



# Intel<sup>®</sup> 820 Chipset Family: 82820 Memory Controller Hub (MCH)

Datasheet

---

*July 2000*

Order Number: [290630-002](#)





Information in this document is provided in connection with Intel products. No license, express or implied, by estoppel or otherwise, to any intellectual property rights is granted by this document. Except as provided in Intel's Terms and Conditions of Sale for such products, Intel assumes no liability whatsoever, and Intel disclaims any express or implied warranty, relating to sale and/or use of Intel products including liability or warranties relating to fitness for a particular purpose, merchantability, or infringement of any patent, copyright or other intellectual property right. Intel products are not intended for use in medical, life saving, or life sustaining applications.

Intel may make changes to specifications and product descriptions at any time, without notice.

Designers must not rely on the absence or characteristics of any features or instructions marked "reserved" or "undefined." Intel reserves these for future definition and shall have no responsibility whatsoever for conflicts or incompatibilities arising from future changes to them.

The Intel 82820 Memory Controller Hub (MCH) may contain design defects or errors known as errata which may cause the product to deviate from published specifications. Current characterized errata are available on request.

Contact your local Intel sales office or your distributor to obtain the latest specifications and before placing your product order.

I<sup>2</sup>C is a 2-wire communications bus/protocol developed by Philips. SMBus is a subset of the I<sup>2</sup>C bus/protocol and was developed by Intel. Implementations of the I<sup>2</sup>C bus/protocol may require licenses from various entities, including Philips Electronics N.V. and North American Philips Corporation.

Alert on LAN is a result of the Intel-IBM Advanced Manageability Alliance and a trademark of IBM

Copies of documents which have an ordering number and are referenced in this document, or other Intel literature, may be obtained from:

Intel Corporation  
www.intel.com  
or call 1-800-548-4725

\*Third-party brands and names are the property of their respective owners.

Copyright © Intel Corporation 2000

# Contents

1.	Overview .....	13
1.1.	Related Documents .....	13
1.2.	The Intel® 820 Chipset System.....	14
1.3.	Intel® 82820 MCH Overview .....	16
2.	Signal Description .....	19
2.1.	Host Interface Signals .....	20
2.2.	Direct RDRAM Interface Signals.....	22
2.3.	The Hub Interface Signals .....	22
2.4.	AGP Interface Signals.....	23
2.4.1.	AGP Addressing Signals.....	23
2.4.2.	AGP Flow Control Signals.....	24
2.4.3.	AGP Status Signals .....	24
2.4.4.	AGP Clocking Signals (Strobes) .....	25
2.4.5.	AGP FRAME# Signals .....	26
2.5.	Clock and Reset Signals .....	28
2.6.	Voltage References, Power, Ground, and Test Signals .....	29
2.7.	Strap Signals.....	29
2.8.	Pin States During Reset.....	30
3.	Register Description .....	31
3.1.	Register Nomenclature and Access Attributes .....	31
3.2.	PCI Configuration Space Access .....	32
3.3.	I/O Mapped Registers .....	34
3.3.1.	CONF_ADDR—Configuration Address Register .....	34
3.3.2.	CONF_DATA—Configuration Data Register .....	35
3.4.	Host-Hub Interface Bridge/DRAM Controller Device Registers (Device 0).....	36
3.4.1.	VID—Vendor Identification Register (Device 0) .....	38
3.4.2.	DID—Device Identification Register (Device 0) .....	38
3.4.3.	PCICMD—PCI Command Register (Device 0).....	39
3.4.4.	PCISTS—PCI Status Register (Device 0) .....	40
3.4.5.	RID—Revision Identification Register (Device 0) .....	41
3.4.6.	SUBC—Sub-Class Code Register (Device 0) .....	41
3.4.7.	BCC—Base Class Code Register (Device 0) .....	41
3.4.8.	MLT—Master Latency Timer Register (Device 0) .....	42
3.4.9.	HDR—Header Type Register (Device 0) .....	42
3.4.10.	APBASE—Aperture Base Configuration Register (Device 0).....	42
3.4.11.	SVID—Subsystem Vendor ID Register (Device 0) .....	44
3.4.12.	SID—Subsystem ID Register (Device 0) .....	44
3.4.13.	CAPPTR—Capabilities Pointer Register (Device 0) .....	44
3.4.14.	GAR0—GAR7—RDRAM Group Architecture Register (Device 0) .....	45
3.4.15.	RDTR—RDRAM Timing Register (Device 0) .....	46
3.4.16.	RDCR—RDRAM Control Register (Device 0) .....	47
3.4.17.	RDRR—RDRAM Refresh Register (Device 0) .....	47
3.4.18.	RPMR—RDRAM Power Management Register (Device 0).....	48
3.4.19.	FDHC—Fixed DRAM Hole Control Register (Device 0) .....	49
3.4.20.	PAM0—PAM6—Programmable Attribute Map Registers (Device 0) .....	49
3.4.21.	GBA0—GBA7—RDRAM Group Boundary Address Register (Device 0).....	53
3.4.22.	DTC—DRAM Throttle Control Register (Device 0).....	55

3.4.23.	DRD—RDRAM Device Register Data Register (Device 0) .....	56
3.4.24.	RICM—RDRAM Initialization Control Management Register (Device 0) ....	57
3.4.25.		59
3.4.25.	SMRAM—System Management RAM Control Register (Device 0) .....	59
3.4.26.	ESMRAMC—Extended System Management RAM Control Register (Device 0)	60
3.4.27.	ACAPID—AGP Capability Identifier Register (Device 0).....	61
3.4.28.	AGPSTAT—AGP Status Register (Device 0) .....	62
3.4.29.	AGPCMD—AGP Command Register (Device 0) .....	63
3.4.30.	AGPCTRL—AGP Control Register (Device 0).....	64
3.4.31.	APSIZE—Aperture Size Register (Device 0).....	65
3.4.32.	ATTBASE—Aperture Translation Table Base Register (Device 0).....	66
3.4.33.	AMTT—AGP Interface Multi-Transaction Timer Register (Device 0) .....	66
3.4.34.	LPTT—Low Priority Transaction Timer Register (Device 0) .....	67
3.4.35.	MCHCFG—MCH Configuration Register (Device 0).....	67
3.4.36.	EAP—Error Address Pointer Register (Device 0) .....	69
3.4.37.	ERRSTS—Error Status Register (Device 0) .....	70
3.4.38.	ERRCMD—Error Command Register (Device 0) .....	72
3.4.39.	SMICMD—SMI Command Register (Device 0) .....	74
3.4.40.	SCICMD—SCI Command Register (Device 0) .....	75
3.4.41.	SKPD—Scratchpad Data Register (Device 0) .....	75
3.4.42.	AGPBCtrl—AGP Buffer Control B Register (Device 0) .....	76
3.4.43.	RTCE—RDRAM Temperature Calibration Enable Register (Device 0).....	76
3.4.44.	AGPAPPEND—AGP Append Disable Register (Device 0).....	76
3.5.	AGP Bridge Registers (Device 1) .....	77
3.5.1.	VID1—Vendor Identification Register (Device 1) .....	78
3.5.2.	DID1—Device Identification Register (Device 1).....	78
3.5.3.	PCICMD1—PCI-PCI Command Register (Device 1) .....	79
3.5.4.	PCISTS1—PCI-PCI Status Register (Device 1).....	80
3.5.5.	RID1—Revision Identification Register (Device 1).....	81
3.5.6.	SUBC1—Sub-Class Code Register (Device 1).....	81
3.5.7.	BCC1—Base Class Code Register (Device 1).....	81
3.5.8.	MLT1—Master Latency Timer Register (Device 1) .....	82
3.5.9.	HDR1—Header Type Register (Device 1).....	82
3.5.10.	PBUSN—Primary Bus Number Register (Device 1) .....	82
3.5.11.	SBUSN—Secondary Bus Number Register (Device 1) .....	83
3.5.12.	SUBUSN—Subordinate Bus Number Register (Device 1).....	83
3.5.13.	SMLT—Secondary Master Latency Timer Register (Device 1) .....	84
3.5.14.	IOBASE—I/O Base Address Register (Device 1) .....	85
3.5.15.	IOLIMIT—I/O Limit Address Register (Device 1) .....	85
3.5.16.	SSTS—Secondary PCI-PCI Status Register (Device 1) .....	86
3.5.17.	MBASE—Memory Base Address Register (Device 1) .....	87
3.5.18.	MLIMIT—Memory Limit Address Register (Device 1).....	87
3.5.19.	PMBASE—Prefetchable Memory Base Address Register (Device 1) .....	88
3.5.20.	PMLIMIT—Prefetchable Memory Limit Address Register (Device 1) .....	88
3.5.21.	BCTRL—PCI-PCI Bridge Control Register (Device 1) .....	89
3.5.22.	ERRCMD1—Error Command Register (Device 1) .....	90

4.	System Address Map.....	91
4.1.	Memory Address Ranges.....	91
4.1.1.	Compatibility Area.....	92
4.1.2.	Extended Memory Area.....	95
4.1.3.	AGP Memory Address Ranges.....	97
4.1.4.	AGP DRAM Graphics Aperture.....	97
4.1.5.	System Management Mode (SMM) Memory Range.....	98
4.1.6.	Memory Shadowing.....	100
4.2.	I/O Address Space.....	100
4.3.	MCH Decode Rules and Cross-Bridge Address Mapping.....	102
4.3.1.	The Hub Interface Decode Rules.....	102
4.3.2.	AGP Interface Decode Rules.....	102
4.3.3.	Legacy VGA Ranges.....	103
5.	Functional Description.....	105
5.1.	Host Interface.....	105
5.1.1.	Host Addresses Above 4 GB.....	107
5.1.2.	Host Bus Cycles.....	107
5.1.3.	Symmetric Multiprocessor (SMP) Protocol Support.....	111
5.1.4.	Frame Buffer Memory Support.....	111
5.2.	AGP Interface.....	112
5.2.1.	AGP Target Operations.....	112
5.2.2.	AGP Transaction Ordering.....	113
5.2.3.	AGP Signal Levels.....	113
5.2.4.	4x AGP Protocol.....	114
5.2.5.	Fast Writes.....	114
5.2.6.	AGP Universal Connector.....	115
5.2.7.	AGP FRAME# Transactions on AGP.....	115
5.3.	DRAM Interface.....	117
5.3.1.	RDRAM.....	117
5.3.1.1.	DRAM Organization, Configuration, and Speed.....	118
5.3.1.1.1.	Rules for Populating RDRAM Devices.....	119
5.3.1.1.2.	RDRAM CMOS Signals Description and Usage.....	120
5.3.1.1.3.	Direct RDRAM Core Refresh.....	121
5.3.1.1.4.	Direct RDRAM Current Calibration.....	122
5.3.1.2.	Direct RDRAM Command Encoding.....	123
5.3.1.2.1.	Row Packet (ROWA/ROWR).....	123
5.3.1.2.2.	Column Packet (COLC/COLX/COLM).....	124
5.3.1.2.3.	Data Packet.....	126
5.3.1.3.	Direct RDRAM Register Programming.....	127
5.3.1.4.	Direct RDRAM Operating States.....	127
5.3.1.5.	RDRAM Operating Pools.....	128
5.3.1.6.	RDRAM Power Management.....	128
5.3.1.7.	Data Integrity.....	129
5.3.1.8.	RDRAM Array Power Management.....	129
5.3.1.8.1.	RDRAM On-die Internal Thermal Sensor Mechanism (Method 1).....	130
5.3.1.8.2.	RDRAM Reads and Writes Counters/Timers Mechanism (Method 2).....	130
5.4.	Power Management.....	132
5.4.1.	Processor Power State Control.....	132
5.4.2.	Sleep State Control.....	132
5.5.	MCH Clocking.....	133



5.6.	MCH System Reset and Power Sequencing .....	134
5.6.1.	MCH Reset .....	135
5.6.2.	MCH Power Sequence Recommendation.....	135
6.	Pinout and Package Information .....	137
6.1.	Pinout Information .....	137
6.2.	Package Specifications .....	143
6.3.	MCH RSL Package Dimensions.....	145
7.	Testability .....	147

## Figures

Figure 1. Intel® 820 Chipset Family System Block Diagram.....	15
Figure 3. PAM Registers.....	51
Figure 4. Detailed Memory System Address Map .....	92
Figure 5. MCH/ Direct RDRAM Diagram .....	117
Figure 6. Sideband CMOS Signal Configuration .....	120
Figure 7. RDRAM Throttle Time (With Counters/Timers Mechanism).....	131
Figure 8. Intel® 82820 MCH Clocking Diagram .....	133
Figure 9. Intel® 82820 MCH Reset Overview .....	134
Figure 10. MCH Power-Up Voltage Sequencing .....	135
Figure 11. MCH Ballout (Top View—Left Side) .....	138
Figure 12. MCH Ballout (Top View—Right Side).....	139
Figure 13. Package Dimensions (324 BGA) – Top and Side Views.....	143
Figure 14. Package Dimensions (324 BGA) – Bottom View .....	144
Figure 15. Initialization Sequence.....	148
Figure 16. Rule for Controlling STB Signal Pairs During XOR Testing .....	148

## Tables

Table 1. Unsupported Host Bus Signals .....	21
Table 2. MCH Strapping Options .....	29
Table 3. Signal States During Reset .....	30
Table 4. MCH PCI Configuration Space (Device 0) .....	36
Table 5. Attribute Bit Assignment.....	50
Table 6. PAM Registers and Associated Memory Segments .....	51
Table 7. MCH Configuration Space (Device 1) .....	77
Table 8. Memory Segments and Attributes.....	93
Table 9. SMM Space Table.....	99
Table 10. SMM Control Table .....	100
Table 11. Types of Transactions Supported by the MCH .....	105
Table 12. Types of Responses Supported by the MCH.....	107
Table 13. Types of Special Cycles Supported by the MCH .....	110
Table 14. AGP Commands Supported by the MCH When Acting as a AGP Target.....	112
Table 15. AGP Data Rates and Signaling Levels.....	113
Table 16. PCI Commands Supported by the MCH When Acting as A FRAME# Target .....	115
Table 17. Maximum Supported Direct RDRAM Configurations .....	118
Table 18. Direct RDRAM Device Configurations .....	118
Table 19. Supported Direct Rambus Channel Speeds .....	119
Table 20. RDRAM Device Grouping .....	119
Table 21. Sideband CMOS Signal Description .....	120
Table 22. CMD Signal Value Decode.....	120
Table 23. RDRAM Refresh Rate.....	121
Table 24. ROWA Packet for Activating (sensing) a Row (i.e., AV = 1).....	123
Table 25. ROWR Packet for Other Operations (i.e., AV = 0).....	123
Table 26. ROWA and ROWR Packet Field Encodings.....	124
Table 27. COLC Packet .....	125
Table 28. COLC Packet Field Encodings.....	125
Table 29. COLM Packet (M = 1) .....	125
Table 30. COLX Packet (M = 0).....	126
Table 31. COLM Packet and COLX Packet Field Encodings .....	126
Table 32. Data Packet.....	126
Table 33. DRAM Operating States.....	127
Table 34. RDRAM Power Management States.....	128
Table 38. MCH Alphabetical Ball List.....	140
Table 39. BGA Package Dimensions (324 BGA).....	144
Table 40. MCH RSL Package Dimensions .....	145
Table 41. XOR Chain #1 Connections .....	149
Table 42. XOR Chain #2 Connections .....	150
Table 43. XOR Chain #3 Connections .....	151
Table 44. XOR Chain #4 Connections .....	152
Table 45. XOR Chain #5 Connections .....	152



## Revision History

Rev.	Description	Date
-001	<ul style="list-style-type: none"><li>• Initial Release</li></ul>	September 1999
-002	<ul style="list-style-type: none"><li>• Added Intel® 820E chipset. This change only affected the Overview chapter.</li><li>• Removed SDRAM support</li><li>• Minor edits throughout for clarity.</li></ul>	July 2000

This page is intentionally left blank.

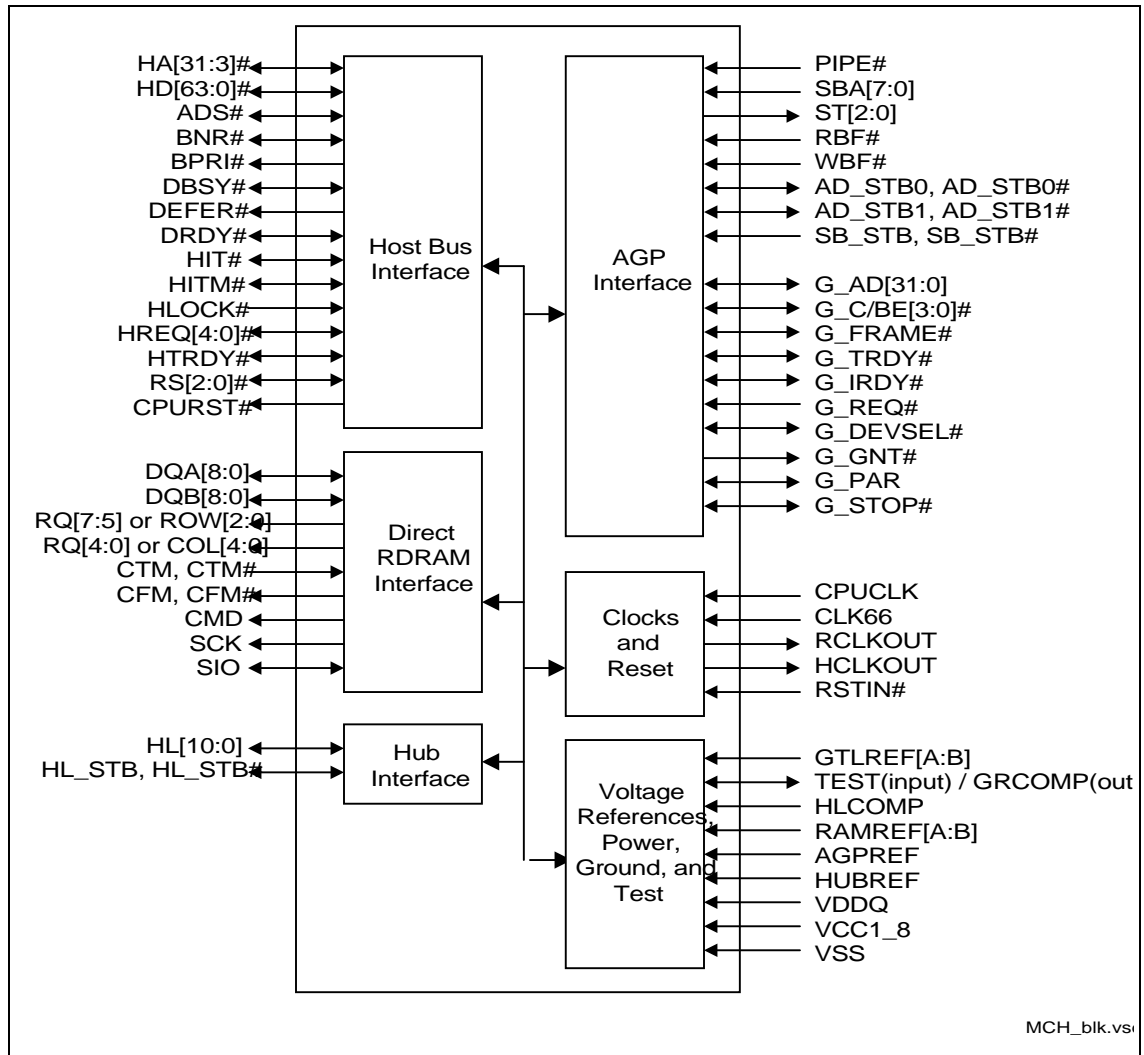
# Intel® 82820 MCH

## Product Features

- Processor/Host Bus Support
  - Supports Intel® Pentium® II or Intel® Pentium® III processor configuration at 100 MHz / 133 MHz
  - Supports Symmetric Multiprocessing Protocol (SMP) for up to two processors
  - APIC related buffer management support
  - Supports 32-bit host bus addressing
  - Supports 6 deep In-Order Queue
  - AGTL+ bus driver technology (gated AGTL+ receivers for reduced power)
  - Supports single-ended AGTL+ termination in uniprocessor configuration
- Direct Rambus Memory Controller
  - Directly support a single Direct Rambus\* Channel
  - Supports PC600, PC700, and PC800 Direct Rambus DRAM modules
  - Maximum memory array size up to 256 MB using 64Mb/72Mb, 512 MB using 128Mb/144Mb, 1 GB using 256Mb/288Mb DRAM technology
  - Supports up to 32 Direct Rambus devices per channel
  - Supports a maximum DRAM address decode space of 4 GB
  - Configurable optional ECC operation:
    - ECC with single bit Error Correction and multiple bit Error Detection
    - Single bit errors corrected and written back to memory (auto-scrubbing)
    - Parity mode not supported
- The Hub Interface
  - Private interconnect between MCH and I/O Controller Hub
- Arbitration
  - Distributed Arbitration Model for Optimum Concurrency Support
  - Concurrent operations of host, hub interface, AGP, and memory buses supported via dedicated arbitration and data buffering logic
- Accelerated Graphics Port (AGP) Interface
  - Supports a single AGP device (either via a connector or on the motherboard)
  - Supports AGP 2.0 including 1x/2x/4x AGP data transfers, 1.5V and 3.3V signaling, and 2x/4x Fast Write protocol
  - AGP Universal Connector support via dual mode buffers
  - AGP PIPE# or SB accesses to DRAM not snooped
  - AGP FRAME# accesses to DRAM are snooped
  - High priority access support
  - Hierarchical PCI configuration mechanism
  - Delayed transaction support for AGP-to-DRAM reads using AGP FRAME# protocol
- Power Management
  - SMRAM space re-mapping to A0000h (128 KB)
  - Supports HSEG and TSEG cacheable extended SMRAM space
    - 128 KB HSEG at 0FEEA0000h remapped to A0000h
    - TSEG is 128 KB/256 KB/512 KB/1 MB at the top of memory
    - No maximum DRAM limit to use extended SMM
  - SMRAM accesses from AGP or the hub interface are not allowed
  - Suspend to DRAM (STR) support
  - ACPI Rev 1.0 compliant power management
  - APM Rev 1.2 compliant power management
- I/O Device Support
  - Intel® 82801AA ICH (Intel® 820 chipset)
  - Intel® 82801BA ICH2 (Intel® 820E chipset)
- Package: 324-pin BGA

The Intel® 82820 MCH may contain design defects or errors known as errata, which may cause the product to deviate from published specifications. Current characterized errata are available on request.

Simplified Block Diagram



# 1. Overview

---

The Intel® 820 chipset family is a high-bandwidth chipset designed as the next generation, high performance, longlife PC platform. The main components of the chipset are the Intel® 82820 Memory Controller Hub (MCH) host bridge and the I/O Controller Hub for the I/O subsystem.

There are two chipsets in the Intel® 820 chipset family:

- Intel® 820 chipset: This chipset contains the 82820 MCH and 82801AA ICH.
- Intel® 820E chipset: This chipset contains the 82820 MCH and 82801BA ICH2.

**Note:** The only component difference between the Intel® 820 chipset and the Intel® 820E chipset is the I/O Controller Hub.

This datasheet provides an overview of the 820 chipset family (see Section 1.2). The remainder of the document describes the Intel® 82820 Memory Controller Hub (MCH).

The Intel® 82820 MCH may contain design defects or errors known as errata which may cause the product to deviate from published specifications. Current characterized errata are available on request.

## 1.1. Related Documents

- *Intel® 82801AA (ICH) and Intel® 82801AB (ICH0) I/O Controller Hub Datasheet* (Doc Ref 290655)
- *Intel® 82801BA I/O Controller Hub (ICH2) Datasheet* (Doc Ref 290687)
- *Intel® 82802AB/AC Firmware Hub (FWH) Datasheet* (Doc Ref 290658)
- *Intel® 820 Chipset Design Guide* (Doc Ref 290631)
- *Intel® 820E Chipset Design Guide* (Doc Ref 298187)


## 1.2. The Intel® 820 Chipset System

Figure 1 show a typical system block diagram based on the Intel® 820 chipset family. The Intel® 820 chipset uses a hub architecture with the MCH as the host bridge hub and the I/O Controller Hub as the I/O hub. The MCH supports processor bus frequencies of 100/133 MHz. The I/O Controller Hub is highly integrated providing many of the functions needed in today's PC platforms; it also provides the interface to the PCI Bus. The MCH and I/O Controller Hub communicate over a dedicated hub interface.

82801AA ICH and 82801BA ICH2 functions include:

- PCI Rev 2.2 compliant with support for 33 MHz PCI operations
- Supports up to 6 Req/Gnt pairs (PCI Slots)
- Power management logic support
- Enhanced DMA controller, interrupt controller, and timer functions
- Integrated IDE controller
  - Ultra ATA/66/33 (ICH)
  - Ultra ATA/100/66/33 (ICH2)
- USB host interface
  - 1 host controller and supports 2 USB ports (ICH)
  - 2 host controllers and supports 4 USB ports (ICH2)
- Integrated LAN controller (ICH2 only)
- System Management Bus (SMBus) compatible with most I<sup>2</sup>C devices
  - ICH has bus master capability
  - ICH2 has both bus master and slave capability
- AC'97 2.1 compliant link for audio and telephony codecs
  - 2 channels (ICH)
  - Up to 6 channels (ICH2)
- Low Pin Count (LPC) interface
- FWH Interface (FWH Flash BIOS support)
  - Intel's FWH component is the 82802: contains Random Number Generator (RNG), five General Purpose Inputs (GPIs), register-based block locking, hardware-based locking, and Flash memory for platform code/data nonvolatile storage
  - FWH component is also available from other suppliers
- Alert on LAN\*
  - AOL (ICH and ICH2)
  - AOL2 (ICH2 only)

Figure 1. Intel® 820 Chipset Family System Block Diagram



## 1.3. Intel® 82820 MCH Overview

The Intel® 82820 Memory Controller Hub (MCH) provides the host interface, DRAM interface, I/O interface, and AGP interface in an Intel® 820/820E chipset platform. The MCH is optimized for both the Intel® Pentium® II processor and the Intel® Pentium® III processor and supports one or two processor configurations. The MCH supports a single channel of Direct Rambus memory technology and the AGP interface is fully compliant with the AGP 2.0 specification. Communication to the I/O Controller Hub is made through a private interface (called “hub interface”). The Intel® 82820 MCH includes the following functions:

- Supports Intel® Pentium® II processor and Intel® Pentium® III processor at 100 MHz/133 MHz
- Supports one or two processor configurations
- Supports an IOQ depth of 6
- Supports 32-bit host addressing using AGTL+ host bus
- Supports PC600, PC700, and PC800 Direct Rambus DRAM modules
- Supports 1GB DRAM with 256Mbit memory technology
- Supports AGP 2.0 compliant 1x/2x/4x data transfer and 2x/4x fast write capability
- Supports hub interface to the I/O Controller Hub
- Features fully optimized data paths and buffering
- Features distributed arbitration for highly concurrent operation
- Compliant with ACPI 1.0 power management

### Host Interface

The MCH supports one or two processor configurations for slot 1. The MCH supports a Processor System Bus frequency of 100 MHz or 133 MHz. Single-ended termination AGTL+ is supported for the single processor configuration. It supports 32-bit host addressing, decoding up to 4 GB of the host memory address space. The MCH has a 6-deep In-Order Queue to support up to six outstanding pipelined address requests on the host bus. Host initiated I/O cycles are decoded to AGP, the hub interface, or MCH configuration space. Host initiated memory cycles are decoded to AGP, the hub interface, or DRAM. Host memory accesses targeting the AGP aperture are translated using the AGP address translation table. All accesses from AGP and the hub interface that hit the graphics aperture are also translated using the AGP address translation table. AGP semantic memory accesses initiated from AGP to DRAM are not snooped on the host bus. Memory accesses using AGP FRAME# protocol from AGP or memory access from the hub interface to DRAM will be snooped on the host bus.



## DRAM Interface

The MCH supports a single channel of Direct RDRAM memory using Rambus Signaling Level (RSL) technology. It also supports 64, 128 and 256Mbit technology along with PC600, PC700, and PC800 Direct Rambus DRAM devices. A maximum of 32 Direct RDRAM devices (64Mbit technology = 256 MB max) are supported. The following table shows the maximum DRAM array size and the minimum increment size for the various DRAM densities supported for MCH.

RDRAM Technology	Increments	Maximum
64Mb/72Mb	8 MB	256 MB
128Mb/144Mb	16 MB	512 MB
256Mb/288Mb	32 MB	1 GB

The MCH provides optional ECC error checking for DRAM data integrity. During DRAM writes, ECC is generated on a QWord (64-bit) basis. Partial QWord writes require a read-modify-write cycle when ECC is enabled. During DRAM reads, the MCH supports detection of single-bit and multiple-bit errors, and will correct single bit errors when correction is enabled. The MCH automatically scrubs single bit errors by writing the corrected value back into DRAM when scrubbing is enabled. ECC can only be enabled when the Direct RDRAMs support the extra two data bits used to store the ECC code.

The MCH provides a maximum DRAM address decode space of 4 GB. The MCH does not remap APIC memory space in hardware. It is the BIOS or system designer's responsibility to limit DRAM population so that adequate PCI, AGP, High BIOS, and APIC memory space can be allocated.

## AGP Interface

A single AGP component or connector (not both) is supported by the MCH AGP interface. To support the AGP Universal Connector, the AGP buffers operate in one of two selectable modes (3.3V mode and 1.5V mode).

- 3.3V mode: buffers are **not** 5 volt tolerant. (AGP 1.0 and 2.0 specifications).
- 1.5V mode: buffers are **not** 3.3 volt tolerant. (AGP 2.0 specification)

The AGP interface supports 4x AGP signaling and 4x Fast Writes. AGP semantic (PIPE# or SBA[7:0]) cycles to DRAM are not snooped on the host bus. AGP FRAME# cycles to DRAM are snooped on the host bus. The MCH supports PIPE# or SBA[7:0] AGP address mechanisms, but not both simultaneously. Either the PIPE# or the SBA[7:0] mechanism must be selected during system initialization. High priority accesses are supported. Only memory writes from the hub interface to AGP are allowed. No transactions from AGP to the hub interface are allowed.

## MCH Clocking

The MCH has two clock input pins: HCLKIN for the host clock and GCLKIN for the AGP clock. Clock Synthesizer chip(s) are responsible for generating the Host clocks, AGP clocks, PCI clocks, and Direct Rambus clocks. These clocks must be synchronous to each other (same clock source). The MCH Processor System Bus interface runs at 100 MHz or 133 MHz. The supported Direct Rambus\* interface speeds are 266 MHz, 300 MHz, 356 MHz, and 400 MHz. (300 MHz at a 100 MHz Processor System Bus frequency, 266 MHz / 356 MHz at 133 MHz Processor System Bus frequency, and 400 MHz at 100 MHz / 133 MHz Processor System Bus frequency.) The AGP interface runs at a constant 66 MHz. The hub interface runs at the same base frequency as the AGP interface.



This page is intentionally left blank.

## 2. Signal Description

---

This section provides a detailed description of MCH signals. The signals are arranged in functional groups according to their associated interface.

The “#” symbol at the end of a signal name indicates that the active, or asserted state occurs when the signal is at a low voltage level. When “#” is not present after the signal name the signal is asserted when at the high voltage level.

The following notations are used to describe the signal type:

<b>I</b>	Input pin
<b>O</b>	Output pin
<b>I/O</b>	Bi-directional Input/Output pin
<b>s/t/s</b>	Sustained Tristate. This pin is driven to its inactive state prior to tri-stating.
<b>as/t/s</b>	Active Sustained Tristate. This applies to some of the hub interface signals. This pin is weakly driven to its last driven value.

The signal description also includes the type of buffer used for the particular signal:

<b>AGTL+</b>	Open Drain AGTL+ interface signal. Refer to the AGTL+ I/O Specification for complete details.
<b>AGP</b>	AGP interface signals. These signals can be programmed to be compatible with AGP 2.0 3.3v or 1.5v Signaling Environment DC and AC Specifications. In 3.3v mode, the buffers are not 5v tolerant. In 1.5v mode, the buffers are not 3.3v tolerant.
<b>CMOS</b>	CMOS buffers.
<b>RSL</b>	Rambus Signaling Level interface signal. Note that RSL channel pins are logically inverted (like AGTL lines). In most cases, a logical value of 0 has a high pin voltage, and a logical value of 1 has a low pin voltage. Refer to the Rambus Specification for complete details.

## 2.1. Host Interface Signals

Name	Type	Description
CPURST#	O AGTL+	<b>CPU Reset:</b> The CPURST# pin is an output from the MCH. The MCH asserts CPURST# while RSTIN# (PCIRST# from the I/O Controller Hub) is asserted and for approximately 1 ms after RSTIN# is deasserted. The MCH also pulses CPURST# for approximately 1 ms when requested via the hub interface. The CPURST# allows the processors to begin execution in a known state.  Note that the I/O Controller Hub must provide processor strap set-up and hold times requirement around the deassertion of CPURST#.
HA[31:3]#	I/O AGTL+	<b>Host Address Bus:</b> HA[31:3]# are connect to the host address bus. During host cycles, HA[31:3]# are inputs. The MCH drives HA[31:3]# during snoop cycles on behalf of the hub interface and AGP FRAME# initiators. Note that the address bus is inverted on the host bus.
HD[63:0]#	I/O AGTL+	<b>Host Data:</b> These signals are connected to the host data bus. Note that the data signals are inverted on the host bus.
ADS#	I/O AGTL+	<b>Address Strobe:</b> The host bus owner asserts ADS# to indicate the first of two cycles of a request phase.
BNR#	I/O AGTL+	<b>Block Next Request:</b> Used to block the current request bus owner from issuing new requests. This signal is used to dynamically control the host bus pipeline depth.
BPRI#	O AGTL+	<b>Priority Agent Bus Request:</b> The MCH is the only Priority Agent on the host bus. It asserts this signal to obtain the ownership of the address bus. This signal has priority over symmetric bus requests and will cause the current symmetric owner to stop issuing new transactions unless the HLOCK# signal was asserted.
DBSY#	I/O AGTL+	<b>Data Bus Busy:</b> Used by the data bus owner to hold the data bus for transfers requiring more than one cycle.
DEFER#	O AGTL+	<b>Defer:</b> MCH generates a deferred response as defined by the rules of the MCH's dynamic defer policy. The MCH also uses the DEFER# signal to indicate a host retry response.
DRDY#	I/O AGTL+	<b>Data Ready:</b> Asserted for each cycle that data is transferred.
HIT#	I/O AGTL+	<b>Hit:</b> Indicates that a caching agent holds an unmodified version of the requested line. Also, driven in conjunction with HITM# by the target to extend the snoop window.
HITM#	I/O AGTL+	<b>Hit Modified:</b> Indicates that a caching agent holds a modified version of the requested line and that this agent assumes responsibility for providing the line. Also, driven in conjunction with HIT# to extend the snoop window.
HLOCK#	I AGTL+	<b>Host Lock:</b> All host bus cycles sampled with the assertion of HLOCK# and ADS#, until the negation of HLOCK# must be atomic (i.e., <i>no hub interface or AGP FRAME# snoopable access</i> to DRAM is allowed when HLOCK# is asserted by the processor).
HREQ[4:0]#	I/O AGTL+	<b>Host Request Command:</b> Asserted during both clocks of request phase. In the first clock, the signals define the transaction type to a level of detail that is sufficient to begin a snoop request. In the second clock, the signals carry additional information to define the complete transaction type. The transactions supported by the MCH Host Bridge are defined in Section 5.1, "Host Interface".
HTRDY#	I/O AGTL+	<b>Host Target Ready:</b> Indicates that the target of the processor transaction is able to enter the data transfer phase.

Name	Type	Description																		
RS[2:0]#	I/O AGTL+	<p><b>Response Signals:</b> Indicates type of response as shown below:</p> <table border="1"> <thead> <tr> <th>RS[2:0]</th> <th>Response type</th> </tr> </thead> <tbody> <tr> <td>000</td> <td>Idle state</td> </tr> <tr> <td>001</td> <td>Retry response</td> </tr> <tr> <td>010</td> <td>Deferred response</td> </tr> <tr> <td>011</td> <td>Reserved (not driven by MCH)</td> </tr> <tr> <td>100</td> <td>Hard Failure (not driven by MCH)</td> </tr> <tr> <td>101</td> <td>No data response</td> </tr> <tr> <td>110</td> <td>Implicit Writeback</td> </tr> <tr> <td>111</td> <td>Normal data response</td> </tr> </tbody> </table>	RS[2:0]	Response type	000	Idle state	001	Retry response	010	Deferred response	011	Reserved (not driven by MCH)	100	Hard Failure (not driven by MCH)	101	No data response	110	Implicit Writeback	111	Normal data response
RS[2:0]	Response type																			
000	Idle state																			
001	Retry response																			
010	Deferred response																			
011	Reserved (not driven by MCH)																			
100	Hard Failure (not driven by MCH)																			
101	No data response																			
110	Implicit Writeback																			
111	Normal data response																			

Table 1. lists the host bus interface signals that are NOT supported by the Intel® 820/820E chipset.

**Table 1. Unsupported Host Bus Signals**

Signal	Description	Function that MCH Does Not Support
A[35:32]#	Address	Extended addressing (over 4 GB)
AERR#	Address Parity Error	Parity protection on address bus
AP[1:0]#	Address Parity	Parity protection on address bus
BINIT#	Bus Initialization	Checking for bus protocol violation and protocol recovery mechanism
DEP[7:0]#	Data Bus ECC/Parity	Enhanced data bus integrity
IERR#	Internal Error	Direct internal error observation via IERR# pin
BREQ0#	Symmetric Agent Bus Request	UP processor strapped, DP processors driven from the I/O Controller Hub GPIO through external driver.
BERR#	Bus Error	Unrecoverable error without a bus protocol violation
RP#	Request Parity	Parity protection on ADS# and REQ[4:0]#
RSP#	Response Parity Signal	Parity protection on RS[2:0]#

## 2.2. Direct RDRAM Interface Signals

Name	Type	Description
DQA[8:0]	I/O RSL	<b>Data Byte A:</b> Nine data pins that carry a byte of read or write data between the MCH and Direct RDRAM. The ninth bit is used to store the ECC error code.
DQB[8:0]	I/O RSL	<b>Data Byte B:</b> Nine data pins that carry a byte of read or write data between the MCH and Direct RDRAM. The ninth bit is used to store the ECC error code.
RQ[7:5] or ROW[2:0]	O RSL	<b>Row Access Control:</b> Three request package pins containing control and address information for row accesses. Note RQ[7:5] can also be named as ROW[2:0].
RQ[4:0] or COL[4:0]	O RSL	<b>Column Access Control:</b> Five request package pins containing control and address information for column accesses. Note RQ[4:0] can also be named as COL[4:0].
CTM	I RSL	<b>Clock To Master:</b> One of the two differential transmit clock signals used for Direct RDRAM operation. It is an input to the MCH and is generated from an external clock synthesizer.
CTM#	I RSL	<b>Clock To Master Compliment:</b> One of the two differential transmit clock signals used for Direct RDRAM operation. It is an input to the MCH and is generated from an external clock synthesizer.
CFM	O RSL	<b>Clock From Master:</b> One of the two differential receive clock signals used for Direct RDRAM operation. It is an output from the MCH.
CFM#	O RSL	<b>Clock From Master Compliment:</b> One of the two differential receive clock signals used for Direct RDRAM operation. It is an output from the MCH.
CMD	O CMOS	<b>Command:</b> Command output to the Direct RDRAM devices used for power mode control, configuring the SIO daisy chain, and framing SIO operations.
SCK	O CMOS	<b>Serial Clock:</b> This signal provides clocking for register accesses and selects Direct RDRAM devices for power management.
SIO	I/O CMOS	<b>Serial Input/Output:</b> Bi-directional serial data signal used for device initialization, register operations, power mode control, and device reset.

## 2.3. The Hub Interface Signals

Name	Type	Description
HL[10:0]	I/O CMOS	<b>Hub interface:</b> Signals other than strobes used for the hub interface.
HL_STB, HL_STB#	I/O CMOS	<b>Hub interface Strobes:</b> Two differential strobes used for transmitting or receiving hub interface information.

## 2.4. AGP Interface Signals

The AGP interface includes addressing, flow control, status, clocking, and FRAME# signals. For more details on the operation of these signals, refer to the AGP Interface Specification Revision 2.0.

### 2.4.1. AGP Addressing Signals

There are two mechanisms by which the AGP master can enqueue AGP requests: PIPE# and SBA (side-band addressing). Upon initialization, one of the methods is chosen. The master may not switch methods without a full reset of the system. When PIPE# is used to enqueue addresses, the master is not allowed to queue addresses using the SBA bus. For example, during configuration time, if the master indicates that it can use either mechanism, the configuration software will indicate which mechanism the master will use. Once this choice has been made, the master will continue to use the mechanism selected until the system is reset (and reprogrammed) to use the other mode. This change of modes is not a dynamic mechanism but rather a static decision when the device is first being configured after reset.

Name	Type	Description
PIPE#	1 AGP	<p><b>Pipeline:</b></p> <p><b>PIPE# Operation:</b> This signal is asserted by the AGP master to indicate a full-width address is to be enqueued by the target using the AD bus. One address is placed in the AGP request queue on each rising clock edge while PIPE# is asserted.</p> <p><b>SBA Operation:</b> This signal is not used if SBA (Side Band Addressing) is selected.</p> <p><b>FRAME# Operation:</b> This signal is not used during AGP FRAME# operation.</p>
SBA[7:0]	1 AGP	<p><b>Side-band Addressing:</b></p> <p><b>PIPE# Operation:</b> These signals are not used during PIPE# operation.</p> <p><b>SBA Operation:</b> These signals (the SBA or Side-Band Addressing bus) are used by the AGP master (graphics component) to place addresses into the AGP request queue. The SBA bus and AD bus operate independently. That is, transaction can proceed on the SBA bus and the AD bus simultaneously.</p> <p><b>FRAME# Operation:</b> These signals are not used during AGP FRAME# operation.</p> <p>The SBA[7:0] signals have weak integrated pull-up resistors.</p>

## 2.4.2. AGP Flow Control Signals

Name	Type	Description
RBF#	I AGP	<p><b>Read Buffer Full:</b></p> <p><b>PIPE# and SBA Operation:</b> Read buffer full indicates if the master is ready to accept previously requested low priority read data. When RBF# is asserted, the MCH is not allowed to initiate the return low priority read data. Thus, the MCH can finish returning the data for the request currently being serviced; however, it can not begin returning data for the next request. RBF# is only sampled at the beginning of a cycle.</p> <p>If the AGP master is always ready to accept return read data, it is not required to implement this signal.</p> <p><b>FRAME# Operation:</b> This signal is not used during AGP FRAME# operation.</p>
WBF#	I AGP	<p><b>Write-Buffer Full:</b></p> <p><b>PIPE# and SBA Operation:</b> Write buffer full indicates if the master is ready to accept Fast Write data from the MCH. When WBF# is asserted the MCH is not allowed to drive Fast Write data to the AGP master. WBF# is only sampled at the beginning of a cycle.</p> <p>If the AGP master is always ready to accept fast write data, then it is not required to implement this signal.</p> <p><b>FRAME# Operation:</b> This signal is <b>not used</b> during AGP FRAME# operation.</p>

## 2.4.3. AGP Status Signals

Name	Type	Description
ST[2:0]	O AGP	<p><b>Status Bus:</b></p> <p><b>PIPE# and SBA Operation:</b> Provides information from the arbiter to an AGP Master on what it may do. ST[2:0] only have meaning to the master when its GNT# is asserted. When GNT# is deasserted, these signals have no meaning and must be ignored. Refer to the AGP Interface Specification revision 2.0 for further explanation of the ST[2:0] values and their meanings.</p> <p><b>FRAME# Operation:</b> ST[2:0] are not used during FRAME# based operation; except that a '111' indicates that the master may begin a FRAME# transaction.</p> <p>Note: Requires 8.2 K<math>\Omega</math> external pull-up resistors.</p>



## 2.4.4. AGP Clcking Signals (Strobes)

Name	Type	Description
AD_STB0	I/O s/t/s AGP	<p><b>AD Bus Strobe-0:</b></p> <p><b>1X Operation:</b> This signal is not used during 1X operation.</p> <p><b>2X Operation:</b> During 2X operation, this signal provides timing for the AD[15:0] and C/BE[1:0]# signals. The agent that is providing the data will drive this signal.</p> <p><b>4X Operation:</b> During 4X operation, this is one-half of a differential strobe pair that provides timing information for the AD[15:0] and C/BE[1:0]# signals.</p>
AD_STB0#	I/O s/t/s AGP	<p><b>AD Bus Strobe-0 Compliment:</b></p> <p><b>1X Operation:</b> This signal is not used during 1X operation.</p> <p><b>2X Operation:</b> During 2X operation, this signal is not used.</p> <p><b>4X Operation:</b> During 4X operation, this is one-half of a differential strobe pair that provides timing information for the AD[15:0] and C/BE[1:0]# signals. The agent that is providing the data will drive this signal.</p>
AD_STB1	I/O s/t/s AGP	<p><b>AD Bus Strobe-1:</b></p> <p><b>1X Operation:</b> This signal is not used during 1X operation.</p> <p><b>2X Operation:</b> During 2X operation, this signal provides timing for the AD[31:16] and C/BE[3:2]# signals. The agent that is providing the data will drive this signal.</p> <p><b>4X Operation:</b> During 4X operation, this is one-half of a differential strobe pair that provides timing information for the AD[31:16] and C/BE[3:2]# signals. The agent that is providing the data will drive this signal.</p>
AD_STB1#	I/O s/t/s AGP	<p><b>AD Bus Strobe-1 Compliment:</b></p> <p><b>1X Operation:</b> This signal is not used during 1X operation.</p> <p><b>2X Operation:</b> During 2X operation, this signal is not used</p> <p><b>4X Operation:</b> During 4X operation, this is one-half of a differential strobe pair that provides timing information for the AD[31:16] and C/BE[3:2]# signals. The agent that is providing the data will drive this signal.</p>
SB_STB	I AGP	<p><b>SBA Bus Strobe:</b></p> <p><b>1X Operation:</b> This signal is not used during 1X operation.</p> <p><b>2X Operation:</b> During 2X operation, this signal provides timing for the SBA bus signals. The agent that is driving the SBA bus will drive this signal.</p> <p><b>4X Operation:</b> During 4X operation, this is one-half of a differential strobe pair that provides timing information for the SBA bus signals. The agent that is driving the SBA bus will drive this signal.</p>
SB_STB#	I AGP	<p><b>SBA Bus Strobe Compliment:</b></p> <p><b>1X Operation:</b> This signal is not used during 1X operation.</p> <p><b>2X Operation:</b> During 2X operation, this signal is not used.</p> <p><b>4X Operation:</b> During 4X operation, this is one-half of a differential strobe pair that provides timing information for the SBA bus signals. The agent that is driving the SBA bus will drive this signal.</p>

## 2.4.5. AGP FRAME# Signals

For transactions on the AGP interface carried using AGP FRAME# protocol these signals operate similar to their semantics in the PCI 2.1 specification (the AGP interface is NOT PCI 2.1 – 66 MHz compliant). The exact role of all AGP FRAME# signals are defined below.

Name	Type	Description <sup>1,2,3</sup>
G_FRAME#	I/O s/t/s AGP	<p><b>FRAME#:</b></p> <p><b>PIPE# and SBA Operation:</b> Not used by AGP SBA and PIPE#, but used during AGP FRAME# .</p> <p><b>Fast Write Operation:</b> G_FRAME# is used to frame transactions as an output from the MCH during Fast Writes.</p> <p><b>FRAME# Operation:</b> G_FRAME# is an output when the MCH acts as an initiator on the AGP Interface. G_FRAME# is asserted by the MCH to indicate the beginning and duration of an access. G_FRAME# is an input when the MCH acts as a FRAME# based AGP target. As a FRAME# based AGP target, the MCH latches the C/BE[3:0]# and the AD[31:0] signals on the first clock edge on which it samples FRAME# active.</p>
G_IRDY#	I/O s/t/s AGP	<p><b>Initiator Ready:</b></p> <p><b>PIPE# and SBA Operation:</b> Not used while enqueueing requests via AGP SBA and PIPE#, but used during the data phase of PIPE# and SBA transactions.</p> <p><b>FRAME# Operation:</b> G_IRDY# is an output when MCH acts as a FRAME# based AGP initiator and an input when the MCH acts as a FRAME# based AGP target. The assertion of G_IRDY# indicates the current FRAME# based AGP bus initiator's ability to complete the current data phase of the transaction.</p> <p><b>Fast Write Operation:</b> G_IRDY# indicates the AGP compliant master is ready to provide <i>all</i> write data for the current transaction. Once G_IRDY# is asserted for a write operation, the master is not allowed to insert wait states. The master is <i>never</i> allowed to insert a wait state during the initial data transfer (32 bytes) of a write transaction. However, it may insert wait states after each 32-byte block is transferred.</p>
G_TRDY#	I/O s/t/s AGP	<p><b>Target Ready:</b></p> <p><b>PIPE# and SBA Operation:</b> Not used while enqueueing requests via AGP SBA and PIPE#, but used during the data phase of PIPE# and SBA transactions.</p> <p><b>FRAME# Operation:</b> G_TRDY# is an input when the MCH acts as an AGP initiator and an output when the MCH acts as a FRAME# based AGP target. The assertion of G_TRDY# indicates the target's ability to complete the current data phase of the transaction.</p> <p>During Fast Write Operation: G_TRDY# indicates the AGP compliant target is ready to receive write data for the entire transaction (when the transfer size is less than or equal to 32 bytes) or is ready to transfer the initial or subsequent block (32 bytes) of data when the transfer size is greater than 32 bytes. The target is allowed to insert wait states after each block (32 bytes) is transferred on write transactions.</p>
G_STOP#	I/O s/t/s AGP	<p><b>Stop:</b></p> <p><b>PIPE# and SBA Operation:</b> This signal is not used during PIPE# or SBA operation.</p> <p><b>FRAME# Operation:</b> STOP# is an input when the MCH acts as a FRAME# based AGP initiator and an output when the MCH acts as a FRAME# based AGP target. STOP# is used for disconnect, retry, and abort sequences on the AGP interface.</p>

Name	Type	Description <sup>1,2,3</sup>
G_DEVSEL#	I/O s/t/s AGP	<p><b>Device Select:</b></p> <p><b>PIPE# and SBA Operation:</b> This signal is not used during PIPE# or SBA operation.</p> <p><b>FRAME# Operation:</b> DEVSEL#, when asserted, indicates that a FRAME#-based AGP target device has decoded its address as the target of the current access. The MCH asserts G_DEVSEL# based on the DRAM address range being accessed by a PCI initiator. As an input, G_DEVSEL# indicates whether any device on the bus has been selected.</p> <p><b>Fast Write Operation:</b> G_DEVSEL# is used when the transaction cannot complete during the block data transfer.</p>
G_REQ#	I AGP	<p><b>Request:</b></p> <p><b>SBA Operation:</b> This signal is not used during PIPE# or SBA operation.</p> <p><b>PIPE# and FRAME# Operation:</b> REQ#, when asserted, indicates that a FRAME# or PIPE# based AGP master is requesting use of the AGP interface. This signal is an input into the MCH.</p>
G_GNT#	O AGP	<p><b>Grant:</b></p> <p><b>SBA, PIPE# and FRAME# Operation:</b> GNT# along with the information on the ST[2:0] signals (status bus) indicates how the AGP interface will be used next. Refer to the AGP Interface Specification Revision 2.0 for further explanation of the ST[2:0] values and their meanings.</p>
G_AD[31:0]	I/O AGP	<p><b>Address/Data Bus:</b></p> <p><b>PIPE# and FRAME# Operation:</b> AD[31:0] are used to transfer both address and data information on the AGP interface.</p> <p><b>SBA Operation:</b> AD[31:0] are used to transfer data on the AGP interface.</p>
G_C/BE[3:0]#	I/O AGP	<p><b>Command/Byte Enable:</b></p> <p><b>FRAME# Operation:</b> During the address phase of a transaction, C/BE[3:0]# define the bus command. During the data phase, C/BE[3:0]# are used as byte enables. The byte enables determine which byte lanes carry meaningful data. The commands issued on the C/BE# signals during FRAME# based AGP are the same C/BE# command described in the PCI 2.1 specification.</p> <p><b>PIPE# Operation:</b> When an address is enqueued using PIPE#, the C/BE# signals carry command information. Refer to the AGP 2.0 Interface Specification Revision 2.0 for the definition of these commands. The command encoding used during PIPE# based AGP is <b>DIFFERENT</b> than the command encoding used during FRAME# based AGP cycles (or standard PCI cycles on a PCI bus).</p> <p><b>SBA Operation:</b> These signals are not used during SBA operation.</p>
G_PAR	I/O AGP	<p><b>Parity:</b></p> <p><b>FRAME# Operation:</b> MCH is driven by the MCH when it acts as a FRAME# based AGP initiator during address and data phases for a write cycle, and during the address phase for a read cycle. PAR is driven by the MCH when it acts as a FRAME# based AGP target during each data phase of a FRAME# based AGP memory read cycle. Even parity is generated across AD[31:0] and C/BE[3:0]#.</p> <p><b>SBA and PIPE# Operation:</b> This signal is not used during SBA and PIPE# operation.</p>

**NOTES:**

1. PCIRST# from the I/O Controller Hub is connected to RSTIN# and is used to reset AGP interface logic within the MCH. The AGP agent also uses PCIRST# provided by the I/O Controller Hub as an input to reset its internal logic.
2. LOCK# signal is NOT supported on the AGP interface, even for FRAME#-based AGP operations.
3. The SERR# and PERR# Signals are NOT supported on the AGP interface.

## 2.5. Clock and Reset Signals

Name	Type	Description
CPUCLK	I CMOS	<b>Host Clock In:</b> CPUCLK receives a buffered host clock from the external clock synthesizer. This clock is used by all of the MCH logic that is in the Host clock domain. Note the clock synthesizer drives this signal to 2.5V.
CLK66	I CMOS	<b>AGP Clock In:</b> CLK66 receives a buffered clock from the clock synthesizer that is synchronously derived from the host clock. This clock is used by all of the MCH logic that is in the AGP clock domain. Note the clock synthesizer drives this to 3.3V, and this input pin is 3.3v tolerant.
RCLKOUT	O CMOS	<b>Direct Rambus Clock Out:</b> This pin provides synchronization data to the DRCG. Note that this pin will only be driven to 1.8V
HCLKOUT	O CMOS	<b>Host Clock Out:</b> This pin provides synchronization data to the DRCG. Note that this pin is only driven to 1.8V.
RSTIN#	I CMOS	<b>Reset In:</b> When asserted, RSTIN# asynchronously resets the MCH logic. This signal is connected to the PCIRST# output of the I/O Controller Hub. All AGP/PCI output and bi-directional signals will also tri-state compliant to PCI Rev 2.0 and 2.1 specifications. Note that this input pin is 3.3v tolerant.

## 2.6. Voltage References, Power, Ground, and Test Signals

Name	Description
GTLREF[A:B]	<b>AGTL Reference:</b> Reference voltage input for the Host AGTL interface.
GRCOMP/ TEST	<p><b>AGP RCOMP Output:</b> GRCOMP is used to calibrate AGP I/O buffers. This signal pin must be connected to a PCB trace representative of the AGP bus data signal traces but sufficiently long to present a long shelf before signal reflection occurs. The AGP buffers are calibrated based on the measured shelf voltage. This signal requires an external 40 ohm pull-down resistor.</p> <p><b>TEST: Test Input:</b> Test is used for manufacturing and board level test purposes. This signal is active high.</p>
HLCOMP	<b>Hub interface (Link) Compensation:</b> This signal is used to calibrate the hub interface I/O buffers. This signal pin must be connected to a PCB trace representative of the hub interface data signal traces but sufficiently long to present a long shelf before signal reflection occurs. The hub interface buffers are calibrated based on the measured shelf voltage.
RAMREF[A:B]	<b>Direct Rambus Reference:</b> Reference voltage input for the Direct Rambus RSL interface.
AGPREF	<b>AGP Reference:</b> Reference voltage input for the AGP interface.
HUBREF	<b>Hub interface Reference:</b> Reference voltage input for the hub interface.
VDDQ	<b>AGP I/O Buffer Supply Voltage.</b>
VCC1_8	<b>1.8V VCC:</b> Provides 1.8V to the MCH core, as well as 1.8V I/O buffers.
VSS	<b>Ground.</b>

## 2.7. Strap Signals

Table 2 indicates the strap options invoked by various MCH signal pins.

**Table 2. MCH Strapping Options**

Name	Definition
HL10	<p><b>Host Bus Frequency:</b> The signal is latched on the rising edge of RSTIN#. It indicate what the Processor System Bus frequency is to select the correct internal frequency ratios. This value can be read from the MCHCFG register.</p> <p>0 = 100 MHz 1 = 133 MHz</p>
HA7#	<p><b>Processor Bus In-Order Queue Depth:</b> The value on HA7# is sampled by all host bus agents (including the MCH) on the rising edge of CPURST#. Its latched value determines the maximum IOQ depth mode supported on the host bus. If HA7# is sampled low the IOQ depth on the bus is one. If HA7# is sampled high the IOQ depth on the bus is the maximum of eight. When the IOQ depth on the bus is set to 8, the MCH limits the number of queued transactions to 6 using BNR#.</p> <p>The MCH does not drive HA7# during CPURST#. If an IOQ depth of 1 is desired, HA7# needs to be driven low during CPURST# by external logic.</p>
ST[2:0]	<b>Manufacturing Test Modes:</b> Requires 8.2 K $\Omega$ external pullup resistors.

## 2.8. Pin States During Reset

Table 3 indicate the MCH signal pin states during reset assertion.

Z	Tristate Outputs
ISO	Isolate Inputs in Inactive State
S	Strap sampled on RSTIN# rising edge
H	Driven High
L	Driven Low
D	Drive Outputs to Functional Logic Level
I	Input Active
U	Undefined - Indeterminate

**Table 3. Signal States During Reset**

Signal Name	During RSTIN#
<b>Host Signals</b>	
HA[31:16]#	Z/I
HA[15]#	L/I
HA[14:8]#	Z/I
HA[7]#	Z/S
HA[6:3]#	Z/I
HD[63:0]#	Z/I
ADS#	Z/I
BNR#	Z/I
BPRI#	Z
DBSY#	Z/I
DEFER#	Z/I
DRDY#	Z/I
HIT#	Z/I
HITM#	Z/I
HLOCK#	I
HREQ[4:0]#	Z/I
HTRDY#	Z/I
RS[2:0]#	Z/I
CPURST#	L
<b>Direct Rambus* Signals</b>	
DQ[15:0]	Z/I
DQP[1:0]	Z/I

Signal Name	During RSTIN#
RQ[7:0]	Z/I
CTM	I
CTM#	I
CFM	Z
CFM#	Z
CMD	Z
SCK	L
SIO	Z/I
<b>Hub Interface, AGP, Miscellaneous Signals</b>	
PD[7:0]	Z
PSTRB	L/I
PSTRB#	H/I
REQM	Z
REQUI	ISO
PSTOP	Z/S
PIPE#	I
SBA[7:0]	I
RBF#	I
WBF#	I
ST[2:0]	I
AD_STB0	Z/I
AD_STB0#	Z/I

Signal Name	During RSTIN#
AD_STB1	Z/I
AD_STB1#	Z/I
SB_STB	Z/I
SB_STB#	Z/I
G_FRAME#	Z/I
G_IRDY#	Z/I
G_TRDY#	Z/I
G_STOP#	Z/I
G_DEVSEL#	Z/I
G_REQ#	I
G_GNT#	Z
G_AD[31:0]	L/I
G_C/BE[3:0]#	L/I
G_PAR	L/I
HCLKIN	I
GCLKIN	I
RSTIN#	I
RCLKOUT	D
HCLKOUT	D
GRCOMP/ TEST	Z/I
ZCOMP	Z/I

## 3. Register Description

---

This chapter describes the MCH PCI configuration registers. A detailed register bit description is provided. The MCH contains two sets of software accessible registers, accessed via the Host I/O address space:

- Control registers I/O mapped into the host I/O space, which controls access to PCI and AGP configuration space (see Section 3.3, “I/O Mapped Registers”).
- PCI configuration registers in the MCH are partitioned into two logical device register sets (“logical” since they reside within a single physical device). The Device 0 register set is dedicated to Host-Hub Interface Bridge/DRAM Controller functionality. This register set controls PCI0 (i.e., DRAM configuration, other chipset operating parameters and optional features). The Device 1 register set is dedicated to Host-AGP Bridge functions. This register set controls AGP interface configurations and operating parameters.

The MCH supports PCI configuration space accesses using PCI Configuration Mechanism #1.

The MCH internal registers (both I/O Mapped and Configuration registers) are accessible by the host. The registers can be accessed as Byte, Word (16-bit), or DWord (32-bit) quantities, with the exception of the CONF\_ADDR register which can only be accessed as a DWord. All multi-byte numeric fields use “little-endian” ordering (i.e., lower addresses contain the least significant parts of the field).

### 3.1. Register Nomenclature and Access Attributes

Symbol	Description
RO	Read Only. If a register is read only, writes to this register have no effect.
R/W	Read/Write. A register with this attribute can be read and written
R/W/L	Read/Write/Lock. A register with this attribute can be read, written, and Lock.
R/WC	Read/Write Clear. A register bit with this attribute can be read and written. However, a write of a 1 clears (sets to 0) the corresponding bit and a write of a 0 has no effect.
R/WO	Read/Write Once. A register bit with this attribute can be written to only once after power up. After the first write, the bit becomes read only.
L	Lock. A register bit with this attribute becomes Read Only after a lock bit is set.
Reserved Bits	Some of the MCH registers described in this section contain reserved bits. These bits are labeled “Reserved”. Software must deal correctly with fields that are reserved. On reads, software must use appropriate masks to extract the defined bits and not rely on reserved bits being any particular value. On writes, software must ensure that the values of reserved bit positions are preserved. That is, the values of reserved bit positions must first be read, merged with the new values for other bit positions and then written back. Note the software does not need to perform read, merge, write operation for the configuration address register.

Symbol	Description
Reserved Registers	In addition to reserved bits within a register, the MCH contains address locations in the configuration space of the Host-Hub Interface Bridge/DRAM Controller and Host-AGP Bridge entities that are marked either "Reserved. When a "Reserved" register location is read, a random value can be returned. ("Reserved" registers can be 8-, 16-, or 32-bit in size). Registers that are marked as "Reserved" must not be modified by system software. Writes to "Reserved" registers may cause system failure.
Default Value Upon Reset	Upon a Full Reset, the MCH sets all of its internal configuration registers to predetermined <b>default</b> states. The default state represents the minimum functionality feature set required to successfully bring up the system. Hence, it does not represent the optimal system configuration. It is the responsibility of the system initialization software (usually BIOS) to properly determine the DRAM configurations, operating parameters and optional system features that are applicable, and to program the MCH registers accordingly.

## 3.2. PCI Configuration Space Access

The MCH and the I/O Controller Hub are physically connected via the hub interface. From a configuration standpoint, the interface connecting the MCH and the I/O Controller Hub is **logically PCI bus #0**. All devices internal to the MCH and I/O Controller Hub appear to be on PCI bus #0. The system primary PCI expansion bus is physically attached to the I/O Controller Hub and, from a configuration standpoint, appears as a hierarchical PCI bus behind a PCI-to-PCI bridge. The primary PCI expansion bus connected to the I/O Controller Hub has a programmable PCI Bus number.

**Note:** Even though the primary PCI bus is referred to as PCI0 in this document, it is not PCI bus #0 from a configuration standpoint.

The MCH contains two PCI devices within a single physical component. The configuration registers for both Device 0 and 1 are mapped as devices residing on PCI bus #0.

- **Device 0: Host-Hub Interface Bridge/DRAM Controller.** Logically this appears as a PCI device residing on PCI bus #0. Physically Device 0 contains the standard PCI registers, AGP capabilities registers, DRAM registers, the Graphics Aperture controller, and other MCH specific registers.
- **Device 1: Host-AGP Bridge.** Logically this appears as a "virtual" PCI-to-PCI bridge residing on PCI bus #0. Physically Device 1 contains the standard PCI-to-PCI bridge registers and the standard AGP configuration registers (including the AGP I/O and memory address mapping).

**Note:** A physical PCI bus #0 does not exist. The interface and the internal devices in the MCH and I/O Controller Hub logically constitute PCI Bus #0 to configuration software.

### PCI Bus Configuration Mechanism

The PCI Bus defines a slot-based "configuration space" that allows each device to contain up to 8 functions with each function containing up to 256 8-bit configuration registers. The PCI specification defines two bus cycles to access the PCI configuration space: Configuration Read and Configuration Write. Memory and I/O spaces are supported directly by the processor. Configuration space is supported by a mapping mechanism implemented within the MCH. The PCI specification defines two mechanisms to access configuration space, Mechanism #1 and Mechanism #2. The MCH supports only Mechanism #1.



The configuration access mechanism makes use of the CONF\_ADDR Register and CONF\_DATA Register. To reference a configuration register a DWord I/O write cycle is used to place a value into CONF\_ADDR that specifies:

- The PCI bus
- The device on that bus
- The function within the device
- A specific configuration register of the device function being accessed.

CONF\_ADDR[31] must be 1 to enable a configuration cycle. CONF\_DATA then becomes a window into the four bytes of configuration space specified by the contents of CONF\_ADDR. Any read or write to CONF\_DATA results in the MCH translating the CONF\_ADDR into the appropriate configuration cycle.

The MCH is responsible for translating and routing the processor I/O accesses to the CONF\_ADDR and CONF\_DATA registers to internal MCH configuration registers, the hub interface, or AGP.

### **Routing Configuration Accesses to Primary PCI (PCI0), or AGP**

The MCH supports two bus interfaces (the hub interface and AGP). PCI configuration cycles are selectively routed to both interfaces. The MCH is responsible for routing PCI configuration cycles to the proper interface. PCI configuration cycles to the I/O Controller Hub internal devices and PCI0 (including downstream devices) are routed to the I/O Controller Hub via the hub interface. AGP configuration cycles are routed to AGP. The AGP interface is treated as a separate PCI bus from the configuration point of view. Routing of configuration accesses to AGP is controlled via the standard PCI-PCI bridge mechanism using information contained within the Primary Bus Number (PBUSN), the Secondary Bus Number (SBUSN), and the Subordinate Bus Number (SUBUSN) registers of the Host-AGP internal “virtual” PCI-PCI bridge device.

### **Logical PCI Bus #0 Configuration Mechanism**

The MCH decodes the Bus Number (bits 23:16) and the Device Number fields of the CONF\_A DDR register. If the Bus Number field of CONF\_ADDR is 0, the configuration cycle is targeting a PCI Bus #0 device.

- The Host-Hub Interface Bridge/DRAM Controller entity within the MCH is hardwired as Device 0 on PCI Bus #0.
- The Host-AGP Bridge entity within the MCH is hardwired as Device 1 on PCI Bus #0.

Configuration cycles to one of the MCH internal devices are confined to the MCH and not sent over the hub interface. Accesses to devices #2 to #31 are forwarded over the hub interface.

### **Primary PCI (PCI0) and Downstream Configuration Mechanism**

If the Bus Number in the CONF\_ADDR register is non-zero, and is less than the value programmed into the MCH Device 1 SBUSN register or greater than the value programmed into the SUBUSN Register, the MCH will generate a configuration cycle over the hub interface. The I/O Controller Hub compares the non-zero Bus Number with the SBUSN and SUBUSN registers of its P2P bridges to determine if the configuration cycle is meant for Primary PCI (PCI0), or a downstream PCI bus.

## AGP Bus Configuration Mechanism

From the chipset configuration perspective, AGP is another PCI bus interface residing on a Secondary Bus side of the “virtual” PCI-PCI bridge referred to as the MCH Host-AGP bridge. On the Primary bus side, the “virtual” PCI-PCI bridge is attached to PCI Bus #0. Therefore, the PBUSN register is hardwired to 0. The “virtual” PCI-PCI bridge entity converts Type #1 PCI Bus Configuration cycles on PCI Bus #0 into Type 0 or Type 1 configuration cycles on the AGP interface. Type 1 configuration cycles on PCI Bus #0 that have a Bus Number that matches the Secondary Bus Number of the MCH Host-AGP bridge will be translated into Type 0 configuration cycles on the AGP interface.

If the Bus Number is non-zero, greater than the value programmed into the SBUSN register, and less than or equal to the value programmed into the SUBUSN register, the MCH will generate a Type 1 PCI configuration cycle on AGP.

## 3.3. I/O Mapped Registers

The MCH contains a set of registers that reside in the host I/O address space – the Configuration Address (CONF\_ADDR) Register and the Configuration Data (CONF\_DATA) Register. The Configuration Address Register enables/disables the configuration space and determines what portion of configuration space is visible through the Configuration Data window.

### 3.3.1. CONF\_ADDR—Configuration Address Register

I/O Address: 0CF8h Accessed as a DWord  
 Default: 00000000h  
 Access: Read/Write  
 Size: 32 bits

CONF\_ADDR is a 32 bit register accessed only when referenced as a Dword. A Byte or Word reference will “pass through” the Configuration Address Register and the hub interface onto the PCI0 bus as an I/O cycle. The CONF\_ADDR register contains the Bus Number, Device Number, Function Number, and Register Number for which a subsequent configuration access is intended.

Bit	Descriptions
31	<b>Configuration Enable (CFGE).</b> 1 = Enable access to PCI configuration space. 0 = Disable access to PCI configuration space.
30:24	Reserved (These bits are read only and have a value of 0).

Bit	Descriptions
23:16	<p><b>Bus Number.</b> When the Bus Number is programmed to 00h the target of the Configuration Cycle is either a hub interface agent MCH or the I/O Controller Hub.</p> <p>The Configuration Cycle is forwarded to the hub interface if the Bus Number is programmed to 00h and the MCH is not the target.</p> <p>If the Bus Number is non-zero and matches the value programmed into the SBUSN Register, a Type 0 PCI configuration cycle will be generated on AGP.</p> <p>If the Bus Number is non-zero, greater than the value in the SBUSN register and less than or equal to the value programmed into the SUBUSN Register, a Type 1 PCI configuration cycle will be generated on AGP.</p> <p>If the Bus Number is non-zero, and is less than the value programmed into the SBUSN or is greater than the value programmed into the SUBUSN Register, a configuration cycle over the hub interface is generated.</p>
15:11	<p><b>Device Number.</b> This field selects one agent on the PCI bus selected by the Bus Number. When the Bus Number field is "00" the MCH decodes the Device Number field. The MCH is always Device Number 0 for the Host-Hub Interface bridge/DRAM Controller entity and Device Number 1 for the Host-AGP entity. Therefore, when the Bus Number =0 and the Device Number=0 or 1 the internal MCH devices are selected.</p> <p>If the Bus Number is non-zero and matches the value programmed into the SBUSN Register a Type 0 PCI configuration cycle will be generated on AGP. The Device Number field is decoded and the MCH asserts one and only one GADxx signal as an IDSEL. GAD11 is asserted to access Device 0, GAD12 for Device 1 and so forth up to Device 20 for which will assert AD31. All device numbers higher than 20 cause a type 0 configuration access with no IDSEL asserted, which will result in a Master Abort reported in the MCH's "virtual" PCI-PCI bridge registers.</p> <p>For Bus Numbers resulting in the hub interface configuration cycles the MCH propagates the Device Number field as A[15:11]. For Bus Numbers resulting in AGP Type 1 Configuration cycles the Device Number is propagated as GAD[15:11].</p>
10:8	<p><b>Function Number.</b> This field is mapped to GAD[10:8] during AGP Configuration cycles and A[10:8] during the hub interface configuration cycles. This allows the configuration registers of a particular function in a multi-function device to be accessed. The MCH ignores configuration cycles to its two internal Devices if the function number is not equal to 0.</p>
7:2	<p><b>Register Number.</b> This field selects one register within a particular Bus, Device, and Function as specified by the other fields in the Configuration Address Register. This field is mapped to GAD[7:2] during AGP Configuration cycles and A[7:2] during the hub interface Configuration cycles.</p>
1:0	Reserved.

### 3.3.2. CONF\_DATA—Configuration Data Register

I/O Address:	0CFCh
Default:	00000000h
Access:	Read/Write
Size:	32 bits

CONF\_DATA is a 32 bit read/write window into configuration space. The portion of configuration space that is referenced by CONF\_DATA is determined by the contents of the CONF\_ADDR register.

Bit	Descriptions
31:0	<p><b>Configuration Data Window (CDW).</b> If bit 31 of CONF_ADDR is 1 any I/O access that to the CONF_DATA register will be mapped to configuration space using the contents of CONF_ADDR.</p>

### 3.4. Host-Hub Interface Bridge/DRAM Controller Device Registers (Device 0)

Table 4 shows the MCH configuration space for Device 0. An “s” in the Default Value field means that the power-up default value for that bit is determined by a strap.

**Table 4: MCH PCI Configuration Space (Device 0)**

Addr. Offset	Symbol	Register Name	Default Value	Access
00–01h	VID	Vendor Identification	8086h	RO
02–03h	DID	Device Identification	2500h	RO
04–05h	PCICMD	PCI Command	0006h	R/W
06–07h	PCISTS	PCI Status	0010h	RO, R/WC
08	RID	Revision Identification	03h	RO
09	—	Reserved	00h	—
0Ah	SUBC	Sub-Class Code	00h	RO
0Bh	BCC	Base Class Code	06h	RO
0Ch	—	Reserved	00h	—
0Dh	MLT	Master Latency Timer	00h	RO
0Eh	HDR	Header Type	00h	RO
10–13h	APBASE	Aperture Base Configuration	0000_0008h	R/W, RO
14–2Bh	—	Reserved	00h	—
2C–2Dh	SVID	Subsystem Vendor Identification	00h	R/WO
2E–2Fh	SID	Subsystem Identification	00h	R/WO
30–33h	—	Reserved	00h	—
34h	CAPPTR	Capabilities Pointer	A0h	RO
35–3Fh	—	Reserved	00h	—
40–47h	GAR0–GAR7	RDRAM Group Architecture (8 registers)	80h	R/W/L
48–4Fh	—	Reserved	00h	—
50h	RDTR	RDRAM Timing	00h	R/W
51h	RDCR	RDRAM Control	00h	R/W
52h	RDRR	RDRAM Refresh	00h	R/W
53h	RPMR	RDRAM Power Management	00h	R/W
54–57h	—	Reserved	—	—
58h	FDHC	Fixed DRAM Hole Control	00h	R/W
59–5Fh	PAM0–PAM6	Programmable Attribute Map (7 registers)	00h	R/W
60–6Fh	GBA0–GBA7	RDRAM Group Boundary Address (8 registers)	(See register description)	R/W/L
70–7Fh	—	Reserved	00h	—
80–87h	DTC	DRAM Throttle Control	0000_0000 0000_0000h	R/W/L
88–8Fh	—	Reserved	00h	—

Addr. Offset	Symbol	Register Name	Default Value	Access
90–91h	DRD	RDRAM Device Register Data	0000h	R/W
92–93h	—	Reserved	00h	—
94–96h	RICM	RDRAM Initialization Control/	000000h	R/W
97–9Ch	—	Reserved	00h	—
9Dh	SMRAM	System Management RAM Control	02h	R/W/L, RO
9Eh	ESMRAMC	Extended System Management RAM Control	38h	R/W/L
9Fh	—	Reserved	00h	—
A0–A3h	ACAPID	AGP Capability Identifier	0020_0002h	RO
A4–A7h	AGPSTAT	AGP Status	1F000207h	RO
A8–ABh	AGPCMD	AGP Command	00000000h	RW
AC–AFh	—	Reserved	00h	—
B0–B3h	AGPCTRL	AGP Control	00000000h	R/W
B4h	APSIZE	Aperture Size	00h	R/W
B5–B7h	—	Reserved	00h	—
B8–BBh	ATTBASE	Aperture Translation Table	00000000h	R/W
BCh	AMTT	AGP MTT Control	00h	R/W
BDh	LPTT	AGP Low Priority Transaction Timer	00h	R/W
BE–BFh	MCHCFG	MCH Configuration	00s0_0000_0000_0s00b	R/W, RO
C0–C3h	—	Reserved	0000_0000h	—
C4–C7h	EAP	Error Address Pointer	XXXX_XXXXh	RO
C8–C9h	ERRSTS	Error Status	0000h	R/WC
CA–CBh	ERRCMD	Error Command	0000h	R/W
CCh	SMICMD	SMI Error Command	00h	R/W
CDh	SCICMD	SCI Error Command	00h	R/W
CE–DDh	—	Reserved	00h	—
DE–DFh	SKPD	Scratchpad Data (not reset)	xxxxh	R/W
E0–E7h	—	Reserved	00h	—
E8–EBh	AGPBCTRL	AGP Buffer Control B	0000 0000h	R/W
EC–F6h	—	Reserved	—	—
F7h	RTCE	RDRAM Temperature Calibration Enable	00h	R/W
F8–FDh	—	Reserved	—	—
FEh	AGPAPPEND	AGP Append Disable	00h	R/W
FFh	—	Reserved	00h	—

### 3.4.1. VID—Vendor Identification Register (Device 0)

Address Offset:	00–01h
Default:	8086h
Access:	Read Only
Size:	16 bits

The VID Register contains the vendor identification number. This 16-bit register combined with the Device Identification Register uniquely identify any PCI device. Writes to this register have no effect.

Bit	Description
15:0	<b>Vendor Identification Number.</b> This is a 16-bit value assigned to Intel. Intel VID = 8086h.

### 3.4.2. DID—Device Identification Register (Device 0)

Address Offset:	02–03h
Default:	2500h
Access:	Read Only
Size:	16 bits

This 16-bit register combined with the Vendor Identification register uniquely identifies any PCI device. Writes to this register have no effect.

Bit	Description
15:0	<b>Device Identification Number.</b> This is a 16 bit value assigned to the MCH Host-Hub interface Bridge/DRAM Controller Function #0.

### 3.4.3. PCICMD—PCI Command Register (Device 0)

Address Offset: 04–05h  
 Default: 0006h  
 Access: Read/Write  
 Size: 16 bits

Since MCH Device 0 does not physically reside on PCI0, many of the bits are not implemented.

Bit	Descriptions
15:10	Reserved.
9	<b>Fast Back-to-Back. (Not Implemented).</b> This bit controls whether or not the master can do fast back-to-back writes. Since Device 0 is strictly a target, this bit is not implemented and is hardwired to 0. Writes to this bit position have no affect.
8	<b>SERR Enable (SERRE).</b> This bit is a global enable bit for Device 0 SERR messaging. The MCH does not have an SERR# signal. The MCH communicates the SERR# condition by sending a SERR message to the I/O Controller Hub. Note that this bit only controls SERR message for Device 0. Device 1 has its own SERRE bit to control error reporting for error conditions occurring on Device 0.  1 = Enable. The MCH is enabled to generate SERR messages over the hub interface for specific Device 0 error conditions that are individually enabled in the ERRCMD register. The error status is reported in the ERRSTS and PCISTS registers.  0 = SERR message is not generated by the MCH for Device 0.
7	<b>Address/Data Stepping. (Not Implemented).</b> Hardwired to 0. Writes to this bit position have no affect.
6	<b>Parity Error Enable (PERRE). (Not Implemented).</b> Hardwired to 0. Writes to this bit position have no affect.
5	<b>VGA Palette Snoop. (Not Implemented).</b> Hardwired to a 0. Writes to this bit position have no affect..
4	<b>Memory Write and Invalidate Enable. (Not Implemented).</b> Hardwired to 0. Writes to this bit position have no affect.
3	<b>Special Cycle Enable. (Not Implemented).</b> Hardwired to a 0. Writes to this bit position have no affect.
2	<b>Bus Master Enable (BME).</b> The MCH is always enabled as a master on hub interface. This bit is hardwired to a 1. Writes to this bit position have no affect.
1	<b>Memory Access Enable (MAE). (Not Implemented).</b> Hardwired to 1. The MCH always allows access to main memory. Writes to this bit position have no affect.
0	<b>I/O Access Enable (IOAE). (Not Implemented).</b> Hardwired to a 0. Writes to this bit position have no affect.

### 3.4.4. PCISTS—PCI Status Register (Device 0)

Address Offset:	06–07h
Default:	0010h
Access:	Read Only, Read/Write Clear
Size:	16 bits

PCISTS is a 16-bit status register that reports the occurrence of error events on Device 0 on the hub interface. Bit 14 is read/write clear. All other bits are Read Only. Since MCH Device 0 is the Host-to-Hub interface bridge, many of the bits are not implemented.

Bit	Descriptions
15	<b>Detected Parity Error (DPE).</b> Hardwired to a 0. Writes to this bit position have no affect.
14	<b>Signaled System Error (SSE).</b> 1 = MCH Device 0 generates a SERR message over the hub interface for any enabled Device 0 error condition. Device 0 error conditions are enabled in the PCICMD and ERRCMD registers. Device 0 error flags are read/reset from the PCISTS or ERRSTS registers. 0 = Software sets SSE to 0 by writing a 1 to this bit.
13	<b>Received Master Abort Status (RMAS).</b> 1 = MCH generates a request over the hub interface and receives a Master Abort completion packet. 0 = Software clears this bit by writing a 1 to it.
12	<b>Received Target Abort Status (RTAS).</b> 1 = MCH generates a request over the hub interface and receives a Target Abort completion packet. 0 = Software clears this bit by writing a 1 to it.
11	<b>Signaled Target Abort Status (STAS). (Not Implemented).</b> Hardwired to 0. The MCH will not generate a Target Abort completion packet over the hub interface. Writes to this bit position have no affect.
10:9	<b>DEVSEL# Timing (DEVT).</b> Hardwired to 00. Hub Interface does not comprehend DEVSEL# protocol. Writes to these bit positions have no affect.
8	<b>Data Parity Detected (DPD).</b> Hardwired to 0. MCH does not support parity on the hub interface. Writes to this bit position have no affect.
7	<b>Fast Back-to-Back (FB2B).</b> Hardwired to 0. The hub interface does not comprehend PCI Fast Back-to-Back protocol. Writes to this bit position have no affect.
6:5	Reserved.
4	<b>Capability List (CLIST).</b> Hardwired to 1. Indicates to the configuration software that this device/function implements a list of new capabilities. A list of new capabilities is accessed via register CAPPTR at configuration address offset 34h. Register CAPPTR contains an offset pointing to the start address within configuration space of this device where the AGP Capability standard register resides.
3:0	Reserved.



### 3.4.5. RID—Revision Identification Register (Device 0)

Address Offset:	08h
Default:	03h
Access:	Read Only
Size:	8 bits

This register contains the revision number of the MCH Device 0. These bits are read only and writes to this register have no effect.

Bit	Description						
7:0	<p><b>Revision Identification Number.</b> This is an 8-bit value that indicates the revision identification number for the MCH Device 0.</p> <table> <thead> <tr> <th>Stepping</th> <th>Value</th> </tr> </thead> <tbody> <tr> <td>B-1</td> <td>03h</td> </tr> <tr> <td>B-2</td> <td>04h</td> </tr> </tbody> </table>	Stepping	Value	B-1	03h	B-2	04h
Stepping	Value						
B-1	03h						
B-2	04h						

### 3.4.6. SUBC—Sub-Class Code Register (Device 0)

Address Offset:	0Ah
Default:	00h
Access:	Read Only
Size:	8 bits

This register contains the Sub-Class Code for the MCH Device 0. This code is 00h indicating a Host Bridge device.

Bit	Description
7:0	<p><b>Sub-Class Code (SUBC).</b> This 8-bit value indicates the category of bridge.</p> <p>00h = Host Bridge.</p>

### 3.4.7. BCC—Base Class Code Register (Device 0)

Address Offset:	0Bh
Default:	06h
Access:	Read Only
Size:	8 bits

This register contains the Base Class Code of the MCH Device 0. This code is 06h indicating a Bridge device. This register is read only.

Bit	Description
7:0	<p><b>Base Class Code (BASEC).</b> This is an 8-bit value that indicates the Base Class Code for the MCH.</p> <p>06h = Bridge device.</p>

### 3.4.8. MLT—Master Latency Timer Register (Device 0)

Address Offset:	0Dh
Default:	00h
Access:	Read Only
Size:	8 bits

The hub interface does not comprehend the concept of Master Latency Timer. Therefore, this register is not implemented.

Bit	Description
7:0	These bits are hardwired to 0. Writes have no effect.

### 3.4.9. HDR—Header Type Register (Device 0)

Offset:	0Eh
Default:	00h
Access:	Read Only
Size:	8 bits

This register identifies the header layout of the configuration space. No physical register exists at this location.

Bit	Descriptions
7:0	This read only field always returns 0 when read, and writes have no affect.

### 3.4.10. APBASE—Aperture Base Configuration Register (Device 0)

Offset:	10–13h
Default:	0000_0008h
Access:	Read/Write, Read Only
Size:	32 bits

The APBASE is a standard PCI Base Address register that is used to set the base of the Graphics Aperture. The standard PCI Configuration mechanism defines the base address configuration register such that only a fixed amount of space can be requested (dependent on which bits are hardwired to 0 or behave as hardwired to 0). To allow for flexibility (of the aperture), an additional register called APSIZE is used as a “back-end” register to control which bits of the APBASE behave as hardwired to 0. This register is programmed by MCH-specific BIOS code that runs before any of the generic configuration software is run.

**Note:** Bit 9 of the MCHCFG register is used to prevent accesses to the aperture range before this register is initialized by the configuration software and the appropriate translation table structure has been established in the main memory.

Bit	Description																																																								
31:28	<b>Upper Programmable Base Address bits—R/W.</b> These bits are used to locate the range size selected via lower bits 27:4. Default = 0000																																																								
27:22	<p><b>Lower “Hardwired”/Programmable Base Address Bits.</b> These bits behave as a “hardwired” or as a programmable depending on the contents of the APSIZE register as defined below:</p> <table border="1" data-bbox="472 478 998 800"> <thead> <tr> <th>27</th> <th>26</th> <th>25</th> <th>24</th> <th>23</th> <th>22</th> <th>Aperture Size</th> </tr> </thead> <tbody> <tr> <td>r/w</td> <td>r/w</td> <td>r/w</td> <td>r/w</td> <td>r/w</td> <td>r/w</td> <td>4 MB</td> </tr> <tr> <td>r/w</td> <td>r/w</td> <td>r/w</td> <td>r/w</td> <td>r/w</td> <td>0</td> <td>8 MB</td> </tr> <tr> <td>r/w</td> <td>r/w</td> <td>r/w</td> <td>r/w</td> <td>0</td> <td>0</td> <td>16 MB</td> </tr> <tr> <td>r/w</td> <td>r/w</td> <td>r/w</td> <td>0</td> <td>0</td> <td>0</td> <td>32 MB</td> </tr> <tr> <td>r/w</td> <td>r/w</td> <td>0</td> <td>0</td> <td>0</td> <td>0</td> <td>64 MB</td> </tr> <tr> <td>r/w</td> <td>0</td> <td>0</td> <td>0</td> <td>0</td> <td>0</td> <td>128 MB</td> </tr> <tr> <td>0</td> <td>0</td> <td>0</td> <td>0</td> <td>0</td> <td>0</td> <td>256 MB</td> </tr> </tbody> </table> <p>Bits 27:22 are controlled by bits 5:0 of the APSIZE register in the following manner:                      If bit APSIZE[5]=0 then APBASE[27]=0 and if APSIZE[5]=1 then APBASE[27]=r/w (read/write). The same applies correspondingly to other bits.</p> <p>Default for APSIZE[5:0]=000000b forces default APBASE[27:22] =000000b (i.e., all bits respond as “hardwired” to 0). This provides a default to the maximum aperture size of 256 MB. The MCH specific BIOS is responsible for selecting smaller size (if required) before PCI configuration software runs and establishes the system address map.</p>	27	26	25	24	23	22	Aperture Size	r/w	r/w	r/w	r/w	r/w	r/w	4 MB	r/w	r/w	r/w	r/w	r/w	0	8 MB	r/w	r/w	r/w	r/w	0	0	16 MB	r/w	r/w	r/w	0	0	0	32 MB	r/w	r/w	0	0	0	0	64 MB	r/w	0	0	0	0	0	128 MB	0	0	0	0	0	0	256 MB
27	26	25	24	23	22	Aperture Size																																																			
r/w	r/w	r/w	r/w	r/w	r/w	4 MB																																																			
r/w	r/w	r/w	r/w	r/w	0	8 MB																																																			
r/w	r/w	r/w	r/w	0	0	16 MB																																																			
r/w	r/w	r/w	0	0	0	32 MB																																																			
r/w	r/w	0	0	0	0	64 MB																																																			
r/w	0	0	0	0	0	128 MB																																																			
0	0	0	0	0	0	256 MB																																																			
21:4	<b>Aperture Size.</b> Hardwired to 0. This forces minimum aperture size selected by this register to be 4 MB.																																																								
3	<b>Prefetchable—RO.</b> This bit is hardwired to 1 to identify the Graphics Aperture range as a prefetchable (i.e., there are no side effects on reads; the device returns all bytes on reads regardless of the byte enables, and the MCH may merge processor writes into this range without causing errors).																																																								
2:1	<b>Type—RO.</b> These bits determine addressing type and they are hardwired to “00” to indicate that address range defined by the upper bits of this register can be located anywhere in the 32-bit address space.																																																								
0	<b>Memory Space Indicator—RO.</b> Hardwired to 0 to identify aperture range as a memory range.																																																								

### 3.4.11. SVID—Subsystem Vendor ID Register (Device 0)

Offset: 2C–2Dh  
 Default: 0000h  
 Access: Read/Write Once  
 Size: 16 bits

This value is used to identify the vendor of the subsystem.

Bit	Description
15:0	<b>Subsystem Vendor ID—R/WO.</b> This field should be programmed during boot-up. After this field is written once, it becomes read only.

### 3.4.12. SID—Subsystem ID Register (Device 0)

Offset: 2E–2Fh  
 Default: 0000h  
 Access: Read/Write Once  
 Size: 16 bits

This value is used to identify a particular subsystem.

Bit	Description
15:0	<b>Subsystem ID—R/WO.</b> This field should be programmed during boot-up. After this field is written once, it becomes read only.

### 3.4.13. CAPPTR—Capabilities Pointer Register (Device 0)

Offset: 34h  
 Default: A0h  
 Access: Read Only  
 Size: 8 bits

The CAPPTR provides the offset that is the pointer to the location where the AGP standard registers are located.

Bit	Description
7:0	<b>Pointer to the Start of AGP Standard Register Block.</b> The value in this field is A0h.

### 3.4.14. GAR0–GAR7—RDRAM Group Architecture Register (Device 0)

Address Offset:	40–47h (GAR0–GAR7)
Default:	80h
Access:	Read/Write/Lock
Size:	8 bits/register

This 8-bit register defines the #of banks and DRAM technology of each device group in the RDRAM channel. There are 8 GAR registers (GAR0–GAR7), which are used to define 8 groups for the Direct RDRAM channel. Also this register can be locked and becomes Read Only when the D\_LCK bit in the SMRAM register is set. This is done to improve the SMM security.

Bit	Description
7:5	Reserved.
4	<b>Device Banks (DB).</b> This field defines #of banks in each device in the group. 0 = 16 dependent banks 1 = 32 dependent banks arranged in two groups of 16 dependent banks (i.e., 2x16)
3	Reserved.
2:1	<b>Device DRAM Technology (DDT).</b> This field defines the DRAM technology of each device in the group. 00 = 64Mbit/72Mbit 01 = 128Mbit/144Mbit 10 = 256Mbit/288Mbit 11 = Reserved
0	Reserved.

### 3.4.15. RDTR—RDRAM Timing Register (Device 0)

Address Offset:	50h
Default:	00h
Access:	Read/Write
Size:	8 bits

This 32-bit register defines the timing parameters for all devices in all channels. The BIOS programs this register with the “least common denominator” values after reading configuration registers of each device in each channel. This register applies to the entire DRAM array.

Bit	Description
7:6	<p><b>Row to Column Delay (tRCD).</b> This field defines the minimum interval between opening a row and column operation on that row in units of RDRAM clocks.</p> <p>00 = Reserved</p> <p>01 = 7 RDRAM clocks</p> <p>10 = 9 RDRAM clocks</p> <p>11 = Reserved</p>
5:4	Reserved
3	<p><b>CAS Access Delay (tCAC).</b> Minimum delay from Read command to Read data in RDRAM clocks</p> <p>0 = Reserved</p> <p>1 = 8 RDRAM clocks</p>
2:0	<p><b>Roundtrip Channel Delay (tRDLY).</b> This field defines the minimum round trip propagation time of the Direct RDRAM channel in units of Direct RDRAM clocks. This value is added to tCAC value to obtain the total Read-to-Data delay.</p> <p>000 = 0</p> <p>001 = 1</p> <p>010 = 2</p> <p>011 = 3</p> <p>100 = 4</p> <p>All other combinations are Reserved</p>

The following table shows the valid tRCD and tCAC combinations for 400 MHz, 356 MHz, 300 MHz, and 266 MHz.

RDRAM Frequency (Rclk)	tRCD in Rclks	tCAC in Rclks
400 MHz	9	8
400 MHz	7	8
356 MHz	9	8
356 MHz	7	8
300 MHz	7	8
266 MHz	7	8

### 3.4.16. RDCR—RDRAM Control Register (Device 0)

Address Offset: 51h  
 Default: 00h  
 Access: Read/Write  
 Size: 8 bits

Bit	Description
7	Reserved.
6	<b>Pool “B” Operation Selection (PBS).</b> This bit selects the operating state of the RDRAM devices in Pool “B”. 0 = All devices in Pool “B” in Standby state 1 = All devices in Pool “B” in Nap state
5:2	Reserved.
1	<b>Aperture Access Global Enable.</b> This bit is used to prevent access to the aperture from any port (processor, PCI0, or AGP/PCI1) before the aperture range is established by the configuration software and appropriate translation table in the main DRAM has been initialized. It must be set after system is fully configured for aperture accesses. 1 = Enable. 0 = Disable (default)
0	Reserved

### 3.4.17. RDRR—RDRAM Refresh Register (Device 0)

Address Offset: 52h  
 Default: 00h  
 Access: Read/Write  
 Size: 8 bits

Bit	Description
7:5	<b>RDRAM Refresh Rate (DRR).</b> The DRAM refresh rate is adjusted according to the frequency selected by this field. Note that refresh is also disabled via this field, and that disabling refresh results in the eventual loss of DRAM data. Note that changing DRR value will reset the refresh request timer. This field is programmed by the BIOS after collecting configuration information from all RDRAM devices in all channels and determining the least common denominator value for refresh. 000 = Refresh Disabled 001 = 1.95 us 010 = 3.9 us 011 = 7.8 us 100 – 111 = Reserved
4:3	Reserved
2:0	<b>RMC Idle Timer (RIT):</b> This field determines the number of host clocks that RMC will remain in the idle state before all open pages are closed. 000 = Infinite (pages are not closed for idle conditions) 001 = 0 (Aggressive page closing. A page is closed if no pending request to that page in the pipeline.) 010 = 2 011 = 4 100 = 8 101 = 16 110 = 32 111 = 64

### 3.4.18. RPMR—RDRAM Power Management Register (Device 0)

Address Offset: 53h  
 Default: 00h  
 Access: Read/Write  
 Size: 8 bits

Bit	Description
7:6	<p><b>Device Napdown Timer (DNT).</b> This field specifies the number of host clocks the RDRAM channel be at idle before the LRU (Least Recent Used) device in Pool A is pushed out to Pool B.</p> <p>00 = 1K (# of clocks the channel must be at idle)            01 = 2K (# of clocks the channel must be at idle)            10 = Reserved            11 = Reserved</p>
5:4	<p><b>Maximum Active Devices in Pool A (ADPA).</b> This field defines the maximum number of RDRAM devices in Pool A that can be in Active Read/Write or Active state at a time. The devices in Pool A that are not in Active Read/write or Active state are in standby state.</p> <p>00 = 1 device            01 = 2 devices            10 = 3 devices            11 = 4 devices</p>
3	<p><b>Device Napdown Enable (DNE).</b></p> <p>1 = Enable. The RDRAM channel inactivity counter is enabled to start counting the continuous inactivity time. When the counter value exceeds the threshold specified by the DNT, the LRU device from Pool A is pushed to Pool B.</p> <p>0 = Disable</p>
2	Reserved
1:0	<p><b>Pool A Capacity (PAC).</b> This field defines the maximum number of RDRAM devices that can reside in Pool A at a time. The devices that are not part of Pool A belong to Pool B.</p> <p>00 = 1 device            01 = 2 devices            10 = 4 devices            11 = 8 devices</p>



### 3.4.19. FDHC—Fixed DRAM Hole Control Register (Device 0)

Address Offset:	58h
Default:	00h
Access:	Read/Write
Size:	8 bits

This 8-bit register controls the fixed DRAM hole: 15-16 MB.

Bit	Description
7	<p><b>Hole Enable (HEN).</b> This field enables a memory hole in DRAM space. Host cycles matching an enabled hole are passed on to I/O Controller Hub through the hub interface. The hub interface cycles matching an enabled hole will be ignored by the MCH. Note that a selected hole is not re-mapped.</p> <p>0 = Disable. No hole</p> <p>1 = 15 MB–16 MB (1 MB)</p>
6:0	Reserved

### 3.4.20. PAM0–PAM6—Programmable Attribute Map Registers (Device 0)

Address Offset:	59–5Fh (PAM0–PAM6)
Default:	00h
Access:	Read/Write
Size:	4 bits/register

The MCH allows programmable memory attributes on 13 *Legacy* memory segments of various sizes in the 640 KBs to 1 MB address range. Seven Programmable Attribute Map (PAM) Registers are used to support these features. Cacheability of these areas is controlled via the MTRR registers in the processor. Two bits are used to specify memory attributes for each memory segment. These bits apply to host accesses, AGP, and the hub interface initiator accesses to the PAM areas. These attributes are:

**RE - Read Enable.** When RE = 1, the host read accesses to the corresponding memory segment are claimed by the MCH and directed to main memory. Conversely, when RE = 0, the host read accesses are directed to PCI0.

**WE - Write Enable.** When WE = 1, the host write accesses to the corresponding memory segment are claimed by the MCH and directed to main memory. Conversely, when WE = 0, the host write accesses are directed to PCI0.


The RE and WE attributes permit a memory segment to be Read Only, Write Only, Read/Write, or disabled. For example, if a memory segment has RE = 1 and WE = 0, the segment is Read Only.

Each PAM Register controls two regions, typically 16 KBs. Each of these regions has a 4-bit field. The four bits that control each region have the same encoding and are defined in the following table.

Table 5. Attribute Bit Assignment

Bits [7, 3] Reserved	Bits [6, 2] Reserved	Bits [5, 1] WE	Bits [4, 0] RE	Description
x	X	0	0	<b>Disabled.</b> DRAM is disabled and all accesses are directed to the hub interface. The MCH does not respond as a PCI target for any read or write access to this area.
X	X	0	1	<b>Read Only.</b> Reads are forwarded to DRAM and writes are forwarded to the hub interface for termination. This write-protects the corresponding memory segment. The MCH responds as an AGP or the hub interface target for read accesses but not for any write accesses.
X	X	1	0	<b>Write Only.</b> Writes are forwarded to DRAM and reads are forwarded to the hub interface for termination. The MCH will respond as an AGP or hub interface target for write accesses but not for any read accesses.
X	X	1	1	<b>Read/Write.</b> This is the normal operating mode of main memory. Both read and write cycles from the host are claimed by the MCH and forwarded to DRAM. The MCH responds as an AGP or the hub interface target for both read and write accesses.

As an example, consider a BIOS that is implemented on the expansion bus. During the initialization process, BIOS can be shadowed in main memory to increase the system performance. When BIOS is shadowed in main memory, it should be copied to the same address location. To shadow the BIOS, the attributes for that address range should be set to write only. The BIOS is shadowed by first doing a read of that address. This read is forwarded to the expansion bus. The host then does a write of the same address, which is directed to main memory. After the BIOS is shadowed, the attributes for that memory area are set to read only so that all writes are forwarded to the expansion bus. Figure 2 and Table 6 show the PAM registers and the associated attribute bits.

**Figure 2. PAM Registers**

**Table 6. PAM Registers and Associated Memory Segments**

PAM Reg	Attribute Bits				Memory Segment	Comments	Offset
PAM0[3:0]	Reserved						
PAM0[7:4]	R	R	WE	RE	0F0000h–0FFFFFFh	BIOS Area	59h
PAM1[3:0]	R	R	WE	RE	0C0000h–0C3FFFh	ISA Add-on BIOS	5Ah
PAM1[7:4]	R	R	WE	RE	0C4000h–0C7FFFh	ISA Add-on BIOS	5Ah
PAM2[3:0]	R	R	WE	RE	0C8000h–0CBFFFh	ISA Add-on BIOS	5Bh
PAM2[7:4]	R	R	WE	RE	0CC000h–0CFFFFh	ISA Add-on BIOS	5Bh
PAM3[3:0]	R	R	WE	RE	0D0000h–0D3FFFh	ISA Add-on BIOS	5Ch
PAM3[7:4]	R	R	WE	RE	0D4000h–0D7FFFh	ISA Add-on BIOS	5Ch
PAM4[3:0]	R	R	WE	RE	0D8000h–0DBFFFh	ISA Add-on BIOS	5Dh
PAM4[7:4]	R	R	WE	RE	0DC000h–0DFFFFh	ISA Add-on BIOS	5Dh
PAM5[3:0]	R	R	WE	RE	0E0000h–0E3FFFh	BIOS Extension	5Eh
PAM5[7:4]	R	R	WE	RE	0E4000h–0E7FFFh	BIOS Extension	5Eh
PAM6[3:0]	R	R	WE	RE	0E8000h–0EBFFFh	BIOS Extension	5Fh
PAM6[7:4]	R	R	WE	RE	0EC000h–0EFFFFh	BIOS Extension	5Fh

**Note:** For details on overall system address mapping scheme see Chapter 4.

### **DOS Application Area (00000h–9FFFh)**

The DOS area is 640 KB in size and it is further divided into two parts. The 512 KB area at 0 to 7FFFFh is always mapped to the main memory controlled by the MCH, while the 128 KB address range from 080000 to 09FFFFh can be mapped to PCI0 or to main DRAM. By default this range is mapped to main memory and can be declared as a main memory hole (accesses forwarded to PCI0) via the FDHC register.

### **Video Buffer Area (A0000h–BFFFFh)**

This 128 KB area is not controlled by attribute bits. The host -initiated cycles in this region are always forwarded to either PCI0 or AGP unless this range is accessed in SMM mode. **Routing of accesses is controlled by the Legacy VGA control mechanism of the “virtual” PCI-PCI bridge device in the MCH.**

This area can be programmed as SMM area via the SMRAM register. When used as a SMM space, this range can not be accessed from the hub interface or AGP.

### **Expansion Area (C0000h–DFFFFh)**

This 128 KB area is divided into eight 16 KB segments which can be assigned with different attributes via PAM control registers.

### **Extended System BIOS Area (E0000h–EFFFFh)**

This 64 KB area is divided into four 16 KB segments which can be assigned with different attributes via PAM control register.

### **System BIOS Area (F0000h–FFFFFh)**

This area is a single 64 KB segment which can be assigned with different attributes via PAM control registers.

### 3.4.21. GBA0–GBA7—RDRAM Group Boundary Address Register (Device 0)

Address Offset:	60–6Fh (GBA0–GBA7)	
Default:	GBA0 (60–61h) = 0001h	GBA4 (60–61h) = 0801h
	GBA1 (62–63h) = 0201h	GBA5 (60–61h) = 0A01h
	GBA2 (64–65h) = 0401h	GBA6 (60–61h) = 0C01h
	GBA3 (66–67h) = 0601h	GBA7 (6E–6Fh) = 0E01h
Access:	Read/Write/Lock	
Size	16 bits/register	

**Note:** This register is locked and becomes Read Only when the D\_LCK bit in the SMRAM register is set. This is done to improve SMM security.

The MCH supports up to 32 RDRAM devices on the Rambus channel. These devices are logically arranged into 8 groups. There are a total of 8 GBA registers. Each group has up to 4 RDRAM devices. Each group requires a separate GBA register. The GBA registers define group ID and the upper and lower addresses for each group in a channel. Bits [0:8] represent the boundary addresses in 8 MB granularity. For example, a value of “01h” indicates 8 MB.

The Group Boundary Address (GBA) Registers are programmed with an 9-bit upper address limit value. Unpopulated groups have a value equal to the previous group (group size = 0). GBA7 reflects the maximum amount of DRAM in the system.

60–61h	GBA0 = Total memory in group0 (in 8 MBs)
62–63h	GBA1 = Total memory in group0 + group1 (in 8 MBs)
64–65h	GBA2 = Total memory in group0 + group1 + group2 (in 8 MBs)
66–67h	GBA3 = Total memory in group0 + group1 + group2 + group3 (in 8 MBs)
....	
6E–6Fh	GBA7 = Total memory in group0 + group1 + group2 + ... + group7 (in 8 MBs)

Bit	Description
15:12	Reserved.
11:9	<p><b>Group ID Field—RO.</b> This field indicates the GBA group (0–7) that a GBA register belongs to. The fields are hardwired in ascending order. The RDRAM Group ID assignments are as follows:</p> <p>000 = GBA0 (60–61h)  001 = GBA1 (62–63h)  010 = GBA2 (64–65h)  011 = GBA3 (66–67h)  100 = GBA4 (68–69h)  101 = GBA5 (6A–6Bh)  110 = GBA6 (6C–6Dh)  111 = GBA7 (6E–6Fh)</p>
8:0	<p><b>Group Boundary Address—R/W/L.</b> This 9-bit value is compared against address lines A[31:23] to determine the upper address limit of a particular group of devices in a channel (i.e., GBA minus previous GBA = group size).</p> <p>This is the Group Boundary Address field. The following restriction is in addition to restrictions that already exist on the GBA registers. The GBA registers must be programmed with the largest RDRAM technology first (starting at GBA0) followed by RDRAM technology that is decreasing in size.</p> <p><b>Example:</b></p> <p>3 - 256 Mbit devices  2 - 128 Mbit devices  6 - 64 Mbit devices</p> <p>GBA0[8:0] = 000001100           (3–256)  GBA1[8:0] = 000010000           (2–128)  GBA2[8:0] = 000010100           (4–64)  GBA3[8:0] = 000010110           (2–64)  GBA4[8:0] = 000010110  GBA5[8:0] = 000010110  GBA6[8:0] = 000010110  GBA7[8:0] = 000010110</p>

### 3.4.22. DTC—DRAM Throttle Control Register (Device 0)

Offset Address: 80–87h  
 Default: 0000\_0000\_0000\_0000h  
 Access: Read/Write/Lock  
 Size: 64 bits

Bits	Description
63	<b>Throttle Lock (TLOCK).</b> This bit secures the DRAM throttling control registers. The bits defaults to '0', then once a '1' is written to either bit, all of the configuration register bits in DTC (including TLOCK) documented below become read-only
62:61	Reserved.
60	<b>Read-Throttle Disable (RDTDIS).</b> 1 = Disable. The MCH ignores memory reads when configured to throttle based on memory bandwidth. This means that MCH will only throttle based on writes. If both this bit and the write-throttle disable below are set, the MCH will NOT throttle based on memory bandwidth. The enable in bits 2:1 must also be set properly to allow bandwidth-based throttling. 0 = Enable (default).
59	<b>Write-Throttle Disable (WRTDIS).</b> 1 = Disable. The MCH ignores memory writes when configured to throttle based on memory bandwidth. This means that MCH will only throttle based on reads. If both this bit and the read-throttle disable above are set, the MCH will NOT throttle based on memory bandwidth. The enable in bits 2:1 must also be set properly to allow bandwidth-based throttling. 0 = Enable (default)
58:46	Reserved.
45:38	<b>Global DRAM Sampling Window (GDSW).</b> This 8-bit value defines the Sampling Windows for the RDRAM counter. This eight-bit value is multiplied by 4 to define the length of time in milliseconds (0–1020) over which the number of RDRAM reads/writes is counted in units of OctWords (16 bytes). If the number of OctWords written/read during this window exceeds the Global DRAM Threshold (GDT) defined below, then the throttling mechanism based on TDM will be invoked to limit DRAM reads/writes to a lower bandwidth checked over smaller time windows.
37:26	<b>Global DRAM Threshold (GDT).</b> This 12-bit value is multiplied by $2^{15}$ to arrive at the number of OctWords that must be written/read within the Global DRAM Sampling Window to cause the Counter-based throttling mechanism to be invoked.
25:20	<b>Throttle Time (TT).</b> This 6-bit value defines the length of time that throttling remains in effect once triggered by the RDRAM counter.  This value provides a multiplier between 0 and 63 which specifies how long RDRAM Counter based throttling remains in effect as a number of Global DRAM Sampling Windows. For example, if GDSW is programmed to 1000_0000b and TT is set to 01_0000b, then throttling will be performed for ~2 seconds once invoked (128 ms * 16).
19:13	<b>Throttle Monitoring Window (TMW).</b> This value defines the replicated time periods within the throttle time (TT) by RDRAM throttling.  The value in this register is padded with 4 0's to specify a window of 0–2047 host clocks with 16 clock granularity. For example, while the throttling mechanism is invoked, DRAM reads/writes are monitored during this window. If the number of OctWords read/written to RDRAM during this window reaches Throttle DRAM Maximum, then read/write from/to DRAM are blocked for the remainder of the window.

Bits	Description
12:3	<b>Throttle DRAM Maximum (TDM).</b> The Throttle DRAM Maximum defines the maximum number of OctWords between 0–1023 which are permitted to be written to and read from DRAM within one Throttle Monitoring Window while the RDRAM throttling mechanism is in effect. RDRAM throttling is invoked with either the START bit or the Counter mechanism.
2:1	<b>RDRAM Throttle Mode (RTMODE).</b> 00 = DRAM Thermal Sensor monitoring enabled without throttling. This setting can be use to generate SERR#, SMI or SCI without automatically starting the throttling. This is the default setting. 01 = DRAM Thermal Sensor monitoring mechanism enabled with throttling. When a RDRAM device's thermal sensor output is detected to be active DRAM throttling begins based on the settings in TT, TMW, and TDM. 10 = DRAM Counter mechanism is enabled. When the threshold set in the GDT field is reached, DRAM throttling begins based on the settings in TT, TMW, and TDM. The DRAM thermal Sensor will still generate SERR#, SMI or SCI with this setting. 11 = Reserved.
0	Reserved

### 3.4.23. DRD—RDRAM Device Register Data Register (Device 0)

Address Offset: 90–91h  
 Default: 0000h  
 Access: Read/Write  
 Size: 16 bits

Bit	Description
15:0	<b>Register Data (RD).</b> This field contains the 16-bit data to be written to a Direct RDRAM register or the data read from a Direct RDRAM register as a result of IOP execution. Data will be valid when the IIO bit of RICHM register has transitioned from 1 to 0



### 3.4.24. RICM—RDRAM Initialization Control Management Register (Device 0)

Address Offset: 94–96h  
 Default: 000000h  
 Access: Read/Write  
 Size: 24 bits

Bit	Description
23	<b>Initiate Initialization Operation (IIO).</b> 1 = Execution of the initialization operation specified by IOP starts. 0 = After the execution is completed, the MCH clears this bit to 0. Software must check to see if this bit is 0 before writing to this bit. The operations that specify register data read from the Direct RDRAM will have the data valid in the DRD register when this is cleared to 0.
22:21	Reserved
20	<b>Initialization Complete (IC).</b> 1 = BIOS sets this bit to 1, after the initialization the RDRAM memory array is complete. This bit is for hardware use.
19	<b>Broadcast Address (BA).</b> 1 = Initialization operation (IOP) command is broadcast to all RDRAM devices in the channel. This field supersedes the SDA field for the appropriate commands.
18	<b>Initialization Opcode (IOP).</b> See description for bits 3:0.
17:9	<b>Device Register Address (DRA).</b> This field specifies the register address for the register read and write operations.
8:4	<b>Serial Device Address (SDA).</b> This 5-bit field specifies the following: <ul style="list-style-type: none"> <li>the <b>serial device ID</b> of the RDRAM device for RDRAM Register Read, RDRAM Register Write, RDRAM Set Reset, RDRAM Clear Reset and RDRAM Set Fast Clock Mode IOP commands.</li> <li>the <b>device ID</b> for Powerdown Entry, Powerdown Exit, Nap Entry, Nap Exit, Current Calibrate and Current Calibrate and Sample IOP commands.</li> <li>the <b>bank address</b> for Refresh and Precharge IOP commands</li> </ul>

Bit	Description																																										
3:0	<p><b>Initialization Opcode (IOP).</b> This field, along with bit 18, specifies the initialization operation to be done on RDRAM device or MCH RAC.</p> <table border="1"> <thead> <tr> <th data-bbox="451 373 573 399">Bits[18, 3:0]</th> <th data-bbox="667 373 865 399">Operation Specified</th> </tr> </thead> <tbody> <tr><td>0, 0000</td><td>RDRAM Register Read</td></tr> <tr><td>0, 0001</td><td>RDRAM Register Write</td></tr> <tr><td>0, 0010</td><td>RDRAM Set Reset</td></tr> <tr><td>0, 0011</td><td>Reserved</td></tr> <tr><td>0, 0100</td><td>RDRAM Set Fast Clock Mode</td></tr> <tr><td>0, 0101 to 0, 1010</td><td>Reserved</td></tr> <tr><td>0, 1011</td><td>RDRAM Clear Reset</td></tr> <tr><td>0, 1100 to 0, 1111</td><td>Reserved</td></tr> <tr><td>1, 0000</td><td>RDRAM Core Initialization (RCI)</td></tr> <tr><td>1, 0001</td><td>RDRAM SIO Reset</td></tr> <tr><td>1, 0010</td><td>RDRAM Powerdown Exit</td></tr> <tr><td>1, 0011</td><td>RDRAM Powerdown Entry</td></tr> <tr><td>1, 0100</td><td>Reserved</td></tr> <tr><td>1, 0101</td><td>Manual Current Calibration of MCH RAC</td></tr> <tr><td>1, 0110</td><td>Load MCH RAC control register with data from DRD register</td></tr> <tr><td>1, 0111</td><td>Initialize MCH RAC</td></tr> <tr><td>1, 1000</td><td>RDRAM Nap Entry</td></tr> <tr><td>1, 1001</td><td>RDRAM Nap Exit</td></tr> <tr><td>1, 1010</td><td>Reserved</td></tr> <tr><td>1, 1011</td><td>RDRAM Precharge</td></tr> </tbody> </table> <p><b>Note:</b> All combinations not shown above are Reserved.</p>	Bits[18, 3:0]	Operation Specified	0, 0000	RDRAM Register Read	0, 0001	RDRAM Register Write	0, 0010	RDRAM Set Reset	0, 0011	Reserved	0, 0100	RDRAM Set Fast Clock Mode	0, 0101 to 0, 1010	Reserved	0, 1011	RDRAM Clear Reset	0, 1100 to 0, 1111	Reserved	1, 0000	RDRAM Core Initialization (RCI)	1, 0001	RDRAM SIO Reset	1, 0010	RDRAM Powerdown Exit	1, 0011	RDRAM Powerdown Entry	1, 0100	Reserved	1, 0101	Manual Current Calibration of MCH RAC	1, 0110	Load MCH RAC control register with data from DRD register	1, 0111	Initialize MCH RAC	1, 1000	RDRAM Nap Entry	1, 1001	RDRAM Nap Exit	1, 1010	Reserved	1, 1011	RDRAM Precharge
Bits[18, 3:0]	Operation Specified																																										
0, 0000	RDRAM Register Read																																										
0, 0001	RDRAM Register Write																																										
0, 0010	RDRAM Set Reset																																										
0, 0011	Reserved																																										
0, 0100	RDRAM Set Fast Clock Mode																																										
0, 0101 to 0, 1010	Reserved																																										
0, 1011	RDRAM Clear Reset																																										
0, 1100 to 0, 1111	Reserved																																										
1, 0000	RDRAM Core Initialization (RCI)																																										
1, 0001	RDRAM SIO Reset																																										
1, 0010	RDRAM Powerdown Exit																																										
1, 0011	RDRAM Powerdown Entry																																										
1, 0100	Reserved																																										
1, 0101	Manual Current Calibration of MCH RAC																																										
1, 0110	Load MCH RAC control register with data from DRD register																																										
1, 0111	Initialize MCH RAC																																										
1, 1000	RDRAM Nap Entry																																										
1, 1001	RDRAM Nap Exit																																										
1, 1010	Reserved																																										
1, 1011	RDRAM Precharge																																										

### 3.4.25. SMRAM—System Management RAM Control Register (Device 0)

Address Offset:	9Dh
Default:	02h
Access:	Read/Write/Lock, Read Only
Size:	8 bits

The SMRAMC register controls how accesses to Compatible and Extended SMRAM spaces are treated. The Open, Close, and Lock bits function only when G\_SMROME bit is set to a 1. Also, the OPEN bit must be reset before the LOCK bit is set.

Bit	Description
7	Reserved
6	<b>SMM Space Open (D_OPEN).</b> When D_OPEN=1 and D_LCK=0, the SMM space DRAM is made visible even when SMM decode is not active. This is intended to help BIOS initialize SMM space. Software should ensure that D_OPEN=1 and D_CLS=1 are not set at the same time. When D_LCK is set to a 1, D_OPEN is reset to 0 and becomes read only.
5	<b>SMM Space Closed (D_CLS).</b> When D_CLS = 1 SMM space DRAM is not accessible to data references, even if SMM decode is active. Code references may still access SMM space DRAM. This will allow SMM software to reference "through" SMM space to update the display even when SMM is mapped over the VGA range. Software should ensure that D_OPEN=1 and D_CLS=1 are not set at the same time.
4	<b>SMM Space Locked (D_LCK).</b> When D_LCK is set to 1 then D_OPEN is reset to 0 and D_LCK, D_OPEN, C_BASE_SEG, H_SMRAM_EN, TSEG_SZ and TSEG_EN, and G_SMROME become "Read Only". GBA[15:0] and GAR[7:0] registers also become "Read Only" after D_LCK is set. D_LCK can be set to 1 via a normal configuration space write but can only be cleared by a Full Reset. The combination of D_LCK and D_OPEN provide convenience with security. The BIOS can use the D_OPEN function to initialize SMM space and then use D_LCK to "lock down" SMM space in the future so that no application software (or BIOS itself) can violate the integrity of SMM space, even if the program has knowledge of the D_OPEN function.  Note: This bit is only reset to '0' on a Full Reset.
3	<b>Global SMRAM Enable (G_SMROME).</b> If set to a 1, then Compatible SMRAM functions is enabled, providing 128 KB of DRAM accessible at the A0000h address while in SMM (ADS# with SMM decode). To enable Extended SMRAM function this bit has be set to 1. Refer to the Section 4.1.5 for more details.  Once D_LCK is set, this bit becomes read only.
2:0	<b>Compatible SMM Space Base Segment (C_BASE_SEG) (RO).</b> This field indicates the location of SMM space. "SMM DRAM" is not remapped. It is simply "made visible" if the conditions are right to access SMM space, otherwise the access is forwarded to the hub interface. C_BASE_SEG is hardwired to 010 to indicate that the MCH supports the SMM space at A0000h–BFFFFh.

### 3.4.26. ESMRAMC—Extended System Management RAM Control Register (Device 0)

Address Offset:	9Eh
Default:	38h
Access:	Read/Write/Lock
Size:	8 bits

The Extended SMRAM register controls the configuration of Extended SMRAM space. The Extended SMRAM (E\_SMRAM) memory provides a write-back cacheable SMRAM memory space that is above 1 MB.

Bit	Description
7	<p><b>H_SMRAM_EN (H_SMRAME).</b> Controls the SMM memory space location (i.e. above 1 MB or below 1 MB) When G_SMRAME is 1 and H_SMRAME bit is set to 1, the high SMRAM memory space is enabled. SMRAM accesses from 0FEEA0000h to 0FEEBFFFFh are remapped to DRAM address 000A0000h to 000BFFFFh.</p> <p>Once D_LCK is set, this bit becomes read only.</p>
6	<p><b>E_SMRAM_ERR (E_SMERR).</b></p> <p>1 = This bit is set when the host accesses the defined memory ranges in Extended SMRAM (High Memory and T-segment) while not in SMM space and with the D-OPEN bit = 0.</p> <p>0 = It is software's responsibility to clear this bit. The software must write a 1 to this bit to clear it</p>
5	<b>SMRAM_Cache (SM_CACHE).</b> This bit is forced to '1' by MCH .
4	<b>SMRAM_L1_EN (SM_L1).</b> This bit is forced to '1' by MCH.
3	<b>SMRAM_L2_EN (SM_L2).</b> This bit is forced to '1' by MCH.
2:1	<p><b>TSEG_SZ[1-0] (T_SZ).</b> Selects the size of the TSEG memory block if enabled. This memory is taken from the top of DRAM space (i.e., TOM minus TSEG_SZ), which is no longer claimed by the memory controller (all accesses to this space are sent to the hub interface, if TSEG_EN is set). This field decodes as follows:</p> <p>00 = (TOM–128 KB) to TOM</p> <p>01 = (TOM–256 KB) to TOM</p> <p>10 = (TOM–512 KB) to TOM</p> <p>11 = (TOM–1 MB) to TOM</p> <p>Once D_LCK is set, this bit becomes read only.</p>
0	<p><b>TSEG_EN (T_EN).</b> Enabling of SMRAM memory (TSEG, 128 KBs, 256 KBs, 512 KBs or 1 MB of additional SMRAM memory) for Extended SMRAM space only. When G_SMRAME =1 and TSEG_EN = 1, the TSEG is enabled to appear in the appropriate physical address space.</p> <p>Once D_LCK is set, this bit becomes read only.</p>

### 3.4.27. ACAPID— AGP Capability Identifier Register (Device 0)

Address Offset:	A0–A3h
Default:	00200002h
Access:	Read Only
Size:	32 bits

This register provides standard identifier for AGP capability.

Bit	Description
31:24	Reserved
23:20	<b>Major AGP Revision Number.</b> These bits provide a major revision number of AGP specification to which this version of MCH conforms. These bits are hardwired to 0010b. 0010b = Indicates AGP Rev. 2.x.
19:16	<b>Minor AGP Revision Number.</b> These bits provide a minor revision number of AGP specification to which this version of MCH conforms. This number is hardwired to 0000. 0000b = Rev x.0). Together with major revision number this field identifies MCH as an AGP REV 2.0 compliant device.
15:8	<b>Next Capability Pointer.</b> AGP capability is the first and the last capability described via the capability pointer mechanism and, therefore, these bits are hardwired to 0 to indicate the end of the capability linked list. 0000_0000b = End of the capability linked list
7:0	<b>AGP Capability ID.</b> This field identifies the linked list item as containing AGP registers. This field has the value 0000_0010b as assigned by the PCI SIG.

### 3.4.28. AGPSTAT—AGP Status Register (Device 0)

Address Offset: A4–A7h  
 Default: 1F000207h  
 Access: Read Only  
 Size: 32 bits

This register reports AGP device capability/status.

Bit	Description
31:24	<b>RQ.</b> Hardwired to 1Fh. 1Fh = Maximum of 32 outstanding AGP command requests can be handled by the MCH.
23:10	Reserved
9	<b>SBA.</b> Hardwired to 1. This bit indicates that the MCH supports side band addressing.
8:6	Reserved
5	<b>4G.</b> Hardwired to 0. This bit indicates that the MCH does not support addresses greater than 4 gigabytes. It is hardwired to 0.
4	<b>FW.</b> The value in this bit indicates the MCH capability of using Fast Writes protocol. This status bit returns a 1 when the Fast Write Capability Enable bit (bit 1 of AGPCTRL) is set to 1. By default, this bit returns a 0, indicating that the MCH is not capable of Fast Writes.  To use Fast Writes, both the Fast Write Capability Enable bit (bit 1 of AGPCTRL) and Fast Write Enable bit (bit 4 of AGPCMD) must be set to 1.
3	Reserved
2:0	<b>RATE.</b> After reset the MCH reports its data transfer rate capability. Bit 0 identifies if AGP device supports 1x data transfer mode; bit 1 identifies if AGP device supports 2x data transfer mode; bit 2 identifies if AGP device supports 4x data transfer mode. 1x, 2x, and 4x data transfer modes are supported by the MCH and, therefore, this bit field has a Default Value = 111.  Note that the selected data transfer mode apply to both AD bus and SBA bus. It also applies to Fast Writes if they are enabled.  111 = 4x mode supported

### 3.4.29. AGPCMD—AGP Command Register (Device 0)

Address Offset: A8–ABh  
 Default: 00000000h  
 Access: Read/Write  
 Size: 32 bits

This register provides control of the AGP operational parameters.

Bit	Description
31:10	Reserved.
9	<b>SBA Enable.</b> 1 = Enable sideband addressing mechanism. 0 = Disable
8	<b>AGP Enable.</b> 1 = Enable. The MCH responds to AGP operations delivered via PIPE# or to operations delivered via SBA, if the <i>AGP Side Band Enable</i> bit is also set to 1. 0 = Disable. The MCH ignores all AGP operations, including the sync cycle. Any AGP operations received while this bit is set to 1 will be serviced even, if this bit is reset to 0. If this bit transitions from a 1 to a 0 on a clock edge in the middle of an SBA command being delivered in 1X mode, the command will be issued.
7:6	Reserved.
5	<b>4G.</b> The MCH, as an AGP target, does not support addressing greater than 4 gigabytes. This bit is hardwired to 0.
4	<b>Fast Write Protocol Enable (FWPE).</b> To use Fast Write Protocol, both this bit and FW Capability Enable (bit 1 of AGPCTRL) must be set to 1. The Data Rate bits must also be set to either 2X or 4X. 1 = Enable. Fast Write protocol is enabled when the Fast Write Capability Enable bit (bit 1 of AGPCTRL) is also set to 1. When both bits are set to 1, the MCH uses the Fast Write protocol for Memory Write transactions from the MCH to the AGP master. Fast Writes occur at the data transfer rate as selected by the data rate bits (2:0) in this register. 0 = Disable. When this bit is 0 or when the data rate bits are set to 1x mode, the Memory Write transactions from the MCH to the AGP master use the standard PCI protocol.
3	Reserved.
2:0	<b>Data Rate.</b> The setting of these bits determines the AGP data transfer rate. One ( <i>and only one</i> ) bit in this field must be set to indicate the desired data transfer rate. Bit 0: 1X, Bit 1: 2X, Bit 2: 4X . The same bit must be set on both master and target. Configuration software updates this field by setting only one bit that corresponds to the capability of AGP master (after that capability has been verified by accessing the same functional register within the AGP masters configuration space.) Note: This field applies to AD and SBA buses. It also applies to Fast Writes if they are enabled. 001 = 1x 010 = 2x 100 = 4x All other bit combinations are illegal

### 3.4.30. AGPCTRL— AGP Control Register (Device 0)

Address Offset: B0–B3h  
 Default: 00000000h  
 Access: Read/Write  
 Size: 32 bits

This register provides for additional control of the AGP interface.

Bit	Description
31:2	Reserved
1	<p><b>Fast Write Capability Enable (FWCE).</b> To use Fast Write Protocol, both this bit and FWPE Enable bit (bit 4 of AGPCMD) must be set to 1.</p> <p>1 = Enable. The MCH is capable of using Fast Write protocol. By default, it is set to 0. When this bit is set to 1, the Fast Write Status Register (bit 4 of AGPSTAT) will also indicate a 1.</p> <p>0 = Disable</p>
0	<p><b>4x Override.</b> When this bit is set to 1 the RATE[2] bit in the AGPSTS register will be read as a 0. This register bit allows BIOS to force 2x mode for 3.3V AGP operation.</p> <p>Note: This bit must be set by the BIOS before AGP configuration.</p>



### 3.4.31. APSIZE— Aperture Size Register (Device 0)

Address Offset:	B4h
Default:	00h
Access:	Read/Write
Size:	8 bits

This register determines the effective size of the Graphics Aperture used for a particular MCH configuration. This register can be updated by the MCH specific BIOS configuration sequence before the PCI standard bus enumeration sequence takes place. If the register is not updated then a default value will select an aperture of maximum size (i.e., 256 MB). The size of the table that corresponds to a 256 MB aperture is not practical for most applications; therefore, these bits must be programmed to a smaller practical value that will force adequate address range to be requested via APBASE register from the PCI configuration software.

Bit	Description																																																								
7:6	Reserved																																																								
5:0	<p><b>Graphics Aperture Size (APSIZE).</b> Each bit in APSIZE[5:0] operates on similarly ordered bits in APBASE[27:22] of the Aperture Base configuration register. When a particular bit of this field is 0 it forces the similarly ordered bit in APBASE[27:22] to behave as “hardwired” to 0. When a particular bit of this field is set to 1 it allows corresponding bit of the APBASE[27:22] to be read/write accessible. Only the following combinations are allowed:</p> <table> <thead> <tr> <th>5</th> <th>4</th> <th>3</th> <th>2</th> <th>1</th> <th>0</th> <th>Aperture Size</th> </tr> </thead> <tbody> <tr> <td>1</td> <td>1</td> <td>1</td> <td>1</td> <td>1</td> <td>1</td> <td>4 MB</td> </tr> <tr> <td>1</td> <td>1</td> <td>1</td> <td>1</td> <td>1</td> <td>0</td> <td>8 MB</td> </tr> <tr> <td>1</td> <td>1</td> <td>1</td> <td>1</td> <td>0</td> <td>0</td> <td>16 MB</td> </tr> <tr> <td>1</td> <td>1</td> <td>1</td> <td>0</td> <td>0</td> <td>0</td> <td>32 MB</td> </tr> <tr> <td>1</td> <td>1</td> <td>0</td> <td>0</td> <td>0</td> <td>0</td> <td>64 MB</td> </tr> <tr> <td>1</td> <td>0</td> <td>0</td> <td>0</td> <td>0</td> <td>0</td> <td>128 MB</td> </tr> <tr> <td>0</td> <td>0</td> <td>0</td> <td>0</td> <td>0</td> <td>0</td> <td>256 MB</td> </tr> </tbody> </table> <p>Default for APSIZE[5:0]=000000b forces default APBASE[27:22] =000000b (i.e., all bits respond as “hardwired” to 0). This provides maximum aperture size of 256 MB. As another example, programming APSIZE[5:0]=111000b hardwires APBASE[24:22]=000b and while enabling APBASE[27:25] as read/write programmable.</p>	5	4	3	2	1	0	Aperture Size	1	1	1	1	1	1	4 MB	1	1	1	1	1	0	8 MB	1	1	1	1	0	0	16 MB	1	1	1	0	0	0	32 MB	1	1	0	0	0	0	64 MB	1	0	0	0	0	0	128 MB	0	0	0	0	0	0	256 MB
5	4	3	2	1	0	Aperture Size																																																			
1	1	1	1	1	1	4 MB																																																			
1	1	1	1	1	0	8 MB																																																			
1	1	1	1	0	0	16 MB																																																			
1	1	1	0	0	0	32 MB																																																			
1	1	0	0	0	0	64 MB																																																			
1	0	0	0	0	0	128 MB																																																			
0	0	0	0	0	0	256 MB																																																			

### 3.4.32. ATTBASE—Aperture Translation Table Base Register (Device 0)

Address Offset:	B8–BBh
Default:	00000000h
Access:	Read/Write
Size:	32 bits

This register provides the starting address of the Graphics Aperture Translation Table Base located in the main DRAM. This value is used by the MCH Graphics Aperture address translation logic (including the GTLB logic) to obtain the appropriate address translation entry required during the translation of the aperture address into a corresponding physical DRAM address. The ATTBASE register may be dynamically changed.

**Note:** The address provided via ATTBASE is 4 KB aligned.

Bit	Description
31:12	<b>Translation Table Base Address Pointer.</b> This field contains a pointer to the base of the translation table used to map memory space addresses in the aperture range to addresses in main memory.
11:0	Reserved

### 3.4.33. AMTT—AGP Interface Multi-Transaction Timer Register (Device 0)

Address Offset:	BCh
Default:	00h
Access:	Read/Write
Size:	8 bits

AMTT controls the amount of time that the MCH arbiter allows the AGP master to perform multiple back-to-back transactions. The MCH AMTT mechanism is used to optimize the performance of the AGP master (using PCI protocol) that performs multiple back-to-back transactions to fragmented memory ranges (and as a consequence it can not use long burst transfers). The AMTT mechanism applies to the host-AGP transactions as well; it guarantees to the processor a fair share of the AGP interface bandwidth.

The number of clocks programmed in the AMTT represents the guaranteed time slice (measured in 66 MHz clocks) allotted to the current agent (either AGP master or Host bridge) after which the AGP arbiter will grant the bus to another agent. The default value of 00h disables this function. The AMTT value can be programmed with 8 clock granularity. For example, if the AMTT is programmed to 18h, then the selected value corresponds to the time period of 24 AGP (66 MHz) clocks.

Bit	Description
7:3	<b>Multi-Transaction Timer Count Value.</b> The number programmed in these bits represents the guaranteed time slice (measured in eight 66 MHz clock granularity) allotted to the current agent (either AGP master or MCH) after which the AGP arbiter will grant the bus to another agent.
2:0	Reserved.

### 3.4.34. LPTT—Low Priority Transaction Timer Register (Device 0)

Address Offset:	BDh
Default:	00h
Access:	Read/Write
Size:	8 bits

LPTT is an 8-bit register similar in a function to AMTT. This register is used to control the minimum tenure on the AGP for low priority data transaction (both reads and writes) issued using PIPE# or SB mechanisms.

The number of clocks programmed in the LPTT represents the guaranteed time slice (measured in 66 MHz clocks) allotted to the current low priority AGP transaction data transfer state. This does not necessarily apply to a single transaction but it can span over multiple low-priority transactions of the same type. After this time expires the AGP arbiter may grant the bus to another agent if there is a pending request. The LPTT does not apply in the case of high-priority request where ownership is transferred directly to high-priority requesting queue. The default value of LPTT is 00h and disables this function. The LPTT value can be programmed with 8 clock granularity. For example, if the LPTT is programmed to 10h, then the selected value corresponds to the time period of 16 AGP (66 MHz) clocks.

Bit	Description
7:3	<b>Low Priority Transaction Timer Count Value.</b> The number of clocks programmed in these bits represents the guaranteed time slice (measured in eight 66 MHz clock granularity) allotted to the current low priority AGP transaction data transfer state.
2:0	Reserved.

### 3.4.35. MCHCFG—MCH Configuration Register (Device 0)

Offset:	BE–BFh
Default:	00s0_0000_0000_0s00b
Access:	Read/Write, Read Only
Size:	16 bits

Bit	Description
15:14	Reserved.
13	<b>Host Frequency—RO.</b> This bit is used to determine the host frequency of the Processor System Bus. It is set by an external strapping option at reset and is Read Only.  0 = 100 MHz 1 = 133 MHz

Bit	Description																					
12:11	<p><b>Rambus Frequency—R/W.</b> These bits are written by the BIOS after polling the Rambus DRAMs and finding the least common denominator speed. The Host-Bus Frequency (Bit 13 of this register) must also be used to determine the exact RDRAM frequency.</p> <table border="1"> <thead> <tr> <th>Bit13</th> <th>Bit[12:11]</th> <th>RDRAM Frequency</th> </tr> </thead> <tbody> <tr> <td>0</td> <td>00</td> <td>300 MHz (Default Mode)</td> </tr> <tr> <td>X</td> <td>01</td> <td>400 MHz</td> </tr> <tr> <td>1</td> <td>00</td> <td>356 MHz</td> </tr> <tr> <td>1</td> <td>10</td> <td>266 MHz</td> </tr> <tr> <td>0</td> <td>10</td> <td>Reserved</td> </tr> <tr> <td>X</td> <td>11</td> <td>Reserved</td> </tr> </tbody> </table> <p>Note: Only when bit 13 is a 1 (indicating 133 MHz Processor System Bus) can BIOS select the 356 or 266 MHz modes.</p>	Bit13	Bit[12:11]	RDRAM Frequency	0	00	300 MHz (Default Mode)	X	01	400 MHz	1	00	356 MHz	1	10	266 MHz	0	10	Reserved	X	11	Reserved
Bit13	Bit[12:11]	RDRAM Frequency																				
0	00	300 MHz (Default Mode)																				
X	01	400 MHz																				
1	00	356 MHz																				
1	10	266 MHz																				
0	10	Reserved																				
X	11	Reserved																				
10:9	Reserved																					
8:7	<p><b>DRAM Data Integrity Mode (DDIM)—R/W.</b> These bits select one of 4 DRAM data integrity modes.</p> <p>00 = Non-ECC (Byte-Wise Writes supported) (DEFAULT)</p> <p>01 = EC Only (Error Checking with No correction)</p> <p>10 = ECC Mode (Generation and Error Checking/Correction)</p> <p>11 = ECC Mode with hardware scrubbing enabled. Same as DDIM="10", plus corrected data written to DRAM</p>																					
6	Reserved.																					
5	<p><b>MDA Present (MDAP)—R/W.</b> This bit works with the VGA Enable bit in the BCTRL register of device 1 to control the routing of host initiated transactions targeting MDA compatible I/O and memory address ranges. This bit should not be set when the VGA Enable bit is not set. If the VGA enable bit is set, then accesses to IO address range x3BCh–x3BFh are forwarded to the hub interface. If the VGA enable bit is not set then accesses to IO address range x3BCh–x3BFh are treated just like any other IO accesses i.e. the cycles are forwarded to AGP if the address is within IOBASE and IOLIMIT and ISA enable bit is not set, otherwise they are forwarded to the hub interface. MDA resources are defined as the following:</p> <p>Memory: 0B0000h–0B7FFFh (Device 1)</p> <p>I/O: 3B4h, 3B5h, 3B8h, 3B9h, 3BAh, 3BFh, (including ISA address aliases, A[15:10] are not used in decode)</p> <p>Any I/O reference that includes the I/O locations listed above, or their aliases, will be forwarded to the hub interface even if the reference includes I/O locations not listed above.</p> <p>The following table shows the behavior for all combinations of MDA and VGA:</p> <table border="1"> <thead> <tr> <th>VGA</th> <th>MDA</th> <th>Behavior</th> </tr> </thead> <tbody> <tr> <td>0</td> <td>0</td> <td>All references to MDA and VGA go to the hub interface</td> </tr> <tr> <td>0</td> <td>1</td> <td>Illegal Combination (DO NOT USE)</td> </tr> <tr> <td>1</td> <td>0</td> <td>All references to VGA go to AGP. MDA-only references (I/O address 3BF and aliases ) will go to the hub interface.</td> </tr> <tr> <td>1</td> <td>1</td> <td>VGA references go to AGP; MDA references go to the hub interface.</td> </tr> </tbody> </table>	VGA	MDA	Behavior	0	0	All references to MDA and VGA go to the hub interface	0	1	Illegal Combination (DO NOT USE)	1	0	All references to VGA go to AGP. MDA-only references (I/O address 3BF and aliases ) will go to the hub interface.	1	1	VGA references go to AGP; MDA references go to the hub interface.						
VGA	MDA	Behavior																				
0	0	All references to MDA and VGA go to the hub interface																				
0	1	Illegal Combination (DO NOT USE)																				
1	0	All references to VGA go to AGP. MDA-only references (I/O address 3BF and aliases ) will go to the hub interface.																				
1	1	VGA references go to AGP; MDA references go to the hub interface.																				
4	Reserved.																					
3	<p><b>AGP I/O Buffer Mode—RO.</b> This status bit reports the AGP I/O buffer mode. It indicates if the AGP interface is operating at 3.3V or 1.5V.</p> <p>1 = 3.3V</p> <p>0 = 1.5V</p>																					

Bit	Description
2	<p><b>In-Order Queue Depth (IOQD)—RO.</b> This bit reflects the value sampled on HA7# on the deassertion of the CPURST#. It indicates the depth of the host bus in-order queue (i.e., level of host bus pipelining). If IOQD is set to 1 (HA7# sampled 1; i.e., undriven on the host bus), then the depth of the host bus in-order queue is configured to the maximum allowed by the host bus protocol (i.e., 8). However, the actual maximum supported by the MCH is 6 and it is controlled by the MCH's host interface logic using the BNR# signaling mechanism. If the IOQD bit is set to 0 (HA7# is sampled asserted; i.e., 0), then depth of the host bus in-order queue is set to 1 (i.e., no pipelining support on the host bus).</p> <p>Note that HA7# is not driven by the MCH during CPURST#. If an IOQ size of 1 is desired, HA7# must be driven low during CPURST# by an external source.</p>
1:0	Reserved.

### 3.4.36. EAP—Error Address Pointer Register (Device 0)

Address Offset:	C4–C7h
Default:	XXXXXXXXh (X=Undefined)
Access	Read Only
Size	32 Bits

This register stores the DRAM address and syndrome when an ECC error occurs.

Bit	Description
31:12	<p><b>Error Address Pointer (EAP).</b> This field is used to store the 4 KB block of main memory where an error (single bit or multi-bit error) has occurred. Note that the value of this bit field represents the address of the first single or the first multiple bit error occurrence after the error flag bits in the ERRSTS register have been cleared by software. A multiple bit error will overwrite a single bit error.</p> <p>Once the error flag bits are set as a result of an error, this bit field is locked and does not change as a result of a new error.</p>
11:8	Reserved.
7:0	<p><b>Error Syndrome (ES).</b> This field is used to store the ECC Syndrome for an error (single bit or multi-bit error). Note that the value of this bit field represents the syndrome of the first single or the first multiple bit error occurrence after the error flags have been cleared by software. A multiple bit error overwrites a single bit error. This register is loaded, locked, and unlocked along with the EAP field.</p>

### 3.4.37. ERRSTS—Error Status Register (Device 0)

Address Offset: C8–C9h  
 Default: 0000h  
 Access: Read/Write Clear  
 Size: 16 bits

This register is used to report various error conditions via the hub interface cycles. An SERR, SMI, or SCI error message may be generated (on a zero to one transition of any of these flags) via the hub interface when enabled in the respective PCICMD/ERRCMD, SMICMD, or SCICMD registers.

Bit	Description
15:12	Reserved.
11	<b>RDRAM Thermal Sensor Flag (RDTSF)—R/WC.</b> 1 = RDRAM Thermal Sensor was detected active. 0 = Software must write a 1 to clear this status bit.
10	Reserved.
9	<b>LOCK to non-DRAM Memory Flag (LCKF)—R/WC.</b> 1 = Host-initiated LOCK cycle targeting non-DRAM memory space occurred. 0 = Software must write a 1 to clear this status bit
8	Reserved.
7	<b>DRAM Throttle Flag (DRTF)—R/WC.</b> 1 = DRAM Read/Write Throttling condition occurred. 0 = Software must write a 1 to clear this status bit.
6	Reserved.
5	<b>Received Unimplemented Special Cycle Completion FLAG (UNSC)—R/WC.</b> 1 = MCH initiated a request via the hub interface and was terminated with a Unimplemented Special Cycle completion. 0 = Software must write a 1 to clear this status bit.
4	<b>AGP Access Outside of Graphics Aperture Flag (OOGF)—R/WC.</b> 1 = AGP access occurred to an address that is outside of the graphics aperture range. 0 = Software must write a 1 to clear this status bit.
3	<b>Invalid AGP Access Flag (IAAF)—R/WC.</b> 1 = AGP access was attempted outside of the graphics aperture and either to the 640 kB – 1 MB range or above the top of memory. 0 = Software must write a 1 to clear this status bit.

Bit	Description
2	<p><b>Invalid Graphics Aperture Translation Table Entry Flag (ITTEF)—R/WC.</b></p> <p>1 = Invalid translation table entry was returned in response to an AGP access to the graphics aperture. Invalid translation table entries include the following:</p> <ul style="list-style-type: none"> <li>— Invalid bit set in table entry.</li> <li>— Uncorrected ECC error (correction turned off or multiple-bit error with correction on).</li> <li>— Translated address hits PAM region.</li> <li>— Translated address hits enabled physical SMM space.</li> </ul> <p>0 = Software must write a 1 to clear this status bit.</p>
1	<p><b>Multiple-bit ECC (uncorrectable) Error Flag (MEF)—R/WC.</b></p> <p>1 = A memory read data transfer had an uncorrectable multiple-bit error. When this bit is set, the address that caused the error and the syndrome are logged in the EAP register. Once this bit is set, the EAP and ES fields are locked until the processor clears this bit by writing a 1. Software uses bits [1:0] to detect whether the logged error address is for Single or Multiple-bit error.</p> <p>0 = Once software completes the error processing, a value of 1 is written to this bit to clear the value (back to 0) and unlock the error logging mechanism.</p>
0	<p><b>Single-bit (correctable) ECC Error Flag (SEF)—R/WC.</b></p> <p>1 = A memory read data transfer had a single-bit correctable error and the corrected data was sent for the access. When this bit is set the address that caused the error and the syndrome are logged in the EAP register. Once this bit is set the EAP and ES fields are locked to further single bit error updates until the processor clears this bit by writing a 1. A multiple bit error that occurs after this bit is set will overwrite the EAP and ES fields with the multiple bit error address and syndrome and the MEF bit will also be set.</p> <p>0 = Software must write a 1 to clear this bit and unlock the error logging mechanism.</p>

### 3.4.38. ERRCMD—Error Command Register (Device 0)

Address Offset:	CA–CBh
Default:	0000h
Access:	Read/Write
Size:	16 bits

This register enables various errors to generate a SERR message via the hub interface. Since the MCH does not have an SERR# signal, SERR messages are passed from the MCH to the I/O Controller Hub over the hub interface. The actual generation of the SERR message is globally enabled for Device 0 via the PCI Command register.

**Note:** An error can generate one and only one error message via the hub interface. It is software's responsibility to make sure that when an SERR error message is enabled for an error condition; SMI and SCI error messages are disabled for that same error condition.

Bit	Description
15:12	Reserved.
11	<p><b>SERR on RDRAM Thermal Sensor Trip Condition.</b></p> <p>1 = MCH generates a SERR error message over the hub interface when the RDRAM Thermal Sensor Flag is set.</p> <p>0 = Disable reporting of this condition.</p>
10	Reserved.
9	<p><b>SERR on LOCK to non-DRAM Memory.</b></p> <p>1 = MCH generates an SERR error message via the hub interface when a host-initiated LOCK transaction targeting non-DRAM memory space occurs.</p> <p>0 = Disable reporting of this condition.</p>
8	Reserved.
7	<p><b>SERR on DRAM Throttle Condition.</b></p> <p>1 = MCH generates an SERR special cycle over the hub interface when a DRAM Read or Write Throttle condition occurs.</p> <p>0 = Disable reporting of this condition.</p>
6	<p><b>SERR on Receiving Target Abort on the Hub Interface.</b></p> <p>1 = MCH generates an SERR special cycle when an MCH originated hub interface cycle is terminated with a Target Abort.</p> <p>0 = Disable reporting of this condition.</p>
5	<p><b>SERR on Receiving Unimplemented Special Cycle Completion.</b></p> <p>1 = MCH generates an SERR special cycle over the hub interface when an MCH initiated request via the hub interface is terminated with a Unimplemented Special Cycle completion.</p> <p>0 = Disable reporting of this condition.</p>
4	<p><b>SERR on AGP Access Outside of Graphics Aperture.</b></p> <p>1 = MCH generates a SERR error message via the hub interface when an AGP access occurs to an address outside of the graphics aperture.</p> <p>0 = Disable reporting of this condition.</p>



Bit	Description
3	<p><b>SERR on Invalid AGP Access.</b></p> <p>1 = Enable. MCH generates an SERR error message via the hub interface when an AGP access occurs to an address outside of the graphics aperture and either to the 640 KB – 1 MB range or above the top of memory.</p> <p>0 = Disable reporting of this condition.</p>
2	<p><b>SERR on Access to Invalid Graphics Aperture Translation Table Entry.</b></p> <p>1 = Enable. MCH generates an SERR hub interface special cycle when an invalid translation table entry was returned in response to a AGP access to the graphics aperture.</p> <p>0 = Disable reporting of this condition via SERR messaging.</p>
1	<p><b>SERR on Multiple-bit ECC Error.</b></p> <p>1 = Enable. A SERR error message via the hub interface will be generated by the MCH when it detects a Multiple-bit ECC error. For systems that do not support ECC this field must be set to 0.</p> <p>0 = Disable reporting of this condition.</p>
0	<p><b>SERR on Single-bit ECC Error.</b></p> <p>1 = Enable. SERR error message via the hub interface will be generated by the MCH when it detects a Single-bit ECC error. For systems that do not support ECC this field must be set to 0.</p> <p>0 = Disable reporting of this condition.</p>

### 3.4.39. SMICMD—SMI Command Register (Device 0)

Address Offset:	CCh
Default:	00h
Access:	Read/Write
Size:	8 bits

This register enables various errors to generate a SMI hub interface special cycle. When an error flag is set in the ERRSTS register it can generate a SERR, SMI, or SCI hub interface special cycle when enabled in the ERRCMD, SMICMD, or SCICMD registers respectively.

**Note:** An error can generate one and only one hub interface error special cycle. It is software's responsibility to make sure that when an SMI error message is enabled for an error condition, SERR and SCI error messages are disabled for that same error condition.

Bit	Description
7:3	Reserved.
2	<p><b>SMI on RDRAM Thermal Sensor Trip.</b></p> <p>1 = Enable. A SMI hub interface special cycle is generated by the MCH when it detects one or more RDRAM devices whose thermal sensor has tripped.</p> <p>0 = Disable reporting of this condition.</p>
1	<p><b>SMI on Multiple-bit ECC Error.</b></p> <p>1 = Enable. A SMI hub interface special cycle is generated by the MCH to the hub interface when it detects a Multiple-bit ECC error. For systems that do not support ECC, this field must be set to 0.</p> <p>0 = Disable reporting of this condition.</p>
0	<p><b>SMI on Single-bit ECC Error.</b></p> <p>1 = Enable. A SMI hub interface special cycle is generated by the MCH to the hub interface when it detects a Single-bit ECC error. For systems that do not support ECC, this field must be set to 0.</p> <p>0 = Disable reporting of this condition.</p>

### 3.4.40. SCICMD—SCI Command Register (Device 0)

Address Offset:	CDh
Default:	00h
Access:	Read/Write
Size:	8 bits

This register enables various errors to generate a SCI hub interface special cycle. When an error flag is set in the ERRSTS register, it can generate a SERR, SMI, or SCI hub interface special cycle when enabled in the ERRCMD, SMICMD, or SCICMD registers respectively.

**Note:** An error can generate one and only one hub interface error special cycle. It is software's responsibility to make sure that when an SCI error message is enabled for an error condition, SERR and SMI error messages are disabled for that same error condition.

Bit	Description
7:3	Reserved.
2	<b>SCI on RDRAM Thermal Sensor Trip.</b> 1 = Enable. A SCI hub interface special cycle via the hub interface will be generated by the MCH when it detects one or more RDRAM devices whose thermal sensor has tripped. 0 = Disable reporting of this condition.
1	<b>SCI on Multiple-bit ECC Error.</b> 1 = Enable. A SCI hub interface special cycle will be generated by the MCH to the hub interface when it detects a Multiple-bit ECC error. For systems that do not support ECC this field must be set to 0. 0 = Disable reporting of this condition.
0	<b>SCI on Single-bit ECC Error.</b> 1 = Enable. A SCI hub interface special cycle will be generated by the MCH to the hub interface when it detects a Single-bit ECC error. For systems that do not support ECC this field must be set to 0. 0 = Disable reporting of this condition.

### 3.4.41. SKPD—Scratchpad Data Register (Device 0)

Address Offset:	DE–DFh
Default:	XXXXh (X=Undefined)
Access:	Read/Write
Size:	16 bits

This register is not reset with RSTIN#. It defaults to an indeterminate value.

Bit	Description
15:0	<b>Scratchpad [15:0].</b> These bits are simply R/W storage bits that have no effect on the MCH functionality.

### 3.4.42. AGPBCTRL—AGP Buffer Control B Register (Device 0)

Address Offset: E8–EBh  
 Default: 0000\_0000h  
 Access: Read/Write  
 Size: 32 bits

This register provides additional control of the 3.3V AGP buffer strength.

Bit	Description
31:16	<b>AGP Buffer Strength Control 1.</b>
15:13	<b>AGP Buffer Strength Control 2.</b>
12:0	Reserved

### 3.4.43. RTCE—RDRAM Temperature Calibration Enable Register (Device 0)

Offset: F7h  
 Default: 00h  
 Access: Read/Write  
 Size: 8 bits

Bit	Description
7	<b>Temperature Calibration Disable.</b> 1 = The MCH does not generate any Temperature Calibration Commands on the Rambus interface. 0 = The MCH generates Temperature Calibration packets as normal (default).
6:0	Reserved

### 3.4.44. AGPAPPEND—AGP Append Disable Register (Device 0)

Address Offset: FEh  
 Default: 00h  
 Access: Read/Write  
 Size: 8 bits

This register provides additional control of the AGP interface.

Bit	Description
7:2	Reserved
1	<b>AGP Write Appending Disable (AGPWAD).</b> 1 = Disables write appending on CPU-to-AGP or hub interface-to-AGP cycles going to AGP. 0 = Enable (default)
0	Reserved

### 3.5. AGP Bridge Registers (Device 1)

Table 7. MCH Configuration Space (Device 1)

Address Offset	Register Symbol	Register Name	Default Value	Access
00–01h	VID1	Vendor Identification	8086h	RO
02–03h	DID1	Device Identification	250Fh	RO
04–05h	PCICMD1	PCI Command Register	0000h	RO, R/W
06–07h	PCISTS1	PCI Status Register	0020h	RO, R/WC
08	RID1	Revision Identification	03h	RO
09	—	Reserved	00h	—
0Ah	SUBC1	Sub-Class Code	04h	RO
0Bh	BCC1	Base Class Code	06h	RO
0Ch	—	Reserved	00h	—
0Dh	MLT1	Master Latency Timer	00h	R/W
0Eh	HDR1	Header Type	01h	RO
0F–17h	—	Reserved	00h	—
18h	PBUSN	Primary Bus Number	00h	RO
19h	SBUSN	Secondary Bus Number	00h	R/W
1Ah	SUBUSN	Subordinate Bus Number	00h	R/W
1Bh	SMLT	Secondary Bus Master Latency Timer	00h	R/W
1Ch	IOBASE	I/O Base Address Register	F0h	R/W
1Dh	IOLIMIT	I/O Limit Address Register	00h	R/W
1E–1Fh	SSTS	Secondary Status Register	02A0h	RO, R/WC
20–21h	MBASE	Memory Base Address Register	0000h	R/W
22–23h	MLIMIT	Memory Limit Address Register	0000h	R/W
24–25h	PMBASE	Prefetchable Memory Base Address Reg.	0000h	R/W
26–27h	PMLIMIT	Prefetchable Memory Limit Address Reg.	0000h	R/W
28–3Dh	—	Reserved	00h	—
3Eh	BCTRL	Bridge Control Register	00h	R/W
3Fh	—	Reserved	00h	—
40h	ERRCMD1	Error Command	00h	R/W
41–FFh	—	Reserved	00h	—

### 3.5.1. VID1—Vendor Identification Register (Device 1)

Address Offset:	00–01h
Default:	8086h
Access:	Read Only
Size:	16 bits

The VID Register contains the vendor identification number. This 16-bit register combined with the Device Identification Register uniquely identify any PCI device. Writes to this register have no effect.

Bit	Description
15:0	<b>Vendor Identification Number.</b> This is a 16-bit value assigned to Intel. Intel VID = 8086h.

### 3.5.2. DID1—Device Identification Register (Device 1)

Address Offset:	02–03h
Default:	250Fh
Access:	Read Only
Size:	16 bits

This 16-bit register combined with the Vendor Identification register uniquely identifies any PCI device. Writes to this register have no effect.

Bit	Description
15:0	<b>Device Identification Number.</b> This is a 16 bit value assigned to the MCH Device 1. MCH Device 1 DID = 250Fh

### 3.5.3. PCICMD1—PCI-PCI Command Register (Device 1)

Address Offset: 04–05h  
 Default: 0000h  
 Access: Read/Write, Read Only  
 Size: 16 bits

Bit	Descriptions
15:10	Reserved.
9	<b>Fast Back-to-Back. (Not implemented).</b> Hardwired to 0.
8	<p><b>SERR Message Enable (SERRE1).</b> This bit is a global enable bit for Device 1 SERR messaging. The MCH does not have an SERR# signal. The MCH communicates the SERR# condition by sending an SERR message to the I/O Controller Hub. Note that this bit only controls SERR messaging for the Device 1. Device 0 has its own SERRE bit to control error reporting for error conditions occurring on Device 0.</p> <p>1 = Enable. The MCH generates SERR messages over the hub interface for specific Device 1 error conditions that are individually enabled in the BCTRL register. The error status is reported in the PCISTS1 register.</p> <p>0 = Disable. The SERR message is not generated by the MCH for Device 1.</p>
7	<b>Address/Data Stepping. (Not implemented).</b> Hardwired to 0.
6	<b>Parity Error Enable (PERRE1). (Not implemented).</b> Hardwired to 0.
5	Reserved.
4	<b>Memory Write and Invalidate Enable—RO. (Not implemented).</b> Hardwired to 0.
3	<b>Special Cycle Enable—RO. (Not implemented).</b> Hardwired to 0.
2	<p><b>Bus Master Enable (BME1)—R/W.</b></p> <p>1 = Enable. AGP Master initiated FRAME# cycles are accepted by the MCH, if they hit a valid address decode range. Note that incoming Configuration cycles still need to be accepted, regardless of the setting of this Bus Master Enable bit. This bit has no affect on AGP Master originated SBA or PIPE# cycles.</p> <p>0 = Disable (default). AGP Master initiated FRAME# cycles are ignored by the MCH resulting in a Master Abort. Ignoring incoming cycles on the secondary side of the P2P bridge effectively disables the bus master on the primary side.</p>
1	<p><b>Memory Access Enable (MAE1)—R/W.</b></p> <p>1 = Enable. This bit must be set to 1 to enable the Memory and Prefetchable memory address ranges defined in the MBASE, MLIMIT, PMBASE, and PMLIMIT registers.</p> <p>0 = Disable. All of Device 1's memory space is disabled.</p>
0	<p><b>I/O Access Enable (IOAE1)—R/W.</b></p> <p>1 = Enable. This bit must be set to 1 to enable the I/O address range defined in the IOBASE, and IOLIMIT registers.</p> <p>0 = Disable. All of Device 1's I/O space is disabled</p>

### 3.5.4. PCISTS1—PCI-PCI Status Register (Device 1)

Address Offset:	06–07h
Default:	0020h
Access:	Read Only, Read/Write Clear
Size:	16 bits

PCISTS1 reports the occurrence of error conditions associated with the primary side of the “virtual” PCI-PCI bridge in the MCH. Since this device does not physically reside on PCI0, it reports the optimum operating conditions so that it does not restrict the capability of PCI0.

Bit	Descriptions
15	<b>Detected Parity Error (DPE1).</b> Hardwired to 0.
14	<b>Signaled System Error (SSE1).</b> 1 = MCH Device 1 generated an SERR message over the hub interface for any enabled Device 1 error condition. Device 1 error conditions are enabled in the PCICMD1 and BCTRL registers. Device 1 error flags are read/reset from the SSTS register. 0 = Software clears this bit by writing a 1 to it.
13	<b>Received Master Abort Status (RMAS1).</b> Hardwired to 0.
12	<b>Received Target Abort Status (RTAS1).</b> Hardwired to 0.
11	<b>Signaled Target Abort Status (STAS1).</b> Hardwired to 0.
10:9	<b>DEVSEL# Timing (DEVT1).</b> Hardwired to 00b.
8	<b>Data Parity Detected (DPD1).</b> Hardwired to 0.
7	<b>Fast Back-to-Back (FB2B1).</b> Hardwired to 0.
6	Reserved.
5	<b>33/30 or 66/60 MHz Capability.</b> Hardwired to 1 (indicates 66/60 MHz).
4:0	Reserved.



### 3.5.5. RID1—Revision Identification Register (Device 1)

Address Offset:	08h
Default:	03h
Access:	Read Only
Size:	8 bits

This register contains the revision number of the MCH Device 1. These bits are read only and writes to this register have no effect.

Bit	Description						
7:0	<p><b>Revision Identification Number.</b> This is an 8-bit value that indicates the revision identification number for the MCH Device 1.</p> <table> <thead> <tr> <th>Stepping</th> <th>Value</th> </tr> </thead> <tbody> <tr> <td>B-1</td> <td>03h</td> </tr> <tr> <td>B-2</td> <td>04h</td> </tr> </tbody> </table>	Stepping	Value	B-1	03h	B-2	04h
Stepping	Value						
B-1	03h						
B-2	04h						

### 3.5.6. SUBC1—Sub-Class Code Register (Device 1)

Address Offset:	0Ah
Default:	04h
Access:	Read Only
Size:	8 bits

This register contains the Sub-Class Code for the MCH Device 1. This code is 04h indicating a PCI-PCI Bridge device. The register is read only.

Bit	Description
7:0	<p><b>Sub-Class Code (SUBC1).</b> This is an 8-bit value that indicates the category of Bridge into which the MCH falls.</p> <p>04h = Host Bridge.</p>

### 3.5.7. BCC1—Base Class Code Register (Device 1)

Address Offset:	0Bh
Default:	06h
Access:	Read Only
Size:	8 bits

This register contains the Base Class Code of the MCH Device 1. This code is 06h indicating a Bridge device. This register is read only.

Bit	Description
7:0	<p><b>Base Class Code (BASEC).</b> This is an 8-bit value that indicates the Base Class Code for the MCH Device 1.</p> <p>06h = Bridge device.</p>

### 3.5.8. MLT1—Master Latency Timer Register (Device 1)

Address Offset:	0Dh
Default:	00h
Access:	Read/Write
Size:	8 bits

This functionality is not applicable. It is described here since these bits should be implemented as a read/write to prevent standard PCI-PCI bridge configuration software from getting “confused” .

Bit	Description
7:3	Not applicable, but support read/write operations. (Reads return previously written data.)
2:0	Reserved.

### 3.5.9. HDR1—Header Type Register (Device 1)

Offset:	0Eh
Default:	01h
Access:	Read Only
Size:	8 bits

This register identifies the header layout of the configuration space. No physical register exists at this location.

Bit	Descriptions
7:0	This read only field always returns 01h when read. Writes have no effect.

### 3.5.10. PBUSN—Primary Bus Number Register (Device 1)

Offset:	18h
Default:	00h
Access:	Read Only
Size:	8 bits

This register identifies that “virtual” PCI-PCI bridge is connected to bus #0.

Bit	Descriptions
7:0	<b>Primary Bus Number.</b> Hardwired to 0s.

### 3.5.11. SBUSN—Secondary Bus Number Register (Device 1)

Offset: 19h  
 Default: 00h  
 Access: Read /Write  
 Size: 8 bits

This register identifies the bus number assigned to the second bus side of the “virtual” PCI-PCI bridge (i.e., to AGP).

Bit	Descriptions
7:0	<b>Secondary Bus Number.</b> This number is programmed by the PCI configuration software to allow mapping of configuration cycles to AGP. (Default = 00h)

### 3.5.12. SUBUSN—Subordinate Bus Number Register (Device 1)

Offset: 1Ah  
 Default: 00h  
 Access: Read /Write  
 Size: 8 bits

This register identifies the subordinate bus (if any) that resides at the level below AGP.

Bit	Descriptions
7:0	<b>Subordinate Bus Number.</b> This number is programmed by the PCI configuration software to allow mapping of configuration cycles to AGP. (Default = 00h)

### 3.5.13. SMLT—Secondary Master Latency Timer Register (Device 1)

Address Offset:	1Bh
Default:	00h
Access:	Read/Write
Size:	8 bits

This register controls the bus tenure of the MCH on AGP. MLT controls the amount of time the MCH as a AGP/PCI bus master, can burst data on the AGP/PCI Bus. The count value is an 8 bit quantity; however, MLT[2:0] are reserved and assumed to be 0 when determining the count value. The MCH's MLT guarantees the AGP master a minimum amount of the system resources. When the MCH begins the first AGP FRAME# cycle after being granted the bus, the counter is loaded and enabled to count from the assertion of FRAME#. If the count expires while the MCH's grant is removed (due to AGP master request), the MCH will lose the use of the bus and the AGP master agent may be granted the bus. If MCH's bus grant is not removed, the MCH continues to own the AGP bus, regardless of the MLT expiration or idle condition. Note that the MCH must always properly terminate an AGP transaction with FRAME# negation prior to the final data transfer.

Bit	Description
7:3	<p><b>Secondary MLT Counter Value.</b> The number of clocks programmed in the MLT represents the guaranteed time slice (measured in 66 MHz AGP clocks) allotted to the MCH, after which it must complete the current data transfer phase and then surrender the bus as soon as its bus grant is removed. For example, if the MLT is programmed to 18h, the value is 24 AGP clocks.</p> <p>00h = Disable. SMLT is disabled. When the MLT is disabled, the burst time for the MCH is unlimited (i.e. the MCH can burst forever).</p>
2:0	Reserved.

### 3.5.14. IOBASE—I/O Base Address Register (Device 1)

Address Offset:	1Ch
Default:	F0h
Access:	Read/Write
Size:	8 bits

This register controls the host to AGP I/O access routing based on the following formula:

$$\text{IO\_BASE} \leq \text{address} \leq \text{IO\_LIMIT}$$

Only upper 4 bits are programmable. For the purpose of address decode, address bits A[11:0] are treated as 0s. Thus, the bottom of the defined I/O address range is aligned to a 4 KB boundary.

**Note:** BIOS must not set this register to 00h; otherwise 0CF8h/0CFCh accesses will be forwarded to AGP.

Bit	Description
7:4	<b>I/O Address Base.</b> Corresponds to A[15:12] of the I/O address. (Default=Fh)
3:0	Reserved.

### 3.5.15. IOLIMIT—I/O Limit Address Register (Device 1)

Address Offset:	1Dh
Default:	00h
Access:	Read/Write
Size:	8 bits

This register controls the host to AGP I/O access routing based on the following formula:

$$\text{IO\_BASE} \leq \text{address} \leq \text{IO\_LIMIT}$$

Only the upper 4 bits are programmable. For the purpose of address decode, address bits A[11:0] are assumed to be FFFh. Thus, the top of the defined I/O address range is at the top of a 4 KB aligned address block.

Bit	Description
7:4	<b>I/O Address Limit.</b> Corresponds to A[15:12] of the I/O address. Default=0h
3:0	Reserved.

### 3.5.16. SSTS—Secondary PCI-PCI Status Register (Device 1)

Address Offset:	1E–1Fh
Default:	02A0h
Access:	Read Only, Read/Write Clear
Size:	16 bits

SSTS is a 16-bit status register that reports the occurrence of error conditions associated with secondary side (i.e., AGP side ) of the “virtual” PCI-PCI bridge embedded within MCH.

Bit	Descriptions
15	<b>Detected Parity Error (DPE1).</b> Note that the function of this bit is not affected by the PERRE1 bit. Also note that PERR# is not implemented in the MCH. 1 = MCH detected a parity error in the address or data phase of AGP bus transactions. 0 = Software sets DPE1 to 0 by writing a 1 to it.
14	<b>Received System Error (SSE1). (Not implemented).</b> Hardwired to 0. The MCH does not have an SERR# signal pin.
13	<b>Received Master Abort Status (RMAS1).</b> 1 = MCH terminated a Host-to-AGP with an unexpected master abort. 0 = Software resets this bit to 0 by writing a 1 to it.
12	<b>Received Target Abort Status (RTAS1).</b> 1 = MCH-initiated transaction on AGP is terminated with a target abort. 0 = Software resets RTAS1 to 0 by writing a 1 to it.
11	<b>Signaled Target Abort Status (STAS1). (Not implemented).</b> Hardwired to 0. The MCH does not generate target abort on AGP.
10:9	<b>DEVSEL# Timing (DEVT1).</b> This 2-bit field indicates the timing of the DEVSEL# signal when the MCH responds as a target on AGP. 01 = Medium timing (hardwired). Indicates the time when a valid DEVSEL# can be sampled by the initiator of the PCI cycle.
8	<b>Data Parity Detected (DPD1). (Not implemented).</b> Hardwired to 0. MCH does not implement G_PERR# function. However, data parity errors are still detected and reported via the hub interface (if enabled by SERRE1 bit of the PCICMD1 register and bit 0 of the BCTRL register).
7	<b>Fast Back-to-Back (FB2B1).</b> This bit is hardwired to 1, since MCH as a target supports fast back-to-back transactions on AGP.
6	Reserved.
5	<b>33/30 or 66/60 MHz Capability.</b> Hardwired to 1 to indicate 66/60 MHz capability.
4:0	Reserved

### 3.5.17. MBASE—Memory Base Address Register (Device 1)

Address Offset:	20–21h
Default:	0000h
Access:	Read/Write
Size:	16 bits

This register controls the host to AGP non-prefetchable memory access routing based on the following formula:

$$\text{MEMORY\_BASE} \leq \text{address} \leq \text{MEMORY\_LIMIT}$$

The upper 12 bits of the register are read/write and correspond to the upper 12 address bits A[31:20] of the 32-bit address. The bottom 4 bits of this register are read-only and return zeroes when read. This register must be initialized by the configuration software. For the purpose of address decode, address bits A[19:0] are assumed to be 0. Thus, the bottom of the defined memory address range will be aligned to a 1 MB boundary.

Bit	Description
15:4	<b>Memory Address Base (MEM_BASE)</b> . Corresponds to A[31:20] of the memory address. (Default=0s)
3:0	Reserved.

### 3.5.18. MLIMIT—Memory Limit Address Register (Device 1)

Address Offset:	22–23h
Default:	0000h
Access:	Read/Write
Size:	16 bits

This register controls the host to AGP non-prefetchable memory access routing based on the following formula:

$$\text{MEMORY\_BASE} \leq \text{address} \leq \text{MEMORY\_LIMIT}$$

The upper 12 bits of the register are read/write and correspond to the upper 12 address bits A[31:20] of the 32-bit address. The bottom 4 bits of this register are read-only and return zeroes when read. This register must be initialized by the configuration software. For the purpose of address decode, address bits A[19:0] are assumed to be FFFFFh. Thus, the top of the defined memory address range will be at the top of a 1 MB aligned memory block.

**Note:** Memory range covered by MBASE and MLIMIT registers are used to map non-prefetchable AGP address ranges (typically, where control/status memory-mapped I/O data structures of the graphics controller will reside). PMBASE and PMLIMIT are used to map prefetchable address ranges (typically, graphics local memory). This segregation allows application of USWC space attribute to be performed in a true plug-and-play manner to the prefetchable address range for improved host-AGP memory access performance.

Bit	Description
15:4	<b>Memory Address Limit (MEM_LIMIT)</b> . Corresponds to A[31:20] of the memory address. (Default=0s)
3:0	Reserved.

### 3.5.19. PMBASE—Prefetchable Memory Base Address Register (Device 1)

Address Offset:	24–25h
Default:	0000h
Access:	Read/Write
Size:	16 bits

This register controls the host to AGP prefetchable memory accesses routing based on the following formula:

$$\text{PREFETCHABLE\_MEMORY\_BASE} \leq \text{address} \leq \text{PREFETCHABLE\_MEMORY\_LIMIT}$$

The upper 12 bits of the register are read/write and correspond to the upper 12 address bits A[31:20] of the 32-bit address. The bottom 4 bits of this register are read-only and return zeroes when read. This register must be initialized by the configuration software. For the purpose of address decode, address bits A[19:0] are assumed to be 0. Thus, the bottom of the defined memory address range is aligned to a 1 MB boundary.

Bit	Description
15:4	<b>Prefetchable Memory Address Base (PMEM_BASE)</b> . Corresponds to A[31:20] of the memory address. (Default=0s)
3:0	Reserved.

### 3.5.20. PMLIMIT—Prefetchable Memory Limit Address Register (Device 1)

Address Offset:	26–27h
Default:	0000h
Access:	Read/Write
Size:	16 bits

This register controls the host to AGP prefetchable memory accesses routing based on the following formula:

$$\text{PREFETCHABLE\_MEMORY\_BASE} = < \text{address} = < \text{PREFETCHABLE\_MEMORY\_LIMIT}$$

The upper 12 bits of the register are read/write and correspond to the upper 12 address bits A[31:20] of the 32-bit address. The bottom 4 bits of this register are read-only and return zeroes when read. This register must be initialized by the configuration software. For the purpose of address decode, address bits A[19:0] are assumed to be FFFFFh. Thus, the top of the defined memory address range is at the top of a 1 MB aligned memory block.

**Note:** Prefetchable memory range is supported to allow segregation by the configuration software between the memory ranges that must be defined as UC and the ones that can be designated as a USWC (i.e., prefetchable) from the processor perspective.

Bit	Description
15:4	<b>Prefetchable Memory Address Limit (PMEM_LIMIT)</b> . Corresponds to A[31:20] of the memory address. (Default=0)
3:0	Reserved.



### 3.5.21. BCTRL—PCI-PCI Bridge Control Register (Device 1)

Address Offset:	3Eh
Default:	00h
Access:	Read/Write
Size	8 bits

This register provides extensions to the PCICMD1 register that are specific to PCI-PCI bridges. The BCTRL provides additional control for the secondary interface (i.e., AGP) as well as some bits that affect the overall behavior of the “virtual” PCI-PCI bridge embedded within MCH (e.g., VGA compatible address ranges mapping).

Bit	Descriptions															
7	<b>Fast Back to Back Enable.</b> Hardwired to 0. Since there is only one target allowed on AGP, this bit is not used. The MCH does not generate FB2B cycles in 1x mode, but does generate FB2B cycles in 2x and 4x Fast Write modes.															
6	<b>Secondary Bus Reset.</b> Hardwired to 0. The MCH does not support generation of reset via this bit on the AGP. Note that the only way to perform a hard reset of the AGP is via the system reset either initiated by software or hardware via the I/O Controller Hub.															
5	<b>Master Abort Mode.</b> Hardwired to 0. This means that when acting as a master on AGP and a master abort occurs, the MCH drops writes and returns all “1s” during reads.															
4	Reserved.															
3	<p><b>VGA Enable.</b> Controls the routing of host initiated transactions targeting VGA compatible I/O, and memory address ranges.</p> <p>1 = The MCH forwards the following host accesses to AGP:</p> <ul style="list-style-type: none"> <li>• Memory accesses in the range 0A0000h–0BFFFFh</li> <li>• I/O addresses where A[9:0] are in the ranges 3B0h–3BBh and 3C0h–3DFh (inclusive of ISA address aliases; A[15:10] are not decoded)</li> </ul> <p>When this bit is 1, forwarding of these accesses issued by the host is independent of the I/O address and memory address ranges defined by the previously defined base and limit registers. Forwarding of these accesses is also independent of the settings of the bit 2 (ISA Enable) of this register, if this bit is 1. If the VGA enable bit is 1, accesses to IO address range x3BCh–x3BFh are forwarded to the hub interface. If the VGA enable bit is 0, accesses to IO address range x3BCh–x3BFh are treated just like any other I/O accesses (i.e., the cycles are forwarded to AGP, if the address is within IOBASE and IOLIMIT and ISA enable bit is not set; otherwise, they are forwarded to the hub interface.</p> <p>0 = VGA compatible memory and I/O range accesses are not forwarded to AGP but rather they are mapped to primary PCI, unless they are mapped to AGP via I/O and memory range registers defined above (IOBASE, IOLIMIT, MBASE, MLIMIT, PMBASE, PMLIMIT). (default)</p> <p>The following table shows the behavior for all combinations of MDA and VGA:</p> <table border="1"> <thead> <tr> <th>VGA</th> <th>MDA</th> <th>Behavior</th> </tr> </thead> <tbody> <tr> <td>0</td> <td>0</td> <td>All references to MDA and VGA go to the hub interface.</td> </tr> <tr> <td>0</td> <td>1</td> <td>Illegal combination (DO NOT USE)</td> </tr> <tr> <td>1</td> <td>0</td> <td>All references to VGA go to AGP. MDA-only references (I/O Address 3BFh and aliases) will go to the hub interface.</td> </tr> <tr> <td>1</td> <td>1</td> <td>VGA references go to AGP; MDA references go to the hub link.</td> </tr> </tbody> </table>	VGA	MDA	Behavior	0	0	All references to MDA and VGA go to the hub interface.	0	1	Illegal combination (DO NOT USE)	1	0	All references to VGA go to AGP. MDA-only references (I/O Address 3BFh and aliases) will go to the hub interface.	1	1	VGA references go to AGP; MDA references go to the hub link.
VGA	MDA	Behavior														
0	0	All references to MDA and VGA go to the hub interface.														
0	1	Illegal combination (DO NOT USE)														
1	0	All references to VGA go to AGP. MDA-only references (I/O Address 3BFh and aliases) will go to the hub interface.														
1	1	VGA references go to AGP; MDA references go to the hub link.														

Bit	Descriptions
2	<p><b>ISA Enable.</b> Modifies the response by the MCH to an I/O access issued by the host that targets ISA I/O addresses. This applies only to I/O addresses that are enabled by the IOBASE and IOLIMIT registers.</p> <p>1 = MCH does not forward to AGP any I/O transactions addressing the last 768 bytes in each 1 KB block, even if the addresses are within the range defined by the IOBASE and IOLIMIT registers. Instead of going to AGP these cycles are forwarded to PCI0 where they can be subtractively or positively claimed by the ISA bridge.</p> <p>0 = All addresses defined by the IOBASE and IOLIMIT for host I/O transactions are mapped to AGP. (default)</p>
1	<p><b>SERR# Enable. (Not Implemented).</b> Hardwired to 0. This bit normally controls forwarding SERR# on the secondary interface to the primary interface. The MCH does not support the SERR# signal on the AGP bus.</p>
0	<p><b>Parity Error Response Enable.</b> This bit controls MCH's response to data phase parity errors on AGP. G_PERR# is not implemented by the MCH.</p> <p>1 = Address and data parity errors on AGP are reported via SERR# mechanism, if enabled by SERRE1 and SERRE.</p> <p>0 = Address and data parity errors on AGP are not reported via the MCH SERR# signal. Other types of error conditions can still be signaled via SERR# independent of this bit's state.</p>

### 3.5.22. ERRCMD1—Error Command Register (Device 1)

Address Offset:	40h
Default:	00h
Access:	Read/Write
Size	8 bits

Bit	Descriptions
7:1	Reserved.
0	<p><b>SERR on Receiving Target Abort on AGP.</b></p> <p>1 = Enable. The MCH generates an SERR special cycle via the hub interface when an MCH originated AGP cycle is terminated with a Target Abort.</p> <p>0 = Disable. Reporting of this condition is disabled.</p>

## 4. System Address Map

---

The 82820 MCH supports 4 GB of addressable memory space and 64 KB+3 of addressable I/O space. Note the supported DRAM technology has limited the total physical memory size to 1 GB. There is a programmable memory address space under the 1 MB region which is divided into regions which can be individually controlled with programmable attributes such as Disable, Read/Write, Write Only, or Read Only. Attribute programming is described in Chapter 3.

Section 4.1, “Memory Address Ranges” focuses on how the memory space is partitioned and the use of these separate memory regions. The I/O address space requirements are described in Section 4.2, “I/O Address Space”.

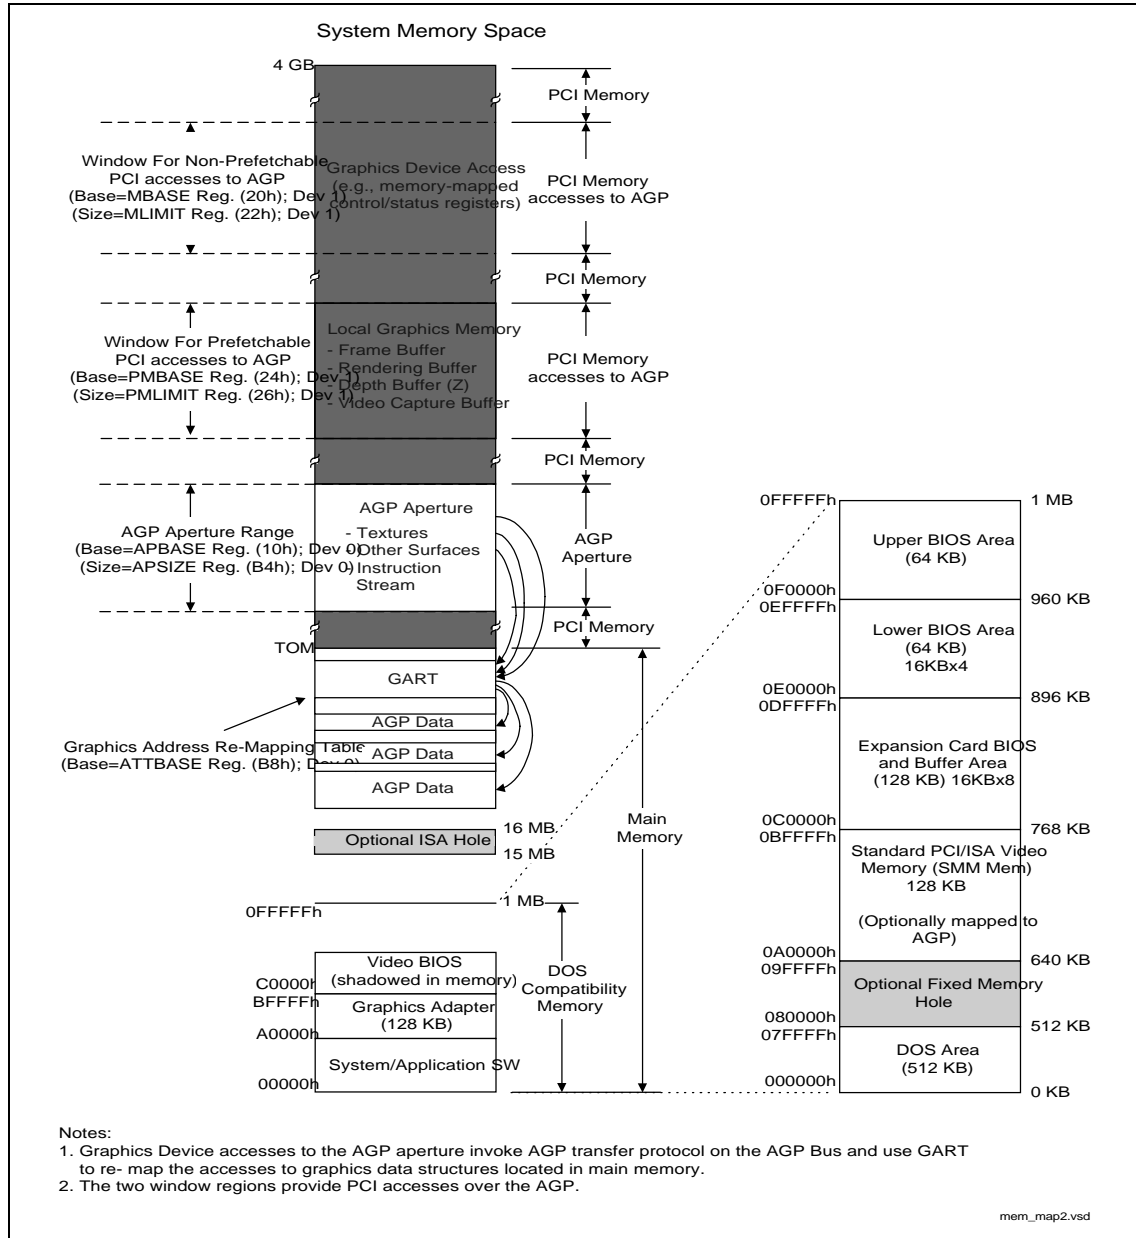
### 4.1. Memory Address Ranges

The MCH claims any host access over 4 GB by terminating the transaction (without forwarding it to the hub interface or AGP). Writes are terminated by dropping the data; for reads, the MCH returns all zeros on the host bus. Note that the Intel® 820 chipset platform does not support the PCI Dual Address Cycle Mechanism (DAC); therefore, the MCH does not allow addressing greater than 4 GB on either the hub interface or AGP interface.

In the following sections, it is assumed that all of the compatibility memory ranges reside on the hub interface. The exceptions to this rule are the VGA ranges which may be mapped to AGP.

Figure 3 provides additional details on mapping specific memory regions as defined and supported by the MCH. The MCH provides a maximum DRAM address decode space of 4 GB. The MCH does not remap APIC memory space. The supported DRAM technologies have limited the total size of memory to 1 GB (256 MB with 64Mbit, 512 MB with 128Mbit, 1 GB with 256Mbit). It is the BIOS or system designer’s responsibility to limit DRAM population so that adequate PCI, AGP, High BIOS, and APIC memory space can be allocated.

Figure 3. Detailed Memory System Address Map



### 4.1.1. Compatibility Area

This area is divided into the following address regions:

- 0–640 KB DOS Area
- 640–768 KB Video Buffer Area
- 768–896 KB in 16 KB sections (total of 8 sections); Expansion Area
- 896–960 KB in 16KB sections (total of 4 sections); Extended System BIOS Area
- 960 KB–1 MB Memory (BIOS Area); System BIOS Area

There are sixteen memory segments in the compatibility area. Thirteen of the memory ranges can be enabled or disabled independently for both read and write cycles.

**Table 8. Memory Segments and Attributes**

Memory Segments	Attributes	Comments
000000h–09FFFFh	Fixed; always mapped to main DRAM	0 to 640K; DOS Region
0A0000h–0BFFFFh	mapped to the hub interface or AGP; configurable as SMM space	Video Buffer (physical DRAM configurable as SMM space)
0C0000h–0C3FFFh	WE RE	Add-on BIOS
0C4000h–0C7FFFh	WE RE	Add-on BIOS
0C8000h–0CBFFFh	WE RE	Add-on BIOS
0CC000h–0CFFFFh	WE RE	Add-on BIOS
0D0000h–0D3FFFh	WE RE	Add-on BIOS
0D4000h–0D7FFFh	WE RE	Add-on BIOS
0D8000h–0DBFFFh	WE RE	Add-on BIOS
0DC000h–0DFFFFh	WE RE	Add-on BIOS
0E0000h–0E3FFFh	WE RE	BIOS Extension
0E4000h–0E7FFFh	WE RE	BIOS Extension
0E8000h–0EBFFFh	WE RE	BIOS Extension
0EC000h–0EFFFFh	WE RE	BIOS Extension
0F0000h–0FFFFFFh	WE RE	BIOS Area

### DOS Area (00000h–9FFFFh)

This 640 KB DOS area is always mapped to the main memory controlled by the MCH.

### Video Buffer Area (A0000h–BFFFFh)

The 128 KB graphics adapter memory region is normally mapped to a legacy video device on PCI0, behind the hub interface (typically, a VGA controller). This area is not controlled by attribute bits and host initiated cycles in this region are forwarded to the hub interface or AGP for termination. This region is also the default region for SMM space.

Accesses to this range are directed to either the hub interface or AGP based on the configuration specified in the PCICMD1 (PCI-PCI Command) and BCTRL (PCI-PCI Bridge Control) registers in MCH Device 1 configuration space. The control is applied for accesses initiated from any of the system interfaces (i.e., host bus, the hub interface, and AGP). Note that for hub interface to AGP accesses, only memory write operations are supported. For more details, see the descriptions in the configuration registers specified above.

The SMRAM Control register controls how SMM accesses to this space are treated.

### **Compatible SMRAM Address Range (A0000h–BFFFFh)**

When compatible SMM space is enabled, SMM-mode processor accesses to this range are routed to physical DRAM at this address. Non-SMM-mode processor accesses to this range are considered to be to the Video Buffer Area as described above. Originated cycles from AGP or the hub interface to enabled SMM space are not allowed and are considered to be to the Video Buffer Area.

### **Monochrome Adapter (MDA) Range (B0000h–B7FFFh)**

Legacy support requires the ability to have a second graphics controller (monochrome) in the system. In an AGP system, accesses in the standard VGA range are forwarded to the AGP bus (depending on configuration bits). Since the monochrome adapter may be on the PCI0 (or ISA) bus behind the hub interface, the MCH must decode cycles in the MDA range and forward them to the hub interface. This capability is controlled by a configuration bit (MDAP bit). In addition to the memory range B0000h–B7FFFh, the MCH decodes IO cycles at 3B4h, 3B5h, 3B8h, 3B9h, 3BAh and 3BFh and forwards them to the hub interface.

### **Expansion Area (C0000h–DFFFFh)**

This 128 KB ISA Expansion region is divided into eight 16 KByte segments. Each segment can be assigned one of four Read/Write states: read-only, write-only, read/write, or disabled. Typically, these blocks are mapped through MCH and are subtractively decoded to ISA space. Memory that is disabled is not remapped.

### **Extended System BIOS Area (E0000h–EFFFFh)**

This 64 KB area is divided into four 16 KByte segments. Each segment can be assigned independent read and write attributes so it can be mapped either to main DRAM or to the hub interface. Typically, this area is used for RAM or ROM. Memory segments that are disabled are not remapped elsewhere.

### **System BIOS Area (F0000h–FFFFFh)**

This area is a single 64 KB segment. This segment can be assigned read and write attributes. The default (after reset) is Read/Write disabled and cycles are forwarded to the hub interface. By manipulating the Read/Write attributes, the MCH can “shadow” BIOS into the main DRAM. When disabled, this segment is not remapped.

### 4.1.2. Extended Memory Area

This memory area covers 100000h (1 MB) to FFFFFFFFh (4 GB -1) address range and it is divided into the following regions:

- Main DRAM Memory from 1 MB to the Top of Memory (maximum of 4 GBs).
- AGP or PCI Memory space from the Top of Memory to 4 GB with two specific ranges:
  - APIC Configuration Space from FEC0\_0000h (4 GB-20 MB) to FECF\_FFFFh and FEE0\_0000h to FEEF\_FFFFh
  - High BIOS area from 4 GB to 4 GB minus 2 MB

#### Main DRAM Address Range (0010\_0000h to Top of Main Memory)

The address range from 1 MB to the top of main memory is mapped to main DRAM address range controlled by the MCH. The Top of Memory (TOM) is limited to 4 GB. All accesses to addresses within this range will be forwarded by the MCH to the DRAM unless a hole in this range is created using the fixed hole as controlled by the FDHC register. Accesses within this hole are forwarded to the hub interface. The range of physical DRAM memory disabled by opening the hole is not remapped to the Top of the Memory.

The MCH provides a maximum DRAM address decode space of 4 GB. The MCH does not remap APIC memory space. The supported DRAM technology have limited the total size of memory to 1 GB(256 MB with 64Mbit, 512 MB with 128Mbit, 1 GB with 256Mbit). It is the BIOS or system designers responsibility to limit DRAM population so that adequate PCI, AGP, High BIOS, and APIC memory space can be allocated.

#### Extended SMRAM Address Range (HSEG and Top of Main Memory - TSEG)

The HSEG and TSEG SMM transaction address spaces reside in this extended memory area.

- **HSEG:** SMM-mode processor accesses to enabled HSEG are remapped to 000A0000h–000BFFFFh. Non-SMM-mode processor accesses to enabled HSEG are considered invalid and the transaction is terminated at the bus interface (i.e., the transaction will *not* be dispatched to DRAM). The exceptions are Non-SMM-mode Write Back cycles that are remapped to SMM space to maintain cache coherency. Originated cycle from AGP or the hub interface to enable SMM space is not allowed. Physical DRAM behind the HSEG transaction address is not remapped and is not accessible.
- **TSEG:** TSEG can be up to 1 MB in size and is at the top of memory. SMM mode processor accesses to enabled TSEG access the physical DRAM at the same address. Non-SMM mode processor accesses to enabled TSEG are considered invalid and the transaction is terminated at the bus interface (i.e., the transaction will *not* be dispatched to DRAM). The exception is Non-SMM mode Write Back cycles that are directed to the physical SMM space to maintain cache coherency. Originated cycle from AGP or the hub interface to enable SMM space is not allowed.

## PCI Memory Address Range (Top of Main Memory to 4 GB)

The address range from the top of main DRAM to 4 GB (top of physical memory space supported by the Intel® 820/820E chipset) is normally mapped to the hub interface. There are two exceptions:

- Addresses decoded to the AGP Memory Window defined by the MBASE, MLIMIT, PMBASE, and PMLIMIT registers are mapped to AGP.
- Addresses decoded to the Graphics Aperture range defined by the APBASE and APSIZE registers are mapped to the main DRAM.

There are two sub-ranges within the hub interface (Memory address range defined as APIC Configuration Space and High BIOS Address Range). *The AGP Memory Window and Graphics Aperture Window MUST NOT overlap with these two ranges.* These ranges are described in the following paragraphs.

### APIC Configuration Space (FEC0\_0000h – FECF\_FFFFh, FEE0\_0000h – FEEF\_FFFFh)

This range is reserved for APIC configuration space which includes the default I/O APIC configuration space. The default Local APIC configuration space is FEE0\_0000h to FEEF\_0FFFh.

Host accesses to the Local APIC configuration space do not result in external bus activity since the Local APIC configuration space is internal to the host. However, a MTRR must be programmed to make the Local APIC range uncacheable (UC). The Local APIC base address in each host should be relocated to the FEC0\_0000h (4 GB - 20 MB) to FECF\_FFFFh range so that one MTRR can be programmed to 64 KB for the Local and I/O APICs. The I/O APIC(s) usually reside in the I/O Controller Hub portion of the chipset or as a standalone component(s).

I/O APIC units are located beginning at the default address FEC0\_0000h. The first I/O APIC is located at FEC0\_0000h. Each I/O APIC unit is located at FEC0\_x000h where *x* is I/O APIC unit number 0 through F(hex). This address range is normally mapped to the hub interface.

The address range between the APIC configuration space and the High BIOS (FED0\_0000h to FFDF\_FFFFh) is always mapped to the hub interface.

### High BIOS Area (FFE0\_0000h – FFFF\_FFFFh)

The top 2 MB of the Extended Memory Region is reserved for System BIOS (High BIOS), extended BIOS for PCI devices, and the A20 alias of the system BIOS. The processor begins execution from the High BIOS after reset. This region is mapped to the hub interface so that the upper subset of this region aliases to 16 MB - 256 KB range. The actual address space required for the BIOS is less than 2 MB but the minimum processor MTRR range for this region is 2 MB so that full 2 MB must be considered.



### 4.1.3. AGP Memory Address Ranges

The MCH can be programmed to direct memory accesses to the AGP bus interface when addresses are within either of two ranges specified via registers in MCH Device 1 configuration space. The first range is controlled via the Memory Base Register (MBASE) and Memory Limit Register (MLIMIT) registers. The second range is controlled via the Prefetchable Memory Base (PMBASE) and Prefetchable Memory Limit (PMLIMIT) registers

Conceptually, address decoding for each range follows the same basic concept. The top 12 bits of the respective Memory Base and Memory Limit registers correspond to address bits A[31:20] of a memory address. For the purpose of address decoding, the MCH assumes that address bits A[19:0] of the memory base are zero and that address bits A[19:0] of the memory limit address are FFFFh. This forces each memory address range to be aligned to 1 MB boundary and to have a size granularity of 1 MB.

The MCH positively decodes memory accesses to AGP memory address space as defined by the following equations:

$$\text{Memory\_Base\_Address} \leq \text{Address} \leq \text{Memory\_Limit\_Address}$$

$$\text{Prefetchable\_Memory\_Base\_Address} \leq \text{Address} \leq \text{Prefetchable\_Memory\_Limit\_Address}$$

The effective size of the range is programmed by the plug-and-play configuration software and depends on the size of memory claimed by the AGP device. Normally these ranges reside above the Top-of-Main-DRAM and below High BIOS and APIC address ranges.

It is essential to support a separate Prefetchable range to apply USWC attribute (from the processor point of view) to that range.

**Note:** MCH Device 1 memory range registers described above are used to allocate memory address space for any devices sitting on AGP bus that require such a window.

### 4.1.4. AGP DRAM Graphics Aperture

Memory-mapped graphics data structures can reside in a *Graphics Aperture* to main DRAM memory. This aperture is an address range defined by the APBASE configuration register of the MCH Device 0. The APBASE register follows the standard base address register template as defined by the PCI 2.1 specification. The size of the range claimed by the APBASE is programmed via “back-end” register APSIZE (programmed by the chipset specific BIOS before plug-and-play session is performed). APSIZE allows selection of the aperture size of 4 MB, 8 MB, 16 MB, 32 MB, 64 MB, 128 MB and 256 MB. By programming APSIZE to a specific size, the corresponding lower bits of APBASE are forced to 0 (behave as hardwired). The default value of APSIZE forces an aperture size of 256 MB. The aperture address range is naturally aligned.

Accesses within the aperture range are forwarded to the main DRAM subsystem. The MCH translates the originally issued addresses via a translation table maintained in main memory. The aperture range should be programmed as non-cacheable in the processor caches.

**Note:** Plug-and-play software configuration model does not allow overlap of different address ranges. Therefore, the AGP Graphics Aperture and AGP Memory Address Range are independent address ranges that may be adjacent but cannot overlap one another.

### 4.1.5. System Management Mode (SMM) Memory Range

The MCH supports the use of main memory as System Management RAM (SMRAM) enabling the use of System Management Mode. The MCH supports two SMRAM options:

- Compatible SMRAM (C\_SMRAM)
- Extended SMRAM (E\_SMRAM).

System Management RAM (SMRAM) space provides a memory area that is available for the SMI handler's and code and data storage. This memory resource is normally hidden from the system operating system so the processor has immediate access to this memory space upon entry to SMM. MCH provides three SMRAM options:

- Below 1 MB option that supports compatible SMI handlers.
- Above 1 MB option that allows new SMI handlers to execute with write-back cacheable SMRAM.
- Optional larger write-back cacheable T\_SEG area from 128KB to 1 MB in size above 1 MB that is reserved from the highest area in system DRAM memory. The above 1 MB solutions require changes to compatible SMRAM handlers code to properly execute above 1 MB.

**Note:** Masters from the hub interface and AGP are not allowed to access the SMM space.

#### SMM Space Definition

SMM space is defined by its **addressed** SMM space and its **DRAM** SMM space. The addressed SMM space is defined as the range of bus addresses used by the processor to access SMM space. DRAM SMM space is defined as the range of physical DRAM memory locations containing the SMM code. SMM space can be accessed at one of three transaction address ranges: Compatible, High and TSEG. The Compatible and TSEG SMM space is not remapped; therefore, the addressed and DRAM SMM space is the same address range. Since the High SMM space is remapped, the addressed and DRAM SMM space is a different address range. Note that the High DRAM space is the same as the Compatible Transaction Address space; therefore, the table below describes three unique address ranges:

- Compatible Transaction Address (Adr C)
- High Transaction Address (Adr H)
- TSEG Transaction Address (Adr T)

These abbreviations are used later in the table describing SMM Space Transaction Handling.

SMM Space Enabled	Transaction Address Space (Adr)	DRAM Space (DRAM)
Compatible (C)	A0000h–BFFFFh	A0000h–BFFFFh
High (H)	0FEEA0000h–0FEEBFFFFh	A0000h–BFFFFh
TSEG (T)	(TOM - TSEG_SZ) to TOM	(TOM - TSEG_SZ) to TOM

#### NOTES:

1. High SMM: Note that this is different than in previous chip sets. In previous chipsets the High segment was the 384 KB region from A0000h to FFFFFh. This has been removed in MCH.
2. TSEG SMM: Note that this is different than in previous chip sets. In previous chip sets the TSEG address space was offset by 256 MB to allow for simpler decoding and the TSEG was remapped to just under the TOM. In the MCH the TSEG region is not offset by 256 MB and it is not remapped.

## SMM Space Restrictions

If any of the following conditions are violated, the results of SMM accesses are unpredictable and may cause the system to hang up:

- The Compatible SMM space **must not** be set up as cacheable.
- High or TSEG SMM transaction address space **must not** overlap address space assigned to system DRAM, the AGP aperture range, or to any “PCI” devices (including hub interface and AGP devices). This is a BIOS responsibility.
- Both D\_OPEN and D\_CLOSE **must not** be set to 1 at the same time.
- When TSEG SMM space is enabled, the TSEG space **must not** be reported to the operating system as available DRAM. This is a BIOS responsibility.
- Any address translated through the AGP Aperture GTLB **must not** target DRAM from 000A0000h to 000FFFFFh.

## SMM Space Combinations

When High SMM is enabled (G\_SMRAME=1 and H\_SMRAM\_EN=1), the Compatible SMM space is effectively disabled. Processor-originated accesses to the Compatible SMM space are forwarded to AGP if VGAEN=1 (also depends on MDAP); otherwise, they are forwarded to hub interface #1. AGP and hub interface #1/#2 originated accesses are **never** allowed to access SMM space.

**Table 9. SMM Space Table**

Global Enable G_SMRAME	High Enable H_SMRAM_EN	TSEG Enable TSEG_EN	Compatible (C) Range	High (H) Range	TSEG (T) Range
0	x	x	Disable	Disable	Disable
1	0	0	Enable	Disable	Disable
1	0	1	Enable	Disable	Enable
1	1	0	Disabled	Enable	Disable
1	1	1	Disabled	Enable	Enable

## SMM Control Combinations

The G\_SMFRAME bit provides a global enable for all SMM memory. The D\_OPEN bit allows software to write to the SMM ranges without being in SMM mode. BIOS can use this bit to initialize SMM code at powerup. The D\_LCK bit limits the SMM range access to only SMM mode accesses. The D\_CLS bit causes SMM data accesses to be forwarded to hub interface #1 or AGP. The SMM software can use this bit to write to video memory while running SMM code out of DRAM.

**Table 10. SMM Control Table**

G_SMFRAME	D_LCK	D_CLS	D_OPEN	Processor in SMM Mode	SMM Code Access	SMM Data Access
0	x	x	x	x	Disable	Disable
1	0	x	0	0	Disable	Disable
1	0	0	0	1	Enable	Enable
1	0	0	1	x	Enable	Enable
1	0	1	0	1	Enable	Disable
1	0	1	1	x	Invalid	Invalid
1	1	x	x	0	Disable	Disable
1	1	0	x	1	Enable	Enable
1	1	1	x	1	Enable	Disable

### 4.1.6. Memory Shadowing

Any block of memory that can be designated as read-only or write-only can be “shadowed” into MCH DRAM memory. Typically, this is done to allow ROM code to execute more rapidly out of main DRAM. ROM is used as a read-only during the copy process while DRAM at the same time is designated write-only. After copying, the DRAM is designated read-only so that ROM is shadowed. Host bus transactions are routed accordingly.

## 4.2. I/O Address Space

The MCH does not support the existence of any other I/O devices beside itself on the host bus. The MCH generates either the hub interface or AGP bus cycles for all host I/O accesses. The MCH contains two internal registers in the host I/O space:

- Configuration Address Register (CONF\_ADDR)
- Configuration Data Register (CONF\_DATA).

These locations are used to implement the PCI configuration space access mechanism as described in the Configuration Register Chapter.

The processor allows 64K+3 bytes to be addressed within the I/O space. The MCH propagates the host I/O address without any translation on to the destination bus and therefore, provides addressability for 64K+3 byte locations. Note that the upper 3 locations can be accessed only during I/O address wrap-around, when host bus A16# address signal is asserted. A16# is asserted on the host bus when an I/O

access is made to 4 bytes from address 0FFFDh, 0FFFEh, or 0FFFFh. A16# is also asserted when an I/O access is made to 2 bytes from address 0FFFFh.

The I/O accesses (other than ones used for configuration space access) are forwarded normally to the hub interface bus, unless they fall within the AGP I/O address range as defined by the mechanisms explained below. The MCH does not post I/O write cycles to IDE. The MCH never responds to I/O cycles initiated on AGP.

### AGP I/O Address Mapping

The MCH can be programmed to direct I/O accesses to the AGP bus when host-initiated I/O cycle addresses are within the AGP I/O address range. This range is controlled via the I/O Base Address (IOBASE) and I/O Limit Address (IOLIMIT) registers in MCH Device 1 configuration space.

Address decoding for this range is based on the following concept. The top 4 bits of the respective I/O Base and I/O Limit registers correspond to address bits A[15:12] of an I/O address. For the purpose of address decoding, the MCH assumes that lower 12 address bits A[11:0] of the I/O base are zero and that address bits A[11:0] of the I/O limit address are FFFh. This forces the I/O address range alignment to a 4 KB boundary and produces a size granularity of 4 KB.

The MCH positively decodes I/O accesses to AGP I/O address space as defined by the following equation:

$$\text{I/O\_Base\_Address} \leq \text{host I/O Cycle Address} \leq \text{I/O\_Limit\_Address}$$

The effective size of the range is programmed by the plug-and-play configuration software and depends on the size of I/O space claimed by the AGP device.

The MCH also forwards accesses to the Legacy VGA I/O ranges as defined and via Device 1 configuration registers BCTRL and PCICMD1, unless a second adapter (monochrome) is present on PCI0 (or ISA). The presence of a second graphics adapter is determined by the MDAP configuration bit. When MDAP is set, the MCH decodes legacy monochrome I/O ranges and forwards them to the hub interface. The I/O ranges decoded for the monochrome adapter are 3B4h, 3B5h, 3B8h, 3B9h, 3BAh and 3BFh.

**Note:** The MCH Device 1 I/O address range registers defined above are used for all I/O space allocation for any devices requiring such a window on AGP.

## 4.3. MCH Decode Rules and Cross-Bridge Address Mapping

The address map described above applies globally to accesses arriving on any of the three interfaces (i.e., Host bus, the hub interface, or AGP).

### 4.3.1. The Hub Interface Decode Rules

The MCH accepts accesses from the hub interface to the following address ranges:

- All memory read and write accesses to Main DRAM (except SMM space).
- All memory write accesses from the hub interface to AGP memory range defined by MBASE, MLIMIT, PMBASE, and PMLIMIT.
- All memory read/write accesses to the Graphics Aperture defined by APBASE and APSIZE.
- Memory writes to VGA range on AGP if enabled.

Memory accesses from the hub interface that fall elsewhere within the memory range and I/O cycles will not be accepted. They are terminated with Master Abort completion.

### 4.3.2. AGP Interface Decode Rules

#### Cycles Initiated Using AGP FRAME# Protocol

The MCH does not support any AGP FRAME# access targeting the hub interface. The MCH claims AGP-initiated memory read and write transactions decoded to the main DRAM range or the Graphics Aperture range. All other memory read and write requests are master-aborted by the AGP initiator as a consequence of MCH not responding to a transaction.

Under certain conditions, the MCH restricts access to the DOS Compatibility ranges governed by the PAM registers by distinguishing access type and destination bus. The MCH accepts AGP FRAME# write transactions to the compatibility ranges, if the PAM designates DRAM as writable. If accesses to a range are not write enabled by the PAM, the MCH does not respond and the cycle results in a master-abort. The MCH accepts AGP FRAME# read transactions to the compatibility ranges, if the PAM designates DRAM as readable. If accesses to a range are not read enabled by the PAM, the MCH does not respond and the cycle results in a master-abort.

If an agent on AGP issues an I/O, PCI Configuration or PCI Special Cycle transaction, the MCH does not respond and the cycle results in a master-abort.

## Cycles Initiated Using AGP PIPE# or SB Protocol

All cycles must reference main memory. That is they must reference the main DRAM address range (excluding PAM) or the Graphics Aperture range (also physically mapped within DRAM but using different address range). AGP accesses to the PAM region from 640 KB – 1 MB are not allowed. AGP accesses to SMM space are not allowed. AGP-initiated cycles that target DRAM are not snooped on the host bus, even if they fall outside of the AGP aperture range.

If a cycle is outside of a valid main memory range, it will terminate as follows:

- Reads: Remap to memory address 0h, return data from address 0h, and set the IAAF error bit in ERRSTS register in the Device 0.
- Writes: Remapped to memory address 0h with BE's deasserted (effectively dropped "on the floor") and set the IAAF error flag in the ERRSTS register in the Device 0.

## AGP Accesses to MCH that Cross Device Boundaries

For AGP FRAME# accesses, when an AGP master gets disconnected, it resumes at the new address that allows the cycle to be routed to or claimed by the new target. Therefore, accesses should be disconnected by the target on potential device boundaries. The MCH disconnects AGP FRAME# transactions on 4 KB boundaries.

AGP PIPE# and SBA accesses are limited to 256 bytes and must hit DRAM. Read accesses crossing a device boundary return invalid data when the access crosses out of DRAM. Write accesses crossing out of DRAM are discarded. The IAAF Error bit is set.

### 4.3.3. Legacy VGA Ranges

The legacy VGA memory range A0000h–BFFFFh is mapped either to the hub interface or to AGP depending on the programming of the VGA Enable bit in the BCTRL configuration register in MCH Device 1 configuration space and the MDAP bit in the MCHCFG configuration register in Device 0 configuration space. The same registers control mapping of VGA I/O address ranges. VGA I/O range is defined as addresses where A[9:0] are in the ranges 3B0h to 3BBh and 3C0h to 3DFh (inclusive of ISA address aliases; A[15:10] are not decoded). The function and interaction of these two bits is described below:

**MDA Present (MDAP):** This bit works with the VGA Enable bit in the BCTRL register of device 1 to control the routing of host-initiated transactions targeting MDA compatible I/O and memory address ranges. This bit should not be set when the VGA Enable bit is not set. If the VGA enable bit is set, accesses to IO address range x3BCh–x3BFh are forwarded to the hub interface. If the VGA enable bit is not set, accesses to IO address range x3BCh–x3BFh are treated just like any other IO accesses (i.e., the cycles are forwarded to AGP, if the address is within IOBASE and IOLIMIT, and ISA enable bit is not set; otherwise, they are forwarded to the hub interface). MDA resources are defined as the following:

- Memory: 0B0000h–0B7FFFh
- I/O: 3B4h, 3B5h, 3B8h, 3B9h, 3BAh, 3BFh,  
(including ISA address aliases, A[15:10] are not used in decode)

Any I/O reference that includes the I/O locations listed above, or their aliases, are forwarded to the hub interface, even if the reference includes I/O locations not listed above.

**VGA Enable:** Controls the routing of host-initiated transactions targeting VGA compatible I/O and memory address ranges. When this bit is set, the MCH forwards the following host accesses to AGP:

- Memory accesses in the range 0A0000h to 0BFFFFh
- I/O addresses where A[9:0] are in the ranges 3B0h to 3BBh and 3C0h to 3DFh (inclusive of ISA address aliases - A[15:10] are not decoded)

When this bit is set, forwarding of these accesses issued by the host is independent of the I/O address and memory address ranges defined by the previously defined base and limit registers. Forwarding of these accesses is also independent of the settings of the bit 2 (ISA Enable) of this register, if this bit is 1. If the VGA enable bit is set, accesses to I/O address range x3BCh–x3BFh are forwarded to the hub interface. If the VGA enable bit is not set, accesses to I/O address range x3BCh–x3BFh are treated just like any other I/O access (i.e. the cycles are forwarded to AGP if the address is within IOBASE and IOLIMIT, and ISA enable bit is not set; otherwise, they are forwarded to the hub interface).

If this bit is 0 (default), VGA compatible memory and I/O range accesses are not forwarded to AGP; rather, they are mapped to the primary the hub interface unless they are mapped to AGP via I/O and memory range registers defined above (IOBASE, IOLIMIT, MBASE, MLIMIT, PMBASE, PMLIMIT)

The following table shows the behavior for all combinations of MDA and VGA.

VGA	MDA	Behavior
0	0	All references to MDA and VGA go to the hub interface.
0	1	Illegal Combination (DO NOT USE)
1	0	All references to VGA go to AGP. MDA-only references (I/O address 3BFh and aliases) go to the hub interface.
1	1	VGA references go to AGP; MDA references go to the hub interface.



## 5. Functional Description

This chapter covers the 82820 MCH functional units including, host interface, AGP interface, power management, clocking, and system reset and power sequencing.

### 5.1. Host Interface

The MCH is optimized for the Intel® Pentium® II and Intel® Pentium® III processors. The MCH supports one or two processor configuration for slot 1. The MCH supports Processor System Bus frequencies of 100 MHz and 133 MHz using AGTL+ signaling. The AGTL+ buffers support single-ended termination for uniprocessor systems. The MCH supports 32-bit host addressing, decoding up to 4 GB of memory address space for the processor. Since the MCH only supports a single RDRAM channel, the maximum physical memory size that can be supported is 1 GB with 256Mbit technology. Host memory writes to address space above 4 GB are immediately terminated and discarded. Host memory reads to address space above 4 GB are immediately terminated and return the value of the pulled-up AGTL+ host bus. The MCH has a 6 deep In-Order Queue to support up to six outstanding pipelined address requests on the host bus. Host-initiated I/O cycles are decoded to AGP, the hub interface, or MCH configuration space. Host-initiated memory cycles are decoded to AGP, the hub interface, or DRAM. Host memory accesses in the AGP aperture are translated using the AGP address translation table. All accesses from AGP and the hub interface that are in the graphics aperture are also translated using the AGP address translation table. AGP PIPE# or SBA[7:0] memory accesses initiated from AGP to DRAM are not snooped on the host bus. Memory accesses initiated from AGP using AGP FRAME# and from the hub interface to DRAM are snooped on the host bus.

The MCH recognizes and supports a large subset of the transaction types that are defined for the Pentium® II processor and Pentium® III processor bus interfaces. However, each of these transaction types has a multitude of response types, some of which are not supported by this controller. All transactions are processed in the order that they are received on the host bus. A summary of transactions supported by the MCH is given in the following table.

**Table 11. Types of Transactions Supported by the MCH**

Transaction	REQa[4:0]#	REQb[4:0]#	MCH Support
Deferred Reply	0 0 0 0 0	X X X X X	The MCH will initiate a deferred reply for a previously deferred transaction.
Reserved	0 0 0 0 1	X X X X X	Reserved
Interrupt Acknowledge	0 1 0 0 0	0 0 0 0 0	Interrupt acknowledge cycles are forwarded to the hub interface. A single byte of data is returned on HD[7:0]#.
Special Transactions	0 1 0 0 0	0 0 0 0 1	See separate table in Special Cycles section.
Reserved	0 1 0 0 0	0 0 0 1 x	Reserved
Reserved	0 1 0 0 0	0 0 1 x x	Reserved
Branch Trace Message	0 1 0 0 1	0 0 0 0 0	The MCH terminates a branch trace message without latching data.
Reserved	0 1 0 0 1	0 0 0 0 1	Reserved

Transaction	REQa[4:0]#	REQb[4:0]#	MCH Support
Reserved	0 1 0 0 1	0 0 0 1 x	Reserved
Reserved	0 1 0 0 1	0 0 1 x x	Reserved
I/O Read	1 0 0 0 0	0 0 x LEN#	I/O read cycles are forwarded to the hub interface or AGP. I/O cycles to the MCH configuration space are not forwarded to AGP or the hub interface.
I/O Write	1 0 0 0 1	0 0 x LEN#	I/O write cycles are forwarded to the hub interface or AGP. I/O cycles to the MCH configuration space are not forwarded to AGP or the hub interface.
Reserved	1 1 0 0 x	0 0 x x x	Reserved
Memory Read and Invalidate	0 0 0 1 0	0 0 x LEN#	Host initiated memory read and invalidate cycles are forwarded to DRAM. The MCH initiates an MRI (LEN=0) cycle to snoop a hub interface or AGP initiated write cycle to DRAM.
Reserved	0 0 0 1 1	0 0 x LEN#	Reserved
Memory Code Read	0 0 1 0 0	0 0 x LEN#	Memory code read cycles are forwarded to DRAM, hub interface, or AGP.
Memory Data Read	0 0 1 1 0	0 0 x LEN#	Host-initiated memory read cycles are forwarded to DRAM, the hub interface, or AGP. The MCH initiates a memory read cycle to snoop a hub interface or AGP initiated read cycle to DRAM.
Memory Write (no retry)	0 0 1 0 1	0 0 x LEN#	This memory write is a writeback cycle and cannot be retried. The MCH forwards the write to DRAM.
Memory Write (can be retried)	0 0 1 1 1	0 0 x LEN#	The standard memory write cycle is forwarded to DRAM, the hub interface, or AGP.

**NOTES:**

- For Memory cycles, REQa[4:3]# = ASZ#. The MCH only supports ASZ# = 00 (32 bit address).
- REQb[4:3]# = DSZ#. For the Intel® Pentium® II and Intel® Pentium® III processors, DSZ# = 00 (64 bit data bus size).
- LEN# = data transfer length as follows:

LEN#	Data length
00	≤ 8 bytes (BE[7:0]# specify granularity)
01	Length = 16 bytes BE[7:0]# all active
10	Length = 32 bytes BE[7:0]# all active
11	Reserved

**Table 12. Types of Responses Supported by the MCH**

RS2#	RS1#	RS0#	Description	MCH Support
0	0	0	Idle	
0	0	1	Retry Response	This response is generated if an access is to a resource that cannot be accessed by the processor at this time and the logic must avoid deadlock. The I/O Controller Hub directed reads and writes, DRAM locked reads, AGP reads and writes can be retried.
0	1	0	Deferred Response	This response can be returned for all transactions that can be executed 'out of order.' Host Memory Reads I/O Reads, Interrupt Acknowledge and I/O Writes cycles to the hub interface, and AGP can be deferred.
0	1	1	Reserved	Reserved
1	0	0	Hard Failure	Not supported.
1	0	1	No Data Response	This is for transactions where the data has already been transferred or for transactions where no data is transferred. Writes and zero length reads receive this response.
1	1	0	Implicit Writeback	This response is given for those transactions where the initial transactions snoop hits on a modified cache line.
1	1	1	Normal Data Response	This response is for transactions where data accompanies the response phase. Reads receive this response.

### 5.1.1. Host Addresses Above 4 GB

Host memory writes to the address space above 4 GB are immediately terminated and discarded. Processor memory reads to address space above 4 GB are immediately terminated and will return the value of the pulled-up AGTL+ Processor System Bus.

### 5.1.2. Host Bus Cycles

The following transaction descriptions illustrate the various operations in their most straightforward representation.

#### Partial Reads

Partial Read transactions include I/O reads and memory read operations of less than or equal to eight bytes (four consecutive bytes for I/O) within an aligned 8 byte span. The byte enable signals (BE[7:0]#) select which bytes in the span to read.

#### Part-Line Read and Write Transactions

The MCH does not support part-line (i.e., 16-byte transactions).

## Cache Line Reads

A read of a full cache line (as indicated by the LEN[1:0]=10 during request phase) requires 32 bytes of data to be transferred, which translates into four data transfers for a given request. If selected as a target, the MCH determines if the address is directed to DRAM, the hub interface, or AGP and provides the corresponding command and control to complete the transaction.

## Partial Writes

Partial Write transactions include I/O and memory write operations of eight bytes or less (maximum of four bytes for I/O) within an aligned 8-byte span. The byte enable signals (BE#[7:0]) select which bytes in the span to write. I/O writes crossing a 4-byte boundary are broken into two separate transactions by the host.

## Cache Line Writes

A write of a full cache line requires 32 bytes of data to be transferred, which translates into four data transfers for a given request.

## Memory Read and Invalidate (length > 0)

A Memory Read and Invalidate (MRI) transaction is functionally equivalent to a cache line read. The purpose of having this special transaction is to support write allocation (write miss case) of cache lines in the processors. When a processor issues a MRI, the cache line is read as in a normal cache line read operation; however, all other caching agents must invalidate this line if they have it in a shared or exclusive state. If a caching agent has this line in the Modified State, then it must be written back to memory and invalidated. The MCH snarfs the write-back data.

## Memory Read and Invalidate (length = 0)

A Memory Read and Invalidate transaction of length zero, MRI(0) does not have an associated Data Response. Executing the transaction informs other agents in the system that the agent issuing this request wants exclusive ownership of a cache line that is in the Shared State (write hit to a shared line). Agents with this cache line invalidate the line. If this line is in the modified state, an implicit write-back cycle is generated and the MCH snarfs the data.

The MCH generates length=0 Memory Read and Invalidate transactions for the hub interface and AGP memory write cycles to DRAM.

## Memory Read (length = 0)

A Memory Read of length zero, MR(0) does not have an associated Data Response. This transaction is used by the MCH to snoop for the hub interface to DRAM and AGP FRAME# snoopable DRAM read accesses. The MCH snoop request policy is identical for the hub interface and AGP FRAME# memory read transactions. Note that the MCH does multiple snoop ahead cycles for hub interface burst reads greater than 32 bytes and for AGP FRAME# master burst reads (i.e., memory read multiple) to DRAM. The MCH performs single MR(0) cycles for the hub interface reads that are  $\leq 32$  bytes and for AGP FRAME# master standard reads or read lines directed to DRAM.

The MCH generates length=0 Memory Read cycles for the hub interface and AGP FRAME# memory read cycles to DRAM.

## Cache Coherency Cycles

The MCH generates an implicit writeback response during host bus read and write transactions when a processor asserts HITM# during the snoop phase. The host-initiated write case has two data transfers; the requesting agents data followed by the snooping agents writeback data.

The MCH performs a memory read and invalidate cycle of length 0 (MRI[0]) on the host bus when a hub interface or AGP FRAME# snoopable DRAM write cycle occurs.

The MCH performs a memory read cycle with length = 0 (MR[0]) on the host bus when a hub interface or AGP FRAME# snoopable DRAM read cycle occurs.

## Interrupt Acknowledge Cycles

A processor agent issues an Interrupt Acknowledge cycle in response to an interrupt from an 8259-compatible interrupt controller. The Interrupt Acknowledge cycle is similar to a partial read transaction, except that the address bus does not contain a valid address. Interrupt Acknowledge cycles are always directed to the hub interface (never to AGP).

## Locked Cycles

The MCH supports resource locking due to the assertion of the LOCK# line on the host bus as follows:

- Host<->DRAM Locked Cycles

The MCH supports host to DRAM locked cycles. The host bus may not execute any other transactions until the locked cycle is complete. The MCH arbiter may grant another hub interface or AGP device but any cycles to DRAM will be blocked.

- Host<-> I/O Controller Hub Locked Cycles

Any host to the I/O Controller Hub locked transaction initiates a locked sequence via the hub interface. The Processor bus implements the bus lock mechanism which means that no change of bus ownership can occur from the time one agent has established a locked transaction (i.e., the initial read cycle of a locked transaction has completed) until the locked transaction is completed. Note that for host transactions to the hub interface, a "LOCK" special cycle is issued to establish the lock prior to the initial read and a "UNLOCK" special cycle is issued via the hub interface after the host lock transaction is completed.

Any concurrent cycle that requires snooping on the host bus is not processed while a LOCK transaction is occurring on the host bus.

Locked cycles from the hub interface to DRAM are not supported.

- Host<->AGP Locked Cycles

The AGP interface does not support locked operations; therefore, both host locked and non-locked transactions destined to AGP are propagated in the same manner. Note, however, that any concurrent cycle that requires snooping on the host bus is not processed while a LOCK transaction is occurring on the host bus.

## Branch Trace Cycles

An agent issues a Branch Trace Cycle for taken branches, if execution tracing is enabled. Address Aa[35:3]# is reserved and can be driven to any value. D[63:32]# carry the linear address of the instruction causing the branch and D[31:0]# carry the target linear address. The MCH responds and retires this transaction but does not latch the value on the data lines or provide any additional support for this type of cycle.

## Special Cycles

A Special Cycle is defined when  $REQa[4:0] = 01000$  and  $REQb[4:0] = xx001$ . In the first address phase,  $Aa[35:3]\#$  is undefined and can be driven to any value. In the second address phase,  $Ab[15:8]\#$  defines the type of Special Cycle issued by the processor. All Host-initiated Special Cycles are routed to the hub interface.

Special Cycles are “posted” in the MCH. The Processor System Bus transaction is terminated immediately. It does not wait for the cycle to propagate or terminate on the hub interface. The table below specifies the cycle type and definition as well as the action taken by the MCH when the corresponding cycles are identified.

**Table 13. Types of Special Cycles Supported by the MCH**

BE[7:0]#	Special Cycle Type	Action Taken
0000 0000	NOP	This transaction has no side-effects.
0000 0001	Shutdown	This transaction is issued when an agent detects a severe software error that prevents further processing. This cycle is claimed by the MCH and propagated as a Shutdown special cycle over the hub interface. This cycle is retired on the host bus after the associated special cycle request is successfully broadcast over the hub interface.
0000 0010	Flush	This transaction is issued when an agent has invalidated its internal caches without writing back any modified lines. The MCH claims this cycle and retires it.
0000 0011	Halt	This transaction is issued when an agent executes a HLT instruction and stops program execution. This cycle is claimed by the MCH and propagated over the hub interface as a Halt special cycle. This cycle is retired on the host bus after the associated special cycle request is successfully broadcast over the hub interface.
0000 0100	Sync	This transaction is issued when an agent has written back all modified lines and has invalidated its internal caches. The MCH claims this cycle and retires it.
0000 0101	Flush Acknowledge	This transaction is issued when an agent has completed a cache sync and flush operation in response to an earlier FLUSH# signal assertion. The MCH claims this cycle and retires it.
0000 0110	Stop Grant Acknowledge	This transaction is issued when an agent enters Stop Grant mode. This cycle is claimed by the MCH and propagated over the hub interface as a Stop Grant special cycle. This cycle is retired on the host bus after the associated special cycle request is successfully broadcast over the hub interface.
0000 0111	SMI Acknowledge	This transaction is first issued when an agent enters the System Management Mode (SMM). $Ab[7]\#$ is also set at this entry point. All subsequent transactions from the host with $Ab[7]\#$ set are treated by the MCH as accesses to the SMM space. No corresponding cycle is propagated to the hub interface. To exit the System Management Mode the host issues another one of these cycles with the $Ab[7]\#$ bit deasserted. The SMM space access is closed by the MCH at this point.
all others	Reserved	

**NOTES:**

- None of the host bus special cycles are propagated to the AGP interface.

### 5.1.3. Symmetric Multiprocessor (SMP) Protocol Support

The MCH is optimized for uniprocessor system but supports the symmetrical multiprocessor configurations of up to two processors on the Processor System Bus for slot 1.

When configured for dual-processor, the 82820 MCH-based platform must integrate I/O APIC functionality and the associated buffer management required with an I/O APIC is handled via the special cycles on the hub interface.

### 5.1.4. Frame Buffer Memory Support

To allow for high-speed write capability for graphics, the Pentium® II and Pentium® III processors support WC (Write-combined memory type, can also be referred as USWC (uncacheable, speculative, write-combining). The USWC memory type provides a write-combining buffering mechanism for write operations. A high percentage of graphics transactions are writes to the memory-mapped graphics region, normally known as the linear frame buffer. Reads and writes to USWC are non-cached and can have no side effects.

In the case of graphics, current 32-bit drivers (without modifications) would use Partial Write protocol to update the frame buffer. The highest performance write transaction on the host bus is the Line Write. By combining several back-to-back Partial write transactions (internal to the processor) into a Line write transaction on the processor bus, the performance of frame buffer accesses would be greatly improved. Writes to USWC memory can be buffered and combined in the processor's write-combining buffers (WCB). The WCB is flushed after executing a serializing, locked, I/O instruction, or the WCB is full (32 bytes). To extend this capability to the current drivers, it is necessary to set up the linear frame buffer address range to be USWC memory type. This can be done by programming the MTRR registers in the processor.

Note that the application of USWC memory attribute is not limited only to the frame buffer support and that the MCH implements write combining for any host to the hub interface or host to AGP posted write.

## 5.2. AGP Interface

The MCH supports 3.3V AGP 1x/2x, and 1.5V AGP 1x/2x/4x devices. The AGP signal buffers have two modes of operation; 3.3V drive/receive (buffers are not 5 volt tolerant), and 1.5V drive/receive (buffers are not 3.3 volt tolerant). The MCH supports 2x/4x clocking transfers for read and write data, and sideband addressing. The MCH also support 2x and 4x clocking for Fast Writes initiated from the MCH (on behalf of the processor).

AGP PIPE# or SBA[7:0] transactions to DRAM do not get snooped and are, therefore, not coherent with the processor caches. AGP FRAME# transactions to DRAM are snooped. AGP PIPE# and SBA[7:0] accesses to and from the hub interface are not supported. AGP FRAME# access from an AGP master to the hub interface are also not supported. Only the AGP FRAME memory writes from the hub interface are supported.

### 5.2.1. AGP Target Operations

As an initiator, the MCH does not initiate cycles using AGP enhanced protocols. The MCH supports AGP cycles targeting interface to main memory only. The MCH supports interleaved AGP PIPE# and AGP FRAME#, or AGP SBA[7:0] and AGP FRAME# transactions.

**Table 14. AGP Commands Supported by the MCH When Acting as a AGP Target**

AGP Command	C/BE[3:0]# Encoding	MCH Host Bridge	
		Cycle Destination	Response as PCiX Target
Read	0000	Main Memory	Low Priority Read
	0000	The Hub interface	Complete with random data
Hi-Priority Read	0001	Main Memory	High Priority Read
	0000	The Hub interface	Complete with random data
Reserved	0010	N/A	No Response
Reserved	0011	N/A	No Response
Write	0100	Main Memory	Low Priority Write
	0100	The Hub interface	Cycle goes to DRAM with BEs inactive
Hi-Priority Write	0101	Main Memory	High Priority Write
	0101	The Hub interface	Cycle goes to DRAM with BEs inactive; does not go to the hub interface
Reserved	0110	N/A	No Response
Reserved	0111	N/A	No Response
Long Read	1000	Main Memory	Low Priority Read
		The Hub interface	Complete locally with random data; does not go to the hub interface
Hi-Priority Long Read	1001	Main Memory	High Priority Read
		The Hub interface	Complete with random data



AGP Command	C/BE[3:0]# Encoding	MCH Host Bridge	
		Cycle Destination	Response as PCIx Target
Flush	1010	MCH	Complete with QW of Random Data
Reserved	1011	N/A	No Response
Fence	1100	MCH	No Response - Flag inserted in MCH request queue
Reserved	1101	N/A	No Response
Reserved	1110	N/A	No Response
Reserved	1111	N/A	No Response

**NOTES:**

1. N/A refers to a function that is not applicable

As a target of an AGP cycle, the MCH supports all the transactions targeted at main memory (summarized in the table above). The MCH supports both normal and high priority read and write requests. The MCH will not support AGP cycles to the hub interface. PIPE# and SBA cycles do not require coherency management and all AGP initiator accesses to main memory using AGP PIPE# or SBA protocol are treated as non-snoopable cycles. These accesses are directed to the AGP aperture in main memory that is programmed as either uncacheable (UC) memory or write combining (WC) in the processor's MTRRs.

### 5.2.2. AGP Transaction Ordering

The MCH observes transaction ordering rules as defined by the AGP Interface Specification Revision 2.0.

### 5.2.3. AGP Signal Levels

The 4x data transfers use 1.5v signaling levels as described in the AGP 2.0 specification. The MCH supports 1x/2x data transfers using either the 3.3v or 1.5v signaling levels. The mechanism to select the data transfer mode is orthogonal to the mechanism to select the signaling level. Table 15 shows the data rates and signaling levels supported by the MCH.

**Table 15. AGP Data Rates and Signaling Levels**

Data Rate <sup>1</sup>	Signaling Level	
	1.5v	3.3v
1x AGP	Yes	Yes
2x AGP	Yes	Yes
4x AGP	Yes	No

**NOTES:**

1. AGP FRAME# data rate and signaling level is the same as 1x AGP.

If the AGP interface is running at 3.3V, see Section 5.6.2, "MCH Power Sequence Recommendation".

## 5.2.4. 4x AGP Protocol

In addition to the 1x and 2x AGP protocol, the MCH supports 4x AGP read and write data transfers and 4x sideband address generation. The 4x operation is compliant with AGP 2.0 specification.

The MCH indicates that it supports 4x data transfers through RATE[2] (bit 2) of the AGP Status Register. When DATA\_RATE[2] of the AGP Command Register is set to 1 during system initialization, the MCH performs AGP read/write data transactions using 4x protocol. This bit is not dynamic. Once this bit is set during initialization, the data transfer rate will not change.

The 4x data rate transfer provides 1.06 GB/s transfer rates. The control signal protocol for the 4x data transfer protocol is identical to 1x/2x protocol. In 4x mode 16 bytes of data are transferred on every 66 MHz clock edge. The minimum throttleable block size remains four 66 MHz clocks which means 64 bytes of data is transferred per block. Three additional signal pins are required to implement the 4x data transfer protocol. These signal pins are complimentary data transfer strobes for the AD bus (2) and the SBA bus (1).

## 5.2.5. Fast Writes

The MCH supports 2x and 4x Fast Writes from the MCH to the graphics controller on AGP. Fast Write operation is compliant with Fast Writes as currently described in AGP 2.0. To use the Fast Write protocol, both AGPCTRL[FWCE] (bit 1 of the AGP Control Register) and AGPCMD[FWPE] (bit 4 of the AGP Command Register) must be set to 1.

AGPCTRL[FWCE] is set to 0 by default. When this bit is set to 1, the MCH indicates that it supports Fast Writes through AGPSTAT[FW] (bit 4 of the AGP Status Register). When both AGPCMD[FWPE] and AGPCTRL[FWCE] are set to 1, the MCH uses Fast Write protocol to transfer memory write data to the AGP master.

Memory writes originating from the host(s) or from the hub interface use the Fast Write protocol when it is both capability enabled and enabled. The data rate used to perform the Fast Writes is dependent on the bits set in the AGP Command Register bits 2:0 (DATA\_RATE). If bit 2 of the AGPCMD[DATA\_RATE] field is 1, the data transfers occur using 4x strobing. If bit 1 of AGPCMD[DATA\_RATE] field is 1, the data transfers occur using 2x strobing. If bit 0 of AGPCMD[DATA\_RATE] field is 1, Fast Writes are disabled and data transfers occur using standard PCI protocol. Note that only one of the three DATA\_RATE bits may be set by initialization software. This is summarized in the following table.

FWCE	FWPE	DATA_RATE [2]	DATA_RATE [1]	DATA_RATE [0]	MCH =>AGP Master Write Protocol
0	0	x	x	x	1x
1	1	0	0	1	1x
1	1	0	1	0	2x Strobing
1	1	1	0	0	4x Strobing

## 5.2.6. AGP Universal Connector

The MCH supports the AGP Universal Connector. The AGP Universal Connector allows either a 1.5V or a 3.3V AGP add-in card to be supported by the system.

## 5.2.7. AGP FRAME# Transactions on AGP

The MCH accepts and generates AGP FRAME# transactions on the AGP bus. The MCH guarantees that AGP FRAME# accesses to DRAM are kept coherent with the processor caches by generating snoops to the host bus. LOCK#, SERR#, and PERR# signals are not supported.

### MCH Initiator and Target Operations

The following table summarizes MCH target operation for AGP FRAME# initiators. The cycles can be either destined to main memory or the hub interface.

**Table 16. PCI Commands Supported by the MCH When Acting as A FRAME# Target<sup>1</sup>**

PCI Command	C/BE[3:0]# Encoding	MCH	
		Cycle Destination	Response as A FRAME# Target
Interrupt Acknowledge	0000	N/A	No Response
Special Cycle	0001	N/A	No Response
I/O Read	0010	N/A	No Response
I/O Write	0011	N/A	No Response
Reserved	0100	N/A	No Response
Reserved	0101	N/A	No Response
Memory Read	0110	Main Memory	Read
	0110	The Hub interface	No Response
Memory Write	0111	Main Memory	Posts Data
	0111	The Hub interface	No Response
Reserved	1000	N/A	No Response
Reserved	1001	N/A	No Response
Configuration Read	1010	N/A	No Response
Configuration Write	1011	N/A	No Response
Memory Read Multiple	1100	Main Memory	Read
	1100	The Hub interface	No Response
Dual Address Cycle	1101	N/A	No Response
Memory Read Line	1110	Main Memory	Read
	1110	The Hub interface	No Response
Memory Write and Invalidate	1111	Main Memory	Posts Data
	1111	The Hub interface	Posts Data

**NOTES:**

1. N/A refers to a function that is not applicable

As a **target** of an AGP FRAME# cycle, the MCH only supports the following transactions:

- *Memory Read, Memory Read Line, and Memory Read Multiple.* These commands are supported identically by the MCH. The MCH does not support reads of the hub interface bus from AGP.
- *Memory Write and Memory Write and Invalidate.* These commands are aliased and processed identically.
- *Other Commands.* Other commands such as I/O R/W and Configuration R/W are not supported by the MCH as a target and result in master abort.
- *Exclusive Access.* The MCH does not support PCI locked cycles as a target.
- *Fast Back-to-Back Transactions.* MCH as a target supports fast back-to-back cycles from an AGP FRAME# initiator.

As an **initiator** of AGP FRAME# cycle, the MCH only supports the following transactions:

- *Memory Read and Memory Read Line.* MCH supports reads from host to AGP. MCH does not support reads from the hub interface to AGP.
- *Memory Read Multiple.* This command is not supported by the MCH as an AGP FRAME# initiator.
- *Memory Write.* MCH initiates AGP FRAME# cycles on behalf of the host or the hub interface. MCH does not issue Memory Write and Invalidate as an initiator. MCH does not support write merging or write collapsing. MCH allows non-snoopable write transactions from the hub interface to the AGP bus.
- *I/O Read and Write.* I/O read and write from the host are sent to the AGP bus. I/O base and limit address range for AGP bus are programmed in AGP FRAME# configuration registers. All other accesses that do not correspond to this programmed address range are forwarded to the hub interface.
- *Exclusive Access.* MCH does not issue a locked cycle on the AGP bus on behalf of either the host or the hub interface. The hub interface and host locked transactions to AGP are initiated as unlocked transactions by the MCH on the AGP bus.
- *Configuration Read and Write.* Host Configuration cycles to AGP are forwarded as Type 1 Configuration Cycles.
- *Fast Back-to-Back Transactions.* MCH as an initiator does not perform fast back-to-back cycles.

## MCH Retry/Disconnect Conditions

The MCH generates retry/disconnect according to the AGP Specification rules when being accessed as a target from the AGP FRAME# device.

## Delayed Transaction

When a AGP FRAME#-to-DRAM read cycle is retried by the MCH, it is processed internally as a Delayed Transaction.

The MCH supports the Delayed Transaction mechanism on the AGP target interface for the transactions issued using AGP FRAME# protocol. This mechanism is compatible with the PCI 2.1 Specification. The process of latching all information required to complete the transaction, terminating with Retry, and completing the request without holding the master in wait-states is called a Delayed Transaction. The MCH latches the Address and Command when establishing a Delayed Transaction. The MCH generates a Delayed Transaction on the AGP only for AGP FRAME# to DRAM read accesses. The MCH does not allow more than one Delayed Transaction access from AGP at any time.

## 5.3. DRAM Interface

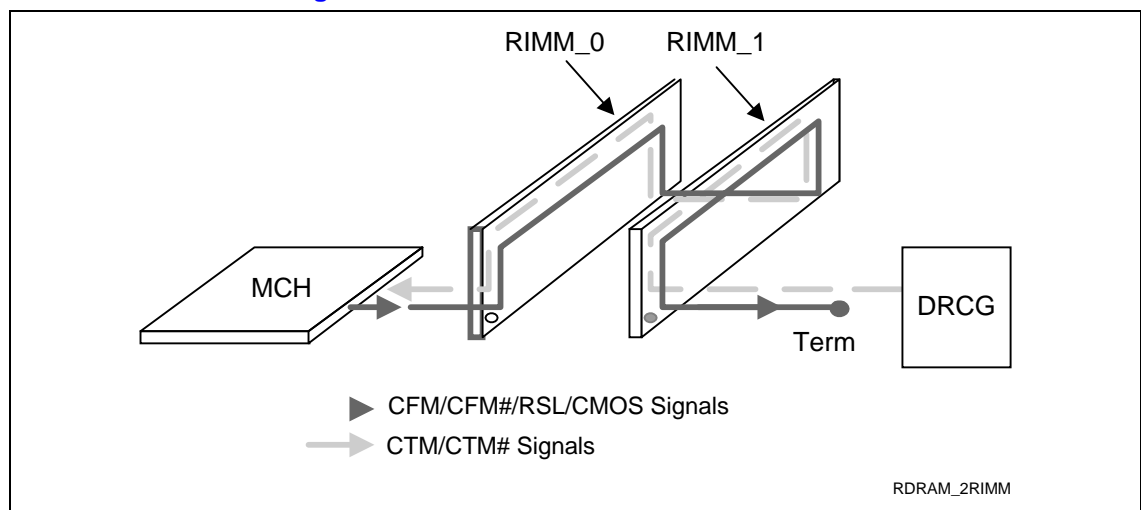
The MCH supports a single Direct RDRAM channel as described in Section 5.3.1, “RDRAM”. The interface between MCH and Direct RDRAM devices is referred to as a “channel.”

### 5.3.1. RDRAM

The Direct RDRAM channel consists of 33 signals including clocks (30 are RSL signals and 3 are CMOS signals). Figure 4 shows the interconnections between the MCH and one Direct RDRAM channel. The MCH supports the following Features:

- Supports a single Direct RDRAM channel without external glue logic
- Supports up to 32 Direct RDRAM devices on a single channel
- Supports 64Mbit, 128Mbit and 256Mbit Direct RDRAM technology
- Supports Standby, NAP and Power Down modes
- Supports up to 8 pages open across all devices at a time
- Supports only the dependent bank type of Direct RDRAM devices
- Supports PC600, PC700, and PC800 Direct RDRAM modules
- Supports active power management of Direct RDRAM devices
- Supports only the page size of 1 KB
- Supports thermal throttling to manage RDRAM memory subsystem power consumption
- Supports ECC

Figure 4. MCH/ Direct RDRAM Diagram



The maximum system memory supported by the MCH depends on the Direct RDRAM device technology. The following table shows the maximum memory supported in various configurations.

**Table 17. Maximum Supported Direct RDRAM Configurations**

Direct RDRAM Technology	Maximum Memory
64Mbit	256 MB
128Mbit	512 MB
256Mbit	1 GB

The row, column, and bank address bits required for the Direct RDRAM device depends on the number of banks and page size of the device. The following table shows the different combinations supported by the MCH.

**Table 18. Direct RDRAM Device Configurations**

CF#	Device Tech	Device Capacity in MB	# of Banks (D= dependent)	Page Size	# of Bank Address Bits	# of Row Address Bits	# of Column Address Bits
1	64Mbit	8	16 (D)	1KB	4	9	6
2	128Mbit	16	2x16 (D)	1KB	5	9	6
3	256Mbit	32	2x16 (D)	1KB	5	10	6
3a	256Mbit	32	2x16 (D)	2KB	5	9	7

A brief overview of the registers that configure the Direct RDRAM interface is provided below.

- **Group Boundary Address Register (GBA):** GBA registers define the upper and lower addresses for a group of Direct RDRAM devices in a channel. Each group requires a separate GBA register. Each group consists of 4 devices. The MCH contains 8 GBA registers.
- **Group Architecture Register (GAR):** GAR registers specify the architecture features of each group of devices in a channel. The architecture features specified are bank type and device core technology. Each GAR represents a group consisting of 4 devices. There is a 1:1 correspondence between GBA and GAR registers.
- **RDRAM Timing Register (RDTR):** The DTR defines the timing parameters for all devices in all channels. BIOS programs this register with “least common denominator” values after reading configuration registers of each device in the channel.
- **RDRAM Control Register (RDCR):** This register provides bits to enable or disable the programmable features available in the RMC.
- **RDRAM Power Management Register (RPMR):** This register provides bits to program Direct RDRAM power management features.
- **RDRAM Initialization Control Register (RICM):** This register provides bits to program MCH to do initialization activities on Direct RDRAM devices.

### 5.3.1.1. DRAM Organization, Configuration, and Speed

The MCH supports 16/18-bit Direct RDRAM configurations. The MCH supports a maximum of 32 devices on a single Direct RDRAM channel. Direct RDRAM channel can be populated with a mix of 64Mbit, 128Mbit, and 256Mbit Direct RDRAM devices.

The MCH supports PC600, PC700, and PC800 Direct RDRAM modules. Each module may contain several devices. The rated speed of those devices should not be confused with the interface speed of the Direct RDRAM channel.

The supported Direct Rambus channel speeds are determined by the RDRAM device speed and the Processor System Bus frequency. The MCH supports 300 MHz Direct Rambus channel at 100 MHz Processor System Bus frequency, 266 MHz and 356 MHz Direct Rambus channel at 133 MHz Processor System Bus frequency, and 400 MHz Direct Rambus channel at 100 MHz or 133 MHz Processor System Bus frequency.

**Table 19. Supported Direct Rambus Channel Speeds**

RDRAM Module	Direct Rambus Channel Frequency / Processor System Bus Frequency (both in MHz)				
	266 / 133	300 / 100	356 / 133	400 / 100	400 / 133
PC600	X	X			
PC700	X	X	X		
PC800	X	X	X	X	X

### 5.3.1.1.1. Rules for Populating RDRAM Devices

A channel can be implemented such that it is fully (32 devices) or partially loaded with RDRAM devices. From the MCH point of view, all devices in the channel are grouped into logical groups. The MCH supports up to 8 groups. All devices populated in a group must be of the same architecture. Thus, all devices in a group must be of the same core technology and have the same number of banks. The following are the rules for populating the groups:

- A group can be partially populated.
- There is no requirement that group members have to be populated in contiguous physical slots.
- There can be a maximum of 8 groups in a channel. A member that does not belong to any of the groups in the channel will not be recognized.

Table 20 provides the device IDs for members in all groups.

**Table 20. RDRAM Device Grouping**

Device IDs for Group Members	Group Name
0, 1, 2, 3	Group#0
4, 5, 6, 7	Group#1
8, 9, 10, 11	Group#2
12, 13, 14, 15	Group#3
16, 17, 18, 19	Group#4
20, 21, 22, 23	Group#5
24, 25, 26, 27	Group#6
28, 29, 30, 31	Group#7

### 5.3.1.1.2. RDRAM CMOS Signals Description and Usage

There are 3 CMOS signal pins on the MCH to support Direct RDRAM device configuration, SIO reset, register accesses, and Nap and PowerDown exits. These signals are SCK, CMD and SIO and are used to perform the following operations:

- SIO pin initialization
- SIO operations (includes register accesses and device reset)
- Device selection for Nap and PowerDown exits

Figure 5. Sideband CMOS Signal Configuration

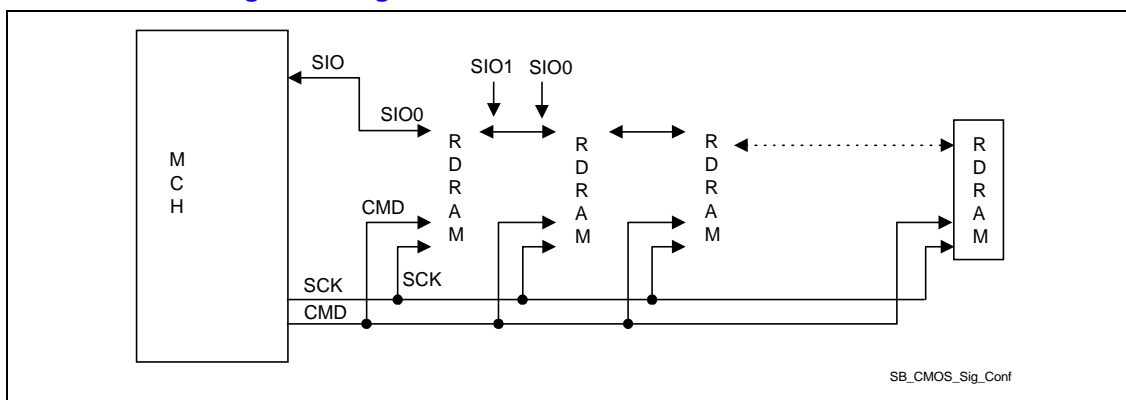


Table 21. Sideband CMOS Signal Description

Signal	Description
SCK	<p><b>Serial Clock:</b> This signal serves as the clock for SIO and CMD signals. SCK is a clock source used for reading from and writing to control register.</p> <ul style="list-style-type: none"> <li>• For SIO operations and pin initialization, <math>SCK \leq 1</math> MHz</li> <li>• For power mode operations, <math>SCK \leq 100</math> MHz</li> </ul>
CMD	<p><b>Command:</b> CMD is a control signal used for power mode transitions, SIO pin configuration during initialization, and framing of SIO operations. CMD is active high. CMD is sampled at both edges of SCK. CMD is a level sensitive signal.</p>
SIO	<p><b>Serial In Out:</b> This bi-directional signal is daisy chained through all Direct RDRAM (Sin to Sout) devices in a channel. This pin carries data used for SIO operations including register accesses, device reset, and device ID initialization. SIO is also used for power mode control. SIO is an active low signal and is sampled on the falling edge of SCK.</p>

Table 22. CMD Signal Value Decode

SIO = 0, CMD Sample Value on 4 SCK Edges				Command	SIO = 1, CMD Sample Value on 4 SCK Edges				Command
Cycle 0		Cycle 1			Cycle 0		Cycle 1		
0	1	X	X	Nap Exit	0	1	x	x	Power-down Exit
1	0	X	X	Reserved	1	0	x	x	Reserved
0	0	X	X	No-op	0	0	x	x	No-op
1	1	1	1	SIO Request Frame	1	1	1	1	SIO Request Frame
1	1	0	0	SIO Reset	1	1	0	0	SIO Reset
1	1	1	0	Reserved	1	1	1	0	Reserved
1	1	0	1	Reserved	1	1	0	1	Reserved



- **SIO Pin Initialization:** Sin and Sout pins on Direct RDRAM devices are bi-directional and their direction needs to be initialized. The “SIO Reset” initializes Sin and Sout pins on all Direct RDRAMs as a daisy chain configuration and is performed with SCK and CMD. Once the SIO daisy chain is fully configured, SIO operations can occur. Note “SIO Reset” does NOT reset the entire device. For a complete description of operation and associated timing diagram, refer to Direct RDRAM data sheet from Rambus.
- **SIO Operations:** SIO operations are also known as Direct RDRAM initialization operations. These operations include Direct RDRAM register accesses and device reset, and are performed using the CMOS pins: SCK, CMD, Sin, and Sout. For a complete description of operation and associated timing diagram, refer to Direct RDRAM data sheet from Rambus.
- **SIO Timing Restriction:** There is a timing restriction between the RDRAM CMOS commands that operate at 1 MHz SCK clock and the commands that operate at 100 MHz SCK clock. After one of the 1 MHz CMOS commands is issued, 1 us of delay is required before issuing one of the 100 MHz CMOS commands.
  - The 1 MHz commands include SRD, SWR, SETR, CLRR, SETF, and SIO Reset.
  - The 100 MHz commands include NAP Exit and PowerDown Exit.
- **Nap and PowerDown Exits:** The Nap and Power Down exits are performed using CMD, SIO and SCK signals. For complete description and timing diagrams associated with Nap and PowerDown exits, refer to Direct RDRAM data sheet from Rambus.

### 5.3.1.1.3. Direct RDRAM Core Refresh

#### RDRAM Refreshes

All rows in a Direct RDRAM device must be refreshed within 32 ms. The refresh rate depends on the device size and page size of a device (Refer to Table 23). RDRAM is refreshed on a bank-by-bank basis, so the MCH keeps track of the bank address of the last refreshed banks. If a channel is populated with devices that have different refresh rate requirements, the fastest refresh rate required device must be used for all devices in the channel.

**Table 23. RDRAM Refresh Rate**

Device Tech	# of Banks (D = dependent)	Device Capacity	Refresh Interval	Page Size	Required Refresh Rate
64Mbit	16 (D)	8 MB	32 ms	1 KB	3.9 us
128Mbit	2x16 (D)	16 MB	64 ms	1 KB	3.9 us
256Mbit	2x16 (D)	32 MB	64 ms	1 KB	3.9 us
256Mbit	2x16 (D)	32 MB	32 ms	2 KB	3.9 us
256Mbit	2x16 (D)	32 MB	64 ms	2 KB	7.8 us

**NOTES:**

1. 256Mbit refresh rates are based upon multi-bank refresh support in the 256Mbit RDRAM generation.

## Refreshes for RDRAM Devices in Nap Mode

When the devices are in the Nap mode, their DLLs need to be refreshed every 10 us to maintain the accurate phase information. Nap exit is performed by sending device ID on the DQ lines along with Rambus CMOS control signals. An RDRAM device can remain in Nap mode for a maximum period of time before it must be brought out of Nap mode, into Standby or Active. Then normal refreshes for core are executed and scheduled along with DLL refreshes.

## Refreshes for RDRAM Devices in Standby, Active Modes, and Power Down

**Active Refresh:** Refresh and precharge after refresh commands are issued from the primary control packet. These commands provide refresh support in Standby/Active modes.

**Self Refresh:** Internal timebase and row/bank address counters in the RDRAM core allow for a self refresh in PowerDown modes without controller support. Appropriate control register(s) in the RDRAM devices must be programmed by BIOS during initialization to enable this feature.

## MCH Support for RDRAM Refreshes

The MCH only does active refreshes every 10 us, which is a DLL refresh cycle. The refresh is burst every 10 us. For example, if the DRR bit is set to 000, the MCH bursts 5 refreshes during each 10 us DLL refresh.

### 5.3.1.1.4. Direct RDRAM Current Calibration

Due to the variation of the operating voltage, operating temperature, and component skew, it is necessary to perform current calibration for all devices (including MCH) residing on the RDRAM channel. The current calibration adjusts the output current (IOL) of devices to their proper range. All devices on the RDRAM interface must be calibrated once every 100 ms. The MCH supports two separate current calibration operations.

- First, the MCH supports internal current calibration for its output buffers on the RDRAM interface.
- Second, the MCH supports the RSL command that is described in the *Rambus RDRAM specification* to perform the current calibration function. The MCH must schedule periodic current calibration operation such that every RDRAM device in the channel is calibrated at least once per 100 ms.

### 5.3.1.2. Direct RDRAM Command Encoding

The operations on a Direct RDRAM channel are performed using control packets. There are two types of command packets: row (ROWA/ROWR) packet and column (COLC/COLM/COLX) packet. Each command packet requires 4 Direct RDRAM clock durations and packet data is transferred on both (leading and falling) edges of the clock. Row packet contains 24 bits and column packet contains 40 bits.

#### 5.3.1.2.1. Row Packet (ROWA/ROWR)

The row packet is defined using three RSL signals RQ[7:5]/ROW[2:0]. It will generally be the first control packet issued to a device. Major characteristics of the row packet are the following:

- The only way to activate (sense) a row within a bank
- Independent of Direct RDRAM's active/standby state
- A non-broadcast row package causes an addressed Direct RDRAM to move to active state

The packet definition of row packet is given below.

**Table 24. ROWA Packet for Activating (sensing) a Row (i.e., AV = 1)**

Row	Cycle 0		Cycle 1		Cycle 2		Cycle 3	
ROW2	DR4T	DR[2]	BR[0]	BR[3]	R[10]	R[8]	R[5]	R[2]
ROW1	DR4F	DR[1]	BR[1]	BR[4]	R[9]	R[7]	R[4]	R[1]
ROW0	DR[3]	DR[0]	BR[2]	REV	AV = 1	R[6]	R[3]	R[0]

**Table 25. ROWR Packet for Other Operations (i.e., AV = 0)**

Row	Cycle 0		Cycle 1		Cycle 2		Cycle 3	
ROW2	DR4T	DR[2]	BR[0]	BR[3]	ROP[10]	ROP[8]	ROP[5]	ROP[2]
ROW1	DR4F	DR[1]	BR[1]	BR[4]	ROP[9]	ROP[7]	ROP[4]	ROP[1]
ROW0	DR[3]	DR[0]	BR[2]	REV	AV = 0	ROP[6]	ROP[3]	ROP[0]

DR4T	DR4F	Device ID
0	0	No row packet
0	1	DR[3:0], DR[4] = 0
1	0	DR[3:0], DR[4] = 1
1	1	Broadcast

DR[4] –DR[0]	Device address
BR[4] –BR[0]	Bank Address
R[10] –R[0]	Row address
AV	Select between ROWA and ROWR, Active Row
ROP[10] – ROP[0]	Opcode for Primary Control Packet
REV	Reserved

Table 26. ROWA and ROWR Packet Field Encodings

AV	Opcode bits									Command Description
	10	9	8	7	6	5	4	3	2: 0	
1	x	x	x	x	x	X	x	x	xxx	Activate Row
0	1	1	0	0	0	0	0	0	000	Precharge
0	1	1	0	0	0	0	0	1	000	Precharge and Relax
0	1	1	0	0	0	1	0	1	000	Precharge and Nap and Relax
0	0	0	0	1	1	0	0	0	000	Refresh
0	1	0	1	0	1	0	0	0	000	Precharge Post-refresh
0	0	0	0	0	0	1	0	0	000	Nap
0	0	0	0	0	0	1	0	1	000	Nap and Relax
0	0	0	0	0	0	1	1	0	000	Conditional Nap
0	0	0	0	0	0	1	1	1	000	Conditional Nap and Relax
0	0	0	0	0	0	0	1	0	000	Power Down
0	0	0	0	0	0	0	0	1	000	Relax
0	0	0	0	0	0	0	0	0	010	Temp Calibration Enable
0	0	0	0	0	0	0	0	0	001	Temp Calibration
0	0	0	0	0	0	0	0	0	000	No Row Operation

**NOTES:**

1. x = Controller drives 0 or 1
2. 0 = Controller drives 0
3. 1 = Controller drives 1

**5.3.1.2.2. Column Packet (COLC/COLX/COLM)**

The column packet is defined using five of the RSL signals RQ[4:0]/COL[4:0]. Major characteristics of column are the following:

- The only way to dispatch column operation for read or write
- Requires the target Direct RDRAM to be in the active state

**Note:** When an Direct RDRAM is in active state, it can receive both row and column packets. When a Direct RDRAM is in Standby state, it can only receive a row packet. So, before sending a column packet, make sure the addressed Direct RDRAM is in active state.

The packet definition of column packet is given below.

**Table 27. COLC Packet**

Column	Cycle 0		Cycle 1		Cycle 2		Cycle 3	
COL4	DC[4]	S = 1					REV	C[4]
COL3	DC[3]						C[5]	C[3]
COL2	DC[2]	COP[1]				REV	BC[2]	C[2]
COL1	DC[1]	COP[0]				BC[4]	BC[1]	C[1]
COL0	DC[0]	COP[2]			COP[3]	BC[3]	BC[0]	C[0]

**NOTES:**

1. DC[4:0] Device ID for Column Operation
2. S Start bit, for framing
3. M Mask bit. Asserted indicates mask format for packet
4. COP[3:0] Column Operation Code
5. C[5:0] Address for Column operation
6. BC[4:0] Bank Address for Column operation
7. REV Reserved

**Table 28. COLC Packet Field Encodings**

S	COP[3]	COP[2]	COP[1]	COP[0]	Command Description
0	X	X	x	X	No operation
1	X	0	0	0	NOCOP. Retire write buffer of this device
1	X	0	0	1	Write
1	X	0	1	1	Read

**NOTES:**

1. All other combination are reserved

**Table 29. COLM Packet (M = 1)**

Column	Cycle 0		Cycle 1		Cycle 2		Cycle 3	
COL4			MA[7]	MA[5]	MA[3]	MA[1]		
COL3		M = 1	MA[6]	MA[4]	MA[2]	MA[0]		
COL2			MB[7]	MB[4]	MB[1]			
COL1			MB[6]	MB[3]	MB[0]			
COL0			MB[5]	MB[2]				

Table 30. COLX Packet (M = 0)

Column	Cycle 0		Cycle 1		Cycle 2		Cycle 3	
COL4			DX[4]	XOP[4]	REV	BX[1]		
COL3		M = 0	DX[3]	XOP[3]	BX[4]	BX[0]		
COL2			DX[2]	XOP[2]	BX[3]			
COL1			DX[1]	XOP[1]	BX[2]			
COL0			DX[0]	XOP[0]				

**NOTES:**

1. DX[4:0] Device ID for Extra operation
2. BX[4:0] Bank Address for Extra operation
3. MA[7:0] Byte Mask (low order)
4. MB[7:0] Byte Mask (high order)
5. XOP[4:0] Opcode for Extra Operation
6. REV Reserved

Table 31. COLM Packet and COLX Packet Field Encodings

M	XOP Bits					Command Description
	4	3	2	1	0	
1	x	x	x	X	X	Bytemasks
0	0	0	0	0	0	NoXop
0	0	1	0	0	0	Calibrate Current
0	0	1	1	0	0	Calibrate Current and Sample

**NOTES:**

1. x = Controller drives 0 or 1
2. 0 = Controller drives 0
3. 1 = Controller drives 1

**5.3.1.2.3. Data Packet**

The definition of the data packet is given Table 32.

Table 32. Data Packet

	Cycle 0		Cycle 1		Cycle 2		Cycle 3	
<b>DQA[8:0]</b>	DA0[8:0]	DA1[8:0]	DA2[8:0]	DA3[8:0]	DA4[8:0]	DA5[8:0]	DA6[8:0]	DA7[8:0]
<b>DQB[8:0]</b>	DB0[8:0]	DB1[8:0]	DB2[8:0]	DB3[8:0]	DB4[8:0]	DB5[8:0]	DB6[8:0]	DB7[8:0]

### 5.3.1.3. Direct RDRAM Register Programming

Software can read and write Direct RDRAM device registers by programming the Direct RDRAM Initialization Control Management (DICM) Register in MCH. The register data returned by the device is available in the Device Register Data (DRD) Register.

### 5.3.1.4. Direct RDRAM Operating States

The Direct RDRAM devices support different operating and idle states to minimize the power consumption and thermal overload. Table 33 provides an overview of the different operating/power states supported by Direct RDRAMs.

**Table 33. DRAM Operating States**

Direct RDRAM State	Functionality	Refresh Scheme	RDRAM Clock State
<b>Inactive States</b>			
PowerDown	No operation allowed except refresh. Direct RDRAM awaits CMOS signals to exit PowerDown state	Self Refresh	stopped
Nap	No operation allowed except refresh. Direct RDRAM awaits Nap exit command to exit Nap	Active Refresh	stopped
<b>Active States</b>			
Standby	Device Ready to receive row packet . with fast clock	Active Refresh	full speed
Active	Device ready to receive any control packet	Active Refresh	full speed
Active-Read	Device ready to receive any control packet. Transmitting data on channel	Active Refresh	full speed
Active-Write	Device ready to receive any control packet. Receiving data from channel	Active Refresh	full speed

- **Active-Read/Write State:** A Direct RDRAM device is in active-Read/Write state when it is transferring data. This state lasts as long as data transfer is occurring. Once the data transfer is done, the Direct RDRAM transitions into the Active or Standby state based on the column command last executed.
- **Active State:** A Direct RDRAM enters into active state immediately after the data transfer from/to that device is done and the last COLC command that caused the data transfer does not have its RC bit set to 1. When a device is in the Active state, it can accept both row and column packets.
- **Standby State:** A Direct RDRAM enters into Standby state either from Active-Read/Write or Active state. Transition from Active-Read/Write to Standby happens if the last column executed has its RC bit set to 1. Transition from Active to Standby happens if COLC or row specifies an operation with Relax. When a device is in Standby mode, it can accept only row packet. Once a device receives any row packet, it transitions into active state and then only it can accept an column packet.
- **Nap State:** A Direct RDRAM enters into Nap state when it receives a row packet which specified an operation with Nap. No operations except refresh is allowed during Nap state.
- **PowerDown State:** A Direct RDRAM enters into PowerDown state when it receives a row packet which specified an operation with PowerDown. No operations except Self-refresh is allowed during PowerDown state.

### 5.3.1.5. RDRAM Operating Pools

To minimize the operating power the RDRAM devices are grouped into two operating pools called Pool “A” and Pool “B”.

#### Pool “A” and Pool “B” Operation

Up to 8 devices can be in Pool “A” at a time. Up to 4 out of 8 in Pool “A” can be in Active Read/Write or Active state at a time. The devices in Pool “A” are in either Active Read/Write, Active, or Standby state. The maximum number of devices in Pool “A” is specified by PAC field of RPMR register.

All devices that do not reside in Pool “A” are members of Pool “B”. All devices in Pool “B” are either in Standby or Nap state. The state of the devices in Pool “B” is specified by PBS field of DRAMC register.

### 5.3.1.6. RDRAM Power Management

Intel® 820/820E chipset systems support ACPI-based power management. The MCH puts all RDRAM devices into a PowerDown (PD) state during S3 power management states. To enter the PowerDown state all RDRAM devices in the channel must be in active or standby state. The MCH then sends a broadcast PowerDown command to that channel.

During PowerDown state, RDRAM devices are put into Self Refresh mode so that external (active) refreshes are not required. During the power-down state, the clocks to RDRAM are shut off.

Exiting the power-down and Nap states are done through CMOS signals. The following table shows the actions taken by MCH during different processor and system power states.

**Table 34. RDRAM Power Management States**

Processor State	System State	State of RDRAMs in Pool “A”	State of RDRAMs in Pool “B”	Refresh Scheme	RDRAM Clock State
C0, C1, C2 (processor in working state)	S0	Active-Read/Write, Active, Standby	Nap	Active	Running
(processor in inactive state)	S1, S3( STR)	No devices in Pool “A”	Power-down	Self	Stopped



### 5.3.1.7. Data Integrity

The MCH supports two data integrity modes for main memory—Error Correcting Code (or Error Checking and Correcting) and Error Checking (EC). In the Error Correcting Code mode, the MCH can optionally be configured to generate the ECC code for writes to memory and check the code for reads from memory. The MCH generates an 8-bit code word for each 64-bit QWord of memory. Since the code word covers a full QWord, writes of less than a QWord require a read-merge-write operation. Consider a DWord write to memory. In this case, when in ECC mode, the MCH reads the QWord where the addressed DWord is written, merges in the new DWord, generates a code covering the new QWord and finally writes the entire QWord and code back to memory. Any correctable (single) errors detected during the initial QWord read are corrected before merging the new DWord.

For the Error Checking mode, the MCH generates and stores a code for each QWord of memory. It then checks the code for reads from memory but does not correct any errors that are found. Thus, the read performance hit associated with ECC is not incurred. Errors are still reported in the Error Status register as they are for ECC mode.

Error scrubbing is supported by the MCH. When enabled, this feature not only corrects single bit error data being returned to the requesting agent but also writes the corrected value back to the DRAM array.

Single bit and multiple bit errors set separate flags in the ERRSTS register. Single bit errors and multiple bit errors can be independently enabled to generate hub interface SERR, SMI, or SCI special cycles to the I/O Controller Hub. The address and syndrome of the first single bit error are latched in the EAP register. Subsequent single bit errors will not overwrite the EAP register unless the single bit error status bit is cleared. A multiple bit error will overwrite the EAP register. Subsequent multiple bit errors will not overwrite the EAP register unless the multiple bit error status bit is cleared.

**Note:** When an 82820-based platform is configured for ECC support, if a multi-bit uncorrectable memory error is detected during a memory read by a system device, an SERR, SCI, or SMI will be generated. This typically results in an NMI; however, bad data may still reach the intended target before the NMI can be generated or before the NMI interrupt handler can service the situation. This may result in bad data being returned to the target and may be permanently stored, resulting in system data corruption. This chipset family does not ensure that targets are protected from this corrupted data in these situations.

### 5.3.1.8. RDRAM Array Power Management

To manage the thermal condition of RDRAM array the MCH implements a choice of two methods to trigger the same RDRAM throttling mechanism:

- Method 1: RDRAM on-die internal thermal sensor mechanism.
- Method 2: RDRAM Read/Write Counters/Timers mechanism. This monitors the amount of DRAM read and write traffic to/from RDRAMs by indirectly estimating DRAM array temperature.

OEMs have a choice of using one or the other triggering mechanism via programming bit [2:1] of the DTC register. The bits available for enabling error reporting are in the ERRSTS and ERRCMD registers.

### 5.3.1.8.1. RDRAM On-die Internal Thermal Sensor Mechanism (Method 1)

With the on-die thermal sensors for each RDRAM device and the Current Calibration protocol specified by Direct RDRAM channel, the MCH can trigger the RDRAM throttling based on the settings of the TEMP register in the RDRAM device and the DTC register in the MCH.

#### Sensor Operation (Method 1)

Current calibration is typically performed on a per-device basis every 100 ms. The “Current Calibration” command is delivered to a RDRAM device through XOP encodings of a Column Packet (COLX). The addressed RDRAM device will respond to the Current Calibration command by emitting a 16 byte data packet. DQA5 of the emitted data packet is the Temperature Sensor Output.

For the DQA5 data bit, **normal device temperature is represented by a "low pin voltage" (driven low)**. When the device's on-die thermal sensor is triggered, DQA5 returns a "high pin voltage". Note that this is a special case that a "high pin voltage" represents a logical 1 (triggered thermal sensor) instead of normally representing a logical 0 in RSL signaling.

#### Thermal Throttling (Method 1)

The MCH performs Current Calibration on each device within the RDRAM array. DQA5 of the emitted data packet from each device is monitored. **A "high pin voltage" on DQA5 represents that the device thermal sensor has been triggered.** (Note that this data bit is a special case where a "high pin voltage" represents a logical 1 instead of the normal logical 0 in RSL signaling.) If the sensor is triggered (DQA5="high pin voltage" and thermal sensor throttling is enabled), RDRAM throttling is invoked according to the TDM, TMW, and TT throttle settings in the DTC register. The RDRAM Thermal Sensor Flag bit is set in the ERRSTS register and an SERR error message is generated via the hub interface (if enabled through the ERRCMD register). RDRAM throttling remains in effect for the duration of time specified by the TT field in the DTC register. Current Calibration continues while throttling is in effect, but the returned DQA5 is ignored. After TT has elapsed, DQA5 is again evaluated for each Current Calibration and throttling is invoked when the sensor is triggered (DQA5="high pin voltage").

### 5.3.1.8.2. RDRAM Reads and Writes Counters/Timers Mechanism (Method 2)

This method estimates the DRAM array temperature indirectly by monitoring the amount of DRAM read and write traffic to/from RDRAMs.

#### Counters/Timers (Method 2)

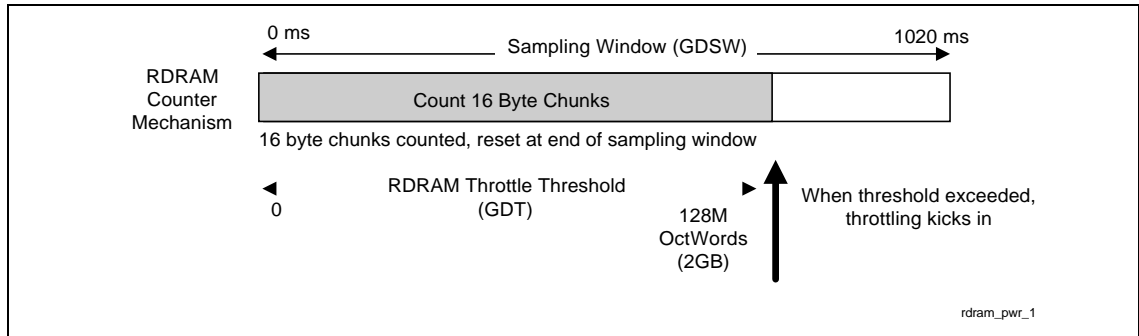
This mechanism uses programmable timers and counters to monitor DRAM traffic for estimating the temperature of the RDRAM array indirectly. Both reads/writes traffic to/from RDRAM are counted together.

The monitor mechanism consists of a timer (GDSW) and a counter (GDT) to measure the amount of RDRAM bandwidth. While the timer (GDSW) is running, the number of RDRAM reads/writes in units of OctWords (16 byte blocks) is counted. A programmable counter threshold (GDT) can be set between 0 and 128M OctWords in 32K increments. If the number of OctWords (read and writes) during the monitoring period exceeds the counter threshold value, the DRAM throttle mechanism is invoked. Note

that a large value is recommended for GDSW (4–1020 ms) to allow short bursts of DRAM activity at peak bandwidth to occur without invoking throttling.

When DRAM throttling is invoked with the Counter mechanism, throttling occurs for a period of time specified by a field in the throttle register. When that period of time expires, throttling stops and the Counter mechanism begins monitoring DRAM traffic again.

When DRAM Read/Write throttling is invoked with the Counter mechanism, a corresponding flag is set in the ERRSTS register. When enabled in the ERRCMD register, a SERR error message is generated over the hub interface.



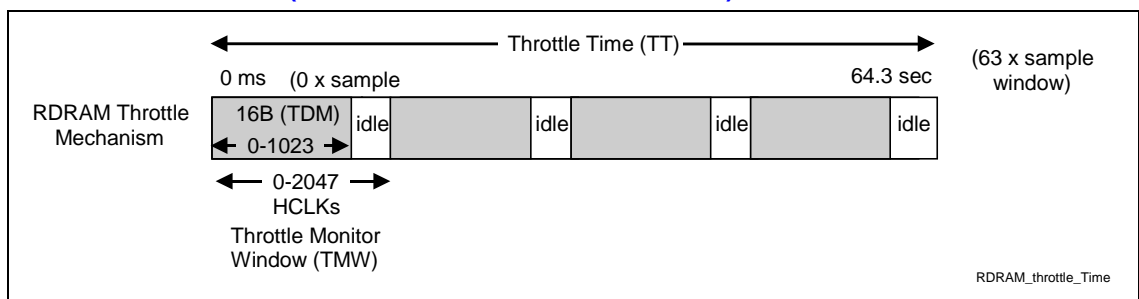
**Note:** The maximum value for GDT windows should be 128M OctWords (2 GB).

## Thermal Throttling (Method 2)

Throttling limits the maximum read and write DRAM bandwidth permitted. A single combined throttling mechanism is provided for reads and writes. The throttling mechanism consists of a timer and a counter. The timer (TMW: Throttle monitor Window) is programmable from 0 to 2047 host clocks in increments of 16 clocks and defines a time window within which DRAM traffic is monitored. While the timer is running the number of DRAM reads (or writes or both) in units of OctWords (16 bytes) are counted. A programmable counter threshold (TDM: Throttle DRAM Maximum) can be set between 0 and 1023. Once the number of OctWords read (or written or both) during the monitoring period reaches the counter threshold value, reads (writes) are no longer allowed to run to the DRAM for the remainder of the throttle time window.

When DRAM throttling is invoked with the Counter mechanism, the throttle monitor windows continuously repeat for a period of time specified by two fields in the throttle register; GDSW and TT. TT (Throttle Time) has a value of 0 to 63 and is multiplied by the value in GDSW (0 to 1020 ms in increments of 4 ms). When that period of time expires, throttling stops and the Counter mechanism begins monitoring DRAM traffic again.

**Figure 6. RDRAM Throttle Time (With Counters/Timers Mechanism)**



## 5.4. Power Management

The platform is compliant with the following specifications:

- APM Rev 1.2
- ACPI Rev 1.0
- PCI Power Management Rev 1.0
- PC'97, Rev 1.0
- PC'98, Rev 1.0

### 5.4.1. Processor Power State Control

- **C0 (Full On):** This is the only state that runs software. All clocks are running, STPCLK# is deasserted and the processor core is active. The processor can service snoops and maintain cache coherency in this state.
- **C1(Auto Halt):** The first level of power reduction occurs when the processor executes an Auto Halt instruction. This stops the execution of the instruction stream and greatly reduces the processors power consumption. The processor can service snoops and maintain cache coherency in this state.
- **C2 (Stop Grant):** The next level of power reduction occurs when the processor is placed into the Stop Grant state by the assertion of STPCLK#. The processor can service snoops and maintain cache coherency in this state.

C2 is entered when software reads the Level 2 Register in the MCH. This is an ACPI defined register but BIOS or APM (via BIOS) can use this facility when entering a low power state.

- **Throttling:** This function can be enabled or disabled via a configuration bit. When this function is enabled, STPCLK# is asserted to place the processor into the C2 state with a programmable duty cycle. This is an ACPI defined function but BIOS or APM (via BIOS) can use this facility.

### 5.4.2. Sleep State Control

- **S0 (Awake):** In this state all power planes are active. All of the ACPI software “C” states are embedded in this state.
- **S1:** The recommended implementation of S1 state is the same as C2 state (Stop Grant), which is entered by the assertion of STPCLK# signal from the I/O Controller Hub to the processor. A further power saving can be achieved by asserting processor SLP# from the I/O Controller Hub. This puts the processor into Sleep State.
- **S2:** ACPI S2 state is not supported in the Intel® 820/820E chipset desktop platform.
- **S3 (Suspend To RAM (STR)):** The next level of power reduction occurs when the clock synthesizers and main power planes (I/O Controller Hub 3.3V plan, MCH 1.8V plan, and the processor) are shut down but the RDRAM memory plane and the I/O Controller Hub resume well remain active. This is the Suspend To RAM (STR) state. All clocks from synthesizers are shut down during the S3 state.
- **S4 and S5 (Suspend To Disk (STD), Soft Off):** The next level of power reduction occurs when the memory power and MCH 1.8V plane are shut down in addition to the clock synthesizers, I/O Controller Hub 3.3V, and the processor power planes. The I/O Controller Hub resume well is still powered.
- **G3 (Mechanical Off):** In this state only the RTC well is powered. The system can only reactivate when the power switch is returned to the On position.


## 5.5. MCH Clocking

The MCH has two clock input pins: HCLKIN for the host clock and GCLKIN for the AGP clock. Two clocks outputs are also implemented in the MCH. RCLKOUT and HCLKOUT provides a feedback to Direct Rambus channel for phase alignment through the Direct Rambus Clock Generator (DRCG). CK133 is a 56-pin SSOP package that generates:

- 4 copies of 2.5V front side bus clock for the processor, and ITP connector
- 3 copies of 2.5V APIC clocks
- 4 copies of 3.3V 66 MHz clocks for AGP, MCH, and the I/O Controller Hub
- 8 copies of 3.3V 33 MHz for PCI, I/O Controller Hub, and LPC devices
- 2 copies of 3.3V 14.318 KHz for I/O Controller Hub, LPC super I/O, or possible ISA devices
- one 3.3V 48 MHz for I/O Controller Hub, or LPC super I/O
- 2 copies of reference clock (half of the front side bus clock) for DRCG

DRCG is a 24-pin SSOP package. DRCG generates the differential clocks for Direct Rambus channel. With 100 MHz host bus clock, DRCG supports 300 or 400 MHz clock speed for the Direct Rambus channel. The clock connections to the various blocks for MCH are shown in the following block diagram.


Figure 7. Intel® 82820 MCH Clocking Diagram



## 5.6. MCH System Reset and Power Sequencing

The reset scheme for the MCH is shown in Figure 8. After PWROK is asserted to indicate the system power is stable, RSTIN# is generated by the I/O Controller Hub and is used as an input to reset the MCH. If RSTIN# is asserted, the MCH always asserts CPURST#. The assertion of CPURST# resets slot 1, as well as ITP. CPURST# is deasserted synchronous to the host bus clock. CPURST# is deasserted approximately 1msec (65536 66 MHz clocks) after detecting the rising edge of RSTIN#.

**Figure 8. Intel® 82820 MCH Reset Overview**



**Note:** This diagram does not represent all the details for schematics connection.

### 5.6.1. MCH Reset

A Full MCH Reset occurs when the power switch is turned on or when exiting a power managed state (Suspend to Disk) where the system power supplies are shut down. It is defined when RSTIN# is asserted. All internal state machines are reset and all registers assume their default values when this reset occurs.


### 5.6.2. MCH Power Sequence Recommendation

For the consideration of component long-term reliability, the following power sequence is strongly recommended while the AGP interface of MCH is running at 3.3V. If the AGP interface is running at 1.5V, the following power sequence recommendation is no longer applicable. The power sequence recommendations are:

- During the power-up sequence, the 1.8V must ramp up to 1.0V BEFORE 3.3V ramps up to 2.2V.
- During the power-down sequence, the 1.8V CAN NOT ramp below 1.0V BEFORE 3.3V ramps below 2.2V.
- The same power sequence recommendation also applies to the entrance and exit of S3 state, since MCH power is complete off during the S3 state.

The following diagram is shown only to illustrate the concept of power ramp sequence requirement between 3.3V and 1.8V. System designers need to be aware of this requirement while designing the voltage regulators and selecting the power supply. For further details on the voltage regulator design, refer to the appropriate chipset design guide (*Intel® 820 Chipset Design Guide* or *Intel® 820E Chipset Design Guide*).

**Figure 9. MCH Power-Up Voltage Sequencing**



This page is intentionally left blank.



## 6. *Pinout and Package Information*

---

### 6.1. Pinout Information

This section lists the MCH ballout assignment. The following two figures show the footprint ballout assignments from a top view of the package. The following table lists the ballout assignment in alphabetical order by signal name.

Figure 10. MCH Ballout (Top View—Left Side)

	1	2	3	4	5	6	7	8	9	10
A	VSS	HCLKOUT	VSS	DQB7	DQB5	DQB3	RQ0	RQ2	RQ5	RQ6
B	RCLKOUT	SCK	CMD	VSS	DQB6	VSS	DQB1	VSS	RQ4	VSS
C	RS1#	SIO	GTLREFA	DQB8	DQB4	DQB2	DQB0	RQ1	RQ3	RQ7
D	DBSY#	ADS#	HIT#	VCC1_8	VSS	VSS	VSS	VSS	VSS	VCC1_8
E	HITM#	RS2#	HREQ3#	HREQ2#	RS0#	VCC1_8	VCC1_8	VCC1_8	VCC1_8	RAMREFB
F	DRDY#	DEFER#	HLOCK#	HTRDY#	BNR#	VCC1_8				
G	BPRI#	HREQ4#	VSS	HREQ1#	HA6#	VCC1_8				
H	HREQ0#	HA9#	HA4#	HA5#	HA8#					
J	HA3#	HA14#	VSS	HA10#	HA12#					
K	HA13#	HA7#	HA16#	HA18#	HA15#					
L	HA11#	HA19#	VSS	HA17#	VSS					
M	HA23#	HA22#	HA21#	HA25#	HA31#					
N	HA27#	HA20#	VSS	HA30#	HA24#					
P	HA26#	HA28#	HA29#	CPURST#	HD3#					
R	HD1#	HD0#	VSS	HD2#	HD5#	VCC1_8	VCC1_8			
T	HD4#	HD10#	HD20#	HD23#	HD25#	HD31#	HD38#	HD36#	HD47#	HD49#
U	HD9#	HD15#	HD13#	HD19#	VSS	HD28#	HD33#	HD40#	HD44#	HD51#
V	HD6#	HCLKIN	VSS	HD18#	HD24#	HD29#	VSS	HD43#	VSS	HD52#
W	HD8#	HD12#	HD17#	HD16#	HD27#	HD30#	HD32#	HD37#	HD39#	HD41#
Y	HD14#	HD7#	HD11#	HD21#	HD22#	HD26#	HD35#	HD34#	HD45#	HD42#

VSS	VSS
VSS	VSS
VSS	VSS
VSS	VSS

Figure 11. MCH Ballout (Top View—Right Side)

	11	12	13	14	15	16	17	18	19	20	
	CTM#	CFM	DQA0	DQA2	DQA6	DQA8	VSS	HLCOMP	VSS	HL7	A
	CTM	CFM#	VSS	DQA4	VSS	VSS	VCC1_8	HL6	HL5	HL4	B
	VSS	VSS	DQA1	DQA3	DQA5	DQA7	VCC1_8	HL9	VCC1_8	HL_STB#	C
	VCC1_8	VSS	VSS	VSS	VSS	VSS	HL8	HL10	HL_STB	VSS	D
	RAMREFA	VCC1_8	VCC1_8	VCC1_8	VCC1_8	VSS	HL2	VSS	HL3	HUBREF	E
				VCC1_8	VDDQ	VSS	G_AD0	HL1	HL0	RSTIN#	F
						G_AD4	G_AD2	G_AD1	G_AD3	G_AD5	G
						G_C/BE0#	G_AD6	G_AD7	VSS	AD_STB0#	H
						G_AD9	VDDQ	G_AD12	AD_STB0	G_AD8	J
						G_PAR	G_AD10	G_AD11	VSS	G_AD14	K
						G_FRAME#	VDDQ	G_AD15	G_AD13	G_C/BE1#	L
						G_AD18	G_AD16	G_STOP#	VSS	G_TRDY#	M
						G_AD20	VDDQ	G_C/BE2#	G_DEVSEL#	G_IRDY#	N
					VCC1_8	G_AD22	G_AD19	G_AD17	VSS	G_AD21	P
					VDDQ	G_C/BE3#	G_AD25	AD_STB1	AD_STB1#	G_AD23	R
	HD54#	HD50#	VSS	VCC1_8	TEST/ GRCOMP	G_AD27	VDDQ	G_AD29	VSS	G_AD24	T
	HD59#	HD60#	HD55#	AGPREF	G_REQ#	SBA7	G_AD26	G_AD28	G_AD31	G_AD30	U
	VSS	GTLREFB	VSS	HD62#	WBF#	RBF#	SBA1	VSS	SBA6	SBA4	V
	HD46#	HD53#	HD63#	HD58#	ST0	PIPE#	SBA3	CLK66	SBA5	SBA0	W
	HD48#	HD57#	HD56#	HD61#	ST1	G_GNT#	ST2	SBA2	SB_STB#	SB_STB	Y

VSS	VSS
VSS	VSS
VSS	VSS
VSS	VSS

Table 35. MCH Alphabetical Ball List

Name	Ball #	Name	Ball #	Name	Ball #	Name	Ball #
AD_STB0	J19	DRDY#	F1	G_C/BE2#	N18	HA26#	P1
AD_STB0#	H20	G_AD0	F17	G_C/BE3#	R16	HA27#	N1
AD_STB1	R18	G_AD1	G18	G_DEVSEL#	N19	HA28#	P2
AD_STB1#	R19	G_AD2	G17	G_FRAME#	L16	HA29#	P3
ADS#	D2	G_AD3	G19	G_GNT#	Y16	HA30#	N4
AGPREF	U14	G_AD4	G16	G_IRDY#	N20	HA31#	M5
BNR#	F5	G_AD5	G20	G_PAR	K16	HCLKIN	V2
BPRI#	G1	G_AD6	H17	G_REQ#	U15	HCLKOUT	A2
CFM	A12	G_AD7	H18	G_STOP#	M18	HD0#	R2
CFM#	B12	G_AD8	J20	G_TRDY#	M20	HD1#	R1
CLK66	W18	G_AD9	J16	GTLREFA	C3	HD2#	R4
CMD	B3	G_AD10	K17	GTLREFB	V12	HD3#	P5
CPURST#	P4	G_AD11	K18	HA3#	J1	HD4#	T1
CTM	B11	G_AD12	J18	HA4#	H3	HD5#	R5
CTM#	A11	G_AD13	L19	HA5#	H4	HD6#	V1
DBSY#	D1	G_AD14	K20	HA6#	G5	HD7#	Y2
DEFER#	F2	G_AD15	L18	HA7#	K2	HD8#	W1
DQA0	A13	G_AD16	M17	HA8#	H5	HD9#	U1
DQA1	C13	G_AD17	P18	HA9#	H2	HD10#	T2
DQA2	A14	G_AD18	M16	HA10#	J4	HD11#	Y3
DQA3	C14	G_AD19	P17	HA11#	L1	HD12#	W2
DQA4	B14	G_AD20	N16	HA12#	J5	HD13#	U3
DQA5	C15	G_AD21	P20	HA13#	K1	HD14#	Y1
DQA6	A15	G_AD22	P16	HA14#	J2	HD15#	U2
DQA7	C16	G_AD23	R20	HA15#	K5	HD16#	W4
DQA8	A16	G_AD24	T20	HA16#	K3	HD17#	W3
DQB0	C7	G_AD25	R17	HA17#	L4	HD18#	V4
DQB1	B7	G_AD26	U17	HA18#	K4	HD19#	U4
DQB2	C6	G_AD27	T16	HA19#	L2	HD20#	T3
DQB3	A6	G_AD28	U18	HA20#	N2	HD21#	Y4
DQB4	C5	G_AD29	T18	HA21#	M3	HD22#	Y5
DQB5	A5	G_AD30	U20	HA22#	M2	HD23#	T4
DQB6	B5	G_AD31	U19	HA23#	M1	HD24#	V5
DQB7	A47	G_C/BE0#	H16	HA24#	N5	HD25#	T5
DQB8	C4	G_C/BE1#	L20	HA25#	M4	HD26#	Y6

Name	Ball #
HD27#	W5
HD28#	U6
HD29#	V6
HD30#	W6
HD31#	T6
HD32#	W7
HD33#	U7
HD34#	Y8
HD35#	Y7
HD36#	T8
HD37#	W8
HD38#	T7
HD39#	W9
HD40#	U8
HD41#	W10
HD42#	Y10
HD43#	V8
HD44#	U9
HD45#	Y9
HD46#	W11
HD47#	T9
HD48#	Y11
HD49#	T10
HD50#	T12
HD51#	U10
HD52#	V10
HD53#	W12
HD54#	T11
HD55#	U13
HD56#	Y13
HD57#	Y12
HD58#	W14
HD59#	U11
HD60#	U12
HD61#	Y14
HD62#	V14
HD63#	W13

Name	Ball #
HIT#	D3
HITM#	E1
HL_STB	D19
HL_STB#	C20
HL0	F19
HL1	F18
HL2	E17
HL3	E19
HL4	B20
HL5	B19
HL6	B18
HL7	A20
HL8	D17
HL9	C18
HL10	D18
HLCOMP	A18
HLOCK#	F3
HREQ0#	H1
HREQ1#	G4
HREQ2#	E4
HREQ3#	E3
HREQ4#	G2
HTRDY#	F4
HUBREF	E20
PIPE#	W16
RAMREFA	E11
RAMREFB	E10
RBF#	V16
RCLKOUT	B1
RQ0	A7
RQ1	C8
RQ2	A8
RQ3	C9
RQ4	B9
RQ5	A9
RQ6	A10
RQ7	C10

Name	Ball #
RS0#	E5
RS1#	C1
RS2#	E2
RSTIN#	F20
SB_STB	Y20
SB_STB#	Y19
SBA0	W20
SBA1	V17
SBA2	Y18
SBA3	W17
SBA4	V20
SBA5	W19
SBA6	V19
SBA7	U16
SCK	B2
SIO	C2
ST0	W15
ST1	Y15
ST2	Y17
TEST/ GRCOMP	T15
VCC1_8	B17
VCC1_8	C17
VCC1_8	C19
VCC1_8	D4
VCC1_8	D10
VCC1_8	D11
VCC1_8	E6
VCC1_8	E7
VCC1_8	E8
VCC1_8	E9
VCC1_8	E12
VCC1_8	E13
VCC1_8	E14
VCC1_8	E15
VCC1_8	F6
VCC1_8	F14

Name	Ball #
VCC1_8	G6
VCC1_8	P15
VCC1_8	R6
VCC1_8	R7
VCC1_8	T14
VDDQ	F15
VDDQ	J17
VDDQ	L17
VDDQ	N17
VDDQ	R15
VDDQ	T17
VSS	A1
VSS	A3
VSS	A17
VSS	A19
VSS	B4
VSS	B6
VSS	B8
VSS	B10
VSS	B13
VSS	B15
VSS	B16
VSS	C11
VSS	C12
VSS	D5
VSS	D6
VSS	D7
VSS	D8
VSS	D9
VSS	D12
VSS	D13
VSS	D14
VSS	D15
VSS	D16
VSS	D20
VSS	E16
VSS	E18



Name	Ball #
VSS	F16
VSS	G3
VSS	H19
VSS	J3
VSS	J9
VSS	J10
VSS	J11
VSS	J12
VSS	K9
VSS	K10
VSS	K11

Name	Ball #
VSS	K12
VSS	K19
VSS	L3
VSS	L5
VSS	L9
VSS	L10
VSS	L11
VSS	L12
VSS	M9
VSS	M10
VSS	M11

Name	Ball #
VSS	M12
VSS	M19
VSS	N3
VSS	P19
VSS	R3
VSS	T13
VSS	T19
VSS	U5
VSS	V3
VSS	V7
VSS	V9

Name	Ball #
VSS	V11
VSS	V13
VSS	V18
WBF#	V15

## 6.2. Package Specifications

This section shows the mechanical dimensions for the MCH. The package is a 324 Ball Grid Array (BGA).

Figure 12. Package Dimensions (324 BGA) – Top and Side Views




Figure 13. Package Dimensions (324 BGA) – Bottom View




Table 36. BGA Package Dimensions (324 BGA)<sup>1,2,3</sup>

Symbol	Min	Nominal	Max	Units	Note
A	2.17	2.38	2.59	mm	
A1	0.50	0.60	0.70	mm	
A2	1.12	1.17	1.22	mm	
D	26.80	27.00	27.20	mm	
D1	23.80	24.00	24.70	mm	
E	26.90	27.00	27.10	mm	
E1	23.80	24.00	24.20	mm	
e	1.27 (solder ball pitch)			mm	
l	1.44 REF.			mm	
J	1.44 REF.			mm	
M	20 x 20 Matrix			mm	
b <sup>2</sup>	0.60	0.75	0.90	mm	
c	0.55	0.61	0.67	mm	

**NOTES:**

1. All dimensions and tolerances conform to ANSI Y14.5-1982
2. Dimension is measured at maximum solder ball diameter parallel to primary datum (-C-)
3. Primary Datum (-C-) and seating plane are defined by the spherical crowns of the solder balls.



### 6.3. MCH RSL Package Dimensions

The package dimensions in Table 37 represent the trace length of the RSL signals in the MCH package. These dimensions must be considered when matching trace lengths. Note that these dimensions are normalized to 0 with RQ7 (the shortest trace on the MCH substrate). They do not represent actual lengths from pad to ball.

**Table 37: MCH RSL Package Dimensions**

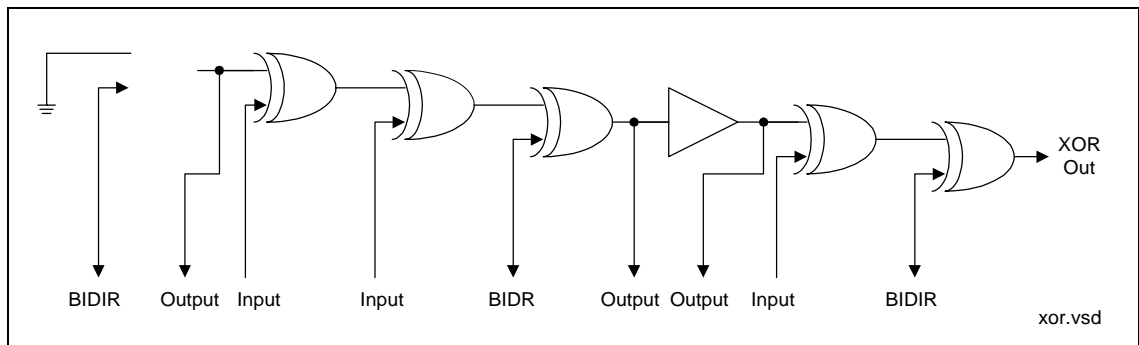
Signal	Ball on MCH	Package Dimension* (mils)	Signal	Ball on MCH	Package Dimension* (mils)
DQA0	A13	138.14	DQB7	A4	237.71
DQA1	C13	19.11	DQB8	C4	138.29
DQA2	A14	163.16	RQ0	A7	179.49
DQA3	C14	39.87	RQ1	C8	27.12
DQA4	B14	97.54	RQ2	A8	162.21
DQA5	C15	62.67	RQ3	C9	5.80
DQA6	A15	186.11	RQ4	B9	71.70
DQA7	C16	95.70	RQ5	A9	133.88
DQA8	A16	230.20	RQ6	A10	122.20
DQB0	C7	39.56	RQ7	C10	0.00
DQB1	B7	95.83			
DQB2	C6	63.49			
DQB3	A6	153.69	CFM	A12	132.37
DQB4	C5	97.33	CFM#	B12	64.63
DQB5	A5	191.43	CTM	B11	56.06
DQB6	B5	152.47	CTM#	A11	126.34

This page is intentionally left blank.

## 7. Testability

For Automated Test Equipment (ATE) board level testing, the MCH uses the new XOR chain for testability.

An XOR-Tree is a chain of XOR gates, with each having one input pin or one bi-directional pin (used as an input pin only) connected to it as shown in the following figure. The output pins (non-bussed pins) should not have additional inversion as shown by the non-inverting buffer between two successive output pins. The first XOR gate should have one pin connected to ground.



The algorithm used for in-circuit test is as follows:

- Drive “0” on to all input and bi-directional pins. This, along with the first XOR gate having one input connected to GND and the Outputs being non-inverting, would consistently produce a “0” at “XOR out”, regardless of how many XOR gates are in the chain.
- Drive each input or bi-directional pin one at a time, first to a 1 and then back to a 0. This causes “XOR out” first to produce a 1 and then a 0.
- The output pins can now be checked by driving a 1 on to the Pin of the XOR gate that has its second input to the XOR gate connected to VCC and then back to a 0. The Output pins will go to a 1 and then back to a 0.

The above algorithm is for all pins properly soldered to the board under test and no pins connected to a power plane. If there is an odd number of total signal pins connected to a power plane or unsoldered and floating to a 1 logic level, The following would happen:


- Drive 0 on to all other input and bi-directional pins. “XOR out” would start out at a 1 level instead of a 0.
- Drive each input or bi-directional pin, one at a time, first to a 1 and then back to a 0. This causes “XOR out” first to produce a 0 and then a 1.

A flexible test model used for in-circuit test needs to determine if “XOR out” initially is a 0 or 1 level. One at a time, each input or bi-directional pin is driven to 1 and then back to a 0. The “XOR out” will first toggle to the opposite state that was initially determined and then toggle back to the initially determined state. Any unsoldered pin will cause “XOR out” not to toggle.

### Initialization Sequence

Five pins need to be controlled to enable XOR chain and “tri-state all pins” test modes as shown in Figure 14.


Figure 14. Initialization Sequence



### Rule for XOR Testing of Complementary STB Pairs

Each complementary strobe pair must have complementary inputs at all the times during XOR chain testing. This adds a constraint that XOR chain #3 and chain #4 cannot be run in parallel. All the affected complementary strobe pairs are highlighted with the note “Dependency on complement signal in Chain #x”, chain #3, and chain #4.

Figure 15. Rule for Controlling STB Signal Pairs During XOR Testing



XOR chain pin assignments are shown in Table 38 through Table 42.

**Table 38. XOR Chain #1 Connections**

NAME	BALL	Chain Element #	Note
CMD	B3	1	Input
SCK	B2	2	
SIO	C2	3	
HCLKOUT	A2	4	
RCLKOUT	B1	5	
HIT#	D3	6	
RS0#	E5	7	
RS1#	C1	8	
ADS#	D2	9	
BNR#	F5	10	
HREQ2#	E4	11	
HREQ1#	G4	12	
HREQ3#	E3	13	
HTRDY#	F4	14	
DBSY#	D1	15	
RS2#	E2	16	
HLOCK#	F3	17	
HA6#	G5	18	
HITM#	E1	19	
DEFER#	F2	20	
HA8#	H5	21	
HA10#	J4	22	
HA12#	J5	23	
DRDY#	F1	24	
HREQ4#	G2	25	
HA4#	H3	26	
HA5#	H4	27	
HA9#	H2	28	
BPRI#	G1	29	
HREQ0#	H1	30	
HA14#	J2	31	
HA3#	J1	32	
HA15#	K5	33	

NAME	BALL	Chain Element #	Note
HA13#	K1	34	
HA18#	K4	35	
HA16#	K3	36	
HA7#	K2	37	
HA11#	L1	38	
HA19#	L2	39	
HA17#	L4	40	
HA31#	M5	41	
HA25#	M4	42	
HA23#	M1	43	
HA22#	M2	44	
HA27#	N1	45	
HA26#	P1	46	
HD1#	R1	47	
HA21#	M3	48	
HA20#	N2	49	
HA28#	P2	50	
HD0#	R2	51	
HD4#	T1	52	
HD10#	T2	53	
HA30#	N4	54	
HA29#	P3	55	
HD9#	U1	56	
CPURST#	P4	57	
HA24#	N5	58	
HD20#	T3	59	
SBA0	W20	60	Output

Table 39. XOR Chain #2 Connections

NAME	BALL	Chain Element #	Note
HD13#	U3	1	Input
HD15#	U2	2	
HD6#	V1	3	
HD23#	T4	4	
HD2#	R4	5	
HD12#	W2	6	
HD19#	U4	7	
HD14#	Y1	8	
HD8#	W1	9	
HD5#	R5	10	
HD25#	T5	11	
HD3#	P5	12	
HD7#	Y2	13	
HD17#	W3	14	
HD18#	V4	15	
HD31#	T6	16	
HD16#	W4	17	
HD11#	Y3	18	
HD28#	U6	19	
HD21#	Y4	20	
HD24#	V5	21	
HD38#	T7	22	
HD27#	W5	23	
HD36#	T8	24	
HD33#	U7	25	
HD22#	Y5	26	
HD29#	V6	27	
HD30#	W6	28	
HD26#	Y6	29	
HD47#	T9	30	
HD40#	U8	31	
HD43#	V8	32	
HD32#	W7	33	
HD35#	Y7	34	

NAME	BALL	Chain Element #	Note
HD37#	W8	35	
HD44#	U9	36	
HD34#	Y8	37	
HD39#	W9	38	
HD45#	Y9	39	
HD41#	W10	40	
HD52#	V10	41	
HD42#	Y10	42	
HD48#	Y11	43	
HD46#	W11	44	
HD57#	Y12	45	
HD53#	W12	46	
HD59#	U11	47	
HD63#	W13	48	
HD49#	T10	49	
HD56#	Y13	50	
HD61#	Y14	51	
HD51#	U10	52	
HD60#	U12	53	
HD50#	T12	54	
HD54#	T11	55	
HD55#	U13	56	
HD58#	W14	57	
HD62#	V14	58	
SBA1	V17	59	Output

**Table 40. XOR Chain #3 Connections**

NAME	BALL	Chain Element #	Note
ST0	W15	1	Input
ST1	Y15	2	
WBF#	V15	3	
G_GNT#	Y16	4	
RBF#	V16	5	
PIPE#	W16	6	
G_REQ#	U15	7	
ST2	Y17	8	
SBA7	U16	9	
CLK66	W18	10	Dependency on complement signal in Chain #4
SBA5	W19	11	
SBA6	V19	12	
G_C/BE3#	R16	13	
G_AD22	P16	14	
G_AD27	T16	15	
G_AD28	U18	16	
G_AD26	U17	17	
G_AD31	U19	18	
G_AD25	R17	19	
G_AD30	U20	20	
G_AD29	T18	21	
AD_STB1	R18	22	Dependency on complement signal in Chain #4
G_AD19	P17	23	
G_AD20	N16	24	
G_AD17	P18	25	
G_AD24	T20	26	
G_AD23	R20	27	
G_AD21	P20	28	
G_AD18	M16	29	
G_AD16	M17	30	
G_C/BE2#	N18	31	
G_DEVSEL#	N19	32	

NAME	BALL	Chain Element #	Note
G_IRDY#	N20	33	
G_STOP#	M18	34	
G_FRAME#	L16	35	
G_TRDY#	M20	36	
G_PAR	K16	37	
G_AD15	L18	38	
G_AD13	L19	39	
G_C/BE1#	L20	40	
G_AD10	K17	41	
G_AD11	K18	42	
G_AD9	J16	43	
G_AD14	K20	44	
G_AD8	J20	45	
G_AD12	J18	46	
AD_STB0	J19	47	Dependency on complement signal in Chain #4
G_AD7	H18	48	
G_AD6	H17	49	
G_AD5	G20	50	
G_AD3	G19	51	
G_AD1	G18	52	
G_AD2	G17	53	
G_C/BE0#	H16	54	
HL_STB	D19	55	Dependency on complement signal in Chain #4
SB_STB#	Y19	56	Output

Table 41. XOR Chain #4 Connections

NAME	BALL	Chain Element #	Note
SBA4	V20	1	Input, Dependency on complement signal in Chain #3
AD_STB1#	R19	2	Dependency on complement signal in Chain #3
AD_STB0#	H20	3	Dependency on complement signal in Chain #3
G_AD4	G16	4	
G_AD0	F17	5	
HL0	F19	6	
HL1	F18	7	
HL2	E17	8	
HL3	E19	9	
HL8	D17	10	
HL10	D18	11	
HL_STB#	C20	12	Dependency on complement signal in Chain #3
HL9	C18	13	
HL4	B20	14	
HL5	B19	15	
HL6	B18	16	
HL7	A20	17	
HLCOMP	A18	18	
SBA2	Y18	19	Output

Table 42. XOR Chain #5 Connections

NAME	BALL	Chain Element #	Note
DQA7	C16	1	Input
DQA8	A16	2	
DQA5	C15	3	
DQA6	A15	4	
DQA3	C14	5	
DQA4	B14	6	
DQA2	A14	7	
DQA1	C13	8	
DQA0	A13	9	
CFM	A12	10	
CFM#	B12	11	
CTM#	A11	12	
CTM	B11	13	
RQ6	A10	14	
RQ7	C10	15	
RQ5	A9	16	
RQ4	B9	17	
RQ3	C9	18	
RQ2	A8	19	
RQ1	C8	20	
RQ0	A7	21	
DQB1	B7	22	
DQB0	C7	23	
DQB3	A6	24	
DQB2	C6	25	
DQB5	A5	26	
DQB6	B5	27	
DQB4	C5	28	
DQB7	A4	29	
DQB8	C4	30	
SBA_4	V20	31	Output