum Packet Length or Bytes to Transfer, whichever is less. The field Bytes To Transfer is decremented by the same amount. And the data toggle bit (dt) is toggled. The host controller will then exit this state for this queue head. The host controller must reload Cerr with maximum value on this response. Advancing the transfer state may cause other process events such as retirement of the qTD and advancement of the queue.
- MDATA. This response will only occur for an IN endpoint. The transaction translator responded with zero or more bytes of data and an MDATA PID. The incremental number of bytes received is accumulated in QH[S-bytes]. The host controller must not adjust Cerr on this response.
- DATA0/1. This response may only occur for an IN endpoint. The number of bytes received is added to the accumulated byte count in QH[S-bytes]. The state of the transfer is advanced by the result and the host controller exits this state for this queue head.
- Advancing the transfer state may cause other processing events such as retirement of the qTD and advancement of the queue.
- If the data sequence PID does not match the expected, the entirety of the data received in this split transaction is ignored, the transfer state is not advanced and this state is exited.
- NAK. The target endpoint Nak'd the full- or low-speed transaction. The state of the transfer is not advanced, and this state is exited. The host controller must reload Cerr with maximum value on this response.
- ERR. There was an error during the full- or low-speed transaction. The ERR status bit is set, Cerr is decremented, the state of the transfer is not advanced, and this state is exited.
- STALL. The queue is halted (an exit condition of the Execute Transaction state). The status field bits: Active bit is cleared and the Halted bit is set and the qTD is retired. Responses which are not enumerated in the list or which are received out of sequence are illegal and may result in undefined host controller behavior. The other possible combinations of tests A, B, C, and D may indicate that data or response was lost. [Table 16-68](#) lists the possible combinations and the appropriate action.

**Table 16-68. Interrupt IN/OUT Do Complete Split State Execution Criteria**

Condition	Action	Description
not(A) not(D)	Ignore QHD	Neither a start nor complete-split is scheduled for the current micro-frame. Host controller should continue walking the schedule.
A not(C)	If PIDCode = IN Halt QHD If PIDCode = OUT Retry start-split	Progress bit check failed. This means a complete-split has been missed. There is the possibility of lost data. If PID Code is an IN, then the Queue head must be halted. If PID Code is an OUT, then the transfer state is not advanced and the state exited (for example, start-split is retried). This is a host-induced error and does not effect Cerr. In either case, set the Missed Micro-frame bit in the status field to a one.
A not(B) C	If PIDCode = IN Halt QHD If PIDCode = OUT Retry start-split	QH.FrameTag test failed. This means that exactly one or more H-Frames have been skipped. This means complete-splits and have missed. There is the possibility of lost data. If PID Code is an IN, then the Queue head must be halted. If PID Code is an OUT, then the transfer state is not advanced and the state exited (for example, start-split is retried). This is a host-induced error and does not effect Cerr. In either case, set the Missed Micro-frame bit in the status field to a one.
A B C not(D)	Execute complete-split	This is the non-error case where the host controller executes a complete-split transaction.
D	If PIDCode = IN Halt QHD If PIDCode = OUT Retry start-split	This is a degenerate case where the start-split was issued, but all of the complete-splits were skipped and all possible intervening opportunities to detect the missed data failed to fire. If PID Code is an IN, then the Queue head must be halted. If PID Code is an OUT, then the transfer state is not advanced and the state exited (for example, start-split is retried). This is a host-induced error and does not effect Cerr. In either case, set the Missed Micro-frame bit in the status field to a one. Note that when executing in the context of a Recovery Path mode, the host controller is allowed to process the queue head and take the actions indicated above, or it may wait until the queue head is visited in the normal processing mode. Regardless, the host controller must not execute a start-split in the context of a executing in a Recovery Path mode.

### 16.6.12.2.8 Managing the QH[FrameTag] Field

The QH[FrameTag] field in a queue head is completely managed by the host controller. The rules for setting QH[FrameTag] are simple:

- Rule 1: If transitioning from Do Start Split to Do Complete Split and the current value of FRINDEX[2–0] is 6, QH[FrameTag] is set to  $\text{FRINDEX}[7–3] + 1$ . This accommodates split transactions whose start-split and complete-splits are in different H-Frames (case 2a, see [Figure 16-53](#)).
- Rule 2: If the current value of FRINDEX[2–0] is 7, QH[FrameTag] is set to  $\text{FRINDEX}[7–3] + 1$ . This accommodates staying in Do Complete Split for cases 2a, 2b, and 2c in [Figure 16-53](#).
- Rule 3: If transitioning from Do\_Start Split to Do Complete Split and the current value of FRINDEX[2–0] is not 6, or currently in Do Complete Split and the current value of (FRINDEX[2–0]) is not 7, FrameTag is set to  $\text{FRINDEX}[7–3]$ . This accommodates all other cases in [Figure 16-53](#).

### 16.6.12.2.9 Rebalancing the Periodic Schedule

System software must occasionally adjust a periodic queue head's S-mask and C-mask fields during operation. This need occurs when adjustments to the periodic schedule create a new bandwidth budget and one or more queue head's are assigned new execution footprints (that is, new S-mask and C-mask values).

It is imperative that system software must not update these masks to new values in the midst of a split transaction. In order to avoid any race conditions with the update, the host controller provides a simple assist to system software. System software sets the Inactivate-on-next-Transaction (I) bit to signal the host controller that it intends to update the S-mask and C-mask on this queue head. System software then waits for the host controller to observe the I-bit is set and transitions the Active bit to a zero. The rules for how and when the host controller clears the Active bit are:

- If the Active bit is cleared, no action is taken. The host controller does not attempt to advance the queue when the I-bit is set.
- If the Active bit is set and the SplitXState is DoStart (regardless of the value of S-mask), the host controller simply clears the Active bit. The host controller is not required to write the transfer state back to the current qTD. Note that if the S-mask indicates that a start-split is scheduled for the current micro-frame, the host controller must not issue the start-split bus transaction; it must clear the Active bit.

System software must save transfer state before setting the I-bit. This is required so that it can correctly determine what transfer progress (if any) occurred after the I-bit was set and the host controller executed its final bus-transaction and cleared the Active bit.

After system software has updated the S-mask and C-mask, it must then reactivate the queue head. Since the Active bit and the I-bit cannot be updated with the same write, system software needs to use the following algorithm to coherently re-activate a queue head that has been stopped using the I-bit.

1. Set the Halted bit, then
2. Clear the I-bit, then
3. Set the Active bit and clear the Halted bit in the same write.

Setting the Halted bit inhibits the host controller from attempting to advance the queue between the time the I-bit is cleared and the Active bit is set.

### 16.6.12.3 Split Transaction Isochronous

Full-speed isochronous transfers are managed using the split-transaction protocol through a USB 2.0 transaction translator in a USB 2.0 hub. The host controller utilizes siTD data structure to support the special requirements of isochronous split-transactions. This data structure uses the scheduling model of isochronous TDs (see [Section 16.6.8, “Managing Isochronous Transfers Using iTDs,”](#) for the operational model of iTDs) with the contiguous data feature provided by queue heads. This simple arrangement allows a single isochronous scheduling model and adds the additional feature that all data received from the endpoint (per split transaction) must land into a contiguous buffer.



### 16.6.12.3.1 Split Transaction Scheduling Mechanisms for Isochronous

Full-speed isochronous transactions are managed through a transaction translator's periodic pipeline. As with full- and low-speed interrupt, system software manages each transaction translator's periodic pipeline by budgeting and scheduling exactly during which micro-frames the start-splits and complete-splits for each full-speed isochronous endpoint occur. The requirements described in Section Split Transaction Scheduling Mechanisms for Interrupt apply. Figure 16-57 illustrates the general scheduling boundary conditions that are supported by the EHCI periodic schedule. The  $S_n$  and  $C_n$  labels indicate micro-frames where software can schedule start- and complete-splits (respectively). The H-Frame boundaries are marked with a large, solid bold vertical line. The B-Frame boundaries are marked with a large, bold, dashed line. The bottom of Figure 16-57 illustrates the relationship of an siTD to the H-Frame.




Figure 16-57. Split Transaction, Isochronous Scheduling Boundary Conditions

When the endpoint is an isochronous OUT, there are only start-splits, and no complete-splits. When the endpoint is an isochronous IN, there is at most one start-split and one to N complete-splits. The scheduling boundary cases are:

- Case 1: The entire split transaction is completely bounded by an H-Frame. For example, the start-splits and complete-splits are all scheduled to occur in the same H-Frame.
- Case 2a: This boundary case is where one or more (at most two) complete-splits of a split transaction IN are scheduled across an H-Frame boundary. This can only occur when the split transaction has the possibility of moving data in B-Frame, micro-frames 6 or 7 (H-Frame micro-frame 7 or 0). When an H-Frame boundary wrap condition occurs, the scheduling of the split transaction spans more than one location in the periodic list.(for example, it takes two siTDs in adjacent periodic frame list locations to fully describe the scheduling for the split transaction). Although the scheduling of the split transaction may take two data structures, all of the complete-splits for each full-speed IN isochronous transaction must use only one data pointer. For this reason, siTDs contain a back pointer. Software must never schedule full-speed isochronous OUTs across an H-Frame boundary.
- Case 2b: This case can only occur for a very large isochronous IN. It is the only allowed scenario where a start-split and complete-split for the same endpoint can occur in the same micro-frame. Software must enforce this rule by scheduling the large transaction first. Large is defined to be anything larger than 579 byte maximum packet size.

A subset of the same mechanisms employed by full- and low-speed interrupt queue heads are employed in siTDs to schedule and track the portions of isochronous split transactions. The following fields are initialized by system software to instruct the host controller when to execute portions of the split transaction protocol:

- SplitXState. This is a single bit residing in the Status field of an siTD (see [Table 16-49](#)). This bit is used to track the current state of the split transaction. The rules for managing this bit are described in [Section 16.6.12.3.3, “Split Transaction Execution State Machine for Isochronous.”](#)
- Frame S-mask. This is a bit-field wherein system software sets a bit corresponding to the micro-frame (within an H-Frame) that the host controller should execute a start-split transaction. This is always qualified by the value of the SplitXState bit. For example, referring to the IN example in [Figure 16-57](#), case 1, the S-mask would have a value of 0b0000\_0001 indicating that if the siTD is traversed by the host controller, and the SplitXState indicates Do Start Split, and the current micro-frame as indicated by FRINDEX[2–0] is 0, then execute a start-split transaction.
- Frame C-mask. This is a bit-field where system software sets one or more bits corresponding to the micro-frames (within an H-Frame) that the host controller should execute complete-split transactions. The interpretation of this field is always qualified by the value of the SplitXState bit. For example, referring to the IN example in [Figure 16-57](#), case 1, the C-mask would have a value of 0b 0011\_1100 indicating that if the siTD is traversed by the host controller, and the SplitXState indicates Do Complete Split, and the current micro-frame as indicated by FRINDEX[2–0] is 2, 3, 4, or 5, then execute a complete-split transaction.
- Back Pointer. This field in a siTD is used to complete an IN split-transaction using the previous H-Frame's siTD. This is only used when the scheduling of the complete-splits span an H-Frame boundary.

There exists a one-to-one relationship between a high-speed isochronous split transaction (including all start- and complete-splits) and one full-speed isochronous transaction. An siTD contains (amongst other things) buffer state and split transaction scheduling information. An siTD's buffer state always maps to one full-speed isochronous data payload. This means that for any full-speed transaction payload, a single siTD's data buffer must be used. This rule applies to both IN and OUTs. An siTD's scheduling information usually also maps to one high-speed isochronous split transaction. The exception to this rule is the H-Frame boundary wrap cases mentioned above.

The siTD data structure describes at most, one frame's worth of high-speed transactions and that description is strictly bounded within a frame boundary. Figure 16-58 illustrates some examples. On the top are examples of the full-speed transaction footprints for the boundary scheduling cases described above. In the middle are time-frame references for both the B-Frames (HS/FS/LS Bus) and the H-Frames. On the bottom is illustrated the relationship between the scope of an siTD description and the time references. Each H-Frame corresponds to a single location in the periodic frame list. The implication is that each siTD is reachable from a single periodic frame list location at a time.




Figure 16-58. siTD Scheduling Boundary Examples

Each case is described below:

- Case 1: One siTD is sufficient to describe and complete the isochronous split transaction because the whole isochronous split transaction is tightly contained within a single H-Frame.
- Case 2a, 2b: Although both INs and OUTs can have these footprints, OUTs always take only one siTD to schedule. However, INs (for these boundary cases) require two siTDs to complete the scheduling of the isochronous split transaction. siTD<sub>X</sub> is used to always issue the start-split and the first N complete-splits. The full-speed transaction (for these cases) can deliver data on the full-speed bus segment during micro-frame 7 of H-Frame<sub>Y+1</sub>, or micro-frame 0 of H-Frame<sub>Y+2</sub>. The complete splits are scheduled using siTD<sub>X+2</sub> (not shown). The complete-splits to extract this data must use the buffer pointer from siTD<sub>X+1</sub>. The only way for the host controller to reach siTD<sub>X+1</sub> from H-Frame<sub>Y+2</sub> is to use siTD<sub>X+2</sub>'s back pointer.

Software must apply the following rules when calculating the schedule and linking the schedule data structures into the periodic schedule:

- Software must ensure that an isochronous split-transaction is started so that it will complete before the end of the B-Frame.
- Software must ensure that for a single full-speed isochronous endpoint, there is never a start-split and complete-split in H-Frame, micro-frame 1. This is mandated as a rule so that case 2a and case 2b can be discriminated. According to the core USB specification, the long isochronous transaction illustrated in Case 2b, could be scheduled so that the start-split was in micro-frame 1 of H-Frame N and the last complete-split would need to occur in micro-frame 1 of H-Frame N+1. However, it is impossible to discriminate between cases 2a and case 2b, which has significant impact on the complexity of the host controller.

### 16.6.12.3.2 Tracking Split Transaction Progress for Isochronous Transfers

Isochronous endpoints do not employ the concept of a halt on error, however the host controller does identify and report per-packet errors observed in the data stream. This includes schedule traversal problems (skipped micro-frames), timeouts and corrupted data received.

In similar kind to interrupt split-transactions, the portions of the split transaction protocol must execute in the micro-frames they are scheduled. The queue head data structure used to manage full- and low-speed interrupt has several mechanisms for tracking when portions of a transaction have occurred. Isochronous transfers use siTDs for their transfers and the data structures are only reachable using the schedule in the exact micro-frame in which they are required (so all the mechanism employed for tracking in queue heads is not required for siTDs). Software has the option of reusing siTD several times in the complete periodic schedule. However, it must ensure that the results of split transaction N are consumed and the siTD re-initialized (activated) before the host controller gets back to the siTD (in a future micro-frame).

Split-transaction isochronous OUTs utilize a low-level protocol to indicate which portions of the split transaction data have arrived. Control over the low-level protocol is exposed in an siTD using the fields Transaction Position (TP) and Transaction Count (T-count). If the entire data payload for the OUT split transaction is larger than 188 bytes, there will be more than one start-split transaction, each of which require proper annotation. If host hold-offs occur, then the sequence of annotations received from the host will not be complete, which is detected and handled by the transaction translator. See [Section 16.6.12.3.1, “Split Transaction Scheduling Mechanisms for Isochronous,”](#) for a description on how these fields are used during a sequence of start-split transactions.

The fields siTD[T-Count] and siTD[TP] are used by the host controller to drive and sequence the transaction position annotations. It is the responsibility of system software to properly initialize these fields in each siTD. Once the budget for a split-transaction isochronous endpoint is established, S-mask, T-Count, and TP initialization values for all the siTD associated with the endpoint are constant. They remain constant until the budget for the endpoint is recalculated by software and the periodic schedule adjusted.

For IN-endpoints, the transaction translator simply annotates the response data packets with enough information to allow the host controller to identify the last data. As with split transaction Interrupt, it is the host controller's responsibility to detect when it has missed an opportunity to execute a complete-split. The

following field in the siTD is used to track and detect errors in the execution of a split transaction for an IN isochronous endpoint.

- **C-prog-mask.** This is an eight-bit bit-vector where the host controller keeps track of which complete-splits have been executed. Due to the nature of the transaction translator periodic pipeline, the complete-splits need to be executed in-order. The host controller needs to detect when the complete-splits have not been executed in order. This can only occur due to system hold-offs where the host controller cannot get to the memory-based schedule. C-prog-mask is a simple bit-vector that the host controller sets a bit for each complete-split executed. The bit position is determined by the micro-frame (FRINDEX[2-0]) number in which the complete-split was executed. The host controller always checks C-prog-mask before executing a complete-split transaction. If the previous complete-splits have not been executed, then it means one (or more) have been skipped and data has potentially been lost. System software is required to initialize this field to zero before setting an siTD's Active bit to a one.

If a transaction translator returns with the final data before all of the complete-splits have been executed, the state of the transfer is advanced so that the remaining complete-splits are not executed. It is important to note that an IN siTD is retired based solely on the responses from the transaction translator to the complete-split transactions. This means, for example, that it is possible for a transaction translator to respond to a complete-split with an MDATA PID. The number of bytes in the MDATA's data payload could cause the siTD[Total Bytes to Transfer] field to decrement to zero. This response can occur, before all of the scheduled complete-splits have been executed. In other interface, data structures (for example, high-speed data streams through queue heads), the transition of Total Bytes to Transfer to zero signals the end of the transfer and results in clearing the Active bit. However, in this case, the result has not been delivered by the transaction translator and the host must continue with the next complete-split transaction to extract the residual transaction state. This scenario occurs because of the pipeline rules for a transaction translator. In summary, the periodic pipeline rules require that on a micro-frame boundary, the transaction translator holds the final two bytes received (if it has not seen an End Of Packet (EOP)) in the full-speed bus pipe stage and gives the remaining bytes to the high-speed pipeline stage. At the micro-frame boundary, the transaction translator could have received the entire packet (including both CRC bytes) but not received the packet EOP. In the next micro-frame, the transaction translator responds with an MDATA and sends all of the data bytes (with the two CRC bytes being held in the full-speed pipeline stage). This could cause the siTD to decrement its Total Bytes to Transfer field to zero, indicating it has received all expected data. The host must still execute one more (scheduled) complete-split transaction in order to extract the results of the full-speed transaction from the transaction translator (for example, the transaction translator may have detected a CRC failure, and this result must be forwarded to the host).

If the host experiences hold-offs that cause the host controller to skip one or more (but not all) scheduled split transactions for an isochronous OUT, then the protocol to the transaction translator is not consistent and the transaction translator detects and reacts to the problem. Likewise, for host hold-offs that cause the host controller to skip one or more (but not all) scheduled split transactions for an isochronous IN, the C-prog-mask is used by the host controller to detect errors. However, if the host experiences a hold-off that causes it to skip all of an siTD, or an siTD expires during a host hold off (for example, a hold-off occurs and the siTD is no longer reachable by the host controller in order for it to report the hold-off event), then system software must detect that the siTDs have not been processed by the host controller (for example, state not advanced) and report the appropriate error to the client driver.

### 16.6.12.3.3 Split Transaction Execution State Machine for Isochronous

In this section, all references to micro-frame are in the context of a micro-frame within an H-Frame.

If the Active bit in the Status byte is a zero, the host controller ignores the siTD and continues traversing the periodic schedule. Otherwise the host controller processes the siTD as specified below. A split transaction state machine is used to manage the split-transaction protocol sequence. The host controller uses the fields defined in Section 16.6.12.3.2, “Tracking Split Transaction Progress for Isochronous Transfers,” plus the variable cMicroFrameBit defined in Section 16.6.12.2.5, “Split Transaction Execution State Machine for Interrupt,” to track the progress of an isochronous split transaction. Figure 16-59 illustrates the state machine for managing an siTD through an isochronous split transaction. Bold, dotted circles denote the state of the Active bit in the Status field of a siTD. The Bold, dotted arcs denote the transitions between these states. Solid circles denote the states of the split transaction state machine and the solid arcs denote the transitions between these states. Dotted arcs and boxes reference actions that take place either as a result of a transition or from being in a state.




Figure 16-59. Split Transaction State Machine for Isochronous

### 16.6.12.3.4 Periodic Isochronous—Do-Start-Split

Isochronous split transaction OUTs use only this state. An siTD for a split-transaction isochronous IN is either initialized to this state, or the siTD transitions to this state from Do Complete Split when a case 2a (IN) or 2b scheduling boundary isochronous split-transaction completes.

Each time the host controller reaches an active siTD in this state, it checks the siTD[S-mask] against cMicroFrameBit. If there is a one in the appropriate position, the siTD executes a start-split transaction.

By definition, the host controller cannot reach an siTD at the wrong time. If the I/O field indicates an IN, then the start-split transaction includes only the extended token plus the full-speed token. Software must initialize the siTD[Total Bytes To Transfer] field to the number of bytes expected. This is usually the maximum packet size for the full-speed endpoint. The host controller exits this state when the start-split transaction is complete.

The remainder of this section is specific to an isochronous OUT endpoint (that is, the I/O field indicates an OUT). When the host controller executes a start-split transaction for an isochronous OUT it includes a data payload in the start-split transaction. The memory buffer address for the data payload is constructed by concatenating siTD[Current Offset] with the page pointer indicated by the page select field (siTD[P]). A zero in this field selects Page 0 and a 1 selects Page 1. During the start-split for an OUT, if the data transfer crosses a page boundary during the transaction, the host controller must detect the page cross, update the siTD[P] bit from a zero to a one, and begin using the siTD Page 1 with siTD[Current Offset] as the memory address pointer. The field siTD[TP] is used to annotate each start-split transaction with the indication of which part of the split-transaction data the current payload represents (ALL, BEGIN, MID, END). In all cases, the host controller simply uses the value in siTD[TP] to mark the start-split with the correct transaction position code.

T-count is always initialized to the number of start-splits for the current frame. TP is always initialized to the first required transaction position identifier. The scheduling boundary case (see [Figure 16-58](#)) is used to determine the initial value of TP. The initial cases are summarized in [Table 16-69](#).

**Table 16-69. Initial Conditions for OUT siTD TP and T-Count Fields**

Case	T-Count	TP	Description
1, 2a	=1	ALL	When the OUT data payload is less than (or equal to) 188 bytes, only one start-split is required to move the data. The one start-split must be marked with an ALL.
1, 2a	!=1	BEGIN	When the OUT data payload is greater than 188 bytes more than one start-split must be used to move the data. The initial start-split must be marked with a BEGIN.

After each start-split transaction is complete, the host controller updates T-count and TP appropriately so that the next start-split is correctly annotated. [Table 16-70](#) illustrates all of the TP and T-count transitions, which must be accomplished by the host controller.

**Table 16-70. Transaction Position (TP)/Transaction Count (T-Count) Transition Table**

TP	T-Count Next	TP Next	Description
ALL	0	N/A	Transition from ALL, to done.
BEGIN	1	END	Transition from BEGIN to END. Occurs when T-count starts at 2.
BEGIN	!=1	MID	Transition from BEGIN to MID. Occurs when T-count starts at greater than 2.
MID	!=1	MID	TP stays at MID while T-count is not equal to 1 (for example, greater than 1). This case can occur for any of the scheduling boundary cases where the T-count starts greater than 3.
MID	1	END	Transition from MID to END. This case can occur for any of the scheduling boundary cases where the T-count starts greater than 2.

The start-split transactions do not receive a handshake from the transaction translator, so the host controller always advances the transfer state in the siTD after the bus transaction is complete. To advance the transfer state the following operations take place:

- The siTD[Total Bytes To Transfer] and the siTD[Current Offset] fields are adjusted to reflect the number of bytes transferred.
- The siTD[P] (page select) bit is updated appropriately.
- The siTD[TP] and siTD[T-count] fields are updated appropriately as defined in [Table 16-70](#).

These fields are then written back to the memory based siTD. The S-mask is fixed for the life of the current budget. As mentioned above, TP and T-count are set specifically in each siTD to reflect the data to be sent from this siTD. Therefore, regardless of the value of S-mask, the actual number of start-split transactions depends on T-count (or equivalently, Total Bytes to Transfer). The host controller must clear the Active bit when it detects that all of the schedule data has been sent to the bus. The preferred method is to detect when T-Count decrements to zero as a result of a start-split bus transaction. Equivalently, the host controller can detect when Total Bytes to Transfer decrements to zero. Either implementation must ensure that if the initial condition is Total Bytes to Transfer is equal to zero and T-count is equal to a one, then the host controller will issue a single start-split, with a zero-length data payload. Software must ensure that TP, T-count and Total Bytes to Transfer are set to deliver the appropriate number of bus transactions from each siTD. An inconsistent combination will yield undefined behavior.

If the host experiences hold-offs that cause the host controller to skip start-split transactions for an OUT transfer, the state of the transfer will not progress appropriately. The transaction translator observes protocol violations in the arrival of the start-splits for the OUT endpoint (that is, the transaction position annotation is incorrect as received by the transaction translator).

Example scenarios are described in [Section 16.6.12.3.7, “Split Transaction for Isochronous—Processing Example.”](#)

The host controller can optionally track the progress of an OUT split transaction by setting appropriate bits in the siTD[C-prog-mask] as it executes each scheduled start-split. The checkPreviousBit() algorithm defined in [Section 16.6.12.3.5, “Periodic Isochronous—Do Complete Split,”](#) can be used prior to executing each start-split to determine whether start-splits were skipped. The host controller can use this mechanism to detect missed micro-frames. It can then clear the siTD's Active bit and stop execution of this siTD. This saves on both memory and high-speed bus bandwidth.

### 16.6.12.3.5 Periodic Isochronous—Do Complete Split

This state is only used by a split-transaction isochronous IN endpoint. This state is entered unconditionally from the Do Start State after a start-split transaction is executed for an IN endpoint. Each time the host controller visits an siTD in this state, it conducts a number of tests to determine whether it should execute a complete-split transaction. The individual tests are listed below. The sequence they are applied depends on which micro-frame the host controller is currently executing which means that the tests might not be applied until after the siTD referenced from the back pointer has been fetched.

- Test A. cMicroFrameBit is bit-wise ANDed with the siTD[C-mask] field. A non-zero result indicates that software scheduled a complete-split for this endpoint, during this micro-frame. This test is always applied to a newly fetched siTD that is in this state.



- Test B. The siTD[C-prog-mask] bit vector is checked to determine whether the previous complete splits have been executed. An example algorithm is given below (this is slightly different than the algorithm used in [Section 16.6.12.2.7, “Periodic Interrupt—Do-Complete-Split”](#)). The sequence in which this test is applied depends on the current value of FRINDEX[2–0]. If FRINDEX[2–0] is 0 or 1, it is not applied until the back pointer has been used. Otherwise it is applied immediately.

```

Algorithm Boolean CheckPreviousBit(siTD.C-prog-mask, siTD.C-mask, cMicroFrameBit)
Begin
    Boolean rvalue = TRUE;
    previousBit = cMicroFrameBit rotate-right(1)
    -- Bit-wise anding previousBit with C-mask indicates whether there
    -- was an intent to send a complete split in the previous micro-
    -- frame. So, if the 'previous bit' is set in C-mask, check
    -- C-prog-mask to make sure it happened.
    if previousBit bitAND siTD.C-mask then
        if not (previousBit bitAND siTD.C-prog-mask) then
            rvalue = FALSE
        End if
    End if
    Return rvalue
End Algorithm

```

If Test A is true and FRINDEX[2–0] is zero or one, then this is a case 2a or 2b scheduling boundary (see [Figure 16-57](#)). See [Section 16.6.12.3.6, “Complete-Split for Scheduling Boundary Cases 2a, 2b,”](#) for details in handling this condition.

If Test A and Test B evaluate to true, then the host controller executes a complete-split transaction using the transfer state of the current siTD. When the host controller commits to executing the complete-split transaction, it updates QH[C-prog-mask] by bit-ORing with cMicroFrameBit. The transfer state is advanced based on the completion status of the complete-split transaction. To advance the transfer state of an IN siTD, the host controller must:

- Decrement the number of bytes received from siTD[Total Bytes To Transfer]
- Adjust siTD[Current Offset] by the number of bytes received
- Adjust the siTD[P] (page select) field if the transfer caused the host controller to use the next page pointer
- Set any appropriate bits in the siTD[Status] field, depending on the results of the transaction.

Note that if the host controller encounters a condition where siTD[Total Bytes To Transfer] is zero, and it receives more data, the host controller must not write the additional data to memory. The siTD[Status-Active] bit must be cleared and the siTD[Status-Babble Detected] bit must be set. The fields siTD[Total Bytes To Transfer], siTD[Current Offset], and siTD[P] are not required to be updated as a result of this transaction attempt.

The host controller accepts (assuming good data packet CRC and sufficient room in the buffer as indicated by the value of siTD[Total Bytes To Transfer]) MDATA and DATA0/1 data payloads up to and including 192 bytes. The host controller may optionally clear siTD[Status-Active] and set siTD[Status-Babble

Detected] when it receives MDATA or DATA0/1 with a data payload of more than 192 bytes. The following responses have the noted effects:

- **ERR.** The full-speed transaction completed with a time-out or bad CRC and this is a reflection of that error to the host. The host controller sets the ERR bit in the siTD[Status] field and clears the Active bit.
- **Transaction Error (XactErr).** The complete-split transaction encounters a Timeout, CRC16 failure, etc. The siTD[Status] field XactErr field is set and the complete-split transaction must be retried immediately. The host controller must use an internal error counter to count the number of retries as a counter field is not provided in the siTD data structure. The host controller will not retry more than two times. If the host controller exhausts the retries or the end of the micro-frame occurs, the Active bit is cleared.
- **DATAx (0 or 1).** This response signals that the final data for the split transaction has arrived. The transfer state of the siTD is advanced and the Active bit is cleared. If the Bytes To Transfer field has not decremented to zero (including the reception of the data payload in the DATAx response), then less data than was expected, or allowed for was actually received. This short packet event does not set the USB interrupt status bit (USBSTS[UI]) to a one. The host controller will not detect this condition.
- **NYET (and Last).** On each NYET response, the host controller also checks to determine whether this is the last complete-split for this split transaction. Last was defined in Section Periodic Interrupt - Do Complete Split. If it is the last complete-split (with a NYET response), then the transfer state of the siTD is not advanced (never received any data) and the Active bit is cleared. No bits are set in the Status field because this is essentially a skipped transaction. The transaction translator must have responded to all the scheduled complete-splits with NYETs, meaning that the start-split issued by the host controller was not received. This result should be interpreted by system software as if the transaction was completely skipped. The test for whether this is the last complete split can be performed by XORing C-mask with C-prog-mask. A zero result indicates that all complete-splits have been executed.
- **MDATA (and Last).** See above description for testing for Last. This can only occur when there is an error condition. Either there has been a babble condition on the full-speed link, which delayed the completion of the full-speed transaction, or software set up the S-mask and/or C-masks incorrectly. The host controller must set the XactErr bit and clear the Active bit.
- **NYET (and not Last).** See above description for testing for Last. The complete-split transaction received a NYET response from the transaction translator. Do not update any transfer state (except for C-prog-mask) and stay in this state.
- **MDATA (and not Last).** The transaction translator responds with an MDATA when it has partial data for the split transaction. For example, the full-speed transaction data payload spans from micro-frame X to X+1 and during micro-frame X, the transaction translator responds with an MDATA and the data accumulated up to the end of micro-frame X. The host controller advances the transfer state to reflect the number of bytes received.

If Test A succeeds, but Test B fails, it means that one or more of the complete-splits have been skipped. The host controller sets the Missed Micro-Frame status bit and clears the Active bit.

### 16.6.12.3.6 Complete-Split for Scheduling Boundary Cases 2a, 2b

Boundary cases 2a and 2b (INs only) (see [Figure 16-57](#)) require that the host controller use the transaction state context of the previous siTD to finish the split transaction. [Table 16-71](#) enumerates the transaction state fields.

**Table 16-71. Summary siTD Split Transaction State**

Buffer State	Status	Execution Progress
Total Bytes To Transfer P (page select) Current Offset TP (transaction position) T-count (transaction count)	All bits in the status field	C-prog-mask

#### NOTE

TP and T-count are used only for Host to Device (OUT) endpoints.

If software has budgeted the schedule of this data stream with a frame wrap case, then it must initialize the siTD[Back Pointer] field to reference a valid siTD and have the T bit in the siTD[Back Pointer] field cleared. Otherwise, software must set the T bit in siTD[Back Pointer]. The host controller's rules for interpreting when to use the siTD[Back Pointer] field are listed below. These rules apply only when the siTD's Active bit is a one and the SplitXState is Do Complete Split.

- When cMicroFrameBit is a 0x1 and the siTD<sub>X</sub>[Back Pointer] T-bit is zero, or
- If cMicroFrameBit is a 0x2 and siTD<sub>X</sub>[S-mask[0]] is zero

When either of these conditions apply, then the host controller must use the transaction state from siTD<sub>X-1</sub>.

In order to access siTD<sub>X-1</sub>, the host controller reads on-chip the siTD referenced from siTD<sub>X</sub>[Back Pointer].

The host controller must save the entire state from siTD<sub>X</sub> while processing siTD<sub>X-1</sub>. This is to accommodate for case 2b processing. The host controller must not recursively walk the list of siTD[Back Pointers].

If siTD<sub>X-1</sub> is active (Active bit is set and SplitXStat is Do Complete Split), then both Test A and Test B are applied as described above. If these criteria to execute a complete-split are met, the host controller executes the complete split and evaluates the results as described above. The transaction state (see [Table 16-71](#)) of siTD<sub>X-1</sub> is appropriately advanced based on the results and written back to memory. If the resultant state of siTD<sub>X-1</sub>'s Active bit is a one, then the host controller returns to the context of siTD<sub>X</sub>, and follows its next pointer to the next schedule item. No updates to siTD<sub>X</sub> are necessary.

If siTD<sub>X-1</sub> is active (Active bit is set and SplitXStat is Do Start Split), then the host controller must clear the Active bit and set the Missed Micro-Frame status bit and the resultant status is written back to memory.

If siTD<sub>X-1</sub>'s Active bit is cleared, (because it was cleared when the host controller first visited siTD<sub>X-1</sub> via siTD<sub>X</sub>'s back pointer, it transitioned to zero as a result of a detected error, or the results of siTD<sub>X-1</sub>'s complete-split transaction cleared it), then the host controller returns to the context of siTD<sub>X</sub> and transitions its SplitXState to Do Start Split. The host controller then determines whether the case 2b start split boundary condition exists (that is, if cMicroframeBit is 1 and siTD<sub>X</sub>[S-mask[0]] is 1). If this criterion

is met the host controller immediately executes a start-split transaction and appropriately advances the transaction state of siTD<sub>X</sub>, then follows siTD<sub>X</sub>[Next Pointer] to the next schedule item. If the criterion is not met, the host controller simply follows siTD<sub>X</sub>[Next Pointer] to the next schedule item. Note that in the case of a 2b boundary case, the split-transaction of siTD<sub>X-1</sub> will have its Active bit cleared when the host controller returns to the context of siTD<sub>X</sub>. Also, note that software should not initialize an siTD with C-mask bits 0 and 1 set and an S-mask with bit 0 set. This scheduling combination is not supported and the behavior of the host controller is undefined.

### 16.6.12.3.7 Split Transaction for Isochronous—Processing Example

There is an important difference between how the hardware/software manages the isochronous split transaction state machine and how it manages the asynchronous and interrupt split transaction state machines. The asynchronous and interrupt split transaction state machines are encapsulated within a single queue head. The progress of the data stream depends on the progress of each split transaction. In some respects, the split-transaction state machine is sequenced using the Execute Transaction queue head traversal state machine.

Isochronous is a pure time-oriented transaction/data stream. The interface data structures are optimized to efficiently describe transactions that need to occur at specific times. The isochronous split-transaction state machine must be managed across these time-oriented data structures. This means that system software must correctly describe the scheduling of split-transactions across more than one data structure.

Then the host controller must make the appropriate state transitions at the appropriate times, in the correct data structures.

For example, [Table 16-72](#) illustrates a few frames worth of scheduling required to schedule a case 2a full-speed isochronous data stream.

**Table 16-72. Example Case 2a—Software Scheduling siTDs for an IN Endpoint**

siTDX		Micro-Frames								InitialSplitXState
#	Masks	0	1	2	3	4	5	6	7	
X	S-Mask					1				Do Start Split
	C-Mask	1	1					1	1	
X+1	S-Mask					1				Do Complete Split
	C-Mask	1	1					1	1	
X+2	S-Mask					1				Do Complete Split
	C-Mask	1	1					1	1	
X+3	S-Mask	Repeats previous pattern								Do Complete Split
	C-Mask									

This example shows the first three siTDs for the transaction stream. Since this is the case-2a frame-wrap case, S-masks of all siTDs for this endpoint have a value of 0x10 (a one bit in micro-frame 4) and C-mask value of 0xC3 (one-bits in micro-frames 0,1, 6 and 7). Additionally, software ensures that the Back Pointer field of each siTD references the appropriate siTD data structure (and the Back Pointer T-bits are cleared).

The initial SplitXState of the first siTD is Do Start Split. The host controller will visit the first siTD eight times during frame X. The C-mask bits in micro-frames 0 and 1 are ignored because the state is Do Start Split. During micro-frame 4, the host controller determines that it can run a start-split (and does) and changes SplitXState to Do Complete Split. During micro-frames 6 and 7, the host controller executes complete-splits. Notice the siTD for frame X+1 has its SplitXState initialized to Do Complete Split. As the host controller continues to traverse the schedule during H-Frame X+1, it will visit the second siTD eight times. During micro-frames 0 and 1 it will detect that it must execute complete-splits.

During H-Frame X+1, micro-frame 0, the host controller detects that siTD<sub>X+1</sub>'s Back Pointer[T] bit is a zero, saves the state of siTD<sub>X+1</sub> and fetches siTD<sub>X</sub>. It executes the complete split transaction using the transaction state of siTD<sub>X</sub>. If the siTD<sub>X</sub> split transaction is complete, siTD's Active bit is cleared and results written back to siTD<sub>X</sub>. The host controller retains the fact that siTD<sub>X</sub> is retired and transitions the SplitXState in siTD<sub>X+1</sub> to Do Start Split. At this point, the host controller is prepared to execute the start-split for siTD<sub>X+1</sub> when it reaches micro-frame 4. If the split-transaction completes early (transaction-complete is defined in Section 16.6.12.3.5, “Periodic Isochronous—Do Complete Split”), that is, before all the scheduled complete-splits have been executed, the host controller changes siTD<sub>X</sub>[SplitXState] to Do Start Split early and naturally skips the remaining scheduled complete-split transactions. For this example, siTD<sub>X+1</sub> does not receive a DATA0 response until H-Frame X+2, micro-frame 1.

During H-Frame X+2, micro-frame 0, the host controller detects that siTD<sub>X+2</sub>'s Back Pointer[T] bit is zero, saves the state of siTD<sub>X+2</sub> and fetches siTD<sub>X+1</sub>. As described above, it executes another split transaction, receives an MDATA response, updates the transfer state, but does not modify the Active bit. The host controller returns to the context of siTD<sub>X+2</sub>, and traverses its next pointer without any state change updates to siTD<sub>X+2</sub>.

During H-Frame X+2, micro-frame 1, the host controller detects siTD<sub>X+2</sub>'s S-mask[0] bit is zero, saves the state of siTD<sub>X+2</sub> and fetches siTD<sub>X+1</sub>. It executes another complete-split transaction, receives a DATA0 response, updates the transfer state and clears the Active bit. It returns to the state of siTD<sub>X+2</sub> and changes its SplitXState to Do Start Split. At this point, the host controller is prepared to execute start-splits for siTD<sub>X+2</sub> when it reaches micro-frame 4.

### 16.6.13 Port Test Modes

EHCI host controllers implement the port test modes Test J\_State, Test K\_State, Test\_Packet, Test Force\_Enable, and Test SEO\_NAK as described in the *USB Specification Revision 2.0*. The required, port test sequence is (assuming the CF-bit in the CONFIGFLAG register is set):

- Disable the periodic and asynchronous schedules by clearing the USBCMD[ASE] and USBCMD[PSE].
- Place all enabled root ports into the suspended state by setting the Suspend bit in the PORTSC register (PORTSC[SUSP]).
- Clear USBCMD[RS] (run/stop) and wait for USBSTS[HCH] to transition to a one. Note that an EHCI host controller implementation may optionally allow port testing with RS set. However, all host controllers must support port testing with RS cleared and HCH set.

- Set the Port Test Control field in the port under test PORTSC register to the value corresponding to the desired test mode. If the selected test is Test\_Force\_Enable, then USBCMD[RS] must then be transitioned back to one, in order to enable transmission of SOFs out of the port under test.
- When the test is complete, system software must ensure the host controller is halted (HCH bit is a one) then it terminates and exits test mode by setting USBCMD[RST].

### 16.6.14 Interrupts

The EHCI host controller hardware provides interrupt capability based on a number of sources. There are several general groups of interrupt sources:

- Interrupts as a result of executing transactions from the schedule (success and error conditions),
- Host controller events (Port change events, etc.), and
- Host controller error events

All transaction-based sources are maskable through the host controller's Interrupt Enable register (USBINTR). Additionally, individual transfer descriptors can be marked to generate an interrupt on completion. This section describes each interrupt source and the processing that occurs in response to the interrupt.

During normal operation, interrupts may be immediate or deferred until the next interrupt threshold occurs. The interrupt threshold is a tunable parameter via the Interrupt Threshold Control field in the USBCMD register. The value of this register controls when the host controller generates an interrupt on behalf of normal transaction execution. When a transaction completes during an interrupt interval period, the interrupt signaling the completion of the transfer will not occur until the interrupt threshold occurs. For example, the default value is eight micro-frames. This means that the host controller will not generate interrupts any more frequently than once every eight micro-frames.

[Section 16.6.14.2.4, “Host System Error”](#) details effects of a host system error.

If an interrupt has been scheduled to be generated for the current interrupt threshold interval, the interrupt is not signaled until after the status for the last complete transaction in the interval has been written back to system memory. This may sometimes result in the interrupt not being signaled until the next interrupt threshold.

Initial interrupt processing is the same, regardless of the reason for the interrupt. When an interrupt is signaled by the hardware, CPU control is transferred to host controller's USB interrupt handler. The precise mechanism to accomplish the transfer is OS specific. For this discussion it is just assumed that control is received. When the interrupt handler receives control, its first action is to read the USBSTS. It then acknowledges the interrupt by clearing all of the interrupt status bits by writing ones to these bit positions. The handler then determines whether the interrupt is due to schedule processing or some other event. After acknowledging the interrupt, the handler (via an OS-specific mechanism), schedules a deferred procedure call (DPC) which will execute later. The DPC routine processes the results of the schedule execution. The precise mechanisms used are beyond the scope of this document.

**NOTE**

The only method software should use for acknowledging an interrupt is by transitioning the appropriate status bits in the USBSTS register from a one to a zero.

**16.6.14.1 Transfer/Transaction Based Interrupts**

These interrupt sources are associated with transfer and transaction progress. They are all dependent on the next interrupt threshold.

**16.6.14.1.1 Transaction Error**

A transaction error is any error that caused the host controller to think that the transfer did not complete successfully. [Table 16-73](#) lists the events/responses that the host can observe as a result of a transaction. The effects of the error counter and interrupt status are summarized in the following paragraphs. Most of these errors set the XactErr status bit in the appropriate interface data structure.

**Table 16-73. Summary of Transaction Errors**

Event/ Result	Queue Head/qTD/iTD/siTD Side Effects		USBSTS[USBERRINT]
	Cerr	Status Field	
CRC	-1	XactErr set	1 <sup>1</sup>
Timeout	-1	XactErr set	1 <sup>1</sup>
Bad PID <sup>2</sup>	-1	XactErr set	1 <sup>1</sup>
Babble	N/A	See <a href="#">Section 16.6.14.1.2, “Serial Bus Babble”</a>	1
Buffer Error	N/A	See <a href="#">Section 16.6.14.1.3, “Data Buffer Error”</a>	

<sup>1</sup> If occurs in a queue head, then USBERRINT is asserted only when Cerr counts down from a one to a zero. In addition the queue is halted.

<sup>2</sup> The host controller received a response from the device, but it could not recognize the PID as a valid PID.

There is a small set of protocol errors that relate only when executing a queue head and fit under the umbrella of a WRONG PID error that are significant to explicitly identify. When these errors occur, the XactErr status bit in the queue head is set and the Cerr field is decremented. When the PID Code indicates a SETUP, the following responses are protocol errors and result in XactErr bit being set and the Cerr field being decremented.

- EPS field indicates a high-speed device and it returns a Nak handshake to a SETUP.
- EPS field indicates a high-speed device and it returns a Nyet handshake to a SETUP.
- EPS field indicates a low- or full-speed device and the complete-split receives a Nak handshake.

**16.6.14.1.2 Serial Bus Babble**

When a device transmits more data on the USB than the host controller is expecting for this transaction, it is defined to be babbling. In general, this is called a Packet Babble. When a device sends more data than the Maximum Length number of bytes, the host controller sets the Babble Detected bit to a one and halts

the endpoint if it is using a queue head. Maximum Length is defined as the minimum of Total Bytes to Transfer and Maximum Packet Size. The Cerr field is not decremented for a packet babble condition (only applies to queue heads). A babble condition also exists if IN transaction is in progress at High-speed EOF2 point. This is called a frame babble. A frame babble condition is recorded into the appropriate schedule data structure. In addition, the host controller must disable the port to which the frame babble is detected.

USBSTS[UEI] (USB error interrupt) is set and if the USBINTR[UEE] (USB error interrupt enable) is set, then a hardware interrupt is signaled to the system at the next interrupt threshold. The host controller must never start an OUT transaction that babbles across a micro-frame EOF.

### NOTE

When a host controller detects a data PID mismatch, it must either: disable the packet babble checking for the duration of the bus transaction or do packet babble checking based solely on Maximum Packet Size. The USB core specification defines the requirements on a data receiver when it receives a data PID mismatch (for example, expects a DATA0 and gets a DATA1 or visa-versa). In summary, it must ignore the received data and respond with an ACK handshake, in order to advance the transmitter's data sequence. The EHCI interface allows system software to provide buffers for a Control, Bulk or Interrupt IN endpoint that are not an even multiple of the maximum packet size specified by the device. Whenever a device misses an ACK for an IN endpoint, the host and device are out of synchronization with respect to the progress of the data transfer. The host controller may have advanced the transfer to a buffer that is less than maximum packet size. The device re-sends its maximum packet size data packet, with the original data PID, in response to the next IN token. In order to properly manage the bus protocol, the host controller must disable the packet babble check when it observes the data PID mismatch.

#### 16.6.14.1.3 Data Buffer Error

This event indicates that an overrun of incoming data or a underrun of outgoing data has occurred for this transaction. This would generally be caused by the host controller not being able to access required data buffers in memory within necessary latency requirements. These conditions are not considered transaction errors, and do not effect the error count in the queue head. When these errors do occur, the host controller records the fact the error occurred by setting the Data Buffer Error bit in the queue head, iTD or siTD.

If the data buffer error occurs on a non-isochronous IN, the host controller will not issue a handshake to the endpoint. This forces the endpoint to resend the same data (and data toggle) in response to the next IN to the endpoint.

If the data buffer error occurs on an OUT, the host controller must corrupt the end of the packet so that it cannot be interpreted by the device as a good data packet. Simply truncating the packet is not considered acceptable. An acceptable implementation option is to 1's complement the CRC bytes and send them. There are other options suggested in the transaction translator section of the *USB Specification Revision 2.0*.



#### 16.6.14.1.4 USB Interrupt (Interrupt on Completion (IOC))

Transfer Descriptors (iTDs, siTDs, and queue heads (qTDs)) contain a bit that can be set to cause an interrupt on their completion. The completion of the transfer associated with that schedule item causes USBSTS[UI] (USB interrupt) to be set. In addition, if a short packet is encountered on an IN transaction associated with a queue head, then this event also causes USBINT to be set. If USBINTR[UE] (USB interrupt enable) is set, a hardware interrupt is signaled to the system at the next interrupt threshold. If the completion is because of errors, USBSTS[UEI] (USB error interrupt) is also set.

#### 16.6.14.1.5 Short Packet

Reception of a data packet that is less than the endpoint's Max Packet size during Control, Bulk or Interrupt transfers signals the completion of the transfer. Whenever a short packet completion occurs during a queue head execution, USBSTS[UI] (USB interrupt bit) is set. If the USB interrupt enable bit is set (USBINTR[UE]), a hardware interrupt is signaled to the system at the next interrupt threshold.

### 16.6.14.2 Host Controller Event Interrupts

These interrupt sources are independent of the interrupt threshold (with the one exception being the Interrupt on Async Advance).

#### 16.6.14.2.1 Port Change Events

Port registers contain status and status change bits. When the status change bits are set, the host controller sets the USBSTS[PCI]. If the port change interrupt enable bit (PCE) in the USBINTR register is set, the host controller issues a hardware interrupt. The port status change bits in PORTSC include:

- Connect change status (CSC)
- Port enable/disable change (PEC)
- Over-current change (OCC)
- Force port resume (FPR)

#### 16.6.14.2.2 Frame List Rollover

This event indicates that the host controller has wrapped the frame list. The current programmed size of the frame list effects how often this interrupt occurs. If the frame list size is 1024, then the interrupt occurs every 1024 milliseconds, if it is 512, then it occurs every 512 milliseconds, etc. When a frame list rollover is detected, the host controller sets the frame list rollover bit, USBSTS[FRI]. If USBINTR[FRE] is set (frame list rollover enable), the host controller issues a hardware interrupt. This interrupt is not delayed to the next interrupt threshold.

#### 16.6.14.2.3 Interrupt on Async Advance

This event is used for deterministic removal of queue heads from the asynchronous schedule. Whenever the host controller advances the on-chip context of the asynchronous schedule, it evaluates the value of USBCMD[IAA]. If it is set, it sets USBSTS[AAI]. If USBINTR[AAE] is set, the host controller issues a hardware interrupt at the next interrupt threshold. A detailed explanation of this feature is described in [Section 16.6.9.2, “Removing Queue Heads from Asynchronous Schedule.”](#)

### 16.6.14.2.4 Host System Error

The host controller is a bus master and any interaction between the host controller and the system may experience errors. The type of host error may be catastrophic to the host controller making it impossible for the host controller to continue in a coherent fashion. Behavior for these types of errors is to halt the host controller. Host-based error must result in the following actions:

- USBCMD[RS] is cleared.
- USBSTS[SEI] and USBSTS[HCH] register are set
- If the host system error enable bit, USBINTR[SEE] is set, the host controller issues a hardware interrupt. This interrupt is not delayed to the next interrupt threshold.

Table 16-74 summarizes the required actions taken on the various host errors.

**Table 16-74. Summary Behavior on Host System Errors**

Cycle Type	Master Abort	Target Abort	Data Phase Parity
Frame list pointer fetch (read)	Fatal	Fatal	Fatal
siTD fetch (read)	Fatal	Fatal	Fatal
siTD status write-back (write)	Fatal	Fatal	Fatal
iTD fetch (read)	Fatal	Fatal	Fatal
iTD status write-back (write)	Fatal	Fatal	Fatal
qTD fetch (read)	Fatal	Fatal	Fatal
qHD status write-back (write)	Fatal	Fatal	Fatal
Data write	Fatal	Fatal	Fatal
Data read	Fatal	Fatal	Fatal

**NOTE**

After a host system error, software must reset the host controller using USBCMD[RST] before re-initializing and restarting the host controller.

## 16.7 Device Data Structures

This section defines the interface data structures used to communicate control, status, and data between device controller driver (DCD) software and the device controller. The data structure definitions in this chapter support a 32-bit memory buffer address space. The interface consists of device queue heads and transfer descriptors.

**NOTE**

Software must ensure that no interface data structure reachable by the device controller spans a 4K-page boundary.

The data structures defined in the section are (from the device controller's perspective) a mix of read-only and read/ writable fields. The device controller must preserve the read-only fields on all data structure writes.

The USB DR module includes DCD software called the USB 2.0 Device API. The device API provides an easy to use Application Program Interface for developing device (peripheral) applications. The device API incorporates and abstracts for the application developer all of the elements of the program interface.

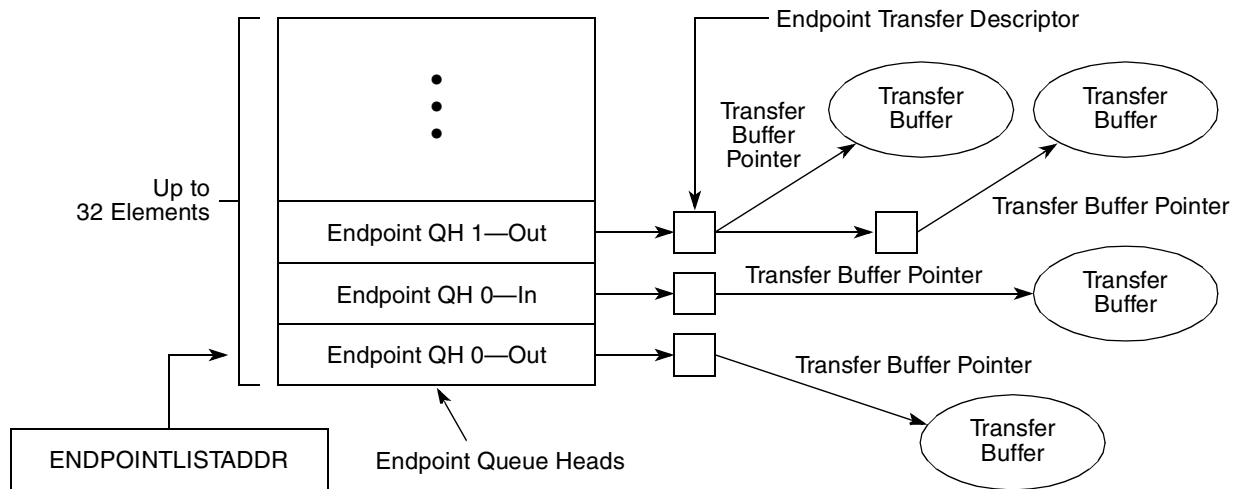


Figure 16-60. Endpoint Queue Head Organization

### 16.7.1 Endpoint Queue Head

The device Endpoint Queue Head (dQH) is where all transfers are managed. The dQH is a 48-byte data structure, but must be aligned on 64-byte boundaries. During priming of an endpoint, the dTD (device transfer descriptor) is copied into the overlay area of the dQH, which starts at the nextTD pointer DWord and continues through the end of the buffer pointers DWords. After a transfer is complete, the dTD status DWord is updated in the dTD pointed to by the currentTD pointer. While a packet is in progress, the overlay area of the dQH is used as a staging area for the dTD so that the Device Controller can access needed information with little minimal latency.

Figure 16-61 shows the Endpoint Queue Head structure.

31	30	29	28	27	26	25	24	23	22	21	20	19	18	17	16	15	14	13	12	11	10	9	8	7	6	5	4	3	2	1	0	offset
Mult		zlt		00		Maximum Packet Length						ios		000_0000_0000_0000														0x00				
Current dTD Pointer <sup>1</sup>																0_0000		0x04														
Next dTD Pointer <sup>1</sup>																0000		T <sup>1</sup>	0x08 <sup>2</sup>													
00		Total Bytes <sup>1</sup>						ioc <sup>1</sup>		000		MultO <sup>1</sup>		00		Status <sup>1</sup>				0x0C <sup>2</sup>												
Buffer Pointer (Page 0) <sup>1</sup>										Current Offset <sup>1</sup>								0x10 <sup>2</sup>														
Buffer Pointer (Page 1) <sup>1</sup>										Reserved								0x14 <sup>2</sup>														
Buffer Pointer (Page 2) <sup>1</sup>										Reserved								0x18 <sup>2</sup>														
Buffer Pointer (Page 3) <sup>1</sup>										Reserved								0x1C <sup>2</sup>														
Buffer Pointer (Page 4) <sup>1</sup>										Reserved								0x20 <sup>2</sup>														
Reserved																0x24																
Setup Buffer Bytes 3–0 <sup>1</sup>																0x28																
Setup Buffer Bytes 7–4 <sup>1</sup>																0x2C																

Figure 16-61. Endpoint Queue Head Layout

- <sup>1</sup> Device controller read/write; all others read-only.
- <sup>2</sup> Offsets 0x08 through 0x20 contain the transfer overlay.

### 16.7.1.1 Endpoint Capabilities/Characteristics

This DWord specifies static information about the endpoint, in other words, this information does not change over the lifetime of the endpoint. Device Controller software should not attempt to modify this information while the corresponding endpoint is enabled.

Table 16-75. Endpoint Capabilities/Characteristics

Bits	Name	Description
31–30	Mult	Mult. This field is used to indicate the number of packets executed per transaction description as given by the following: 00 - Execute N Transactions as demonstrated by the USB variable length packet protocol where N is computed using the Maximum Packet Length (dQH) and the Total Bytes field (dTD) 01 Execute 1 Transaction. 10 Execute 2 Transactions. 11 Execute 3 Transactions. <b>Note:</b> Non-ISO endpoints must set Mult = 00. <b>Note:</b> ISO endpoints must set Mult = 01, 10, or 11 as needed.
29	zlt	Zero length termination select. This bit is used to indicate when a zero length packet is used to terminate transfers where to total transfer length is a multiple. This bit is not relevant for Isochronous transfers. 0 Enable zero length packet to terminate transfers equal to a multiple of the Maximum Packet Length. (default). 1 Disable the zero length packet on transfers that are equal in length to a multiple Maximum Packet Length.
28–27	—	Reserved, should be cleared. These bit reserved for future use and should be cleared.

**Table 16-75. Endpoint Capabilities/Characteristics (continued)**

Bits	Name	Description
26–16	Maximum Packet Length	Maximum packet length. This directly corresponds to the maximum packet size of the associated endpoint (wMaxPacketSize). The maximum value this field may contain is 0x400 (1024).
15	ios	Interrupt on setup (IOS). This bit is used on control type endpoints to indicate if USBINT is set in response to a setup being received.
14–0		Reserved, should be cleared. Bits reserved for future use and should be cleared.

### 16.7.1.2 Transfer Overlay

The seven DWords in the overlay area represent a transaction working space for the device controller. The general operational model is that the device controller can detect whether the overlay area contains a description of an active transfer. If it does not contain an active transfer, then it will not read the associated endpoint.

After an endpoint is readied, the dTD will be copied into this queue head overlay area by the device controller. Until a transfer is expired, software must not write the queue head overlay area or the associated transfer descriptor. When the transfer is complete, the device controller will write the results back to the original transfer descriptor and advance the queue.

See dTD for a description of the overlay fields.

### 16.7.1.3 Current dTD Pointer

The current dTD pointer is used by the device controller to locate the transfer in progress. This word is for USB\_DR (hardware) use only and should not be modified by DCD software.

**Table 16-76. Current dTD Pointer**

Bits	Description
31–5	Current dtd. This field is a pointer to the dTD that is represented in the transfer overlay area. This field will be modified by the Device Controller to next dTD pointer during endpoint priming or queue advance.
4–0	Reserved, should be cleared. Bit reserved for future use and should be cleared.

### 16.7.1.4 Setup Buffer

The setup buffer is dedicated storage for the 8-byte data that follows a setup PID.

#### NOTE

Each endpoint has a TX and an RX dQH associated with it, and only the RX queue head is used for receiving setup data packets.

**Table 16-77. Multiple Mode Control**

DWord	Bits	Description
1	31–0	Setup Buffer 0. This buffer contains bytes 3 to 0 of an incoming setup buffer packet and is written by the device controller to be read by software.
2	31–0	Setup Buffer 1. This buffer contains bytes 7 to 4 of an incoming setup buffer packet and is written by the device controller to be read by software.

### 16.7.2 Endpoint Transfer Descriptor (dTD)

The dTD describes to the device controller the location and quantity of data to be sent/received for given transfer. The DCD should not attempt to modify any field in an active dTD except the Next Link Pointer, which should only be modified as described in section Managing Transfers with Transfer Descriptors.

31	30	29	28	27	26	25	24	23	22	21	20	19	18	17	16	15	14	13	12	11	10	9	8	7	6	5	4	3	2	1	0	offset
Next Link Pointer																												0000	T	0x00		
00		Total Bytes <sup>1</sup>										ioc	000		MultO	00		Status <sup>1</sup>										0x04				
Buffer Pointer (Page 0)											Current Offset <sup>1</sup>																	0x08				
Buffer Pointer (Page 1)											0	Frame Number <sup>1</sup>																	0x0C			
Buffer Pointer (Page 2)											0000_0000_0000																	0x10				
Buffer Pointer (Page 3)											0000_0000_0000																	0x14				
Buffer Pointer (Page 4)											0000_0000_0000																	0x18				

**Figure 16-62. Endpoint Transfer Descriptor (dTD)**

<sup>1</sup> Device controller read/write; all others read-only.

**Table 16-78. Next dTD Pointer**

Bits	Description
31–5	Next transfer element pointer. This field contains the physical memory address of the next dTD to be processed. The field corresponds to memory address signals [31:5], respectively.
4–1	Reserved, should be cleared. Bits reserved for future use and should be cleared.
0	Terminate (T). 1=pointer is invalid. 0=Pointer is valid (points to a valid Transfer Element Descriptor). This bit indicates to the Device Controller that there are no more valid entries in the queue.

Table 16-79. dTD Token

Bits	Description												
31	Reserved, should be cleared. Bit reserved for future use and should be cleared.												
30–16	<p>Total Bytes. This field specifies the total number of bytes to be moved with this transfer descriptor. This field is decremented by the number of bytes actually moved during the transaction and only on the successful completion of the transaction.</p> <p>The maximum value software may store in the field is 5*4K(5000H). This is the maximum number of bytes 5 page pointers can access. Although it is possible to create a transfer up to 20K this assumes the 1st offset into the first page is 0. When the offset cannot be predetermined, crossing past the 5th page can be guaranteed by limiting the total bytes to 16K**. Therefore, the maximum recommended transfer is 16K(4000H).</p> <p>If the value of the field is zero when the host controller fetches this transfer descriptor (and the active bit is set), the device controller executes a zero-length transaction and retires the transfer descriptor.</p> <p>It is not a requirement for IN transfers that Total Bytes To Transfer be an even multiple of Maximum Packet Length. If software builds such a transfer descriptor for an IN transfer, the last transaction will always be less than Maximum Packet Length.</p>												
15	Interrupt On Complete (IOC). This bit is used to indicate if USBINT is to be set in response to device controller being finished with this dTD.												
14–12	Reserved, should be cleared. Bits reserved for future use and should be cleared.												
11–10	<p>Multiplier Override (MultiO). This field can be used for transmit ISO's (that is, ISO-IN) to override the multiplier in the QH. This field must be zero for all packet types that are not transmit-ISO.</p> <p>Example:</p> <p>if QH.multiplier = 3; Maximum packet size = 8; Total bytes = 15; MultiO = 0 [default]  Three packets are sent: {Data2(8); Data1(7); Data0(0)}</p> <p>if QH.multiplier = 3; Maximum packet size = 8; Total bytes = 15; MultiO = 2  Two packets are sent: {Data1(8); Data0(7)}</p> <p>For maximal efficiency, software should compute MultiO = greatest integer of (Total Bytes/Max. Packet Size) except for the case when Total bytes = 0; then MultiO should be 1.</p> <p><b>Note:</b> Non-ISO and non-TX endpoints must set MultiO = 00.</p>												
9–8	Reserved, should be cleared. Bits reserved for future use and should be cleared.												
7–0	<p>Status. This field is used by the Device Controller to communicate individual command execution states back to the Device Controller software. This field contains the status of the last transaction performed on this qTD. The bit encodings are:</p> <table border="1"> <thead> <tr> <th>Bit</th> <th>Status Field Description</th> </tr> </thead> <tbody> <tr> <td>7</td> <td>Active</td> </tr> <tr> <td>6</td> <td>Halted</td> </tr> <tr> <td>5</td> <td>Data Buffer Error</td> </tr> <tr> <td>3</td> <td>Transaction Error</td> </tr> <tr> <td>4,2,0</td> <td>Reserved, should be cleared</td> </tr> </tbody> </table>	Bit	Status Field Description	7	Active	6	Halted	5	Data Buffer Error	3	Transaction Error	4,2,0	Reserved, should be cleared
Bit	Status Field Description												
7	Active												
6	Halted												
5	Data Buffer Error												
3	Transaction Error												
4,2,0	Reserved, should be cleared												

Table 16-80. Buffer Pointer Page 0

Bits	Description
31–12	Buffer Pointer. Selects the page offset in memory for the packet buffer. Non virtual memory systems will typically set the buffer pointers to a series of incrementing integers.
11–0	Current Offset. Offset into the 4kb buffer where the packet is to begin.

**Table 16-81. Buffer Pointer Page 1**

Bits	Description
31–12	Buffer Pointer. Selects the page offset in memory for the packet buffer. Non virtual memory systems will typically set the buffer pointers to a series of incrementing integers.
11	Reserved
10–0	Frame Number. Written by the device controller to indicate the frame number in which a packet finishes. This is typically be used to correlate relative completion times of packets on an ISO endpoint.

**Table 16-82. Buffer Pointer Pages 2–4**

Bits	Description
31–12	Buffer Pointer. Selects the page offset in memory for the packet buffer. Non virtual memory systems will typically set the buffer pointers to a series of incrementing integers.
11–0	Reserved

## 16.8 Device Operational Model

The function of the device operation is to transfer a request in the memory image to and from the Universal Serial Bus. Using a set of linked list transfer descriptors, pointed to by a queue head, the device controller will perform the data transfers. The following sections explain the use of the device controller from the device controller driver (DCD) point-of-view and further describe how specific USB bus events relate to status changes in the device controller programmer's interface.

### 16.8.1 Device Controller Initialization

After hardware reset, the USB DR module is disabled until the run/stop bit (USBCMD[RS]) is set to a '1'. In the disabled state, the pull-up on the USB D+ is not active which prevents an attach event from occurring. At a minimum, it is necessary to have the queue heads setup for endpoint zero before the device attach occurs. Shortly after the device is enabled, a USB reset will occur followed by setup packet arriving at endpoint 0. A queue head must be prepared so that the device controller can store the incoming setup packet.

The MPC8313E USB PHY and clock must be configured prior to initialization of the USB controller. Initialization of the MPC8313E USB PHY interface is performed via software control following a power-on reset.

In order to initialize a device, the software should perform the following steps:

1. Set the controller mode to device mode. Optionally set USBMODE[SDIS] (streaming disable).

#### NOTE

Transitioning from host mode to device mode requires a device controller reset before modifying USBMODE.

2. Optionally modify the BURSTSIZE register.
3. Program PORTSC[PTS] if using a non-ULPI PHY.



4. Set CONTROL[USB\_EN]
5. Allocate and initialize device queue heads in system memory Minimum: Initialize device queue heads 0 Tx and 0 Rx.

#### NOTE

All device queue heads must be initialized for control endpoints before the endpoint is enabled. Device queue heads for non-control endpoints must be initialized before the endpoint can be used.

For information on device queue heads, refer to [Section 16.7, “Device Data Structures.”](#)

6. Configure the ENDPOINTLISTADDR pointer.  
For additional information on ENDPOINTLISTADDR, refer to the register table.
7. Enable the microprocessor interrupt associated with the USB DR module and optionally change setting of USBCMD[ITC].  
Recommended: enable all device interrupts including: USBINT, USBERRINT, Port Change Detect, USB Reset Received, DCSuspend.  
For a list of available interrupts refer to the USBINTR and the USBSTS register tables.
8. Set USBCMD[RS] to run mode.  
After the run bit is set, a device reset will occur. The DCD must monitor the reset event and set the DEVICEADDR register, set the ENDPTCTRLx registers, and adjust the software state as described in the Bus Reset section of the following Port State and Control section below.

#### NOTE

Endpoint 0 is designed as a control endpoint only and does not need to be configured using ENDPTCTRL0 register.

It is also not necessary to initially prime Endpoint 0 because the first packet received will always be a setup packet. The contents of the first setup packet will require a response in accordance with USB device framework command set.

## 16.8.2 Port State and Control

From a chip or system reset, the USB\_DR enters the powered state. A transition from the powered state to the attach state occurs when the run/stop bit (USBCMD[RS]) is set to a ‘1’. After receiving a reset on the bus, the port will enter the defaultFS or defaultHS state in accordance with the protocol reset described in Appendix C.2 of the *USB Specification Rev. 2.0*. The following state diagram depicts the state of a USB 2.0 device.




Figure 16-63. USB 2.0 Device States

States powered, attach, defaultFS/HS, suspendFS/HS are implemented in the USB\_DR and are communicated to the DCD using the following status bits:

Table 16-83. Device Controller State Information Bits

Bits	Register
DCSuspend (SLI)	USBSTS
USB Reset Received (URI)	USBSTS
Port Change Detect (PCI)	USBSTS
High-Speed Port	PORTSC

It is the responsibility of the DCD to maintain a state variable to differentiate between the DefaultFS/HS state and the Address/Configured states. Change of state from Default to Address and the Configured states is part of the enumeration process described in the device framework section of the USB 2.0 Specification.

As a result of entering the Address state, the device address register (DEVICEADDR) must be programmed by the DCD.

Entry into the Configured indicates that all endpoints to be used in the operation of the device have been properly initialized by programming the ENDPTCTRL $n$  registers and initializing the associated queue heads.

### 16.8.2.1 Bus Reset

A bus reset is used by the host to initialize downstream devices. When a bus reset is detected, the USB\_DR controller will renegotiate its attachment speed, reset the device address to 0, and notify the DCD by interrupt (assuming the USB reset interrupt enable bit, USBINTR[URE], is set). After a reset is received, all endpoints (except endpoint 0) are disabled and any primed transactions will be canceled by the device controller. The concept of priming will be clarified below, but the DCD must perform the following tasks when a reset is received:

Clear all setup token semaphores by reading the ENDPTSETUPSTAT register and writing the same value back to the ENDPTSETUPSTAT register.

Clear all the endpoint complete status bits by reading the ENDPTCOMPLETE register and writing the same value back to the ENDPTCOMPLETE register.

Cancel all primed status by waiting until all bits in the ENDPTPRIME are 0 and then writing 0xFFFF\_FFFF to ENDPTFLUSH.

Read the reset bit in the PORTSC register (PORTSC[PR]) and make sure that it is still active. A USB reset will occur for a minimum of 3 ms and the DCD must reach this point in the reset cleanup before end of the reset occurs, otherwise a hardware reset of the device controller is recommended (rare.)

- A hardware reset can be performed by writing a one to the USB\_DR reset bit in (USBCMD[RST]). Note that a hardware reset will cause the device to detach from the bus by clearing USBCMD[RS] bit. Thus, the DCD must completely re-initialize the USB\_DR after a hardware reset.

Free all allocated dTDs because they will no longer be executed by the device controller. If this is the first time the DCD is processing a USB reset event, then it is likely that no dTDs have been allocated.

At this time, the DCD may release control back to the OS because no further changes to the device controller are permitted until a Port Change Detect is indicated.

After a Port Change Detect, the device has reached the default state and the DCD can read the PORTSC to determine if the device is operating in FS or HS mode. At this time, the device controller has reached normal operating mode and DCD can begin enumeration according to the USB Chapter 9, Device Framework.

#### NOTE

The device DCD may use the FS/HS mode information to determine the bandwidth mode of the device.

In some applications, it may not be possible to enable one or more pipes while in FS mode. Beyond the data rate issue, there is no difference in DCD operation between FS and HS modes.

## 16.8.2.2 Suspend/Resume

### 16.8.2.2.1 Suspend Description

In order to conserve power, USB\_DR automatically enters the suspended state when no bus traffic has been observed for a specified period. When suspended, the USB\_DR maintains any internal status, including its address and configuration. Attached devices must be prepared to suspend at any time they are powered, regardless of if they have been assigned a non-default address, are configured, or neither. Bus activity may cease due to the host entering a suspend mode of its own. In addition, a USB device shall also enter the suspended state when the hub port it is attached to is disabled.

The USB\_DR exits suspend mode when there is bus activity. It may also request the host to exit suspend mode or selective suspend by using electrical signaling to indicate remote wake-up. The ability of a device to signal remote wake-up is optional. The USB\_DR is capable of remote wake-up signaling. When the USB\_DR is reset, remote wake-up signaling must be disabled.

### 16.8.2.2.2 Suspend Operational Model

The USB\_DR moves into the suspend state when suspend signaling is detected or activity is missing on the upstream port for more than a specific period. After the device controller enters the suspend state, the DCD is notified by an interrupt (assuming DC Suspend Interrupt is enabled). When the USBSTS[SLI] (device controller suspend) is set, the device controller is suspended.

DCD response when the device controller is suspended is application specific and may involve switching to low power operation.

Information on the bus power limits in suspend state can be found in USB 2.0 specification.

### 16.8.2.2.3 Resume

If the USB\_DR is suspended, its operation is resumed when any non-idle signaling is received on its upstream facing port. In addition, the USB\_DR can signal the system to resume operation by forcing resume signaling to the upstream port. Resume signaling is sent upstream by writing a '1' to the PORTSC[FPR] (resume bit) while the device is in suspend state. Sending resume signal to an upstream port should cause the host to issue resume signaling and bring the suspended bus segment (one more devices) back to the active condition.

#### NOTE

Before resume signaling can be used, the host must enable it by using the Set Feature command defined in device framework (Chapter 9) of the USB 2.0 Specification.

## 16.8.3 Managing Endpoints

The USB 2.0 specification defines an endpoint, also called a device endpoint or an address endpoint as a uniquely addressable portion of a USB device that can source or sink data in a communications channel between the host and the device. The endpoint address is specified by the combination of the endpoint number and the endpoint direction.

The channel between the host and an endpoint at a specific device represents a data pipe. Endpoint 0 for a device is always a control type data channel used for device discovery and enumeration. Other types of endpoints supported by USB include bulk, interrupt, and isochronous. Each endpoint type has specific behavior related to packet response and error handling. More detail on endpoint operation can be found in the USB 2.0 specification.

The USB\_DR supports up to three endpoint specified numbers. The DCD can enable, disable, and configure each endpoint.

Each endpoint direction is essentially independent and can be configured with differing behavior in each direction. For example, the DCD can configure endpoint 1-IN to be a bulk endpoint and endpoint 1-OUT to be an isochronous endpoint. This helps to conserve the total number of endpoints required for device operation. The only exception is that control endpoints must use both directions on a single endpoint number to function as a control endpoint. Endpoint 0 is, for example, is always a control endpoint and uses the pair of directions.

Each endpoint direction requires a queue head allocated in memory. If the maximum of 6 endpoint numbers, one for each endpoint direction are being used by the device controller, then 12 queue heads are required. The operation of an endpoint and use of queue heads are described later in this document.

### 16.8.3.1 Endpoint Initialization

After hardware reset, all endpoints except endpoint zero are uninitialized and disabled. The DCD must configure and enable each endpoint by writing to configuration bit in the `ENDPTCTRLn` register. Each 32-bit `ENDPTCTRLn` is split into an upper and lower half. The lower half of `ENDPTCTRLn` is used to configure the receive or OUT endpoint and the upper half is likewise used to configure the corresponding transmit or IN endpoint. Control endpoints must be configured the same in both the upper and lower half of the `ENDPTCTRLn` register otherwise the behavior is undefined. The following table shows how to construct a configuration word for endpoint initialization.

**Table 16-84. Device Controller Endpoint Initialization**

Field	Value
Data Toggle Reset	1
Data Toggle Inhibit	0
Endpoint Type	00 Control 01 Isochronous 10 Bulk 11 Interrupt
Endpoint Stall	0

#### 16.8.3.1.1 Stalling

There are two occasions where the USB\_DR may need to return to the host a STALL.

The first occasion is the functional stall, which is a condition set by the DCD as described in the USB 2.0 device framework (Chapter 9). A functional stall is only used on non-control endpoints and can be enabled in the device controller by setting the endpoint stall bit in the `ENDPTCTRLn` register associated with the

given endpoint and the given direction. In a functional stall condition, the device controller will continue to return STALL responses to all transactions occurring on the respective endpoint and direction until the endpoint stall bit is cleared by the DCD.

A protocol stall, unlike a function stall, is used on control endpoints is automatically cleared by the device controller at the start of a new control transaction (setup phase). When enabling a protocol stall, the DCD should enable the stall bits (both directions) as a pair. A single write to the ENDPTCTRL $n$  register can ensure that both stall bits are set at the same instant.

**NOTE**

Any write to the ENDPTCTRL $n$  register during operational mode must preserve the endpoint type field (that is, perform a read-modify-write).

**Table 16-85. Device Controller Stall Response Matrix**

USB Packet	Endpoint Stall Bit.	Effect on STALL Bit.	USB Response
SETUP packet received by a non-control endpoint	N/A	None	STALL
IN/OUT/PING packet received by a non-control endpoint	'1	None	STALL
IN/OUT/PING packet received by a non-control endpoint	'0	None	ACK/NAK/NYET
SETUP packet received by a control endpoint	N/A	Cleared	ACK
IN/OUT/PING packet received by a control endpoint	'1	None	STALL
IN/OUT/PING packet received by a control endpoint	'0	None	ACK/NAK/NYET

**16.8.3.2 Data Toggle**

Data toggle is a mechanism to maintain data coherency between host and device for any given data pipe. For more information on data toggle, refer to the *Universal Serial Bus Revision 2.0 Specification*.

**16.8.3.2.1 Data Toggle Reset**

The DCD may reset the data toggle state bit and cause the data toggle sequence to reset in the device controller by writing a '1' to the data toggle reset bit in the ENDPTCTRL $n$  register. This should only be necessary when configuring/initializing an endpoint or returning from a STALL condition.

**16.8.3.2.2 Data Toggle Inhibit**

This feature is for test purposes only and should never be used during normal device controller operation.

Setting the data toggle Inhibit bit active ('1') causes the USB\_DR to ignore the data toggle pattern that is normally sent and accept all incoming data packets regardless of the data toggle state.

In normal operation, the USB\_DR checks the DATA0/DATA1 bit against the data toggle to determine if the packet is valid. If Data PID does not match the data toggle state bit maintained by the device controller for that endpoint, the Data toggle is considered not valid. If the data toggle is not valid, the device controller assumes the packet was already received and discards the packet (not reporting it to the DCD).

To prevent the USB\_DR from re-sending the same packet, the device controller will respond to the error packet by acknowledging it with either an ACK or NYET response.

### 16.8.3.3 Device Operational Model For Packet Transfers

All transactions on the USB bus are initiated by the host and in turn, the device must respond to any request from the host within the turnaround time stated in the *Universal Serial Bus Revision 2.0 Specification*.

A USB host will send requests to the USB\_DR in an order that can not be precisely predicted as a single pipeline, so it is not possible to prepare a single packet for the device controller to execute. However, the order of packet requests is predictable when the endpoint number and direction is considered. For example, if endpoint 2 (transmit direction) is configured as a bulk pipe, then we can expect the host will send IN requests to that endpoint. This USB\_DR prepares packets for each endpoint/direction in anticipation of the host request. The process of preparing the device controller to send or receive data in response to host initiated transaction on the bus is referred to as ‘priming’ the endpoint. This term will be used throughout the following documentation to describe the USB\_DR operation so the DCD can be architected properly use priming. Further, note that the term ‘flushing’ is used to describe the action of clearing a packet that was queued for execution.

#### 16.8.3.3.1 Priming Transmit Endpoints

Priming a transmit endpoint will cause the device controller to fetch the device transfer descriptor (dTD) for the transaction pointed to by the device queue head (dQH). After the dTD is fetched, it will be stored in the dQH until the device controller completes the transfer described by the dTD. Storing the dTD in the dQH allows the device controller to fetch the operating context needed to handle a request from the host without the need to follow the linked list, starting at the dQH when the host request is received.

After the device has loaded the dTD, the leading data in the packet is stored in a FIFO in the device controller. This FIFO is split into virtual channels so that the leading data can be stored for any endpoint up to the maximum number of endpoints configured at device synthesis time.

After a priming request is complete, an endpoint state of primed is indicated in the ENDPTSTATUS register. For a primed transmit endpoint, the device controller can respond to an IN request from the host and meet the stringent bus turnaround time of High Speed USB.

Since only the leading data is stored in the device controller FIFO, it is necessary for the device controller to begin filling in behind leading data after the transaction starts. The FIFO must be sized to account for the maximum latency that can be incurred by the system memory bus.

#### 16.8.3.3.2 Priming Receive Endpoints

Priming receive endpoints is identical to priming of transmit endpoints from the point of view of the DCD. At the device controller the major difference in the operational model is that there is no data movement of the leading packet data simply because the data is to be received from the host.

Note as part of the architecture, the FIFO for the receive endpoints is not partitioned into multiple channels like the transmit FIFO. Thus, the size of the RX FIFO does not scale with the number of endpoints.

### 16.8.3.4 Interrupt/Bulk Endpoint Operational Model

The behaviors of the device controller for interrupt and bulk endpoints are identical. All valid IN and OUT transactions to bulk pipes will handshake with a NAK unless the endpoint had been primed. Once the endpoint has been primed, data delivery will commence.

A dTD will be retired by the device controller when the packets described in the transfer descriptor have been completed. Each dTD describes N packets to be transferred according to the USB Variable Length transfer protocol. The formula and table on the following page describe how the device controller computes the number and length of the packets to be sent/received by the USB vary according to the total number of bytes and maximum packet length.

With Zero Length Termination (ZLT) = 0

$$N = \text{INT}(\text{number of bytes}/\text{max. packet length}) + 1$$

With Zero Length Termination (ZLT) = 1

$$N = \text{MAXINT}(\text{number of bytes}/\text{max. packet length})$$

**Table 16-86. Variable Length Transfer Protocol Example (ZLT = 0)**

Bytes (dTD)	Max. Packet Length (dQH)	N	P1	P2	P3
511	256	2	256	255	
512	256	3	256	256	0
512	512	2	512	0	

**Table 16-87. Variable Length Transfer Protocol Example (ZLT = 1)**

Bytes (dTD)	Max. Packet Length (dQH)	N	P1	P2	P3
511	256	2	256	255	
512	256	2	256	256	
512	512	1	512		

**NOTE**

The MULT field in the dQH must be set to '00' for bulk, interrupt, and control endpoints.

TX-dTD is complete when:

- All packets described dTD were successfully transmitted. \*\*\* Total bytes in dTD will equal zero when this occurs.

RX-dTD is complete when:

- All packets described in dTD were successfully received. \*\*\* Total bytes in dTD will equal zero when this occurs.



- A short packet (number of bytes < maximum packet length) was received. \*\*\* This is a successful transfer completion; DCD must check Total Bytes in dTD to determine the number of bytes that are remaining. From the total bytes remaining in the dTD, the DCD can compute the actual bytes received.
- A long packet was received (number of bytes > maximum packet size) OR (total bytes received > total bytes specified). \*\*\* This is an error condition. The device controller will discard the remaining packet, and set the Buffer Error bit in the dTD. In addition, the endpoint will be flushed and the USBERR interrupt will become active.

On the successful completion of the packet(s) described by the dTD, the active bit in the dTD will be cleared and the next pointer will be followed when the Terminate bit is clear. When the Terminate bit is set, the USB\_DR will flush the endpoint/direction and cease operations for that endpoint/direction.

On the unsuccessful completion of a packet (see long packet above), the dQH will be left pointing to the dTD that was in error. In order to recover from this error condition, the DCD must properly re-initialize the dQH by clearing the active bit and update the nextTD pointer before attempting to re-prime the endpoint.

### NOTE

All packet level errors such as a missing handshake or CRC error will be retried automatically by the device controller.

There is no required interaction with the DCD for handling such errors.

#### 16.8.3.4.1 Interrupt/Bulk Endpoint Bus Response Matrix

Table 16-88. Interrupt/Bulk Endpoint Bus Response Matrix

	Stall	Not Primed	Primed	Underflow	Overflow
<b>Setup</b>	Ignore	Ignore	Ignore	N/A	N/A
<b>In</b>	STALL	NAK	Transmit	BS Error <sup>1</sup>	N/A
<b>Out</b>	STALL	NAK	Receive + NYET/ACK <sup>2</sup>	N/A	NAK
<b>Ping</b>	STALL	NAK	ACK	N/A	N/A
<b>Invalid</b>	Ignore	Ignore	Ignore	Ignore	Ignore

<sup>1</sup> Force Bit Stuff Error.

<sup>2</sup> NYET/ACK—NYET unless the Transfer Descriptor has packets remaining according to the USB variable length protocol then ACK.  
SYSERR—System error should never occur when the latency FIFOs are correctly sized and the DCD is responsive.

## 16.8.3.5 Control Endpoint Operation Model

### 16.8.3.5.1 Setup Phase

All requests to a control endpoint begin with a setup phase followed by an optional data phase and a required status phase. The USB\_DR will always accept the setup phase unless the setup lockout is engaged.

The setup lockout will engage so that future setup packets are ignored. Lockout of setup packets ensures that while software is reading the setup packet stored in the queue head, that data is not written as it is being read potentially causing an invalid setup packet.

The setup lockout mechanism can be disabled and a tripwire type semaphore will ensure that the setup packet payload is extracted from the queue head without being corrupted by an incoming setup packet. This is the preferred behavior because ignoring repeated setup packets due to long software interrupt latency would be a compliance issue.

#### Setup Packet Handling

- Disable Setup Lockout by writing '1' to Setup Lockout Mode (SLOM) in USBMODE (once at initialization). Setup lockout is not necessary when using the tripwire as described below.

#### NOTE

Leaving the Setup Lockout Mode as '0' will result in a potential compliance issue.

- After receiving an interrupt and inspecting ENDPTSETUPSTAT to determine that a setup packet was received on a particular pipe:
  - Write '1' to clear corresponding bit ENDPTSETUPSTAT.
  - Write '1' to Setup Tripwire (SUTW) in USBCMD register.
  - Duplicate contents of dQH.SetupBuffer into local software byte array.
  - Read Setup TripWire (SUTW) in USBCMD register. (if set—continue; if cleared—goto 2)
  - Write '0' to clear Setup Tripwire (SUTW) in USBCMD register.
  - Process setup packet using local software byte array copy and execute status/handshake phases.

#### NOTE

After receiving a new setup packet the status and/or handshake phases may still be pending from a previous control sequence. These should be flushed and de-allocated before linking a new status and/or handshake dTD for the most recent setup packet.

### 16.8.3.5.2 Data Phase

Following the setup phase, the DCD must create a device transfer descriptor for the data phase and prime the transfer.

After priming the packet, the DCD must verify a new setup packet has not been received by reading the ENDPTSETUPSTAT register immediately verifying that the prime had completed. A prime will complete

when the associated bit in the ENDPTPRIME register is zero and the associated bit in the ENDPTSTATUS register is a one. If a prime fails, that is, The ENDPTPRIME bit goes to zero and the ENDPTSTATUS bit is not set, then the prime has failed. This can only be due to improper setup of the dQH, dTD or a setup arriving during the prime operation. If a new setup packet is indicated after the ENDPTPRIME bit is cleared, then the transfer descriptor can be freed and the DCD must reinterpret the setup packet.

Should a setup arrive after the data stage is primed, the device controller will automatically clear the prime status (ENDPTSTATUS) to enforce data coherency with the setup packet.

#### NOTE

The MULT field in the dQH must be set to '00' for bulk, interrupt, and control endpoints.

#### NOTE

Error handling of data phase packets is the same as bulk packets described previously.

#### 16.8.3.5.3 Status Phase

Similar to the data phase, the DCD must create a transfer descriptor (with byte length equal zero) and prime the endpoint for the status phase. The DCD must also perform the same checks of the ENDPTSETUPSTAT as described above in the data phase.

#### NOTE

The MULT field in the dQH must be set to '00' for bulk, interrupt, and control endpoints.

#### NOTE

Error handling of data phase packets is the same as bulk packets described previously.

#### 16.8.3.5.4 Control Endpoint Bus Response Matrix

Shown in the following table is the device controller response to packets on a control endpoint according to the device controller state.

**Table 16-89. Control Endpoint Bus Response Matrix**

Token Type	Endpoint State					Setup Lockout
	Stall	Not Primed	Primed	Underflow	Overflow	
Setup	ACK	ACK	ACK	N/A	SYSERR <sup>1</sup>	
In	STALL	NAK	Transmit	BS Error <sup>2</sup>	N/A	N/A
Out	STALL	NAK	Receive + NYET/ACK <sup>3</sup>	N/A	NAK	N/A

**Table 16-89. Control Endpoint Bus Response Matrix (continued)**

Token Type	Endpoint State					Setup Lockout
	Stall	Not Primed	Primed	Underflow	Overflow	
Ping	STALL	NAK	ACK	N/A	N/A	N/A
Invalid	Ignore	Ignore	Ignore	Ignore	Ignore	Ignore

<sup>1</sup> SYSERR—System error should never occur when the latency FIFOs are correctly sized and the DCD is responsive.

<sup>2</sup> Force Bit Stuff Error.

<sup>3</sup> NYET/ACK—NYET unless the Transfer Descriptor has packets remaining according to the USB variable length protocol then ACK.

### 16.8.3.6 Isochronous Endpoint Operational Model

Isochronous endpoints are used for real-time scheduled delivery of data and their operational model is significantly different than the host throttled Bulk, Interrupt, and Control data pipes. Real time delivery by the USB\_DR will be accomplished by the following:

- Exactly MULT Packets per (micro)Frame are transmitted/received. Note that MULT is a two-bit field in the device Queue Head. The variable length packet protocol is not used on isochronous endpoints.
- NAK responses are not used. Instead, zero length packets are sent in response to an IN request to unprimed endpoints. For unprimed RX endpoints, the response to an OUT transaction is to ignore the packet within the device controller.
- Prime requests always schedule the transfer described in the dTD for the next (micro)frame. If the ISO-dTD is still active after that frame, then the ISO-dTD will be held ready until executed or canceled by the DCD.

The USB\_DR in host mode uses the periodic frame list to schedule data exchanges to Isochronous endpoints. The operational model for device mode does not use such a data structure. Instead, the same dTD used for Control/Bulk/Interrupt endpoints is also used for isochronous endpoints. The difference is in the handling of the dTD.

The first difference between bulk and ISO-endpoints is that priming an ISO-endpoint is a delayed operation such that an endpoint will become primed only after a SOF is received. After the DCD writes the prime bit, the prime bit will be cleared as usual to indicate to software that the device controller completed a priming the dTD for transfer. Internal to the design, the device controller hardware masks that prime start until the next frame boundary. This behavior is hidden from the DCD but occurs so that the device controller can match the dTD to a specific (micro)frame.

Another difference with isochronous endpoints is that the transaction must wholly complete in a (micro)frame. Once an ISO transaction is started in a (micro)frame it will retire the corresponding dTD when MULT transactions occur or the device controller finds a fulfillment condition.

The transaction error bit set in the status field indicates a fulfillment error condition. When a fulfillment error occurs, the frame after the transfer failed to complete wholly, the device controller will force retire the ISO-dTD and move to the next ISO-dTD.

It is important to note that fulfillment errors are only caused due to partially completed packets. If no activity occurs to a primed ISO-dTD, the transaction will stay primed indefinitely. This means it is up to software discard transmit ISO-dTDs that pile up from a failure of the host to move the data.

Finally, the last difference with ISO packets is in the data level error handling. When a CRC error occurs on a received packet, the packet is not retried similar to bulk and control endpoints. Instead, the CRC is noted by setting the Transaction Error bit and the data is stored as usual for the application software to sort out.

- TX Packet Retired
  - MULT counter reaches zero.
  - Fulfillment Error [Transaction Error bit is set]
  - #Packets Occurred > 0 AND # Packets Occurred < MULT

#### NOTE

For TX-ISO, MULT Counter can be loaded with a lesser value in the dTD Multiplier Override field. If the Multiplier Override is zero, the MULT Counter is initialized to the Multiplier in the QH.

- RX Packet Retired:
  - MULT counter reaches zero.
  - Non-MDATA Data PID is received
  - Overflow Error:
    - Packet received is > maximum packet length. [Buffer Error bit is set]
    - Packet received exceeds total bytes allocated in dTD. [Buffer Error bit is set]
  - Fulfillment Error [Transaction Error bit is set]
  - # Packets Occurred > 0 AND # Packets Occurred < MULT
  - CRC Error [Transaction Error bit is set]

#### NOTE

For ISO, when a dTD is retired, the next dTD is primed for the next frame. For continuous (micro)frame to (micro)frame operation the DCD should ensure that the dTD linked-list is out ahead of the device controller by at least two (micro)frames.

### 16.8.3.6.1 Isochronous Pipe Synchronization

When it is necessary to synchronize an isochronous data pipe to the host, the (micro)frame number (FRINDEX register) can be used as a marker. To cause a packet transfer to occur at a specific (micro)frame number [N], the DCD should interrupt on SOF during frame N-1. When the FRINDEX=N-1, the DCD must write the prime bit. The USB\_DR will prime the isochronous endpoint in (micro)frame N-1 so that the device controller will execute delivery during (micro)frame N.

**CAUTION**

Priming an endpoint towards the end of (micro)frame N-1 will not guarantee delivery in (micro)frame N. The delivery may actually occur in (micro)frame N+1 if device controller does not have enough time to complete the prime before the SOF for packet N is received.

**16.8.3.6.2 Isochronous Endpoint Bus Response Matrix**

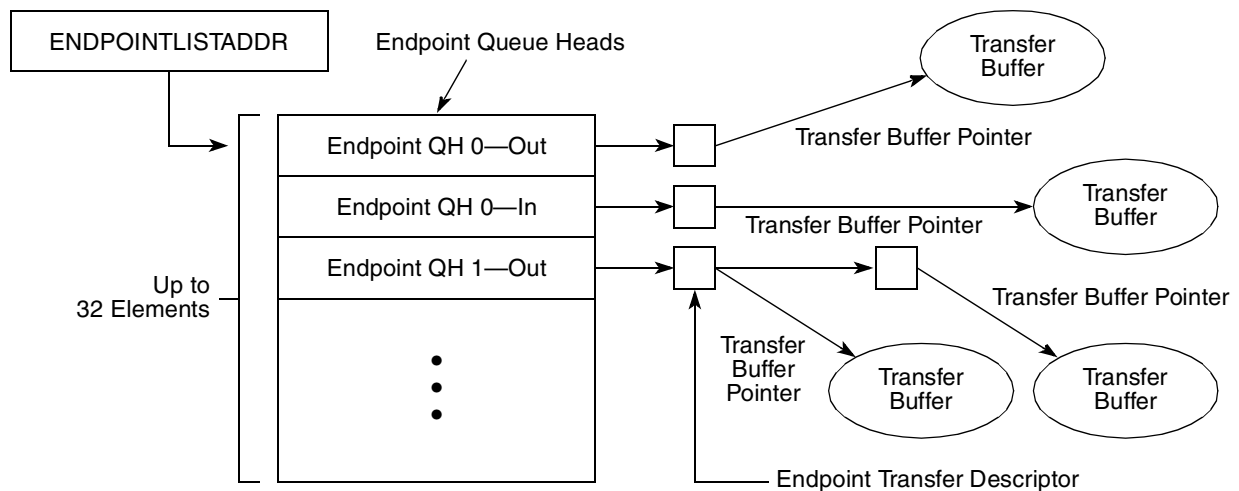
**Table 16-90. Isochronous Endpoint Bus Response Matrix**

	Stall	Not Primed	Primed	Underflow	Overflow
<b>Setup</b>	STALL	STALL	STALL	N/A	N/A
<b>In</b>	NULL <sup>1</sup> Packet	NULL Packet	Transmit	BS Error <sup>2</sup>	N/A
<b>Out</b>	Ignore	Ignore	Receive	N/A	Drop Packet
<b>Ping</b>	Ignore	Ignore	Ignore	Ignore	Ignore
<b>Invalid</b>	Ignore	Ignore	Ignore	Ignore	Ignore

<sup>1</sup> Zero Length Packet.

<sup>2</sup> Force Bit Stuff Error.

**16.8.4 Managing Queue Heads**



**Figure 16-64. Endpoint Queue Head Diagram**

The device queue head (dQH) points to the linked list of transfer tasks, each depicted by the device Transfer Descriptor (dTD). An area of memory pointed to by ENDPOINTLISTADDR contains a group of all dQH's in a sequential list as shown in Figure 16-64. The even elements in the list of dQH's are used for receive endpoints (OUT/SETUP) and the odd elements are used for transmit endpoints (IN/INTERRUPT). Device transfer descriptors are linked head to tail starting at the queue head and ending at a terminate bit. Once the dTD has been retired, it will no longer be part of the linked list from the queue head. Therefore,

software is required to track all transfer descriptors since pointers will no longer exist within the queue head once the dTD is retired (see section Software Link Pointers).

In addition to the current and next pointers and the dTD overlay examined in section Operational Model For Packet Transfers, the dQH also contains the following parameters for the associated endpoint: Multiplier, Maximum Packet Length, Interrupt On Setup. The complete initialization of the dQH including these fields is demonstrated in the next section.

#### 16.8.4.1 Queue Head Initialization

One pair of device queue heads must be initialized for each active endpoint. To initialize a device queue head:

- Write the wMaxPacketSize field as required by the USB Chapter 9 or application specific protocol.
- Write the multiplier field to 0 for control, bulk, and interrupt endpoints. For ISO endpoints, set the multiplier to 1, 2, or 3 as required bandwidth and in conjunction with the USB Chapter 9 protocol. Note that in FS mode, the multiplier field can only be 1 for ISO endpoints.
- Write the next dTD Terminate bit field to '1.'
- Write the Active bit in the status field to '0.'
- Write the Halt bit in the status field to '0.'

#### NOTE

The DCD must only modify dQH if the associated endpoint is not primed and there are no outstanding dTDs.

#### 16.8.4.2 Operational Model for Setup Transfers

As discussed in [Section 16.8.3.5, "Control Endpoint Operation Model,"](#) setup transfer requires special treatment by the DCD. A setup transfer does not use a dTD but instead stores the incoming data from a setup packet in an 8-byte buffer within the dQH.

Upon receiving notification of the setup packet, the DCD should handle the setup transfer as demonstrated here:

1. Copy setup buffer contents from dQH - RX to software buffer.
2. Acknowledge setup backup by writing a '1' to the corresponding bit in ENDPTSETUPSTAT.

#### NOTE

The acknowledge must occur before continuing to process the setup packet.

#### NOTE

After the acknowledge has occurred, the DCD must not attempt to access the setup buffer in the dQH - RX. Only the local software copy should be examined.

3. Check for pending data or status dTD's from previous control transfers and flush if any exist as discussed in section Flushing/De-priming an Endpoint.

**NOTE**

It is possible for the device controller to receive setup packets before previous control transfers complete. Existing control packets in progress must be flushed and the new control packet completed.

4. Decode setup packet and prepare data phase [optional] and status phase transfer as required by the USB Chapter 9 or application specific protocol.

## 16.8.5 Managing Transfers with Transfer Descriptors

### 16.8.5.1 Software Link Pointers

It is necessary for the DCD software to maintain head and tail pointers to the for the linked list of dTDs for each respective queue head. This is necessary because the dQH only maintains pointers to the current working dTD and the next dTD to be executed. The operations described in next section for managing dTD will assume the DCD can use reference the head and tail of the dTD linked list.

**NOTE**

To conserve memory, the reserved fields at the end of the dQH can be used to store the Head and Tail pointers but it still remains the responsibility of the DCD to maintain the pointers.

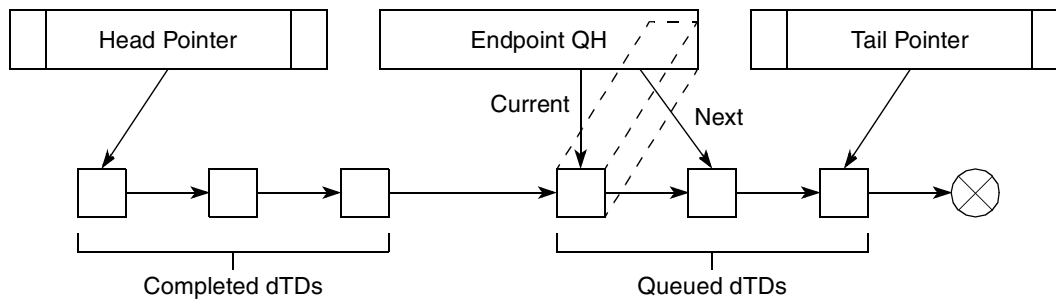


Figure 16-65. Software Link Pointers

### 16.8.5.2 Building a Transfer Descriptor

Before a transfer can be executed from the linked list, a dTD must be built to describe the transfer. Use the following procedure for building dTDs.

Allocate 8-DWord dTD block of memory aligned to 8-DWord boundaries. Example: bit address 4–0 would be equal to '00000'.

Write the following fields:

1. Initialize first 7 DWords to '0'.
2. Set the terminate bit to '1'.
3. Fill in total bytes with transfer size.
4. Set the interrupt on complete if desired.
5. Initialize the status field with the active bit set to '1' and all remaining status bits set to '0'.



6. Fill in buffer pointer page 0 and the current offset to point to the start of the data buffer.
7. Initialize buffer pointer page 1 through page 4 to be one greater than each of the previous buffer pointer.

### 16.8.5.3 Executing a Transfer Descriptor

To safely add a dTD, the DCD must follow this procedure which will handle the event where the device controller reaches the end of the dTD list at the same time a new dTD is being added to the end of the list.

Determine whether the link list is empty:

Check DCD driver to see if pipe is empty (internal representation of linked-list should indicate if any packets are outstanding).

Case 1: Link list is empty

1. Write dQH next pointer AND dQH terminate bit to '0' as a single DWord operation.
2. Clear active and halt bit in dQH (in case set from a previous error).
3. Prime endpoint by writing '1' to correct bit position in ENDPTPRIME.

Case 2: Link list is not empty

1. Add dTD to end of linked list.
2. Read correct prime bit in ENDPTPRIME—if '1' DONE.
3. Set ATDTW bit in USBCMD register to '1'.
4. Read correct status bit in ENDPTSTATUS. (store in tmp. variable for later).
5. Read ATDTW bit in USBCMD register.
  - If '0' goto 3.
  - If '1' continue to 6.
6. Write ATDTW bit in USBCMD register to '0'.
7. If status bit read in (3) is '1' DONE.
8. If status bit read in (3) is '0' then Goto Case 1: Step 1.

### 16.8.5.4 Transfer Completion

After a dTD has been initialized and the associated endpoint primed the device controller will execute the transfer upon the host-initiated request. The DCD will be notified with a USB interrupt if the Interrupt On Complete bit was set or alternately, the DCD can poll the endpoint complete register to find when the dTD had been executed. After a dTD has been executed, DCD can check the status bits to determine success or failure.

#### **CAUTION**

Multiple dTD can be completed in a single endpoint complete notification. After clearing the notification, DCD must search the dTD linked list and retire all dTDs that have finished (Active bit cleared).

By reading the status fields of the completed dTDs, the DCD can determine if the transfers completed successfully. Success is determined with the following combination of status bits:

- Active = 0
- Halted = 0
- Transaction Error = 0
- Data Buffer Error = 0

Should any combination other than the one shown above exist, the DCD must take proper action. Transfer failure mechanisms are indicated in the Device Error Matrix.

In addition to checking the status bit the DCD must read the Transfer Bytes field to determine the actual bytes transferred. When a transfer is complete, the Total Bytes transferred is by decremented by the actual bytes transferred. For Transmit packets, a packet is only complete after the actual bytes reaches zero, but for receive packets, the host may send fewer bytes in the transfer according the USB variable length packet protocol.

### 16.8.5.5 Flushing/De-Priming an Endpoint

It is necessary for the DCD to flush to de-prime one more endpoints on a USB device reset or during a broken control transfer. There may also be application specific requirements to stop transfers in progress. The following procedure can be used by the DCD to stop a transfer in progress:

1. Write a ‘1’ to the corresponding bit(s) in ENDPTFLUSH.
2. Wait until all bits in ENDPTFLUSH are ‘0’.
3. Software note: this operation may take a large amount of time depending on the USB bus activity. It is not desirable to have this wait loop within an interrupt service routine.
4. Read ENDPTSTATUS to ensure that for all endpoints commanded to be flushed, that the corresponding bits are now ‘0’” If the corresponding bits are ‘1’ after step #2 has finished, then the flush failed as described in the following:

Explanation: In very rare cases, a packet is in progress to the particular endpoint when commanded flush using ENDPTFLUSH. A safeguard is in place to refuse the flush to ensure that the packet in progress completes successfully. The DCD may need to repeatedly flush any endpoints that fail to flush be repeating steps 1–3 until each endpoint is successfully flushed.

### 16.8.5.6 Device Error Matrix

The following table summarizes packet errors that are not automatically handled by the USB\_DR

**Table 16-91. Device Error Matrix**

Error	Direction	Packet Type	Data Buffer Error Bit	Transaction Error Bit
Overflow **	RX	Any	1	0
ISO Packet Error	RX	ISO	0	1
ISO Fulfillment Error	Both	ISO	0	1

Notice that the device controller handles all errors on Bulk/Control/Interrupt Endpoints except for a data buffer overflow. However, for ISO endpoints, errors packets are not retried and errors are tagged as indicated.

**Table 16-92. Error Descriptions**

<b>Overflow</b>	Number of bytes received exceeded max. packet size or total buffer length. ** This error will also set the Halt bit in the dQH and if there are dTDs remaining in the linked list for the endpoint, then those will not be executed.
<b>ISO Packet Error</b>	CRC Error on received ISO packet. Contents not guaranteed to be correct.
<b>ISO Fulfillment Error</b>	Host failed to complete the number of packets defined in the dQH mult field within the given (micro)frame. For scheduled data delivery the DCD may need to readjust the data queue because a fulfillment error will cause Device Controller to cease data transfers on the pipe for one (micro)frame. During the 'dead' (micro)frame, the Device Controller reports error on the pipe and primes for the following frame.

## 16.8.6 Servicing Interrupts

The interrupt service routine must consider that there are high-frequency, low-frequency operations, and error operations and order accordingly.

### 16.8.6.1 High-Frequency Interrupts

High frequency interrupts in particular should be handed in the order below. The most important of these is listed first because the DCD must acknowledge a setup buffer in the timeliest manner possible.

**Table 16-93. Interrupt Handling Order**

Execution Order	Interrupt	Action
1a	USB Interrupt <sup>1</sup> ENDPTSETUPSTATUS	Copy contents of setup buffer and acknowledge setup packet (as indicated in section Managing Queue Heads). Process setup packet according to USB 2.0 Chapter 9 or application specific protocol.
1b	USB Interrupt ENDPTCOMPLETE	Handle completion of dTD as indicated in section Managing Queue Heads.
2	SOF Interrupt	Action as deemed necessary by application. This interrupt may not have a use in all applications.

<sup>1</sup> It is likely that multiple interrupts to stack up on any call to the Interrupt Service Routine AND during the Interrupt Service Routine.

### 16.8.6.2 Low-Frequency Interrupts

The low frequency events include the following interrupts. These interrupt can be handled in any order since they don't occur often in comparison to the high-frequency interrupts.

**Table 16-94. Low Frequency Interrupt Events**

Interrupt	Action
Port Change	Change software state information.
Sleep Enable (Suspend)	Change software state information. Low power handling as necessary.
Reset Received	Change software state information. Abort pending transfers.

### 16.8.6.3 Error Interrupts

Error interrupts will be least frequent and should be placed last in the interrupt service routine.

**Table 16-95. Error Interrupt Events**

Interrupt	Action
USB Error Interrupt	This error is redundant because it combines USB Interrupt and an error status in the dTD. The DCD will more aptly handle packet-level errors by checking dTD status field upon receipt of USB Interrupt (w/ ENDPTCOMPLETE).
System Error	Unrecoverable error. Immediate Reset of core; free transfers buffers in progress and restart the DCD.

## 16.9 Deviations from the EHCI Specifications

The host mode operation of the USB DR module is nearly EHCI-compatible with few minor differences. For the most part, the module conforms to the data structures and operations described in Section 3, “Data Structures,” and Section 4, “Operational Model,” in the EHCI specification. The particulars of the deviations occur in the following areas:

- Embedded transaction translator—Allows direct attachment of FS and LS devices in host mode without the need for a companion controller.
- Device operation—In host mode, the device operational registers are generally disabled and thus device mode is mostly transparent when in host mode. However, there are a couple exceptions documented in the following sections.
- Embedded design interface—The module does not have a PCI interface and therefore the PCI configuration registers described in the EHCI specification are not applicable.

For the purposes of the DR implementing dual-role host/device controller with support for OTG applications, it is necessary to deviate from the EHCI specification. Device operation and OTG operation are not specified in the EHCI and thus the implementation supported in the DR module is proprietary.

### 16.9.1 Embedded Transaction Translator Function

The DR module supports directly connected full and low speed devices without requiring a companion controller by including the capabilities of a USB 2.0 high speed hub transaction translator. Although there is no separate transaction translator block in the system, the transaction translator function normally associated with a high speed hub has been implemented within the DMA and Protocol engine blocks. The embedded transaction translator function is an extension to EHCI interface, but makes use of the standard

data structures and operational models that exist in the EHCI specification to support full and low speed devices.

### 16.9.1.1 Capability Registers

The following additions have been added to the capability registers to support the embedded transaction translator Function:

- N\_TT added to HSCPARAMS—Host Controller Structural Parameters
- N\_PTT added to HSCPARAMS—Host Controller Structural Parameters

See [Section 16.3.1.3, “Host Controller Structural Parameters \(HSCPARAMS\),”](#) for usage information.

### 16.9.1.2 Operational Registers

The following additions have been added to the operational registers to support the embedded TT:

- ASYNCTTSTS is a new register.
- Addition of two-bit Port Speed (PSPD) to the PORTSC register.

### 16.9.1.3 Discovery

In a standard EHCI controller design, the EHCI host controller driver detects a Full speed (FS) or Low speed (LS) device by noting if the port enable bit is set after the port reset operation. The port enable will only be set in a standard EHCI controller implementation after the port reset operation and when the host and device negotiate a High-Speed connection (that is, Chirp completes successfully).

The module will always set the port enable after the port reset operation regardless of the result of the host device chirp result and the resulting port speed will be indicated by the PSPD field in PORTSC. Therefore, the standard EHCI host controller driver requires an alteration to handle directly connected Full and Low speed devices or hubs. The change is a fundamental one in that is summarized in [Table 16-96](#).

**Table 16-96. Functional Differences Between EHCI and EHCI with Embedded TT**

Standard EHCI	EHCI with Embedded Transaction Translator
After port enable bit is set following a connection and reset sequence, the device/hub is assumed to be HS.	After port enable bit is set following a connection and reset sequence, the device/hub speed is noted from PORTSC.
FS and LS devices are assumed to be downstream from a HS hub thus, all port-level control is performed through the Hub Class to the nearest Hub.	FS and LS device can be either downstream from a HS hub or directly attached. When the FS/LS device is downstream from a HS hub, then port-level control is done using the Hub Class through the nearest Hub. When a FS/LS device is directly attached, then port-level control is accomplished using PORTSC.
FS and LS devices are assumed to be downstream from a HS hub with HubAddr=X. [where HubAddr > 0 and HubAddr is the address of the Hub where the bus transitions from HS to FS/LS (that is, Split target hub)]	FS and LS device can be either downstream from a HS hub with HubAddr = X [HubAddr > 0] or directly attached [where HubAddr = 0 and HubAddr is the address of the Root Hub where the bus transitions from HS to FS/LS (that is, Split target hub is the root hub)]

### 16.9.1.4 Data Structures

The same data structures used for FS/LS transactions through a HS hub are also used for transactions through the Root Hub. Here it is demonstrated how the Hub Address and Endpoint Speed fields should be set for directly attached FS/LS devices and hubs:

1. QH (for direct attach FS/LS)—Async. (Bulk/Control Endpoints) Periodic (Interrupt)
  - Hub Address = 0
  - Transactions to direct attached device/hub.
    - QH.EPS = Port Speed
  - Transactions to a device downstream from direct attached FS hub.
    - QH.EPS = Downstream Device Speed

#### NOTE

When QH.EPS = 01 (LS) and PORTSC[PSPD] = 00 (FS), a LS-pre-pid will be sent before the transmitting LS traffic.

Maximum Packet Size must be less than or equal 64 or undefined behavior may result.

2. siTD (for direct attach FS)—Periodic (ISO Endpoint)
  - All FS ISO transactions:
    - Hub Address = 0
    - siTD.EPS = 00 (full speed)

Maximum Packet Size must less than or equal to 1023 or undefined behavior may result.

### 16.9.1.5 Operational Model

The operational models are well defined for the behavior of the transaction translator (see *Universal Serial Bus Revision 2.0 Specification*) and for the EHCI controller moving packets between system memory and a USB-HS hub. Since the embedded transaction translator exists within the DR module there is no physical bus between EHCI host controller driver and the USB FS/LS bus. These sections will briefly discuss the operational model for how the EHCI and transaction translator operational models are combined without the physical bus between. The following sections assume the reader is familiar with both the EHCI and USB 2.0 transaction translator operational models.

#### 16.9.1.5.1 Microframe Pipeline

The EHCI operational model uses the concept of H-frames and B-frames to describe the pipeline between the Host (H) and the Bus (B). The embedded transaction translator shall use the same pipeline algorithms specified in the *Universal Serial Bus Revision 2.0 Specification* for a Hub-based transaction translator.

All periodic transfers always begin at B-frame 0 (after SOF) and continue until the stored periodic transfers are complete. As an example of the microframe pipeline implemented in the embedded transaction translator, all periodic transfers that are tagged in EHCI to execute in H-frame 0 will be ready to execute on the bus in B-frame 0.

It is important to note that when programming the S-mask and C-masks in the EHCI data structures to schedule periodic transfers for the embedded transaction translator, the EHCI host controller driver must

follow the same rules specified in EHCI for programming the S-mask and C-mask for downstream Hub-based transaction translators.

Once periodic transfers are exhausted, any stored asynchronous transfer will be moved. Asynchronous transfers are opportunistic in that they shall execute whenever possible and their operation is not tied to H-frame and B-frame boundaries with the exception that an asynchronous transfer can not babble through the SOF (start of B-frame 0).

### 16.9.1.5.2 Split State Machines

The start and complete split operational model differs from EHCI slightly because there is no bus medium between the EHCI controller and the embedded transaction translator. Where a start or complete-split operation would occur by requesting the split to the HS hub, the start/complete split operation is simple an internal operation to the embedded transaction translator. [Table 16-97](#) summarizes the conditions where handshakes are emulated from internal state instead of actual handshakes to HS split bus traffic.

**Table 16-97. Emulated Handshakes**

Condition	Emulate TT Response
<b>Start-Split:</b> All asynchronous buffers full	NAK
<b>Start-Split:</b> All periodic buffers full	ERR
<b>Start-Split:</b> Success for start of Async. Transaction	ACK
<b>Start-Split:</b> Start Periodic Transaction	No Handshake (Ok)
<b>Complete-Split:</b> Failed to find transaction in queue	Bus Time Out
<b>Complete-Split:</b> Transaction in Queue is Busy	NYET
<b>Complete-Split:</b> Transaction in Queue is Complete	[Actual Handshake from FS/LS device]

### 16.9.1.5.3 Asynchronous Transaction Scheduling and Buffer Management

The following *Universal Serial Bus Revision 2.0 Specification* items are implemented in the embedded transaction translator:

- USB 2.0–11.17.3
  - Sequencing is provided and a packet length estimator ensures no full-speed/low-speed packet babbles into SOF time.
- USB 2.0–11.17.4
  - • Transaction tracking for 2 data pipes.
- USB 2.0–11.17.5
  - • Clear\_TT\_Buffer capability provided

### 16.9.1.5.4 Periodic Transaction Scheduling and Buffer Management

The following *Universal Serial Bus Revision 2.0 Specification* items are implemented in the embedded transaction translator:

- USB 2.0–11.18.6.[1-2]

- Abort of pending start-splits
  - EOF (and not started in microframes 6)
  - Idle for more than 4 microframes
- Abort of pending complete-splits
  - EOF
  - Idle for more than 4 microframes

#### NOTE

There is no data schedule mechanism for these transactions other than the microframe pipeline. The embedded TT assumes the number of packets scheduled in a frame does not exceed the frame duration (1 msec) or else undefined behavior may result.

#### 16.9.1.5.5 Multiple Transaction Translators

The maximum number of embedded transaction translators that is currently supported is one as indicated by the N\_TT field in the HCSPARAMS register. See [Section 16.3.1.3, “Host Controller Structural Parameters \(HCSPARAMS\),”](#) for more information.

### 16.9.2 Device Operation

The co-existence of a device operational controller within the DR module has little effect on EHCI compatibility for host operation except as noted in this section.

### 16.9.3 Non-Zero Fields the Register File

Some of the reserved fields and reserved addresses in the capability registers and operational registers have use in device mode, the following must be adhered to:

- Write operations to all EHCI reserved fields (some of which are device fields in the DR module) in the operation registers should always be written to zero. This is an EHCI requirement of the device controller driver that must be adhered to.
- Read operations by the module must properly mask EHCI reserved fields (some of which are device fields in the DR module registers).

### 16.9.4 SOF Interrupt

The SOF interrupt is a free running 125  $\mu$ sec interrupt for host mode. EHCI does not specify this interrupt, but it has been added for convenience and as a potential software time base. Note that the free running interrupt is shared with the device-mode start-of-frame interrupt. See [Section 16.3.2.2, “USB Status Register \(USBSTS\),”](#) and [Section 16.3.2.3, “USB Interrupt Enable Register \(USBINTR\),”](#) for more information.



## 16.9.5 Embedded Design

This is an Embedded USB Host Controller as defined by the EHCI specification and thus does not implement the PCI configuration registers.

### 16.9.5.1 Frame Adjust Register

Given that the optional PCI configuration registers are not included in this implementation, there is no corresponding bit level timing adjustments like those provided by the Frame Adjust register in the PCI configuration registers. Starts of microframes are timed precisely to 125  $\mu$ sec using the transceiver clock as a reference clock. That is, 60 MHz transceiver clock for 8-bit physical interfaces and full-speed serial interfaces or 30 MHz transceiver clock for 16-bit physical interfaces.

## 16.9.6 Miscellaneous Variations from EHCI

### 16.9.6.1 Programmable Physical Interface Behavior

The modules support multiple physical interfaces which can operate in different modes when the module is configured with the software programmable Physical Interface Modes. The control bits for selecting the PHY operating mode have been added to the PORTSC register providing a capability that is not defined by the EHCI specification.

### 16.9.6.2 Discovery

#### 16.9.6.2.1 Port Reset

The port connect methods specified by EHCI require setting the port reset bit in the register for a duration of 10 msec. Due to the complexity required to support the attachment of devices that are not high speed there are counter already present in the design that can count the 10 msec reset pulse to alleviate the requirement of the software to measure this duration. Therefore, the basic connection is then summarized as the following:

- [Port Change Interrupt] Port connect change occurs to notify the host controller driver that a device has attached.
- Software shall write a '1' to the reset the device.
- Software shall write a '0' to the reset the device after 10 msec.
  - This step, which is necessary in a standard EHCI design, may be omitted with this implementation. Should the EHCI host controller driver attempt to write a '0' to the reset bit while a reset is in progress the write will simple be ignored and the reset will continue until completion.
- [Port Change Interrupt] Port enable change occurs to notify the host controller that the device in now operational and at this point the port speed has been determined.

### 16.9.6.2.2 Port Speed Detection

After the port change interrupt indicates that a port is enabled, the EHCI stack should determine the port speed. Unlike the EHCI implementation which will re-assign the port owner for any device that does not connect at High-Speed, this host controller supports direct attach of non-HS devices. Therefore, the following differences are important regarding port speed detection:

- Port owner is read-only and always reads 0.
- A 2-bit port speed indicator has been added to PORTSC to provide the current operating speed of the port to the host controller driver.
- A 1-bit high-speed indicator has been added to PORTSC to signify that the port is in HS vs. FS/LS

## 16.10 Timing Diagrams

This section contains diagrams showing the basic operation of the ULPI interface. For a more detailed description refer to the ULPI Specifications.






Figure 16-66. ULPI Timing

Table 16-98. ULPI Timing


Parameter	Symbol	Min	Max	Units
Control signal setup time	T <sub>SC</sub>	—	4	ns
Data setup time	T <sub>SD</sub>	—	4	ns
Control signal hold time	T <sub>HC</sub>	0	—	ns
Data hold time	T <sub>HD</sub>	0	—	ns
Control output delay	T <sub>DC</sub>	2	7	ns
Data output delay	T <sub>DD</sub>	2	7	ns



**Figure 16-67. Sending of RX CMD**



**Figure 16-68. ULPI Data Transmit (NOPID)**



**Figure 16-69. ULPI Data Transmit (PID)**




Figure 16-70. ULPI Data Receive




Figure 16-71. ULPI Register Write




Figure 16-72. ULPI Register Read

# Chapter 17

## I<sup>2</sup>C Interfaces

This chapter describes the two inter-IC (IIC or I<sup>2</sup>C) bus interfaces implemented on this device. Note that for most intents, the I<sup>2</sup>C interfaces are identical and are described as a single generic controller. Where necessary, differences between the two controllers are noted.

### 17.1 Introduction

The inter-IC (IIC or I<sup>2</sup>C) bus is a two-wire—serial data (SDA) and serial clock (SCL)—bidirectional serial bus that provides a simple, efficient method of data exchange between this device and other devices, such as microcontrollers, EEPROMs, real-time clock devices, A/D converters, and LCDs. Figure 17-1 shows a block diagram of the I<sup>2</sup>C interface.




Figure 17-1. I<sup>2</sup>C Block Diagram

The two-wire I<sup>2</sup>C bus minimizes interconnections between devices. The synchronous, multiple-master I<sup>2</sup>C bus allows the connection of additional devices to the bus for expansion and system development. The bus

includes collision detection and arbitration that prevent data corruption if two or more masters attempt to control the bus simultaneously.

### 17.1.1 Features

Each I<sup>2</sup>C interface includes the following features:

- Two-wire interface
- Multiple-master operational
- Arbitration lost interrupt with automatic mode switching from master to slave
- Calling address identification interrupt
- START and STOP signal generation/detection
- Acknowledge bit generation/detection
- Bus busy detection
- Software-programmable clock frequency
- Software-selectable acknowledge bit
- On-chip filtering for spikes on the bus

### 17.1.2 Modes of Operation

The I<sup>2</sup>C unit on this device can operate in one of the following modes:

- Master mode. The I<sup>2</sup>C initiates a transfer, generates clock signals, and terminates a transfer. It cannot use its own slave address as a calling address. The I<sup>2</sup>C cannot be a master and a slave simultaneously.
- Slave mode. The I<sup>2</sup>C is addressed by an I<sup>2</sup>C master. The module must be enabled before a START condition from an I<sup>2</sup>C master is detected.
- Interrupt-driven byte-to-byte data transfer. When successful slave addressing is achieved (and SCL<sub>n</sub> returns to zero), the data transfer can proceed on a byte-to-byte basis in the direction specified by the R/W bit sent by the calling master. Each byte of data must be followed by an acknowledge bit, which is signaled from the receiving device. Several bytes can be transferred during a data transfer session.
- Boot sequencer mode. I<sup>2</sup>C1 controller supports boot sequencer mode. This mode can be used to initialize the configuration registers in the device after the I<sup>2</sup>C module is initialized. Boot sequencer mode is selected using the BOOTSEQ field in the reset configuration word high. Note that the hard-coded reset configuration word high value is boot sequencer mode disabled. This mode is not supported by I<sup>2</sup>C2 controller.
- Reset configuration load (I<sup>2</sup>C1 only). In this mode, the I<sup>2</sup>C1 interface loads the reset configuration words from an EEPROM at a specific calling address while the rest of the device is in the reset state ( $\overline{\text{HRESET}}$  asserted). Once the reset configuration words are latched inside the device, I<sup>2</sup>C1 is reset until  $\overline{\text{HRESET}}$  is negated. After  $\overline{\text{HRESET}}$  is negated, the device may be initialized using boot sequence mode according to the BOOTSEQ field in the reset configuration word. See [Section 17.4.5, “Boot Sequencer Mode.”](#)

Additionally, the following three I<sup>2</sup>C–specific states are defined for the I<sup>2</sup>C interface:

- **START condition.** This condition denotes the beginning of a new data transfer (each data transfer contains several bytes of data) and awakens all slaves.
- **Repeated START condition.** A START condition that is generated without a STOP condition to terminate the previous transfer.
- **STOP condition.** The master can terminate the transfer by generating a STOP condition to free the bus.

## 17.2 External Signal Descriptions

The following sections give an overview of signals and provide detailed signal descriptions.

### 17.2.1 Signal Overview

The I<sup>2</sup>C interface uses the SDA<sub>*n*</sub> and SCL<sub>*n*</sub> signals, described in [Table 17-1](#), for data transfer. Note that the signal patterns driven on SDA<sub>*n*</sub> represent address, data, or read/write information at different stages of the protocol.

**Table 17-1. I<sup>2</sup>C Interface Signal Descriptions**

Signal Name	Idle State	I/O	State Meaning
Serial Clock (SCL1, SCL2)	High	I	When the I <sup>2</sup> C module is idle or acts as a slave, SCL <sub><i>n</i></sub> defaults as an input. The unit uses SCL <sub><i>n</i></sub> to synchronize incoming data on SDA <sub><i>n</i></sub> . The bus is assumed to be busy when SCL <sub><i>n</i></sub> is detected low.
		O	As a master, the I <sup>2</sup> C module drives SCL <sub><i>n</i></sub> along with SDA <sub><i>n</i></sub> when transmitting. As a slave, the I <sup>2</sup> C module drives SCL <sub><i>n</i></sub> negates for data pacing.
Serial Data (SDA1, SDA2)	High	I	When the I <sup>2</sup> C module is idle or in a receiving mode, SDA <sub><i>n</i></sub> defaults as an input. The unit receives data from other I <sup>2</sup> C devices on SDA <sub><i>n</i></sub> . The bus is assumed to be busy when SDA <sub><i>n</i></sub> is detected low.
		O	When writing as a master or slave, the I <sup>2</sup> C module drives data on SDA <sub><i>n</i></sub> synchronous to SCL <sub><i>n</i></sub> .

### 17.2.2 Detailed Signal Descriptions

SDA<sub>*n*</sub> and SCL<sub>*n*</sub>, described in [Table 17-2](#), serve as a communication interconnect with other devices. All devices connected to these signals must have open-drain or open-collector outputs. The logic AND function is performed on both of these signals with external pull-up resistors. Refer to the hardware specifications for electrical characteristics.

Table 17-2. I<sup>2</sup>C Interface Signals—Detailed Signal Descriptions

Signal	I/O	Description
SCL1, SCL2	I/O	Serial clock. Performs as an input when the device is programmed as an I <sup>2</sup> C slave. SCL <sub>n</sub> also performs as an output when the device is programmed as an I <sup>2</sup> C master.
	O	As outputs for the bidirectional serial clock, these signals operate as described below.
		<b>State Meaning</b>
	I	As inputs for the bi-directional serial clock, these signals operate as described below.
<b>State Meaning</b>		Asserted/Negated—The I <sup>2</sup> C unit uses this signal to synchronize incoming data on SDA <sub>n</sub> . The bus is assumed to be busy when this signal is detected low.
SDA1, SDA2	I/O	Serial data. Performs as an input when the device is in a receiving mode. SDA <sub>n</sub> also performs as an output signal when the device is transmitting (as an I <sup>2</sup> C master or a slave).
	O	As outputs for the bi-directional serial data, these signals operate as described below.
		<b>State Meaning</b>
	I	As inputs for the bi-directional serial data, these signals operate as described below.
<b>State Meaning</b>		Asserted/Negated—Used to receive data from other devices. The bus is assumed to be busy when SDA <sub>n</sub> is detected low.

## 17.3 Memory Map/Register Definition

Table 17-3 lists the I<sup>2</sup>C–specific registers and their addresses.

Table 17-3. I<sup>2</sup>C Memory Map

Address	I <sup>2</sup> C Register	Access	Reset	Section/Page
0x0_3000	I2C1ADR—I <sup>2</sup> C1 address register	R/W	0x00	<a href="#">17.3.1.1/17-5</a>
0x0_3004	I2C1FDR—I <sup>2</sup> C1 frequency divider register	R/W	0x00	<a href="#">17.3.1.2/17-6</a>
0x0_3008	I2C1CR—I <sup>2</sup> C1 control register	R/W	0x00	<a href="#">17.3.1.3/17-7</a>
0x0_300C	I2C1SR—I <sup>2</sup> C1 status register	R/W	0x81	<a href="#">17.3.1.4/17-8</a>
0x0_3010	I2C1DR—I <sup>2</sup> C1 data register	R/W	0x00	<a href="#">17.3.1.5/17-9</a>
0x0_3014	I2C1DFSRR—I <sup>2</sup> C1 digital filter sampling rate register	R/W	0x10	<a href="#">17.3.1.6/17-9</a>
0x0_301C– 0x0_30FF	Reserved, should be cleared	—	—	—
0x0_3100	I2C2ADR—I <sup>2</sup> C2 address register	R/W	0x00	<a href="#">17.3.1.1/17-5</a>
0x0_3104	I2C2FDR—I <sup>2</sup> C2 frequency divider register	R/W	0x00	<a href="#">17.3.1.2/17-6</a>
0x0_3108	I2C2CR—I <sup>2</sup> C2 control register	R/W	0x00	<a href="#">17.3.1.3/17-7</a>
0x0_310C	I2C2SR—I <sup>2</sup> C2 status register	R/W	0x81	<a href="#">17.3.1.4/17-8</a>
0x0_3110	I2C2DR—I <sup>2</sup> C2 data register	R/W	0x00	<a href="#">17.3.1.5/17-9</a>



Table 17-3. I<sup>2</sup>C Memory Map (continued)

Address	I <sup>2</sup> C Register	Access	Reset	Section/Page
0x0_3114	I2C2DFSRR—I <sup>2</sup> C2 digital filter sampling rate register	R/W	0x10	17.3.1.6/17-9
0x0_311C– 0x0_31FF	Reserved, should be cleared	—	—	—

### 17.3.1 Register Descriptions

This section describes the I<sup>2</sup>C registers in detail. Note that reserved bits should always be written with the value they return when read. That is, the register should be programmed by reading the value, modifying appropriate fields, and writing back the value. The return value of the reserved fields should not be assumed, even though the reserved fields return zero. This does not apply to the I<sup>2</sup>C<sub>n</sub> data register (I2C<sub>n</sub>DR).

#### 17.3.1.1 I<sup>2</sup>C<sub>n</sub> Address Register (I2C<sub>n</sub>ADR)

Figure 17-2 shows the I2C<sub>n</sub>ADR register, which contains the address to which the I<sup>2</sup>C interface responds when addressed as a slave. Note that this is not the address that is sent on the bus during the address-calling cycle when the I<sup>2</sup>C module is in master mode.

Figure 17-2. I<sup>2</sup>C<sub>n</sub> Address Register (I2C<sub>n</sub>ADR)


Table 17-4 describes the bit settings of I2C<sub>n</sub>ADR.

Table 17-4. I2C<sub>n</sub>ADR Field Descriptions

Bits	Name	Description
0–6	ADDR	Slave address. Contains the specific slave address that is used by the I <sup>2</sup> C interface. Note that the default mode of the I <sup>2</sup> C interface is slave mode for an address match. Note that an address match is one of the conditions that can cause I2C <sub>n</sub> SR[MIF] to be set, signaling an interrupt pending condition.
7	—	Reserved, should be cleared

### 17.3.1.2 I<sup>2</sup>C<sub>n</sub> Frequency Divider Register (I2CnFDR)

Figure 17-3 shows the bits of the I<sup>2</sup>C<sub>n</sub> frequency divider register.



**Figure 17-3. I<sup>2</sup>C<sub>n</sub> Frequency Divider Register (I2CnFDR)**


Table 17-5 describes the bit settings of I2CnFDR. It also maps I2CnFDR[FDR] to the clock divider values. Although it describes the ratio between the I<sup>2</sup>C controller internal clock and SCL, the default ratio of I<sup>2</sup>C controller clock and CSB is 1:1 (I<sup>2</sup>C controller clock frequency is three times slower than CSB clock frequency). This ratio is set in SCCR[ENCCM]. Clock ratios of I<sup>2</sup>C1 are controllable but clock ratio for I<sup>2</sup>C2 is not and it is always 1:1 with CSB. Consider this factor when selecting an FDR value.

**Table 17-5. I2Cn FDR Field Descriptions**

Bits	Name	Description																																																																																																																																										
0-1	—	Reserved, should be cleared																																																																																																																																										
2-7	FDR	<p>Frequency divider ratio. Used to prescale the clock for bit-rate selection. The serial bit clock frequency of SCL<sub>n</sub> is equal to the I<sup>2</sup>C<sub>n</sub> controller clock divided by the divider. The serial bit clock frequency divider selections are described as follows:</p> <table border="1" style="width: 100%; border-collapse: collapse;"> <thead> <tr> <th>FDR</th> <th>Divider (Decimal)</th> <th>FDR</th> <th>Divider (Decimal)</th> <th>FDR</th> <th>Divider (Decimal)</th> </tr> </thead> <tbody> <tr><td>0x00</td><td>384</td><td>0x16</td><td>12288</td><td>0x2B</td><td>1024</td></tr> <tr><td>0x01</td><td>416</td><td>0x17</td><td>15360</td><td>0x2C</td><td>1280</td></tr> <tr><td>0x02</td><td>480</td><td>0x18</td><td>18432</td><td>0x2D</td><td>1536</td></tr> <tr><td>0x03</td><td>576</td><td>0x19</td><td>20480</td><td>0x2E</td><td>1792</td></tr> <tr><td>0x04</td><td>640</td><td>0x1A</td><td>24576</td><td>0x2F</td><td>2048</td></tr> <tr><td>0x05</td><td>704</td><td>0x1B</td><td>30720</td><td>0x30</td><td>2560</td></tr> <tr><td>0x06</td><td>832</td><td>0x1C</td><td>36864</td><td>0x31</td><td>3072</td></tr> <tr><td>0x07</td><td>1024</td><td>0x1D</td><td>40960</td><td>0x32</td><td>3584</td></tr> <tr><td>0x08</td><td>1152</td><td>0x1E</td><td>49152</td><td>0x33</td><td>4096</td></tr> <tr><td>0x09</td><td>1280</td><td>0x1F</td><td>61440</td><td>0x34</td><td>5120</td></tr> <tr><td>0x0A</td><td>1536</td><td>0x20</td><td>256</td><td>0x35</td><td>6144</td></tr> <tr><td>0x0B</td><td>1920</td><td>0x21</td><td>288</td><td>0x36</td><td>7168</td></tr> <tr><td>0x0C</td><td>2304</td><td>0x22</td><td>320</td><td>0x37</td><td>8192</td></tr> <tr><td>0x0D</td><td>2560</td><td>0x23</td><td>352</td><td>0x38</td><td>10240</td></tr> <tr><td>0x0E</td><td>3072</td><td>0x24</td><td>384</td><td>0x39</td><td>12288</td></tr> <tr><td>0x0F</td><td>3840</td><td>0x25</td><td>448</td><td>0x3A</td><td>14336</td></tr> <tr><td>0x10</td><td>4608</td><td>0x26</td><td>512</td><td>0x3B</td><td>16384</td></tr> <tr><td>0x11</td><td>5120</td><td>0x27</td><td>576</td><td>0x3C</td><td>20480</td></tr> <tr><td>0x12</td><td>6144</td><td>0x28</td><td>640</td><td>0x3D</td><td>24576</td></tr> <tr><td>0x13</td><td>7680</td><td>0x29</td><td>768</td><td>0x3E</td><td>28672</td></tr> <tr><td>0x14</td><td>9216</td><td>0x2A</td><td>896</td><td>0x3F</td><td>32768</td></tr> <tr><td>0x15</td><td>10240</td><td></td><td></td><td></td><td></td></tr> </tbody> </table> <p><b>Note:</b> The value's shown in the table are applicable only for the default value of DFSRR. Refer to AN2919.</p>	FDR	Divider (Decimal)	FDR	Divider (Decimal)	FDR	Divider (Decimal)	0x00	384	0x16	12288	0x2B	1024	0x01	416	0x17	15360	0x2C	1280	0x02	480	0x18	18432	0x2D	1536	0x03	576	0x19	20480	0x2E	1792	0x04	640	0x1A	24576	0x2F	2048	0x05	704	0x1B	30720	0x30	2560	0x06	832	0x1C	36864	0x31	3072	0x07	1024	0x1D	40960	0x32	3584	0x08	1152	0x1E	49152	0x33	4096	0x09	1280	0x1F	61440	0x34	5120	0x0A	1536	0x20	256	0x35	6144	0x0B	1920	0x21	288	0x36	7168	0x0C	2304	0x22	320	0x37	8192	0x0D	2560	0x23	352	0x38	10240	0x0E	3072	0x24	384	0x39	12288	0x0F	3840	0x25	448	0x3A	14336	0x10	4608	0x26	512	0x3B	16384	0x11	5120	0x27	576	0x3C	20480	0x12	6144	0x28	640	0x3D	24576	0x13	7680	0x29	768	0x3E	28672	0x14	9216	0x2A	896	0x3F	32768	0x15	10240				
FDR	Divider (Decimal)	FDR	Divider (Decimal)	FDR	Divider (Decimal)																																																																																																																																							
0x00	384	0x16	12288	0x2B	1024																																																																																																																																							
0x01	416	0x17	15360	0x2C	1280																																																																																																																																							
0x02	480	0x18	18432	0x2D	1536																																																																																																																																							
0x03	576	0x19	20480	0x2E	1792																																																																																																																																							
0x04	640	0x1A	24576	0x2F	2048																																																																																																																																							
0x05	704	0x1B	30720	0x30	2560																																																																																																																																							
0x06	832	0x1C	36864	0x31	3072																																																																																																																																							
0x07	1024	0x1D	40960	0x32	3584																																																																																																																																							
0x08	1152	0x1E	49152	0x33	4096																																																																																																																																							
0x09	1280	0x1F	61440	0x34	5120																																																																																																																																							
0x0A	1536	0x20	256	0x35	6144																																																																																																																																							
0x0B	1920	0x21	288	0x36	7168																																																																																																																																							
0x0C	2304	0x22	320	0x37	8192																																																																																																																																							
0x0D	2560	0x23	352	0x38	10240																																																																																																																																							
0x0E	3072	0x24	384	0x39	12288																																																																																																																																							
0x0F	3840	0x25	448	0x3A	14336																																																																																																																																							
0x10	4608	0x26	512	0x3B	16384																																																																																																																																							
0x11	5120	0x27	576	0x3C	20480																																																																																																																																							
0x12	6144	0x28	640	0x3D	24576																																																																																																																																							
0x13	7680	0x29	768	0x3E	28672																																																																																																																																							
0x14	9216	0x2A	896	0x3F	32768																																																																																																																																							
0x15	10240																																																																																																																																											

### 17.3.1.3 I<sup>2</sup>C<sub>n</sub> Control Register (I2C<sub>n</sub>CR)

Figure 17-4 shows the I<sup>2</sup>C<sub>n</sub> control register.



**Figure 17-4. I<sup>2</sup>C<sub>n</sub> Control Register (I2C<sub>n</sub>CR)**

Table 17-6 describes the I2C<sub>n</sub>CR bit settings.

**Table 17-6. I2C<sub>n</sub>CR Field Descriptions**

Bits	Name	Description
0	MEN	Module enable. Controls the software reset of the I <sup>2</sup> C module. 0 The module is reset and disabled. The interface is held in reset, but the registers can still be accessed. 1 The I <sup>2</sup> C module is enabled. MEN must be set before any other control register bits have any effect. All I <sup>2</sup> C registers for slave receive or master START can be initialized before setting this bit.
1	MIEN	Module interrupt enable 0 Interrupts from the I <sup>2</sup> C module are disabled. This does not clear any pending interrupt conditions. 1 Interrupts from the I <sup>2</sup> C module are enabled. An interrupt occurs provided I2C <sub>n</sub> SR[MIF] is also set.
2	MSTA	Master/slave mode START 0 On a transition to zero, a STOP condition is generated and the mode changes from master to slave. Cleared without generating a STOP condition when the master loses arbitration. 1 When MSTTA changes from zero to one, a START condition is generated on the bus and master mode is selected.
3	MTX	Transmit/receive mode select. Selects the direction of the master and slave transfers. When configured as a slave, this bit should be set by software according to I2C <sub>n</sub> SR[SRW]. In master mode, the bit should be set according to the type of transfer required. Therefore, for address cycles, this bit will always high. MTX is cleared when the master loses arbitration. 0 Receive mode 1 Transmit mode
4	TXAK	Transfer acknowledge. Specifies the value driven onto the SDA <sub>n</sub> line during acknowledge cycles for both master and slave receivers. The value of this bit applies only when the I <sup>2</sup> C module is configured as a receiver, not a transmitter. It also does not apply to address cycles; when the device is addressed as a slave, an acknowledge is always sent. 0 An acknowledge signal (low value on SDA <sub>n</sub> ) is sent out to the bus at the 9th clock bit after receiving one byte of data. 1 No acknowledge signal response (high value on SDA <sub>n</sub> ) is sent.
5	RSTA	Repeated START. Note that this bit is not readable, which means if a read is performed to RSTA, a zero value is returned. 0 No START condition is generated 1 Setting this bit always generates a repeated START condition on the bus, provides the device with the current bus master. Attempting a repeated START at the wrong time (or if the bus is owned by another master), results in loss of arbitration.
6	—	Reserved, should be cleared
7	BCST	Broadcast 0 Disables the broadcast accept capability 1 Enables the I <sup>2</sup> C to accept broadcast messages at address zero

### 17.3.1.4 I<sup>2</sup>C<sub>n</sub> Status Register (I2C<sub>n</sub>SR)

I2CSR is shown in Figure 17-5.




Figure 17-5. I<sup>2</sup>C<sub>n</sub> Status Register (I2C<sub>n</sub>SR)

Table 17-7 describes the bit settings of the I2C<sub>n</sub>SR.

Table 17-7. I2C<sub>n</sub>SR Field Descriptions

Bits	Name	Description
0	MCF	Data transfer. When one byte of data is transferred, the bit is cleared. It is set by the falling edge of the 9th clock of a byte transfer. 0 Byte transfer in progress. MCF is cleared under the following conditions: <ul style="list-style-type: none"> <li>• When I2C<sub>n</sub>DR is read in receive mode or when I2C<sub>n</sub>DR is written in transmit mode.</li> <li>• After a start sequence is recognized by the I<sup>2</sup>C controller in slave mode.</li> </ul> 1 Byte transfer is completed
1	MAAS	Addressed as a slave. When the value in I2C <sub>n</sub> ADR matches the calling address or when the calling address is the broadcast address and broadcast mode is enabled (I2C <sub>n</sub> CR[BCST] is set), this bit is set. The processor is interrupted if I2C <sub>n</sub> CR[MIE] is set. Next, the processor must check the SRW bit and set I2C <sub>n</sub> CR[MTX] accordingly. Writing to the I2C <sub>n</sub> CR automatically clears this bit. 0 Not addressed as a slave 1 Addressed as a slave
2	MBB	Bus busy. Indicates the status of the bus. When a START condition is detected, MBB is set. If a STOP condition is detected, it is cleared. 0 I <sup>2</sup> C bus is idle 1 I <sup>2</sup> C bus is busy
3	MAL	Arbitration lost. Automatically set when the arbitration procedure is lost. Note that the device does not automatically retry a failed transfer attempt. 0 Arbitration is not lost. Can only be cleared by software 1 Arbitration is lost
4	BCSTM	Broadcast match. Writing to the I2C <sub>n</sub> CR automatically clears this bit. 0 There has not been a broadcast match. 1 The calling address matches with the broadcast address and broadcast mode is enabled. This is also set if this I <sup>2</sup> C drives an address of all 0s.
5	SRW	Slave read/write. When MAAS is set, SRW indicates the value of the R/W command bit of the calling address, which is sent from the master. 0 Slave receive, master writing to slave 1 Slave transmit, master reading from slave. This bit is valid only when both of the following conditions are true: <ul style="list-style-type: none"> <li>• A complete transfer occurred and no other transfers have been initiated.</li> <li>• The I<sup>2</sup>C interface is configured as a slave and has an address match.</li> </ul> By checking SRW, the processor can select slave transmit/receive mode according to the command of the master.

Table 17-7. I2CnSR Field Descriptions (continued)

Bits	Name	Description
6	MIF	Module interrupt. The MIF bit is set when an interrupt is pending, causing a processor interrupt request (provided I2CnCR[MIEN] is set). 0 No interrupt is pending. Can be cleared only by software. 1 Interrupt is pending. MIF is set when one of the following events occurs: <ul style="list-style-type: none"> <li>• One byte of data is transferred (set at the falling edge of the 9th clock).</li> <li>• The value in I2CnADR matches with the calling address in slave-receive mode.</li> <li>• Arbitration is lost.</li> </ul>
7	RXAK	Received acknowledge. The value of SDA <sub>n</sub> during the reception of acknowledge bit of a bus cycle. If the received acknowledge bit (RXAK) is low, it indicates that an acknowledge signal has been received after the completion of eight bits of data transmission on the bus. If RXAK is high, it means no acknowledge signal has been detected at the 9th clock. 0 Acknowledge received 1 No acknowledge received

### 17.3.1.5 I<sup>2</sup>Cn Data Register (I2CnDR)

The I2Cn data register is shown in Figure 17-6.


Figure 17-6. I<sup>2</sup>Cn Data Register (I2CnDR)

Table 17-8 shows the bit descriptions for I2CnDR.

Table 17-8. I2CnDR Field Descriptions

Bits	Name	Description
0–7	DATA	Transmission starts when an address and the R/W bit are written to the data register and the I <sup>2</sup> C interface performs as the master. A data transfer is initiated when data is written to the I2CnDR. The most-significant bit is sent first in both cases. In master receive mode, reading the data register allows the read to occur, but also allows the I <sup>2</sup> C module to receive the next byte of data on the I <sup>2</sup> C interface. In slave mode, the same function is available after it is addressed. Note that in both master receive and slave receive modes, the very first read is always a dummy read.

### 17.3.1.6 Digital Filter Sampling Rate Register (I2CnDFSRR)

I2CnDFSRR is shown in Figure 17-7.

Figure 17-7. I<sup>2</sup>Cn Digital Filter Sampling Rate Register (I2CnDFSRR)

Table 17-9 shows the I2CnDFSRR field descriptions.

**Table 17-9. I2CnDFSRR Field Descriptions**

Bits	Name	Description
0–1	—	Reserved, should be cleared
2–7	DFSRR	Digital filter sampling rate. To assist in filtering out signal noise, the sample rate is programmed. DFSRR is used to prescale the frequency at which the digital filter takes samples from the I <sup>2</sup> C bus. The resulting sampling rate is calculated by dividing the platform frequency by the non-zero value of DFSRR. If I2CnDFSRR is cleared, the I <sup>2</sup> C bus sample points default to the reset divisor 0x10.

## 17.4 Functional Description

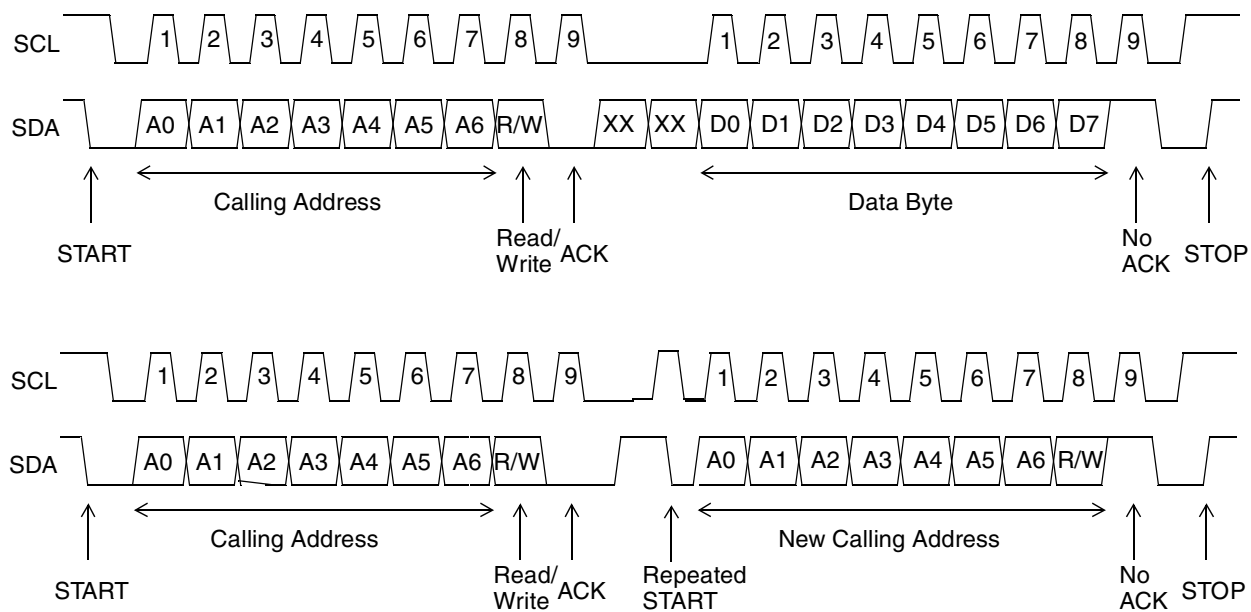
The I<sup>2</sup>C unit always performs as a slave receiver as a default, unless explicitly programmed to be a master or slave transmitter. If boot sequencer mode is selected, the I<sup>2</sup>C interface performs as a slave receiver after the boot sequence has completed.

### 17.4.1 Transaction Protocol

A standard I<sup>2</sup>C transfer consists of the following:

- START condition
- Slave target address transmission
- Data transfer
- STOP condition

Figure 17-8 shows the interaction of these four parts with the calling address, data byte, and new calling address components of the I<sup>2</sup>C protocol. The details of the protocol are described in the following subsections.



**Figure 17-8. I<sup>2</sup>C Interface Transaction Protocol**

### 17.4.1.1 START Condition

When the I<sup>2</sup>C bus is not engaged (both SDA<sub>n</sub> and SCL<sub>n</sub> lines are at logic high), a master can initiate a transfer by sending a START condition. As shown in Figure 17-8, a START condition is defined as a high-to-low transition of SDA<sub>n</sub> while SCL<sub>n</sub> is high. This condition denotes the beginning of a new data transfer. Each data transfer can contain several bytes and awakens all slaves. The START condition is initiated by a software write that sets I2CnCR[MSTA].

### 17.4.1.2 Slave Address Transmission

The first byte of data transferred by the master immediately after the START condition is the slave address. This is a seven-bit calling address followed by a R/ $\overline{W}$  bit, which indicates the direction of the data transferred to the slave. Each slave in the system has a unique address. When the I<sup>2</sup>C module is operating as a master, it must not transmit an address that is the same as its slave address. An I<sup>2</sup>C device cannot be master and slave at the same time.

Only the slave with a calling address that matches the one transmitted by the master responds by returning an acknowledge bit (negating the SDA<sub>n</sub> signal at the 9th clock) as shown in Figure 17-8. If no slave acknowledges the address, the master should generate a STOP condition or a repeated START condition.

When slave addressing is successful (and SCL<sub>n</sub> returns to zero), the data transfer can proceed on a byte-to-byte basis in the direction specified by the R/ $\overline{W}$  bit sent by the calling master.

The I<sup>2</sup>C module responds to a general call (broadcast) command when I2CnCR[BCST] is set. A broadcast address is always zero; however the I<sup>2</sup>C module does not check the R/W bit. The second byte of the broadcast message is the master address. Because the second byte is automatically acknowledged by hardware, the receiver device software must verify that the broadcast message is intended for itself by reading the second byte of the message. If the master address is for another receiver device and the third byte is a write command, software can ignore the third byte during the broadcast. If the master address is for another receiver device and the third byte is a read command, software must write 0xFF to I2CnDR with I2CnCR[TXAK] = 1, so that it does not interfere with the data written from the addressed device.

Each data byte is 8 bits long. Data bits can be changed only while SCL<sub>n</sub> is low and must be held stable while SCL<sub>n</sub> is high, as shown in Figure 17-8. There is one clock pulse on SCL<sub>n</sub> for each data bit, and the most significant bit (msb) is transmitted first. Each byte of data must be followed by an acknowledge bit, which is signaled from the receiving device by pulling SDA<sub>n</sub> low at the 9th clock. Therefore, one complete data byte transfer takes 9 clock pulses. Several bytes can be transferred during a data transfer session.

If the slave receiver does not acknowledge the master, the SDA<sub>n</sub> line must be left high by the slave. The master can then generate a stop condition to abort the data transfer or a START condition (repeated START) to begin a new calling.

If the master receiver does not acknowledge the slave transmitter after a byte of transmission, the slave interprets that the end-of-data has been reached. Then the slave releases the SDA<sub>n</sub> line for the master to generate a STOP or a START condition.

### 17.4.1.3 Repeated START Condition

Figure 17-8 shows a repeated START condition, which is generated without a STOP condition that can terminate the previous transfer. The master uses this method to communicate with another slave or with the same slave in a different mode (transmit/receive mode) without releasing the bus.

### 17.4.1.4 STOP Condition

The master can terminate the transfer by generating a STOP condition to free the bus. A STOP condition is defined as a low-to-high transition of the SDA<sub>n</sub> signal while SCL<sub>n</sub> is high. For more information, see Figure 17-8. Note that a master can generate a STOP even if the slave has transmitted an acknowledge bit, at which point the slave must release the bus. The STOP condition is initiated by a software write that clears I2CnCR[MSTA].

As described in Section 17.4.1.3, “Repeated START Condition,” the master can generate a START condition followed by a calling address without generating a STOP condition for the previous transfer. This is called a repeated START condition.

### 17.4.1.5 Protocol Implementation Details

The following sections give details of how aspects of the protocol are implemented in the I<sup>2</sup>C module.

#### 17.4.1.5.1 Transaction Monitoring—Implementation Details

The different conditions of the I<sup>2</sup>C data transfers are monitored as follows (see Figure 17-8):

- START conditions are detected when an SDA<sub>n</sub> fall occurs while SCL<sub>n</sub> is high.
- STOP conditions are detected when an SDA<sub>n</sub> rise occurs while SCL<sub>n</sub> is high.
- Data transfers in progress are canceled when a STOP condition is detected or if there is a slave address mismatch. Cancellation of data transactions resets the clock module.
- The bus is detected to be busy upon the detection of a START condition and idle upon the detection of a STOP condition.

#### 17.4.1.5.2 Control Transfer—Implementation Details

The I<sup>2</sup>C module contains logic that controls the output to the serial data (SDA) and serial clock (SCL) lines of the I<sup>2</sup>C. The SCL<sub>n</sub> output is pulled low as determined by the internal clock generated in the clock module. The SDA<sub>n</sub> output can change only at the midpoint of a low cycle of the SCL<sub>n</sub>, unless it is performing a START, STOP, or repeated START condition. Otherwise, the SDA<sub>n</sub> output is held constant.

SDA<sub>n</sub> is negated when one or more of the following conditions are true:

- Master mode
  - Data bit (transmit)
  - ACK bit (receive)
  - START condition
  - STOP condition
  - Repeated START condition



- Slave mode
  - Acknowledging address match
  - Data bit (transmit)
  - ACK bit (receive)

The  $SCL_n$  signal corresponds to the internal  $SCL_n$  signal when one or more of the following conditions are true in either master or slave mode:

- Master mode
  - Bus owner
  - Lost arbitration
  - START condition
  - STOP condition
  - Repeated START condition begin
  - Repeated START condition end
- Slave mode
  - Address cycle
  - Transmit cycle
  - ACK cycle

#### 17.4.1.6 Address Compare—Implementation Details

The address compare block determines whether a slave has been properly addressed, either by its slave address or by the general broadcast address (which addresses all slaves). The following address comparisons are performed:

- Whether a broadcast message has been received, to update  $I2CnSR$
- Whether the module has been addressed as a slave, to update  $I2CnSR$  and to generate an interrupt
- Whether the address transmitted by the current master matches the general broadcast address

#### 17.4.2 Arbitration Procedure

The I<sup>2</sup>C interface is a true multiple-master bus. If two or more masters simultaneously try to control the bus, each master's clock synchronization procedure (including the I<sup>2</sup>C module) determines the bus clock—the low period is equal to the longest clock-low period and the high is equal to the shortest one among the masters. A bus master loses arbitration if it transmits a logic 1 on  $SDAn$  while another master transmits a logic 0. The losing masters immediately switch to slave-receive mode and stop driving the  $SDAn$  line. In this case, the transition from master to slave mode does not generate a STOP condition. Meanwhile, the I<sup>2</sup>C unit sets  $I2CnSR[MAL]$  to indicate the loss of arbitration and, as a slave, services the transaction if it is directed to itself.

If the I<sup>2</sup>C module is enabled in the middle of an ongoing byte transfer, the interface behaves as follows:

- Slave mode—the I<sup>2</sup>C module ignores the current transfer on the bus and starts operating whenever a subsequent START condition is detected.

- Master mode—the I<sup>2</sup>C module cannot tell whether the bus is busy; therefore, if a START condition is initiated, the current bus cycle can be corrupted. This ultimately causes in the current bus master to lose arbitration, after which bus operations return to normal.

### 17.4.2.1 Arbitration Control

The arbitration control block controls the arbitration procedure of the master mode. A loss of arbitration occurs whenever the master detects a 0 on the external SDA<sub>n</sub> line while attempting to drive a 1, tries to generate a START or repeated START at an inappropriate time, or detects an unexpected STOP request on the line.

In master mode, arbitration by the master is lost (and I2CnSR[MAL] is set) under the following conditions:

- SDA<sub>n</sub> samples low when the master drives high during an address or data-transmit cycle (transmit).
- SDA<sub>n</sub> samples low when the master drives high during a data-receive cycle of the acknowledge (ACK) bit (receive).
- A START condition is attempted when the bus is busy.
- A repeated START condition is requested in slave mode.
- A repeated START condition is attempted when the requesting device is not the bus owner
- Unexpected STOP condition detected

Note that the I<sup>2</sup>C module does not automatically retry a failed transfer attempt.

### 17.4.3 Handshaking

The clock synchronization mechanism can be used as a handshake in data transfer. Slave devices can hold SCL<sub>n</sub> low after completion of a 1-byte transfer (9 bits). In such cases, it halts the bus clock and forces the master clock into wait states until the slave releases the SCL<sub>n</sub> line.

### 17.4.4 Clock Control

The clock control block handles requests from the clock signal for transferring and controlling data for multiple tasks.

A 9-cycle data transfer clock is requested for the following conditions:

- Master mode
  - Transmit slave address after START condition
  - Transmit slave address after repeated START condition
  - Transmit data
  - Receive data
- Slave mode
  - Transmit data
  - Receive data
  - Receive slave address after START or repeated START condition

### 17.4.4.1 Clock Synchronization

Due to the wire AND logic on the SCL<sub>n</sub> line, a high-to-low transition on the SCL<sub>n</sub> line affects all devices connected on the bus. The devices begin counting their low period when the master negates the SCL<sub>n</sub> line. After a device has negated SCL<sub>n</sub>, it holds the SCL<sub>n</sub> line low until the clock high state is reached. However, the change of low-to-high in a device clock may not change the state of SCL<sub>n</sub> if another device is still within its low period. Therefore, SCL<sub>n</sub> is held low by the device with the longest low period. Devices with shorter low periods enter a high wait state during this time. When all devices concerned have counted off their low periods, SCL<sub>n</sub> is released and asserted. Then there is no difference between the devices' clocks and the state of SCL<sub>n</sub>, and all the devices begin counting their high periods. The first device to complete its high period negates SCL<sub>n</sub> again.

### 17.4.4.2 Input Synchronization and Digital Filter

The following sections describes synchronization of the input signals and the filtering of SCL<sub>n</sub> and SDA<sub>n</sub> in detail.

#### 17.4.4.2.1 Input Signal Synchronization

The input synchronization block synchronizes the input SCL<sub>n</sub> and SDA<sub>n</sub> signals to the system clock and detects transitions of these signals.

#### 17.4.4.2.2 Filtering of SCL<sub>n</sub> and SDA<sub>n</sub> Lines

The SCL<sub>n</sub> and SDA<sub>n</sub> inputs are filtered to eliminate noise. Three consecutive samples of the SCL<sub>n</sub> and SDA<sub>n</sub> lines are compared to a pre-determined sampling rate. If they are all high, the output of the filter is high. If they are all low, the output is low. If they are any combination of highs and lows, the output is whatever the value of the line was in the previous clock cycle.

The sampling rate is equal to a binary value stored in the frequency register I2CDFSRR. The duration of the sampling cycle is controlled by a down counter. This allows a software write to the I2CDFSRR to control the filtered sampling rate.

### 17.4.4.3 Clock Stretching

Slaves can use the clock synchronization mechanism to slow down the transfer bit rate. After the master has driven SCL<sub>n</sub> low, the slave can drive SCL<sub>n</sub> low for the required period and then release it. If the slave SCL<sub>n</sub> low period is greater than the master SCL<sub>n</sub> low period, the resulting SCL<sub>n</sub> low period is extended.

## 17.4.5 Boot Sequencer Mode

Boot sequencer mode is selected at power-on reset by the BOOTSEQ field of the high-order reset configuration word. If boot sequencer mode is selected, the I<sup>2</sup>C module communicates with one or more EEPROMs through the I<sup>2</sup>C interface. EEPROMs can be programmed to initialize one or more configuration registers. Note that as described in [Section 4.3.2.2.3, "Boot Sequencer Configuration,"](#) the default value for BOOTSEQ is 0b00, which corresponds to the I<sup>2</sup>C boot sequencer being disabled at power-up.

Boot sequencer mode also supports an extension of the standard I<sup>2</sup>C interface that uses more address bits to allow for EEPROM devices that have more than 256 bytes. This extended addressing mode is selectable using a different encoding in the BOOTSEQ field of the high-order reset configuration word (see [Section 4.3.2.2.3, “Boot Sequencer Configuration.”](#)) In this mode, only one EEPROM device can be used and the maximum number of registers is limited by the size of the EEPROM.

If the standard I<sup>2</sup>C interface is used, the I<sup>2</sup>C module addresses the first EEPROM, and reads 256 bytes. Then it issues a repeated start and addresses the next EEPROM address. This sequence continues until the CONT bit is cleared. If the last register is not detected before wrapping back to the first address, an error condition is detected. In other words, if the CONT bit for not cleared on the final 7 bytes, an error condition is detected, causing the I<sup>2</sup>C controller to hang. The I<sup>2</sup>C module continues to read from the EEPROM as long as the continue (CONT) bit is set in the EEPROM. The CONT bit resides in the address/attributes field that is transferred from the EEPROM, as described in [Section 17.4.5.2, “EEPROM Calling Address.”](#) There should be no other I<sup>2</sup>C traffic when the boot sequencer is active.

### 17.4.5.1 Using the Boot Sequencer for Reset Configuration

The reset configuration word can be loaded by using the I<sup>2</sup>C boot sequencer. See [Section 4.3.2.2.3, “Boot Sequencer Configuration.”](#)

Note that this usage does not prevent using the I<sup>2</sup>C boot sequencer to initiate the device in the normal functional mode, after reset state has completed. However, an I<sup>2</sup>C serial EEPROM of extended addressing type must be used and the first two EEPROM data structures must contain dedicated reset information.

### 17.4.5.2 EEPROM Calling Address

The EEPROM calling address is 0b101\_0000. The first EEPROM to be addressed must be programmed to respond to this address, or an error is generated. Any additional EEPROMs are addressed in sequential order.

### 17.4.5.3 EEPROM Data Format

The I<sup>2</sup>C module expects a particular format for data to be programmed in the EEPROM. [Figure 17-9](#) shows an example of the EEPROM contents, including the preamble, data format, and CRC.

0	1	2	3	4	5	6	7	
1	0	1	0	1	0	1	0	Preamble
0	1	0	1	0	1	0	1	
1	0	1	0	1	0	1	0	
ACS	BYTE_EN			1	ADDR[12:13]			First Configuration Preload Command
ADDR[14:21]								
ADDR[22:29]								
DATA[0:7]								
DATA[8:15]								
DATA[16:23]								
DATA[24:31]								
ACS	BYTE_EN			1	ADDR[12:13]			Second Configuration Preload Command
ADDR[14:21]								
ADDR[22:29]								
DATA[0:7]								
DATA[8:15]								
DATA[16:23]								
DATA[24:31]								
.....								
ACS	BYTE_EN			1	ADDR[12:13]			Last Configuration Preload Command
ADDR[14:21]								
ADDR[22:29]								
DATA[0:7]								
DATA[8:15]								
DATA[16:23]								
DATA[24:31]								
0	0	0	0	0	0	0	0	End Command
0	0	0	0	0	0	0	0	
0	0	0	0	0	0	0	0	
CRC[0:7]								
CRC[8:15]								
CRC[16:23]								
CRC[24:31]								

Figure 17-9. EEPROM Contents

- A preamble should be the first 3 bytes programmed into the EEPROM. It should have a value of 0xAA55AA. The I<sup>2</sup>C checks to ensure that this preamble is correctly detected before proceeding.
- Following the preamble, there should be a series of configuration registers (known as register preloads). Each configuration register should be programmed according to a particular format, as shown in Figure 17-10.

0	1	2	3	4	5	6	7
ACS	BYTE_EN			CONT	ADDR[12:13]		
ADDR[14:21]							
ADDR[12:29]							
DATA[0:7]							
DATA[8:15]							
DATA[16:23]							
DATA[24:31]							

**Figure 17-10. EEPROM Data Format for One Register Preload Command**

- The first byte holds alternate configuration space (ACS), byte enables, and continue (CONT) attributes.
- The 2 least-significant bits of the address are derived from the byte enables. address offset. Therefore, the address offset programmed into the EEPROM preload should be a word offset.
- The most significant 16 bits (assuming 36-bit addressing) of the address are prepended from either IMMRRBAR or alternate configuration space.
- After the first 3 bytes, 4 bytes of data should hold the desired value of the configuration register, regardless of the size of transaction.

Byte enables should be asserted for any byte that will be written, and they should be asserted contiguously, creating a 1, 2, or 4 byte write to a register. The boot sequencer assumes that a big-endian address is stored in the EEPROM. In addition, byte enable bit 0 (bit 1 of the byte) corresponds to the most-significant byte of data (data[0:7]), and byte enable bit 3 (bit 4 of the byte) corresponds to the least-significant byte of data (data[24:31]).

By asserting ACS, an alternate configuration space address is prepended to the write request from the boot sequencer according to the value in the ALTCBAR register. This will allow for external memories to be configured. Otherwise, IMMRBAR is prepended to the EEPROM address.

If the CONT bit is cleared, the first 3 bytes, including ACS, the byte enables, and the address, should be cleared 0. Also, the data contains the final CRC. A CRC-32 algorithm is used to check the integrity of the data. The following polynomial is used:

$$1 + x^1 + x^2 + x^4 + x^5 + x^7 + x^8 + x^{10} + x^{11} + x^{12} + x^{16} + x^{22} + x^{23} + x^{26} + x^{32}$$

The CRC should cover all bytes stored in the EEPROM before the CRC. This includes the preamble, all register preloads, and the first 3 bytes of the last 7-byte preload (which should be all zeros).

#### 17.4.5.4 Boot Sequencer Done Indication

Dedicated hardware is not provided to indicate whether the boot sequencer operation completed successfully. It is recommended to use one of the GPIO signals for that purpose. To do this, the last register preload programmed into the EEPROM should contain the address of the appropriate GPIO register and data that causes the setting of the required GPIO signal. The GPIO signal may be used for an external device or for debug purposes.

### 17.5 Initialization/Application Information

This section describes some programming guidelines recommended for the I<sup>2</sup>C interface. [Figure 17-11](#) is a recommended flowchart for I<sup>2</sup>C interrupt service routines.

A **sync** assembly instruction must be executed after each I<sup>2</sup>C register read/write access to guarantee that register accesses occur in order.

The I<sup>2</sup>C controller does not guarantee its recovery from all illegal I<sup>2</sup>C bus activity. In addition, a malfunctioning device may hold the bus captive. A good programming practice is for software to rely on a watchdog timer to help recover from I<sup>2</sup>C bus hangs. The recovery routine should also handle the case when the illegal I<sup>2</sup>C bus behavior causes the status bits returned after an interrupt to be inconsistent with what was expected.

#### 17.5.1 Interrupt Service Routine Flowchart

[Figure 17-11](#) shows an example algorithm for an I<sup>2</sup>C interrupt service routine. Deviation from the flowchart may result in unpredictable I<sup>2</sup>C bus behavior. However, in the slave receive mode (not shown), the interrupt service routine may need to set I2CnCR[TXAK] when the next-to-last byte is to be accepted. It is recommended that a **sync** instruction follow each I<sup>2</sup>C register read or write to guarantee that register accesses occur in order.




Figure 17-11. Example I<sup>2</sup>C Interrupt Service Routine Flowchart



## 17.5.2 Initialization Sequence

A hard reset initializes all of the I<sup>2</sup>C registers to their default states. The following initialization sequence initializes the I<sup>2</sup>C unit:

1. All I<sup>2</sup>C registers must be located in a cache-inhibited page.
2. Update I2CnFDR[FDR] and select the required division ratio to obtain the SCLn frequency from the CSB (platform) clock.
3. Update I2CnADR to define the slave address for this device.
4. Modify I2CnCR to select master/slave mode, transmit/receive mode, and interrupt-enable or disable.
5. Set the I2CnCR[MEN] to enable the I<sup>2</sup>C interface.

## 17.5.3 Generation of START

After initialization, the following sequence can be used to generate START:

1. If the device is connected to a multimaster I<sup>2</sup>C system, check whether the serial bus is free (I2CnSR[MBB] = 0) before switching to master mode.
2. Select master mode (set I2CnCR[MSTA]) to transmit serial data and select transmit mode (set I2CnCR[MTX]) for the address cycle.
3. Write the slave address being called into I2CnDR. The data written to I2CnDR[0–6] comprises the slave calling address. I2CnCR[MTX] indicates the direction of transfer (transmit/receive) required from the slave.

The scenario above assumes that the I<sup>2</sup>C interrupt bit (I2CnSR[MIF]) is cleared. If MIF is set at any time, an I<sup>2</sup>C interrupt is generated (provided interrupt reporting is enabled with I2CnCR[MIEN] = 1).

## 17.5.4 Post-Transfer Software Response

Transmission or reception of a byte automatically sets the data transferring bit (I2CnSR[MCF]), which indicates that one byte has been transferred. The I<sup>2</sup>C interrupt bit (I2CnSR[MIF]) is also set and an interrupt is generated to the processor if the interrupt function is enabled during the initialization sequence (I2CnCR[MIEN] is set). In the interrupt handler, software must take the following steps:

1. Clear I2CnSR[MIF]
2. Read the I2CnDR in receive mode or write to I2CnDR in transmit mode. Note that this causes I2CnSR[MCF] to be cleared, as shown in [Figure 17-11](#).
3. When an interrupt occurs at the end of the address cycle, the master remains in transmit mode. If master receive mode is required, I2CnCR[MTX] must be toggled at this stage (see [Figure 17-11](#)).

If the interrupt function is disabled, software can service the I2CnDR in the main program by monitoring I2CnSR[MIF]. In this case, I2CnSR[MIF] must be polled rather than I2CnSR[MCF] because MCF behaves differently when arbitration is lost. Note that interrupt or other bus conditions may be detected before the I<sup>2</sup>C signals have time to settle. Thus, when polling I2CnSR[MIF] (or any other I2CnSR bits), software delays may be needed to give the I<sup>2</sup>C signals sufficient time to settle.

During slave-mode address cycles (I2CnSR[MAAS] is set), I2CnSR[SRW] should be read to determine the direction of the subsequent transfer and I2CnCR[MTX] should be programmed accordingly. For slave-mode data cycles (MAAS is cleared), I2CnSR[SRW] is not valid and I2CnCR[MTX] must be read to determine the direction of the current transfer (see Figure 17-11).

### 17.5.5 Generation of STOP

A data transfer ends with a STOP condition generated by the master device. A master transmitter can generate a STOP condition after all the data has been transmitted.

If a master receiver wants to terminate a data transfer, it must inform the slave transmitter by not acknowledging the last byte of data (by setting the transmit acknowledge bit (I2CnCR[TXAK])) before reading the next-to-last byte of data. At this time, the next-to-last byte of data has been transferred on the I<sup>2</sup>C interface, so the last byte does not receive the data acknowledge (because I2CnCR[TXAK] is set). Before the interrupt service routine reads the last byte of data, a STOP condition must first be generated.

### 17.5.6 Generation of Repeated START

At the end of a data transfer, if the master still wants to communicate on the bus, it can generate another START condition followed by another slave address without first generating a STOP condition. This is accomplished by setting I2CnCR[RSTA].

### 17.5.7 Generation of SCLn When SDA<sub>n</sub> is Negated

It is sometimes necessary to force the I<sup>2</sup>C module to become the I<sup>2</sup>C bus master out of reset and drive SCL<sub>n</sub> (even though SDA<sub>n</sub> may already be driven, which indicates that the bus is busy). This can occur when a system reset does not cause all I<sup>2</sup>C devices to be reset. Thus, SDA<sub>n</sub> can be negated low by another I<sup>2</sup>C device while this I<sup>2</sup>C module is coming out of reset and will stay low indefinitely. The following procedure can be used to force this I<sup>2</sup>C module to generate SCL<sub>n</sub> so that the device driving SDA<sub>n</sub> can finish its transaction:

1. Disable the I<sup>2</sup>C module and set the master bit by setting I2CnCR to 0x20.
2. Enable the I<sup>2</sup>C module by setting I2CnCR to 0xA0.
3. Read I2CnDR.
4. Return the I<sup>2</sup>C module to slave mode by setting I2CnCR to 0x80.

### 17.5.8 Slave Mode Interrupt Service Routine

In the slave interrupt service routine, the module addressed as a slave should be tested to check if a calling of its own address has been received. If I2CnSR[MAAS] is set, software should set the transmit/receive mode select bit (I2CnCR[MTX]) according to the R $\bar{W}$  command bit (I2CnSR[SRW]). Writing to I2CnCR clears MAAS automatically. MAAS is read as set only in the interrupt handler at the end of that address cycle where an address match occurred; interrupts resulting from subsequent data transfers clear MAAS. A data transfer can then be initiated by writing to I2CnDR for slave transmits or dummy reading from I2CnDR in slave-receive mode. The slave negates SCL<sub>n</sub> between byte transfers. SCL<sub>n</sub> is released when I2CnDR is accessed in the required mode.

### 17.5.8.1 Slave Transmitter and Received Acknowledge

In the slave transmitter routine, the received acknowledge bit (I2CnSR[RXAK]) must be tested before sending the next byte of data. The master signals an end-of-data by not acknowledging the data transfer from the slave. When no acknowledge is received (I2CnSR[RXAK] is set), the slave transmitter interrupt routine must clear I2CnCR[MTX] to switch the slave from transmitter to receiver mode. A dummy read of I2CnDR then releases SCL<sub>n</sub> so that the master can generate a STOP condition. See [Figure 17-11](#).

### 17.5.8.2 Loss of Arbitration and Forcing of Slave Mode

When a master loses arbitration the following conditions all occur:

- I2CnSR[MAL] is set
- I2CnCR[MSTA] is cleared (changing the master to slave mode)
- An interrupt occurs (if enabled) at the falling edge of the 9th clock of this transfer

Thus, the slave interrupt service routine should first test I2CnSR[MAL] and software should clear it if it is set. See [Section 17.4.2.1, “Arbitration Control.”](#)



# Chapter 18

## DUART

This chapter describes the two (dual) universal asynchronous receiver/transmitters (UARTs) of the device. It describes the functional operation, the DUART initialization sequence, and the programming details for the DUART registers and features.

### 18.1 Overview

The DUART consists of two (dual) universal asynchronous receiver/transmitters (UARTs). The UARTs act independently; all references to UART refer to one of these receiver/transmitters. Each UART is clocked by the system clock. The DUART programming model is compatible with the PC16552D.

The UART interface is point-to-point, meaning that only two UART devices are attached to the connecting signals. As shown in [Figure 18-1](#), each UART module consists of the following:

- Receive and transmit buffers
- Clear to send ( $\overline{\text{CTS}}$ ) input port and request to send ( $\overline{\text{RTS}}$ ) output port for data-flow control.
- 16-bit counter for baud rate generation
- Interrupt control logic




Figure 18-1. UART Block Diagram

### 18.1.1 Features

The DUART includes these features:

- Full-duplex operation
- Programming model compatible with the original PC16450 UART and the PC16550D (an improved version of the PC16450 that also operates in FIFO mode)
- PC16450 register reset values
- FIFO mode for both transmitter and receiver, providing 16-byte FIFOs
- Serial data encapsulation and decapsulation with standard asynchronous communication bits (START, STOP, and parity)
- Maskable transmit, receive, line status, and MODEM status interrupts
- Software-programmable baud generators that divide the system clock by 1 to  $(2^{16}-1)$  and generate a 16x clock for the transmitter and receiver engines
- Clear-to-send ( $\overline{\text{CTS}}$ ) and ready-to-send ( $\overline{\text{RTS}}$ ) MODEM control functions
- Software-selectable serial interface data format (data length, parity, 1/1.5/2 STOP bit, baud rate)
- Line and MODEM status registers
- Line-break detection and generation
- Internal diagnostic support, local loopback, and break functions
- Prioritized interrupt reporting
- Overrun, parity, and framing error detection

## 18.1.2 Modes of Operation

The communication channel provides a full-duplex asynchronous receiver and transmitter using an operating frequency derived from the system clock.

The transmitter accepts parallel data from a write to the transmitter holding register (UTHR). In FIFO mode, the data is placed directly into an internal transmitter shift register of the transmitter FIFO. The transmitter converts the data to a serial bit stream, inserting the appropriate START, STOP, and optional parity bits. Finally, it outputs a composite serial data stream on the channel transmitter serial data output signal (SOUT). The transmitter status may be polled or interrupt driven.

The receiver accepts serial data bits on the channel receiver serial data input signal (SIN), converts it to parallel format, checks for a START bit, parity (if any), STOP bits, and transfers the assembled character (with START, STOP, parity bits removed) from the receiver buffer (or FIFO) in response to a read of the UART's receiver buffer register (URBR). The receiver status may be polled or interrupt driven.

## 18.2 External Signal Descriptions

This section contains a signal overview and detailed signal descriptions.

### 18.2.1 Signal Overview

Table 18-1 summarizes the DUART signals. Note that although the actual device signal names are prepended with the 'UART\_' prefix as shown in the table, the functional (abbreviated) signal names are often used throughout this chapter.

Table 18-1. DUART Signal Overview

Signal Name	I/O	Pins	Reset Value	State Meaning
UART_SIN[1:2]	I	2	1	Serial in data UART1 and UART2
UART_SOUT[1:2]	O	2	1	Serial out data UART1 and UART2
$\overline{\text{UART\_CTS}}[1:2]$	I	2	1	Clear to send UART1 and UART2
$\overline{\text{UART\_RTS}}[1:2]$	O	2	1	Request to send UART1 and UART2

### 18.2.2 Detailed Signal Descriptions

The DUART signals are described in detail in Table 18-2.

Table 18-2. DUART Signals—Detailed Signal Descriptions

Signal	I/O	Description
UART_SIN[1:2]/DSP_UART_SIN	I	Serial data in. Data is received on the receivers of UART1, UART2, or DSP_UART through its respective serial data input signal, with the least significant bit received first.
		<b>State Meaning</b> Asserted/Negated—Represents the data being received on the UART interface.
		<b>Timing</b> Assertion/Negation—An internal logic sample signal, <i>rxcnt</i> , uses the frequency of the baud-rate generator to sample the data on SIN.

Table 18-2. DUART Signals—Detailed Signal Descriptions (continued)

Signal	I/O	Description	
UART_SOUT[1:2]/ DSP_UART_SOUT	O	Serial data out. The serial data output signals for the UART1, UART2, or DSP_UART are set (mark condition) when the transmitter is disabled, idle, or operating in the local loopback mode. Data is shifted out on these signals, with the least significant bit transmitted first.	
		<b>State Meaning</b>	Asserted/Negated—Represents the data transmitted on the respective UART interface.
		<b>Timing</b>	Assertion/Negation—An internal logic sample signal, <i>rxcnt</i> , uses the frequency of the baud-rate generator to update and drive the data on SOUT.
UART_CTS[1:2]	I	Clear to send. Connected to the respective $\overline{\text{RTS}}$ outputs of the other UART devices on the bus. They can be programmed to generate an interrupt on change-of-state of the signal.	
		<b>State Meaning</b>	Asserted/Negated—Represent the clear to send condition for their respective UART.
		<b>Timing</b>	Assertion/Negation—Sampled at the rising edge of every system clock.
UART_RTS[1:2]	O	Request to send. Can be programmed to be automatically negated and asserted by either the receiver or transmitter. When connected to the $\overline{\text{CTS}}$ input of a transmitter, this signal can be used to control serial data flow.	
		<b>State Meaning</b>	Asserted/Negated—Represents the data being transmitted on the respective UART interface.
		<b>Timing</b>	Assertion/Negation—Updated and driven at the rising edge of every system clock.

### 18.3 Memory Map/Register Definition

There are two complete sets of DUART registers (one for UART1 and one for UART2). The two UARTs are identical, except that the registers for UART1 are located at offsets 0x0\_4500 (local), and the registers for UART2 are located at offsets 0x0\_4600 (local). Throughout this chapter, the registers are described by a singular acronym: for example, LCR represents the line control register for either UART1 or UART2.

The registers in each UART interface are used for configuration, control, and status. The divisor latch access bit, ULCR[DLAB], is used to access the divisor latch least- and most-significant bit registers and the alternate function register. Refer to [Section 18.3.1.7, “Line Control Registers \(ULCR1 and ULCR2\),”](#) for more information on ULCR[DLAB].

All DUART registers are one byte wide; reads and writes to these registers must be byte-wide operations. [Table 18-3](#) provides a register summary with references to the section and page that contain detailed information about each register. Undefined byte address spaces within offset 0x4000–0x4FFF are reserved.

Table 18-3. DUART Register Summary

Offset	Register	Access	Reset	Section/Page
0x0_4500	URBR—ULCR[DLAB] = 0 UART1 receiver buffer register	R	0x00	<a href="#">18.3.1.1/18-5</a>
	UTHR—ULCR[DLAB] = 0 UART1 transmitter holding register	W	0x00	<a href="#">18.3.1.2/18-6</a>
	UDLB—ULCR[DLAB] = 1 UART1 divisor least significant byte register	R/W	0x00	<a href="#">18.3.1.3/18-6</a>



Table 18-3. DUART Register Summary (continued)

Offset	Register	Access	Reset	Section/Page
0x0_4501	UIER—ULCR[DLAB] = 0 UART1 interrupt enable register	R/W	0x00	<a href="#">18.3.1.4/18-8</a>
	UDMB—ULCR[DLAB] = 1 UART1 divisor most significant byte register	R/W	0x00	<a href="#">18.3.1.3/18-6</a>
0x0_4502	UIIR—ULCR[DLAB] = 0 UART1 interrupt ID register	R	0x01	<a href="#">18.3.1.5/18-9</a>
	UFCR—ULCR[DLAB] = 0 UART1 FIFO control register	W	0x00	<a href="#">18.3.1.6/18-10</a>
	UAFR—ULCR[DLAB] = 1 UART1 alternate function register	R/W	0x00	<a href="#">18.3.1.12/18-16</a>
0x0_4503	ULCR—ULCR[DLAB] = x UART1 line control register	R/W	0x00	<a href="#">18.3.1.7/18-11</a>
0x0_4504	UMCR—ULCR[DLAB] = x UART1 MODEM control register	R/W	0x00	<a href="#">18.3.1.8/18-13</a>
0x0_4505	ULSR—ULCR[DLAB] = x UART1 line status register	R	0x60	<a href="#">18.3.1.9/18-14</a>
0x0_4506	UMSR—ULCR[DLAB] = x UART1 MODEM status register	R	0x00	<a href="#">18.3.1.10/18-15</a>
0x0_4507	USCR—ULCR[DLAB] = x UART1 scratch register	R/W	0x00	<a href="#">18.3.1.11/18-16</a>
0x0_4510	UDSR—ULCR[DLAB] = x UART1 DMA status register	R	0x01	<a href="#">18.3.1.13/18-17</a>
0x0_4600	URBR—ULCR[DLAB] = 0 UART2 receiver buffer register	R	0x00	<a href="#">18.3.1.1/18-5</a>
	UTHR—ULCR[DLAB] = 0 UART2 transmitter holding register	W	0x00	<a href="#">18.3.1.2/18-6</a>
	UDLB—ULCR[DLAB] = 1 UART2 divisor least significant byte register	R/W	0x00	<a href="#">18.3.1.3/18-6</a>
0x0_4601	UIER—ULCR[DLAB] = 0 UART2 interrupt enable register	R/W	0x00	<a href="#">18.3.1.4/18-8</a>
	UDMB—ULCR[DLAB] = 1 UART2 divisor most significant byte register	R/W	0x00	<a href="#">18.3.1.3/18-6</a>
0x0_4602	UIIR—ULCR[DLAB] = 0 UART2 interrupt ID register	R	0x01	<a href="#">18.3.1.5/18-9</a>
	UFCR—ULCR[DLAB] = 0 UART2 FIFO control register	W	0x00	<a href="#">18.3.1.6/18-10</a>
	UAFR—ULCR[DLAB] = 1 UART2 alternate function register	R/W	0x00	<a href="#">18.3.1.12/18-16</a>
0x0_4603	ULCR—ULCR[DLAB] = x UART2 line control register	R/W	0x00	<a href="#">18.3.1.7/18-11</a>
0x0_4604	UMCR—ULCR[DLAB] = x UART2 MODEM control register	R/W	0x00	<a href="#">18.3.1.8/18-13</a>
0x0_4605	ULSR—ULCR[DLAB] = x UART2 line status register	R	0x60	<a href="#">18.3.1.9/18-14</a>
0x0_4606	UMSR—ULCR[DLAB] = x UART2 MODEM status register	R	0x00	<a href="#">18.3.1.10/18-15</a>
0x0_4607	USCR—ULCR[DLAB] = x UART2 scratch register	R/W	0x00	<a href="#">18.3.1.11/18-16</a>
0x0_4610	UDSR—ULCR[DLAB] = x UART2 DMA status register	R	0x01	<a href="#">18.3.1.13/18-17</a>


## 18.3.1 Register Descriptions

The following sections describe the UART1 and UART2 registers.

### 18.3.1.1 Receiver Buffer Registers (URBR1 and URBR2)

These registers contain the data received from the transmitter on the UART buses. In FIFO mode, when read, they return the first byte received. For FIFO status information, refer to the UDSR[RXRDY] description.

Except for the case when there is an overrun, URBR returns the data in the order it was received from the transmitter. Refer to the ULSR[OE] description, [Section 18.3.1.9, “Line Status Registers \(ULSR1 and ULSR2\).”](#) [Figure 18-2](#) shows the receiver buffer registers. Note that these registers have same offset as the UTHR<sub>s</sub>.



**Figure 18-2. Receiver Buffer Registers (URBR1 and URBR2)**

[Table 18-4](#) describes URBR.


**Table 18-4. URBR Field Descriptions**

Bits	Name	Description
0–7	DATA	Data received from the transmitter on the UART bus [read only]

### 18.3.1.2 Transmitter Holding Registers (UTHR1 and UTHR2)

A write to these 8-bit registers causes the UART devices to transfer 5 to 8 data bits on the UART bus in the format set up in the ULCR (line control register). In FIFO mode, data written to UTHR is placed into the FIFO. The data written to UTHR is the data sent onto the UART bus, and the first byte written to UTHR is the first byte onto the bus. UDSR[TXRDY] indicates when the FIFO is full. Refer to [Table 18-21](#) and [Table 18-22](#).

[Figure 18-3](#) shows the bits in the UTHR<sub>s</sub>.



**Figure 18-3. Transmitter Holding Registers (UTHR1 and UTHR2)**

[Table 18-5](#) describes the UTHR.

**Table 18-5. UTHR Field Descriptions**


Bits	Name	Description
0–7	DATA	Data that is written to UTHR [Write only]

### 18.3.1.3 Divisor Most and Least Significant Byte Registers (UDMB and UDLB)

UDLB is concatenated with the divisor most significant byte register (UDMB) to create the divisor used to divide the input clock into the DUART. The output frequency of the baud generator is 16 times the baud rate; therefore, the desired baud rate = platform clock frequency ÷ (16 × [UDMB||UDLB]). Equivalently,

$[UDMB||UDLB:0b0000] = \text{platform clock frequency}/\text{desired baud rate}$ . Baud rates that can be generated by specific input clock frequencies are shown in Table 18-8.

Figure 18-4 shows the bits in the UDMBs.




**Figure 18-4. Divisor Most Significant Byte Registers (UDMB1 and UDMB2)**

Table 18-6 describes the UDMB.

**Table 18-6. UDMB Field Descriptions**

Bits	Name	Description
0–7	UDMB	Divisor most significant byte

Figure 18-5 shows the bits in the UDLBs.



**Figure 18-5. Divisor Least Significant Byte Registers (UDLB1 and UDLB2)**

Table 18-7 describes the UDLB.

**Table 18-7. UDLB Field Descriptions**

Bits	Name	Description
0–7	UDLB	Divisor least significant byte. This is concatenated with UDMB.

Table 18-8 shows baud rate for a variety of input clock frequencies.

**Table 18-8. Baud Rate Examples**

Baud Rate (Decimal)	Divisor		Input Clock (System Clock) Frequency (MHz)	Percent Error (Decimal)
	Decimal	Hex		
9,600	866	362	133	0.013
19,200	433	1B1	133	0.013
38,400	216	D8	133	0.218
56,000	148	94	133	0.300
128,000	65	41	133	0.090

Table 18-8. Baud Rate Examples (continued)

Baud Rate (Decimal)	Divisor		Input Clock (System Clock) Frequency (MHz)	Percent Error (Decimal)
	Decimal	Hex		
256,000	32	20	133	1.471
9,600	1087	43F	167	0.022
19,200	544	220	167	0.070
38,400	272	110	167	0.070
56,000	186	BA	167	0.206
128,000	82	52	167	0.557
256,000	41	29	167	0.557

To get the percent error value, the following three steps are taken:

1. The input clock frequency (ICF) is divided by the actual frequency input (AFI) to get the correct divisor value (ICF/AFI, where AFI = baud rate  $\times$  16  $\times$  divisor).
2. The divisor value is subtracted from 1.
3. The result from the step two is multiplied by 100 to calculate the final percent error. The result is calculated in absolute value (no negative numbers).

These steps can be described with the following equation:

$$\text{Percent error value} = (1 - \text{AFI/ICF}) \times 100$$

### 18.3.1.4 Interrupt Enable Registers (UIER1 and UIER2)

The UIER gives the user the ability to mask specific UART interrupts to the programmable interrupt controller (PIC).

Figure 18-6 shows the bits in the UIER.




Figure 18-6. Interrupt Enable Registers (UIER1 and UIER2)

Table 18-9 describes the UIER fields.

**Table 18-9. UIER Field Descriptions**

Bits	Name	Description
0–3	—	Reserved
4	EMSI	Enable MODEM status interrupt 0 Mask interrupts caused by UMSR[DCTS] being set. 1 Enable and assert interrupts when UMSR[CTS] changes state.
5	ERLSI	Enable receiver line status interrupt 0 Mask interrupts when ULSR's overrun, parity error, framing error, or break interrupt bits are set. 1 Enable and assert interrupts when ULSR's overrun, parity error, framing error or break interrupt bits are set.
6	ETHREI	Enable transmitter holding register empty interrupt 0 Mask interrupt when ULSR[THRE] is set. 1 Enable and assert interrupts when ULSR[THRE] is set.
7	ERDAI	Enable received data available interrupt 0 Mask interrupt when new receive data is available or receive data time-out has occurred. 1 Enable and assert interrupts when a new data character is received from the external device and/or a time-out interrupt occurs in FIFO mode.

### 18.3.1.5 Interrupt ID Registers (UIIR1 and UIIR2)

The UIIRs indicate when an interrupt is pending from the corresponding UART and what type of interrupt is active. They also indicate if the FIFOs are enabled.


The DUART prioritizes interrupts into four levels and records these in the corresponding UIIR. The four levels of interrupt conditions in order of priority are as follows:

1. Receiver line status
2. Received data ready/character time-out
3. Transmitter holding register empty
4. MODEM status

See Table 18-11 for more details.

When the UIIR is read, the associated DUART serial channel freezes all interrupts and indicates the highest priority pending interrupt. While this read transaction is occurring, the associated DUART serial channel records new interrupts, but does not change the contents of UIIR until the read access is complete.

Figure 18-7 shows the bits in the UIIR.



**Figure 18-7. Interrupt ID Registers (UIIR1 and UIIR2)**

Table 18-10 describes the fields of the UIIR.

**Table 18-10. UIIR Field Descriptions**

Bits	Name	Description
0–1	FE	FIFOs enabled. Reflects the setting of UFCR[FEN].
2–3	—	Reserved
4	IID3	Interrupt ID bits identify the highest priority interrupt that is pending as indicated in Table 18-11. IID3 is set along with IID2 only when a time out interrupt is pending for FIFO mode.
5–6	IID2–IID1	Interrupt ID bits identify the highest priority pending interrupt as indicated in Table 18-11.
7	IID0	IID0 indicates when an interrupt is pending. 0 The UART has an active interrupt ready to be serviced. 1 No interrupt is pending.

The bits contained in the UIIR registers are described in Table 18-11.

**Table 18-11. UIIR IID Bits Summary**

IID3–IID0	Priority Level	Interrupt Type	Interrupt Description	How To Reset Interrupt
0001	—	—	—	—
0110	Highest	Receiver line status	Overrun error, parity error, framing error, or break interrupt	Reading the line status register
0100	Second	Received data available	Receiver data available or trigger level reached in FIFO mode.	Reading the receiver buffer register or if the number of bytes in the receiver FIFO drops below the trigger level.
1100	Second	Character time-out	No characters were removed from or input to the receiver FIFO during the last four character times and at least one character is in the receiver FIFO.	Reading the receiver buffer register
0010	Third	UTHR empty	Transmitter holding register is empty.	Reading UIIR or writing to UTHR
0000	Fourth	MODEM status	$\overline{\text{CTS}}$ input value changed since last read of UMSR.	Reading UMSR


### 18.3.1.6 FIFO Control Registers (UFCR1 and UFCR2)

UFCR is used to enable and clear the receiver and transmitter FIFOs, set a receiver FIFO trigger level to control the received data available interrupt, and select the type of DMA signaling.

UFCR bits cannot be programmed unless FIFO enable bits are set. When changing from FIFO mode to 16450 mode (non-FIFO mode) and vice versa, data is automatically cleared from the FIFOs.

After all of the bytes in the receiver FIFO are cleared, the receiver internal shift register is not cleared. Similarly, the bytes are cleared in the transmitter FIFO, but the transmitter internal shift register is not cleared. Both TFR and RFR are self clearing.

Figure 18-8 shows the bits in the UFCRs.



**Figure 18-8. FIFO Control Registers (UFCR1 and UFCR2)**

Table 18-12 describes the fields of the UFCRs.

**Table 18-12. UFCR Field Descriptions**

Bits	Name	Description
0–1	RTL	Receiver trigger level. A received data available interrupt occurs when UIER[ERDAI] is set and the number of bytes in the receiver FIFO equals RTL value. 00 1 byte 01 4 bytes 10 8 bytes 11 14 bytes
2–3	—	Reserved
4	DMS	DMA mode select. See <a href="#">Section 18.4.5.2, “DMA Mode Select”</a> 0 UDSR[RXRDY] and UDSR[TXRDY] bits are in mode 0. 1 UDSR[RXRDY] and UDSR[TXRDY] bits are in mode 1 if UFCR[FEN] = 1.
5	TFR	Transmitter FIFO reset 0 No action 1 Clears all bytes in the transmitter FIFO and resets the FIFO counter/pointer to 0
6	RFR	Receiver FIFO reset 0 No action 1 Clears all bytes in the receiver FIFO and resets the FIFO counter/pointer to 0
7	FEN	FIFO enable 0 FIFOs are disabled and cleared 1 Transmitter and receiver FIFOs are enabled.

### 18.3.1.7 Line Control Registers (ULCR1 and ULCR2)


The ULCRs specify the data format for the UART bus and set the divisor latch access bit ULCR[DLAB], which controls the ability to access the divisor latch least and most significant bit registers and the alternate function register.

After initializing ULCR, the software should not rewrite the ULCR while valid transfers on the UART bus are active. The software should not rewrite the ULCR until the last STOP bit is received and no new characters are being transferred on the bus.

The stick parity bit, ULCR[SP], assigns a set parity value for the parity bit time slot sent on the UART bus. The set value is defined as mark parity (logic 1) or space parity (logic 0). ULCR[PEN] and ULCR[EPS] help determine the set parity value. See [Table 18-14](#). ULCR[NSTB] defines the number of STOP bits to be sent at the end of the data transfer. The receiver checks only the first STOP bit, regardless of the number

of STOP bits selected. The word length select bits (1 and 0) define the number of data bits transmitted or received as a serial character. The word length does not include START, parity, and STOP bits.

Figure 18-9 shows the bits in the ULCRs.



**Figure 18-9. Line Control Register (ULCR1 and ULCR2)**

Table 18-13 describes the ULCR fields.

**Table 18-13. ULCR Field Descriptions**

Bits	Name	Description
0	DLAB	Divisor latch access bit 0 Access to all registers except UDLB, UAFR, and UDMB. 1 Ability to access UDMB, UDLB, and UAFR.
1	SB	Set break 0 Send normal UTHR data onto the SOUT signal. 1 Force logic 0 to be on SOUT. Data in the UTHR is not affected.
2	SP	Stick parity 0 Stick parity is disabled. 1 If PEN = 1 and EPS = 1, space parity is selected; if PEN = 1 and EPS = 0, mark parity is selected.
3	EPS	Even parity select. See <a href="#">Table 18-14</a> . 0 If PEN = 1 and SP = 0 then odd parity is selected. 1 If PEN = 1 and SP = 0 then even parity is selected.
4	PEN	Parity enable 0 No parity generation and checking. 1 Generate parity bit as a transmitter, and check parity as a receiver.
5	NSTB	Number of STOP bits 0 One STOP bit is generated in the transmitted data. 1 When a 5-bit data length is selected, 1 1/2 STOP bits are generated. When either a 6-, 7-, or 8-bit word length is selected, two STOP bits are generated.
6–7	WLS	Word length select. Number of bits that comprise the character length. 00 5 bits 01 6 bits 10 7 bits 11 8 bits




**Table 18-14. Parity Selection Using ULCR[PEN], ULCR[SP], and ULCR[EPS]**

PEN	SP	EPS	Parity Selected
0	0	0	No parity
0	0	1	No parity
0	1	0	No parity
0	1	1	No parity
1	0	0	Odd parity
1	0	1	Even parity
1	1	0	Mark parity
1	1	1	Space parity

### 18.3.1.8 MODEM Control Registers (UMCR1 and UMCR2)

The UMCRs, shown in [Figure 18-10](#), control the interface with the external peripheral device on the UART bus.

**Figure 18-10. Modem Control Register (UMCR1 and UMCR2)**


[Table 18-15](#) describes the UMCR fields.

**Table 18-15. UMCR Field Descriptions**

Bits	Name	Description
0–2	—	Reserved, should be cleared
3	LOOP	Local loopback mode 0 Normal operation. 1 Functionally, the data written to UTHR can be read from URBR of the same UART, and UMCR[RTS] is tied to UMSR[CTS].
4–5	—	Reserved
6	RTS	Ready to send 0 Negates corresponding $\overline{\text{UART\_RTS}}$ output. 1 Assert corresponding $\overline{\text{UART\_RTS}}$ output. Informs external MODEM or peripheral that the UART is ready for sending/receiving data.
7	—	Reserved

### 18.3.1.9 Line Status Registers (ULSR1 and ULSR2)

The ULSRs, shown in [Figure 18-11](#), monitor the status of the data transfer on the UART buses. To isolate the status bits from the proper character received through the UART bus, software should read the ULSR and then the URBR.



**Figure 18-11. Line Status Register (ULSR1 and ULSR2)**

[Table 18-16](#) describes the ULSR fields.

**Table 18-16. ULSR Field Descriptions**

Bits	Name	Description
0	RFE	Receiver FIFO error. 0 Cleared when there are no errors in the receiver FIFO or on a read of the ULSR with no remaining receiver FIFO errors. 1 Set when one of the characters in the receiver FIFO encounters an error (framing, parity, or break interrupt).
1	TEMT	Transmitter empty 0 Either or both the UTHR or the internal transmitter shift register has a data character. In FIFO mode, a data character is in the transmitter FIFO or the internal transmitter shift register. 1 Both the UTHR and the internal transmitter shift register are empty. In FIFO mode, both the transmitter FIFO and the internal transmitter shift register are empty.
2	THRE	Transmitter holding register empty 0 UTHR is not empty. 1 A data character has transferred from the UTHR into the internal transmitter shift register. In FIFO mode, the transmitter FIFO contains no data character.
3	BI	Break interrupt 0 Cleared when the ULSR is read or when a valid data transfer is detected (that is, STOP bit is received). 1 Received data of logic 0 for more than START bit + Data bits + Parity bit + one STOP bits length of time. A new character is not loaded until SIN returns to the mark state (logic 1) and a valid START is detected. In FIFO mode, a zero character is encountered in the FIFO (the zero character is at the top of the FIFO). In FIFO mode, only one zero character is stored.
4	FE	Framing error 0 Cleared when ULSR is read or when a new character is loaded into the URBR from the receiver shift register. 1 Invalid STOP bit for receive data (only the first STOP bit is checked). In FIFO mode, FE is set when the character that detected a framing error is encountered in the FIFO (that is the character at the top of the FIFO). An attempt to resynchronize occurs after a framing error. The UART assumes that the framing error (due to a logic 0 being read when a logic 1 (STOP) was expected) was due to a STOP bit overlapping with the next START bit, so it assumes this logic 0 sample is a true START bit and then will receive the following new data.
5	PE	Parity error 0 Cleared when ULSR is read or when a new character is loaded into URBR. 1 Unexpected parity value encountered when receiving data. In FIFO mode, the character with the error is at the top of the FIFO.

Table 18-16. ULSR Field Descriptions (continued)

Bits	Name	Description
6	OE	Overrun error 0 Cleared when ULSR is read 1 Before URBR was read, it was overwritten with a new character. The old character is lost. In FIFO mode, the receiver FIFO is full (regardless of the receiver FIFO trigger level setting) and a new character has been received into the internal receiver shift register. The old character was overwritten by the new character. Data in the receiver FIFO was not overwritten.
7	DR	Data ready 0 Cleared when URBR is read or when all of the data in the receiver FIFO is read. 1 A character was received in the URBR or the receiver FIFO.

### 18.3.1.10 MODEM Status Registers (UMSR1 and UMSR2)

The UMSRs, shown in Figure 18-12, track the status of the MODEM (or external peripheral device)  $\overline{\text{CTS}}$ , set for the corresponding UART.




Figure 18-12. Modem Status Register (UMSR1 and UMSR2)


Table 18-17 describes UMSR fields.

Table 18-17. UMSR Field Descriptions

Bits	Name	Description
0–2	—	Reserved, should be cleared
3	CTS	Clear to send. Represents the inverted value of the $\overline{\text{CTS}}$ input pin from the external peripheral device. 0 Corresponding $\overline{\text{CTS}}_n$ is negated. 1 Corresponding $\overline{\text{CTS}}_n$ is asserted. The MODEM or peripheral device is ready for data transfers.
4–6	—	Reserved, should be cleared
7	DCTS	Delta clear to send 0 No change on the corresponding $\overline{\text{CTS}}_n$ signal since the last read of UMSR[CTS]. 1 $\overline{\text{CTS}}_n$ changed since the last read of UMSR[CTS]. Causes an interrupt if UIER[EMSI] is set to detect this condition.

### 18.3.1.11 Scratch Registers (USCR1 and USCR2)

USCR, shown in [Figure 18-13](#), are for debugging software or the DUART hardware. The USCRs do not affect the operation of the DUART.



**Figure 18-13. Scratch Register (USCR)**


[Table 18-18](#) describes USCR fields.

**Table 18-18. USCR Field Descriptions**

Bits	Name	Description
0–7	DATA	Data

### 18.3.1.12 Alternate Function Registers (UAFR1 and UAFR2)

The UAFRs, shown in [Figure 18-14](#), allow software to write to both UART1 and UART2 registers simultaneously with the same write operation. The UAFRs also provide a means for the device's performance monitor to track the baud clock.



**Figure 18-14. Alternate Function Register (UAFR)**

[Table 18-19](#) describes UAFR fields.

**Table 18-19. UAFR Field Descriptions**

Bits	Name	Description
0–5	—	Reserved
6	BO	Baud clock select 0 The baud clock is not gated off. 1 The baud clock is gated off.
7	CW	Concurrent write enable 0 Disables writing to both UART1 and UART2. 1 Enables concurrent writes to corresponding UART registers. A write to a register in UART1 is also a write to the corresponding register in UART2 and vice versa.

### 18.3.1.13 DMA Status Registers (UDSR1 and UDSR2)

The DMA status registers (UDSRs), shown in [Figure 18-15](#), return transmitter and receiver FIFO status and provide the ability to assist DMA data operations to and from the FIFOs.



**Figure 18-15. DMA Status Register (UDSR)**

[Table 18-20](#) describes the fields of the UDSRs.

**Table 18-20. UDSR Field Descriptions**

Bits	Name	Description
0–5	—	Reserved
6	TXRDY	Transmitter ready. Reflects the status of the transmitter FIFO or the UTHR. The status depends on the DMA mode selected, which is determined by UFCR[DMS] and UFCR [FEN]. 0 The bit is cleared, as shown in <a href="#">Table 18-22</a> . 1 This bit is set, as shown in <a href="#">Table 18-21</a> .
7	RXRDY	Receiver ready. This read-only bit reflects the status of the receiver FIFO or URBR. The status depends on the DMA mode selected, which is determined by UFCR[DMS] and UFCR [FEN]. 0 The bit is cleared, as shown in <a href="#">Table 18-24</a> . 1 This bit is set, as shown in <a href="#">Table 18-23</a> .

**Table 18-21. UDSR[TXRDY] Set Conditions**

DMS	FEN	DMA Mode	Meaning
0	0	0	TXRDY is set after the first character is loaded into the transmitter FIFO or UTHR.
0	1	0	
1	0	0	
1	1	1	TXRDY is set when the transmitter FIFO is full.

**Table 18-22. UDSR[TXRDY] Cleared Conditions**

DMS	FEN	DMA Mode	Meaning
0	0	0	TXRDY is cleared when there are no characters in the transmitter FIFO or UTHR.
0	1	0	
1	0	0	
1	1	1	TXRDY is cleared when there are no characters in the transmitter FIFO or UTHR. TXRDY remains clear while the transmitter FIFO is not yet full.

**Table 18-23. UDSR[RXRDY] Set Conditions**

DMS	FEN	DMA Mode	Meaning
0	0	0	RXRDY is set when there are no characters in the receiver FIFO or URBR.
0	1	0	
1	0	0	
1	1	1	RXRDY is set when the trigger level has not been reached and there has been no time out.

**Table 18-24. UDSR[RXRDY] Cleared**

DMS	FEN	DMA Mode	Meaning
0	0	0	RXRDY is cleared when there is at least one character in the receiver FIFO or URBR.
0	1	0	
1	0	0	
1	1	1	RXRDY is cleared when the trigger level or a time-out has been reached. RXRDY remains cleared until the receiver FIFO is empty.

## 18.4 Functional Description


The communication channel provides a full-duplex asynchronous receiver and transmitter using an operating frequency derived from the system clock signal.

The transmitter accepts parallel data with a write access to UTHR. In FIFO mode, the data is placed directly into an internal transmitter shift register, or into the transmitter FIFO—see [Section 18.4.5, “FIFO Mode.”](#) The transmitting registers convert the data to a serial bit stream by inserting the appropriate START, STOP, and optional parity bits. Finally, the registers output a composite serial data stream on the channel transmitter serial data output (SOUT). The transmitter status may be polled or interrupt driven.

The receiver accepts serial data on the channel receiver serial data input (SIN), converts the data into parallel format, and checks for START, STOP, and parity bits. In FIFO mode, the receiver removes the START, STOP, and parity bits and then transfers the assembled character from the receiver buffer, or receiver FIFO. This transfer occurs in response to a read of the UART receiver buffer register (URBR). The receiver status may be polled or interrupt-driven.

## 18.4.1 Serial Interface

The UART bus is a serial, full-duplex, point-to-point bus as shown in Figure 18-16. Therefore, only two devices are attached to the same signals and there is no need for address or arbitration bus cycles.



**Figure 18-16. UART Bus Interface Transaction Protocol Example**

A standard UART bus transfer is composed of either three or four parts:

- START bit
- Data transfer (least significant bit is first data bit on the bus)
- Parity bit (optional)
- STOP bits

An internal logic sample signal, *rxcnt*, uses the frequency of the baud-rate generator to drive the bits on SOUT.

The following sections describe the four components of the serial interface, the baud-rate generator, local loopback mode, different errors, and FIFO mode.

### 18.4.1.1 START Bit

A write to UTHR generates a START bit on the SOUT signal. Figure 18-16 shows that the START bit is defined as a logic 0. The START bit denotes the beginning of a new data transfer which is limited to the bit length programmed in ULCR. When the bus is idle, SOUT is high.

### 18.4.1.2 Data Transfer

Each data transfer contains 5, 6, 7, or 8 bits of data. The ULCR data bit length for the transmitter and receiver UART devices must agree before a transfer begins; otherwise, a parity or framing error may occur. A transfer begins when UTHR is written. At that time, a START bit is generated followed by 5 to 8 of the data bits previously written to the UTHR. The data bits are driven from the least- to the most-significant bits. After the parity and STOP bits, a new data transfer can begin if new data is written to UTHR.

### 18.4.1.3 Parity Bit

The user has the option of using even, odd, no parity, or stick parity (see Section 18.3.1.7, “Line Control Registers (ULCR1 and ULCR2).” Both the receiver and transmitter parity definitions must agree before

transferring data. When receiving data, a parity error can occur if an unexpected parity value is detected (see Section 18.3.1.9, “Line Status Registers (ULSR1 and ULSR2)”).

#### 18.4.1.4 STOP Bit

The transmitter device ends the write transfer by generating a STOP bit. The STOP bit is always high. The user can program the length of the STOP bit(s) in the ULCR. Both the receiver and transmitter STOP bit length must agree before attempting to transfer data. A framing error can occur if an invalid STOP bit is detected.

### 18.4.2 Baud-Rate Generator Logic

Each UART contains an independent programmable baud-rate generator, that is capable of taking the system clock input and dividing the input by any divisor from 1 to  $2^{16} - 1$ .

5. The baud rate is defined as the number of bits per second that can be sent over the UART bus. The formula for calculating baud rate is as follows:

$$\text{Baud rate} = (1/16) \times (\text{system clock frequency/divisor value})$$

Therefore, the output frequency of the baud-rate generator is 16 times the baud rate.

1. The divisor value is determined by the following two 8-bit registers to form a 16-bit binary number:
  - UART divisor most significant byte register (UDMB)
  - UART divisor least significant byte register (UDLB)

Upon loading either of the divisor latches, a 16-bit baud-rate counter is loaded.

The divisor latches must be loaded during initialization to ensure proper operation of the baud-rate generator. Both UART devices on the same bus must be programmed for the same baud rate before starting a transfer.

The baud clock can be passed to the performance monitor by enabling UAFR[BO]. This can be used to determine baud-rate errors.

### 18.4.3 Local Loopback Mode

Local loopback mode is provided for diagnostic testing. The data written to UTHR can be read from the receiver buffer register (URBR) of the same UART. In this mode, the MODEM control register UMC[R]RTS is internally tied to the MODEM status register UMSR[CTS]. The transmitter SOUT is set to a logic 1 and the receiver SIN is disconnected. The output of the transmitter shift register is looped back into the receiver shift register input. The  $\overline{\text{CTS}}$  (input signal) is disconnected,  $\overline{\text{RTS}}$  is internally connected to CTS, and the RTS (output signal) becomes inactive. In this diagnostic mode, data that is transmitted is immediately received. In local loopback mode the transmit and receive data paths of the DUART can be verified. Note that in local loopback mode, the transmit/receive interrupts are fully operational and can be controlled by the interrupt enable register (UIER).



## 18.4.4 Errors

The following sections describe framing, parity, and overrun errors which may occur while data is transferred on the UART bus. Each of the error bits are usually cleared, as described below, when the line status register (ULSR) is read.

### 18.4.4.1 Framing Error

When an invalid STOP bit is detected, a framing error occurs and ULSR[FE] is set. Note that only the first STOP bit is checked. In FIFO mode, ULSR[FE] is set when the character at the top of the FIFO detects a framing error. An attempt to re-synchronize occurs after a framing error. The UART assumes that the framing error (due to a logic 0 being read when a logic 1 (STOP) was expected) was due to a STOP bit overlapping with the next START bit. ULSR[FE] is cleared when ULSR is read or when a new character is loaded into the URBR from the receiver shift register.

### 18.4.4.2 Parity Error

When unexpected parity values are encountered while receiving data, a parity error occurs and ULSR[PE] is set. In FIFO mode, ULSR[PE] is set when the character with the error is at the top of the FIFO. ULSR[PE] is cleared when ULSR is read or when a new character is loaded into the URBR.

### 18.4.4.3 Overrun Error

When a new (overwriting character) STOP bit is detected and the old character is lost, an overrun error occurs and ULSR[OE] is set. In FIFO mode, ULSR[OE] is set after the receiver FIFO is full (despite the receiver FIFO trigger level setting) and a new character has been received into the internal receiver shift register. Data in the FIFO is not overwritten; only the shift register data is overwritten. Therefore, the interrupt occurs immediately. ULSR[OE] is cleared when ULSR is read.

## 18.4.5 FIFO Mode

The UARTs use an alternate mode (FIFO mode) to relieve the processor core from excessive software overhead. The FIFO control register (UFCR) is used to enable and clear the receiver and transmitter FIFOs and set the FIFO receiver trigger level UFCR[RTL] to control the received data available interrupt UIER[ERDAI].

The UFCR also selects the type of DMA signaling. The UDSR[RXRDY] indicates the status of the receiver FIFO. UDSR[TXRDY] indicate when the transmitter FIFO is full. When in FIFO mode, data written to UTHR is placed into the transmitter FIFO. The first byte written to UTHR is the first byte onto the UART bus.

### 18.4.5.1 FIFO Interrupts

In FIFO mode, the UIER[ERDAI] is set when a time-out interrupt occurs. A receive data time-out generates a maskable interrupt condition (through UIER[ERDAI]). See [Section 18.3.1.4, “Interrupt Enable Registers \(UIER1 and UIER2\)”](#).

UIIR indicates whether the FIFOs are enabled. UIIR[IID3] is set only for FIFO mode interrupts. The character time-out interrupt occurs when no characters have been removed from or input to the receiver FIFO during the last four character times and at least one character is in the receiver FIFO. The character time-out interrupt (controlled by UIIR[IID]) is cleared when URBR is read. See [Section 18.3.1.5, “Interrupt ID Registers \(UIIR1 and UIIR2\).”](#)

UIIR[FE] indicates whether FIFO mode is enabled.

### 18.4.5.2 DMA Mode Select

UDSR[RXRDY] reflects the status of the receiver FIFO or URBR. In mode 0 (UFCR[DMS] is cleared), UDSR[RXRDY] is cleared when at least one character is in the receiver FIFO or URBR; it is set when there are no more characters in the receiver FIFO or URBR. This occurs regardless of the UFCR[FEN] setting. In mode 1 (UFCR[DMS] and UFCR[FEN] are set), UDSR[RXRDY] is cleared when the trigger level or a time-out has been reached; it is set when there are no more characters in the receiver FIFO.

UDSR[TXRDY] reflects the status of the transmitter FIFO or UTHR. In mode 0 (UFCR[DMS] is cleared), UDSR[TXRDY] is cleared when there are no characters in the transmitter FIFO or UTHR; it is set after the first character is loaded into the transmitter FIFO or UTHR. This occurs regardless of the UFCR[FEN] setting. In mode 1 (UFCR[DMS] and UFCR[FEN] are set), UDSR[TXRDY] is cleared when there are no characters in the transmitter FIFO or UTHR; it is set when the transmitter FIFO is full.

See [Section 18.3.1.13, “DMA Status Registers \(UDSR1 and UDSR2\),”](#) for a complete description of the UDSR[RXRDY] and UDSR[TXRDY] bits.

### 18.4.5.3 Interrupt Control Logic

An interrupt is active when DUART interrupt ID register bit 0 (UIIR[0]), is cleared. UIER is used to mask specific interrupt types. See [Section 18.3.1.4, “Interrupt Enable Registers \(UIER1 and UIER2\).”](#)

When the interrupts are disabled in UIER, polling software can not use UIIR[0] to determine whether the UART is ready for service. Software must monitor the appropriate ULSR and UMSR bits. UIIR[0] can be used for polling if the interrupts are enabled in UIER.

## 18.5 DUART Initialization/Application Information

The following requirements must be met for DUART accesses:

- All DUART registers must be mapped to a cache-inhibited and guarded area. (That is, the WIMG setting in the MMU needs to be 0b01x1.)
- All DUART registers are 1 byte wide. Reads and writes to these registers must be byte-length operations.

A system reset puts the DUART registers to a default state. Before the interface can transfer serial data, the following initialization steps are recommended:

1. Update the programmable interrupt controller (PIC) DUART channel interrupt vector source registers.
2. Set data attributes and control bits in the ULCR, UFCR, UAFR, UMCR, UDLB, and UDMB.

3. Set the data attributes and control bits of the external MODEM or peripheral device.
4. Set the interrupt enable register (UIER).
5. To start a write transfer, write to the UTHR.
6. Poll UIIR if the interrupts generated by the DUART are masked.



# Chapter 19

## Serial Peripheral Interface

### 19.1 Overview

The serial peripheral interface (SPI) allows the device to exchange data between other PowerQUICC® family chips, the MC68360, MC68302, M68HC11, and M68HC05 microcontroller families, and other family devices. The SPI can be used to communicate with peripheral devices such as EEPROMs, real-time clocks, A/D converters, and ISDN devices.

The SPI is a full-duplex, synchronous, character-oriented channel that supports a four-wire interface (receive, transmit, clock, and slave select). The SPI block consists of transmitter and receiver sections, an independent baud-rate generator, and a control unit. The transmitter and receiver sections use the same clock, which is derived from the SPI baud rate generator in master mode or externally in slave mode. During an SPI transfer, data is sent and received simultaneously.

The SPI receiver and transmitter are double-buffered, as shown in [Figure 19-1](#), giving an effective FIFO size (latency) of 2 characters. The SPI's MSB/LSB is shifted out first. When the SPI is disabled in the SPI mode register (SPMODE[EN] = 0), it consumes little power.

## 19.2 Introduction

The SPI block diagram is shown in Figure 19-1.




Figure 19-1. SPI Block Diagram

### 19.2.1 Features

The major features of the SPI are listed as follows:

- Four-signal interface (SPIMOSI, SPIMISO, SPICLK, and  $\overline{\text{SPISEL}}$ )
- Full-duplex operation
- Works with 32-bit data characters or with a range from 4-bit to 16-bit data characters
- Supports back-to-back character transmission and reception
- Supports reverse data mode for 8/16/32 character length
- Supports master SPI mode
- Supports multiple-master environment
- Maximum clock rate is (input clock rate/4) in master mode; (input clock rate/2) in slave mode
- Independent programmable baud rate generator
- Programmable clock phase and polarity
- Local loopback capability for testing
- Open-drain outputs support multiple-master configuration

## 19.2.2 SPI Transmission and Reception Process

Because the SPI is a character-oriented communication unit, the core is responsible for packing and unpacking the receive and transmit frames. A frame consists of all of the characters transmitted or received during a completed SPI transmission session, from the first character written to the SPITD register to the last character transmitted following the setting of SPCOM[LST]. See [Section 19.4.1.4, “SPI Command Register \(SPCOM\),”](#) for more information.

The core receives data by reading the SPI receive data hold register (SPIRD). The SPI then clears the not empty SPIE[NE] to free up the SPIRD register for the next receive operation. The core transmits data by writing it into the SPI transmit data hold register (SPITD). The SPI then clears the not full (NF) bit in the SPI event register (SPIE) to indicate that the SPITD register contains a character for transmission. When the next character to be transmitted is going to be the final one in the current frame, the core sets SPCOM[LST], and then writes the final character to SPITD.

The SPI core handshake protocol can be implemented by either using polling or interrupts. When using a polling, the core reads the SPIE in a predefined frequency and acts according to the value of the SPIE bits. The polling frequency depends on the SPI serial channel frequency. When using the interrupt mechanism, setting either the not full (NF) or not empty (NE) bits of SPIE causes an interrupt to the processor core. The core then reads SPIE and acts accordingly. The three basic modes of operation for transmitting and receiving are master, slave and multiple-master.

### NOTE

When both NE and NF bits are set, the processor core should read the received data before transmitting new data.

The SPMODE[LEN] determines the character length sent by the hardware. The core is responsible for any bit manipulation to pack/unpack data into the appropriate character length. See the SPMODE[LEN] description in [Table 19-4](#) for more information.

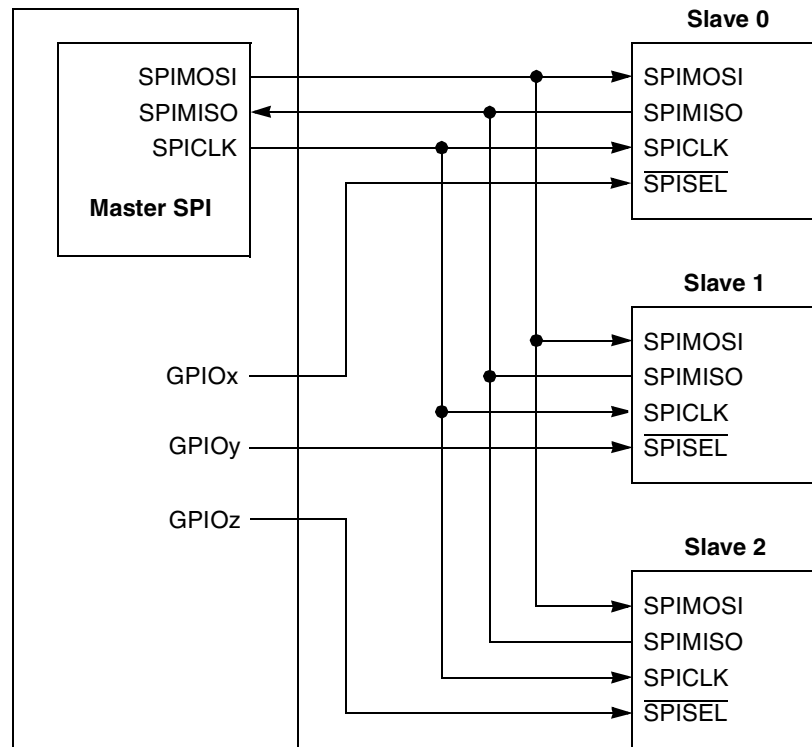
## 19.2.3 Modes of Operation

The SPI can be programmed to work in a single- or multiple-master environment. This section describes SPI master and slave operations in a single-master configuration. It also discusses the multiple master environment.

The following sections summarize the main modes of operation that the SPI supports.

### 19.2.3.1 SPI as a Master Device

In master mode, the SPI sends a message to the slave peripheral, which sends back a simultaneous reply. A single master device with multiple slaves can use general-purpose parallel I/O signals to selectively enable slaves, as shown in [Figure 19-2](#). To eliminate the multi-master error in a single-master environment, the master's  $\overline{\text{SPISEL}}$  input should be forced inactive by an external pull up.



**Figure 19-2. Single-Master/Multi-Slave Configuration**

To start exchanging data, the processor core writes the data to be sent into the SPITD register. The SPI then generates programmable clock pulses on SPICLK for each character. It shifts Tx data out on the SPI master-out slave-in (SPIMOSI) and Rx data in on the SPI master-in slave-out (SPIMISO) simultaneously. During transmission, the core is responsible for supplying the data whenever the SPI requests it to ensure smooth operation. After the last data (LST command and data afterwards), the first character written to SPITD acts as a start command for the SPI.

The SPI continues transmitting and receiving characters until SPCOM[LST] is set or an error occurs.

The SPI sets SPIE[NF] to issue a maskable interrupt to the interrupt controller whenever its transmit buffer is not full. It also sets the NF bit after sending the last word. In response, the core should read the exception flags that relate to the last word. The SPI sets SPIE[NE] to issue a maskable interrupt to the interrupt controller whenever the receiver buffer has been filled with data.

### 19.2.3.2 SPI as a Slave Device

In slave mode, the SPI receives messages from an SPI master and sends a simultaneous reply. The slave's  $\overline{\text{SPISEL}}$  must be asserted before Rx clocks are recognized. Once  $\overline{\text{SPISEL}}$  is asserted, SPICLK becomes an input from the master to the slave. SPICLK can be any frequency from DC to input clock/2.


To prepare for data transfers, the core writes data to be sent into the SPITD register. Once  $\overline{\text{SPISEL}}$  is asserted, the slave shifts data out from SPIMISO and in through SPIMOSI. The SPI sets the NF bit of the SPIE register and a maskable interrupt is issued when a full buffer finishes receiving and sending or after an error. The SPI continues reception until  $\overline{\text{SPISEL}}$  is negated.



Transmission continues until no more data is available or  $\overline{\text{SPISEL}}$  is negated. Transmission continues once  $\overline{\text{SPISEL}}$  is reasserted and  $\text{SPICLK}$  begins toggling. After the characters in the buffer are sent, the SPI sends one as long as  $\overline{\text{SPISEL}}$  remains asserted.

### 19.2.3.3 SPI in Multiple-Master Operation

The SPI can operate in a multiple-master environment in which all SPI devices are connected to the same bus. In this configuration, the  $\text{SPIMOSI}$ ,  $\text{SPIMISO}$ , and  $\text{SPICLK}$  signals of all SPIs are shared; but the  $\overline{\text{SPISEL}}$  inputs are connected separately, as shown in [Figure 19-3](#). Only one SPI device can act as master at a time—all others must be slaves. When a SPI is configured as a master, if its  $\overline{\text{SPISEL}}$  input is asserted, a multiple-master error occurs because more than one SPI device is a bus master. The SPI sets  $\text{SPIE}[\text{MME}]$  in the SPI event register and a maskable interrupt is issued to the core. It also disables SPI operation and the output drivers of the SPI signals. The core must clear  $\text{SPMODE}[\text{EN}]$ , correct the problems, and clear  $\text{SPIE}[\text{MME}]$  before the SPI can be used again.



**Notes:**

1. All signals are open-drain.
2. For a multiple-master configuration with more than two masters,  $\overline{\text{SPISEL}}$  and SPIE[MME] do not detect all possible conflicts.
3. It is the responsibility of software to arbitrate for the SPI bus (with token passing, for example).
4.  $\text{SELOUT}_x$  signals are implemented in software with general-purpose I/O signals.

**Figure 19-3. Multiple-Master Configuration**

The maximum sustained data rate that the SPI supports is input clock/50. However, the SPI can transfer a single character at much higher rates—input clock/4 in master mode and input clock/2 in slave mode. Gaps should be inserted between multiple characters to keep from exceeding the maximum sustained data rate.

### 19.3 External Signal Descriptions

The SPI's four wire interface consists of transmit, receive, clock, and slave select.

## 19.3.1 Overview

Table 19-1 lists signal properties.

**Table 19-1. Signal Properties**

Name	Function	Reset	Pull Up
SPIMISO	Master input slave output	—	Required in open drain mode
SPIMOSI	Master output slave input	—	Required in open drain mode
SPICLK	Input/output serial clock connected to the other SPICLK	—	Required in open drain mode
SPISEL	SPI slave select	—	Required in open drain mode

## 19.3.2 Detailed Signal Descriptions

Table 19-2 describes the signals in detail.

**Table 19-2. Detailed Signal Descriptions**

Signal	I/O	Description	
SPIMISO	I/O	Master input slave output	
		<b>State Meaning</b>	Asserted—The data that has been transmitted/received from/to the SPI (depends if master or slave mode) is high Negated—The data that has ben transmitted/received from/to the SPI (depends if master or slave mode) is low
		<b>Timing</b>	Assertion—According to the SPICLK assertion/negation/in the middle of phase (depends on SPMODE) Negation—According to the SPICLK assertion/negation/in the middle of phase (depends on SPMODE)
SPIMOSI	I/O	Master output slave input	
		<b>State Meaning</b>	Asserted—The data that has been transmitted/received from/to the SPI (depends if master or slave mode) is high Negated—The data that has ben transmitted/received from/to the SPI (depends if master or slave mode) is low
		<b>Timing</b>	Assertion—According to the SPICLK assertion/negation/in the middle of phase (depends on SPMODE) Negation—According to the SPICLK assertion/negation/in the middle of phase (depends on SPMODE)

Table 19-2. Detailed Signal Descriptions (continued)

Signal	I/O	Description	
SPICLK	I/O	Serial clock in or serial clock out for slave or master mode respectively	
		<b>State Meaning</b>	Assertion/Negation according to SPMODE[PM, DIV16] register rate configuration
		<b>Timing</b>	Assertion/Negation—during frame reception/transmission
$\overline{\text{SPISEL}}$	I	SPI slave select	
		<b>State Meaning</b>	Asserted—In slave mode declares the slave has been selected for the coming frame. In master mode assertion causes MME multiple-master error. Negated—In slave mode means the specific SPI has not been selected. In master mode needs to be negated for regular operation.
		<b>Timing</b>	Assertion—In slave mode along with the data from the slave Negation—In slave mode with the end of the frame (according to SPMODE[LEN]). In master mode before data is first written to SPITD and remains constant.

The SPI can be configured as a slave or a master in single- or multiple-master environments mode. The master SPI generates the transfer clock SPICLK using the SPI baud rate generator (BRG). The SPI BRG takes its input from input clock, which is generated in the device clock synthesizer.

SPICLK is a gated clock, active only during data transfers. Four combinations of SPICLK phase and polarity can be configured with the clock invert (SPMODE[CI]) and clock phase (SPMODE[CP]) register bits. SPI signals can also be configured as open-drain to support a multiple-master configuration in which a shared SPI signal is driven by the device or an external SPI device.

The SPI master-in slave-out SPIMISO signal acts as an input for master devices and as an output for slave devices. Conversely, the master-out slave-in SPIMOSI signal is an output for master devices and an input for slave devices. The dual functionality of these signals allows the SPIs in a multiple-master environment to communicate with one another using a common hardware configuration.

- When the SPI is a master, SPICLK is the clock output signal that shifts received data in from SPIMISO and transmitted data out to SPIMOSI. SPI masters must output a slave select signal to enable SPI slave devices by using a separate general-purpose I/O signal. Assertion of the  $\overline{\text{SPISEL}}$  while the SPI is configured as a master causes an error.
- When the SPI is a slave, SPICLK is the clock input that shifts received data in from SPIMOSI and transmitted data out through SPIMISO.  $\overline{\text{SPISEL}}$  is the enable input to the SPI slave. In a multiple-master environment,  $\overline{\text{SPISEL}}$  (always an input) is also used to detect an error when more than one master is operating.

## 19.4 Memory Map/Register Definition

Table 19-3 shows the memory mapped registers of the SPI and their offsets. It lists the offset, name, and a cross-reference to the complete description of each register. Note that the full register address is comprised of IMMRBAR together with the SPI block base address and offset listed in Table 19-3. Undefined 4-byte address spaces within offset 0x000–0xFFF are reserved.

Table 19-3. SPI Register Summary

Offset	Register	Access	Reset Value	Section/Page
0x000–0x01F	Reserved	—	—	—
0x020	SPI mode register (SPMODE)	R/W	0x0000_0000	19.4.1.1/1919-9
0x024	SPI event register (SPIE)	Mixed	0x0000_0000	19.4.1.2/1919-12
0x028	SPI mask register (SPIM)	R/W	0x0000_0000	19.4.1.3/1919-13
0x02C	SPI command register (SPCOM)	W	0x0000_0000	19.4.1.4/1919-14
0x030	SPI transmit register (SPITD)	W	0x0000_0000	19.4.1.5/1919-14
0x034	SPI receive register (SPIRD)	R	0xFFFF_FFFF	19.4.1.6/1919-15
0x038–0xFFFF	Reserved	—	—	—

## 19.4.1 Register Descriptions

### 19.4.1.1 SPI Mode Register (SPMODE)

SPMODE, shown in [Figure 19-4](#), controls both the SPI operation mode and clock source.

Offset 0x020

Access: Read/write




Figure 19-4. SPMODE-SPI Mode Register Definition

[Table 19-4](#) describes the SPMODE fields.

Table 19-4. SPMODE Field Descriptions

Bits	Name	Description
0	—	Reserved. Should be cleared.
1	LOOP	Loop mode. Enables local loopback operation. 0 Normal operation. 1 Loopback mode. Used to test the SPI controller internal functionality, the transmitter output is internally connected to the receiver input. The receiver and transmitter operate normally, except that received data is ignored. The SPI acts normally in loop back mode; therefore, negating SPISEL in slave mode stops transmission, negating it in master mode and causing an MME error.
2	CI	Clock invert. Inverts SPI clock polarity. See <a href="#">Figure 19-5</a> and <a href="#">Figure 19-6</a> for more information 0 The inactive state of SPICLK is low. 1 The inactive state of SPICLK is high.
3	CP	Clock phase. Selects the transfer format. See <a href="#">Figure 19-5</a> and <a href="#">Figure 19-6</a> for more information. 0 SPICLK starts toggling at the middle of the data transfer. 1 SPICLK starts toggling at the beginning of the data transfer.

Table 19-4. SPMODE Field Descriptions (continued)

Bits	Name	Description
4	DIV16	Divide by 16. Selects the clock source for the SPI baud rate generator (SPI BRG) when configured as an SPI master. In slave mode, SPICLK is the clock source. 0 The SPI block input clock is the input to the SPI BRG. 1 The SPI block input clock/16 is the input to the SPI BRG. In slave mode this bit must be cleared.
5	REV	Reverse data mode for 8-/16-/32-bit character length only (see Section 19.4.1.6.1, “Reverse Mode SPMODE[REV] Examples.”) 0 LSB sent/received first (for data LEN < 32 the data is located at the lower half-word LSB) 1 MSB sent/received first
6	M/S	Master/slave. Selects master or slave mode. 0 The SPI is a slave. 1 The SPI is a master.
7	EN	Enable SPI. Any other bits in SPMODE must not change when EN is set. 0 The SPI is disabled. The SPI is in a idle state and consumes minimal power. The SPI BRG is not functioning and the input clock is disabled. 1 The SPI is enabled. <b>Note:</b> The SPI controller requires a minimal gap of at least 10 input clocks between disabling the SPI and re-enabling. This minimal gap is sufficient provided that SPMODE[PM] and SPMODE[DIV16] are cleared during the time in which SPMODE[EN] is cleared.
8–11	LEN	Character length in bits per character. LEN can be either 32-bits, or 4- to 16-bits that are shown as follows: 0000 32-bit characters 0001–0010 Reserved, causes erratic behavior. 0011 4-bit characters ... 1111 16-bit characters The TX and RX registers (SPITD, SPIRD) hold 32 bits at a time. A character length of 32 bits fills the TX and RX registers; therefore, all of the bits in these registers are valid. However, if the character length selected by LEN is equal or less than 16 bits, then the valid bits will reside in the lower half-word of the transmit and receive registers. For example, if the character length is set to 16 bits then the valid bits will be 16–31, if the character length is set to 5 bits that the valid bits will be 16–20. Note that the transmit and receive registers each can hold only one character regardless of the character length.
12–15	PM	Prescale modulus select. Specifies the divide ratio of the prescale divider in the SPI clock generator. The SPI baud rate generator clock source (either input clock or input clock divided by 16, depending on DIV16 bit) is divided by $4 * ([PM] + 1)$ , a range from 4 to 64. The clock has a 50% duty cycle. For example, if the prescale modulus is set to PM = 0011 and DIV16 is set, the system/SPICLK clock ratio will be $16 * (4 * (0011 + 1)) = 256$ . In slave mode this field must be cleared.
16–18	—	Reserved. Should be cleared.
19	OD	Open drain mode. 0 All output pins are configured to normal mode. 1 All output pins are configured to open drain mode.
20–31	—	<b>Note:</b> Reserved. Should be cleared.

Figure 19-5 shows the SPI transfer format in which SPICLK starts toggling in the middle of the transfer (SPMODE[CP] = 0).




Figure 19-5. SPI Transfer Format with SPMODE[CP] = 0

Figure 19-6 shows the SPI transfer format in which SPICLK starts toggling at the beginning of the transfer (SPMODE[CP] = 1).




Figure 19-6. SPI Transfer Format with SPMODE[CP] = 1


### 19.4.1.2 SPI Event Register (SPIE)

The SPI event register (SPIE) generates interrupts and reports events recognized by the SPI. When an event is recognized, the SPI sets the corresponding SPIE bit. Most SPIE bits can be cleared by writing a '1'. Writing '0' has no effect. Setting a bit in the SPI mask register (SPIM) enables, and clearing a bit masks the corresponding interrupt. Unmasked SPIE bits must be cleared before the core clears internal interrupt requests. Figure 19-7 shows SPI event register.

## Serial Peripheral Interface

Offset 0x024

Access: Mixed



**Figure 19-7. SPIE—SPI Event Register Definition**

Table 19-5 describes the SPIE fields.

**Table 19-5. SPIE Field Descriptions**

Bits	Name	Description
0–16	—	Reserved, should be cleared.
17	LT	Last character was transmitted. The last character is transmitted and new data can be written to SPID for further transmission.
18	DNR	<b>Note:</b> Data not ready. In slave mode only when $\overline{\text{SPISEL}}$ is asserted before data is ready in the SPI, IDLE is sent on the line and UN bit is also asserted the SPI should be disable to restart its operation.
19	OV	Slave/master overrun. Indicates whether an overrun has occurred during reception. In case of overrun the SPI continues transmission/reception process while reporting overrun for the missing characters.
20	UN	Slave underrun. Indicates whether the SPI transmitter did not have data to transmit on time, and therefore, whether IDLE was sent on the line. Valid only in slave mode ( $\text{SPMODE}[\text{M/S}] = 0$ ). In master mode ( $\text{SPMODE}[\text{M/S}] = 1$ ) if the SPI's transmitter has no valid data to transmit the SPICLK stop toggling and transmission/reception is frozen (no underrun is reported), when data is written to the SPITD the transmission resumes.
21	MME	Multiple-master error. Set when $\overline{\text{SPISEL}}$ is asserted externally while the SPI is in master mode. Note that the MME error can occur in loopback mode.
22	NE	Not empty. When set Indicates that SPIRD contains a received character. 0 The receiver is empty 1 The receiver has valid received data and indications about LST (command register) and OV (SPIE).The core is free to read the content of the receiver. Reading the receiver SPIRD clears NE if no more data is available.
23	NF	Not full. Indicates whether SPITD is not in use and a new character can be written to it by the core. 0 The transmitter is full. 1 The transmitter is not full. The core is free to write to the transmitter. NF must be clear to enable the transmission of another character (writing to the transmitter clears NF)
24–31	—	Reserved. Should be cleared.

### 19.4.1.3 SPI Mask Register (SPIM)


The SPI mask register (SPIM), shown in Figure 19-8, enables/masks interrupts for events that are recognized by the SPI. When an event is recognized, the SPI sets the corresponding SPIE bit. Setting a



SPIM bit enables and clearing a SPIM bit masks the corresponding interrupt. Unmasked SPIE bits must be cleared before the core clears its internal interrupt requests.

Offset 0x028

Access: Read/write



**Figure 19-8. SPIM—SPI Mask Register Definition**

Table 19-6 describes the SPIM fields.

**Table 19-6. SPIM Field Descriptions**


Bits	Name	Description
0–16	—	Reserved, should be cleared.
17	LT	Last character transmitted 0 LT event will not cause an SPI interrupt 1 LT event causes an SPI interrupt
18	DNR	In slave mode data not ready 0 Slave DNR event will not cause an SPI interrupt <b>Note:</b> 1 Slave DNR event causes an SPI interrupt
19	OV	Slave/Master Overrun interrupt mask 0 Slave/Master Overrun event will not cause an SPI interrupt 1 Slave/Master Overrun event causes an SPI interrupt
20	UN	Slave Underrun interrupt mask 0 Slave Underrun event will not cause an SPI interrupt 1 Slave Underrun event causes an SPI interrupt
21	MME	Multimaster error interrupt mask 0 Multimaster error event will not cause an SPI interrupt 1 Multimaster error event causes an SPI interrupt
22	NE	Not Empty interrupt mask 0 Not Empty event will not cause an SPI interrupt 1 Not Empty event causes an SPI interrupt
23	NF	Not Full interrupt mask 0 Not Full event will not cause an SPI interrupt 1 Not Full event causes an SPI interrupt
24–31	—	Reserved, should be cleared.

### 19.4.1.4 SPI Command Register (SPCOM)

The SPI command register (SPCOM), shown in [Figure 19-9](#), is used to end SPI operation.

Offset 0x02C

Access: Write only



**Figure 19-9. SPI Command Register Definition**

[Table 19-7](#) describes the SPCOM fields.

**Table 19-7. SPCOM Field Descriptions**


Bits	Name	Description
0–8	—	Reserved, should be cleared.
9	LST	This bit represents the last character. Should be set before the last character is written to the SPITD. This results in SPIE[LT] being set when the character is fully transmitted and by that gives indication about the frame being fully transmitted. 0 This character is not the last character of the frame 1 This character is the last character of the frame
10–31	—	Reserved, should be cleared.

### 19.4.1.5 SPI Transmit Data Hold Register (SPITD)

SPITD holds the character to be transmitted. The number of bits in each character is specified by SPMODE[LEN]. Each time SPIE[NF] is set, the core can write another character of data to SPITD, if there is no error indication in the SPIE. At the end of the frame the core should set SPCOM[LST] and prepare the last character of data. [Figure 19-10](#) shows the SPI transmit data hold register.

Offset 0x030

Access: Write only



**Figure 19-10. SPI Transmit Data Hold Register Definition**

[Table 19-8](#) shows the field descriptions of the SPI transmit data hold register.

**Table 19-8. SPI Transmit Data Hold Field Descriptions**

Bits	Name	Description
0–31	DATA	These bits are the data to be sent.

### 19.4.1.6 SPI Receive Data Hold Register (SPIRD)

SPIRD, shown in Figure 19-11, is used to receive a character of data from the SPI channel. Each time SPIE[NE] is set, the core can read SPIRD.

Offset 0x034

Access: Read-only




Figure 19-11. SPI Receive Data Hold Register Definition

Table 19-9 shows the field descriptions of the SPI receive data hold register.

Table 19-9. SPI Receive Data Hold Field Descriptions

Bits	Name	Description
0–31	DATA	Received data. These bits are the received data from the SPI bus.

#### 19.4.1.6.1 Reverse Mode SPMODE[REV] Examples

In reverse data mode (SPMODE[REV] = 1) and regular data mode (SPMODE[REV] = 0) the data is placed in the SPIRD after reception is completed as described below for character length of 8 bits (SPMODE[LEN] = 7).

Offset 0x036

Access: Read-only




Figure 19-12. Example SPMODE[REV] = 0 SPMODE[LEN] = 7 LSB Sent First

Offset 0x036

Access: Read-only






Figure 19-13. Example SPMODE[REV] = 1 SPMODE[LEN] = 7 MSB Sent First



**Figure 19-14. Example SPMODE[REV] = 1 SPMODE[LEN] = 15 MSB Sent First**



**Figure 19-15. Example SPMODE[REV] = 0 SPMODE[LEN] = 15 LSB Sent First**

## 19.5 Initialization/Application Information

The following sections describe programming examples of the SPI master and slave.

### 19.5.1 SPI Master Programming Example

The following sequence initialize the SPI to run at a high speed in master mode:

1. Configure a parallel I/O signal to operate as the SPI select output signal if needed.
2. Write 0xFFFFFFFF to SPIE to clear any previous events. Configure SPIM to enable all desired SPI interrupts.
3. Configure SPMODE to enable normal operation (not loopback), master mode, SPI enabled, character length, and the fastest speed possible.
4. Write the first character to be sent to SPITD.

### 19.5.2 SPI Slave Programming Example

The following is an example initialization sequence to follow when the SPI is in slave mode. It is very similar to the SPI master example, except that  $\overline{\text{SPIS\!E\!L}}$  is used instead of a general-purpose I/O signal.

1. Write 0xFFFFFFFF to SPIE to clear any previous events.
2. Configure SPIM to enable all desired SPI interrupts.
3. Configure SPMODE to enable normal operation (not loopback), slave mode, SPI enabled, and characters length.
4. Write the first data to be sent to SPITD, to enable the SPI to be ready once the master begins to transfer.

# Chapter 20

## JTAG/Testing Support

### 20.1 Overview

The device provides a JTAG (Joint Test Action Group) interface to facilitate boundary-scan testing. The JTAG interface complies to the IEEE 1149.1 boundary-scan specification. For additional information about JTAG operations, refer to the IEEE 1149.1 specification.

The JTAG interface consists of a set of five signals, three JTAG registers (see [Section 20.3, “JTAG Registers and Scan Chains,”](#)) and a test access port (TAP) controller, described in the following sections. A block diagram of the JTAG interface is shown in [Figure 20-1](#).




Figure 20-1. JTAG Interface Block Diagram

### 20.2 JTAG Signals

The device provides the following five dedicated JTAG signals:

- Test data input (TDI)
- Test data output (TDO)
- Test mode select (TMS)

- Test reset ( $\overline{\text{TRST}}$ )
- Test clock (TCK)

The TDI and TDO signals input and output all instructions and data to the JTAG scan registers. JTAG operations are controlled by the TAP controller through the TMS and TCK signals. Boundary-scan data is latched by the TAP controller on the rising edge of the TCK signal. The  $\overline{\text{TRST}}$  signal is specified as optional by the IEEE 1149.1 specification, and is used to reset the TAP controller asynchronously. The assertion of the  $\overline{\text{TRST}}$  signal at power-on reset ensures that the JTAG logic does not interfere with the normal operation of the device.

## 20.2.1 External Signal Descriptions

The JTAG signals are summarized in [Table 20-1](#).

**Table 20-1. JTAG Test Signals Summary**

Name	Description	Functional Block	Function	Reset Value	I/O
TCK	Test clock	Debug	Clock for JTAG testing.	—	I
TDI	Test data input		Serial input for instructions and data to the JTAG test subsystem. Internally pulled up.	—	I
TDO	Test data output		Serial data output for the JTAG test subsystem. High impedance except when scanning out data.	High impedance	O
TMS	Test mode select		Carries commands to the TAP controller for boundary scan operations. Internally pulled up.	—	I
$\overline{\text{TRST}}$	Test reset		Resets the TAP controller asynchronously. Internally pulled up.	—	I

[Table 20-2](#) shows detailed descriptions of the JTAG test signals.

**Table 20-2. JTAG Test—Detailed Signal Descriptions**

Signal	I/O	Description	
TCK	I	JTAG test clock.	
		<b>State Meaning</b>	Asserted/Negated—Should be driven by a free-running clock signal with a 30–70% duty cycle. Input signals to the TAP are clocked in on the rising edge. Changes to the TAP output signals occur on the falling edge. The test logic allows TCK to be stopped.
		<b>Timing</b>	See IEEE 1149.1 specification for more details.
TDI	I	JTAG test data input.	
		<b>State Meaning</b>	Asserted/Negated—The value present on the rising edge of TCK is clocked into the selected JTAG test instruction or data register. An unterminated input appears as a high signal level to the test logic due to an internal pull-up resistor.
		<b>Timing</b>	See IEEE 1149.1 specification for more details.

Table 20-2. JTAG Test—Detailed Signal Descriptions (continued)

Signal	I/O	Description	
TDO	O	JTAG test data output.	
		<b>State Meaning</b>	Asserted/Negated—The contents of the selected internal instruction or data register are shifted out on this signal on the falling edge of TCK. Remains in a high-impedance state except when scanning data.
		<b>Timing</b>	See IEEE 1149.1 specification for more details.
TMS	I	JTAG test mode select.	
		<b>State Meaning</b>	Asserted/Negated—Decoded by the internal JTAG TAP controller to distinguish the primary operation of the test support circuitry. An unterminated input appears as a high signal level to the test logic due to an internal pull-up resistor.
		<b>Timing</b>	See IEEE 1149.1 specification for more details.
$\overline{\text{TRST}}$	I	JTAG test reset.	
		<b>State Meaning</b>	Asserted—Causes asynchronous initialization of the internal JTAG TAP controller. Must be asserted during power-on reset in order to properly initialize the JTAG TAP and for normal operation of the device. An unterminated input appears as a high signal level to the test logic due to an internal pull-up resistor. Negated— Normal operation.
		<b>Timing</b>	See IEEE 1149.1 specification for more details.

## 20.3 JTAG Registers and Scan Chains

The bypass, boundary-scan, and instruction JTAG registers and their associated scan chains are mandatory for compliance with the IEEE 1149.1 specification.

- Bypass register. The bypass register is a single-stage register used to bypass the boundary-scan latches of the device during board-level boundary-scan operations involving components other than the device. The use of the bypass register reduces the total scan string size of the boundary-scan test.
- Boundary-scan registers. The JTAG interface provides a chain of registers dedicated to boundary-scan operations. To be JTAG-compliant, these registers cannot be shared with any functional registers of the device. The boundary-scan register chain includes registers controlling the direction of the input/output drivers, in addition to the registers reflecting the signal value received or driven.

The boundary-scan registers capture the input or output state of the device's signals during a Capture\_DR TAP controller state. When a data scan is initiated following the Capture\_DR state, the sampled values are shifted out through the TDO output while new boundary-scan register values are shifted in through the TDI input. At the end of the data scan operation, the boundary-scan registers are updated with the new values during an update\_DR TAP controller state.

- Instruction register. The 8-bit JTAG instruction register serves as an instruction and status register. As TAP controller instructions are scanned in through the TDI input, the TAP controller status bits are scanned out through the TDO output.

- TAP controller. The device provides a standard JTAG TAP controller that controls instruction and data scan operations. The TMS signal controls the state transitions of the TAP controller.



# Chapter 21

## General Purpose I/O (GPIO)

### 21.1 Introduction

This chapter describes the general-purpose I/O module, including pin descriptions, register settings, and interrupt capabilities. Figure 21-1 shows the block diagram of the GPIO module.




Figure 21-1. GPIO Module Block Diagram

#### 21.1.1 Overview

The GPIO module supports 32 general-purpose I/O ports. Each port can be configured as an input or as an output. If a port is configured as an input, it can optionally generate an interrupt on detection of a change. If a port is configured as an output, it can be individually configured as an open-drain or a fully active output.

#### 21.1.2 Features

The GPIO unit implements the following features:

- 32 input/output ports
- Some ports have dedicated processor signals. Others are multiplexed together with other functional signals. See Chapter 3, “Signal Descriptions.”
- All signals are configured as inputs when the device comes out of reset and also when  $\overline{\text{HRESET}}$  is asserted.

- Open-drain capability on all ports
- All ports can optionally generate an interrupt upon changing their state.

## 21.2 External Signal Description

The following section provides information about GPIO signals.

### 21.2.1 Signals Overview

Table 21-1 provides detailed descriptions of the external GPIO signals.

**Table 21-1. IPIC External Signals—Detailed Signal Descriptions**

Signal	I/O	Description
GPIO[0:31]	I/O	General purpose I/O. Each signal can be set individually to act as input or output, according to application needs.
		<b>State Meaning</b> Asserted/Negated—Defined per application.
		<b>Timing</b> Assertion/Negation—Inputs can be asserted completely asynchronously. Outputs are asynchronous to any externally visible clock

## 21.3 Memory Map/Register Definition

The GPIO has programmable registers that occupy 24 bytes of memory-mapped space. Note that reading undefined portions of the memory map returns all zeros and writing has no effect.

All GPIO registers are 32 bits wide and are located on 32-bit address boundaries. All addresses used in this chapter are offsets from the address held in IMMRBAR as defined in Chapter 2, “Memory Map.”


Table 21-2 shows the memory map of GPIO.

**Table 21-2. GPIO Register Address Map**

Offset	Register	Access	Reset Value	Section/Page
0xC00	GPIO direction register (GPDIR)	R/W	0x0000_0000	<a href="#">21.3.1/21-3</a>
0xC04	GPIO open drain register (GPODR)	R/W	0x0000_0000	<a href="#">21.3.2/21-3</a>
0xC08	GPIO data register (GPDAT)	R/W	0x0000_0000	<a href="#">21.3.3/21-4</a>
0xC0C	GPIO interrupt event register (GPIER)	w1c	Undefined	<a href="#">21.3.4/21-4</a>
0xC10	GPIO interrupt mask register (GPIMR)	R/W	0x0000_0000	<a href="#">21.3.5/21-4</a>
0xC14	GPIO external interrupt control register (GPICR)	R/W	0x0000_0000	<a href="#">21.3.6/21-5</a>

### 21.3.1 GPIO Direction Register (GPDIR)

The GPIO direction registers (GPDIR), shown in [Figure 21-2](#), defines the direction of the individual ports.



**Figure 21-2. GPIO Direction Register (GPDIR)**


[Table 21-3](#) defines the bit fields of GPDIR.

**Table 21-3. GPDIR Bit Settings**

Bits	Name	Description
0–31	<i>Dn</i>	Direction. Indicates whether a signal is used as an input or an output. 0 The corresponding signal is an input. 1 The corresponding signal is an output.

### 21.3.2 GPIO Open Drain Register (GPODR)

The GPIO open drain register (GPODR), shown in [Figure 21-3](#), defines the way individual ports drive their output.



**Figure 21-3. GPIO Open Drain Register (GPODR)**


[Table 21-4](#) defines the bit fields of GPODR.

**Table 21-4. GPODR Bit Settings**

Bits	Name	Description
0–31	<i>Dn</i>	Open-drain configuration. Indicates whether a signal is actively driven as an output or an open-drain driver. This register has no effect on signals programmed as inputs in the corresponding GPDIR. 0 The I/O signal is actively driven as an output. 1 The I/O signal is an open-drain driver. As an output, the signal is driven active-low, otherwise it is three-stated.

### 21.3.3 GPIO Data Register (GPDAT)

The GPIO data register (GPDAT), shown in [Figure 21-4](#), carries the data in/out for the individual ports.



**Figure 21-4. GPIO Data Register (GPDAT)**


[Table 21-5](#) defines the bit fields of GPDAT.

**Table 21-5. GPnDAT Bit Settings**

Bits	Name	Description
0–31	$D_n$	Data. Write data is latched and presented on external signals if GPDIR has configured the port as an output. Read operation always returns the data at the signal.

### 21.3.4 GPIO Interrupt Event Register (GPIER)

The GPIO interrupt event register (GPIER), shown in [Figure 21-5](#), carries information of the events that caused an interrupt. Each bit in GPIER, corresponds to an interrupt source. GPIER bits are cleared by writing ones. However, writing zero has no effect.



**Figure 21-5. GPIO Interrupt Event Register (GPIER)**


[Table 21-6](#) defines the bit fields of GPIER.

**Table 21-6. GPIER Bit Settings**

Bits	Name	Description
0–31	$D_n$	Interrupt events. Indicates whether an interrupt event occurred on the corresponding GPIO signal. 0 No interrupt event occurred on the corresponding GPIO signal. 1 Interrupt event occurred on the corresponding GPIO signal.

### 21.3.5 GPIO Interrupt Mask Register (GPIMR)

The GPIO interrupt mask register (GPIMR), shown in [Figure 21-6](#), defines the interrupt masking for the individual ports. When a masked interrupt request occurs, the corresponding GPIER bit is set, regardless of the GPIMR state. When one or more non-masked interrupt events occur, the GPIO module issues an interrupt to the on chip interrupt controller.



**Figure 21-6. GPIO Interrupt Mask Register (GPIMR)**


Table 21-7 defines the bit fields of GPIMR.

**Table 21-7. GPIMR Bit Settings**

Bits	Name	Description
0–31	<i>Dn</i>	Interrupt mask. Indicates whether an interrupt event is masked or not masked. 0 The input interrupt signal is masked (disabled). 1 The input interrupt signal is not masked (enabled).

### 21.3.6 GPIO Interrupt Control Register (GPICR)

The GPIO interrupt control register (GPICR), shown in Figure 21-7, determines whether the corresponding port line asserts an interrupt request on either a high-to-low change or any change on the state of the signal.



**Figure 21-7. GPIO Interrupt Control Register (GPICR)**

Table 21-8 defines the bit fields of GPICR.

**Table 21-8. GPICR Bit Settings**

Bits	Name	Description
0–31	<i>Dn</i>	Edge detection mode. The corresponding port line asserts an interrupt request according to the following: 0 Any change on the state of the port generates an interrupt request. 1 High-to-low change on the port generates an interrupt request.



# Appendix A

## Revision History

This appendix provides a list of the major differences between revisions of the *MPC8313E PowerQUICC II Pro Integrated Processor Family Reference Manual*.

### A.1 Changes From Revision 1 to Revision 2

Major changes to the *MPC8313E PowerQUICC II Pro Integrated Processor Family Reference Manual*, from Revision 1 to Revision 2 are as follows:

Section, Page	Changes
	Throughout book—Combined USB_PLL_GND0 and USB_PLL_GND1 signals into 1 signal: USB_PLL_GND.
	Throughout book—changed all instances of E <sup>2</sup> PROM to EEPROM.
	Throughout book—Added signals THERM0 and THERM1 signals.
	Throughout book—Replaced register name USBGP with CONTROL.
	Throughout book—Added overbars to $\overline{LCSn}$ , $\overline{CAS}$ , and $\overline{RAS}$ ; and deleted overbars from $LGPLn$ .
1.1, 1-3	Replaced text in last sub-bullet under DDR SDRAM memory controller with the following: <ul style="list-style-type: none"><li>— 2.5-V SSTL2 compatible I/O for DDR1, 1.8-V SSTL_18 compatible I/O for DDR2</li></ul>
1.2.6.1, 1-14	Replaced the text in the section with the following: <ul style="list-style-type: none"><li>• Designed to comply with <i>Universal Serial Bus Revision 2.0 Specification</i></li><li>• Supports operation as a stand-alone USB host controller<ul style="list-style-type: none"><li>— Supports USB root hub with one downstream-facing port</li><li>— Enhanced host controller interface (EHCI) compatible</li></ul></li><li>• Supports operation as a stand-alone USB device<ul style="list-style-type: none"><li>— Supports one upstream-facing port</li><li>— Supports three programmable bidirectional USB endpoints</li></ul></li><li>• Supports high-speed (480-Mbps), full-speed (12-Mbps), and low-speed (1.5-Mbps) operations. Low speed is only supported in host mode.</li><li>• Supports USB on-the-go mode when using an external ULPI (UTMI+ low-pin interface) PHY, which includes both device and host functionality</li><li>• On-chip USB-2.0 full-/high-speed PHY with ULPI (UTMI+ low-pin interface) and serial interface</li></ul>

## Revision History

- Supports Wake-on-USB, a method for bringing the device from standby mode to full operating mode
- Host and device support

2.3, 2-2

Deleted Table 2-1 and renumbered rest of tables in section.

2.3, 2-29

In Table 2-2, changed the following row to read:

0x3_1BF8	IP block revision register	R	0x0000_0000 _0002_00A0	14.6.4.6/14-73
----------	----------------------------	---	---------------------------	----------------

Chapter 3, Throughout Updated signals for Rev 2.1 silicon.

3.1, 3-2

In Figure 3-1, added LB\_POR\_CFG\_BOOT\_ECC\_DIS to the TSEC\_MDC signal.

3.1, 3-7

In Table 3-1, changed TSEC\_MDC row and added a footnote as follows:

TSEC_MDC	Ethernet management data clock	Ethernet management	1	O	15-2/15-8	LB_POR_CFG_BOOT_ECC_DIS <sup>1</sup>	—
----------	--------------------------------	---------------------	---	---	-----------	--------------------------------------	---

<sup>1</sup> The LB\_POR\_CFG\_BOOT\_ECC\_DIS function will be selected on the TSEC\_MDC pin whenever  $\overline{\text{HRESET}}$  is asserted; the pin will act as TSEC\_MDC at all other times. The reset block will sample this signal on  $\overline{\text{PORESET}}$  negation only; the sampled value is then passed to the eLBC controller to enable/disable ECC checking during boot time.

3.1, 3-7

In Table 3-1, added the following rows after TSEC\_MDIO:

TSEC_TMR_CLK	1588 external timer reference clock input	eTSEC	1	I	15-2/15-8	UART_SOUT2/LA10	
TSEC_TMR_GCLK	1588 timers	eTSEC	1	O	15-2/15-8	UART_SIN2/LA11	
TSEC_TMR_PP1	1588 timer pulse per period 1	eTSEC	1	O	15-2/15-8	$\overline{\text{UART\_CTS}}[2]/\text{LA12}$	
TSEC_TMR_PP2	1588 timer pulse per period 2	eTSEC	1	O	15-2/15-8	$\overline{\text{UART\_RTS}}[2]/\text{LA13}$	
TSEC_TMR_PP3	1588 timer pulse per period 3	eTSEC	1	O	15-2/15-8	GPIO14/LA9	
TSEC_TMR_TRIG1	1588 external timer trigger input 1	eTSEC	1	I	15-2/15-8	IIC1_SDA/LA14	
TSEC_TMR_TRIG2	1588 external timer trigger input 2	eTSEC	1	I	15-2/15-8	LA7	
TSEC_TMR_ALARM1	1588 timer; current time is equal to or greater than alarm time comparator register	eTSEC	1	O	15-2/15-8	LA8	
TSEC_TMR_ALARM2	1588 timer; current time is equal to or greater than alarm time comparator register	eTSEC	1	O	15-2/15-8	IIC1_SCL/LA15	



3.1, 3-9 In Table 3-1, added the following row after LCLK[0:1]:

LB_POR_CFG_BOOT_ECC_DIS	Boot time ECC checking	eLBC	1	1	—	TSEC_MDC	
-------------------------	------------------------	------	---	---	---	----------	--

Chapter 4 Throughout chapter—Editing done for consistency.

4.1.1, 4-1 In Table 4-1, in the Description column for SRESET, changed the Timing, Negation statement to read: ‘Occurs after being serviced.’

4.3.2.2, 4-14 In Table 4-11, in the Description column for bits 9–11, added a sentence after the first sentence as follows:

This bit combined with bit RLEXT determines where the device boots from.

4.3.2.2.4, 4-17 Second paragraph, first sentence, was changed to read as follows:

The boot ROM location reset configuration word field, shown in Table 4-5, establishes the location of boot ROM.

4.3.2.2.4, 4-17 Combined Table 4-15 and Table 4-16, creating a new Table 4-15 as follows:

**Table 4-15. Boot ROM Location**

RCWHR Bits	Field Name	Value (Binary)	Meaning	
			Legacy Mode (RLEXT = 00)	NAND Flash Mode (RLEXT = 01)
9–11	ROMLOC	000	DDR SDRAM	Reserved
		001	PCI	Local bus NAND Flash—8-bit small page ROM
		010	Reserved, should be cleared	Reserved
		011	Reserved, should be cleared	Reserved
		100	Reserved	Reserved
		101	Local bus GPCM—8-bit ROM	Local bus NAND Flash—8-bit large page ROM
		110	Local bus GPCM—16-bit ROM	Reserved
		111	Reserved	Reserved

Deleted Table 4-16 and renumbered the rest of the tables in the chapter.

4.3.3.1.1, 4-23 In Table 4-23, in the NOR Flash row, changed the setting for BR0[PS] to 10.

4.4.4, 4-32 Second paragraph, last sentence, changed to read as follows:

When using the single crystal option, the frequency for SYS\_CLK\_IN must be chosen such that the USB reference will be 24 or 48 MHz when utilizing the divide by 1 or 2 option, that is, the SYS\_CLK\_IN must be 24 or 48 MHz.

4.4.7, 4-32 Removed Section 4.4.7, “Clock Summary,” and Table 4-28. Renumbered the rest of the tables in the chapter.

4.5.1, 4-33 In Table 4-29, added a ‘Reset’ column.

4.5.1.6, 4-37 In Table 4-33, bit 31, changed the description to read as follows:

Reserved, this bit should never be set.



- 5.3.2.7.1, 5-26      Rewrote the first paragraph as follows:  
The DDR debug configuration enables a DDR memory controller to enter debug mode in which the DDR SDRAM source ID field and data valid strobe are driven onto one of two optional sets of pins:
- 5.3.2.8, 5-27      In Table 5-30, added the following note in the Description column for DDR\_cfg:  
**Note:** DDR\_cfg must be set according to the logical type of the DDR memory devices, as it effects logic behavior of the DDR controller as well as the physical parameters of the DDR I/O pads.
- 5.7.2, 5-50      Changed the 5th thru 8th bullets to read:
- Maximum period of ~412.3 seconds (at 167-MHz bus clock and prescaler = 256) for 16-bit timer
  - Maximum period of ~825 seconds (at 167-MHz bus clock and prescaler = 256) for 32-bit timer
  - Maximum period of thousands of years (at 167-MHz bus clock and prescaler = 256) for 64-bit timer
  - 3-nanosecond timer resolution (at 167-MHz bus clock and no prescaler)
  - Resolution and maximum period can be traded off by selecting prescaler divisor
- 5.7.6.1, 5-62      In the third paragraph, changed the last sentence to read:  
The maximum period (when the reference value is all ones and the prescaler divides by 256) for one 16-bit timer is ~50 ms at 333 MHz.
- 5.8.1, 5-65      Deleted the section number and title only and renumbered the remaining sections.
- 5.8.3.2, 5-67      In Table 5-67, bits 23–30, added the following footnote to each bit description:  
**Note:** This bit will not be affected by the wake-up event if the corresponding mask bit in PMCMR is cleared.  
Table 5-67, bit 31, changed the description to the following:
- |    |      |   |
|----|------|---|
| 31 | PMCI | <p>Power management controller interrupt.<br/>When set, indicates that one of the following events has occurred:</p> <ul style="list-style-type: none"> <li>• One of the unmasked wake-up events (bits 23–30) occurred and PMCCR1[PME_vPEN] is cleared, or</li> <li>• PM current state (as indicated in PMCCR1[CURR_STATE]) is different than PM next state (as written to PCIPMR1[Power_State] and indicated in PMCCR1[NEXT_STATE]) and PMCCR1[USE_STATE] is set, or</li> <li>• CSB platform is in low power mode and a new CSB bus request is detected</li> </ul> <p>If PMCMR[PMCI] is set, the PMC interrupt request to the PowerPC core is driven, causing the PowerPC core to exit its low power state. PMCI can be cleared by writing a 1 to it (writing zero has no effect).</p> |
|----|------|---|
- 5.8.3.4, 5-70      In Figure 5-54, changed the bit field 'NEXT\_STATE' to read only.
- 5.8.3.7.1, 5-79      In the text box of Figure 5-57, changed 'Uninitialized' to 'Not Initialized'.  
In note 6, changed the Word 'uninitialized' to 'non-initialized'.  
In the Notes, changed the second note 3 to 4 and renumbered the rest of the notes.
- 5.8.3.7.2, 5-84      Added the initial two steps.
- 5.8.3.7.3, 5-88      In Figure 5-58, changed the signal name 'PWR\_EN' to 'EXT\_PWR\_CTRL'.  
In Figure 5-59, changed the signal name 'PWR\_EN' to 'EXT\_PWR\_CTRL'.

## Revision History

- 6.2.1, 6-2 In Table 6-2, reserved fields changed from “Write reserved, read = 0” to “Reserved, write should preserve reset value.”
- 8.4.2, 8-6 In table 8-2, signal IRQ[0:40:7], in the Description column (State Meaning), changed the first sentence to read:  
When an external interrupt request signal is asserted, the priority is checked by the IPIC unit, and the interrupt is conditionally passed to the processor.
- 9.4.1.2, 9-10 In Table 9-7, Bits 9–11 row, changed 011–111 Reserved to 011 Reserved.  
In Table 9-7, Bits 13–15 row, changed 011–111 Reserved to 011 Reserved.
- 9.5.5.1, 9-34 In Table 9-26, added the following row:

2 Gbits	256Mbits x 8	15 x 10 x 2	1 Gbytes	2 Gbytes
---------	--------------	-------------	----------	----------

Changed ‘Row x Column x Sub-Bank Bits’ column for the 1 Gbits, 2 Gbits, and 4 Gbits rows to: ... x 2.

- 9.6.1, 9-50 In Table 9-35, for ODT\_PD\_EXIT, changed it to be set to 0001 for DDR1; for FOUR\_ACT, changed it to be set for 00001 for DDR1.
- Chapter 10 Reformatted registers throughout Chapter 10, “Enhanced Local Bus Controller.”  
Modified frequencies 333 and 666 MHz to 33.3 and 66.6 MHz, respectively.  
Added eLBC IP Rev. 1.1 features.  
Removed references to PLL.
- 10.1.3, 10-3 In the first paragraph, second sentence replace with the following:  
The internal transaction address is limited to 32 bits, so all chip selects must fall within the 4-Gbyte window addressed by the internal transaction address. When a memory transaction is dispatched to the eLBC, the internal transaction address is compared with the address information of each bank (chip select).  
In the first paragraph, last sentence, replaced with the following:  
Thus, with the eLBC in GPCM or FCM, or UPM mode, only one of the four chip selects is active at any time for the duration of the transaction except in the case of UPM refresh where all UPM machines that are enabled for refresh have concurrent chip select assertion.
- 10.2, 10-4 In Table 10-2, LGPL0/LFCLE row, changed the first sentence in the State Meaning to read:  
Asserted/Negated—In UPM mode, LGPL0 is one of six general purpose signals; it is driven with a value programmed into the UPM array.  
In the LGPL1/LFALE row, changed the first sentence in the State Meaning to read:  
Asserted/Negated—In UPM mode, LGPL1 is one of six general purpose signals; it is driven with a value programmed into the UPM array.  
In the  $\overline{\text{LOE}}/\overline{\text{LGPL2}}/\overline{\text{LFRE}}$  row, changed the second and third sentences in the State Meaning to read:

- In UPM mode, LGPL2 is one of six general purpose signals; it is driven with a value programmed into the UPM array.  
 $\overline{\text{LFRE}}$  enables data read cycles from NAND Flash EEPROMs controlled by FCM.
- In the LGPL3/ $\overline{\text{LFWP}}$  row, changed the first sentence in the State Meaning to read: Asserted/Negated—In UPM mode, LGPL3 is one of six general purpose signals; it is driven with a value programmed into the UPM array.
- 10.2, 10-4 In Table 10-2, LBCTL signal, changed the first sentence in the signal description to the following:  
 The memory controller activates LBCTL for the local bus when a GPCM-, UPM-, or FCM-controlled bank is accessed.  
 For the LA[0:25] signal, changed the State Meaning to read:  
 Asserted/Negated—LA is the address bus used to transmit addresses to external RAM devices. Refer to Section 10.5, “Initialization/Application Information,” for address signal multiplexing.
- 10.3.1.1, 10-10 In Table 10-4, bits 19–20, replaced the second sentence in the description to:  
 For BR0, PS is configured from the field in reset configuration word as loaded during reset.
- 10.3.1.2, 10-12 Second paragraph, changed to the following:  
 The OR $n$  registers are interpreted differently depending on which of the three machine types is selected for that bank. Because bank 0 can be used to boot, the reset value of OR0 may be different depending on power-on configuration options. Table 10-4 shows the reset values for OR0.
- 10.3.1.2.1, 10-12 Moved Table 10-6 under Section 10.3.1.2, following the second paragraph and renumbered as Table 10-4. Renumbered the following tables.
- 10.3.1.7, 10-24 In the first paragraph, fourth sentence and on, replaced text with:  
 To avoid race conditions between software and a busy eLBC, registers that affect currently running special operation and LSOR must not be re-written before a pending special operation has been completed. The UPM and FCM have different indications of when such special operations are completed. The behavior of eLBC is unpredictable if special operation modes are altered between LSOR being written and the relevant memory controller completing that access.
- 10.3.1.9, 10-26 Replaced the first paragraph, bullets, and second paragraph with the following:  
 The transfer error status register (LTESR) indicates the cause of an error or event. LTESR, shown in Figure 10-13, is a write-1-to-clear register. Reading LTESR occurs normally; however, write operations can clear but not set bits. A bit is cleared whenever the register is written, and the data in the corresponding bit location is a 1. For example, to clear only the write protect error bit (LTESR[WP]) without affecting other LTESR bits, 0x0400\_0000 should be written to the register. After any error/event reported by LTESR, LTEATR[V] must be cleared for LTESR to updated again.

**Revision History**

10.3.1.15, 10-33 In Table 10-22, for bit 0, changed the first sentence in the Description column to read:


This bit should be set when using low bus clock frequencies (see device hardware specifications for applicable frequencies).

In Table 10-22, for bits 27-31, changed the first paragraph to the following:

System clock divider. Sets the frequency ratio between the system clock and the local bus clock. The system clock is equivalent to `csb_clk` or twice `csb_clk` (if `RCWL[LBIUCM]` is set). Only the values shown below are allowed.

10.3.1.16, 10-33 In Table 10-23, bits 16–19, changed the first sentence to read: ‘... high (CW0, CW1, RBW and RSW), FCM ...’.

10.4.2, 10-45 Replaced Figure 10-31 with the following:



10.4.2.1, 10-46 Replaced Table 10-30 with the following:

Option Register Attributes				Signal Timing (LCLK Clock Cycles) <sup>1</sup>				
TRLX	EHTR	XACS	ACS	t <sub>ARCS</sub>	t <sub>CSRP</sub>	t <sub>AOE</sub>	t <sub>OEN</sub>	t <sub>RC</sub>
0	0	0	0X	0	2+SCY	1	0	2+SCY
0	0	0	10	¼ (½)	1¾+SCY (2+SCY)	1	0	2+SCY
0	0	0	11	½	1½+SCY	1	0	2+SCY
0	0	1	0X	0	2+SCY	1	0	2+SCY
0	0	1	10	1	1+SCY	1	0	2+SCY
0	0	1	11	2	1+SCY	2	0	3+SCY
0	1	0	0X	0	2+SCY	1	1	3+SCY
0	1	0	10	¼ (½)	1¾+SCY (1½+SCY)	1	1	3+SCY
0	1	0	11	½	1½+SCY	1	1	3+SCY
0	1	1	0X	0	2+SCY	1	1	3+SCY
0	1	1	10	1	1+SCY	1	1	3+SCY

Option Register Attributes				Signal Timing (LCLK Clock Cycles) <sup>1</sup>				
TRLX	EHTR	XACS	ACS	t <sub>ARCS</sub>	t <sub>CSRP</sub>	t <sub>AOE</sub>	t <sub>OEN</sub>	t <sub>RC</sub>
0	1	1	11	2	1+SCY	2	1	4+SCY
1	0	0	0X	0	2+2×SCY	1	4	6+2×SCY
1	0	0	10	1¼ (1½)	1¾+2×SCY (1½+2×SCY)	2	4	7+2×SCY
1	0	0	11	1½	1½+2×SCY	2	4	7+2×SCY
1	0	1	0X	0	2+2×SCY	1	4	6+2×SCY
1	0	1	10	2	1+2×SCY	2	4	7+2×SCY
1	0	1	11	3	1+2×SCY	3	4	8+2×SCY
1	1	0	0X	0	2+2×SCY	1	8	10+2×SCY
1	1	0	10	1¼ (1½)	1¾+2×SCY (1½+2×SCY)	2	8	11+2×SCY
1	1	0	11	1½	1½+2×SCY	2	8	11+2×SCY
1	1	1	0X	0	2+2×SCY	1	8	10+2×SCY
1	1	1	10	2	1+2×SCY	2	8	11+2×SCY
1	1	1	11	3	1+2×SCY	3	8	12+2×SCY

<sup>1</sup> Times in parentheses are specific for the case LCRR[CLKDIV] = 2; other times apply to all CLKDIV values.

## 10.4.2.2, 10-49

Replace Table 10-31 with the following:

Option Register Attributes				Signal Timing (LCLK Clock Cycles) <sup>1</sup>				
TRLX	XACS	ACS	CSNT	t <sub>AWCS</sub>	t <sub>CSWP</sub>	t <sub>AWE</sub>	t <sub>WEN</sub>	t <sub>wc</sub>
0	0	00	0	0	2+SCY	1	0	2+SCY
0	0	10	0	¼ (½)	1¾+SCY(2 +SCY)	1	0	2+SCY
0	0	11	0	½	1½+SCY	1	0	2+SCY
0	1	00	0	0	2+SCY	1	0	2+SCY
0	1	10	0	1	1+SCY	1	0	2+SCY
0	1	11	0	2	1+SCY	2	0	3+SCY
0	0	00	1	0	2+SCY	1	¼ (0)	2+SCY
0	0	10	1	¼ (½)	1½+SCY	1	0	1¾+SCY (1½+SCY)
0	0	11	1	½	1¼+SCY (1+SCY)	1	0	1¾+SCY (1½+SCY)

Option Register Attributes				Signal Timing (LCLK Clock Cycles) <sup>1</sup>				
TRLX	XACS	ACS	CSNT	t <sub>AWCS</sub>	t <sub>CSWP</sub>	t <sub>AWE</sub>	t <sub>WEN</sub>	t <sub>wc</sub>
0	1	00	1	0	2+SCY	1	¼ (0)	2+SCY
0	1	10	1	1	¾+SCY (½+SCY)	1	0	1¾+SCY (1½+SCY)
0	1	11	1	2	¾+SCY (½+SCY)	2	0	2¾+SCY (2½+SCY)
1	0	00	0	0	2+2×SCY	1	0	2+2×SCY
1	0	10	0	1¼ (1½)	1¾+2×SCY (2+2×SCY)	2	0	3+2×SCY
1	0	11	0	1½	1½+2×SCY	2	0	3+2×SCY
1	1	00	0	0	2+2×SCY	1	0	2+2×SCY
1	1	10	0	2	1+2×SCY	2	0	3+2×SCY
1	1	11	0	3	1+2×SCY	3	0	4+2×SCY
1	0	00	1	0	3+2×SCY	1	1¼ (1)	3+2×SCY
1	0	10	1	1¼ (1½)	1½+2×SCY	2	0	2¾+2×SCY (2½+2×SCY)
1	0	11	1	1½	1¼+2×SCY (1+2×SCY)	2	0	2¾+2×SCY (2½+2×SCY)
1	1	00	1	0	3+2×SCY	1	1¼ (1)	3+2×SCY
1	1	10	1	2	¾+2×SCY (½+2×SCY)	2	0	2¾+2×SCY (2½+2×SCY)
1	1	11	1	3	¾+2×SCY (½+2×SCY)	3	0	3¾+2×SCY (3½+2×SCY)

<sup>1</sup> Times in parentheses are specific for the case LCRR[CLKDIV] = 2; other times apply to all CLKDIV values.

- 10.4.2.3.2, 10-51 In the second paragraph, third sentence, changed to read: ‘OR<sub>n</sub>[CSNT], along with OR<sub>n</sub>[TRLX], control ...’.
- 10.4.3.4.1, 10-70 In Table 10-36, changed the setting for field SCY to read: From por\_cfg\_scy[1:3].
- 10.4.3.4.2, 10-71 Changed step 4, first paragraph, to read:
4. If ECC checking is enabled, the FCM recovers from the spare region the stored ECC for each 512-byte block of boot data. The boot block must be prepared with ECC protection. During ECC generation, software should use FMR[ECCM] = 0 for small-page devices, and FMR[ECCM] = 1 for large-page devices.
- 10.4.4.1.2, 10-74 Second paragraph, replaced with the following:



By default, all local bus refreshes are performed using the refresh pattern of UPMA. This means that if refresh is required, MAMR[RFEN] must be set. It also means that only one refresh routine should be programmed and be placed in UPMA, which serves as the refresh executor. Any banks assigned to a UPM are provided with the common UPMA refresh pattern if the RFEN bit of the corresponding UPM is set, concurrently. UPMA assigned banks, therefore, always receive refresh services when MAMR[RFEN] is set, while UPMB and UPMC assigned banks also receive (the same) refresh services if the corresponding MxMR[RFEN] bits are set. In this scenario, more than one chip select may assert at the same time, as refresh pattern runs for all banks assigned to UPM with RFEN bit set.

## 10.4.4.4.1. 10-78

In Table 10-38, bit 24 and bits 26–27, added the following at the end of each description:

**Note:** AMX must not change values in any RAM word which begins a loop.

In bits 24 and 31, added the following at the end of each description:

In case of UPM writes, program UTA and LAST in same RAM word.

In case of UPM reads, program UTA and LAST in consecutive or same RAM words.

## 10.4.4.4.5, 10-83

In the second paragraph, replace the last sentence and add two bullets as follows:

Also, special care must be taken:

- LAST and LOOP must not be set together.
- Loop start word should not have an AMX change with regard to the previous word.

## 10.4.4.4.7, 10-84

Replaced the first two paragraphs and Table 10-40 with the following:

The address lines can be controlled by the pattern the user provides in the UPM. The address multiplex (AMX) bits in the RAM word can choose between driving the transaction address (AMX = 00), driving it according to the multiplexing specified by the MxMR[AM] field (AMX = 10), or driving the contents of MAR (AMX = 11) on the address signals. In all cases, LA[21:25] of the eLBC are driven by the five lsbs of the address selected by AMX, regardless of whether the next address (NA) bit of the RAM word is used to increment the current address. The effect of NA = 1 is visible only when AMX = 00 chooses the column address.

Table 10-40 shows how the RAM word AMX bits and MxMR[AM] settings can be used to affect row × column address multiplexing on the LA[10:25] signals.

**Table 10-40. UPM Address Multiplexing**

	msb		Internal Transaction Address																												lsb						
	0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20	21	22	23	24	25	26	27	28	29		30	31				
AMX = 10 MxMR[AM] = 000 (Row)									LAD						LA																						
AMX = 00 (Col)																											LA										



In this mode, as soon as UPM samples LUPWAIT negated on the rising edge of the bus clock, it immediately generates an internal transfer acknowledge, which allows a data transfer one bus clock cycle later. The generation of transfer acknowledge is early because LUPWAIT is not re-synchronized. The acknowledge occurs early or normally depending on whether the UPM was already frozen in WAIT cycles or not. This feature allows the synchronous negation of LUPWAIT to affect a data transfer, even if UTA, WAEN, and LAST are set simultaneously.

- 10.4.4.4.9, 10-85 Added paragraph to end of section as follows:  
 In case of UPM writes, program UTA and LAST in same RAM word. In case of UPM reads, program UTA and LAST in consecutive or same RAM words.
- 10.5.1.1, 10-87 In the first paragraph changed the first sentence to read: ‘... on LAD[0:15] (with zero bits on LAD[16:31]) during address ...’.
- 10.5.1.2, 10-87 Changed the title to read, ‘Non-Multiplexed Address and Data Buses’.
- 10.5.1.3, 10-88 Added new Section 10.5.1.3 as follows:

### 10.5.1.3 Multiplexed Address and Data to Save Maximum Pins in 8- to 16-Bit Addressing

With the use of a feature called address byte swap by setting LBCR[ABSWP], data and address muxing can be swapped from the default available. Currently, LAD[0:31] carries A[0:31]. In case of 8-bit interface with 8-bit addressing we do not get benefit of pin reduction with the available muxing. This is because, the MSB of data is muxed with MSB of address. While 8-bit data required is LAD[0:7] and address required is lower order bits A[24:31] we need to pull out all the 16 bits out of the device. For pin limited devices, this feature can be used where LAD[0:7] is mapped to the lsb's of A[24:31] and LAD[8:15] carries **lsb+1** [16:23]. As a result while interfacing with 8 bit address and 8-bit data only LAD[0:7] is suffice with LALE pin. The only drawback of this feature is that it does not support burst as all the address is latched from LAD bus.

- 10.5.1.4, 10-88 Second paragraph, second sentence, changed ‘... 166-MHz bus frequency ...’ to ‘... 133-MHz bus frequency ...’.
- 10.5.4.5, 10-94 After the first paragraph, added the following paragraph:  
 Note that operations specified by OP3 and OP4 (status read) should never be skipped while erasing a NAND Flash device, because, in case that happens, contention may arise on LGPL4. A possible case is that the next transaction from eLBC may try to use that pin as an output and since the NAND Flash device might already be driving it, contention will occur. In case OP3 and OP4 operations are skipped, it may also happen that a new command is issued to the NAND Flash device even when the device has not yet finished processing the previous request. This may also result in unpredictable behavior.
- 10.5.4.6, 10-94 After the first paragraph, added the following paragraph:  
 Note that operations specified by OP5 and OP6 (status read) should never be skipped while programming a NAND Flash device, because, in case that happens,

contention may arise on LGPL4. A possible case is that the next transaction from eLBC may try to use that pin as an output and since the NAND Flash device might already be driving it, contention will occur. In case OP5 and OP6 operations are skipped, it may also happen that a new command is issued to the NAND Flash device even when the device has not yet finished processing the previous request. This may also result in unpredictable behavior.

10.5.6, 10-100

Added the following text at the end of the section:

If a UPM device has  $\overline{OE}$ , it should not be asserted in the same RAM word as the TA signal. If  $\overline{OE}$  and TA are both asserted in the same RAM word, then the eLBC may not be able to sample the correct data during reads. Therefore,  $\overline{OE}$  must be asserted earlier than TA.

10.5.7, 10-102

Deleted Sections, 10.5.7, 10.5.7.1, and 10.5.7.1.1. Deleted Table 10-48 and Figures 10-80 through 10-82.

11.4.1, 11-3

Table 11-2, in TA bit field description, added the following sentence: ‘The translation address must be aligned based on the window’s size.’

12.2, 12-2

Table 12-1, in 0x0\_8030 row, Access column, changed to Mixed.

12.3.1, 12-3

Figure 12-2, changed Access to ‘mixed’.

12.3.2, 12-5

First paragraph, second sentence changed to read:

OMIMR can be read from the CSB or the PCI bus, but it can be written only from the PCI bus.

12.3.8.1, 12-11

In Table 12-10, changed bits 11-10 row to the following:

11–10	PRC	PCI read command. This field indicates the type of PCI read command to use. 00 Reserved 01 PCI read line 10 PCI read multiple 11 Reserved
-------	-----	---

13.3, 13-11

In Tables 13-4 and 13-5 added a ‘Reset’ column.

13.3.2.11, 13-23

In the first paragraph, added the following:

Inbound and outbound windows for the same bus should not overlap. Therefore, situations where an inbound window translation points back into an outbound window, or where an outbound translation window points back into an inbound window, are not allowed.

13.3.2.11, 13-23

In Table 13-18, in the Description column for the TA bit added the following sentence:

The specified address must be aligned to the window size, as defined by  $PIWAR_n[IWS]$ .

13.3.2.12, 13-24

In Table 13-19, in the Description column for the BA bit added the following sentence:

The specified address must be aligned to the window size, as defined by  $PIWAR_n[IWS]$ .

13.3.3, 13-25	In Table 13-22, updated RID value from 10 to 21.
14.2, 14-9	In Table 14-3, changed the ‘Section’ column to ‘Section/Page’ and added the page numbers.
14.6.4.5, 14-74	In the first paragraph, second sentence changed to read: The value of ID is always 0x0000_0000_0002_00A0, indicating that this is the first version of the SEC 2.2.
14.6.4.5, 14-74	In Figure 14-45, changed the Reset value to: 0x0000_0000_0002_00A0.
Chapter 15, 15-1	Throughout the chapter remove text and references describing extraction of data to allocate in the L2 cache. The MPC8313E does not support extraction to L2 cache.
15.4, 15-7	In Table 15-1, for signal TSEC <sub>n</sub> _RX_ER, in the Description column changed to: RGMII, RTBI—Unused. In Table 15-1, added signal TSEC_1588_PP3 as follows:

TSEC_1588_PP3	1588—Pulse out 3 Timer pulse per period 3. It is phase aligned with 1588 timer clock (chip external output pin).	0
---------------	---	---

15.4.1, 15-8	In Table 15-2, TSEC <sub>n</sub> _CRS, State Meaning, changed reference from TSEC <sub>n</sub> _TX_CLK to TSEC <sub>n</sub> _CRS. In TEC_GTX-CLK125, Description column, added the following statement: This input is not used in these modes: <ul style="list-style-type: none"> <li>• RMII</li> <li>• SGMII</li> <li>• MII</li> </ul> In Table 15-2, added Signal TSEC_1588_PP3, as follows:
--------------	--

TSEC_1588_PP3	0	1588 pulse out 3. Timer pulse per period 3. It is phase aligned with 1588 timer clock (chip external output pin)
---------------	---	--

15.5.2, 15-15	In Table 15-4, corrected the values for RQFCR and RQFPR as follows: <i>0xnnnn_nnnn</i>
15.5.3.1.3, 15-24	After item #3, added a fourth sub-bullet as follows: — Special function interrupts are: MSRO, MMRD, and MMRW
15.5.3.1.3, 15-24	In Figure 15-4, changed bit 27 to FGPI. In Table 15-7, add the following for bit 27, FGPI:

27	FGPI	Filer generated general purpose interrupt on a set of filer rule match. This bit will be set upon reception of a frame that matches a GPI rule sequence that is specified in the filer. It is synchronized with the setting of RXF. 0 No filer generated interrupt has occurred. 1 The filer has accepted a frame via a matching rule that the RQFCR[GPI] bit set.
----	------	--

15.5.3.1.4, 15-28	In Figure 15-5, bit 26, changed to FGPIEN. In Table 15-8, changed bits 25–27 to the following:
-------------------	---

## Revision History

25–26	—	Reserved
27	FGPIEN	Filer general purpose interrupt enable

15.5.3.1.6, 15-32 In Table 15-10, changed the description in bit 28 to the following:

28	R100M	RGMII 100 mode. This bit is ignored unless RPM are set and MACCFG2[I/F Mode] is assigned to 10/100 (01). 0 RGMII is in 10 Mbps mode; 1 RGMII is in 100 Mbps mode;
----	-------	---

15.5.3.3.1, 15-49 In Table 15-27, bit 29, in the Description column, added the following sentence:  
Note that frames less than or equal to 16B in length are always silently dropped.

15.5.3.3.7, 15-57 In Figure 15-29, changed bit 0 to GPI and the reset value to ‘undefined’.

In Table 15-33, add the following for bit 0, GPI:

0	GPI	General purpose interrupt. When a property matches the value in the RQPROP entry at this index, and REJ = 0 and AND = 0, the filer will instruct the Rx descriptor controller to set IEVENT[FGPI] when the corresponding receive frame is written to memory. If the timer is enabled (TMR_CTRL[TE] = 1), then TMR_PEVENT[RXP] will also be set.
---	-----	--

15.5.3.3.8, 15-58 In Figure 5-29, changed the reset value to ‘undefined’.

In Table 15-34, bits 16-31, in the Description column, replaced the fourth paragraph with the following:

A value in the length/type field greater than 1500 and less than 1536 is treated as a type encoding by the parser. Since no recognized types exist in this range, the controller will not parse beyond the length/type field of any such frame.

Replaced item 4 and added the following text:

4. The MPLS tagged packets—In this case, one can use arbitrary extraction bytes to compare to the actual ethertype if a filer rule is intending to file based on an MPLS label existence.

**Note:** Users of the eTSEC parser/filer should be aware of a difference in behavior between rev 1 and rev 2 silicon in cases where the Ethernet type/length field contains a value between 1500 and 1536.

In rev 2 silicon, values between 1500 and 1536 are interpreted as a type. Since there are currently no valid types in this range publicly defined by IANA, the controller will not parse beyond the length/type field of any such frame.

If the same packet is encountered with rev 1 silicon, parser/filer behavior is different. With rev 1 silicon, such packets are treated as payload length. S/W must confirm the parser and filer results by checking the type/length field after the packet has been written to memory to see if it falls in this range.

15.5.3.5.1, 15-67 In Table 15-39, bits 26 and 27, add the following statement in the Description column after the first paragraph: Must be 0 if MACCFG2[Full Duplex] = 0.

15.5.3.5.2, 15-67 In Table 15-40, bits 16–19, in the Description column, added the following sentence: ‘Values from 0x3 to 0xF are supported by the controller.’

For bit 29, in the Description column, changed the first paragraph to read:

Pad and append CRC. This bit is cleared by default. This bit must be set when in half-duplex mode (MACCFG2[Full Duplex] is cleared).

- 15.5.3.5.4, 15-70 In Figure 15-39 and Table 15-42, changed the Collision Window to bits 26–31 and bits 20–25 to Reserved.
- 15.5.3.5.5, 15-71 In Table 15-43, for bits 16–31, in the Description column, changed the first paragraph to read:  
By default this field is set to 0x0600 (1536 bytes). It sets the maximum Ethernet frame size in both the transmit and receive directions. (Refer to MACCFG2[Huge Frame].)
- 15.5.3.5.9, 15-74 In Figure 15-44 corrected the register to reflect that this is a write-only register.
- 15.5.3.6.25, 15-90 In Table 15-79, for bits 0–31, changed the description to read:

0–31	TBYT	Transmit byte counter. Increments by the number of bytes that were put on the wire including fragments of frames that were involved with collisions. This count does not include preamble/SFD or jam bytes, except for half-duplex flow control (back-pressure triggered by TCTRL[THDF] = 1). For THDF, the sum total of 'phantom' preamble bytes transmitted for flow control purposes is included in the TBYT increment value of the next frame to be transmitted, up to 65,535 bytes of frame and phantom preamble. Note that the value of TBYT may be greater than the actual number of bytes transmitted if the frame is truncated because it exceeds MAXFRM.
------	------	---

- 15.5.3.6.26, 15-91 In Table 15-80, changed bits 0–31 to 0–9: Reserved and 10–31: TPKT.
- 15.5.3.3.2, 15-106 In Figure 14-103, changed the size of ATTRELI[EI] to 18:25
- 15.5.3.3.2, 15-106 In Table 14-106, in the Bits column, changed the size of ATTRELI[EI] to 18–25 and Reserved to 26–31.
- 15.5.3.9.1, 15-107 Replaced the first paragraph with the following:  
The RQPRM<sub>n</sub> registers specify the minimum number of BDs required to prevent flow control being asserted and the total number of Rx BDs in their respective ring. Whenever the free BD count calculated by the eTSEC for any active ring drops below the value of RQPRM<sub>n</sub>[FBTHR] for that ring, link level flow control will be asserted. Software must not write to RQPRM<sub>n</sub> while LFC is enabled and the eTSEC is actively receiving frames. However, software may modify these registers after disabling LFC by clearing RCTRL[LFC]. Note that packets may be lost due to lack of RxBDs while RCTRL[LFC] is clear. Software can prevent packet loss by manually generating pause frames (via TCTRL[TFC\_PAUSE]) to cover the time when RCTRL[LFC] is clear. Figure 15-104 describes the definition for the RQPRM<sub>n</sub> register.
- 15.5.3.10, 15-108 Added the following paragraph:  
IEEE Standard 1588 compliant timestamping on this device is accomplished using the per-port transmit timestamping registers within each Ethernet controller memory space (See Section 15.5.3.2.11, “Transmit Time Stamp Identification Register (TMR\_TXTS1–2\_ID),” and Section 15.5.3.2.12, “Transmit Time Stamp Register (TMR\_TXTS1–2\_H/L).”) in conjunction with the following common registers, which are located within the memory space for eTSEC1. Because the common IEEE Std. 1588 timestamping registers exist within the eTSEC1 memory space, the eTSEC1 controller must remain enabled in order to use IEEE Std. 1588 timestamping for any Ethernet port.

## Revision History

15.5.3.10.1, 15-108 In Figure 15-106, bit 27, made bit 'Reserved'.

In Table 15-109, in bit 6–15 row, changed to the following:

6–15	TCLK_PERIOD	1588 timer reference clock period. The timer clock counter will increment by TCLK_PERIOD every time the accumulator register overflows. This clock period must be larger than the clock period of the timer reference clock. For applications where user does not want the clock period to be added, they can program this field to 1 to count the clock ticks. This field defaulted to 1 to count overflow ticks. For nanosecond granularity on 1588 timer counter rate, the TCLK_PERIOD should be calculated using the following equation: $TCLK\_PERIOD = 10^9 / \text{Nominal\_Frequency}$
------	-------------	---

In Table 15-109, in Bit 25 row, added the following to the Description column: 'Note that this setting is reserved if CKSEL=01'.

In Bit 26 row, Description column, added the following note:

**Note:** Prior to initiating timer reset (setting TMSR), must gracefully stop receiver (see MACCFG1[RX\_EN] description).

In Bit row 27, deleted DBG and made bit 'Reserved'.

In Bit row 30–31, added the following:

11 RTC clock input.

15.5.3.10.9, 15-115 In Figure 15-112, changed access from read only to read/write.

15.5.3.10.12, 15-117 In Figure 15-115, changed access from mixed to read/write.

15.5.3.10.12, 15-117 In Table 15-122, replaced the Description column with the following:

0–63	ALARM_H/L	Alarm time comparator register. The corresponding alarm event in TMR_TEVENT is set when the current time counter becomes equal to or greater than this alarm time compare value in TMR_ALARM $n$ _L/H. Writing the TMR_ALARM $n$ _L register deactivates the alarm event after it has fired. Writing the TMR_ALARM $n$ _L followed by the TMR_ALARM $n$ _H register rearms the alarm function with the new compare value. The value programmed in this register must be an integer multiple of TMR_CTRL[TCLK_PERIOD] in order to get correct result. This register is reset to all ones to avoid false alarm after reset. In FS mode the alarm trigger is used as an indication to the fiber start down counting. Only alarm 1 supports this mode. In FS mode, alarm polarity bit should be configured to 0 (rising edge).
------	-----------	---

15.5.3.10.16, 15-118 In Figure 15-116, changed access from mixed to read/write.

15.6.2.8, 15-151 In the second paragraph, last sentence, changed to the following:

Only frames addressed specifically to the MAC's station address or a valid multicast or broadcast address can be examined for the Magic Packet sequence.

15.6.2.10, 15-152 The first three sub-bullets were changed to read:

- Receive data frame interrupts, when bits RXB or RXF in IEVENT are set
- Transmit data frame interrupts, when bits TXB or TXF in IEVENT are set
- Error, diagnostic, and special interrupts (all bits in IEVENT other than RXB, RXF, TXB, or TXF)



15.6.2.13, 15-155	In Table 15-152, in the ‘Parser error’ row, deleted the note at the end of the description.
15.6.4.1.1, 15-162	Added the following bullet: <ul style="list-style-type: none"> <li>• The GPI field offers the user the ability to interrupt the core upon matching a rule that causes a frame to be filed to memory. Once the last RxBD corresponding to that frame is written to memory, the IEVENT[FGPI] event will be asserted. This bit will be set regardless of any interrupt coalescing that may be set.</li> </ul>
15.6.4.1.4, 15-164	Added the following two sentences to end of section: <p>A functional interrupt is provided via use of the general purpose interrupt (GPI) bit in the filer table. When a property matches the value in the RQPROP entry at this index, and REJ = 0 and AND = 0, the filer will set IEVENT[FGPI] when the corresponding receive frame is written to memory. This allows the user to set up a filer rule where the core will be interrupted upon the reception of ‘special’ frames.</p> <p>If the timer is enabled (TMR_CTRL[TE] = 1), then the interrupt dedicated for timer events (in addition to the usual receive, transmit and error interrupts) will be asserted.</p>
15.6.5.2.1, 15-171	In the second paragraph, last sentence, changed to the following: <p>As soon as the hardware consumes a BD (by writing it back to memory), RBPTR<sub>n</sub> will advance and the free BD count will reflect the correct number of available free BDs.</p>
15.6.6, 15-171	In the first paragraph, deleted the last sentence. In the second paragraph, first sentence, changed to the following: ‘IEEE 1588 ... nodes to a master clock ... .’ In the third paragraph, first sentence changed to the following: ‘The eTSEC includes a new timer ... .’
15.6.6.2, 15-173	In Figure 15-141, removed the Parser/Data Extraction Logic unit.
15.6.6.4.1, 15-175	Deleted Section 15.6.6.4.1 and renumbered the following section.
15.6.6.4.2, 15,175	Replaced the first paragraph with the following: <p>The eTSEC receive filer has been enhanced with the addition of a general-purpose event bit. This event bit can be used in conjunction with filing table rules to identify 1588 packets and indicate these packets by setting special timer status register bits (TMR_STAT). Additionally, 1588 packets can be easily identified by upper-layer software by using the filer to queue all PTP packets to one or more predefined virtual queues. See Section 15.6.4.1.1, “Filing Rules,” for further information.</p>
15.6.6.5, 15-176	Added to new sections, including figures and tables, as follows:
15.6.7.3, 15-182	In Table 15-163, offset 4–7, bits 0–31, in the Description column, added the following sentence: ‘For best performance, use 64-byte aligned receive buffer pointer addresses’.

## Revision History

- 15.7.1.6, 15-203 Replaced the first paragraph with the following: 15.7.1.6, 15-203 In Table 15-179, removed references to TBICON[Enable Wrap] and TBICON[Comma Detect].
- 15.7.1.6, 15-204 In Table 15-179, changed the following row to read:

Perform an MII Mgmt write cycle to TBI.  
Writing to MII Mgmt Control with 16-bit data intended for TBI's AN Advertisement register,  
MIIMCON[0000\_0000\_0000\_0000\_0000\_0001\_1010\_0000]  
This advertises to the Link Partner that the TBI supports PAUSE and Full Duplex mode and does not support Half Duplex mode.

- 16.3, 16-5 In Table 16-3, Section/Page column, corrected page numbers.
- 16.3, 16-5 In Table 16-3, added the following row:

0x2_3170	ULPI VIEWPORT—ULPI Register Access	Mixed	0x0000_0000	16.3.2.12/16-24
----------	------------------------------------	-------	-------------	-----------------

- 16.3, 16-5 In Table 16-3, changed the reset value for the following addresses:

Address Offset	Value
0x2_3104	0x0001_0011
0x2_3124	0x0000_0183
0x2_3140	0x0008_0000
0x2_3164	0x0000_0000
0x2_3184	0x1000_0000
0x2_31A4	0x0000_0C20

- 16.3.1.3, 16-9 In Table 16-6, “HCSPARAMS Register Field Descriptions,” removed “The reset value of this field is always 0 after the USBDR controller is configured as a host by writing 0x3 to USBMODE; else, the reset value is always 1” from HCSPARAMS[N\_PCC] bit field description.
- 16.3.2.12, 16-23 Added the following new section, figure, and table. Renumbered the following sections, figures, and tables.

### 16.3.2.12 ULPI Register Access (ULPI VIEWPORT)

The register provides indirect access to the ULPI PHY register set. Although the controller modules perform access to the ULPI PHY register set, there may be extraordinary circumstances where software may need direct access. Be advised that writes to the ULPI through the ULPI viewport can substantially harm standard USB operations. Currently no usage model has been defined where software should need to execute writes directly to the ULPI. Note that executing read operations through the ULPI viewport should have no harmful side effects to standard USB operations. Also note that if the ULPI interface is not enabled, this register will always read zeros.



operations will not be able execute. Undefined behavior results if a read or write operation is performed when ULPISS is cleared. To execute a wakeup operation, write all 32-bits of the ULPI Viewport where ULPIPORT is constructed appropriately and the ULPIWU bit is set and the ULPIRUN bit is cleared. Poll the ULPI Viewport until ULPIWU is cleared for the operation to complete.

To execute a read or write operation, write all 32-bits of the ULPI Viewport where ULPIDATWR, ULPIADDR, ULPIPORT, ULPIRW are constructed appropriately and the ULPIRUN bit is set. Poll the ULPI Viewport until ULPIRUN is cleared for the operation to complete. For read operations, ULPIDATRD is valid once ULPIRUN is cleared.

The polling method above can be replaced with interrupts using the ULPI interrupt defined in the USBSTS and USBINTR registers. When a wakeup or read/write operation completes, the ULPI interrupt is set.

16.3.2.13, 16-25 In Table 16-22, changed the description for bit 23 with the following:

23	PHCD	<p>PHY low power suspend. This bit is not defined in the EHCI specification.</p> <p>Host mode:</p> <ul style="list-style-type: none"> <li>The PHY can be put into low power suspend —when the downstream device has been put into suspend mode or when no downstream device is connected. Low power suspend is completely under the control of software.</li> </ul> <p>Device mode:</p> <ul style="list-style-type: none"> <li>The PHY can be put into low power suspend—when the device is not running (USBCMD[RS] = 0b) or suspend signaling is detected on the USB. Low power suspend will be cleared automatically when the resume signaling has been detected or when forcing port resume.</li> </ul> <p>0 Normal PHY operation. 1 Signal the PHY to enter low power suspend mode</p> <p>Reading this bit indicates the status of the PHY.</p> <p><b>Note:</b> If there is no clock connected to the USBDR_CLK signals, PHCD must be set and the following registers should not be written: DEVICE_ADDR/PERIODICLISTBASE, PORTSC, ENDPTCTRL0, ENDPTCTRL1, ENDPTCTRL2.</p>
----	------	--

16.3.2.15, 16-31 In Table 16-24, “OTGSC Register Field Descriptions,” corrected bit field description of field 7–5 from “Reserved, should be cleared” to “Reserved, writes should preserve reset value.”

16.3.2.24, 16-38 In fourth paragraph, second sentence, changed as follows:  
If AGE\_CNT\_THRESH is equal to zero, priority state one is always chosen.

16.3.2.26, 16-40 In Table 16-35, bit 31, changed the description to read as follows:

31	rd_prefetch_val	<p>Selects whether 32 bytes or 64 bytes are fetched during burst read transactions at the system interface. When this input is LOW 64 bytes are fetched and when it is HIGH 32 bytes are fetched. The setting of rd_prefetch_val must match the setting of the larger of TXPBURST and RXPBURST fields in the BURSTSIZE register. If either of these fields is 64 bytes, then rd_prefetch_val must be left cleared. Otherwise, this value should be set.</p> <p>0 64-byte fetch 1 32-byte fetch</p>
----	-----------------	--

16.3.2.27, 16-41 In Table 16-36, bit 29, changed the description to read as follows:

29	USB_EN	<p>UTMI mode: This bit is used to enable the USB interface. It must be set before setting RS bit in USB CMD register.</p> <p>1 Enable 0 Disable</p> <p>ULPI mode: In safe mode, all USB interface signals are put into input mode or driven inactive, except for SUSPEND_STP which is driven high. Also, the input signal DIR is forced to appear high to the controller. This prevents any start-up problems that otherwise could occur if the PHY and the controller take significantly different times to complete power-on reset.</p> <p>1 Normal operation 0 Safe mode</p>
----	--------	---

16.5.6, 16-59 In Figure 16-40, for RL, changed to bits 31–28, C to bit 27, and Maximum Packet Length to bits 26–16.

16.6.1, 16-65 In first paragraph, first sentence, removed ‘as illustrated in Table 25’. Replaced the second with the following:  
After a hardware reset, only the operational registers will be at their default values.  
Deleted Table 16-63 and renumbered the rest of the tables in the chapter.

17.3.1.5, 17-9 Table 17-8, changed the last sentence in the DATA description to read:  
Note that in both master receive and slave receive modes, the very first read is always a dummy read.

17.5.5, 17-22 In second paragraph, removed the the sentence: ‘For 1-byte transfers, a dummy read should be performed by the interrupt service routine (see Figure 17-11).’

18.3.1.3, 18-7 In Table 18-8, replaced the table with 133 and 167 MHz information.

18.3.1.3, 18-8 In the last paragraph, changed item 1 as follows:  
1. The input clock frequency (ICF) is divided by the actual frequency input (AFI) to get the correct divisor value ( $ICF/AFI$ , where  $AFI = \text{baud rate} \times 16 \times \text{divisor}$ ).

## A.2 Changes From Revision 0 to Revision 1

Major changes to the *MPC8313E PowerQUICC II Pro Integrated Processor Family Reference Manual*, from Revision 0 to Revision 1 are as follows:

Section, Page	Changes
	Throughout book, removed ‘MPC8313E and MPC8313 specific,’ text. Removed LSYNC_OUT and LSYNC_IN.
	Throughout book—removed internal signal LBC_PM_REF_10 references.
	Throughout book—fixed the direction of signals in figures and tables.
1.1, 1-4	Throughout chapter, replaced ‘On-chip USB-2.0 full-speed/high-speed PHY with ULPI (UTMI + low-pin interface),’ with the following: On-chip USB-2.0 full-speed/high-speed PHY with UTMI

## Revision History

- 1.1, 1-6 Replaced ‘Supports wake-up from Ethernet Magic Packet, USB, GPIO, and PCI (PME input as host),’ with the following:  
Supports wake-up from Ethernet Magic Packet, USB, GPIO, PCI (PME input as host), timer, and external interrupts
- 1.2.6, 1-14 Removed Section 1.2.6, “Serial ATA (SATA) Controller.”
- 2.3, 2-5 In Table 2-2, changed the reset value of SWSRR from 0x0000\_0000 to 0x0000, as it is a 16-bit register.
- 3.1, 3-1 Throughout book—added signals: USB\_PHY\_PWR, USB\_PHY\_GND, USB\_VDDA, USB\_VSSA, USB\_PLL\_PWR1, USB\_PLL\_PWR3, USB\_PLL\_GND0, USB\_PLL\_GND1, USB\_VSSA\_BIAS, and USB\_VDDA\_BIAS.  
Removed signals: SD\_PLL\_TPA\_ANA and SD\_PLL\_TPD.
- 3.1, 3-20 Removed Table 3-3, because of inconsistencies between MPC8313 hardware specification document and the reference manual.
- 4.3.1.1, 4-10 In Table 4-5, for offset 0100, changed PCI frequency range from 25–66.666 MHz to 24–66.666 MHz.
- 4.3.1.3, 4-11 In Table 4-6, modified the description for CFG\_CLKIN\_DIV.
- 4.3.2, 14-13 In Table 4-8, added more description to DDRCM bit.
- 4.3.3.3.2, 4-27 Table 4-25, for bits 1, 3, 12–15, and 20–27, removed ‘Reserved, should be cleared,’ and added ‘Reserved.’ Also reformatted the table.
- 4.4.3, 4-31 Removed overbars for the signal CFG\_CLKIN\_DIV.
- 4.4.5, 4-33 Modified Ethernet clocking description.
- 4.4.7, 4-34 Deleted Figure 4-8 and replaced it with Table 4-8, “System Clock Frequencies.”
- 4.5.2.3, 4-42 Table 4-37, changed the description for [4:5], [12:14], and [16:31] by removing ‘should be cleared.’
- 4.5.2.3, 4-41 Table 4-38, in ENCCM description, replaced ‘Encryption core clock mode,’ with ‘Encryption core and I2C1 clock mode.’
- 5.3.2.5, 5-21 In Table 5-27, replaced incomplete set of 1588 signals (TSEC\_TMR\_CLK, TSEC\_TMR\_CLK\_OUT, TSEC\_TMR\_TRIG\_IN, TSEC\_TMR\_TRIG\_OUT, TSEC\_TMR\_PULSE\_OUT1, TSEC\_TMR\_PULSE\_OUT2), with the complete set ( TSEC\_1588\_CLK, TSEC\_1588\_GCLK, TSEC\_1588\_TRIG1, TSEC\_1588\_TRIG2, TSEC\_1588\_PP1, TSEC\_1588\_PP2, TSEC\_1588\_PP3, TSEC\_1588\_ALARM1, TSEC\_1588\_ALARM2).
- 5.3.2.6, 5-25 In Table 5-29 added footnote for bit SICRH[30]:  
If RCWH[ETSEC1M] is RGMII/RTBI then the reset value is 1, otherwise it is 0.
- 5.3.2.8, 5-26 In Table 5-30, rewrote DSO\_EN and Q\_DRN description:  
DDR driver software override is disable when DSO\_EN is 0 and is enable when DSO\_EN is 1. Drain queue before sleep is disable when Q\_DRN is 0 and enable when Q\_DRN is 1.

5.4.3, 5-28	Under the bullet, 'WDT reset/interrupt output mode,' replaced 'soft reset' with 'hard reset.'
5.4.4, 5-29	In Table 5-32, changed the reset value of SWSRR to 0x0000, as it is a 16-bit register.
5.7.1, 5-48	Throughout this section, removed references to internal clock signal ipg_clock.
5.8.4.7.2, 5-84	Modified item 8, by adding 'PMC interrupt clearing sequence.'
7.3.1, 7-15	In Figure 7-2, removed Performance Monitor section (it is not available for the e300c2 and earlier core).
7.3.4.2, 7-31	Table 7-7, under System Reset, the description should read: Caused by the assertion of $\overline{hreset}$
9.4.1.6, 9-17	In Table 9-11, for CPO, changed 11111 to Reserved, and removed Automatic calibration.
9.4.1.6, 9-17	Fixed spacing on bit settings of FOUR_ACT.
9.4.1.6, 9-17	Removed text from description of TIMING_CFG_2[FOUR_ACT]: 'This field is concatenated with TIMING_CFG_3[EXT_FOUR_ACT] to obtain a 5-bit value for the total activate to precharge time. Note that the decode of 000000-000011 is equal to 16-19 clocks when TIMING_CFG_3[EXT_ACTTOPRE] = 0, but it is equal to 03 clocks when TIMING_CFG_3[EXT_ACTTOPRE] = 1.'
9.5, 9-33	Modified Figure 9-22 to correctly represent the DDR SDRAM configuration.
9.5.1.1, 9-34	Fixed cross-references to the DDR SDRAM tables.
9.5.1.1, 9-35	In Tables 9-25 and 9-26, fixed the row for 2-Gbit and 4-Gbit memory configurations.
9.5.1.1, 9-35	In Table 9-26, added a new memory configuration for the 4-Gbits row as follows: 4 Gbits   512 Mbits × 8   15 × 11 × 3   2 Gbytes   4 Gbytes
10.1.2, 10-2	Modified the feature list by removing the, 'Up to 256-byte bursts, arbitrarily aligned,' statement.
10.2, 10-4	In Table 10-1, for signal LAD, changed the number of signals from 32 to 16.
10.2, 10-5	In Table 10-2, removed text from the LALE state meaning description.
10.2, 10-10	In Table 10-3, changed footnote to read: FMR[BOOT] is set during reset if RCWH[RLEXT] selects FCM as the boot controller.
10.3.1.1, 10-11	In Figure 10-2, modified the first sentence of the footnote to read: BR0 has its valid bit (V) set for RCWH[ROMLOC] = LBC.
10.3.1.2.2, 10-13	Added a new table, Table 10-6. Reset Value of OR0 Register, and renumbered the remaining tables.
10.3.1.2.3, 10-16	In Figure 10-3, filled in value for TBD, 'bit P is configured from the value of RCWH[ROMLOC].' Removed field SSS from reset and replaced it with 010 and reference in footnote.

## Revision History

10.3.1.9, 10-26	Added the following statement at the end of the current description: After any error/event reported by LTESR, LTEATR[V] must be cleared for LTESR to update again.
10.3.1.12, 10-31	Added LTESR in the list of registers to be cleared by software once LTEATR[V] is set.
10.3.1.15, 10-34	In Table 10-21, for bit LCRR[PBYP, removed ‘if the PLL is unable to lock. Clarified for bit LCRR[EADC], LCLK.’ Added line to LCRR[CLKDIV], ‘The system clock can be equal to csb_clk or twice csb_clk ( if RCWL[LBIUCM] is set.’
10.4.2.5, 10-58	In Table 10-31, Register BR0, corrected the source to be from RCWH[ROMLOC], not POR cfg pin.
10.4.3.4.1, 10-71	In Table 10-35, Register OR0, corrected the source to be from RCWH[ROMLOC], not POR cfg pin. Also corrected the setting of OR0 SCY to 010 from 011.
10.4.3.4.3, 10-74	Removed this section.
12.3.8.1, 12-11	In Table 12-10, added the following note to fields DAHE and SAHE: The DMA does not support address hold when the external trigger mode is selected (EMSEN = 1).
13.3.3.13, 13-33	Modified PIMMR from 28 to 12 writeable bits.
13.3.3.5, 13-30	In Table 13-27 replaced the chip revision ID with 8’h10.
14.1, 14-1	Removed SEC frequency information from the bulleted list of features in the introduction.
14.3.2, 14-12	Second paragraph, last line should read, ‘and described in Table 14-6.’
14.3.2.2, 14-15	In Table 14-7, added a note to AESU CTR non-snooping.
14.3.4, 14-17	Third paragraph, replaced ‘36-bit’ with ‘32-bit.’
14.3.4, 14-18	In Table 14-10, changed bits 24–31 to ‘Reserved.’ In second paragraph, replaced reference to Figure 14-7 with Figure 14-6.
14.4, 14-20	Replaced ‘are used in SEC,’ with ‘are used in SEC 2.2.’
14.4.2.11, 14-40	Figure 14-25, added SHA-224 to the Name column, between SHA-1 and SHA-256.
14.4.3.1, 14-41	Change the third paragraph to read: Table 14-25 describes AESUMR fields.
14.4.3.6, 14-47	Change the second paragraph to read: Table 14-29 describes AESUISR fields.
14.5.1.2, 14-62	Table 14-35 changed the title of the table to read: ‘Crypto-Channel Pointer Status Register Error Field Definitions.’
14.5.1.2, 14-63	Change the first paragraph after Table 14-36 with the following: Table 14-36 shows the possible values of the PAIR_PTR field in the CCPSR.



14.5.1.2, 14-63	Table 14-36 changed the title of the table to read: ‘Crypto-Channel Pointer Status Register PAIR_PTR Field Values.’
14.5.2.1, 14-65	Replaced ‘channel configuration register NT and CDIE bits in the CCR,’ with ‘crypto-channel configuration register NT and CDIE bits in the CCCR.’
15.2, 15-2	Added ‘and RTBI’ to physical interface sub-bullet, ‘1000 Mbps full-duplex RGMII and RTBI.’
15.5.2, 15-12	Removed register sections (and memory map rows) from MPC8313 product which only have a 32-bit address space.
15.5.3.1.1, 15-21	In Figure 15-2, corrected value of TSEC_ID[TSEC_REV_MN] to be 00.
15.5.3.1.1, 15-22	In Table 15-5, corrected value of TSEC_ID[TSEC_REV_MN] to be 00.
15.5.3.1.3, 15-24	Updated register to show that they are w1c (all fields).
15.5.3.1.6, 15-30	Updated cross-references.
15.5.3.1.6, 15-31	In Table 15-10, changed bit 16 to ‘Reserved’ and removed the rest of the text.
15.5.3.1.6, 15-32	In Table 15-11, changed TBIM, SGMII fields from 0 to 1.
15.5.3.1.8, 15-34	In Table 15-13, changed the description of DMACTRL[GTS] to the following: If this bit is set, the Ethernet controller stops transmission after all frames that are currently in the Tx FIFO or scheduled have been transmitted, and the GTSC interrupt in the IEVENT register is asserted. A frame that has started reading buffer descriptors or data from memory will be read to completion and be transmitted before the GTSC interrupt occurs. However, if no frame has been scheduled for transmission and the Tx FIFO is empty, the GTSC interrupt is asserted immediately. Once transmission has completed, clearing GTS will restart transmit.
15.5.3.3.2, 15-40	Updated register to show that they are w1c (all fields).
15.5.3.3.3, 15-42	In Table 15-21, add the following note to DFVLAN[TAG]: Note that, if using DFVLAN to set a custom ethertype (that is, using a value other than 0x8100), packets received with a custom tag are not counted by any of the RMON counters. Affected counters include TRMGV, RMCA, RBCA, RXCF, RXPF, RXUO, RALN, RFLR, ROVR, RJBR, TMCA, TBCA, TXPF, TXCF.
15.5.3.3.8, 15-46	Removed register sections (and memory map rows) from MPC8313 product which only have a 32-bit address space.
15.5.3.4.1, 15-50	In Table 15-32, add the following note to RCTRL[VLEX] and RCTRL[PRSDEP]: If PRSDEP is cleared, VLEX must be cleared as well. (VLAN tag extraction is only supported when the parser is enabled.)
15.5.3.4.1, 15-50	In Table 15-32, add the following note to RCTRL[FILREN] and RCTRL[PRSDEP]: If PRSDEP is cleared, FILREN must be cleared as well.
15.5.3.4.1, 15-50	Changed first sentence of RCTRL[PRSDEP] setting 00 to;

	Parser disabled. Receive frame filer must also be disabled by clearing RCTRL[FILREN].
15.5.3.4.1, 15-50	Changed RCTRL[FILREN] description (but not bit settings) to the following: Filer enable, when set, the receive frame filer is enabled. This will file accepted frames to a particular RxBD ring according to rules defined in the filer table. In this case, PRSDEP must not be cleared.
15.5.3.4.2, 15-51	Updated register to show that they are w1c (all fields).
15.5.3.4.6, 15-57	Added note to RQFPR[ETY].
15.5.3.4.8, 15-59	In Table 15-39, added the following note for bit TOS: IPv6 the Traffic Class field is extracted using the IP header definition in RFC 2460. IPv6 headers formed using the earlier RFC 1883 have a different format and must be handled with software.
15.5.3.4.12, 15-63	Removed register sections (and memory map rows) from MPC8313 product which only have a 32-bit address space.Using the filer to match ETY does not work in the case of PPPoe packets, beacuse the PPPoe ethertype in the original packet, 0x8864, is always overwritten with the PPP protocol field. thus, matches on ETY == 0x8864 always fail, Instead, software should use PID=1 fields IP4(ETy = 0x0021) and IP6 (ETy = 0x0057) to distinguish PPPoe session packets carrying IPv4 and IPv6 datagrams. Other PPP protocols are encoded in the ETY field, but many of them overlap with real ethertype definitions. Consult IANA and IEEE for possible ambiguities.
15.5.3.6.2, 15-69	In Table 15-47, added cross-reference to buffer descriptors section to description of MACCFG2[Huge Frame].
15.5.3.6.5, 15-73	Added the following note to section, Maximum Frame Length Register (MAXFRM): If MACCFG[Huge Frame] = 0, the value of this field must be less than or equal to MRBLR[MRBL] x (minimum number of RxBDs per ring). See Section 15.5.3.6.2, "MAC Configuration 2 Register (MACCFG2)," Section 15.5.3.4.9, "Maximum Receive Buffer length Register (MRBLR)," and Section 15.6.6.3, "Receive Buffer Descriptors (RXBD)."
15.5.3.6.6, 15-73	Added a note regarding eTSEC system clock to MgmtClk bit description.
15.5.3.7.19, 15-89	Modified RCSE bit description.
15.5.3.7.25, 15-92	In Table 15-86, added the following note to TBYT[TBYT]: The value of TBYT may be greater than the actual number of bytes transmitted if the frame is truncated because it exceeds MAXFRM.
15.5.3.9, 15-107	Removed text related to L2 cache.
15.6.2.10.3, 15-157	In Table 15-149, "Interrupt Coalescing Timing Threshold Ranges," replaced eTSEC frequencies of 266 and 333 MHz with 133 and 166 MHz information/rows.

15.6.5.2, 15-170	Modified Figure 15-145 and changed the name from ‘Alcatel MAC’ to ‘Ethernet MAC.’
15.6.6.2, 15-178	In Table 15-164, modified bit TR to clarify when truncation occurs on transmit.
115.6.6.3, 15-182	In Table 15-165, corrected description of RxBD TR (truncation) field to state that TR can also be set if a frame length equal to the maximum frame length is received. Modified descriptions of bit LG.
15.7.1.5, 15-199	In Table 15-178, changed SerDes and SGMII signal frequency from 625 MHz to 1250 MHz, TXD and RXD to TXDp/TXDn and RXDp/RXDn, respectively.
Chapter 16	Throughout this chapter changed references from endpoint 5 to endpoint 2 and removed offsets for endpoints 3, 4, and 5.
16.1, 16-1	Changed ‘Access in Memory Map’ to match register figure; USBSTS, PORTSC, OTGSC, ENDRPT Complete, and ENDPTCTRL. In figure, changed access to ‘Mixed from R/W,’ in memory map, changed reset values to match register figures: USBCMD, FRIDEX, and PERIODICLISTBASE.
16.1, 16-2	In Figure 16-1, updated diagram.
16.2, 16-3	Table 16-1: removed mentioning of UTMI signals, as they are internal signals; added UTMI PHY external signals.
16.2.2, 16-4	Removed the note, ‘The ULPI signals are multiplexed with UTMI interface.’
16.3, 16-7	Removed ENDPTCTRL3–5 registers.
16.3.1.6, 16-11	In Figure 16-7, changed DEN reset value to 0011.
16.3.2.13, 16.25	Corrected bit description by switching PORTSC[LS] bit 01 (J-state) with 10 (K-state). Also modified the order to 00,10,01,11.
16.3.2.13, 16-27	In Table 16-22, for bits 11-10, modified the order of 00,01,10,11 to 00,10,01,11. Changed PORTSC[PTX] to ‘Reserved.’
16.3.2.27, 16-41	In Table 16-36, added the following note for bits 30 and 31: ‘PORTSC[PHCD] bit must be set.’ Also for RefSel[1:0] changed bits 01 to 24 MHz instead of 16 MHz and changed 00 to ‘Reserved’ instead of 12 MHz.
16.10, 16-155	In Table 16-98, replaced all TBDs with actual values.
17.3.1.2, 17-5	In Table 17-5, added the following notes: <ul style="list-style-type: none"> <li>1. The values shown in the table are applicable only for the default value of DFSRR, refer to AN2919.</li> <li>2. I<sup>2</sup>C controller clock of I2C1 is derived from csb_clk/SCCR[SDHCCM].</li> </ul>
17.3.1.2, 17-7	Replaced the clock ratio for controller clock : csb clock to 1:1.
17.3.1.3, 17-7	In Table 17-6, replaced TSEC2CM with ENCCM and stated that I2C1 is controllable through this, and clock ratios for I2C2 are not controllable and is always 1:1 with CSB.
20.1, 20-1	Removed the pull-up register for TCK.
20.2.1, 20-2	In Table 20-1, removed reset values for all input signals.

## Revision History

20.2.1, 20-2	Removed the statement, 'An unterminated input appears as a high signal level to the test logic due to an internal pull up resistor,' from TCK description.
Chapter 21	Throughout this chapter replaced IPIC with GPIO.
21.1, 21-1	Removed the section before Introduction.
21.3, 21-2	Replaced the statement, 'All addresses used in this chapter are offsets from the GPIO starting address,' with the following: All addresses used in this chapter are offsets from the address held in IMMRBAR.

# Glossary

The glossary contains an alphabetical list of terms, phrases, and abbreviations used in this reference manual.

---

## A

**Architecture.** A detailed specification of requirements for a processor or computer system. It does not specify details of how the processor or computer system must be implemented; instead it provides a template for a family of compatible *implementations*.

**Atomic access.** A bus access that attempts to be part of a read-write operation to the same address uninterrupted by any other access to that address (the term refers to the fact that the transactions are indivisible). The Power Architecture technology implements atomic accesses through the **lwarx/stwcx** instruction pair.

**Autobaud.** The process of determining a serial data rate by timing the width of a single bit.

---

## B

**Beat.** A single state on the bus interface that may extend across multiple bus cycles. A transaction can be composed of multiple address or data *beats*.

**Big-endian.** A byte-ordering method in memory where the address *n* of a word corresponds to the *most-significant byte*. In an addressed memory word, the bytes are ordered (left to right) 0, 1, 2, 3, with 0 being the *most-significant byte*. See *Little-endian*.

**Boundedly undefined.** A characteristic of certain operation results that are not rigidly prescribed by the Power Architecture technology. Boundedly-undefined results for a given operation may vary among implementations and between execution attempts in the same implementation.

Although the architecture does not prescribe the exact behavior for when results are allowed to be boundedly undefined, the results of executing instructions in contexts where results are allowed to be boundedly undefined are constrained to ones that could have been achieved by executing an arbitrary sequence of defined instructions, in valid form, starting in the state the machine was in before attempting to execute the given instruction.

**Breakpoint.** A programmable event that forces the core to take a breakpoint exception.

**Burst.** A multiple-beat data transfer whose total size is typically equal to a cache block.

**Bus clock.** Clock that causes the bus state transitions.

**Bus master.** The owner of the address or data bus; the device that initiates or requests the transaction.

## C

**Cache.** High-speed memory containing recently accessed data or instructions (subset of main memory).

**Cache block.** A small region of contiguous memory that is copied from memory into a *cache*. The size of a cache block may vary among processors; the maximum block size is one *page*. In Power Architecture processors, *cache coherency* is maintained on a cache-block basis. Note that the term ‘cache block’ is often used interchangeably with ‘cache line.’

**Cache coherency.** An attribute wherein an accurate and common view of memory is provided to all devices that share the same memory system. Caches are coherent if a processor performing a read from its cache is supplied with data corresponding to the most recent value written to memory or to another processor’s cache.

**Cache flush.** An operation that removes from a cache any data from a specified address range. This operation ensures that any modified data within the specified address range is written back to main memory. This operation is generated typically by a Data Cache Block Flush (**dcbf**) instruction.

**Caching-inhibited.** A memory update policy in which the *cache* is bypassed and the load or store is performed to or from main memory.

**Cast out.** A *cache block* that must be written to memory when a cache miss causes a cache block to be replaced.

**Changed bit.** One of two *page history bits* found in each *page table entry* (PTE). The processor sets the changed bit if any store is performed into the *page*. See also *Page access history bits* and *Referenced bit*.

**Clean.** An operation that causes a cache block to be written to memory, if modified, and then left in a valid, unmodified state in the cache.

**Clear.** To cause a bit or bit field to register a value of zero. See also *Set*.

**Context synchronization.** An operation that ensures that all instructions in execution complete past the point where they can produce an *exception*, that all instructions in execution complete in the context in which they began execution, and that all subsequent instructions are *fetched* and executed in the new context. Context synchronization may result from executing specific instructions (such as **isync** or **rfi**) or when certain events occur (such as an exception).

**Copy-back operation.** A cache operation in which a cache line is copied back to memory to enforce cache coherency. Copy-back operations consist of snoop push-out operations and cache cast-out operations.

- 
- D**
- Direct-mapped cache.** A cache in which each main memory address can appear in only one location within the cache; operates more quickly when the memory request is a cache hit.
- Double data rate.** Memory that allows data transfers at the start and end of a clock cycle, thereby doubling the data rate.
- 
- E**
- Effective address (EA).** The 32-bit address specified for a load, store, or an instruction fetch. This address is then submitted to the MMU for translation to either a *physical memory* address or an I/O address.
- Exclusive state.** MEI state (E) in which only one caching device contains data that is also in system memory.
- 
- F**
- Fetch.** Retrieving instructions from either the cache or main memory and placing them into the instruction queue.
- Flush.** An operation that causes a cache block to be invalidated and the data, if modified, to be written to memory.
- Frame-check sequence (FCS).** Specifies the standard 32-bit cyclic redundancy check (CRC) obtained using the standard CCITT-CRC polynomial on all fields except the preamble, SFD, and CRC.
- 
- G**
- General-purpose register (GPR).** Any of the 32 registers in the general-purpose register file. These registers provide the source operands and destination results for all integer data manipulation instructions. Integer load instructions move data from memory to GPRs and store instructions move data from GPRs to memory.
- Guarded.** The guarded attribute pertains to out-of-order execution. When a page is designated as guarded, instructions and data cannot be accessed out-of-order.
- 
- H**
- Harvard architecture.** An architectural model featuring separate caches and other memory management resources for instructions and data.
- 
- I**
- Illegal instructions.** A class of instructions that are not implemented for a particular processor. These include instructions not defined by the architecture. In addition, for 32-bit implementations, instructions that are defined only for 64-bit implementations are considered to be illegal instructions. For 64-bit implementations instructions that are defined only for 32-bit implementations are considered to be illegal instructions.
- Implementation.** A particular processor that conforms to the architecture, but may differ from other architecture-compliant implementations for example in design, feature set, and implementation of *optional* features.

**Inbound ATMU windows.** Mappings that perform address translation from the external address space to the local address space, attach attributes and transaction types to the transaction, and map the transaction to its target interface.

**In-order.** An aspect of an operation that adheres to a sequential model. An operation is said to be performed in-order if, at the time that it is performed, it is known to be required by the sequential execution model.

**Integer unit.** An execution unit in the core responsible for executing integer instructions.

**Inter-packet gap.** The gap between the end of one Ethernet packet and the beginning of the next transmitted packet.

**Instruction latency.** The total number of clock cycles necessary to execute an instruction and make ready the results of that instruction.

---

**K** **Kill.** An operation that causes a *cache block* to be invalidated without writing any modified data to memory.

---

**L** **L2 cache.** Level-2 cache. See *Secondary cache*.

**Latency.** The number of clock cycles necessary to execute an instruction and make ready the results of that execution for a subsequent instruction.

**Least-significant bit (lsb).** The bit of least value in an address, register, field, data element, or instruction encoding.

**Least-significant byte (LSB).** The byte of least value in an address, register, data element, or instruction encoding.

**Little-endian.** A byte-ordering method in memory where the address  $n$  of a word corresponds to the *least-significant byte*. In an addressed memory word, the bytes are ordered (left to right) 3, 2, 1, 0, with 3 being the *most-significant byte*. See *Big-endian*.

**Local access window.** Mapping used to translate a region of memory to a particular target interface, such as the DDR SDRAM controller or the PCI controller. The local memory map is defined by a set of eight local access windows. The size of each window can be configured from 4 Kbytes to 2 Gbytes.

---

**M** **Media access control (MAC) sublayer.** Sublayer that provides a logical connection between the MAC and its peer station. Its primary responsibility is to initialize, control, and manage the connection with the peer station.

**Media-independent interface (MII) sublayer.** Sublayer that provides a standard interface between the MAC layer and the physical layer for 10/100-Mbps operations. It isolates the MAC layer and the physical layer, enabling the MAC layer to be used with various implementations of the physical layer.



**Medium-dependent interface (MDI) sublayer.** Sublayer that defines different connector types for different physical media and PMD devices.

**Memory access ordering.** The specific order in which the processor performs load and store memory accesses and the order in which those accesses complete.

**Memory-mapped accesses.** Accesses whose addresses use the page or block address translation mechanisms provided by the MMU and that occur externally with the bus protocol defined for memory.

**Memory coherency.** An aspect of caching in which it is ensured that an accurate view of memory is provided to all devices that share system memory.

**Memory consistency.** Refers to agreement of levels of memory with respect to a single processor and system memory (for example, on-chip cache, secondary cache, and system memory).

**Memory management unit (MMU).** The functional unit that is capable of translating an *effective (logical) address* to a physical address, providing protection mechanisms, and defining caching methods.

**Modified/exclusive/invalid (MEI).** *Cache coherency* protocol used to manage caches on different devices that share a memory system. Note that neither the PowerPC ISA nor the Power ISA definitions specifies the implementation of an MEI protocol to ensure cache coherency.

**Modified state.** MEI state (M) in which one, and only one, caching device has the valid data for that address. The data at this address in external memory is not valid.

**Most-significant bit (msb).** The highest-order bit in an address, registers, data element, or instruction encoding.

**Most-significant byte (MSB).** The highest-order byte in an address, registers, data element, or instruction encoding.

---

## N

**NaN.** An abbreviation for not a number; a symbolic entity encoded in floating-point format. There are two types of NaNs—signaling NaNs and quiet NaNs.

**No-op.** No-operation. A single-cycle operation that does not affect registers or generate bus activity.

- 
- O**
- OCeaN.** (On-chip network) Non-blocking crossbar switch fabric. Enables full duplex port connections at 128Gb/s concurrent throughput and independent per port transaction queuing and flow control. Permits high bandwidth, high performance, as well as the execution of multiple data transactions.
- Outbound ATMU windows.** Mappings that perform address translations from local 32-bit address space to the address spaces of, which may be much larger than the local space. Outbound ATMU windows also map attributes such as transaction type or priority level.
- 
- P**
- Packet.** A unit of binary data that can be routed through a network. Sometimes packet is used to refer to the frame plus the preamble and start frame delimiter (SFD).
- Page.** A region in memory. The OEA defines a page as a 4-Kbyte area of memory aligned on a 4-Kbyte boundary.
- Page access history bits.** The *changed* and *referenced* bits in the PTE keep track of the access history within the page. The referenced bit is set by the MMU whenever the page is accessed for a read or write operation. The changed bit is set when the page is stored into. See [Changed bit](#) and [Referenced bit](#).
- Page fault.** A page fault is a condition that occurs when the processor attempts to access a memory location that does not reside within a *page* not currently resident in *physical memory*. A page fault exception condition occurs when a matching, valid *page table entry* (PTE[V] = 1) cannot be located.
- Page table.** A table in memory is comprised of *page table entries*, or PTEs. It is further organized into eight PTEs per PTEG (page table entry group). The number of PTEGs in the page table depends on the size of the page table (as specified in the SDR1 register).
- Page table entry (PTE).** Data structures containing information used to translate *effective address* to physical address on a 4-Kbyte page basis. A PTE consists of 8 bytes of information in a 32-bit processor and 16 bytes of information in a 64-bit processor.
- Physical coding sublayer (PCS).** Sublayer responsible for encoding and decoding data stream to and from the MAC sublayer.
- Physical medium attachment (PMA) sublayer.** Sublayer responsible for serializing code groups into a bit stream suitable for serial bit-oriented physical devices (SERDES) and vice versa. Synchronization is also performed for proper data decoding in this sublayer. The PMA sits between the PCS and the PMD sublayers.
- Physical medium dependent (PMD) sublayer.** Sublayer responsible for signal transmission. The typical PMD functionality includes amplifier, modulation, and wave shaping. Different PMD devices may support different media.

**Physical memory.** The actual memory that can be accessed through the system's memory bus.

**Pipelining.** A technique that breaks operations, such as instruction processing or bus transactions, into smaller distinct stages or tenures (respectively) so that a subsequent operation can begin before the previous one has completed.

**Primary opcode.** The most-significant 6 bits (bits 0–5) of the instruction encoding that identifies the type of instruction.

**Program order.** The order of instructions in an executing program. More specifically, this term is used to refer to the original order in which program instructions are fetched into the instruction queue from the cache.

**Protection boundary.** A boundary between *protection domains*.

**Protection domain.** A protection domain is a segment, a virtual page, a BAT area, or a range of unmapped effective addresses. It is defined only when the appropriate relocate bit in the MSR (IR or DR) is 1.

---

## Q

**Quad word.** A group of 16 contiguous locations starting at an address divisible by 16.

**Quiesce.** To come to rest. The processor is said to quiesce when an exception is taken or a **sync** instruction is executed. The instruction stream is stopped at the decode stage and executing instructions are allowed to complete to create a controlled context for instructions that may be affected by out-of-order, parallel execution. See [Context synchronization](#).

---

## R

**rA.** The rA instruction field is used to specify a GPR to be used as a source or destination.

**rB.** The rB instruction field is used to specify a GPR to be used as a source.

**rD.** The rD instruction field is used to specify a GPR to be used as a destination.

**rS.** The rS instruction field is used to specify a GPR to be used as a source.

**Record bit.** Bit 31 (or the Rc bit) in the instruction encoding. When it is set, updates the condition register (CR) to reflect the result of the operation.

**Reconciliation sublayer.** Sublayer that maps the terminology and commands used in the MAC layer into electrical formats appropriate for the physical layer entities.

**Reduced instruction set computing (RISC).** An *architecture* characterized by fixed-length instructions with nonoverlapping functionality and by a separate set of load and store instructions that perform memory accesses.

**Referenced bit.** One of two *page history bits* found in each *page table entry*. The processor sets the *referenced bit* whenever the page is accessed for a read or write. See also [Page access history bits](#).

**Reservation.** The processor establishes a reservation on a *cache block* of memory space when it executes an **lwarx** instruction to read a memory semaphore into a GPR.

**Reservation station.** A buffer between the dispatch and execute stages that allows instructions to be dispatched even though the results of instructions on which the dispatched instruction may depend are not available.

---

## S

**Secondary cache.** A cache memory that is typically larger and has a longer access time than the primary cache. A secondary cache may be shared by multiple devices. Also referred to as L2, or level-2, cache.

**Set (v).** To write a nonzero value to a bit or bit field; the opposite of *clear*. The term ‘set’ may also be used to generally describe the updating of a bit or bit field.

**Set (n).** A subdivision of a *cache*. Cacheable data can be stored in a given location in one of the sets, typically corresponding to its lower-order address bits. Because several memory locations can map to the same location, cached data is typically placed in the set whose *cache block* corresponding to that address was used least recently. See *Set-associative*.

**Set-associative.** Aspect of cache organization in which the cache space is divided into sections, called *sets*. The cache controller associates a particular main memory address with the contents of a particular set, or region, within the cache.

**Slave.** The device addressed by a master device. The slave is identified in the address tenure and is responsible for supplying or latching the requested data for the master during the data tenure.

**Snooping.** Monitoring addresses driven by a bus master to detect the need for coherency actions.

**Snoop push.** Response to a snooped transaction that hits a modified cache block. The cache block is written to memory and made available to the snooping device.

**Stall.** An occurrence when an instruction cannot proceed to the next stage.

**Sticky bit.** A bit that when *set* must be cleared explicitly.

**Superscalar machine.** A machine that can issue multiple instructions concurrently from a conventional linear instruction stream.

**Supervisor mode.** The privileged operation state of a processor. In supervisor mode, software, typically the operating system, can access all control registers and can access the supervisor memory space, among other privileged operations.

**Synchronization.** A process to ensure that operations occur strictly *in order*. See *Context synchronization*.

**System memory.** The physical memory available to a processor.

- 
- T**
- Tenure.** The period of bus mastership. There can be separate address bus tenures and data bus tenures.
- Throughput.** The measure of the number of instructions that are processed per clock cycle.
- Time-division multiplex (TDM).** A single serial channel used by several channels taking turns.
- Transaction.** A complete exchange between two bus devices. A transaction is typically comprised of an address tenure and one or more data tenures, which may overlap or occur separately from the address tenure. A transaction may be minimally comprised of an address tenure only.
- Transfer termination.** Signal that refers to both signals that acknowledge the transfer of individual beats (of both single-beat transfer and individual beats of a burst transfer) and to signals that mark the end of the tenure.
- Translation lookaside buffer (TLB).** A cache that holds recently-used *page table entries*.
- 
- U**
- User mode.** The operating state of a processor used typically by application software. In user mode, software can access only certain control registers and can access only user memory space. No privileged operations can be performed. Also referred to as problem state.
- 
- V**
- Virtual address.** An intermediate address used in the translation of an *effective address* to a physical address.
- Virtual memory.** The address space created using the memory management facilities of the processor. Program access to *virtual memory* is possible only when it coincides with *physical memory*.
- 
- W**
- Way.** A location in the cache that holds a cache block, its tags, and status bits.
- Word.** A 32-bit data element.
- Write-back.** A cache memory update policy in which processor write cycles are directly written only to the cache. External memory is updated only indirectly, for example, when a modified cache block is *cast out* to make room for newer data.
- Write-through.** A cache memory update policy in which all processor write cycles are written to both the cache and memory.



# Index

## A

- Address broadcast enable, 7-23
- Address mask (LBC), 10-11
- AFEU
  - status register, 14-34, 14-44
- Alignment
  - overview, 7-32
- Application
  - examples, 1-18
  - information, *see* Initialization/application information
- Arbiter, *see* CSB arbiter and bus monitor
- Arbitration
  - I<sup>2</sup>C interface
    - arbitration control, 17-13
    - loss of arbitration—forcing of slave mode, 17-23
    - procedure for arbitration, 17-13
- Architecture, overview of device, 1-7
- Architecture, PowerPC, 7-13

## B

- Big-endian, 12-3, 12-19, 13-25, 16-7
- Block address translation (BAT), 7-3
  - see also* Memory Management Unit (MMU), 7-3
- Block diagrams
  - clock subsystem, 4-29
  - DDR controller, 9-2, 9-29
  - DMA/messaging unit, 12-1
    - DMA controller, 12-16
  - DUART, 18-2
  - e300 core, 1-9
  - eTSEC, 15-2
  - general purpose timers, 5-49
  - GPIO<sub>n</sub> module, 21-1
  - I/O sequencer, 11-1
  - I<sup>2</sup>C interface, 17-1
  - IPIC interrupt sources, 8-3
  - JTAG interface, 20-1
  - local bus controller (LBC), 10-1
  - PCI, 13-1
  - periodic interval timer, 5-43, 5-47
  - real time clock module, 5-36, 5-41
  - security engine, 1-10
  - SPI, 19-2
  - timer pair-cascaded mode, 5-63

- timers super-cascaded mode, 5-63
- watchdog timer, 5-29, 5-34
- Boot sequencer
  - I<sup>2</sup>C interface, 4-23–4-25, 17-2, 17-15
- Boundary-scan testing, *see* JTAG interface
- Branch processing unit (BPU), 7-1
  - overview, 7-7
- Branch trace enable (BE), 7-18
- Breakpoints
  - signaling, 7-38
- Buffer descriptors, *see* eTSEC, buffer descriptors
- Bus interface
  - I<sup>2</sup>C, 1-16
    - PCI bus arbitration unit, 1-13
- Bus interface unit (BIU), 7-10
- Bus monitor, *see* CSB arbiter and bus monitor

## C

- Caches
  - cache locking
    - way locking, 7-30
  - operations, 7-10
  - way-locking, 7-30
- Channel reset, 14-64
- cint* (critical interrupt signal), 8-2
- Clock multiplier, 7-12
- Clocks
  - DDR clock distribution, 9-42
  - eTSEC
    - inputs and outputs, 15-8
    - management clock out (EC\_MDC), 15-9, 15-73
  - I<sup>2</sup>C
    - clock stretching, 17-15
    - clock synchronization, 17-15
    - input synchronization and digital filter, 17-15
  - introduction, 4-28
  - LBC bus clocks and clock ratios, 10-3
    - clock ratio register (LCRR), 10-33
  - PCI agent mode, 4-30
  - PCI host mode, 4-30
    - PCI clock outputs (PCI\_CLK\_OUT[0:7]), 4-30
  - signals, 4-3–4-4
    - PCI\_CLK, 4-4
    - PCI\_CLK\_OUT[0:7], 4-4, 4-30

- PCI\_SYNC\_IN, 4-4
- PCI\_SYNC\_OUT, 4-4
- SYS\_CLK\_IN, 4-3
- USB\_CLK\_IN, 4-3
- subsystem block diagram, 4-29
- system
  - domains, 4-30
  - registers
    - configuration, 4-37–4-41
- Coherent system bus (CSB)
  - arbiter, *see also* CSB arbiter and bus monitor
  - overview, 6-1
- Completion unit, overview, 7-8
- Condition register (CR), 7-16
- Configuration
  - boot sequencer, 4-17
  - DDR, 9-9–9-29, 9-31
  - eTSEC interfaces, 15-192
  - LBC
    - configuration register (LBCR), 10-31
  - PCI
    - host/agent mode, 13-3
    - PCI arbiter, 13-4
  - reset, 4-9
    - sampled signals, 3-12
    - see also* Reset, configuration
- Controller registers, 14-67
- Core interface
- Core, *see* e300 core
- CR (condition register)
  - overview, 7-16
- Critical input (*cint*) interrupt, 7-33
- Critical interrupt
  - exception enable (G2\_LE only), 7-18
- Crypto-channel
  - configuration register, 14-55
- Crypto-channel registers, 14-55
- CSB arbiter and bus monitor
  - coherent system bus, 6-1
  - error handling sequence, 6-16
  - features, 6-1
  - functionality
    - arbitration policy, 6-10
      - address bus arbitration after  $\overline{\text{ARTRY}}$ , 6-13
      - address bus arbitration with  $\overline{\text{PRIORITY}}[0:1]$ , 6-11
      - address bus arbitration with  $\overline{\text{REPEAT}}$ , 6-12
      - address bus parking, 6-13
      - data bus arbitration, 6-13
    - bus error detection, 6-13
      - address only transaction type, 6-14
      - address time out, 6-13
      - data time out, 6-14

- illegal (ECIW $\overline{\text{X}}$ /ECOW $\overline{\text{X}}$ ) transaction type, 6-15
- reserved transaction type, 6-15
- transfer error, 6-14
- initialization sequence, 6-16
- memory map/register definition, 6-2
- overview, 6-1
- registers, 6-2–6-10

## D

- Data address translation, 7-18
- Data cache enable, 7-22
- Data cache flash invalidate, 7-22
- Data cache lock, 7-22
- Data cache way lock, 7-25
- Data encryption standard execution units (DEU), 14-19, 14-40
- Data TLB miss on load interrupt, 7-33
- Data TLB miss on store interrupt, 7-33
- DBAT, *see* Block address translation (BAT)
- DDR controller
  - address signal mappings, 9-4
  - block diagram, 9-2, 9-29
  - clock distribution, 9-42
  - configuration, example, 9-31
  - data beat ordering, 9-48
  - features, 9-2
  - functional description, 9-29
  - initialization/application information, 9-49
    - programming different memory types, 9-50
  - memory map/register definition, 9-8
  - modes of operation, 9-3
  - on-die termination for CSs, 9-7
  - page mode and logical bank retention, 9-48
  - register descriptions, 9-9
    - by acronym, *see* Register Index
    - configuration registers, 9-9–9-29
- SDRAM operation, 9-32
  - address multiplexing, 9-34
  - initialization sequence, 9-53
  - JEDEC standard interface commands, 9-37
  - mode-set command timing, 9-42
  - organizations supported, 9-33
  - refresh operation, 9-44
    - power-saving modes, 9-45
    - timing, 9-45
  - registered DIMM mode, 9-43
  - timing, 9-39
  - write timing adjustments, 9-43
- self-refresh
  - operation in sleep mode, 9-47
- signals summary, 9-3
  - see also* Signals, DDR



- DDR memory controller
  - debug configuration, 5-26
  - overview, 1-10
- Debug configuration, 5-26
  - DDR, 5-26
  - local bus, 5-27
- Debug facilities, 7-38
  - performance monitor uses, 7-12
- Debug modes
  - LBC source ID debug mode, 10-4
- Decrementer, 7-33
- Descriptor structure, 14-11
- DEU
  - FIFOs, 14-28
  - interrupt control register, 14-25, 14-47
  - interrupt status register, 14-24, 14-45
  - IV register, 14-27
  - key registers, 14-28
  - key size register, 14-20, 14-21, 14-42
  - mode register, 14-19, 14-40
  - reset control register, 14-22, 14-43
- DMA controller, *see* DMA/messaging unit, DMA controller
- DMA engine, 16-46
- DMA/messaging unit
  - block diagram, 12-1
  - DMA controller
    - block diagram, 12-16
    - descriptors, 12-18
      - big-endian mode, 12-19
      - DMA chain, 12-19
      - little-endian mode, 12-20
    - halt and error conditions, 12-17
    - operation, 12-16
      - coherency, 12-17
    - overview, 1-17, 12-16
  - features, 12-1
  - functional description, 12-15
  - initialization steps
    - in chaining mode, 12-20
    - in direct mode, 12-20
  - memory map/register definition, 12-2
  - message unit, 12-15
    - doorbell registers, 12-16
    - messaging registers, 12-15
  - registers, 12-3–12-15
    - by acronym, *see* Register Index
    - configuration, control, and status registers, 12-3–12-9
- Do-Complete-Split state
  - asynchronous, 16-91
  - periodic interrupt, 16-100
- Doorbell registers, 12-6–12-7
- Do-Start-Split state
  - asynchronous, 16-91
  - periodic interrupt, 16-99
- DSI (data storage interrupt), 7-32
- DTLB, 7-3
- Dual universal asynchronous receiver/transmitters, *see* DUART
- DUART
  - asynchronous communication bits, 18-2
  - parity bit, 18-19
  - START bit, 18-19
  - STOP bit, 18-20
  - baud-rate generator logic, 18-20
  - block diagram, 18-2
  - divisor latch access bit (ULCR<sub>n</sub>[DLAB]), 18-4, 18-11
  - error handling, 18-21
    - framing error, 18-9, 18-14, 18-19, 18-20, 18-21
    - overrun error, 18-21
    - parity error, 18-21
  - features, 18-2
  - functional description, 18-18
  - initialization/application information, 18-22
  - interrupt handling
    - interrupt control logic, 18-22
    - interrupt enable and control registers, 18-8–18-10
  - memory map/register definition, 18-4–18-5
  - modes of operation, 18-3
    - DMA mode selection, 18-22
    - FIFO mode, 18-21
    - interrupts, 18-21
    - local loop-back mode, 18-20
  - overview, 18-1
  - PC16450 UART compatibility, 18-2
  - registers, 18-5–18-18
    - by acronym, *see* Register Index
  - serial interface data format, 18-2
  - serial interface operation, 18-19–18-20
    - data transfer, 18-19
    - START bit, 18-19
    - STOP bit, 18-20
    - transaction protocol example, 18-19
  - signals, 18-3–18-4
    - UART\_CTS[0:1] (DUART clear to send), 18-1, 18-3, 18-4
    - UART\_RTS[0:1] (DUART request to send), 18-1, 18-3, 18-4
    - UART\_SIN [0:1] (DUART transmitter serial data in), 18-3
    - UART\_SOUT [0:1] (DUART transmitter serial data out), 18-3, 18-4
- Dynamic power management enable, 7-21

## E

- e300 core
  - block diagram, 1-9
  - overview, 1-7
- e300 core, differences between cores, 7-39
- EC\_GTX\_CLK125 (eTSEC gigabit transmit 125 MHz source) signal, 15-9
- EC\_MDC (eTSEC management data clock) signal, 15-9
- EC\_MDIO (eTSEC management data input/output, BIDI) signal, 15-9
- Error handling
  - CSB arbiter and bus monitor, 6-16
  - DMA/messaging unit, 12-17
  - DUART, 18-21
    - framing error, 18-9, 18-14, 18-19, 18-20, 18-21
    - overrun error, 18-21
    - parity error, 18-21
  - eTSEC, 15-158–15-160
  - I<sup>2</sup>C interface
    - boot sequencer mode, 4-26, 17-16
  - LBC
    - transfer error registers, 10-25–10-30
- Ethernet controller, *see* TSEC
- eTSEC
  - block diagram, 15-2
  - buffer descriptors, 15-185–15-192
    - receive buffer descriptors (RxBD), 15-190
    - transmit buffer descriptors (TxBD), 15-186
  - clocks
    - inputs and outputs, 15-8
    - management clock out (EC\_MDC), 15-9, 15-73
  - configuration of interfaces, 15-192
    - MAC configuration, 15-64
    - MII interface mode, 15-193
    - RGMIi interface mode, 15-196
    - RMII interface mode, 15-200
    - RTBI interface mode, 15-204
  - error-handling, 15-158–15-160
  - features, 15-2
  - functional description, 15-134
  - gigabit Ethernet channel operation, 15-143
    - flow control, 15-154
    - frame reception, 15-147
    - frame recognition, 15-150
    - frame transmission, 15-145
    - initialization sequence, 15-143
      - soft reset and reconfiguring procedure, 15-145
    - internal and external loop back, 15-158
    - inter-packet gap time, 15-158
    - Magic Packet mode, 15-154
    - preamble customization, 15-148
    - RMON support, 15-150
  - hash function
    - algorithm, 15-152
    - registers, 15-106–15-108
  - initialization/application information, 15-192
    - gigabit Ethernet channel, 15-143
    - soft reset and reconfiguring procedure, 15-145
  - see also* eTSEC, configuration
  - interrupts, 15-155–15-158
    - interrupt coalescing, 15-156
      - by frame count threshold, 15-156
      - by timer threshold, 15-157
    - interrupt registers, 15-24–15-29
  - lossless flow control, 15-175
    - back pressure determination and free buffers, 15-175
    - software use of hardware-initiated back pressure, 15-177
  - MAC functionality, 15-64–15-79
    - configuration, 15-64
    - CSMA/CD control, 15-64
    - handling packet collisions, 15-65
    - packet flow control, 15-65
    - PHY links control, 15-66
    - registers, 15-67–15-79
  - memory map/register definition, 15-11
    - detailed memory map, 15-12–15-22
    - eTSEC2–4 controller offsets, 15-22
    - top-level module map, 15-11
  - modes of operation, 15-4
    - RMON support, 15-79
  - overview, 15-1
  - physical interface connections, 15-134
    - media-independent interface (MII), 15-135
    - reduced gigabit media-independent interface (RGMIi), 15-136
    - reduced media-independent interface (RMII), 15-135
    - reduced ten-bit interface (RTBI), 15-137
  - quality of service (QoS) support, 15-165–15-175
    - receive queue filer, 15-167
    - transmission scheduling, 15-173
  - register descriptions, 15-22–15-134
    - by acronym, *see* Register Index
    - DMA attribute registers, 15-108
    - general control and status registers, 15-22
    - hash function registers, 15-106–15-108
    - lossless flow control registers, 15-109–15-110
    - MAC registers, 15-67–15-79
    - MIB registers, 15-79–15-106
    - receive control and status registers, 15-48–15-63
    - ten-bit interface registers, 15-123–15-134
    - transmit control and status registers, 15-35–15-46
  - signals
    - see also* Signals, eTSEC
    - summary, 15-6

- signals<\$startmode, 15-8
- TCP/IP off-load, 15-160–15-165
  - frame control blocks, 15-161
  - receive path off-load, 15-163
  - transmit path off-load, 15-161

## EU

- access, 14-65, 14-66
- assignment status register, 14-67
- Exception little-endian mode, 7-17
- Exception prefix, 7-18
- Exceptions
  - overview, 7-30
- External interrupt enable, 7-17

## F

### Features

- overview of device features, 1-2
- Fetch register, 14-62
- FIFO RAM controller, 16-46
- Floating-point available, 7-18
- Floating-point exception mode 0, 7-18
- Floating-point exception mode 1, 7-18
- Floating-point model
  - FP registers (FPR $n$ ), 7-16
  - FPR $n$  (floating-point registers 0–31), 7-16
- Force branch indirect on bus, 7-23
- FPR $n$  (floating-point registers 0–31), 7-16
- FPSCR (floating-point status and control reg.), 7-16
- FSTNs
  - host controller operational model, 16-95
  - software operational model, 16-97

## G

### G2

- overview, 7-13
- General Purpose I/O module (GPIO), *see* GPIO
- General purpose timers (GTM), 5-48
  - block diagram, 5-49
  - external signal description, 5-51
  - features, 5-49
  - functional description, 5-61
    - capture modes, 5-61
    - cascaded modes, 5-62
    - general-purpose timer units, 5-61
    - reference modes, 5-61
  - initialization/application information, 5-64
  - memory map/register definition, 5-52
  - modes of operation, 5-50
    - capture, 5-51
    - cascaded, 5-50

- clock source, 5-50
- reference, 5-50
- overview, 5-48
- registers, 5-54–5-60
- GPCM (LBC general-purpose chip-select machine), 10-45
  - see also* Local bus controller (LBC)

## GPIO

- block diagram, 21-1
- features, 21-1
- memory map/register definition, 21-2
- overview, 21-1
- registers, 21-3–21-5
- signals, 21-2
- GPR $n$  (general-purpose registers 0–31), 7-16

## H

- Hash function, *see* eTSEC, hash function
- HID $n$  (hardware implementation registers 0–2)
  - PLL configuration, 7-23
- High BAT enable, 7-24
- HRESET, 4-8

## I

- I/O sequencer, 13-1
  - block diagram, 11-1
  - features, 11-2
  - functional description, 11-6
    - PCI outbound address translation, 11-7
    - transaction forwarding, 11-6
      - from the CSB port, 11-7
      - from the DMA port, 11-7
      - from the PCI ports, 11-7
    - transaction ordering, 11-8
  - memory map/register definition, 11-2–11-3
  - overview, 11-1
  - registers, 11-3–11-6
- I<sup>2</sup>C interface
  - arbitration
    - arbitration control, 17-13
    - loss of arbitration—forcing of slave mode, 17-23
    - procedure for arbitration, 17-13
  - block diagram, 17-1
  - boot sequencer mode, 4-23–4-25, 17-2, 17-15
    - error condition behavior, 4-26, 17-16
  - calling address match condition, 17-5
  - clock control, 17-14
    - clock stretching, 17-15
    - clock synchronization, 17-15
    - input synchronization and digital filter, 17-15
    - master mode, 17-14

- slave mode, 17-14
- data transfer, 17-11
- error handling
  - boot sequencer mode, 4-26, 17-16
- features, 17-2
- frequency divider
  - frequency divider register (I2CFDR), 17-6
- functional description, 17-10
- handshaking, 17-14
- implementation details, 17-12
  - address compare, 17-13
  - control transfer, 17-12
  - transaction monitoring, 17-12
- initialization/application information, 17-19–17-20
  - boot sequencer mode, *see* I<sup>2</sup>C interface, boot sequencer mode
  - generation of SCL when SDA low, 17-22
  - initialization sequence, 17-21
  - post-transfer software response, 17-21
  - repeated START generation, 17-22
  - START generation, 17-11, 17-21
  - STOP generation, 17-12, 17-22
- interrupts
  - calling address match condition, 17-5
  - flowchart for interrupt service routine, 17-19
  - interrupt after transfer, 17-21
  - interrupt enable bit (I2CCR[MIEN]), 17-7
  - interrupt on START, 17-21
  - interrupt pending status bit (I2CSR[MIF]), 17-9
  - interrupt-driven byte-to-byte transfers, 17-2
  - read of last byte, 17-22
  - slave mode interrupt service routine guidelines, 17-22
    - for slave transmitter routine, 17-23
    - loss of arbitration, 17-23
- memory map/register definition, 17-4
- modes of operation, 17-2
  - boot sequencer mode, 17-2, 17-15
  - interrupt-driven byte-to-byte data transfer, 17-2
  - master mode, 17-2
  - slave mode, 17-2
- overview, 1-16
- register descriptions, 17-5–17-10
  - by acronym, *see* Register Index
- serial data/clock wires, 17-1
- signals, 17-3–17-4
  - see also* Signals, I<sup>2</sup>C
- transaction protocol, 17-10
  - handshaking, 17-14
  - repeated START condition, 17-3, 17-12
  - slave address transmission, 17-11
  - START condition, 17-3, 17-11, 17-21
  - STOP condition, 17-3, 17-12, 17-22
- IBATnU/L (instruction block address translation regs. 0–7, upper/lower), 7-3
- ID register, 14-73
- IEEE 1149.1 specifications
  - specification compliance, 20-3
- Initialization
  - DDR (initialization and application information), 9-49
    - programming different memory types, 9-50
  - eTSEC (initialization and application information), 15-143, 15-192
    - see also* eTSEC, configuration
  - I<sup>2</sup>C interface (initialization and application information)
    - STOP generation, 17-22
- Initialization/application information
  - CSB arbiter and bus monitor, 6-16
  - DMA/messaging unit, 12-20
  - GTM registers, 5-64
  - host controller, 16-68
  - I<sup>2</sup>C interface, 17-19–17-20
    - boot sequencer mode, *see* I<sup>2</sup>C interface, boot sequencer mode
    - generation of SCL when SDA low, 17-22
    - initialization sequence, 17-21
    - post-transfer software response, 17-21
    - repeated START generation, 17-22
    - START generation, 17-11, 17-21
    - STOP generation, 17-12
- LBC, 10-88
- PCI, 13-60
- power management control, 5-91
- real time clock module
  - RTC programming guidelines, 5-42
- SPI, 19-16
  - master programming example, 19-16
  - slave programming example, 19-16
- watchdog timer, 5-35
- Initiator write, 14-66
- Instruction address translation, 7-18
- Instruction cache enable, 7-21
- Instruction cache flash invalidate, 7-22
- Instruction cache lock, 7-22
- Instruction cache way lock, 7-25
- Instruction timing
  - overview, 7-35–7-36
  - see also* Execution timing
- int* (internal interrupt signal), 8-2
- INTA, 12-15
- Integrated programmable interrupt controller, *see* IPIC
- Interfaces
  - I<sup>2</sup>C, 1-16
  - JTAG
    - block diagram, 20-1

- TAP controller, 20-4
- PCI interface
  - bus arbitration unit, 1-13
- USB, 16-2
- Interrupt
  - clear register, 14-71
  - mask register, 14-68
  - status register, 14-70
- Interrupt handling
  - interrupt types
    - performance monitor interrupt, 7-12
- Interrupts
  - channel done, 14-64
  - channel error, 14-64
  - DUART
    - interrupt control logic, 18-22
    - interrupt enable and control registers, 18-8–18-10
  - eTSEC, 15-155–15-158
    - interrupt registers, 15-24–15-29
  - general, 14-64
  - I<sup>2</sup>C interface
    - calling address match condition, 17-5
    - flowchart for interrupt service routine, 17-19
    - interrupt after transfer, 17-21
    - interrupt enable bit (I2CCR[MIEN]), 17-7
    - interrupt on START, 17-21
    - interrupt pending status bit (I2CSR[MIF]), 17-9
    - interrupt-driven byte-to-byte transfers, 17-2
    - read of last byte, 17-22
    - slave mode interrupt service routine guidelines, 17-22
      - for slave transmitter routine, 17-23
    - loss of arbitration, 17-23
  - LBC interrupt register, 10-28
  - SEC, 14-75
  - split transaction, 16-92
- IPIC
  - block diagram, 8-3
  - features, 8-4
  - functional description, 8-28
  - interrupts
    - configuration, 8-29
    - highest priority, 8-31
    - internal
      - group relative priority, 8-30
    - machine check, 8-35
    - masking sources, 8-34
    - mixed
      - group relative priority, 8-30
    - request masking, 8-35
    - source priorities, 8-31
      - levels, 8-31
    - types, 8-28
      - vector generation and calculation, 8-35
      - memory map/register definition, 8-6–8-7
      - modes of operation, 8-4
        - core disable mode, 8-5
        - core enable mode, 8-4
      - overview, 8-1
      - registers, 8-7–8-28
        - by acronym, *see* Register Index
      - signals, 8-5–8-6
        - cint* (critical interrupt), 8-2
        - int* (internal interrupt), 8-2
        - mcp* (machine check processor), 8-2
        - overview, 8-5
        - smi* (system management interrupt), 8-2
  - ISI (instruction storage interrupt), 7-32
  - Isochronous transfer descriptor (iTD), *see* USB interface, isochronous (high-speed) transfer descriptor (iTd)

## J

- Joint test action group, *see* JTAG interface
- JTAG interface
  - block diagram, 20-1
  - registers, 20-3
  - signals, 20-1–20-3
  - TAP (test access port) controller, 20-4
  - TAP controller, 20-4
- JTAG test and debug interface, 7-12, 7-38

## L

- LA[27:31] (LBC non-multiplexed address) signals, 10-6
- LAD[0:31] (LBC multiplexed address/data) signals, 10-7
- LALE (LBC external address latch enable) signal, 10-5, 10-41
- LBC, *see* Local bus controller (LBC)
- LBCTL (LBC data buffer control) signal, 10-6, 10-44
- $\overline{\text{LBS}}[0:3]$  (LBC UPM byte select) signals, 10-5
- LCK[0:2] (LBC clock) signals, 10-7
- $\overline{\text{LCS}}[0:7]$  (LBC chip select) signals, 10-5
- $\overline{\text{LCS}}_0$  (LBC chip select 0) signal, 10-56, 10-69
- LGPL0 (LBC GP line 0) signal, 10-6
- LGPL1 (LBC GP line 1) signal, 10-6
- LGPL2 (LBC GP line 2) signal, 10-6
- LGPL3 (LBC GP line 3) signal, 10-6
- LGPL4 (LBC GP line 4) signal, 10-6
- LGPL5 (LBC GP line 5) signal, 10-6
- $\overline{\text{LGT}}_A$  (LBC GPCM transfer acknowledge) signal, 10-6, 10-55
- Little-endian, 16-7
- Little-endian mode enable, 7-18
- Load/store unit (LSU), 7-1
  - overview, 7-8

- Local bus controller (LBC)
- address and address space checking, 10-41
  - address mask field—option registers, 10-11
  - atomic bus operations, 10-44
  - block diagram, 10-1
  - boot chip-select operation, 10-56, 10-69
  - bus monitor, 10-44
  - bus turnaround, 10-91
    - additional address phases (UPM cycles), 10-92
    - address following read, 10-91
    - read data following address, 10-91
    - read-modify-write cycle (parity), 10-92
  - clocks and clock ratios, 10-3
    - clock ratio register (LCRR), 10-33
  - configuration
    - LBC configuration register (LBCR), 10-31
  - error handling
    - transfer error registers, 10-25–10-30
  - external access termination ( $\overline{\text{LGTA}}$ ), 10-55
  - features, 10-2
  - functional description, 10-40
  - general-purpose chip-select machine (GPCM), 10-45
    - chip-select and write enable negation timing, 10-50
    - chip-select assertion timing, 10-49
    - extended hold time on read accesses, 10-54
    - GPCM mode
      - registers, 10-13, 10-15
    - output enable timing, 10-54
    - programmable wait state configuration, 10-50
    - relaxed timing, 10-51
    - timing configuration, 10-46, 10-47, 10-65
  - initialization/application information, 10-88
  - interrupts
    - transfer error interrupt enable register (LTEIR), 10-28
  - memory map/register definition, 10-7
  - memory refresh timer prescaler, 10-22
  - modes of operation, 10-3
    - bus clock and clock ratios, 10-3
    - GPCM mode, registers, 10-13, 10-15
    - source ID debug mode, 10-4
    - UPM mode, registers, 10-18
  - overview, 1-14, 10-2
  - peripherals, 10-88
    - GPCM timing, 10-90
    - hierarchy for very high speeds, 10-89
    - multiplexed address/data, 10-88
  - port sizes, 10-92
  - register descriptions, 10-9
    - by acronym, *see* Register Index
  - signals, 10-4
    - see also* Signals, LBC
  - UPM interfaces, 10-72–10-102
    - block diagram, 10-72
    - extended hold time (reads), 10-87
    - programming the UPMs, 10-75
    - RAM array, 10-78
      - address multiplexing, 10-84
      - byte select signal timing, 10-83
      - chip select signal timing, 10-82
      - data timing, 10-85
      - general purpose signal timing, 10-83
      - LGPL[0:5] timing (LAST), 10-86
      - loop control, 10-83
      - RAM word definition, 10-79
      - REDO, 10-84
      - wait mechanism (WAEN), 10-86
    - signal timing, 10-77
    - synchronous UPWAIT (early transfer acknowledge), 10-87
    - UPM mode
      - registers, 10-18, 10-19
    - UPM requests, 10-73
      - exception requests, 10-75
      - memory access requests, 10-74
      - refresh timer requests, 10-74
      - software requests, 10-75
    - ZBT SRAM interface, 10-97, 10-102
  - $\overline{\text{LOE}}$  (LBC GPCM output enable) signal, 10-6
  - $\overline{\text{LWE}}$ [0:3] (LBC GPCM write enable) signals, 10-5
- ## M
- MA[0:14] (DDR address bus) signals, 9-6
  - MAC functionality, *see* eTSEC, MAC functionality
  - Machine check enable, 7-18
  - Master control register, 14-74
  - MBA[0:1] (DDR logical bank address) signals, 9-6
  - $\overline{\text{MCAS}}$  (DDR column address strobe) signal, 9-6
  - $\overline{\text{MCK}}$ [0:5] (DDR clock output complement) signals, 9-7
  - $\overline{\text{MCK}}$ [0:5] (DDR clock output) signals, 9-7
  - $\overline{\text{MCKE}}$ [0:3] (DDR clock enable) signals, 9-8
  - mcp* (machine check processor signal), 8-2
  - $\overline{\text{MCS}}$ [0:3] (DDR chip select) signals, 9-7
  - MDEU
    - context registers, 14-38
    - data size register, 14-32, 14-37
    - FIFOs, 14-40
    - interrupt control register, 14-36
    - interrupt status register, 14-32, 14-35
    - key registers, 14-39
    - key size register, 14-32
    - mode register, 14-28
    - reset control register, 14-33

MDM[0:8] (DDR SDRAM data output mask) signals, 9-7  
 MDQ[0:8] (DDR data bus strobe) signals, 9-5, 9-30  
 MDVAL (DDR/LBC debug mode data valid) signal, 10-7  
 Memory accesses, 7-37  
 Memory management unit (MMU)  
   overview, 7-9, 7-34  
 Memory maps  
   accessing CCSR memory from external masters, 5-16  
   address translation and mapping, 5-3  
   complete IMMR map, 2-1  
   configuring local access windows, 5-14  
   CSB arbiter and bus monitor, 6-2  
   DDR controller, 9-8  
   distinguishing local access windows from other mapping functions, 5-14  
   DMA/messaging unit, 12-2  
   DUART, 18-4–18-5  
   eTSEC, 15-11  
   general purpose timers, 5-52  
   GPIO, 21-2  
   I/O sequencer, 11-2–11-3  
   I<sup>2</sup>C, 17-4  
   inbound address translation and mapping windows, 5-15  
   IPIC, 8-6–8-7  
   LBC, 10-7  
   local access register, 5-4  
   local access windows  
     introduction, 5-4  
     precedence, 5-14  
   outbound address translation and mapping windows, 5-15  
 PCI, 13-11  
   periodic interval timer, 5-44  
   power management control, 5-65  
   real time clock module, 5-37  
   reset, 4-32  
   SPI, 19-8  
   USB interface, 16-5  
   watchdog timer, 5-30  
   window into configuration space, 5-4  
 Message digest execution unit (MDEU), 14-28  
 MODT[0:3] (DDR on-die termination) signals, 9-7  
 MPC603e core, *see* e300 core  
 MRAS (DDR row address strobe) signal, 9-6  
 MSR (machine state register), 7-17  
 MSRCID[0:4] (DDR/LBC debug source ID) signals, 10-7  
 MWE (DDR write enable) signal, 9-7

## N

No-op the data cache touch instructions, 7-23

## P

### PCI

address map  
   address translation  
     PCI outbound, 11-7  
 block diagram, 13-1  
 bridge  
   arbitration example, 13-45  
   PCI parity operation, 13-58  
 bus arbitration, 13-43  
   algorithm, 13-44  
   broken master lock-out, 13-45  
   master latency timer, 13-45  
   parking, 13-44  
 bus arbitration unit, 1-13  
 bus commands, 13-46  
 bus error functions, 13-57  
   parity, 13-57  
   reporting, 13-57  
 bus operations, 13-53  
   agent mode configuration access, 13-55  
   data streaming, 13-54  
   dual address cycles, 13-54  
   fast back-to-back transactions, 13-53  
   host mode configuration access, 13-54  
   interrupt acknowledge, 13-56  
   special cycle command, 13-55  
 CompactPCI Hot Swap specification support, 13-60  
 DMA controller, *see* DMA/messaging unit, DMA controller  
   features, 13-3  
   functional description, 13-43  
   inbound address translation, 13-59  
   inbound windows, 5-15  
   initialization/application information, 13-60  
     sequence for agent mode, 13-60  
     sequence for host mode, 13-60  
   memory map/register definition, 13-11  
 modes of operation, 13-3  
   host/agent mode configuration, 13-3  
   PCI arbiter configuration, 13-4  
 output hold configuration, 4-21  
 overview, 1-12  
 protocol fundamentals, 13-47  
   addressing, 13-47  
   basic transfer control, 13-47  
   bus driving and turnaround, 13-48  
   bus transactions, 13-49  
   byte enable signals, 13-48  
   device selection, 13-48

- read and write transactions, 13-49
  - burst read example, 13-50
  - burst write example, 13-51
  - single beat read example, 13-49
  - single beat write example, 13-50
- transaction termination, 13-51
  - target-initiated terminations, 13-52
- registers, 13-12–13-40
  - configuration access, 13-12–13-15
  - configuration space, 13-25–13-40
    - base class code, 13-31
    - BIST control, 13-33
    - cache line size, 13-32
    - capabilities pointer, 13-36
    - device ID, 13-27
    - GPL base address register 0, 13-34
    - GPL base address registers 1,2, 13-34
    - GPL extended base address registers 1,2, 13-35
    - header type, 13-33
    - hot swap register block, 13-40
    - interrupt line, 13-37
    - interrupt pin, 13-37
    - latency timer, 13-32
    - MAX LAT, 13-38
    - MIN GNT, 13-37
    - PCI arbiter control, 13-39
    - PCI command, 13-28
    - PCI function, 13-38
    - PCI power management register 0 (PCIPMR0), 13-41
    - PCI power management register 1 (PCIPMR1), 13-42
    - PCI status, 13-29
    - PIMMR base address, 13-33
    - revision ID, 13-30
    - standard programming interface, 13-30
    - subclass code, 13-31
    - sub-system device ID, 13-36
    - sub-system vendor ID, 13-36
    - vendor ID, 13-27
  - control and status, 13-15–13-25
  - signals, 13-4–13-11
- PCI\_C/BE[7:0] (PCI command/byte enable) signals, 13-6
- PCI\_CLK, 4-4
- PCI\_CLK\_OUT[0:7], 4-4, 4-30
- PCI\_PERR (PCI parity error) signal, 13-9
- PCI\_PME (PCI PME) signal, 13-9
- PCI\_REQ[4:0] (PCI bus request) signals, 13-9
- PCI\_SERR (PCI system error) signal, 13-10
- PCI\_STOP (PCI stop) signal, 13-10
- PCI\_SYNC\_IN, 4-4
- PCI\_SYNC\_OUT, 4-4
- PCI\_TRDY (PCI target ready) signal, 13-11
- Performance
  - characterizing through performance monitor event counting, 7-12
- Performance monitor, 7-12
- Performance monitor APU
  - interrupt triggered by events, 7-12
- Periodic interval timer (PIT), 5-42
  - block diagram, 5-43
  - external signal description, 5-43
  - features, 5-43
  - functional block diagram, 5-47
  - functional description, 5-47
  - memory map/register definition, 5-44
  - modes of operation, 5-43
  - operational modes, 5-48
  - overview, 5-42
  - registers, 5-44–5-47
- PHY clocks, 16-5
- PHY interface, 16-46
- PIC, *see* IPIC
- PKEU
  - EU\_GO register, 14-27, 14-38
  - status register, 14-23
- Power management
  - DDR interface, 9-45
  - modes, 7-11
- Power management control, 5-64
  - external signal description, 5-65
  - functional description, 5-72
    - dynamic power management, 5-72
    - exiting low power states, 5-78
    - shutting down unused blocks, 5-73
    - software-controlled power-down states, 5-73
  - initialization/application information, 5-91
  - memory map/register definition, 5-65
  - registers, 5-66–5-69
- Power saving mode
  - SEC, 14-75
- Power-on reset (POR)
  - flow, 4-6
  - output signal states during reset, 3-13
  - reset configuration signals, 3-12
  - timing diagram, 4-8
- PowerPC architecture
  - levels of implementation, 7-13
  - overview, 7-13
- Privilege level (PR), 7-17
- Processor core, *see* e300 processor core
- Program interrupt, 7-33
- Programmable interrupt controller, *see* IPIC



**Q**

- Quality of service (QoS), *see* eTSEC
- Queue element transfer descriptor (qTD), *see* USB interface, queue element transfer descriptor (qTD)
- Queue heads, *see* USB interface, queue heads

**R**

- Real time clock module (RTC), 5-35
  - block diagram, 5-36, 5-41
  - external signals description, 5-36
  - features, 5-36
  - functional description, 5-41
  - initialization/application information
    - RTC programming guidelines, 5-42
  - memory map/register definition, 5-37
  - modes of operation, 5-36
  - operational modes, 5-41
  - overview, 5-35
  - registers, 5-37–5-40
- Recoverable exception, 7-18
- Registers
  - AFEU
    - status, 14-34, 14-44
  - by acronym, *see* Register Index
  - clock
    - configuration, 4-37–4-41
  - configuration registers
    - hardware implementation registers (HID<sub>n</sub>)
      - PLL configuration, 7-23
  - crypto-channel
    - configuration, 14-55
    - general, 14-55
  - CSB arbiter and bus monitor, 6-2–6-10
  - DDR
    - configuration registers, 9-9–9-29
  - DEU
    - interrupt control, 14-25, 14-47
    - interrupt status, 14-24, 14-45
    - IV, 14-27
    - key, 14-28
    - key size, 14-20, 14-21, 14-42
    - mode, 14-19, 14-40
    - reset control, 14-22, 14-43
  - DMA/messaging unit, 12-3–12-15
    - configuration, control, and status registers, 12-3–12-9
  - DUART, 18-5–18-18
  - eTSEC, 15-22–15-134
    - DMA attribute registers, 15-108
    - general control and status registers, 15-22
    - hash function registers, 15-106–15-108
    - lossless flow control registers, 15-109–15-110
    - MAC registers, 15-67–15-79
    - MIB registers, 15-79–15-106
    - receive control and status registers, 15-48–15-63
    - ten-bit interface registers, 15-123–15-134
    - transmit control and status registers, 15-35–15-46
  - EU assignment status, 14-67
  - fetch, 14-62
  - GPIO, 21-3–21-5
  - I/O sequencer, 11-3–11-6
  - I<sup>2</sup>C interface, 17-5
  - ID, 14-73
  - interrupt
    - clear, 14-71
    - mask, 14-68
    - status, 14-70
  - IPIC, 8-7–8-28
  - JTAG
    - boundary-scan registers, 20-3
  - JTAG interface, 20-3
    - bypass register, 20-3
    - instruction register, 20-3
    - status register, 20-3
  - LBC, 10-9
  - master control, 14-74
  - MDEU
    - context registers, 14-38
    - data size, 14-32, 14-37
    - interrupt control, 14-36
    - interrupt status, 14-32, 14-35
    - key, 14-39
    - key size, 14-32
    - mode, 14-28
    - reset control, 14-33
  - PCI, 13-12–13-40
    - configuration access, 13-12–13-15
    - configuration space, 13-25–13-40
    - control and status, 13-15–13-25
  - PKEU
    - EU\_GO, 14-27, 14-38
    - status, 14-23
  - reset
    - configuration, 4-13, 4-32–4-37
  - SPI, 19-9–19-16
  - time base facility (TBL/TBU)
    - for reading, 7-17
  - USB interface
    - capability registers, 16-7
    - operational registers, 16-11
- Reset
  - actions, 4-5
  - causes, 4-5
  - configuration, 4-9

- boot memory space, 4-17
- boot ROM location, 4-18
- boot sequencer, 4-17
- CLKIN division, 4-11
- default words, 4-26
  - hard coded reset configuration words usage examples, 4-27
- e300 core true little endian mode, 4-20
- LALE, 4-21
- loading from I<sup>2</sup>C EEPROM, 4-23
  - boot sequencer, 4-23
  - calling address, 4-23
  - data format in reset configuration mode, 4-23
  - load fail, 4-25
- loading words, 4-21
  - from local bus EEPROM, 4-21
- PCI host/agent, 4-16
- selecting input signals, 4-11
- signals, 4-9
- TSEC1 mode, 4-19
- TSEC2 mode, 4-20
- words, 4-12
- words source, 4-10
- DDR debug configuration, 5-26
- functional device description, 5-33, 5-61, 5-72
- hard reset (HRESET)
  - flow, 4-8
- memory map/register definition, 4-32
- operations, 4-4
- PCI output hold configuration, 4-21
- power-on reset (POR)
  - flow, 4-6
  - timing diagram, 4-8
- registers
  - configuration, 4-13, 4-32–4-37
- signals, 4-1–4-2
- soft reset (SRESET)
  - flow, 4-9
- soft reset actions
  - and reconfiguring the eTSEC, 15-145
- Reset and clocking blocks
  - boot ROM location, 4-18
  - boot sequencer configuration, 4-17
  - CPU boot configuration, 4-15
  - external signal description, 5-43, 5-51, 5-65
  - functional description, 4-4, 5-33, 5-61, 5-72
  - host/agent configuration, 4-16
  - memory map/register definition, 5-30, 5-37, 5-44, 5-52
  - system PLL ratio, 4-30
  - TSEC width, 4-19
- RMON support, *see* eTSEC, modes of operation

## S

- SCL (I<sup>2</sup>C serial clock) signal, 17-3, 17-4
- SDA (I<sup>2</sup>C serial data) signal, 17-3, 17-4
- SEC
  - bus interface
  - interrupts, 14-75
  - power saving mode, 14-75
- Security engine
  - block diagram, 1-10
  - overview, 1-10
- Segment registers (SR<sub>n</sub>), 7-19
- Serial data/clock wires, 17-1
- Serial peripheral interface, *see* SPI
- Signals
  - clock, 4-3–4-4
    - PCI\_CLK, 4-4
    - PCI\_CLK\_OUT[0:7], 4-4, 4-30
    - PCI\_SYNC\_IN, 4-4
    - PCI\_SYNC\_OUT[0:7], 4-4
    - SYS\_CLK\_IN, 4-3
    - USB\_CLK\_IN, 4-3
  - complete signal listing
    - reference by functional block, 3-3–3-12
  - configuration, sampled at reset, 3-12
  - control, description, 5-36, 5-44, 5-52, 5-65
- DDR
  - MA[0:14] (address bus), 9-6
  - MBA[0:1] (logical bank address), 9-6
  - MCAS (column address strobe), 9-6
  - MCK[0:5] (DDR clock output complements), 9-7
  - MCK[0:5] (DDR clock outputs), 9-7
  - MCKE[0:3] (DDR clock enables), 9-8
  - MCS[0:3] (chip selects), 9-7
  - MDM[0:8] (SDRAM data output mask), 9-7
  - MDQS[0:8] (data bus strobes), 9-5, 9-30
  - MODT[0:3] (on-die termination), 9-7
  - MRAS (row address strobe), 9-6
  - MWE (write enable), 9-7
- DUART, 18-3–18-4
  - UART\_CTS[0:1] (DUART clear to send), 18-1, 18-3, 18-4
  - UART\_RTS[0:1] (DUART request to send), 18-1, 18-3, 18-4
  - UART\_SIN [0:1] (DUART transmitter serial data in), 18-3
  - UART\_SOUT [0:1] (DUART transmitter serial data out), 18-3, 18-4
- eset and clocking blocks, 5-43, 5-65
- eTSEC
  - EC\_GTX\_CLK125 (eTSEC gigabit transmit 125 MHz source), 15-9

- EC\_MDC (eTSEC management data clock), 15-9
- EC\_MDIO (eTSEC management data input/output, BIDI), 15-9
- TSECn\_COL (eTSEC 1-4 collision input), 15-8
- TSECn\_CRS (eTSEC 1-4 carrier sense input/FIFO receiver flow control), 15-8
- TSECn\_GTX\_CLK (eTSEC 1-4 gigabit transmit clock), 15-8
- TSECn\_RX\_CLK (eTSEC 1-4 receive clock), 15-9
- TSECn\_RX\_DV (eTSEC 1-4 receive data valid), 15-9
- TSECn\_RX\_ER (eTSEC 1-4 receive error), 15-10
- TSECn\_RXD[7:0] (eTSEC 1-4 receive data in), 15-9
- TSECn\_TX\_CLK (eTSEC 1-4 transmit clock in), 15-10
- TSECn\_TX\_EN (eTSEC 1-4 transmit data valid), 15-10
- TSECn\_TX\_ER (eTSEC 1-4 transmit error), 15-10
- TSECn\_TXD[7:0] (eTSEC 1-4 transmit data out), 15-10
- GPIO, 21-2
- groupings figure, 3-3
- I<sup>2</sup>C
  - SCL (serial clock), 17-3, 17-4
  - SDA (serial data), 17-3, 17-4
- IPIC, 8-5-8-6
  - cint (critical interrupt), 8-2
  - int (internal interrupt), 8-2
  - mcp (machine check processor), 8-2
  - smi (system management interrupt), 8-2
- JTAG interface, 20-1-20-3
- LBC
  - LA[27:31] (non-multiplexed address), 10-6
  - LAD[0:31] (multiplexed address/data), 10-7
  - LALE (external address latch enable), 10-5, 10-41
  - LBCTL (data buffer control), 10-6, 10-44
  - LBS[0:3] (UPM byte select), 10-5
  - LCK[0:2] (clock), 10-7
  - LCS[0:7] (chip select), 10-5
  - LCS0 (LBC chip select 0), 10-56, 10-69
  - LGPL0 (GP line 0), 10-6
  - LGPL1 (GP line 1), 10-6
  - LGPL2 (GP line 2), 10-6
  - LGPL3 (GP line 3), 10-6
  - LGPL4 (GP line 4), 10-6
  - LGPL5 (GP line 5), 10-6
  - LGTA (GPCM transfer acknowledge), 10-6, 10-55
  - LOE (GPCM output enable), 10-6
  - LWE[0:3] (GPCM write enable), 10-5
  - MDVAL (debug mode data valid), 10-7
  - MSRCID[0:4] (debug source ID), 10-7
  - TA (data transfer acknowledge), 10-43
  - UPWAIT (UPM wait), 10-6, 10-73
- muxing, 3-14
- output, states during reset, 3-13
- overview, 3-1, 7-37
- PCI, 13-4-13-11
  - INTA, 12-15
  - PCI\_C/BE[7:0] (command/byte enable), 13-6
  - PCI\_PERR (parity error), 13-9
  - PCI\_PME (PME signal), 13-9
  - PCI\_REQ[4:0] (bus request), 13-9
  - PCI\_SERR (system error), 13-10
  - PCI\_STOP (stop), 13-10
  - PCI\_TRDY (target ready), 13-11
- reset, 4-1-4-2
  - configuration, 4-9
  - reset and clocking blocks, 5-51
  - USB interface, 16-3-16-5
    - see also* USB interface, signals
- Single-step
  - trace enable (SE), 7-18
- smi (system management interrupt signal), 8-2
- Snapshot arbiters, 14-65
- Software watchdog timer, *see* Watchdog timer (WDT)
- SPI, 1-18
  - block diagram, 19-2
  - features, 19-2
  - initialization/application information, 19-16
    - master programming example, 19-16
    - slave programming example, 19-16
  - memory map/register definition, 19-8
  - modes of operation, 19-3
    - master, 19-3
    - multiple-master, 19-5
    - slave, 19-4
  - registers, 19-9-19-16
    - by acronym, *see* Register Index
  - signals, 19-6-19-8
    - descriptions, 19-7
    - overview, 19-7
  - transmission and reception process, 19-3
- Split transaction isochronous transfer descriptor (siTD), *see* USB interface, split transaction isochronous transfer descriptor (siTD)
- Split transactions, *see* USB interface, split transactions
- SPRG0-SPRG7, 7-3
- SPRs (special purpose registers), 7-16-7-20
- SRESET, 4-9
- SRn (segment registers 0-15), 7-19
- Supervisor-level SPRs, 7-19
- SYS\_CLK\_IN signal, 4-3
- System call (sc), 7-33
- System configuration
  - local memory map overview, 5-1
  - overview, 5-1
  - registers, 5-17
- System interface, 16-45

System management interrupt ( $\overline{smi}$ ), 7-34  
 System timers  
   overview, 1-18

## T

TA (LBC data transfer acknowledge) signal, 10-43  
 TAP (test access port) controller, 20-4  
 TBL/TBU (time base facility)  
   for reading, 7-17  
   time base/decrementer, 7-11  
 True little-endian, 7-24  
 TSEC  
   overview, 1-11  
 TSEC<sub>n</sub>\_COL (eTSEC 1–4 collision input) signals, 15-8  
 TSEC<sub>n</sub>\_CRS (eTSEC 1–4 carrier sense input/FIFO receiver flow control) signals, 15-8  
 TSEC<sub>n</sub>\_GTX\_CLK (eTSEC 1–4 gigabit transmit clock) signals, 15-8  
 TSEC<sub>n</sub>\_RX\_CLK (eTSEC 1–4 receive clock) signals, 15-9  
 TSEC<sub>n</sub>\_RX\_DV (eTSEC 1–4 receive data valid) signals, 15-9  
 TSEC<sub>n</sub>\_RX\_ER (eTSEC 1–4 receive error) signals, 15-10  
 TSEC<sub>n</sub>\_RXD[7:0] (eTSEC 1–4 receive data in) signals, 15-9  
 TSEC<sub>n</sub>\_TX\_CLK (eTSEC 1–4 transmit clock in) signals, 15-10  
 TSEC<sub>n</sub>\_TX\_EN (eTSEC 1–4 transmit data valid) signals, 15-10  
 TSEC<sub>n</sub>\_TX\_ER (eTSEC 1–4 transmit error) signals, 15-10  
 TSEC<sub>n</sub>\_TXD[7:0] (eTSEC 1–4 transmit data out) signals, 15-10

## U

Universal serial bus, *see* USB interface  
 UPWAIT (LBC UPM wait) signal, 10-6, 10-73  
 USB interface  
   asynchronous schedule, 16-47, 16-80  
   adding queue heads, 16-81  
   empty detection, 16-84  
   list queue head pointer, 16-49  
   organization, 16-49  
   reclamation status bit, 16-85  
   removing queue heads, 16-82  
   traversal (start event), 16-85  
 dual-role controller  
   overview, 1-14  
 features, 16-2  
 frame list  
   link pointer format, 16-48  
 functional descriptions  
   DMA engine, 16-46

FIFO RAM controller, 16-46  
 PHY interface, 16-46  
   system interface, 16-45  
 host data structures, 16-47  
 host operations, 16-67  
   behavior during wake-up events, 16-71  
   host controller initialization, 16-68  
   periodic schedule, 16-75  
   periodic schedule frame boundaries vs. bus frame boundaries, 16-72  
   port power control (PPC), 16-69  
   port suspend/resume, 16-70  
   reporting over-current, 16-69  
   schedule traversal rules, 16-71  
   split transactions, 16-90  
     execution state machine for isochronous, 16-110  
     isochronous, 16-104  
     periodic interrupt—Do-Complete-Split, 16-100  
     periodic interrupt—Do-Start-Split, 16-99  
     scheduling mechanisms for isochronous, 16-105  
     *see also* USB interface, split transactions  
     tests for execution, 16-100  
     tracking progress for isochronous transfers, 16-108  
   suspend/resume, 16-69  
 isochronous (high-speed) transfer descriptor (iTD), 16-47, 16-49  
   buffer page pointer list (plus), 16-51  
   host controller operational model, 16-76  
   managing transfers, 16-76  
   next link pointer, 16-50  
   periodic scheduling threshold, 16-79  
   software operational model, 16-78  
   transaction status and control list, 16-51  
 memory map/register definitions, 16-5  
 modes of operation, 16-3  
 module, 16-2  
 modules  
   dual-role (DR), 16-2  
 overview, 1-13, 16-2  
 periodic frame list, 16-47  
 periodic frame span traversal node (FSTN), 16-66  
   back path link pointer, 16-67  
   normal path pointer, 16-67  
 periodic schedule, 16-47  
   organization, 16-48  
   rebalancing, 16-104  
 queue element transfer descriptor (qTD), 16-47, 16-56  
   alternate next pointer, 16-57  
   buffer page pointer list, 16-61  
   next pointer, 16-57  
   token, 16-58

- queue heads, 16-47, 16-62
  - endpoint capabilities/characteristics, 16-63
  - horizontal link pointer, 16-62
  - managing control/bulk/interrupt transfers, 16-85
    - buffer pointer list use for data streaming with qTDs, 16-86
    - ping control, 16-89
    - to the periodic schedule, 16-88
    - transfer complete interrupts, 16-88
  - managing the QH[FrameTag] field, 16-103
  - transfer overlay, 16-65
- registers
  - by acronym, *see* Register Index
  - capability registers, 16-7
  - operational registers, 16-11
- signals, 16-3–16-5
  - ULPI interface, 16-4
  - UTMI interface, 16-4
- split transaction isochronous transfer descriptor (siTD), 16-47, 16-53
  - back link pointer, 16-56
  - buffer pointer list (plus), 16-55
  - endpoint capabilities/characteristics, 16-53
  - next link pointer, 16-53
  - transfer state, 16-54

- split transactions, 16-90
  - asynchronous transfers, 16-90
    - Do-Complete-Split state, 16-91
    - Do-Start-Split state, 16-91
  - interrupt, 16-92
    - execution state machine, 16-98
    - host controller operational model for FSTNs, 16-95
    - scheduling mechanisms, 16-92
    - software operational model for FSTNs, 16-97
    - tracking progress for interrupt transfers, 16-98
- USB\_CLK\_IN signal, 4-3
- User-level SPRs, 7-16

## W

- Watchdog timer (WDT), 5-29
  - block diagram, 5-29
  - features, 5-29
  - functional block diagram, 5-34
  - functional description, 5-33
  - memory map/register definition, 5-30
  - modes of operation, 5-30, 5-34
  - overview, 5-29
  - registers, 5-31–5-33

## Z

- ZBT SRAM interface (LBC), 10-97, 10-102

