

AMD-8131™ HyperTransport™ PCI-X® Tunnel Data Sheet

1 Overview

The AMD-8131™ HyperTransport™ PCI-X® Tunnel (referred to as *the IC* in this document) is a HyperTransport™ technology (referred to as *link* in this document) tunnel developed by AMD that provides two PCI-X bridges.

1.1 Device Features

- HyperTransport technology tunnel with side A and side B.
 - Side A is 16 bits (input and output); side B is 8 bits.
 - Either side may connect to the host or to a downstream HyperTransport technology compliant device.
 - Each side supports HyperTransport technology-defined reduced bit widths: 8-bit, 4-bit, and 2-bit.
 - Each side supports transfer rates of 1600, 1200, 800, and 400 mega-transfers per second.
 - Maximum bandwidth is 6.4 gigabytes per second across side A (half upstream and half downstream) and 3.2 gigabytes per second across side B.
 - Independent transfer rate and bit width selection for each side.
 - Link disconnect protocol supported.
- Two PCI-X (rev. 1.0) bridges: bridge A and bridge B.
 - Each bridge supports a 64-bit data bus.
 - Each bridge supports operational modes of PCI-X and legacy PCI revision 2.2 protocol.
 - Bridges support 133, 100, and 66 MHz transfer rates in PCI-X mode.
 - Bridges support 66 and 33 MHz transfer rates in PCI mode.
 - Independent transfer rates and operational modes for each bridge.
 - Each bridge includes support for up to 5 PCI masters with clock, request, and grant signals.
 - Each bridge includes an IOAPIC with four redirection registers. Legacy interrupt controller and IOAPIC modes supported.
 - SHPC-compliant hot plug controller and support.
- 37.5 x 37.5 millimeter, 829-pin BGA package.
- 3.3 volt PCI-X signaling; 1.2 volt link signaling; 1.8 volt core.

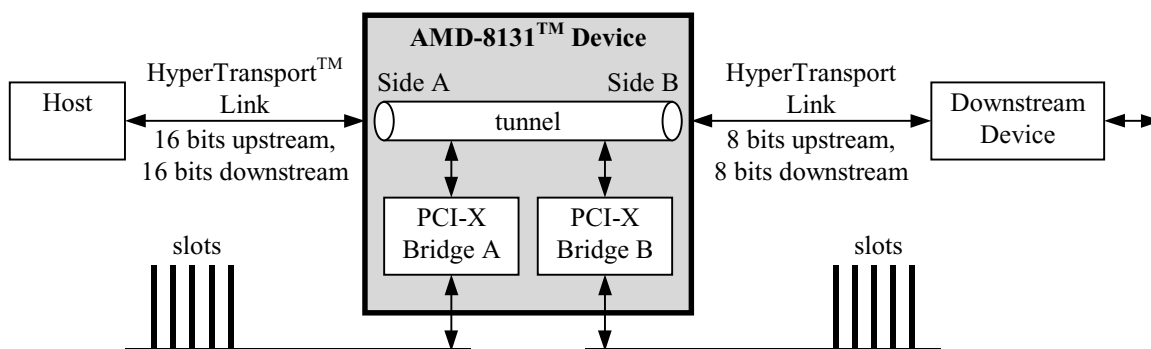


Figure 1: System block diagram.

© 2004 Advanced Micro Devices, Inc.

All rights reserved. The contents of this document are provided in connection with Advanced Micro Devices, Inc. ("AMD") products. AMD makes no representations or warranties with respect to the accuracy or completeness of the contents of this publication and reserves the right to make changes to specifications and product descriptions at any time without notice. No license, whether express, implied, arising by estoppel or otherwise, to any intellectual property rights is granted by this publication. Except as set forth in AMD's Standard Terms and Conditions of Sale, AMD assumes no liability whatsoever, and disclaims any express or implied warranty, relating to its products including, but not limited to, the implied warranty of merchantability, fitness for a particular purpose, or infringement of any intellectual property right. AMD's products are not designed, intended, authorized or warranted for use as components in systems intended for surgical implant into the body, or in other applications intended to support or sustain life, or in any other application in which the failure of AMD's product could create a situation where personal injury, death, or severe property or environmental damage may occur. AMD reserves the right to discontinue or make changes to its products at any time without notice.

Trademarks

AMD, the AMD Arrow logo, and combinations thereof, and AMD-8131 are trademarks of Advanced Micro Devices, Inc.

HyperTransport is a licensed trademark of the HyperTransport Technology Consortium.

PCI-X is a registered trademark of the PCI-Special Interest Group (PCI-SIG).

Other product names used in this publication are for identification purposes only and may be trademarks of their respective companies.

Table of Contents

1	Overview	1
1.1	Device Features	1
2	Ordering Information	7
3	Signal Descriptions	7
3.1	Terminology	7
3.2	Tunnel Link Signals	8
3.3	PCI-X® Signals	9
3.4	Test and Miscellaneous Signals	12
3.5	Power and Ground	14
3.5.1	Power Plane Sequencing	14
4	Functional Operation	14
4.1	Overview	14
4.2	Reset And Initialization	14
4.2.1	Non-Hot Plug Initialization	15
4.2.2	Hot Plug Initialization	15
4.3	Clocking	16
4.3.1	Systemboard Requirements	16
4.3.2	Characterization	17
4.3.3	Clock Gating	18
4.4	Tunnel Links	19
4.4.1	Link PHY	19
4.5	PCI-X® Bridges	19
4.5.1	Tags, UnitIDs, SeqIDs And Ordering	19
4.5.2	Interrupt Controllers	19
4.5.2.1	Error NMI Interrupts	21
4.5.3	Hot Plug	21
4.5.3.1	Multi-slot Hot Plug	21
4.5.3.2	Single-Slot Hot Plug	24
4.5.3.3	Serial Interface	27
4.5.3.3.1	Serial Data From The Power Controllers To The IC	28
4.5.3.3.2	Serial Data From The IC To The Power Controllers	29
4.5.3.4	SHPC Interrupts, Events, And Errors	30
4.5.3.5	Reset To Hot Plug Slots	30
4.5.4	PCI-X, PHY Compensation Update	30
4.5.5	Transactions Claimed By The Bridges	31
4.5.6	Various Behaviors	31
4.5.7	Error Conditions And Handling	33
4.6	Performance-Related Information	37
4.6.1	Bandwidth Percentage	37
4.6.2	Latency	39
5	Registers	40
5.1	Register Overview	40
5.1.1	Configuration Space	40
5.1.2	Register Naming and Description Conventions	40

5.2	PCI-X, Bridge Configuration Registers	41
5.3	PCI-X IOAPIC Configuration Registers	63
5.4	IOAPIC Registers	65
5.5	SHPC Working Registers	67
6	Electrical Data	74
6.1	Absolute Ratings	74
6.2	Operating Ranges	74
6.3	DC Characteristics	75
6.4	AC Characteristics	76
7	Ball Designations	77
8	Package Specification	82
9	Test	83
9.1	High Impedance Mode	83
9.2	NAND Tree Mode	83
10	Appendix	86
10.1	Revision History	86

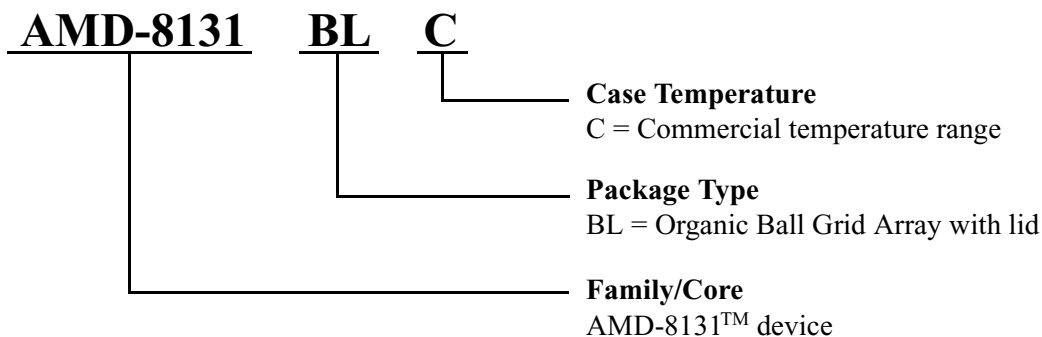
List of Figures

Figure 1:	System block diagram.....	1
Figure 2:	Systemboard clocking.....	17
Figure 3:	Correction for characterization.....	18
Figure 4:	System diagram for multiple hot plug slots on a bridge.....	22
Figure 5:	System diagram of PME# signals.....	23
Figure 6:	System diagram of M66EN signals.....	23
Figure 7:	Multi-slot hot plug enable/disable sequence.....	24
Figure 8:	Single-slot hot plug system diagram.....	25
Figure 9:	Single-slot hot plug enable/disable sequence.....	26
Figure 10:	Single-slot hot plug M66EN connections.....	26
Figure 11:	Hot plug serial interface connections.....	27
Figure 12:	Configuration space.....	40
Figure 13:	Ball designations.....	77
Figure 14:	Package mechanical drawing.....	82
Figure 15:	NAND tree.....	83

List of Tables

Table 1:	IO signal types.	7
Table 2:	Signal isolation groups.	22
Table 3:	Channel 00b, interrupt capable serial hot plug data to the IC.	28
Table 4:	Channel 01b, non-interrupt capable serial hot plug data to the IC.	29
Table 5:	Serial hot plug data from the IC to the power controller.	29
Table 6:	Bandwidth percentages.	38
Table 7:	Some latencies.	39
Table 8:	Configuration spaces.	41
Table 9:	Memory mapped address spaces.	41
Table 10:	Register attributes.	41
Table 11:	Absolute maximum ratings.	74
Table 12:	Operating ranges.	74
Table 13:	Current and power consumption.	75
Table 14:	DC characteristics for signals on the VDD33 power plane.	75
Table 15:	AC requirements for REFCLK.	76
Table 16:	AC data for PCI clocks.	76
Table 17:	AC data for PCI bus.	76
Table 18:	Alphabetical listing of signals A_ACK64# to B_PRESET#.	78
Table 19:	Alphabetical listing of signals B_REQ0# to VDD18.	79
Table 20:	Alphabetical listing of signals VDD18 to VSS.	80
Table 21:	Test modes.	83

2 Ordering Information



3 Signal Descriptions

3.1 Terminology

See section 5.1.2 for a description of the register naming convention used in this document. See the AMD-8131™ HyperTransport™ PCI-X® Tunnel Design Guide for additional information.

Signals with a # suffix are active low.

Signals described in this chapter utilize the following IO cell types:

Name	Notes
Input	Input signal only.
Output	Output signal only. This includes outputs that are capable of being in the high-impedance state.
OD	Open drain output. These signals are driven low and expected to be pulled high by external circuitry.
IO	Input or output signal.
IOD	Input or open-drain output.
Analog	Analog signal.
w/PU	With pullup. The signal includes a pullup resistor to the signal's power plane. The resistor value is nominally 8K ohms.

Table 1: IO signal types.

The following provides definitions and reference data about each of the IC pins. “During Reset” provides the state of the pin while RESET# is asserted. “After Reset” provides the state of the pin immediately after RESET# is deasserted. “Func.” means that the pin is functional and operating per its defined function.

3.2 Tunnel Link Signals

The following are signals associated with the HyperTransport links. [B, A] in the signal names below refer to the A and B sides of the tunnel. [P, N] are the positive and negative sides of differential pairs.

Pin name and description	IO cell type	Power plane*	During reset	After reset												
LDTCOMP[3:0] . Link compensation pins for both sides of the tunnel. These are designed to be connected through resistors as follows: <table border="0"> <thead> <tr> <th><u>Bit</u></th> <th><u>Function</u></th> <th><u>External Connection</u></th> </tr> </thead> <tbody> <tr> <td>[0]</td> <td>Positive receive compensation</td> <td>Resistor to VDD12B</td> </tr> <tr> <td>[1]</td> <td>Negative receive compensation</td> <td>Resistor to VSS</td> </tr> <tr> <td>[3, 2]</td> <td>Transmit compensation</td> <td>Resistor from bit [2] to bit [3]</td> </tr> </tbody> </table> These resistors are used by the compensation circuit. The output of this circuit is combined with DevA:0x[E8, E4, E0] to determine compensation values that are passed to the link PHYs.	<u>Bit</u>	<u>Function</u>	<u>External Connection</u>	[0]	Positive receive compensation	Resistor to VDD12B	[1]	Negative receive compensation	Resistor to VSS	[3, 2]	Transmit compensation	Resistor from bit [2] to bit [3]	Analog	VDD-12A		
<u>Bit</u>	<u>Function</u>	<u>External Connection</u>														
[0]	Positive receive compensation	Resistor to VDD12B														
[1]	Negative receive compensation	Resistor to VSS														
[3, 2]	Transmit compensation	Resistor from bit [2] to bit [3]														
LRACAD_[P, N][15:0]; LRBCAD_[P, N][7:0] . Receive link command-address-data bus.	Link input	VDD12														
LRACLK[1, 0]_[P, N]; LRCLK0_[P, N] . Receive link clock.	Link input	VDD12														
LR[B, A]CTL_[P, N] . Receive link control signal.	Link input	VDD12														
LTACAD_[P, N][15:0]; LTBCAD_[P, N][7:0] . Transmit link command-address-data bus.	Link output	VDD12	Diff High**	Func.												
LTACLK[1, 0]_[P, N]; LTCLK0_[P, N] . Transmit link clock.	Link output	VDD12	Func.	Func.												
LT[B, A]CTL_[P, N] . Transmit link control signal.	Link output	VDD12	Diff Low**	Func.												

* The signals connected to the A side of the tunnel are powered by VDD12A and the signals connected to the B side of the tunnel are powered by VDD12B.

** Diff High and Diff Low for these link pins specify differential high and low; e.g., Diff High specifies that the _P signal is high and the _N signal is low.

If one of the sides of the tunnel is not used on a platform, then the unconnected link should be treated as follows, for every 10 differential pairs: connect all of the _P differential inputs together and through a resistor to VSS; connect all the _N differential inputs together and through a resistor to VDD12; leave the differential outputs unconnected. If there are unused link signals on an active link (because the IC is connected to a device with a reduced bit width), then the unused differential inputs and outputs should also be connected in this way.

3.3 PCI-X® Signals

[B, A] in the following signal names is used to differentiate between PCI-X bridge A and PCI-X bridge B. “Normal” refers to non-hot plug mode (DevA:0x48[HPENB, HPENA]) or hot plug mode with external isolation switches (Dev[B, A]:0x40[HPSSS#]); “Single slot HP” refers to hot plug mode while Dev[B, A]:0x40[HPSSS#] is low. When in “normal” mode, PCI-X signals that typically connect to the slot are placed into the “During reset” state when [B, A]_PRESET# is asserted; the other signals are placed into the “During reset” state only when RESET# is asserted. The “Single slot HP”, “During reset” column refers to while RESET# is asserted for all the signals.

Pin name and description	IO cell type	Power plane	Normal		Single slot HP	
			During reset	After reset	During reset	After reset
[B, A]_ACK64#. PCI-X® acknowledge for 64-bit transfers.	IO	VDD33	3-State	3-State	Low	Low
[B, A]_AD[63:0]. PCI-X® address-data bus.	IO	VDD33	Low	Parked	Low	Low
[B, A]_CBE_L[7:0]. PCI-X® command-byte enable bus.	IO	VDD33	Low	Parked	Low	Low
A_COMPAT. Strapping option to specify if PCI-X bridge A is the default bus in the system. This may only be associated with PCI-X® bridge A. See also DevA:0x48[COMPAT].	Input	VDD33				
[B, A]_DEVSEL#. PCI-X device select signal. During reset, these signals may be 3-state or they may be driven, based on the requirements of the PCI-X initialization pattern.	IO	VDD33	See left	3-State	Low	Low
[B, A]_FRAME#. PCI-X® frame signal.	IO	VDD33	3-State	3-State	Low	Low
[B, A]_GNT[4:0]#. PCI-X® master grant signals. Some of these signals are an input while PWROK is deasserted; all other times, they are outputs. [B, A]_GNT[4:3]# each require a strapping resistor to help specify the bridge operating frequency after a PWROK reset; see section 4.2. Note: A_GNT[2:1] require weak pull-up resistors to VDD33. B_GNT2# is an input while PWROK is low and an output at all other times. As an input, it is used to specify the default state of DevB:0x40[HPSSS#]. HPSSS# specifies if the IC supports a single hot plug slot on the bridge without external isolation switches. A weak resistor should be tied from this signal to VDD33 or to ground. [B, A]_HPSORLC. Hot plug serial output reset latch clock output (alternate functions to [B, A]_GNT4# selected by hot plug mode, DevA:0x48[HPENA, HPENB]). These are used in support of the hot plug serial interface. See section 4.5.3.3 for details.	IO (See left)	VDD33	High	High	[0] is Low; [4:1] are high	[0] is Low; [4:1] are high
[B, A]_IRDY#. PCI-X® master ready signal.	IO	VDD33	3-State	3-State	Low	Low

Pin name and description	IO cell type	Power plane	Normal		Single slot HP	
			During reset	After reset	During reset	After reset
[B, A]_M66EN. Frequency select input for [B, A]_PCLK while in conventional PCI mode. When not in hot plug mode, the state of this signal is captured at the rising edge [B, A]_PRESET# and (see section 4.2.1). After the corresponding [B, A]_PRESET# signal goes high, the state of [B, A]_M66EN is ignored. In hot plug mode, this signal may be driven low as an output after initialization.	IOD	VDD33	3-State	3-State	Low	Low
[B, A]_PAR. PCI-X® parity signal.	IO	VDD33	Low	Func.	Low	Low
[B, A]_PAR64. PCI-X® upper 32-bit parity signal for 64-bit transfers.	IO	VDD33	Low	Func.	Low	Low
[B, A]_PCIXCAP. PCI-X® frequency capabilities selection. The state of this signal is captured at the rising edge [B, A]_PRESET# and used to determine the bus mode (see section 4.2.1). After the corresponding [B, A]_PRESET# signal goes high, the state of [B, A]_PCIXCAP is ignored.	Input	VDD33				
[B, A]_PCLK[4:0]. Up to 133 MHz PCI-X® clock outputs. [B, A]_HPSID. Hot plug serial input data (alternative function to [B, A]_PCLK4# selected by hot plug mode, DevA:0x48[HPENA, HPENB]). See section 4.5.3.3 for details.	IO (See left)	VDD33	Func.	Func.	[4:1]: Func. [0]: Low	[4:1]: Func. [0]: Low
[B, A]_PERR#. PCI-X® parity error. This signal is only applicable to parity errors on the secondary PCI-X® bus interface.	IO	VDD33	3-State	3-State	Low	Low
[B, A]_PIRQ[A, B, C, D]#. PCI-X® interrupt requests. [B, A]_PIRQA# may be an input or an open-drain output in support of the hot plug controller. [B, A]_PIRQ[B, C, D]# are inputs only.	IOD; input	VDD33	3-State	3-State	Low	Low
[B, A]_PLLCLKO. PLL clock output. See section 4.3 for details.	Output	VDD33	Func.	Func.	Func.	Func.
[B, A]_PLLCLKI. PLL clock input. See section 4.3 for details.	Input	VDD33				
[B, A]_PME#. Power management event interrupt. The IC asserts this signal when SHPC-defined power management events occur. This signal is typically connected to the system Southbridge, where it may be used to initiate system state transitions.	OD	VDD33	3-State	3-State	3-State	3-State
[B, A]_PRESET#. Secondary PCI bus reset. This is asserted whenever RESET# is asserted or when programmed by Dev[B, A]:0x3C[SBRST]. Assertion of this pin does not reset any logic internal to the IC. Note: the PCI requirement for a delay between the rising edge of RST# and the first configuration access is not enforced by the IC with hardware; it is expected that this requirement be enforced through software. In hot plug mode, this is connected to [B, A]_HPSORR# of the power controller. See section 4.5.3.3 for details.	Output	VDD33	Low	High	Low	High

Pin name and description	IO cell type	Power plane	Normal		Single slot HP	
			During reset	After reset	During reset	After reset
<p>[B, A]_REQ[4:0]#. PCI-X® master request input signals. A_REQ[2:0]# are used for test-mode selection; see section 9.</p> <p>[B, A]_HPSOLC. Hot plug serial output latch clock (alternative function to [B, A]_REQ4 selected by hot plug mode, DevA:0x48[HPENA, HPENB]). See section 4.5.3.3 for details.</p>	IO (See left)	VDD33	REQs are inputs; HP-SOLC: High	REQs are inputs; HP-SOLC: High	[0]: Low; [3:1]: inputs; HP-SOLC: High	[0]: Low; [3:1]: inputs; HP-SOLC: High
<p>[B, A]_REQ64#. PCI-X® request for 64-bit transfers. The IC drives this signal to the asserted state while [B, A]_PRESET# is asserted.</p>	IO	VDD33	Low	3-State	Low	Low
<p>[B, A]_SERR#. PCI-X® system error signal.</p>	Input	VDD33			Low	Low
<p>[B, A]_STOP#. PCI-X® target abort signal. During reset, these signals may be 3-state or they may be driven, based on the requirements of the PCI-X® initialization pattern.</p>	IO	VDD33	See left	3-State	Low	Low
<p>[B, A]_TRDY#. PCI-X® target ready signal. During reset, these signals may be 3-state or they may be driven, based on the requirements of the PCI-X® initialization pattern.</p>	IO	VDD33	See left	3-State	Low	Low
<p>HPSIC. Hot plug serial input clock; see section 4.5.3.3 for details. This signal is an input only while PWROK is low and an output at all other times. As an input, it is used to specify the default state of DevA:0x40[HPSSS#]. HPSSS# specifies if the IC supports a single hot plug slot on the bridge without external isolation switches. A weak resistor should be tied from this signal to VDD33 or to ground.</p>	IO (See left)	VDD33	High	High	High	High
<p>HPSIL#. Hot plug serial input load; see section 4.5.3.3 for details. This signal is an input while PWROK is low and an output at all other times. As an input, it is used to specify if bridge B of the IC is in hot plug mode or not; the latched state is available in DevA:0x48[HPENB]. To specify that bridge B is in hot plug mode, a weak resistor to VDD33 should be placed on this signal. To specify that bridge B is not in hot plug mode, a weak pulldown resistor to ground should be placed on this node. When neither bridge A nor B are in hot plug mode, this signal is always driven high.</p>	IO (See left)	VDD33	High	High	High	High
<p>HPSOC. Hot plug serial output clock; see section 4.5.3.3 for details.</p>	Output	VDD33	High	High	High	High

Pin name and description	IO cell type	Power plane	Normal		Single slot HP	
			During reset	After reset	During reset	After reset
HPSOD. Hot plug serial output data; see section 4.5.3.3 for details. This signal is an input while PWROK is low and an output at all other times. As an input, it is used to specify if bridge A of the IC is in hot plug mode or not; the latched state is available in DevA:0x48[HPENA]. To specify that bridge A is in hot plug mode, a weak resistor to VDD33 should be placed on this signal. To specify that bridge A is not in hot plug mode, a weak pulldown resistor to ground should be placed on this node. When neither bridge A nor B are in hot plug mode, this signal is always driven low.	IO (See left)	VDD33	Low	High	Low	High
NIOAIRQ[A, B, C, D]#. Non-IOAPIC interrupt requests. Each of these signals require a weak pullup resistor to VDD33. In particular, if the state of NIOAIRQC# is low during the rising edge of PWROK, then the IC will enter a production test mode that results in undefined behavior in [B, A]_PLLCLKO and [B, A]_PLLCLKI. See section 4.5.2 and Dev[B, A]:0x40[NIOAMODE] for details about the function of these pins.	OD	VDD33	3-State	3-State	3-State	3-State
P_CAL, P_CAL#. PCI-X® PHY calibration pins. These are designed for the following external circuit: P_CAL should be connected through a resistor to ground; P_CAL# should be connected through a resistor to VDD33. The calculated calibration values associated with these resistors are provided DevA:0x[54, 50][CALCCOMP].	Input	VDD33				

If a bridge is to be left unused, the signals associated with that bridge should be connected as follows:

- The following signals do not require any connection: [B, A]_AD[63:0], [B, A]_CBE_L[7:0], [B, A]_PAR, [B, A]_PAR64, [B, A]_PCLK[4:0], [B, A]_PRESET#.
- The following signals should be tied high through resistors: [B, A]_ACK64#, [B, A]_DEVSEL#, [B, A]_FRAME#, [B, A]_IRDY#, [B, A]_PERR#, [B, A]_PIRQ[D:A]#, [B, A]_REQ[4:0]#, [B, A]_SERR#, [B, A]_STOP#, [B, A]_TRDY#, [B, A]_GNT[4:0]#, [B, A]_PME#, [B, A]_REQ64#.
- The following signals should be grounded: [B, A]_PCIXCAP, [B, A]_M66EN.
- [B, A]_PLLCLKO should be connected to [B, A]_PLLCLKI.

3.4 Test and Miscellaneous Signals

Pin name and description	IO cell type	Power plane	During reset	After reset
CMPOVR. Link automatic compensation override. 0=Link automatic compensation is enabled. 1=The compensation values stored in DevA:0x[E0, E4, E8] control the compensation circuit. The state of this signal determines the default value for DevA:0x[E0, E4, E8][ACTL and BCTL] at the rising edge of PWROK.	Input	VDD33		
FREE[22:1]. These pins should be left unconnected.				
LDTSTOP#. Link disconnect control signal. This pin is also used for test-mode selection; see section 9.	Input	VDD33		
NC[3:0]. These pins should be left unconnected.				

PWROK. Power OK. 1=All power planes are valid. The rising edge of this signal is deglitched; it is not observed internally until it is high for more than 6 consecutive REFCLK cycles. See section 4.2 for more details about this signal.	Input	VDD33		
REFCLK. 66 MHz reference clock. This is required to be operational and valid for a minimum of 200 microseconds prior to the rising edge of PWROK and always while PWROK is high.	Input	VDD33		
RESET#. Reset input. See section 4.2 for details. Note: RESET# is also used as the hot plug [B, A]_HPSOR# reset. When RESET# is asserted, the hot plug shift register and control latches are reset.	Input	VDD33		
RSVD[22:0]. These pins should be left unconnected.				
STRAPL[3:2]. Strapping options to be tied low. These pins should be tied to ground.	Input	VDD33		
TEST. This pin is required to be tied low for functional operation. See section 9 for details.	Input	VDD33		

3.5 Power and Ground

VDD12[B, A]. 1.2 volt power plane for the HyperTransport technology pins. VDD12A provides power to the A side of the tunnel. VDD12B provides power to the B side of the tunnel.

VDD18. 1.8-volt power plane for the core of the IC.

VDDA18. Analog 1.8-volt power plane for the PLLs in the core of the IC. This power plane is required to be filtered from digital noise.

VDD33. 3.3-volt power plane for IO.

VSS. Ground.

3.5.1 Power Plane Sequencing

The following are power plane requirements that may imply power supply sequencing requirements.

- VDD33 is required to always be higher than VDD18, VDDA18, and VDD12[B, A].
- VDD18 and VDDA18 are required to always be higher than VDD12[B, A].

4 Functional Operation

4.1 Overview

The IC connects to the host through either the side A or side B HyperTransport™ link interface. The other side of the tunnel may or may not be connected to another device. Host-initiated transactions that do not target the IC or the bridge flow through the tunnel to the downstream device. Transactions claimed by the device are passed to internal registers or to one of the PCI-X® bridges.

See section 5.1 for details about the software view of the IC. See section 5.1.2 for a description of the register naming convention. See the AMD-8131™ HyperTransport™ PCI-X® Tunnel Design Guide for additional information.

4.2 Reset And Initialization

RESET# and PWROK are both required to be low while the power planes to the IC are invalid and for at least 1 millisecond after the power planes are valid. Deassertion of PWROK is referred to as a *cold reset*. After PWROK is brought high, RESET# is required to stay low for at least 1 additional millisecond. After RESET# is brought high, the links go through the initialization sequence.

After a cold reset, the IC can be reset by asserting RESET# while PWROK remains high. This is referred to as a *warm reset*. RESET# must be asserted for no less than 1 millisecond during a warm reset.

4.2.1 Non-Hot Plug Initialization

The operational mode (conventional PCI or PCI-X) and frequency of the PCI-X bridges ([B, A]_PCLK[4:0]) are determined after a cold reset by the [B, A]_PCIXCAP signals, the [B, A]_M66EN signals, and by strapping resistors on [B, A]_GNT[4:3] (the A_ signals specify bridge A and the B_ signals specify bridge B). For [B, A]_GNT[4:3], to select a 1, a pullup resistor to VDD33 is placed on the signal; to select a 0, a pulldown resistor to ground is placed on the signal. The mode and frequency is determined while PWROK is low and held after PWROK goes high. The options and associated selects are:

<u>GNT[4:3]#</u>	<u>PCIXCAP</u>	<u>M66EN</u>	<u>Mode</u>	<u>Frequency</u>	<u>Supported R/G/P</u>
XXb	Grounded	0	Conventional PCI	33.33 MHz	[4:0]
XXb	Grounded	1	Conventional PCI	66.67 MHz	[3:0]
XXb	Middle voltage	X	PCI-X®	66.67 MHz	[3:0]
01b	Pullup to VDD33	X	PCI-X®	100.00 MHz	[3:0]
00b	Pullup to VDD33	X	PCI-X®	133.33 MHz	[2:0]

Notes to the above table:

- X means that the state does not matter to the IC.
- The GNT[4:3]# column shows the value latched by the IC while PWROK is low.
- The PCIXCAP column indicates how the IC observes PCIXCAP. *Grounded* indicates PCIXCAP is tied to ground. *Middle voltage* indicates PCIXCAP is tied to a pullup resistor to VDD33 and between 1 and 5 parallel pulldown resistors to ground; these pulldown resistors may come from the systemboard and from cards located in up to four slots. *Pullup to VDD33* indicates PCIXCAP is tied to a pullup resistor to VDD33 and no pulldown resistors.
- The Supported R/G/P column indicates the sets of REQ#, GNT#, and PCLK signals that are supported by the IC's bridge in that mode (there may be other constraints such as electrical requirements that further limit the number of external devices supported).
- If a bridge from the IC supports 5 slots, then only 33 MHz conventional PCI mode is supported. In this situation, M66EN and PCIXCAP should be grounded on the systemboard to the slots so that all PCI cards properly initialize.
- The state of the straps is reflected in Dev[B, A]:0xA0[SCF] and Dev[B, A]:0x40[CPCI66] after a cold reset.
- If the systemboard supports PCI-X mode operation for a bridge, then a pullup resistor to VDD33 must be placed on the bridge's PCIXCAP pin. To limit the frequency of a PCI-X-capable bridge to 66 MHz on a systemboard, the systemboard must also include a pulldown resistor from the bridge's PCIXCAP pin to ground. The strapping options on GNT[4:3]# are used to distinguish between systems that support 100MHz and 133 MHz; in either of these two cases, the system board should include no pulldown resistors on PCIXCAP.

4.2.2 Hot Plug Initialization

Bridges in hot plug mode are always placed into 33 MHz, conventional PCI mode after RESET# is asserted. The operational speed and mode are then initialized by software through the following steps: (1) initialization of write once registers in the SHPC[B, A]:XX register block, (2) optional execution of Power Only All Slots SHPC command, (3) acquisition of the capabilities and presence information for each slot by observing the RST#, M66EN, PCIXCAP, PRSNT1#, and PRSNT2# signals, (4) determination of the highest common bus frequency and mode that may be selected, (5) execution of Set Bus Segment Speed/Mode SHPC command for the selected speed and mode, and (6) execution of Enable All Slots SHPC command.

This sequence is the same when hot plug single-slot support is selected (Dev[B, A]:0x40[HPSSS#]). However,

all of the slot signals are forced low until the slot is powered.

4.3 Clocking

It is required that REFCLK be valid in order for the IC to operate. Also, the LR[B, A]CLK inputs from the operation links must also be valid at the frequency defined DevA:0xCC[FREQA] and DevA:0xD0[FREQB]. The IC provides [B, A]_PCLK[4:0] as the clocks to the secondary bus devices.

4.3.1 Systemboard Requirements

The IC provides the PCI clocks for secondary-bus devices, [B, A]_PCLK[4:0], and a PLL feedback for itself, [B, A]_PLLCLKO to [B, A]_PLLCLKI. [B, A]_PLLCLKO is fixed at 66 MHz while [B, A]_PCLK[4:0] varies based on the specified secondary-bus frequency.

The systemboard is required to include a loopback connection from [B, A]_PLLCLKO to [B, A]_PLLCLKI. The length of this connection is required to be approximately the same as the length of the [B, A]_PCLK traces from the IC to the external PCI devices (the length of the connection from A_PLLCLKO to A_PLLCLKI should be the same as the length of the A_PCLK signals and the length of the connection from B_PLLCLKO to B_PLLCLKI should be the same as the length of the B_PCLK signals) such that the flight time of the [B, A]_PCLK signals is the same as the flight time of the PLL feedback. Flight time is defined as the time difference between the rising edge of the clock as observed at the source of the systemboard trace ([B, A]_PLLCLKO and [B, A]_PCLK at the IC) and the rising edge of the clock as observed at the destination of the systemboard trace ([B, A]_PLLCLKI at the IC and [B, A]_PCLK at the external device), as shown in Figure 2.

The IC is designed such that, for the purposes of meeting the IC AC timing requirements, if the PCLK flight time matches the PLL feedback flight time, then PCLK as observed at the destination is equivalent to the PCI-defined PCLK signal to the IC. Accordingly, the PLL feedback flight time is required to be the same as any of the PCLK trace flight times (for a bridge), within the skew limits specified by the PCI specifications for PCLK to different devices (2 ns for conventional PCI 33 MHz; 1 ns for conventional PCI 66 MHz; 0.5 ns for all PCI-X mode frequencies).

To improve the correlation between PCLK and the PLL feedback flight time, the delay of [B, A]_PLLCLKO and [B, A]_PCLK[4:0], relative to each other, may be altered through Dev[B, A]:0x40[PCLKDEL, PLLODEL]. However, the delay created by each increment of Dev[B, A]:0x40[PCLKDEL and PLLODEL] varies from device to device. Therefore, DevA:0x48[BDCV] provides the approximate time differential for each increment of Dev[B, A]:0x40[PCLKDEL and PLLODEL] on a given device. Thus, the values programmed into Dev[B, A]:0x40[PCLKDEL and PLLODEL] should be determined using DevA:0x48[BDCV] to adjust timing on a platform as follows:

- If the PLL feedback flight time is greater than the PCLK flight time by n picoseconds, then Dev[B, A]:0x40[PCLKDEL] should be set to: $n / (1250 / \text{DevA:0x48[BDCV]})$.
- If the PLL feedback flight time is less than the PCLK flight time by n picoseconds, then Dev[B, A]:0x40[PLLODEL] should be set to: $n / (1250 / \text{DevA:0x48[BDCV]})$.

The result of the above equations should be rounded to the nearest integer for best accuracy. Note that only one of Dev[B, A]:0x40[PCLKDEL] or Dev[B, A]:0x40[PLLODEL] should be set to a value other than zero, never both.

The PCLK flight time may vary with different [B, A]_PCLK frequencies (this may be a consequence of imperfect signal integrity of [B, A]_PCLK). These differences in flight time may be accounted for on a platform by: (1) for each supported [B, A]_PCLK frequency, determine the different time adjustment values required, and (2) use system BIOS to program Dev[B, A]:0x40[PCLKDEL and PLLODEL] to correct by these time adjustment values based on the [B, A]_PCLK frequency and DevA:0x48[BDCV].

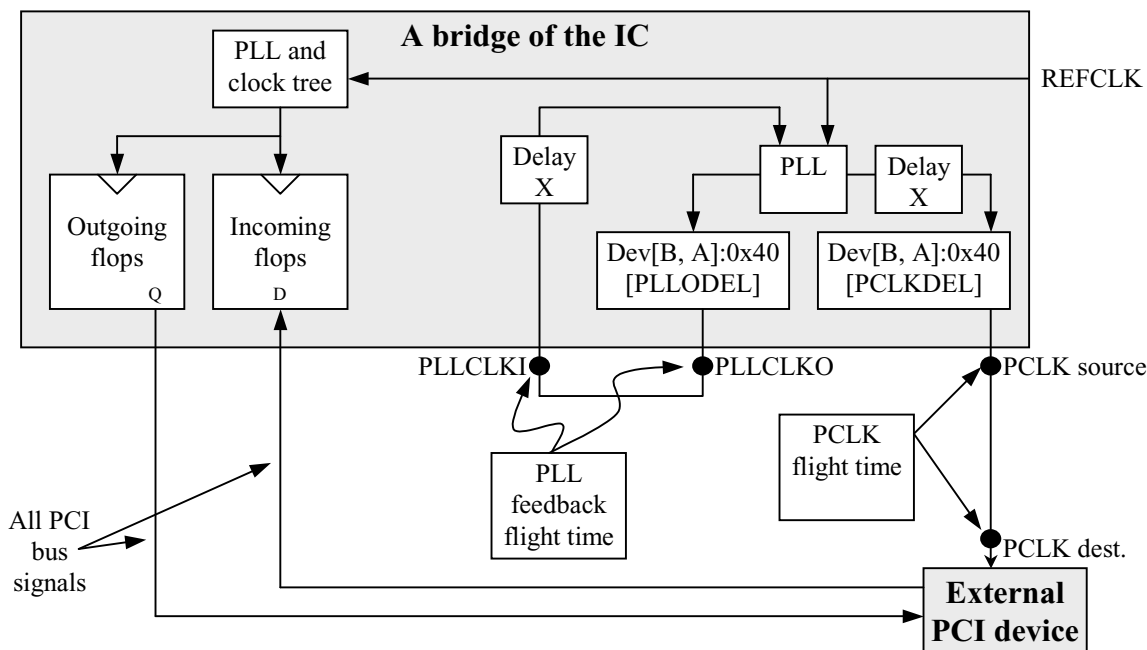


Figure 2: Systemboard clocking.

4.3.2 Characterization

For the purposes of characterization, there is no PCI-defined PCLK signal into the IC, such as is typically used to measure setup, hold, and output valid delay times of PCI-bus signals. As shown in Figure 2, there is an unspecified skew between the PCLK outputs and PLLCLKO (Delay X). Because of this, along with the fact that the PLL feedback frequency may not be the same as PCLK, the PLL feedback signal cannot be used to characterize the IC PCI bus signals. Instead, the IC should be characterized as follows:

- Make sure Dev[B, A]:0x40[PCLKDEL and PLLODEL] = 0h.
- Measure the PLL feedback flight time, PLLFT, and the flight time of a [B, A]_PCLK signal connected to an external device, PCLKFT. Algebraically calculate the difference as follows: $DIFFT = PLLFT - PCLKFT$. Note that DIFFT may be a positive or a negative value.
- Characterize the PCI-bus signals using, as a reference clock, the destination of the [B, A]_PCLK used to calculate PCLKFT. Then, adjust the measurements to obtain the correct values as follows:
 - **Output valid delay.** Algebraically subtract DIFFT from the measured output valid delay time of PCI bus signals driving out of the IC. Corrected output valid delay = Measured output valid delay - DIFFT.
 - **Setup.** Algebraically add DIFFT to the measured setup time of PCI bus signals being driven into the IC. Corrected setup time = Measured setup time + DIFFT.
 - **Hold.** Algebraically subtract DIFFT from the measured hold time of PCI bus signals being driven into the IC. Corrected hold time = Measured hold time - DIFFT.

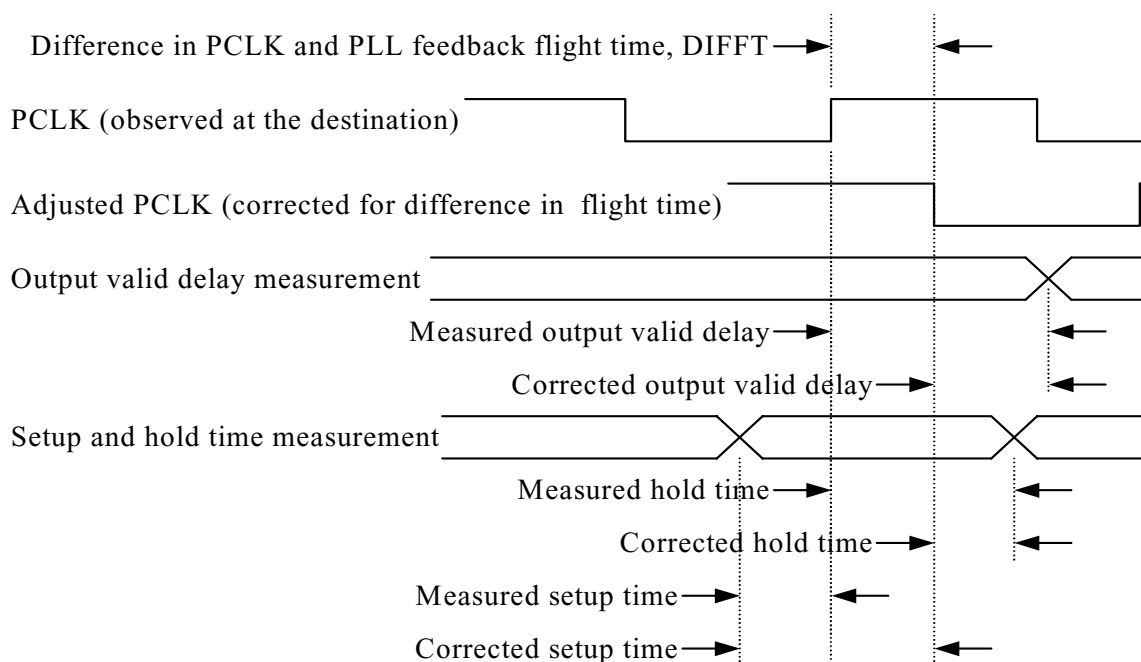


Figure 3: Correction for characterization.

4.3.3 Clock Gating

Internal clocks may be disabled during power-managed system states such as power-on suspend. It is required that all upstream requests initiated by the IC be suspended while in this state.

To enable clock gating, DevA:0xF0[ICGSMAF] is programmed to the values in which clock gating will be enabled. Stop Grant cycles and STPCLK deassertion link broadcasts interact to define the window in which the IC is enabled for clock gating during LDTSTOP# assertions. The system is placed into power managed states by steps that include a broadcast over the links of the Stop Grant cycle that includes the System Management Action Field (SMAF) followed by the assertion of LDTSTOP#. When the IC detects the Stop Grant broadcast which is enabled for clock gating, it enables clock gating for the next assertion of LDTSTOP#. While exiting the power-managed state, the system is required to broadcast a STPCLK deassertion message. The IC uses this message to disable clock gating during LDTSTOP# assertions. This is important because an LDTSTOP# assertion is not guaranteed to occur after the Stop Grant broadcast is received. The clock gating window must be closed to ensure that clock gating does not occur during Stop Grant for LDTSTOP# assertions that are not associated with the power states specified by DevA:0xF0[ICGSMAF].

In summary, Stop Grant broadcasts with SMAF fields specified by DevA:0xF0[ICGSMAF] enable the clock gating window and STPCLK deassertion broadcasts disable the window. If LDTSTOP# is asserted while the clock gating window is enabled, then clock gating occurs.

It is expected that clock gating is only employed during power-on suspend. Therefore, OS and driver software ensure that no DMA or interrupt activity occurs. In addition, it is required that there be no host accesses to the bridges or internal registers in progress from the time that LDTSTOP# is asserted for clock gating until the link reconnects after LDTSTOP# is deasserted.

4.4 Tunnel Links

Each HyperTransport link supports CLK receive and transmit frequencies of 200, 400, 600, 800 MHz. The side A and side B frequencies are independent of each other.

4.4.1 Link PHY

The PHY includes automatic compensation circuitry and a software override mechanism, as specified by DevA:0x[E8, E4, E0]. The IC only implements synchronous mode clock forwarding FIFOs. So only the link receive and transmit frequencies specified in DevA:0x[D0, CC][FREQB, FREQA] are allowed.

4.5 PCI-X® Bridges

The IC includes two 64-bit PCI-X bridges, bridge A and bridge B. Each independently support PCI-X mode or conventional PCI mode, clocks speeds of 33, 66, 100, and 133 MHz, and SHPC-compatible hot plug. Each include an IOAPIC register set. Each support 64-bit addressing in PCI-X and legacy PCI modes.

4.5.1 Tags, UnitIDs, SeqIDs And Ordering

The IC requires two HyperTransport technology-defined UnitIDs. The first UnitID applies to bridge A and the second UnitID applies to bridge B. It is contained in the following transactions:

- External master requests associated with the bridge.
- IOAPIC interrupt requests associated with the bridge.
- Responses to host-initiated requests that enter the address space of the bridge including configuration registers (DevA registers for bridge A and DevB registers for bridge B), IO and memory space windows defined in the configuration registers of the bridge, and the base address register spaces defined by the bridge.

In addition, the UnitID associated with the bridge is returned in the response to upstream requests and is used to determine the destination of the response (bridge A or bridge B).

The assigned SrcTag value increments with each non-posted request from 0 to 28 and then rolls over to 0 again; the first SrcTag assigned after reset is 0. Up to 29 non-posted requests to the link may be outstanding at a time per bridge.

Based on the state of Dev[B, A]:0x40[NZSEQID], the IC may or may not generate a non-zero SeqID values in the upstream link requests that result from external PCI master read requests.

All bridge-sourced transactions are compliant to PCI ordering rules. As PCI transactions are converted to link transactions, they are translated as described in the link specification.

Downstream non-posted link requests to a bridge that contain non-zero SeqID values are required to complete on that bus prior to initiating subsequent non-posted requests to that bus with the same SeqID value. Thus, only one downstream non-posted request with each non-zero SeqID value can be outstanding to a bridge at a time.

4.5.2 Interrupt Controllers

Each bridge supports the four PCI-defined interrupt signals, [B, A]_PIRQ[D, C, B, A]#. Assertion of these interrupt signals may be converted to link interrupt request messages as specified by the IOAPIC register space. Also, the interrupts from both bridges may be combined and output to the respective NIOAIRQ[D, C, B, A]# pins, based on Dev[B, A]:0x40[NIOAMODE].

It is expected that system BIOS sets both Dev[B, A]:0x40[NIOAMODE] bits and that the interrupt type is determined by the way the operating system programs the interrupt mask bit (RDR[IM]; see section 5.4) of the redirection registers (non-IOAPIC-capable operating systems set the mask bits resulting in NIOAIRQ[D:A]# signal assertions; IOAPIC-capable operating systems clear the mask bits resulting in interrupt request messages to the host). The NIOAIRQ[D:A]# signals from all instances of the IC on a platform may be connected together (respectively: A to A, B to B, etc.). These four nodes are expected to be passed to the system's legacy interrupt controller to generate interrupts on behalf of the IC's bridges when IOAPIC interrupts are not supported.

There is a set of IOAPIC registers associated with each bridge. They include a standard PCI function header (function 1 of each bridge) and memory mapped registers. In addition, expanded programmability of these registers is included in Dev[B, A]:0x[BC, B8]. The IOAPIC registers specify the conversion of PIRQ pin assertions to link interrupt request messages.

Typically, for PCI interrupts, the redirection register (RDR; see section 5.4) is set up as follows: MT=fixed; DM=physical mode; POL=active low; TM=level sensitive; and IM=not masked. The RDR fields are mapped into link interrupt request messages as follows:

<u>RDR field</u>	<u>Field in link packet</u>
IV[7:0] (interrupt vector)	Vector (bit time 5)
MT[2:0] (message type)	MT[2:0] (bits[4:2] of bit time 3); MT[3] (bit[7] of bit time 3) should always be low; note that the encoding of these bits changes between the value in the RDR and the value placed into the link packet.
DM (destination mode)	DM (bit[6] of bit time 3)
TM (trigger mode)	RQEOI (bit[5] of bit time 3)
DEST[7:0] (destination)	IntrInfo[15:8] (bit time 4); IntrInfo[55:16] should always be low.

DS, POL, IRR, and IM from the RDR are not included in the link interrupt packet.

The state of PASSPW and INTRINFO[55:24, 7] from the IDRDR register (see Dev[B, A]:0x[BC, B8]) are also passed along in the link interrupt packet.

If RDR[TM]=level sensitive for the interrupt request, then the IRR register is set when the interrupt is detected. After the interrupt request message is sent to the host, the host is required to generate an EOI broadcast message when finished with that interrupt. IRR is cleared in any RDRs (in either bridge) with IDRDR/RDR fields that match the IntrInfo fields of the EOI broadcast as follows:

<u>IntrInfo[15:8]</u>	<u>Match fields</u>
00h	IntrInfo[31:16] = {IDRDR[31:24], RDR/IDRDR[IV]};
01h-FFh	IntrInfo[31:8] = {IDRDR[31:24], RDR/IDRDR[IV], RDR/IDRDR[DEST]};

If the interrupt signal is still asserted when the corresponding RDR logic receives an IRR-clearing EOI, then IRR is set again immediately and a new interrupt request message is sent. If the interrupt signal is deasserted near the time the corresponding IRR-clearing EOI is received, then it is undefined whether an additional interrupt request message is sent. If RDR[TM]=edge sensitive, then the state of the IRR bit is not specified and the RDR logic for that interrupt does not observe EOIs.

Each RDR in the IC operates independently. If interrupts are received simultaneously by two RDR controllers, then the corresponding interrupt request messages from each are transmitted in an unspecified order.

If LDTSTOP# is asserted near the time that an interrupt is asserted, then the corresponding interrupt request message may or may not be sent before the disconnect sequence completes. If it is not sent before the disconnect sequence completes, then it is not dropped; it is sent after the link is re-connected.

External devices are required to assert PIRQ[D:A]# for at least 3 PCLK cycles in order to guarantee that the IC detects the assertion, regardless of the state of the corresponding RDR[TM] field.

4.5.2.1 Error NMI Interrupts

NMI interrupts may be generated as a result of assertions of the [B, A]_PERR# and [B, A]_SERR# signals (regardless of the source of the assertion), as enabled by Dev[B, A]:0x44[NMIEN]. These interrupt requests are generated with the following link format: PassPW=0; INTRINFO[55:24]=0000_00F8h; IV=00h; DEST=FFh; DM=0 (physical); TM=0 (edge); MT=0011b (NMI).

4.5.3 Hot Plug

Each PCI-X Bridge includes an SHPC-compliant hot plug controller that may be used to support hot plug capable conventional PCI or PCI-X slots. Strapping options on HPSOD and HPSIL# specify if hot plug is supported on bridge A and/or B. If hot plug is supported on a bridge, then all slots connected to that bridge are required to include hot plug support circuitry. With the exception of a single-slot hot plug implementation, the hot plug support circuitry includes one or more Texas Instruments TPS2340 hot plug power controllers, power switches, and associated slot isolation switches to provide electrical isolation for most of the slot signals. For a single-slot hot plug implementation, the IC provides the bus isolation function; therefore only the TI TPS2340 hot plug power controller and the power switches are required. Each bridge supports a maximum of 4 slots when hot plug mode is enabled.

The IC's hot plug controller is designed to interface with the TI TPS2340 hot plug power controller. Each TI TPS2340 controls two slots and provides two separate sets of isolation switch controls. TI TPS2340 controllers may be cascaded to support additional slots. A single TI TPS2340 hot plug power controller cannot be shared across bridge A and bridge B. The IC is connected to the power controller via a serial bus. One serial interface supports the power controllers for bridge A and bridge B.

4.5.3.1 Multi-slot Hot Plug

If multiple hot plug slots are supported on a bridge, isolation switches are required for each slot to provide electrical isolation. Each TI TPS2340 hot plug power controller provides two pairs of isolation switch control signals, BUSENx# and CLKENx#, to control the state of the switches.

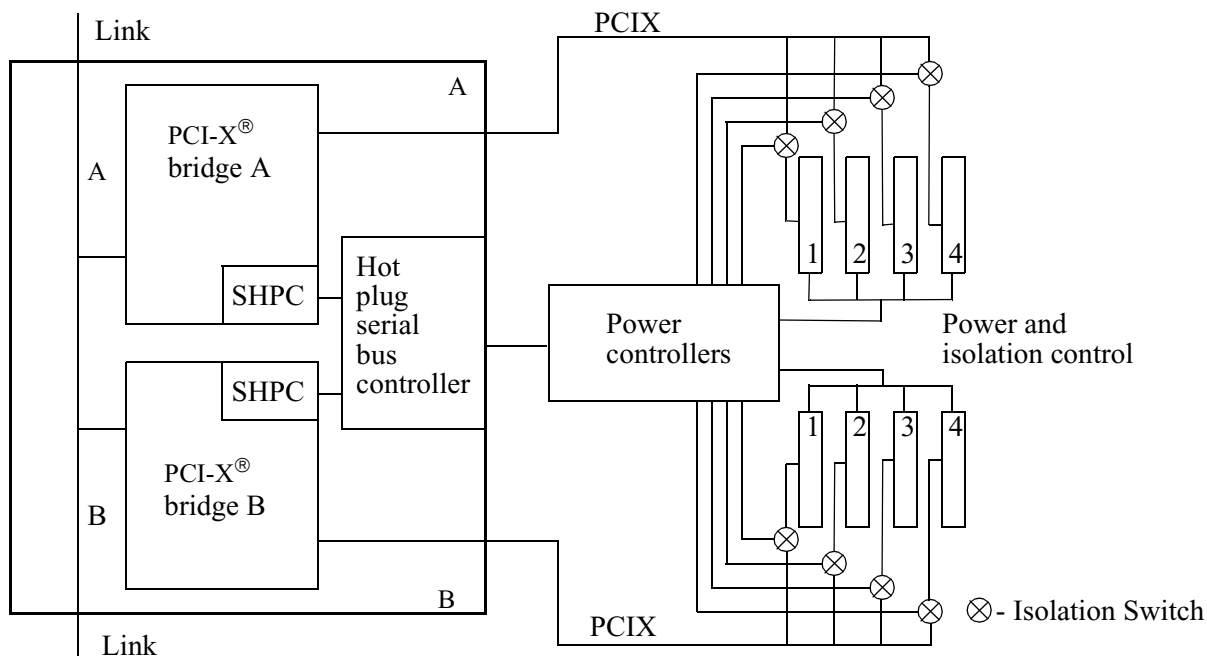


Figure 4: System diagram for multiple hot plug slots on a bridge.

The following table associates the hot plug power controller’s isolation switch control signal with the IC’s slot signals.

Power controller signal	Slot signals isolated
BUSENx#	[B, A]_ACK64#, [B, A]_AD[63:0], [B, A]_CBE_L[7:0], [B, A]_DEVSEL#, [B, A]_FRAME#, [B, A]_GNT#[3:0], [B, A]_IRDY#, [B, A]_PAR, [B, A]_PAR64, [B, A]_PERR#, [B, A]_PIRQ[A, B, C, D]#, [B, A]_REQ#[3:0], [B, A]_REQ64#, [B, A]_SERR#, [B, A]_STOP#, [B, A]_TRDY#.
CLKENx#	[B, A]_PCLK[3:0], [B, A]_M66EN.

Table 2. Signal isolation groups.

The TI TPS2340 hot plug power controller controls RESET# to the slot. The IC’s [B, A]_PRESET# signals are connected to the TI TPS2340 hot plug power controller’s serial interface control signal, SORR#.

Some operating systems require that each configuration-space bus number provide a separate PME# signal to a general-purpose set of PME# status bits provided by the platform system management logic. The IC’s [B, A]_PME# signals are associated with the power management configuration registers Dev[B, A]:0x[9C:98], both of which are observed by software on the primary side of the PCI bridges and are therefore on the same bus number. Therefore, the IC’s two [B, A]_PME# pins may be connected together and passed to the platform system management logic.

The slots are observed by software on the IC’s secondary bus, which is a different bus number from the primary side. Therefore, the each bridge should provide a separate PME# signal to the platform system management logic, that logically connects to all the slots behind the bridge. The TI TPS2340 hot plug power controller’s PME# inputs connects to the PME# signal of each hot plug slot. Its PME# outputs for one bridge should be connected together and passed to the platform system management logic.

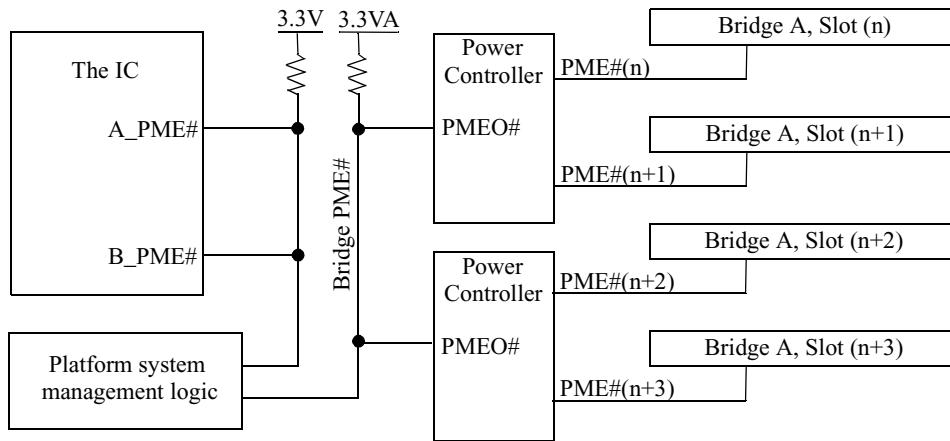


Figure 5: System diagram of PME# signals.

The slot signals that are used to communicate the speed, capability (M66EN and PCIXCAP), and presence of an adapter card (PRSNT[1:2]#) are isolated from the other slots in a hot plug implementation. These signals are directly connected from the slot connector to their associated TI TPS2340. The state of these signals is provided to the IC through the serial interface. The IC, in turn, makes the state of these signals available to system software. The [B, A]_PCIXCAP and [B, A]_M66EN pins on the IC are not used for sensing speed and mode. The IC's [B, A]_PCIXCAP pins are left unconnected.

The connection and function of the M66EN signal is unique in a hot plug implementation for two reasons: (1) M66EN is driven as an output of the IC; (2) its isolation switch control is driven by CLKEN# rather than BUSEN# (unlike other PCI/PCI-X control signals). In a hot plug configuration, the IC's [B, A]_M66EN pin is configured as an open-drain output. It is driven low by the IC if it is determined that the bus is to run at 33MHz (conventional PCI mode), as indicated in SHPC[B, A]:x10[MODE].

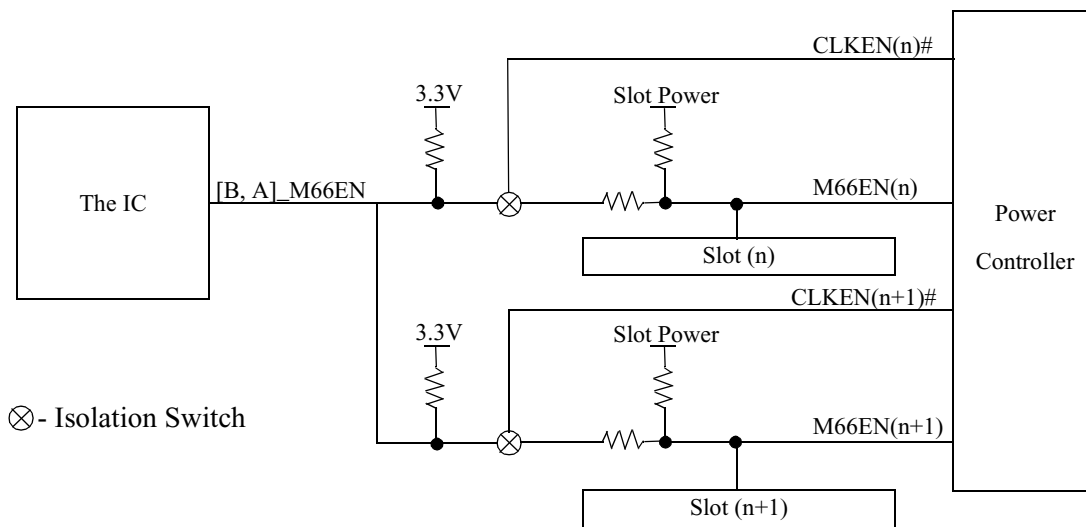
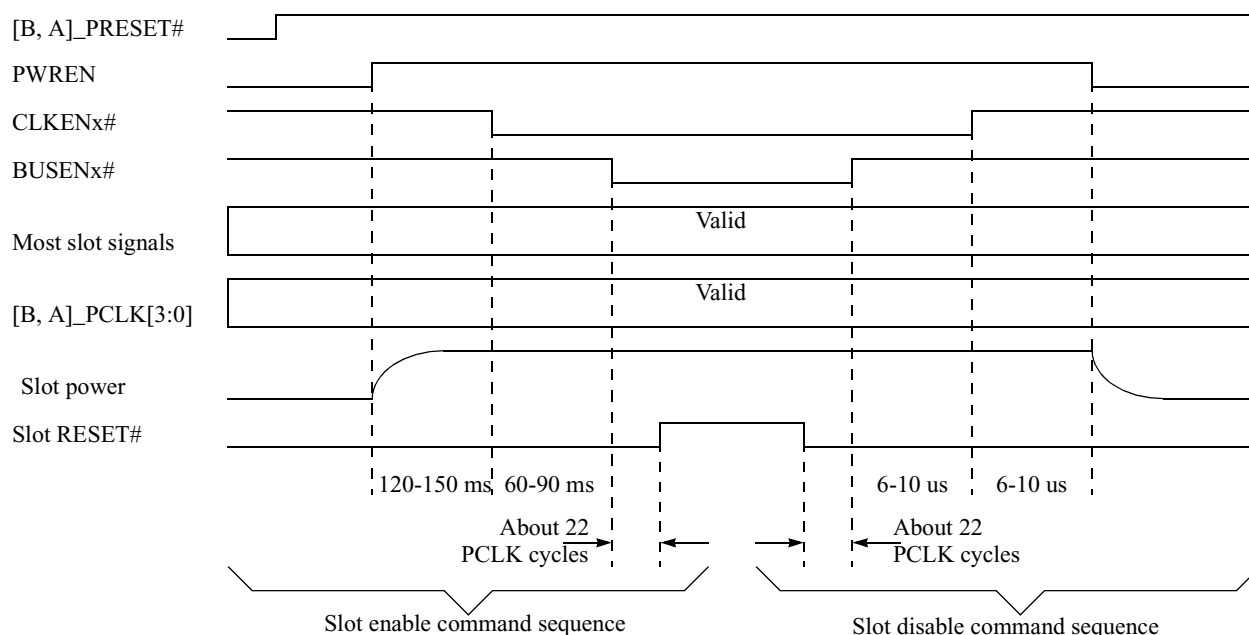


Figure 6: System diagram of M66EN signals.

The process of setting the state of [B, A]_M66EN when the bridge is initialized is: (1) the TI TPS2340 hot plug power controller is programmed to apply power to the slots via the SHPC[B, A]:14 power-only all slots command; (2) software observes the state of the speed capability signals for the slots by reading SHPC[B, A]:[30:24][M66_CAP]; (3) software issues the SHPC[B, A]:14 set bus segment speed/mode command, which immediately places the appropriate state on [B, A]_M66EN out of the IC; (4) software issues the SHPC[B, A]:14 slot enable command, which results in the assertion of CLKENx# so that [B, A]_M66EN out of the IC is enabled to the slot.



- Signal states are shown from the perspective of the pins of the IC; the perspective from the slot is different due to the isolation switches controlled by CLKENx# and BUSENx#.
- “Most slot signals” includes the signals controlled by BUSENx#.
- M66EN is driven low after RESET# is asserted; after that, its state is determined by the bus speed and mode.

Figure 7: Multi-slot hot plug enable/disable sequence.

4.5.3.2 Single-Slot Hot Plug

Isolation switches are not required if the bridge supports a single hot plug slot. The IC provides the isolation function by controlling the slot signals appropriately. The TI TPS2340 hot plug power controller is still required.

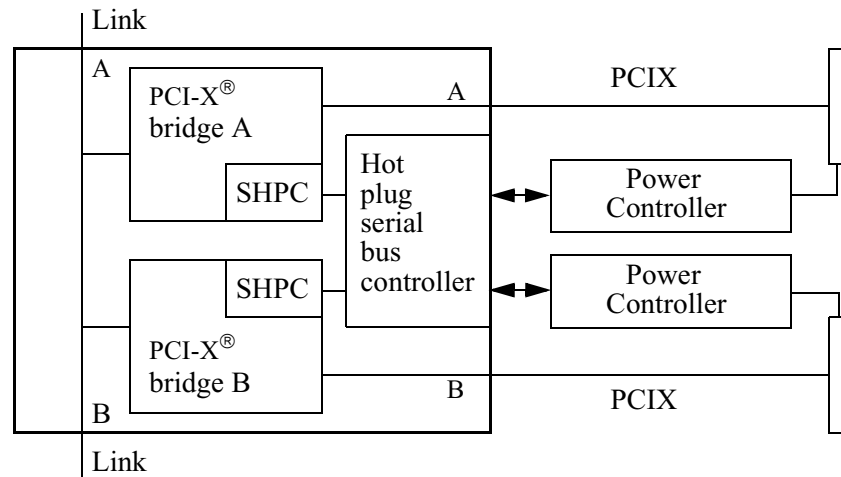
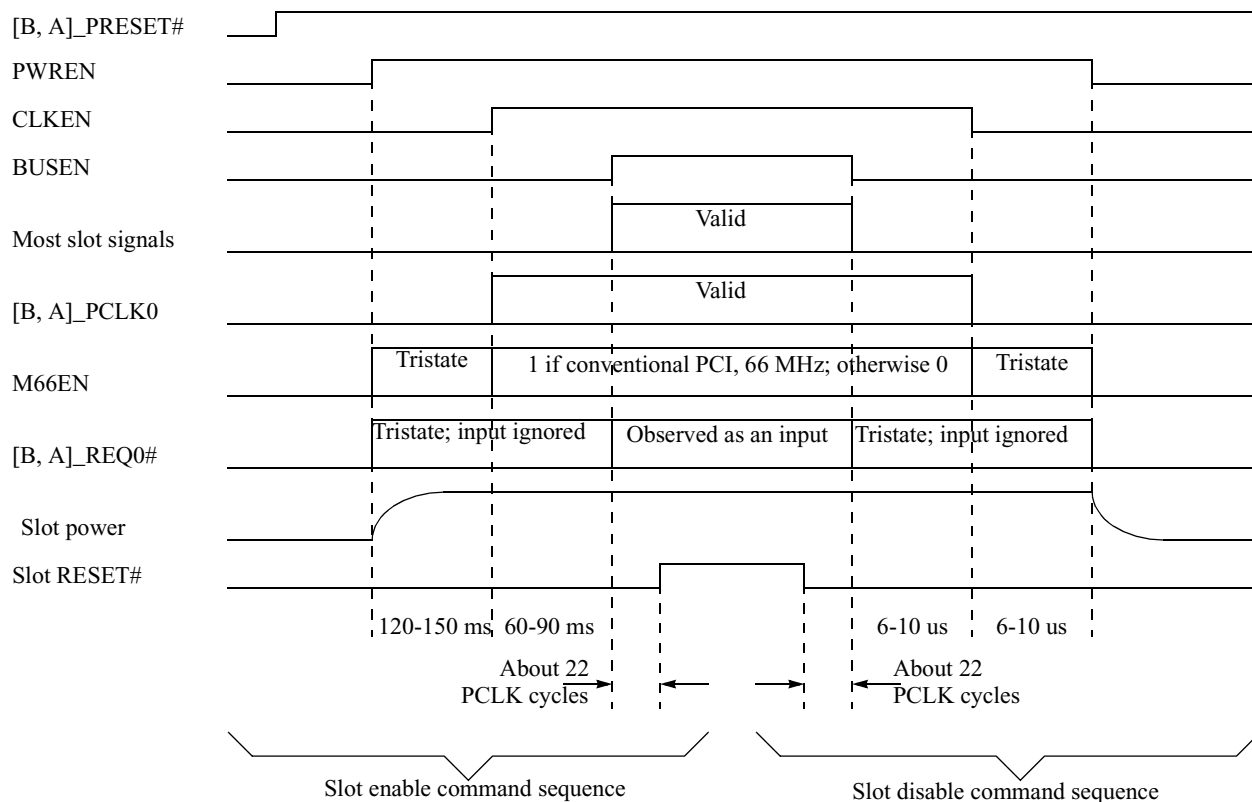


Figure 8: Single-slot hot plug system diagram.

Single-slot hot plug support is enabled for each bridge through strapping options on HPSIC and B_GNT2#. The state of these strapping options are observable through Dev[B, A]:0x40[HPSSS#].

The IC drives all slot signals low throughout the duration of a cold reset and continues to do so until after the TI TPS2340 hot plug power controller applies power to the adapter. The IC interprets the SHPC commands to control the signals in the power-only, slot enable, and slot disable sequences. The following figure shows how signals are controlled by the IC during the sequence initiated by the slot enable command and the slot disable command.



- CLKEN and BUSEN represent the approximate times in which the TI TPS2340 change the state of its CLKENx# and BUSENx# signals. However, these signals out of the TI TPS2340 do not directly control the above signals.
- PWREN represents the times in which the TI TPS2340 enables power to the slot.
- “Most slot signals” includes the signals controlled by BUSENx# in section 4.5.3.1.

Figure 9: Single-slot hot plug enable/disable sequence.

The IC’s [B, A]_PCIXCAP pins are left unconnected. PCIXCAP and PRSNT[1:2]# from the slot are connected to the TI TPS2340 hot plug power controller.

The IC’s [B, A]_M66EN pins are connected directly to the slot with a pull-up resistor to the slot power plane. This pin remains tri-stated until the SHPC enables the slot so the state provided by the card in the slot may be observed. This pin is driven low by the IC if it is determined that the bus is to run at 33MHz in conventional PCI mode.




Figure 10: Single-slot hot plug M66EN connections.

The process of setting the state of [B, A]_M66EN when the bridge is initialized is: (1) the TI TPS2340 hot plug power controller applies power to the slot via the SHPC[B, A]:14 power-only command; (2) software observes the state of the speed capability signals for the slot by reading SHPC[B, A]:[24][M66_CAP]; (3) software issues the SHPC[B, A]:14 set bus segment speed/mode command which determines the state of [B, A]_M66EN (but the state of the pins does not change); (4) software issues the SHPC[B, A]:14 slot enable command and [B, A]_M66EN may be driven low at the same time that CLKENx# is asserted out of the TI TPS2340.

The IC is designed such that only active-low interrupts (from [B, A]_PIRQ[D:A]#) are supported when in single-slot support mode, while the slot is not enabled. If the IOAPIC is programmed for active high interrupts in this mode, then spurious interrupt requests are generated.

4.5.3.3 Serial Interface

The hot plug serial interface operates at 8.33 MHz. It converts SHPC commands to a serial format to communicate with the TI TPS2340 hot plug power controllers. In addition, it is used to read status information from the TI TPS2340 hot plug power controllers and update the IC's SHPC status registers accordingly. The following table identifies two different groups of serial interface signals. *Common Serial* signals are connections between the IC and all TI TPS2340 hot plug power controllers (shared across both PCI/PCI-X bridges). *Bridge specific* signals are connections between the IC and only those TI TPS2340 hot plug power controllers connected to a particular bridge. (For additional information, see *TI TPS2340A Dual-slot PCI Hot-Plug Power Controller*.)

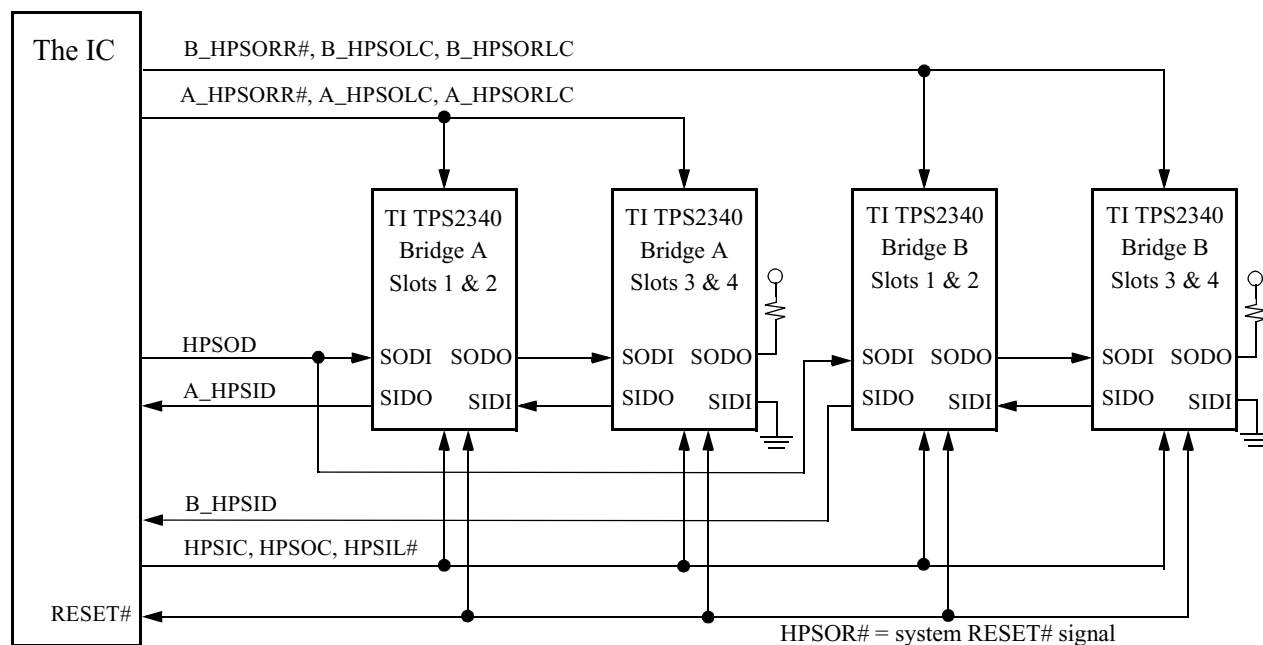


Figure 11: Hot plug serial interface connections.

Common Serial Signals

HPSIC. Output. This 8.33 MHz clock is used to shift serial data from the TI TPS2340 to the IC over [B, A]_HPSID. Each bit is captured by the IC on the rising edge of this clock.

HPSIL#. Output. This signal is used to specify the start of an input data frame and the data type--or channel number--being shifted into the IC over [B, A]_HPSID.

HPSOC. Output. This 8.33 MHz clock is used to shift serial data from the IC to the TI TPS2340 over HPSOD. Each bit is captured by the TI TPS2340s on the rising edge of this clock.

HPSOD. Output. This is the serial data shifted from the IC to the TI TPS2340s, clocked by HPSOC.

HPSOR#. The TI TPS2340's SOR# input is connected to the same platform reset signal as is used for the IC's RESET#. When SOR# is asserted, all of the slot control outputs of the TI TPS2340 are reset except the slot reset signal, RESETx# (which is reset by [B, A]_HPSORR#).

Bridge Specific Serial Signals

[B, A]_HPSID. Input (multiplexed with [B, A]_PCLK4). This is the serial data shifted into the IC from the TI TPS2340s, clocked by HPSIC.

[B, A]_HPSOLC. Output (multiplexed with [B, A]_REQ4#). This signal is used to load the state of the serial data shifted into the TI TPS2340 over HPSOD into output latches. All outputs of the TI TPS2340 are updated on the rising edge of [B, A]_HPSOLC, except the slot reset signal, RESETx# (which is updated by [B, A]_HPSORLC).

[B, A]_HPSORLC. Output (multiplexed with [B, A]_GNT4#). This signal is used to load the state of the slot reset signals shifted into the TI TPS2340 over HPSOD into output latches. RESETx# out of the TI TPS2340 is updated on the rising edge of [B, A]_HPSORLC.

[B, A]_HPSORR#. Output (multiplexed with [B, A]_PRESET#). This signal is used to reset the state of the TI TPS2340's output latches that drive RESETx# (to low).

4.5.3.3.1 Serial Data From The Power Controllers To The IC

Channel 00b interrupt-capable data and channel 01b non-interrupt-capable data is shifted into the IC from the TI TPS2340 over [B, A]_HPSID using HPSIC as the clock. This data is continuously shifted into the IC, toggling between channels 00b and 01b. HPSIL# controls the start of each block and specifies the channel number. HPSIL# transitions after the falling edge of HPSIC. The tables below show how the data for 4 slots are transferred into the IC. However, the actual number of slots transferred is limited to the maximum of the number of slots on bridge A or bridge B.

Clock #	HPSIL#	Data out of TI TPS2340 over [B, A]_HPSID after rising edge of HPSIC
0	1	
1	0 (start)	
2	0 (chan[0])	
3	0 (chan[1])	
4	1	
5	1	First TPS2340 SWA signal (1=MRL sensor is open) passed to SHPC[B, A]:24[MRLS].
6	1	First TPS2340 BUTTONA# signal (1=attention button is being pressed; this is an inversion of the state as it is placed onto the TPS2340 BUTTONA# pin) passed to SHPC[B, A]:24[AB].
7	1	First TPS2340 power fault state (0=power fault is detected) passed through an inverter to SHPC[B, A]:24[PF].
8	1	First TPS2340 PRSNT2A# signal passed to SHPC[B, A]:24[PRSNT1_2].

Table 3: Channel 00b, interrupt capable serial hot plug data to the IC.

9	1	First TPS2340 PRSNT1A# signal passed to SHPC[B, A]:24[PRSNT1_2].
12:10	1	Reserved.
17:13	1	First TPS2340 slot B signals passed to SHPC[B, A]:28[PRSNT1_2, PF, AB, MRLS].
20:18	1	Reserved.
25:21	1	Second TPS2340 slot A signals passed to SHPC[B, A]:2C[PRSNT1_2, PF, AB, MRLS].
28:26	1	Reserved.
33:29	1	Second TPS2340 slot B signals passed to SHPC[B, A]:30[PRSNT1_2, PF, AB, MRLS].

Table 3: Channel 00b, interrupt capable serial hot plug data to the IC.

Clock #	HPSIL#	Data out of TI TPS2340 over [B, A]_HPSID after rising edge of HPSIC
1	0 (start)	
2	1 (chan[0])	
3	0 (chan[1])	
4	1	
5	1	First TPS2340 M66ENA signal passed to SHPC[B, A]:24[M66_CAP].
6	1	First TPS2340 PCIXCAPA# signal passed to SHPC[B, A]:24[PCIX_CAP].
7	1	First TPS2340 PCIXCAPA# signal passed to SHPC[B, A]:24[PCIX_CAP].
8	1	First TPS2340 auxiliary power fault state (not observable).
12:9	1	Reserved.
16:13	1	First TPS2340 slot B signals passed to SHPC[B, A]:28[M66_CAP, PCIX_CAP].
20:17	1	Reserved.
24:21	1	Second TPS2340 slot A signals passed to SHPC[B, A]:2C[M66_CAP, PCIX_CAP].
28:25	1	Reserved.
32:29	1	Second TPS2340 slot B signals passed to SHPC[B, A]:30[M66_CAP, PCIX_CAP].

Table 4: Channel 01b, non-interrupt capable serial hot plug data to the IC.

4.5.3.3.2 Serial Data From The IC To The Power Controllers

Serial data is transferred over HPSOD to the TI TPS2340s, where it is stored, using HPSOC as the clock. The state of the outputs and control signals stored in the power controller do not change until a rising edge of [B, A]_HPSORLC, in the case of RESETx#, and [B, A]HPSOLC, in the case of the rest of the signals. The data is shifted whenever there is a need to change the state of these signals, normally as a result of a command to SHPC[B, A]:14. The IC shifts out four slots worth of data, regardless of how many slots are actually attached to the bridge, followed by pulses on [B, A]_HPSORLC and [B, A]HPSOLC. HPSOD transitions after the falling edge of HPSOC.

Clock #	Data out of the IC over HPSOD during the rising edge of HPSOC
1	Second TPS2340 slot B power enable state. 1=Power enabled to the slot.
2	Second TPS2340 CLKENB state. 1=CLKEN signals enabled.
3	Second TPS2340 BUSENB state. 1=BUSEN signals enabled.
4	Second TPS2340 RESETB# state.
5	Second TPS2340 PWRLEDB# state. 1=Turn on power LED.

Table 5: Serial hot plug data from the IC to the power controller.

6	Second TPS2340 ATTLEDB# state. 1=Turn on attention LED.
8:7	Reserved.
14:9	Second TPS2340 slot A control signals and outputs.
16:15	Reserved.
22:17	First TPS2340 slot B control signals and outputs.
24:23	Reserved.
30:25	First TPS2340 slot A control signals and outputs.

Table 5: Serial hot plug data from the IC to the power controller.

4.5.3.4 SHPC Interrupts, Events, And Errors

Under the conditions described by SHPC[B, A]:20, the IC may assert [B, A]_PIRQA#, [B, A]_PME#, or indicate a system error on the links.

4.5.3.5 Reset To Hot Plug Slots

The state of reset for each hot plug slot is passed from the SHPC controller, through the serial bus, to the TI TPS2340 hot plug power controllers, where it is driven to the slots. The PCI requirement for a delay between the rising edge of RST# and the first configuration access is not enforced by the IC with hardware; it is expected that this requirement be enforced through software.

The hot plug software driver may be used to inhibit configuration accesses to a slot after commands that result in deassertions of RST# to the slot are executed. The set bus speed/mode command and the enable slot command result deassertions of RST#. Each of these commands complete in less than 250 milliseconds after being received by the IC. The PCI requirement results in a 0.5 to 1.0 second period (depending on the bus frequency) after the deassertion of RST# during which configuration accesses to the slot are not allowed. Therefore, if the hot plug driver inhibits configuration accesses to the slot for 1.25 seconds after these commands are sent to the IC, the PCI requirement should be satisfied.

4.5.4 PCI-X® PHY Compensation Update

The PCI-X PHY calculated compensation values may change at any time. These may be altered, based on DevA:0x[54, 50], before being passed on to the PHY. The IC ensures that the PCI-X bridges are idle when new values are passed to the PHY. The following logic is implemented to accomplish this:

- Shortly after reset, or whenever a new value is written to DevA:0x[54, 50], and every 16 milliseconds thereafter, the IC determines if any of the values that are to be present to the PHY have changed and therefore need to be updated. If the values have not changed, no action takes place until the next 16 millisecond period passes and the values are checked again.
- If the compensation values are to be updated, then the logic request control over the bridges for the compensation update.
- The logic ensures that the bridges are idle for at least two PCLK cycles before and 4 PCLK cycles after the new values are passed to the PHY. During these six or more PCLK cycles, the IC drives the value of 0123_4567_89AB_CDEFh onto [B, A]_AD[63:0].

4.5.5 Transactions Claimed By The Bridges

The bridges claim no upstream transactions. They claim the following downstream transactions:

- All memory and IO space specified by Dev[B, A]:0x[30:1C].
- All configuration cycles to the implemented functions of DevA or DevB (see section 5.1.2).
- All configuration cycles to buses behind bridge A and bridge B.
- All EOI broadcasts are passed to the IOAPIC.
- All Stop Grant and STPCLK broadcasts are observed for clock gating (see section 4.3.3).
- If DevA:0x48[COMPAT]=1, then all memory space, IO space, and interrupt acknowledge packets in which the COMPAT bit is set are claimed and passed to bridge A.

Per the link protocol, when the COMPAT bit is set in the transaction and DevA:0x48[COMPAT]=0, then the IC never claims the transaction. Such transactions are automatically passed to the other side of the tunnel (or master aborted if the IC is at the end of the chain).

4.5.6 Various Behaviors

- Cacheline-wrap mode is not supported. If a transaction is initiated that indicates this protocol, it is disconnected at the first data phase.
- Downstream special cycles that are encoded in configuration cycles to device 31 of the bridge's secondary bus number (per the PCI-to-PCI bridge specification) are translated to special cycles on the bridge.
- Secondary-bus configuration cycles are never claimed by the IC (including configuration cycles to device 31 in which special cycles are encoded per the PCI-to-PCI bridge specification).
- In the translation from type 1 link configuration cycles to secondary bus type 0 configuration cycles, the IC converts the device number to an IDSEL AD signal as follows: device 0 maps to AD[16]; device 1 maps to AD[17]; and so forth. Device numbers 16 through 31 are not valid.
- Transactions that cross address space boundaries, as defined by the window configuration registers, Dev[B, A]:0x[30:1C], result in undefined behavior.
- If the bridge is in PCI-X mode and an upstream memory read or write request is issued with the No Snoop bit of the attribute field set high, then the coherent bit of the corresponding link read sized or write sized requests is low. The coherent bit is high for all other link requests, including all conventional PCI transactions and IO commands. The No Snoop field bit of the attribute field is always low in downstream requests to the PCI-X bridge.
- The following tables show the relationship between PCI-X transactions in which the relaxed ordering bit is set and link packets:

Downstream link transaction	Corresponding PCI-X® transaction
A read request in which bit[3] of the command field (response may pass posted write) is set.	Relaxed ordering bit of the attribute field is set.
A response in which PassPW is set.	Relaxed ordering bit of the attribute field is set.
A posted memory write in which the PassPW bit is set.	No effect; the relaxed ordering bit is zero regardless of the state of the PassPW bit.

Upstream PCI-X® transaction	Corresponding link transaction
A read request in which the relaxed ordering bit of the attribute field is set.	Bit[3] of the command field (response may pass posted write) in the read request is set.
A split completion in which the relaxed ordering bit of the attribute field is set.	PassPW is set in the response.
An immediate response to a downstream link read request in which bit[3] of the command field (response may pass posted write) is set.	PassPW is set in the response (even though there is no attribute field associated with the PCI-X response).
A posted memory write in which the relaxed ordering bit of the attribute field is set.	No effect; PassPW is zero regardless of the state of the relaxed ordering bit.

- If there is a downstream PCI-X request that results in a device-specific error in the completion message, then the response passed to the link indicates a target abort (error bit set; NXA clear).
- When the IC asserts [B, A]_DEVSEL#, it does so using the medium decoding clock in conventional PCI mode and the “B” decoding clock in PCI-X mode.
- If there is a link transaction to IO-space that targets a bridge and that crosses a naturally aligned DWORD boundary, then the IC does not send the transaction to the bus and the link response is a master abort (error bit set; NXA set).

4.5.7 Error Conditions And Handling

The tables below describe how the IC responds to error conditions.

Error handling for secondary bus responses to the IC.

Response from the secondary bus	Behavior of the IC
Signal master abort (DEVSEL not asserted); Dev[B, A]:0x3C[MARSP]=0.	<ul style="list-style-type: none"> • Link response with data of all F's for non-posted transaction. • Posted transaction is discarded. • Dev[B, A]:0x1C[RMA]=1.
Signal master abort (DEVSEL not asserted); Dev[B, A]:0x3C[MARSP]=1; non-posted request.	<ul style="list-style-type: none"> • Link response with data of all F's and the error bit set. • Dev[B, A]:0x04[STA]=1. • Dev[B, A]:0x1C[RMA]=1.
Signal master abort (DEVSEL not asserted); Dev[B, A]:0x3C[MARSP]=1; posted request.	<ul style="list-style-type: none"> • Posted transaction is discarded. • Dev[B, A]:0x1C[RMA]=1. • If Dev[B, A]:0x04[SERREN]=1, then: <ul style="list-style-type: none"> • Dev[B, A]:0x04[SSE]=1; and • Outgoing links are flooded with sync packets.
Signal master abort (DEVSEL not asserted) for a split completion message to the secondary bus.	<ul style="list-style-type: none"> • Split completion transaction is discarded. • Dev[B, A]:0x1C[RMA]=1. • Dev[B, A]:0xA0[SCD]=1. • If Dev[B, A]:0x04[SERREN]=1, then: <ul style="list-style-type: none"> • Dev[B, A]:0x04[SSE]=1; and • Outgoing links are flooded with sync packets.
Signal target abort; non-posted request.	<ul style="list-style-type: none"> • Link response with data of all F's and the error bit set. • Dev[B, A]:0x04[STA]=1. • Dev[B, A]:0x1C[RTA]=1.
Signal target abort; posted request.	<ul style="list-style-type: none"> • Remainder of posted transaction is discarded. • Dev[B, A]:0x1C[RTA]=1. • If Dev[B, A]:0x04[SERREN]=1, then: <ul style="list-style-type: none"> • Dev[B, A]:0x04[SSE]=1; and • Outgoing links are flooded with sync packets. Link response with data of all F's and the error bit set.
Signal target abort for a split completion message to the secondary bus.	<ul style="list-style-type: none"> • Split completion transaction is discarded. • Dev[B, A]:0x1C[RTA]=1. • Dev[B, A]:0xA0[SCD]=1. • If Dev[B, A]:0x04[SERREN]=1, then: <ul style="list-style-type: none"> • Dev[B, A]:0x04[SSE]=1; and • Outgoing links are flooded with sync packets.
Split completion error; PCI-X bridge error; master abort Dev[B, A]:0x3C[MARSP]=0.	<ul style="list-style-type: none"> • Link response with data of all F's for non-posted transaction. • Dev[B, A]:0x1C[RMA]=1.
Split completion error; PCI-X bridge error; master abort Dev[B, A]:0x3C[MARSP]=1.	<ul style="list-style-type: none"> • Link response with data of all F's and the error bit set. • Dev[B, A]:0x04[STA]=1. • Dev[B, A]:0x1C[RMA]=1.
Split completion error; PCI-X bridge error; target abort.	<ul style="list-style-type: none"> • Link response with data of all F's and the error bit set. • Dev[B, A]:0x04[STA]=1. • Dev[B, A]:0x1C[RTA]=1.

Split completion error; PCI-X bridge error; write data parity error.	<ul style="list-style-type: none"> • A TargetDone response is returned with no error indicated. • If Dev[B, A]:0x3C[PEREN]=1, then: <ul style="list-style-type: none"> • Dev[B, A]:0x1C[MDPE]=1.
Split completion error; completer error; byte count out of range.	<ul style="list-style-type: none"> • Link response with data of all F's and the error bit set. • Dev[B, A]:0x04[STA]=1.
Split completion error; completer error; split write data parity error.	<ul style="list-style-type: none"> • A TargetDone response is returned with no error indicated. • If Dev[B, A]:0x3C[PEREN]=1, then: <ul style="list-style-type: none"> • Dev[B, A]:0x1C[MDPE]=1.
Split completion error; completer error; device specific error.	<ul style="list-style-type: none"> • Link response with data of all F's and the error bit set. • Dev[B, A]:0x04[STA]=1.

Error handling error conditions detected by the IC.

Error condition to IC	Behavior of the IC
Link response to the bridge indicates a master abort; PCI-X mode.	<ul style="list-style-type: none"> • Split completion indicates a master abort PCI-X bridge error. • Dev[B, A]:0x04[RMA]=1.
Link response to the bridge indicates a master abort; conventional PCI mode; Dev[B, A]:0x3C[MARSP]=0.	<ul style="list-style-type: none"> • Data of all F's is returned. • Dev[B, A]:0x04[RMA]=1.
Link response to the bridge indicates a master abort; conventional PCI mode; Dev[B, A]:0x3C[MARSP]=1.	<ul style="list-style-type: none"> • Target abort signaled on the PCI bus. • Dev[B, A]:0x04[RMA]=1. • Dev[B, A]:0x1C[STA]=1.
Link response to the bridge indicates a target abort in conventional PCI mode.	<ul style="list-style-type: none"> • Signal target abort on PCI bus; the rest of the data associated with transaction is discarded. • Dev[B, A]:0x04[RTA]=1. • Dev[B, A]:0x1C[STA]=1.
Link response to the bridge indicates a target abort in PCI-X mode.	<ul style="list-style-type: none"> • Split completion indicates a target abort PCI-X bridge error; the rest of the data associated with the PCI-X transaction is discarded. • Dev[B, A]:0x04[RTA]=1. • Dev[B, A]:0x1C[STA]=1.
Parity error is detected in address or attribute phase of a transaction from the secondary bus to the IC.	<ul style="list-style-type: none"> • Dev[B, A]:0x1C[DPE]=1. • If [B, A]:0x3C[PEREN]=1, then: <ul style="list-style-type: none"> • The IC claims the transaction and terminates with a target abort; and • Dev[B, A]:0x1C[STA]=1; and • If Dev[B, A]:0x04[SERREN]=1, then: <ul style="list-style-type: none"> • Dev[B, A]:0x04[SSE]=1; and • Both links are flooded with sync packets.
Parity error detected in the data phase of a split completion message from the secondary bus to the IC.	<ul style="list-style-type: none"> • Dev[B, A]:0x1C[DPE]=1. • If [B, A]:0x3C[PEREN]=1, then: <ul style="list-style-type: none"> • The IC claims the transaction and terminated with a target abort; and • Dev[B, A]:0x1C[MDPE]=1; and • If Dev[B, A]:0x04[SERREN]=1, then: <ul style="list-style-type: none"> • Dev[B, A]:0x04[SSE]=1; and • Both links are flooded with sync packets.

Parity error detected in the data phase of a split completion transaction from the secondary bus to the IC.	<ul style="list-style-type: none"> • The data is passed to the link as read by the IC. • Dev[B, A]:0x1C[DPE]=1. • If [B, A]:0x3C[PEREN]=1, then: <ul style="list-style-type: none"> • Dev[B, A]:0x1C[MDPE]=1; and • [B,A] PERR# is asserted; and • If Dev[B, A]:0x44[NMIEN]=1, then: <ul style="list-style-type: none"> • NMI is generated.
Parity error is detected in the data phase of a read response from the secondary bus to the IC.	<ul style="list-style-type: none"> • The data is passed to the link as read by the IC. • Dev[B, A]:0x1C[DPE]=1. • If Dev[B, A]:0x3C[PEREN]=1, then: <ul style="list-style-type: none"> • Dev[B, A]:0x1C[MDPE]=1; and • [B, A]_PERR# is asserted; and • If Dev[B, A]:0x44[NMIEN]=1, then: <ul style="list-style-type: none"> • NMI is generated.
[B, A]_PERR# is detected asserted after a data phase of a read split completion or read immediate response from the IC to the secondary bus.	<ul style="list-style-type: none"> • If Dev[B, A]:0x44[NMIEN]=1, then: <ul style="list-style-type: none"> • NMI is generated.
[B, A]_PERR# is detected asserted after the data phase of a non-posted write from the IC to the secondary bus.	<ul style="list-style-type: none"> • A TargetDone response is returned with no error indicated. • If Dev[B, A]:0x3C[PEREN]=1, then: <ul style="list-style-type: none"> • Dev[B, A]:0x1C[MDPE]=1. • If Dev[B, A]:0x44[NMIEN]=1, then: <ul style="list-style-type: none"> • NMI is generated.
Parity error is detected in the data phase of a posted or non-posted write from the secondary bus to the IC.	<ul style="list-style-type: none"> • The data is passed to the link as received by the IC. • Dev[B, A]:0x1C[DPE]=1. • If Dev[B, A]:0x3C[PEREN]=1, then: <ul style="list-style-type: none"> • [B, A]_PERR# is asserted; and • If Dev[B, A]:0x44[NMIEN]=1, then: <ul style="list-style-type: none"> • NMI is generated.
[B, A]_PERR# is detected asserted after the data phase of a posted write from the IC to the secondary bus.	<ul style="list-style-type: none"> • If Dev[B, A]:0x3C[PEREN]=1, then: <ul style="list-style-type: none"> • Dev[B, A]:0x1C[MDPE]=1; and • If Dev[B, A]:0x04[SERREN]=1, then: <ul style="list-style-type: none"> • Dev[B, A]:0x04[SSE]=1; and • Both links are flooded with sync packets. • If Dev[B, A]:0x44[NMIEN]=1, then: <ul style="list-style-type: none"> • NMI is generated.
[B, A]_SERR# or [B, A]_SHPC_SERR assertion is detected.	<ul style="list-style-type: none"> • Dev[B, A]:0x1C[RSE]=1. • If both Dev[B, A]:0x04[SERREN]=1 and Dev[B, A]:0x3C[SERREN]=1, then: <ul style="list-style-type: none"> • Dev[B, A]:0x04[SSE]=1; and • Both links are flooded with sync packets. • If Dev[B, A]:0x44[NMIEN]=1, then: <ul style="list-style-type: none"> • NMI is generated.

Link CRC error is detected.	<ul style="list-style-type: none"> • Appropriate bit of DevA:0x[C8:C4][CRCERR]=1. • If both Dev[B, A]:0x04[SERREN]=1 and DevA:0x[C8:C4][CRCFEN]=1, then: <ul style="list-style-type: none"> • Dev[B, A]:0x04[SSE]=1; and • Outgoing links are flooded with sync packets.
Link incoming sync flood is detected.	<ul style="list-style-type: none"> • Outgoing links are flooded with sync packets.
Discard timer time out (only in conventional PCI mode).	<ul style="list-style-type: none"> • Transaction is thrown out. • Dev[B, A]:0x3C[DTS]=1. • If both Dev[B, A]:0x3C[DTSE]=1 and Dev[B, A]:0x04[SERREN]=1, then: <ul style="list-style-type: none"> • Dev[B, A]:0x04[SSE]=1; and • Outgoing links are flooded with sync packets.
The IC detects an illegal address/byte enable combination during the address phase of a transaction from the secondary bus to the IC.	<ul style="list-style-type: none"> • The transaction is terminated by signaling a target abort. • Dev[B, A]:0x1C[STA]=1.

4.6 Performance-Related Information

4.6.1 Bandwidth Percentage

One important measure of the IC performance is the ratio of data bandwidth that can pass over the PCI bus to the theoretical maximum. This is called the bandwidth percentage or BWP. The theoretical maximum is defined as the product of the bus width and the bus frequency. This does not account for clocks that are required by the protocol for purposes other than data transfer, e.g., address phases. The IC is implemented in such a way that when external PCI devices generate long data length transactions, the clocks that are lost to the protocol become less significant to the BWP. If external PCI devices generate many small data length transactions, then this acts to reduce the BWP.

Memory read requests and write requests from external PCI masters have different BWP characteristics. Generally speaking, memory writes are handled with high BWP values because they use the link posted channel. Write requests are converted from PCI protocol to link protocol by the IC efficiently. Memory reads, however, use the link non-posted channel. The response to each read request cannot be provided to the PCI bus until it is available. The time required for the response to reach the PCI bus includes the time for the request to pass from the PCI bus to the IC and from the IC to the host, and for the response to pass from the host back to the IC and from the IC out to the PCI bus. If there are other tunnel devices between the IC and the host, these will act to further increase this latency. If, for example, a single PCI master generates a pattern of (1) a single-cacheline (64 bytes) memory read request, (2) a burst of the response data to the PCI bus, (3) repeat, and that is the only activity on the bus, then the BWP is low; there will be many idle clocks, waiting for the response, between each cacheline burst on the PCI bus.

BWP may be improved by support for multiple, simultaneous read requests. In conventional PCI mode and PCI-X mode, the IC supports up to 8 independent PCI read requests simultaneously, per bridge. If, as in the example above, two masters generate a pattern of (1) a single-cacheline memory read request, (2) a burst of the response data to the PCI bus, (3) repeat, then, the total BWP may roughly double.

In PCI-X mode, long data length reads result in a high BWP because the time spent transferring data on the PCI bus becomes large compared to the time waiting for the response data.

In conventional PCI mode, the size of the read requests is not provided by the protocol. Instead, the memory read command code provides hints as to the amount of data that may be required by the external PCI master. Dev[B, A]:0x4C specifies the number of cachelines that are prefetched from the host based on the PCI command code. The initial prefetch value in this register should be balanced based on the requirements of the system. If it is too high (such that the master does not use all the data), then unnecessary memory read requests are generated and the corresponding data is thrown out. If it is too low, then the response latency is not well covered, resulting in a low BWP.

The logic for each bridge can generate up to 29 link read requests, each request for up to 64 bytes (one cacheline) of data, in support of read commands generated by external PCI-bus devices. These link requests are generated by the IC in the order in which they are received from the PCI bus. However, the responses to these PCI read requests may be provided on the PCI bus in a different order from the order in which the requests were received by the IC; this reordering depends on a number of factors including when the data for each request is provided by the host and if the master is ready to accept the data.

The following table provides some bandwidth percentages measured in an ideal-model simulation environment. This data is provided for guidance, with no guarantee that it represents the exact behavior of a real-world system design. For these measurements:

- All simulations are run while multiple, simultaneous requests are being processed. Read measurements are taken on the PCI bus, after host latency has passed, such that host latency does not affect the results.
- No other traffic is present except the requests associated with the measurements.
- Conventional PCI numbers presume that prefetching is enabled in continuous mode and the initial prefetch value is programmed to be high enough to cover the round-trip latency of the response. Thus, all cachelines of the response to a request can be burst onto the PCI bus without interruption.
- These results apply to all PCI-bus frequencies supported by the IC.
- *Number of cachelines* specifies the number of 64-byte cachelines requested by the PCI master. Note that the overhead is the same, regardless of the number of cachelines requested.
- *Total clocks* specifies the number of clocks from the end of one burst, associated with one request, to the end of the next burst, associated with another request. It is the sum of the burst clocks and overhead clocks.
- *Overhead clocks* specifies the number of PCI bus clocks during which there is no data transfer due to PCI protocol overhead and overhead inherent to the IC.
- *Data burst clocks* specifies the number of PCI bus clocks during which data is transferred.
- *BWP* is bandwidth percentage. $BWP = (\text{BurstClks} / \text{TotalClks}) * 100$.

Type	Number of cachelines	PCI bus width	Total clocks	Over-head clocks	Data burst clocks	BWP
Conventional PCI writes to host memory	2	32 bit	37	5	32	86
		64 bit	21	5	16	76
	8	32 bit	133	5	128	96
		64 bit	69	5	64	93
	32	32 bit	517	5	512	99
		64 bit	261	5	256	98
Conventional PCI reads from host memory	2	32 bit	39	7	32	82
		64 bit	23	7	16	70
	8	32 bit	135	7	128	95
		64 bit	71	7	64	90
	32	32 bit	519	7	512	99
		64 bit	263	7	256	97
PCI-X writes to host memory	2	32 bit	41	9	32	78
		64 bit	25	9	16	64
	8	32 bit	137	9	128	93
		64 bit	73	9	64	88
	32	32 bit	521	9	512	98
		64 bit	265	9	256	97
PCI-X reads from host memory	2	32 bit	41	9	32	78
		64 bit	25	9	16	64
	8	32 bit	137	9	128	93
		64 bit	73	9	64	88
	32	32 bit	521	9	512	98
		64 bit	265	9	256	97

Table 6: Bandwidth percentages.

4.6.2 Latency

The following table provides some latency values measured in an ideal-model simulation environment. This data is provided for guidance, with no guarantee that it represents the exact behavior of a real-world system design. For these measurements: link A is 16 bits, 800 MHz; link B is 8 bits, 800 MHz; the host is connected to link A; no other traffic is present except the transactions described; latency is measured from the first clock in which the transaction is clocked into the IC until the first clock in which the transaction is clocked out of the IC.

Latency (ns)	Description
75-85	Read response from link A to link B.
65-75	Read response from link B to link A.
225-235	Read request from conventional PCI, 33 MHz, any width (32 or 64 bit), to link A.
300-310	Read response from link A to conventional PCI, 33 MHz, any width.
135-145	Read request from conventional PCI, 66 MHz, any width, to link A.
180-190	Read response from link A to conventional PCI, 66 MHz, any width.
180-190	Read request from PCI-X®, 66 MHz, any width, to link A.
180-190	Read response from link A to PCI-X®, 66 MHz, any width.
130-140	Read request from PCI-X®, 100 MHz, any width, to link A.
145-155	Read response from link A to PCI-X®, 100 MHz, any width.
110-120	Read request from PCI-X®, 133 MHz, any width, to link A.
135-145	Read response from link A to PCI-X®, 133 MHz, any width.

Table 7: Some latencies.

5 Registers

5.1 Register Overview

The IC includes several sets of registers accessed through a variety of address spaces. IO address space refers to register addresses that are accessed through x86 IO instructions such as IN and OUT. PCI configuration space is typically accessed by the host through IO cycles to CF8h and CFCh. There is also memory space and indexed address space in the IC.

5.1.1 Configuration Space

The address space for the IC configuration registers is broken up into *busses*, *devices*, *functions*, and, *offsets*, as defined by the link specification. It is accessed by HyperTransport™ technology-defined type 0 configuration cycles. The device number is mapped into bits[15:11] of the configuration address. The function number is mapped into bits[10:8] of the configuration address. The offset is mapped to bits[7:2] of the configuration address.

The following diagram shows the devices in configuration space as viewed by software.

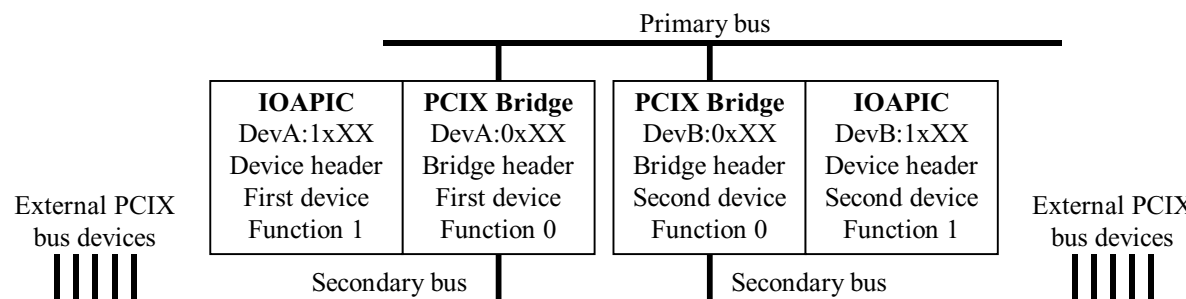


Figure 12: Configuration space.

Device A, above, is programmed to be the link base UnitID and device B is the link base UnitID plus 1.

5.1.2 Register Naming and Description Conventions

Configuration register locations are referenced with mnemonics that take the form of Dev[A|B]:[7:0]x[FF:0], where the first set of brackets contain the device number, the second set of brackets contain the function number, and the last set of brackets contain the offset.

Other register locations (e.g., memory mapped registers) are referenced with an assigned mnemonic that specifies the address space and offset. These mnemonics start with two or three characters that identify the space followed by characters that identify the offset within the space.

Register fields within register locations are also identified with a name or bit group in brackets following the register location mnemonic.

The following are configuration spaces:

Device	Function	Mnemonic	Registers
"A"	0	DevA:0xXX	PCI-PCI bridge A registers; link and PCI-X® capabilities block
"A"	1	DevA:1xXX	IOAPIC for PCI-X® bridge A.
"B"	0	DevB:0xXX	PCI-PCI bridge B registers; PCI-X® capabilities block
"B"	1	DevB:1xXX	IOAPIC for PCI-X® bridge B.

Table 8: Configuration spaces.

The IC does not claim configuration-register accesses to unimplemented functions within its devices (they are forwarded to the other side of the tunnel). Accesses to unimplemented register locations within implemented functions are claimed; such writes are ignored and reads always respond with all zeros.

The following are memory mapped spaces:

Base address register	Size (bytes)	Mnemonic	Registers
Dev[B,A]:1x10/48	4K	IOAXX	IOAPIC registers. Base address register at offset 10h enabled by Dev[B, A]:1x44[OSVISBAR].
Dev[B,A]:0x10	4K	SHPC[B,A]:XX	Standard hot plug controller register set. Access to these registers is enabled by DevA:0x48[HPENB, HPENA]. Access to these registers is provided through both memory space and configuration space; to access through configuration space, Dev[B, A]:0x90[SELECT] specifies the DWORD offset and Dev[B, A]:0x94 provides the DWORD data port.

Table 9: Memory mapped address spaces.

The following are register attributes found in the register descriptions.

Type	Description
Read or read-only	Capable of being read by software. Read-only implies that the register cannot be written to by software.
Write	Capable of being written by software.
Set by hardware	Register bit is set high by hardware.
Write once	After RESET#, these registers may be written to once. After being written, they become read only until the next RESET# assertion. The write-once control is byte based. So, for example, software may write each byte of a write-once DWORD as four individual transactions. As each byte is written, that byte becomes read only.
Write 1 to clear	Software must write a 1 to the bit in order to clear it. Writing a 0 to these bits has no effect.
Write 1 only	Software can set the bit high by writing a 1 to it. However subsequent writes of 0 will have no effect. RESET# must be asserted in order to clear the bit.

Table 10: Register attributes.

5.2 PCI-X® Bridge Configuration Registers

These registers are located in PCI configuration space, in the first device (device A) and second device (device B), function 0. See section 5.1.2 for a description of the register naming convention.

PCI-X Bridge Vendor And Device ID Register**Dev[B, A]:0x00**

Default: 7450 1022h

Attribute: Read only.

Bits	Description
31:16	PCI bridge device ID.
15:0	Vendor ID.

PCI-X Bridge Status And Command Register**Dev[B, A]:0x04**

Default: 0230 0000h

Attribute: See below.

Bits	Description
31	DPE: detected parity error. Read only. This bit is fixed in the low state.
30	SSE: signaled system error. Read; set by hardware; write 1 to clear. 1=A system error was signaled (both links were flooded with sync packets). This bit cannot be set unless Dev[B, A]:0x04[SERREN] is high. Note: this bit is cleared by PWROK reset but not by RESET#.
29	RMA: received master abort. Read; set by hardware; write 1 to clear. 1=A request sent to the host bus received a master abort (an NXA error response). Note: this bit is cleared by PWROK reset but not by RESET#.
28	RTA: received target abort. Read; set by hardware; write 1 to clear. 1=A request sent to the host bus received a target abort (a non-NXA error response). Note: this bit is cleared by PWROK reset but not by RESET#.
27	STA: Signaled target abort. Read; set by hardware; write 1 to clear. 1=A target abort was signaled to the host (a non-NXA error response). Note: this bit is cleared by PWROK reset but not by RESET#.
26:21	Read only. These bits are fixed in their default state.
20	Capabilities pointer. Read only. This bit is fixed in the high state.
19:9	Reserved
8	SERREN: SERR# enable. Read-write. 1=Dev[B, A]:0x04[SSE] is enabled to be set high in response to detected system errors. 0=Dev[B, A]:0x04[SSE] cannot be set high and the IC does not flood the links with sync packets.
7	Reserved.
6	PERSP: Parity error response. Read-write. This bit controls no hardware. It is provided for compatibility with the PCI-PCI bridge specification.
5	Reserved.
4	MWIEN: Memory write and invalidate enable. Read-write. This bit does not control any internal hardware; it is provided for compatibility with the PCI-PCI bridge specification.
3	Special cycle enable. Read only. This bit is hardwired low.
2	MASEN: PCI master enable. Read-write. 1=Enables secondary bus masters to initiate cycles to the host.
1	MEMEN: memory enable. Read-write. 1=Enables access to the secondary bus memory space and to the SHPC register space through the memory-space BAR, Dev[B, A]:0x10 (this bit does not affect access to SHPC registers through configuration space, Dev[B, A]:0x[94:90]).
0	IOEN: IO enable. Read-write. 1=Enables access to the secondary bus IO space.

PCI-X Bridge Revision and Class Code Register**Dev[B, A]:0x08**

Default: 0604 0?11h

Attribute: Read only.

Bits	Description
31:8	CLASSCODE. Provides the bridge class code as defined in the PCI specification. Bits[3:1] of this register are zero. DevA:0x08[8] is the same as DevA:0x48[COMPAT]. DevB:0x08[8] is zero.
7:0	REVISION.

PCI-X Bridge BIST-Header-Latency-Cache Register**Dev[B, A]:0x0C**

Default: 0081 ??00h

Attribute: See below.

Bits	Description
31:24	BIST. Read only. These bits are fixed at their default values.
23:16	HEADER. Read only. These bits are fixed at their default values.
15:8	LATENCY. Read-write. These bits control no hardware. The default value after the deassertion of RESET# is 00h when the bridge is in conventional PCI mode and 40h when the bridge is in PCI-X® mode.
7:0	CACHE. Read only. These bits are fixed at their default values.

PCI-X SHPC Base Address Register**Dev[B, A]:0x10**

This register is reserved if DevA:0x48[HPENB, HPENA] is low.

Default: 0000 0000 0000 0004h

Attribute: See below.

Bits	Description
63:12	SHPCBAR: SHPC base address register. Read-write. These bits specify the memory address space of the SHPC register set, SHPC[B, A]:xx. Note: bits[63:40] are required to be programmed low; setting any of these bits high results in undefined behavior.
11:0	Hardwired. Read only. These bits are all hardwired to their default state to indicate a 4K byte block of 64-bit, non-prefetchable memory space.

PCI-X Bridge Bus Numbers And Secondary Latency Register**Dev[B, A]:0x18**

Default: ??00 0000h

Attribute: See below.

Bits	Description
31:27	SECLAT[7:3] . Read-write. Secondary latency timer. The default value of SECLAT[7:0] after the deassertion of RESET# is 00h when the bridge is in conventional PCI mode and 40h when the bridge is in PCI-X® mode.
26:24	SECLAT[2:0] . Read only, 000b. Secondary latency timer.
23:16	SUBBUS . Read-write. Subordinate bus number.
15:8	SECBUS . Read-write. Secondary bus number.
7:0	PRIBUS . Read-write. Primary bus number.

PCI-X Bridge Memory Base-Limit Registers**Dev[B, A]:0x[30:1C]**

These registers specify the IO-space (Dev[B, A]:0x1C and Dev[B, A]:0x30), non-prefetchable memory-space (Dev[B, A]:0x20), and prefetchable memory-space (Dev[B, A]:0x24, Dev[B, A]:0x28, and Dev[B, A]:0x2C) address windows for transactions that are mapped from the 40-bit link address space to the secondary PCI bus.

The links support 25 bits of IO space. PCI-X supports 32 bits of IO space. Host accesses to the link-defined IO region are mapped to the PCI-X IO window with the 7 MSB always zero. PCI-X IO accesses in which any of the 7 MSBs are other than zero are ignored. The PCI-X IO space window is defined as follows:

```
PCI-X IO window =
  { 7'h00, Dev[B,A]:30[24:16], Dev[B,A]:0x1C[15:12], 12'hFFF } >= address >=
  { 7'h00, Dev[B,A]:30[8:0], Dev[B,A]:0x1C[7:4], 12'h000 };
```

The links and PCI-X support 40 bits of memory space. The PCI-X non-prefetchable memory space window is defined as follows:

```
PCI-X non-prefetchable memory window =
  { 24'h00, DevA:0xD8[15:8], Dev[B,A]:0x20[31:20], 20'hF_FFFF } >= address >=
  { 24'h00, DevA:0xD8[7:0], Dev[B,A]:0x20[15:4], 20'h0_0000 };
```

The links support 40 bits of memory space. PCI-X supports 64 bits of prefetchable memory space. All link memory mapped IO space may be within the PCI-X prefetchable memory window. PCI-X memory accesses in which any of bits[63:40] are other than zero are ignored. The PCI-X prefetchable memory space window is defined as follows:

```
PCI-X prefetchable memory window =
  { 24'h0, Dev[B,A]:2C[7:0], Dev[B,A]:0x24[31:20], 20'hF_FFFF } >= address >=
  { 24'h0, Dev[B,A]:28[7:0], Dev[B,A]:0x24[15:4], 20'h0_0000 };
```

These windows may also be altered by Dev[B, A]:0x3C[VGAEN, ISAEN]. When the address (from either the host or from a secondary bus master) is inside one of the windows, then the transaction is assumed to be intended for a target that sits on the secondary bus. Therefore, the following transactions are possible:

- Host-initiated transactions inside the windows are routed to the secondary bus.
- Secondary PCI-initiated transactions inside the windows are not claimed by the IC.

- Host initiated transactions outside the windows are passed through the tunnel or master aborted if the IC is at the end of a chain.
- Secondary PCI-initiated transactions outside the windows are claimed by the IC using medium decoding and passed to the host.

So, for example, if IOBASE > IOLIM, then no host-initiated IO-space transactions are forwarded to the secondary bus and all secondary-PCI-bus-initiated IO-space (not configuration) transactions are forwarded to the host. If MEMBASE > MEMLIM and PMEMBASE > PМЕMLIM, then no host-initiated memory-space transactions are forwarded to the secondary bus and all secondary-PCI-bus-initiated memory-space transactions are forwarded to the host.

The window may be affected by DevA:0x48[COMPAT] as well.

Dev[B, A]:0x1C. Default: 0220 01F1h

Attribute: See below.

Bits	Description
31	DPE: detected parity error. Read; set by hardware; write 1 to clear. 1=The IC detected an address parity error as the target of a secondary bus cycle or a data parity error during a data phase of an upstream transaction.
30	RSE: received system error. Read; set by hardware; write 1 to clear. 1=The IC detected that either [B, A]_SERR# or [B, A]_SHPC_SERR is asserted. In order to clear this bit, these signals must be deasserted. Note: this bit is cleared by PWROK reset but not by RESET#.
29	RMA: received master abort. Read; set by hardware; write 1 to clear. 1=The IC received a master abort as a master on the secondary bus. Note: this bit is cleared by PWROK reset but not by RESET#.
28	RTA: received target abort. Read; set by hardware; write 1 to clear. 1=The IC received a target abort as a master on the secondary PCI bus. Note: this bit is cleared by PWROK reset but not by RESET#.
27	STA: signaled target abort. Read; set by hardware; write 1 to clear. 1=The IC generated a target abort as a target on the secondary PCI bus. Note: this bit is cleared by PWROK reset but not by RESET#.
26:25	Device select timing. Read only. These bits are hard wired to indicate medium decoding.
24	MDPE: master data parity error. Read; set by hardware; write 1 to clear. 1=The IC detected a parity error during a data phase of a read or detected [B, A]_PERR# asserted during a write as a master on the secondary bus and Dev[B, A]:0x3C[PEREN] is set.
23	FBBEN: fast back to back enable. Read only. This bit is fix in the low state to indicate that the IC does not support fast back to back transactions from different masters.
22:16	Read only. These bits are fixed in their default state.
15:12	IOLIM. IO limit address bits[15:12]. See Dev[B, A]:0x[30:1C] above.
11:8	Reserved.
7:4	IOBASE. IO base address bits[15:12]. See Dev[B, A]:0x[30:1C] above.
3:0	Reserved.

Dev[B, A]:0x20. Default: 0000 FFF0h

Attribute: Read-write.

Bits	Description
31:20	MEMLIM. Non-prefetchable memory limit address bits[31:20]. See Dev[B, A]:0x[30:1C] above.

19:16	Reserved.
15:4	MEMBASE. Non-prefetchable memory base address bits[31:20]. See Dev[B, A]:0x[30:1C] above.
3:0	Reserved.

Dev[B, A]:0x24. Default: 0001 FFF1h Attribute: Read-write.

Bits	Description
31:20	PMEMLIM. Prefetchable memory limit address bits[31:20]. See Dev[B, A]:0x[30:1C] above.
19:16	Reserved.
15:4	PMEMBASE. Prefetchable memory base address bits[31:20]. See Dev[B, A]:0x[30:1C] above.
3:0	Reserved.

Dev[B, A]:0x28. Default: 0000 0000h Attribute: Read-write.

Bits	Description
31:0	PMEMBASE. Prefetchable memory base address bits[63:32]. See Dev[B, A]:0x[30:1C] above.

Dev[B, A]:0x2C. Default: 0000 0000h Attribute: Read-write.

Bits	Description
31:0	PMEMLIM. Prefetchable memory limit address bits[63:32]. See Dev[B, A]:0x[30:1C] above.

Dev[B, A]:0x30. Default: 0000 FFFFh Attribute: Read-write.

Bits	Description
31:16	IOLIM. IO limit address bits[31:16]. See Dev[B, A]:0x[30:1C] above.
15:0	IOBASE. IO base address bits[31:16]. See Dev[B, A]:0x[30:1C] above.

PCI-X Bridge Capabilities Pointer Register

Dev[B, A]:0x34

Default: 0000 00A0h Attribute: Read only.

Bits	Description
31:8	Reserved.
7:0	CAPABILITIES_PTR. Specifies the offset to standard PCI-X® registers.

PCI-X Bridge Interrupt and Bridge Control Register

Dev[B, A]:0x3C

Default: 0000 0?FFh Attribute: See below.

Bits	Description
31:28	Reserved.

27	DTSE: discard timer sync flood enable. Read-write. 1=If both Dev[B, A]:0x04[SERREN] and Dev[B, A]:0x3C[DTS] are high, then Dev[B, A]:0x04[SSE] is set and the links are flooded with sync packets.
26	DTS: discard timer status. Read; set by hardware; write 1 to clear. 1=The 15-bit discard timer timed out. This bit is not capable of being set when the secondary bus is in PCI-X® mode. Note: this bit is cleared by PWROK reset but not by RESET#.
25:23	Reserved.
22	SBRST: secondary bus reset. Read-write. 1=[B, A]_PRESET# asserted; secondary PCI bus placed into reset state. 0=[B, A]_PRESET# not asserted. Note: the PCI requirement for a delay between the rising edge of RST# and the first configuration access is not enforced by the IC with hardware; it is expected that this requirement be enforced through software.
21	MARSP: master abort response. Read-write. 1=The response to non-posted requests that come from the host bus or secondary bus that results in a master aborts will indicate a target abort to the initiating bus (through PCI bus protocol or link protocol); posted requests that are master aborted result in assertion of Dev[B, A]:0x04[SSE]. 0=Master aborts result in normal responses; read responses are sent with the appropriate amount of data, which are all 1s, and writes are ignored.
20	Reserved.
19	VGAEN: VGA decoding enable. Read-write. 1=Route host-initiated commands targeting VGA-compatible address ranges to the secondary bus. These include memory accesses from A0000h to BFFFFh (within the bottom megabyte of memory space only), IO accesses in which address bits[9:0] range from 3B0h to 3BBh or 3C0h to 3DFh (address bits[15:10] are not decoded, regardless of Dev[B, A]:0x3C[ISAEN]; also this only applies to the first 64K of IO space; i.e., address bits[31:16] must be low). 0=The IC does not decode VGA-compatible address ranges.
18	ISAEN: ISA decoding enable. Read-write. 1=The IO address window specified by Dev[B, A]:0x1C[15:0] and Dev[B, A]:0x30 is limited to the first 256 bytes of each 1K byte block specified; this only applies to the first 64K bytes of IO space. 0=The PCI IO window is the whole range specified by Dev[B, A]:0x1C[15:0] and Dev[B, A]:0x30.
17	SERREN: system error enable. Read-write. If Dev[B, A]:0x04[SERREN] and Dev[B, A]:0x3C[SERREN] are both high and if [B, A]_SERR# or [B, A]_SHPC_SERR is detected asserted (Dev[B, A]:0x1C[RSE] = 1), then the IC responds by flooding the outgoing link with sync packets and sets Dev[B, A]:0x04[SSE]. If either Dev[B, A]:0x04[SERREN] or Dev[B, A]:0x3C[SERREN] are low, then Dev[B, A]:0x1C[RSE] does not stop link operation or cause Dev[B, A]:0x04[SSE] to be set.
16	PEREN: parity error response enable. Read-write. 1=Enable parity error detection on secondary PCI interface (see Dev[B, A]:0x1C[MDPE]); [B, A]_PERR# signal enabled to set status bit or be driven. 0=Dev[B, A]:0x1C[MDPE] cannot be set; [B, A]_PERR# signal is ignored and it is not driven by the IC.
15:8	INTERRUPT_PIN. Read only. If DevA:0x48[HPENB, HPENA] is low, then Dev[B, A]:0x3C[INTERRUPT_PIN] is 00h. If DevA:0x48[HPENB, HPENA] is high, then Dev[B, A]:0x3C[INTERRUPT_PIN] is 01h. When hot plug mode is enabled, [B, A]_PIRQA# can be asserted for hot plug events.
7:0	INTERRUPT_LINE. Read-write. These bits control no internal logic.

PCI-X Miscellaneous Register

Dev[B, A]:0x40

Default: 001F 0001h.

Attribute: See below.

Bits	Description
31:28	PCLKDEL: [B, A]_PCLK[4:0] delay. Read-write. Relative to [B, A]_PLLCLKO, [B, A]_PCLK[4:0] are delayed by PCLKDEL*TAP_DELAY, where TAP_DELAY is defined by DevA:0x48[BDCV]. When a new value is written to this field, there may be a glitch on the output clocks. If [B, A]_PCLK[4:0] are used as the reference clock to PLLs external to the IC, then changing this field may cause these PLLs to lose synchronization. For this reason, it is recommended that [B, A]_PRESET# be asserted (through Dev[B, A]:0x3C[SBRST]) while this field is updated and that [B, A]_PRESET# be deasserted at least 1 millisecond after this field is updated. See section 4.3 for more information about [B, A]_PCLK[4:0].
27:24	PLLODEL: PLLCLKO delay. Read-write. Relative to [B, A]_PCLK[4:0], [B, A]_PLLCLKO is delayed by PLLODEL*TAP_DELAY, where TAP_DELAY is defined by DevA:0x48[BDCV]. When a new value is written to this field, the clock delay is shifted such that there is no glitch on the output signal. Changing this field may cause the internal PLL that generates [B, A]_PCLK[4:0] to lose synchronization for a period of no more than 100 microseconds. For this reason, it is recommended that software alter this value by only one at a time. E.g., to change from 0 to 3, software should write a 1, then write a 2, then write a 3 to this field. See section 4.3 for more information about [B, A]_PLLCLKO.
23:21	Reserved.
20:16	PCLKEN: PCLK enable. Read-write. Each of these bits controls a [B, A]_PCLK[4:0] signal. Bit 16 controls PCLK 0, bit 17 controls PCLK1, and so forth. 1=The PCLK signal is enabled to toggle. 0=The PCLK signal is forced low. It is intended that this be used to disable PCLK signals that correspond to unimplemented PCI-X® devices or slots.
15:13	Reserved.
12:8	PFEN[4:0]#: prefetch enables (active low). Read-write. Each of these bits apply to one [B, A]_REQ#/GNT# pair (Dev[B, A]:0x40[PFEN0#] applies to [B, A]_REQ0#/GNT0# and so forth). 0=Prefetching is enabled for the specified external master in conventional PCI mode. 1=Prefetching is not enabled. When prefetching is not enabled, memory read requests from external masters are allowed to burst from the transaction starting address up to the 64-byte cacheline boundary, at which point the transaction is disconnected with data. If prefetching is enabled, while the burst is taking place, the next cacheline is read from the host such that the burst may be continued for multiple cachelines. This field is ignored when the bridge is in PCI-X® mode (Dev[B, A]:0xA0[SCF] is not zero). It is expected that these bits are normally left at 0.
7:5	Reserved.
4	NZSEQID: non-zero sequence ID. Read-write. 0=The SeqID value is 0h in the upstream link requests that result from the bridge's PCI master memory read requests. 1=The SeqID value is not zero in the upstream link requests that result from the bridge's PCI master memory read requests; if DevA:0x40[NZSEQID] is high, then a SeqID of 1h is generated for these transactions from bridge A; if DevB:0x40[NZSEQID] is high, then a SeqID of 2h is generated for these transactions from bridge B. This applies only to memory read requests from external masters. Setting these bits high may reduce host memory efficiency and bandwidth. It is not expected that these bits will need to be set; the order in which requests are delivered to destinations does not matter in most cases.
3	Must be low. Read-write. This bit is required to be low at all times; setting it high results in undefined behavior.

2	<p>HPSSS#: hot plug single-slot support (active low). Read only. The default state of DevA:0x40[HPSSS#] is captured off of HPSIC at the rising edge of PWROK. The default state of DevB:0x40[HPSSS#] is captured off of B_GNT2# at the rising edge of PWROK. 0=If the bridge is in hot plug mode (as specified by DevA:0x48[HPENB or HPENA]), then the bridge supports a single hot plug slot without external isolation switches. In this mode, external isolation switches between the IC and the slot are not required. See section 4.5.3 for details. 1=External isolation switches are required for all hot plug slots. This bit is required to be high if more than one external device is supported by the bridge. If HPSSS# is low while the DevA:0x48[HPENB, HPENA] bit corresponding to the same bridge is low, then undefined behavior results.</p>
1	<p>CPCI66: conventional PCI mode frequency. Read only. Dev[B, A]:0xA0[SCF]=0h, then the bridge is in conventional PCI mode and this bit is valid; otherwise its state is unknown. 0=[B, A]_PCLK[4:0] toggle at 33 MHz. 1=[B, A]_PCLK[4:0] toggle at 66 MHz. The default state for this field is determined by strapping options described in section 4.2.</p>
0	<p>NIOAMODE: non-IOAPIC mode. Read-write. This is used to enable [B, A]_PIRQ[D, C, B, A]# to the NIOAIRQ[D, C, B, A]# pins. 0=The state of the PIRQ[D:A]# pin is passed to the NIOAIRQ[D:A]# pin as if the PIRQ[D:A]# pin were always high. 1=The state of the PIRQ[D:A]# pin from the bridge is ANDed with the state from the other bridge and passed to the NIOAIRQ[D:A]# pin. This is shown in the following equations:</p> $\begin{aligned} \text{NIOAIRQA\#} &= \sim(\text{DevA:0x40[NIOAMODE]} \& \sim\text{A_PIRQA\#} \& \text{RDRA0[IM]} \\ & \quad \text{DevB:0x40[NIOAMODE]} \& \sim\text{B_PIRQA\#} \& \text{RDRB0[IM]}); \\ \text{NIOAIRQB\#} &= \sim(\text{DevA:0x40[NIOAMODE]} \& \sim\text{A_PIRQB\#} \& \text{RDRA1[IM]} \\ & \quad \text{DevB:0x40[NIOAMODE]} \& \sim\text{B_PIRQB\#} \& \text{RDRB1[IM]}); \end{aligned}$ <p>And similarly for NIOAIRQ[C and D]#. Were RDR[B, A][3:0][IM] is the interrupt mask field of the redirection register (see section 5.4); [B, A] = the bridge letter; [3:0] = the redirection register index. Note that the NIOAIRQ[D:A]# pins are open drain outputs. So a high on the PIRQ input is translated to the high-impedance state on the NIOAIRQ output. See section 4.5.2 for more details about interrupt routing. It is expected that this bit is normally left high by system BIOS.</p>

PCI-X Miscellaneous II Register
Dev[B, A]:0x44

Default: 0000 0000h.

Attribute: Read-write.

Bits	Description
31:1	<p>RW. Read-write. These bits control no hardware. These bits are reserved and should be left in the default state.</p>
0	<p>NMIEN: NMI on error enable. Read-write. 1=Assertions of the [B, A]_PERR# and [B, A]_SERR# signals (regardless of the source of the assertion) result in NMI interrupts; see section 4.5.2.1.</p>

Pins Latched At Boot Register**DevA:0x48**

The default value for bits in this register is latched at the rising edge of PWROK.

Default: 0000 000?h.

Attribute: See below.

Bits	Description
31:12	Reserved.
11:8	BDCV: buffer delay calculation value. Read only. This provides the approximate buffer delay calculation value used to determine the delay value of each tap in Dev[B, A]:0x40[PLLODEL and PCLKDEL] as follows: TAP_DELAY \approx 1250/BDCV picoseconds. Expected values for BDCV are from 5h to Fh.
7:4	Reserved.
3	HPENB: bridge B hot plug enable. Read only. This bit captures the state of HPSIL# at the rising edge of PWROK. 0=Hot plug mode is not enabled on bridge B. 1=Hot plug mode is enabled on bridge B. See section 4.5.3 for details. If this bit is low while DevB:0x40[HPSSS#] is low, then undefined behavior results.
2	HPENA: bridge A hot plug enable. Read only. This bit captures the state of HPSOD at the rising edge of PWROK. 0=Hot plug mode is not enabled on bridge A. 1=Hot plug mode is enabled on bridge A. See section 4.5.3 for details. If this bit is low while DevA:0x40[HPSSS#] is low, then undefined behavior results.
1	Reserved.
0	COMPAT: compatibility bus. Read-write. 1=The IC routes all host initiated accesses in which the link-defined compat bit is set to the secondary bus. The default state of this bit is latched off of A_COMPAT at the trailing edge of PWROK reset.

Prefetch Control Register**Dev[B, A]:0x4C**

This register specifies the prefetching policy when a bridge is in conventional PCI mode. This register is ignored when in PCI-X mode. It includes three sets of initial prefetch registers (IPF_x) and three corresponding sets of continuous prefetch enable registers (CPFEN_x): one for memory read multiple requests (x=MRM); one for memory read line requests (x=MRL); and one for memory read requests (x=MR). When a PCI master, for which prefetching is enabled through Dev[B, A]:0x40[PFEN#], initiates a host read with a command code of MRM, MRL, or MR, then the IC sends link read requests as follows: (1) an initial up-to-one cacheline request from the initial address to the end of the cacheline and (2) additional cachelines as specified by the IPF_x field that corresponds to the command code issued by the PCI master. When each cacheline of data starts to be transferred to the master over the PCI bus, an additional cacheline of data may be requested, as specified by the appropriate CPFEN_x bit. An *unrequested prefetch*, as specified in some of the fields of this register, is a speculative link request--or prefetch--generated by the IC for a cacheline of data beyond the cacheline address generated by the conventional PCI master; a *requested prefetch* is up-to-one cacheline of memory read data at an address generated by the conventional PCI master. See section 4.6.1 for information about how prefetching is related to performance.

Default: 0000 2C00h.

Attribute: Read-write.

Bits	Description
31:14	Reserved.

13	<p>CPFEN_MRM. Continuous prefetch enable for memory read multiple request. When a master initiates a burst read to the host with a memory read multiple command code, this specifies if continuous prefetching is enabled.</p> <p>1 = One new request for a cacheline of prefetch data is sent to the host by the IC when data from an earlier cacheline starts to be transferred to the requesting master over the PCI bus.</p> <p>0 = There are no new requests for prefetch data after the initial batch specified by IPF_MRM.</p>
12:10	<p>IPF_MRM: Initial prefetch for memory read multiple request. This specifies the number of additional cacheline prefetches (after the initial prefetch of up to one cacheline) when a PCI master initiates a burst read to the host with a memory read multiple command code. If prefetching is disabled in Dev[B, A]:0x40[PFEN#], then the value of this register is ignored. 0=no additional prefetches; 1=1 additional prefetch; 2=2 additional prefetches, and so forth.</p>
9	<p>CPFEN_MRL. Continuous prefetch enable for memory read line request. See CPFEN_MRM.</p>
8:6	<p>IPF_MRL: Initial prefetch for memory read line request. See IPF_MRM.</p>
5	<p>CPFEN_MR. Continuous prefetch enable for memory read request. See CPFEN_MRM.</p>
4:2	<p>IPF_MR: Initial prefetch for memory read request. See IPF_MRM.</p>
1	<p>DPDM: Discard unrequested prefetch data upon master request. 1=No further prefetching occurs and all unrequested prefetches are discarded when another master requests the PCI bus; also, unrequested prefetches are discarded if the discard timer reaches 16 PCLKs (requested prefetches are discarded if the discard timer reaches 32K PCLKs). 0=Requests from other masters do not affect prefetching; requested and unrequested prefetches are discarded if the discard timer reaches 32K PCLKs. This bit is typically programmed low by system BIOS.</p>
0	<p>DPDH: Discard unrequested prefetch data upon host request. 1=If the IC receives a host request to the PCI bus, then:</p> <ol style="list-style-type: none"> (1) if there is not an outstanding requested prefetch for a given previously-established PCI read request, then all of the unrequested prefetches associated with that PCI read request are discarded; or (2) if there is an outstanding requested prefetch for a given previously-established PCI read request, then the data for that requested prefetch, along with the data for subsequent unrequested prefetches, is allowed to burst onto the PCI bus until the burst is disconnected (either by the PCI master or by the IC because it does not possess the data necessary to continue the burst); when the burst is disconnected, any remaining unrequested prefetches associated with the PCI read request are discarded. (3) if there is a burst in progress on the PCI bus, the IC disconnects the burst at the next convenient cacheline boundary and discards any outstanding unrequested prefetches associated with the transaction. <p>0=Host requests do not affect prefetching.</p> <p>Programming of this bit may vary based on platform requirements. DPDH is typically programmed high by system BIOS to protect against stale prefetch-data scenarios, as described in the PCI specification, revision 2.3, section 3.10, point 6; scenarios similar to this have been observed, albeit rarely. However, if the secondary PCI bus includes a device that is accessed frequently as a target, then setting this bit may result in reduced memory read bandwidth. In such cases, it may be preferable to program this bit low.</p>

PCI-X PHY Compensation Control Registers**DevA:0x[54, 50]**

The PCI-X PHY circuitry includes automatic compensation that is used to adjust the drive strength of the PCI-X and other IO cells, including all output signals of the IC that are on the VDD33 power plane. The compensation circuits calculate the drive strength for the rising edge and falling edge of the outputs. These registers provide visibility into the calculated output of the compensation circuits, the ability to override the calculated value with software-controlled values, and the ability to offset the calculated values with a fixed difference. The overrides and difference values may be different between bridges A and B. These registers specify the compensation parameters as follows:

- DevA:0x50: output rising edge (P) drive strength compensation; associated with the P_CAL pin.
- DevA:0x54: output falling edge (N) drive strength compensation; associated with the P_CAL# pin.

Higher values in these registers represent higher drive strength; the values range from 0h to Fh (16 steps).

Default: 0000 0000.

Attribute: See below.

Bits	Description										
31	Must be low. Read-write. This bit is required to be low at all times; setting it high results in undefined behavior.										
30:20	Reserved.										
19:16	CALCCOMP: calculated compensation value. Read only. This provides the calculated value from the auto compensation circuitry. The default value of this field is not predictable. DevA:0x50[CALCCOMP] is affected by the value of the resistor connected to P_CAL and DevA:0x54[CALCCOMP] is affected by the value of the resistor connected to P_CAL#. In both cases, larger the values of resistors (measured in ohms) result in smaller CALCCOMP values.										
15	Reserved.										
14:13	BCTL: Bridge B PHY control value. Read-write. These two bits combine to specify the PHY compensation value that is applied to bridge B outputs as follows: <table border="1" style="margin-left: 20px;"> <thead> <tr> <th>BCTL</th> <th>Description</th> </tr> </thead> <tbody> <tr> <td>00b</td> <td>Apply CALCCOMP directly as the compensation value.</td> </tr> <tr> <td>01b</td> <td>Apply BDATA directly as the compensation value.</td> </tr> <tr> <td>10b</td> <td>Apply the sum of CALCCOMP and BDATA as the compensation value. If the sum exceeds Fh, then Fh is applied.</td> </tr> <tr> <td>11b</td> <td>Apply the difference of CALCCOMP minus BDATA as the compensation value. If the difference is less than 0h, then 0h is applied.</td> </tr> </tbody> </table>	BCTL	Description	00b	Apply CALCCOMP directly as the compensation value.	01b	Apply BDATA directly as the compensation value.	10b	Apply the sum of CALCCOMP and BDATA as the compensation value. If the sum exceeds Fh, then Fh is applied.	11b	Apply the difference of CALCCOMP minus BDATA as the compensation value. If the difference is less than 0h, then 0h is applied.
BCTL	Description										
00b	Apply CALCCOMP directly as the compensation value.										
01b	Apply BDATA directly as the compensation value.										
10b	Apply the sum of CALCCOMP and BDATA as the compensation value. If the sum exceeds Fh, then Fh is applied.										
11b	Apply the difference of CALCCOMP minus BDATA as the compensation value. If the difference is less than 0h, then 0h is applied.										
12	Reserved.										
11:8	BDATA: Bridge B data value. Read-write. This value is applied to the bridge B PHY compensation as described in BCTL.										
7	Reserved.										

6:5	ACTL: Bridge A PHY control value. Read-write. These two bits combine to specify the PHY compensation value that is applied to bridge A outputs as follows: <table border="1"> <thead> <tr> <th><u>ACTL</u></th> <th><u>Description</u></th> </tr> </thead> <tbody> <tr> <td>00b</td> <td>Apply CALCCOMP directly as the compensation value.</td> </tr> <tr> <td>01b</td> <td>Apply ADATA directly as the compensation value.</td> </tr> <tr> <td>10b</td> <td>Apply the sum of CALCCOMP and ADATA as the compensation value. If the sum exceeds Fh, then Fh is applied.</td> </tr> <tr> <td>11b</td> <td>Apply the difference of CALCCOMP minus ADATA as the compensation value. If the difference is less than 0h, then 0h is applied.</td> </tr> </tbody> </table>	<u>ACTL</u>	<u>Description</u>	00b	Apply CALCCOMP directly as the compensation value.	01b	Apply ADATA directly as the compensation value.	10b	Apply the sum of CALCCOMP and ADATA as the compensation value. If the sum exceeds Fh, then Fh is applied.	11b	Apply the difference of CALCCOMP minus ADATA as the compensation value. If the difference is less than 0h, then 0h is applied.
<u>ACTL</u>	<u>Description</u>										
00b	Apply CALCCOMP directly as the compensation value.										
01b	Apply ADATA directly as the compensation value.										
10b	Apply the sum of CALCCOMP and ADATA as the compensation value. If the sum exceeds Fh, then Fh is applied.										
11b	Apply the difference of CALCCOMP minus ADATA as the compensation value. If the difference is less than 0h, then 0h is applied.										
4	Reserved										
3:0	ADATA: Bridge A data value. Read-write. This value is applied to the bridge A PHY compensation as described in ACTL.										

SHPC Capabilities Register**Dev[B, A]:0x90**

This register is reserved if DevA:0x48[HPENB, HPENA] is low.

Default: 0000 980Ch

Attribute: See below.

Bits	Description
31	CIP: Controller Interrupt Pending. Read only. 1=One or more bits in SHPC[B, A]:18 is set. 0=All bits in SHPC[B, A]:18 are cleared.
30	CSERRP: Controller System Error Pending. Read only. 1=One or more bits in SHPC[B, A]:1C is set. 0=All bits in SHPC[B, A]:1C are cleared.
29:24	Reserved.
23:16	SELECT: DWORD Select. Read-write. Specifies the DWORD from the SHPC[B, A]:XX register set that is accessible through Dev[B, A]:0x94. 00h selects SHPC[B, A]:00; 01h selects SHPC[B, A]:04; and so on.
15:8	Next Capability Pointer. Read only. Points to the next capability block.
7:0	Capabilities ID. Read only. Specifies the capabilities ID for SHPC.

SHPC Data Register**Dev[B, A]:0x94**

This register is reserved if DevA:0x48[HPENB, HPENA] is low.

Default: 0000 0000h

Attribute: Read-write.

Bits	Description
31:0	DATA: SHPC data port. Accesses to this port access the register of the SHPC[B, A]:XX register set that is indexed by Dev[B, A]:0x90[SELECT].

Power Management Capabilities Register**Dev[B, A]:0x98**

This register is reserved if DevA:0x48[HPENB, HPENA] is low.

Default: 480A ??01h

Attribute: Read only.

Bits	Description
31:27	PMES: PME support. Indicates PME# support in device state D0 (system state S0) and device state D3 hot (system state S1).
26	D2S: D2 support. Indicates that D2 device power state is not supported.
25	D1S: D1 support. Indicates that D1 device power state is not supported.
24:22	AUXCR: auxiliary current requirements. Indicates that there is no requirement for auxiliary current since the D3 cold device power state is not supported.
21	DSI: Device specific initialization. Indicates that there is no special initialization requirement.
20	Reserved.
19	PMECLK: PME clock. Indicates that the PCI clock is required for PME# generation.
18:16	Version. Specifies that the PCI function complies with Revision 1.1 of the <i>PCI Power Management Interface Specification</i> .
15:8	Next Capability Pointer. Read only. Points to the next capability block. DevA:0x98[15:8]=C0h. DevB:0x98[15:8]=00h.
7:0	Capabilities ID. Specifies the Capabilities ID for PCI Power Management.

Power Management Status and Control Register**Dev[B, A]:0x9C**

This register is reserved if DevA:0x48[HPENB, HPENA] is low.

Default: 0000 0000h

Attribute: See below

Bits	Description
31:24	Reserved.
23	BPCC_EN: bus power/clock control enable. Read only. Indicates that the bus power/clock control policies defined in Section 4.7.1 of the <i>PCI Bus Power Management Interface Specification Rev. 1.1</i> have been disabled.
22:16	Reserved.
15	PME_STS: PME# status. Read; set by hardware; write 1 to clear. Set when [B, A]_PME# is asserted as a result of an SHPC PME event (see SHPC[B, A]:20).
14:9	Reserved.
8	PME_EN: PME enable. Read-write. 1=Enables [B, A]_PME# assertion if Dev[B, A]:0x9C[PME_STS] is set.
7:2	Reserved.
1:0	PWRS: power state. Read-write. Indicates the current power state of the function. 00b = D0. 11b = D3 hot. If software attempts to write unsupported state to this field (01b = D1 or 10b = D2), the write operation completes normally on the bus; however, the data is discarded and no state change occurs.

PCI-X Secondary Status Register

Dev[B, A]:0xA0

Default: 0??3 B807h (bits[21:20] reset to low; see below for bits[24:22]). Attribute: See below.

Bits	Description
31:25	Reserved.
24:22	SCF: secondary clock frequency. Read only. This specifies the frequency of the secondary bus the last time [B, A]_PRESET# was asserted. 0h=conventional PCI mode; 1h=66 MHz PCI-X® mode; 2h=100 MHz PCI-X® mode; 3h=133 MHz PCI-X® mode; 4h-7h are reserved. The default state for this field is determined by strapping options described in section 4.2.
21	SRD: split request delayed. Read only; hardwired low. The IC automatically limits the number of upstream link read requests to the number of downstream buffers available; so there is no reason to limit the number of ADQs in read requests accepted by the IC.
20	SCO: split completion overrun. Read only; hardwired low. The IC automatically limits the number of downstream PCI-X® read requests to the number of upstream response buffers available; so there is no reason to terminate a split completion for this reason.
19	USC: unexpected split completion. Read; set by hardware; write 1 to clear. 1=An unexpected Split Completion with a Requester ID equal to the bridge's secondary bus number, device number 00h, and function number 0 was received on the secondary interface.
18	SCD: split completion discarded. Read; set by hardware; write 1 to clear. 1=The bridge discarded a split completion moving toward the secondary bus because the requester would not accept it.
17	133 MHz capable. Read only. This bit is hardwired high to indicate support for 133 MHz.
16	64-bit device. Read only. This bit is hardwired high to indicate a 64-bit secondary bus.
15:8	Next capability pointer. Read only. Points to the next capability block.
7:0	Capabilities ID. Read only. Specifies the capabilities ID for PCI-X® configuration space.

PCI-X Bridge Status Register

Dev[B, A]:0xA4

Default: 0003 0000h

Attribute: See below.

Bits	Description
31:22	Reserved.
21	SRD: split request delayed. Read only; hardwired low. The IC automatically limits the number of downstream PCI-X® read requests to the number of upstream buffers available; so there is no reason to limit the number of ADQs in read requests accepted by the IC.
20	Split completion overrun. Read only. This bit is hardwired low.
19	Unexpected split completion. Read only. This bit is hardwired low.
18	Split completion discarded. Read only. This bit is hardwired low.
17	133 MHz capable. This bit is set high arbitrarily. It has no meaning since the primary bus is not PCI-X®.
16	64-bit device. Read only. This bit is set high arbitrarily. It has no meaning since the primary bus is not PCI-X®.
15:8	Bus number. Read only. These bits reflect the state of Dev[B, A]:0x18[PRIBUS].

7:3	Device number. Read only. For DevA, these bits reflect the state of DevA:0xC0[BUID]. For DevB, these bits reflect the state of DevA:0xC0[BUID] plus 1.
2:0	Function number. Read only. This is 0h to reflect the value of this function.

PCI-X Upstream Split Transaction Register

Dev[B, A]:0xA8

Default: FFFF 000Eh

Attribute: See below.

Bits	Description
31:16	USTCL: upstream split transaction commitment limit. Read-write. This register controls no hardware. The IC automatically limits the number of upstream link read requests to the number of downstream buffers available; so there is no reason to limit the number of ADQs in read requests accepted by the IC. This field is required to be greater than or equal to Dev[B, A]:0xA8[USTC]. A value of FFFFh specifies that there is no limit. It is expected that this register will be left at its default value by software.
15:0	USTC: upstream split transaction capacity. Read only. This field specifies the number of downstream response ADQs that can be stored for completion on the secondary bus.

PCI-X Downstream Split Transaction Register

Dev[B, A]:0xAC

Default: FFFF 0002h

Attribute: See below.

Bits	Description
31:16	DSTCL: downstream split transaction commitment limit. Read-write. This register controls no hardware. The IC automatically limits the number of downstream PCI-X® read requests to the number of upstream buffers available; so there is no reason to limit the number of ADQs in read requests generated by the IC. This field is required to be greater than or equal to Dev[B, A]:0xAC[DSTC]. A value of FFFFh specifies that there is no limit. It is expected that this register will be left at its default value by software.
15:0	DSTC: downstream split transaction capacity. Read only. This field specifies the number of upstream response ADQs that can be stored for completion to the link.

Interrupt Discovery Configuration Registers

Dev[B, A]:0x[BC, B8]

These two locations duplicate access to the IOAPIC register space defined in section 5.4. Dev[B, A]:0xB8[INDEX] provides the index and Dev[B, A]:0xBC provides the data port. The definition of the indexed registers is as described in section 5.4. Some fields of the IDRDR register are identical to RDR fields (IM, POL, TM, DM, DEST, IRR); these represent duplicate access to the same physical registers (not duplicate registers). Other IDRDR fields (INTRINFO, PASSPW) represent new functionality.

See section 4.5.2 for more information about interrupts.

Dev[B, A]:0xB8. Default: 8000 ??08h

Attribute: See below.

Bits	Description
31:24	Capability type. Read only. This field is hardwired to indicate the link-defined interrupt discovery and configuration block.
23:16	INDEX. Read-write. Specifies the register accessed through the Dev[B, A]:0xBC dataport. This is the same as IOA00 described in section 5.4.
15:8	Next capability pointer. Read only. Points to the next capability block. The value of this register varies as follows: <ul style="list-style-type: none"> • If DevA:0x48[HPENA]=0 then DevA:0xB8[18:5]=C0h (HT capability block). • If DevA:0x48[HPENA]=1 then DevA:0xB8[18:5]=90h (hot plug capability block). • If DevA:0x48[HPENB]=0 then DevB:0xB8[18:5]=00h (last capability block). • If DevA:0x48[HPENB]=1 then DevB:0xB8[18:5]=90h (hot plug capability block).
7:0	Capabilities ID. Read only. Specifies the capabilities ID for link configuration space.

IDRDR. Default: 0000 0000 F800 0001h

Attribute: See below.

Bits	Description
63	IRR. Read; set by hardware; cleared by hardware or write 1 to clear. This provides duplicate access to RDR[IRR] described in section 5.4. However, writing a 1 to this bit clears this register; this is not the case with RDR[IRR].
62	PASSPW. Read-write. The state of this bit is reflected in the PassPW bit of the link interrupt request packet. It is expected to be programmed low in all cases.
61:56	Reserved.
55:24	INTRINFO[55:24]. Read-write. IntrInfo[55:24] in the link interrupt request packet.
23:16	IV. Read-write. IntrInfo[23:16] in the link interrupt request packet. This provides duplicate access to RDR[IV] described in section 5.4.
15:8	DEST. Read-write. IntrInfo[15:8] in the link interrupt request packet. This provides duplicate access to RDR[DEST] described in section 5.4.
7	INTRINFO[7]. Read-write. IntrInfo[7] in the link interrupt request packet.
6	DM. Read-write. IntrInfo[6] in the link interrupt request packet. This provides duplicate access to RDR[DM] described in section 5.4.
5	TM. Read-write. IntrInfo[5] in the link interrupt request packet. This provides duplicate access to RDR[TM] described in section 5.4.
4:2	MT. Read-write. IntrInfo[4:2] in the link interrupt request packet. Accesses to RDR[MT] described in section 5.4. result in translated accesses to this field; see RDR[MT] for details.
1	POL. Read-write. This provides duplicate access to RDR[POL] described in section 5.4.
0	IM. Read-write. This provides duplicate access to RDR[IM] described in section 5.4.

Link Command Register**DevA:0xC0**

Default: 0040 0008h

Attribute: See below.

Bits	Description
31:29	Slave/primary interface type. Read only.
28	DOUI: drop on uninitialized link. Read-write. This specifies the behavior of transactions that are sent to uninitialized links. 0=Transactions that are received by the IC and forwarded to a side of the tunnel, when DevA:0x[C4/C8][INITCPLT and ENDOCH] for that side of the tunnel are both low, remain in buffers awaiting transmission indefinitely (waiting for INITCPLT to be set high). 1=Transactions that are received by the IC and forwarded to a side of the tunnel, when DevA:0x[C4/C8][INITCPLT and ENDOCH] for that side of the tunnel are both low, behave as if ENDOCH were high. Note: this bit is cleared by PWROK reset but not by RESET#.
27	DEFDIR: default direction. Read-write. 0=Send secondary PCI bus master requests to the master link host as specified by DevA:0xC0[MASHST]. 1=Send secondary PCI bus master requests to the opposite side of the tunnel.
26	MASHST: master host. Read; set and cleared by hardware. This bit indicates which link is the path to the master (or only) host bridge on the HyperTransport technology chain. 1=The hardware set this bit as a result of a write command from the B side of the tunnel to any of the bytes of DevA:0xC0[31:16]. 0=The hardware cleared this bit as a result of a write command from the A side of the tunnel to any of the bytes of DevA:0xC0[31:16]. This bit, along with DEFDIR, is used to determine the side of the tunnel to which secondary PCI bus master requests are sent.
25:21	UnitID count. Read only. Specifies the number of UnitIDs used by the IC (two).
20:16	BUID: base UnitID. Read-write. This specifies the link-protocol base UnitID. The IC's logic uses this value to determine the UnitIDs for link request and response packets. When a new value is written to this field, the response includes a UnitID that is based on the new value in this register.
15:8	Reserved.
7:0	Capabilities ID. Read only. Specifies the capabilities ID for link configuration space.

Link Configuration And Control Register**DevA:0xC4 and DevA:0xC8**

DevA:0xC4 applies side A of the tunnel and DevA:0xC8 applies to side B of the tunnel. The default value for bit[5] may vary (see the definition).

Default: ??11 0020h for DevA:0xC4 and ??00 0020h for DevA:0xC8. Attribute: See below.

Bits	Description
31	Reserved.
30:28	LWO: link width out. Read-write. Specifies the operating width of the outgoing link. Legal values are 001b (16 bits; DevA:0xC4 only), 000b (8 bits), 101b (4 bits), 100b (2 bits), and 111b (not connected). Note: this field is cleared by PWROK reset but not by RESET#; the default value of this field depends on the widths of the links of the connecting device, per the link specification. Note: after this field is updated, the link width does not change until either RESET# is asserted or a link disconnect sequence occurs through or LDTSTOP#.
27	Reserved.

26:24	LWI: link width in. Read-write. Specifies the operating width of the incoming link. Legal values are 001b (16 bits; DevA:0xC4 only), 000b (8 bits), 101b (4 bits), 100b (2 bits), and 111b (not connected). Note: this field is cleared by PWROK reset but not by RESET#; the default value of this field depends on the widths of the links of the connecting device, per the link specification. Note: after this field is updated, the link width does not change until either RESET# is asserted or a link disconnect sequence occurs through an LDTSTOP# assertion.
23	Reserved.
22:20	Max link width out. Read only. This specifies the width of the outgoing link to be 16 bits wide for side A and 8 bits wide for side B.
19	Reserved.
18:16	Max link width in. Read only. This specifies the width of the incoming link to be 16 bits wide for side A and 8 bits wide for side B.
15	Reserved.
14	EXTCTL: extended control time during initialization. Read-write. This specifies the time in which LT[B, A]CTL is held asserted during the initialization sequence that follows an LDTSTOP# deassertion, after LR[B, A]CTL is detected asserted. 0=At least 16 bit times. 1=About 50 microseconds. Note: this bit is cleared by PWROK reset but not by RESET#.
13	LDT3SEN: link three-state enable. Read-write. 1=During the LDTSTOP# disconnect sequence, the link transmitter signals are placed into the high impedance state and the receivers are prepared for the high impedance mode. For the receivers, this includes cutting power to the receiver differential amplifiers and ensuring that there are no resultant high-current paths in the circuits. 0=During the LDTSTOP# disconnect sequence, the link transmitter signals are driven, but in an undefined state, and the link receiver signals are assumed to be driven. Note: this bit is cleared by PWROK reset but not by RESET#. AMD recommends that this bit be set high in single-processor systems and be low in multi-processor systems.
12:10	Reserved.
9:8	CRCERR: CRC Error. Read; set by hardware; write 1 to clear. Bit[9] applies to the upper byte of the link (DevA:0xC4 only) and bit[8] applies to the lower byte. 1=The hardware detected a CRC error on the incoming link. Note: this bit is cleared by PWROK reset but not by RESET#.
7	TXOFF: transmitter off. Read; write 1 only. 1=No output signals on the link toggle; the input link receivers are disabled and the pins may float.
6	ENDPOCH: end of chain. Read; write 1 only or set by hardware. 1=The link is not part of the logical HyperTransport technology chain; packets which are issued or forwarded to this link are either dropped or result in an NXA error response, as appropriate; packets received from this link are ignored and CRC is not checked; if the transmitter is still enabled (TXOFF), then it drives only NOP packets with good CRC. ENDOCH may be set by writing a 1 to it or it may be set by hardware if the link is determined to be disconnected at the rising edge of RESET#.
5	INITCPLT: initialization complete. Read only. This bit is set by hardware when low-level link initialization has successfully completed. If there is no device on the other end of the link, or if the device on the other side of the link is unable to properly perform link initialization, then the bit is not set. This bit is cleared when RESET# is asserted or after the link disconnect sequence completes after the assertion of LDTSTOP#.
4	LKFAIL: link failure. Read; set by hardware; write 1 to clear. This bit is set high by the hardware when a CRC error is detected on the link (if enabled by CRCFEN) or if the link is not used in the system. Note: this bit is cleared by PWROK reset, not by RESET#.

3	CRCERRCMD: CRC error command. Read-write. 1=The link transmission logic generates erroneous CRC values. 0=Transmitted CRC values match the values calculated per the link specification. This bit is intended to be used to check the CRC failure detection logic of the device on the other side of the link.
2	Reserved.
1	CRCFEN: CRC flood enable. Read-write. 1=CRC errors (in link A for DevA:0xC4[CRCFEN]; in link B for DevA:0xC8[CRCFEN]) result in sync packets to both outgoing links and the LKFAIL bit is set. 0=CRC errors do not result in sync packets or setting the LKFAIL bit.
0	Reserved.

Link Frequency Capability 0 Register**DevA:0xCC**

Default: 0035 0022h.

Attribute: See below.

Bits	Description
31:16	FREQCAPA: link A frequency capability. Read only. These bits indicate that A side of the tunnel supports 200, 400, 600, and 800 MHz link frequencies.
15:12	Reserved.
11:8	FREQA: link A frequency. Read-write. Specifies the link side A frequency. Legal values are 0h (200 MHz), 2h (400 MHz), 4h (600 MHz), and 5h (800 MHz). Note: this bit is cleared by PWROK reset, not by RESET#. Note: after this field is updated, the link frequency does not change until either RESET# is asserted or a link disconnect sequence occurs through LDTSTOP#.
7:0	REVISION. Read only. The IC is designed to version 1.02 of the link specification.

Link Frequency Capability 1 Register**DevA:0xD0**

Default: 0035 0002h.

Attribute: See below.

Bits	Description
31:16	FREQCAPB: link B frequency capability. Read only. These bits indicate that that B side of the tunnel supports 200, 400, 600, and 800 MHz link frequencies.
15:12	Reserved.
11:8	FREQB: link B frequency. Read-write. Specifies the link side B frequency. Legal values are 0h (200 MHz), and 2h (400 MHz), 4h (600 MHz), and 5h (800 MHz). Note: this bit is cleared by PWROK reset, not by RESET#. Note: after this field is updated, the link frequency does not change until either RESET# is asserted or a link disconnect sequence occurs through LDTSTOP#.
7:0	Link device feature capability indicator. Read only. These bits are set to indicate that the IC supports LDTSTOP#.

Link Enumeration Scratchpad Register**DevA:0xD4**

Default: 0000 0000h.

Attribute: See below.

Bits	Description
31:16	Reserved.
15:0	ESP: enumeration scratchpad. Read-write. This field controls no hardware within the IC. Note: this bit is cleared by PWROK reset, not by RESET#.

Link Non-Prefetchable Memory Space Extension Register**DevA:0xD8**

Default: 0000 0000h.

Attribute: Read-write.

Bits	Description
31:16	Reserved.
15:8	NPUML: non-prefetchable upper memory limit. This field provides bits[39:32] of the non-prefetchable memory space address limit specified by Dev[B, A]:0x20[MEMLIM]. See Dev[B,A]:0x1C for details.
7:0	NPUMB: non-prefetchable upper memory base. This field provides bits[39:32] of the non-prefetchable memory space address base specified by Dev[B, A]:0x20[MEMBASE]. See Dev[B,A]:0x1C for details.

Link PHY Compensation Control Registers**DevA:0x[E8, E4, E0]**

The link PHY circuitry includes automatic compensation that is used to adjust the electrical characteristics for the link transmitters and receivers on both sides of the tunnel. There is one compensation circuit for the receivers and one for each polarity of the transmitters. These registers provide visibility into the calculated output of the compensation circuits, the ability to override the calculated value with software-controlled values, and the ability to offset the calculated values with a fixed difference. The overrides and difference values may be different between sides A and B of the tunnel. These registers specify the compensation parameters as follows:

- DevA:0xE0: transmitter rising edge (P) drive strength compensation.
- DevA:0xE4: transmitter falling edge (N) drive strength compensation.
- DevA:0xE8: receiver impedance compensation.

For DevA:0x[E4, E0], higher values represent higher drive strength; the values range from 01h to 13h (19 steps). For DevA:0xE8, higher values represent lower impedance; the values range from 00h to 1Fh (32 steps).

Note: the default state of these registers is set by PWROK reset; assertion of RESET# does not alter any of the fields.

Default: See below.

Attribute: See below.

Bits	Description
31	Must be low. Read-write. This bit is required to be low at all times; setting it high results in undefined behavior.
30:21	Reserved.
20:16	CALCCOMP: calculated compensation value. Read only. This provides the calculated value from the auto compensation circuitry. The default value of this field is not predictable.

15	Reserved.										
14:13	<p>BCTL: link side B PHY control value. Read-write. These two bits combine to specify the PHY compensation value that is applied to side B of the tunnel as follows:</p> <table border="0"> <thead> <tr> <th><u>BCTL</u></th> <th><u>Description</u></th> </tr> </thead> <tbody> <tr> <td>00b</td> <td>Apply CALCCOMP directly as the compensation value.</td> </tr> <tr> <td>01b</td> <td>Apply BDATA directly as the compensation value.</td> </tr> <tr> <td>10b</td> <td>Apply the sum of CALCCOMP and BDATA as the compensation value. In DevA:0x[E4, E0], if the sum exceeds 13h, then 13h is applied. In DevA:0x[E8], if the sum exceeds 1Fh, then 1Fh is applied.</td> </tr> <tr> <td>11b</td> <td>Apply the difference of CALCCOMP minus BDATA as the compensation value. If the difference is less than 01h, then 01h is applied.</td> </tr> </tbody> </table> <p>The default value of this field (from PWROK reset) is controlled by the CMPOVR signal. If CMPOVR = 0, the default is 00b. If CMPOVR = 1, the default is 01b.</p>	<u>BCTL</u>	<u>Description</u>	00b	Apply CALCCOMP directly as the compensation value.	01b	Apply BDATA directly as the compensation value.	10b	Apply the sum of CALCCOMP and BDATA as the compensation value. In DevA:0x[E4, E0], if the sum exceeds 13h, then 13h is applied. In DevA:0x[E8], if the sum exceeds 1Fh, then 1Fh is applied.	11b	Apply the difference of CALCCOMP minus BDATA as the compensation value. If the difference is less than 01h, then 01h is applied.
<u>BCTL</u>	<u>Description</u>										
00b	Apply CALCCOMP directly as the compensation value.										
01b	Apply BDATA directly as the compensation value.										
10b	Apply the sum of CALCCOMP and BDATA as the compensation value. In DevA:0x[E4, E0], if the sum exceeds 13h, then 13h is applied. In DevA:0x[E8], if the sum exceeds 1Fh, then 1Fh is applied.										
11b	Apply the difference of CALCCOMP minus BDATA as the compensation value. If the difference is less than 01h, then 01h is applied.										
12:8	<p>BDATA: link side B data value. Read-write. This value is applied to the side B of the tunnel PHY compensation as described in BCTL. The default for DevA:0x[E4, E0] is 08h. The default for DevA:0xE8 is 0Fh.</p>										
7	Reserved.										
6:5	<p>ACTL: link side A PHY control value. Read-write. These two bits combine to specify the PHY compensation value that is applied to side A of the tunnel as follows:</p> <table border="0"> <thead> <tr> <th><u>ACTL</u></th> <th><u>Description</u></th> </tr> </thead> <tbody> <tr> <td>00b</td> <td>Apply CALCCOMP directly as the compensation value.</td> </tr> <tr> <td>01b</td> <td>Apply ADATA directly as the compensation value.</td> </tr> <tr> <td>10b</td> <td>Apply the sum of CALCCOMP and ADATA as the compensation value. In DevA:0x[E4, E0], if the sum exceeds 13h, then 13h is applied. In DevA:0x[E8], if the sum exceeds 1Fh, then 1Fh is applied.</td> </tr> <tr> <td>11b</td> <td>Apply the difference of CALCCOMP minus ADATA as the compensation value. If the difference is less than 01h, then 01h is applied.</td> </tr> </tbody> </table> <p>The default value of this field (from PWROK reset) is controlled by the CMPOVR signal. If CMPOVR = 0, the default is 00b. If CMPOVR = 1, the default is 01b.</p>	<u>ACTL</u>	<u>Description</u>	00b	Apply CALCCOMP directly as the compensation value.	01b	Apply ADATA directly as the compensation value.	10b	Apply the sum of CALCCOMP and ADATA as the compensation value. In DevA:0x[E4, E0], if the sum exceeds 13h, then 13h is applied. In DevA:0x[E8], if the sum exceeds 1Fh, then 1Fh is applied.	11b	Apply the difference of CALCCOMP minus ADATA as the compensation value. If the difference is less than 01h, then 01h is applied.
<u>ACTL</u>	<u>Description</u>										
00b	Apply CALCCOMP directly as the compensation value.										
01b	Apply ADATA directly as the compensation value.										
10b	Apply the sum of CALCCOMP and ADATA as the compensation value. In DevA:0x[E4, E0], if the sum exceeds 13h, then 13h is applied. In DevA:0x[E8], if the sum exceeds 1Fh, then 1Fh is applied.										
11b	Apply the difference of CALCCOMP minus ADATA as the compensation value. If the difference is less than 01h, then 01h is applied.										
4:0	<p>ADATA: link side A data value. Read-write. This value is applied to the side A of the tunnel PHY compensation as described in ACTL. The default for DevA:0x[E4, E0] is 08h. The default for DevA:0xE8 is 0Fh.</p>										

Clock Control Register**DevA:0xF0**

See section 4.3.3 for details on clock gating. AMD system recommendations for System Management Action Field (SMAF) codes are: 0=ACPI C2; 1=ACPI C3; 2=FID/VID change; 3=ACPI S1; 4=ACPI S3; 5=Throttling; 6=ACPI S4/S5. AMD recommends setting this register to 0004_0008h (to gate clocks during S1).

Default: 0000 0000h.

Attribute: Read-write.

Bits	Description
31:19	Reserved.
18	CGEN: clock gate enable. 1=Internal clock gating, as specified by bits[7:0] of this register, is enabled.
17	Must be low. This bit is required to be low at all times; setting it high results in undefined behavior.
16	Must be low. This bit is required to be low at all times; setting it high results in undefined behavior.
15:8	Reserved.
7:0	ICGSMAF: internal clock gating system management action fields. Each of the bits of this field correspond to SMAF values that are captured in Stop Grant cycles from the host. For each bit, 1=When LDTSTOP# is asserted prior to a Stop Grant cycle in which the SMAF field matches the ICGSMAF bit that is asserted, then the IC power is reduced through gating of internal clocks. 0=No power reduction while LDTSTOP# is asserted. For example, if clock gating is required for SMAF values of 3 and 5, then ICGSMAF[3, 5] must be high. See section 4.3.3 for details.

5.3 PCI-X IOAPIC Configuration Registers

These registers are located in PCI configuration space, in the first device (device A) and second device (device B), function 1. See section 5.1.2 for a description of the register naming convention.

IOAPIC Vendor And Device ID Register**Dev[B, A]:1x00**

Default: 7451 1022h

Attribute: Read only.

Bits	Description
31:16	IOAPIC device ID.
15:0	Vendor ID.

IOAPIC Status And Command Register**Dev[B, A]:1x04**

Default: 0200 0000h

Attribute: See below.

Bits	Description
31:3	Read only. These bits are fixed in their default state.
2	MASEN: PCI master enable. Read-write. 1=Enables IOAPIC to initiate interrupt requests to the host. Note: if Dev[B, A]:1x44[OSVISBAR]=0, then the state of this bit is ignored. Note: Dev[B, A]:1x44[IOAEN] must be high to enable interrupt requests, regardless of the state of this bit.

1	MEMEN: memory enable. Read-write. 1=Enables access to the memory space specified by DevA:1x10. Note: if Dev[B, A]:1x44[OSVISBAR]=0, then the state of this bit is ignored. Note: Dev[B, A]:1x44[IOAEN] must be high to enable access to the register space, regardless of the state of this bit.
0	IO enable. Read only. This bit is fixed in the low state.

IOAPIC Revision and Class Code Register**Dev[B, A]:1x08**

Default: 0800 1001h

Attribute: Read only.

Bits	Description
31:8	CLASSCODE. Provides the IOAPIC class code.
7:0	REVISION.

IOAPIC Device BIST-Header-Latency-Cache Register**Dev[B, A]:1x0C**

Default: 0000 0000h

Attribute: Read only.

Bits	Description
31:24	BIST. These bits are fixed at their default values.
23:16	HEADER. These bits are fixed at their default values.
15:8	LATENCY. These bits are fixed at their default values.
7:0	CACHE. These bits are fixed at their default values.

IOAPIC Base Address Register**Dev[B, A]:1x10 and Dev[B, A]:1x48**

Offsets 10h and 48h provide access to the same 8-byte register. Offset 48h is always accessible. However, offset 10h can be disabled from read and write access through Dev[B, A]:1x44[OSVISBAR].

Default: 0000 0000 0000 0004h

Attribute: See below.

Bits	Description
63:12	IOABAR: IOAPIC base address register. Read-write. These bits specify the address space of the IOAPIC register set, IOAxx. Note: bits[63:40] are required to be programmed low; setting any of these bits high results in undefined behavior.
11:0	Hardwired. Read only. These bits are all hardwired to their default state to indicate a 4K byte block of 64-bit, non-prefetchable memory space.

IOAPIC Device Subsystem ID and Subsystem Vendor ID Register**Dev[B, A]:1x2C**

Default: 0000 0000h

Attribute: Read; write once.

Bits	Description
31:16	Subsystem ID. This field controls no hardware.
15:0	Subsystem vendor ID. This field controls no hardware.

IOAPIC Control Register**Dev[B, A]:1x44**

Default: 0000 0000h

Attribute: Read-write.

Bits	Description
31:2	Reserved.
1	IOAEN: IOAPIC enable. 1=Access to the IOAPIC registers pointed to by Dev[B, A]:1x10/48 is enabled and the IOAPIC is enabled to generate interrupt requests.
0	OSVISBAR: operating system visible base address register. 0=Dev[B, A]:1x10 is not visible; reads provide all zeros and writes are ignored. Also, the state of Dev[B, A]:1x04[MASEN, MEMEN] are ignored. 1=The IOAPIC BAR is read-write accessible through Dev[B, A]:1x10 and Dev[B, A]:1x04[MASEN, MEMEN] function as specified.

IOAPIC Base Address Register**Dev[B, A]:1x48**

Offsets 10h and 48h provide access to the same 8-byte register. Offset 48h is always accessible. However, offset 10h can be disabled from read and write access through Dev[B, A]:1x44[OSVISBAR]. See offset 10h for the register specification.

5.4 IOAPIC Registers

These registers are located in IOAxx memory space. The base address register for these registers is Dev[B, A]:1x10/48. See section 5.1.2 for a description of the register naming convention. See also section 4.5.2 for more details about interrupt operation. See also Dev[B, A]:0x[BC, B8] for a description of alternative access to these registers and expanded programmability.

The IOAPIC register set supports 4 interrupts and corresponding redirection registers. The space is indexed through two memory-mapped ports: IOA00 (IOAxx at offset 00h) provides the 8-bit index register; IOA10h (IOAxx at offset 10h) provides the 32-bit data port. Writes to IOA10h, the 32-bit data port, must be 32-bit, aligned accesses; other than 32-bit writes result in undefined behavior. Reads provide all four bytes regardless of the byte enables.

The index written to IOA00 selects one of the following:

IOA00[7:0]	Description	Default
00h	APIC ID register. Bits[27:24] are read-write; they control no hardware. All other bits are reserved.	0000 0000h
01h	IOAPIC version register. Read only. These bits are fixed in their default state.	0003 0011h
02h	IOAPIC arbitration ID register. Bits[27:24] are read-write; they control no hardware. All other bits are reserved.	0000 0000h

10h-17h	<p>RDR: redirection registers. Each of the 4 redirection registers utilizes two indexes. Bits[63:32] are accessed through the odd indexes and bits[31:0] are accessed through the even indexes. They are mapped to the PCI interrupt pins as follows:</p> <table border="1"> <thead> <tr> <th><u>Pin</u></th> <th><u>IOA00 for bits[31:0]</u></th> <th><u>IOA00 for bits[63:32]</u></th> </tr> </thead> <tbody> <tr> <td>[B, A]_PIRQA#</td> <td>10h</td> <td>11h</td> </tr> <tr> <td>[B, A]_PIRQB#</td> <td>12h</td> <td>13h</td> </tr> <tr> <td>[B, A]_PIRQC#</td> <td>14h</td> <td>15h</td> </tr> <tr> <td>[B, A]_PIRQD#</td> <td>16h</td> <td>17h</td> </tr> </tbody> </table>	<u>Pin</u>	<u>IOA00 for bits[31:0]</u>	<u>IOA00 for bits[63:32]</u>	[B, A]_PIRQA#	10h	11h	[B, A]_PIRQB#	12h	13h	[B, A]_PIRQC#	14h	15h	[B, A]_PIRQD#	16h	17h	bits[63:32] = 0000 0000h. bits[31:0] = 0001 0000h.
<u>Pin</u>	<u>IOA00 for bits[31:0]</u>	<u>IOA00 for bits[63:32]</u>															
[B, A]_PIRQA#	10h	11h															
[B, A]_PIRQB#	12h	13h															
[B, A]_PIRQC#	14h	15h															
[B, A]_PIRQD#	16h	17h															
18h-FFh	Reserved.																

RDR: the redirection registers are defined as follows:

Bits	Description
63:56	DEST: destination. Read-write. IntrInfo[15:8] in the link interrupt request packet. In physical mode, bits[59:56] specify the APIC ID of the target processor. In logical mode bits[63:56] specify a set of processors.
55:17	Reserved.
16	IM: interrupt mask. Read-write. 1=Interrupt is masked. When the interrupt is specified to be in edge-sensitive mode and this bit transitions from 1 to 0, then no interrupt request is generated regardless of the state of the interrupt line. When the interrupt is specified to be in level-sensitive mode and the interrupt line is in the asserted state, then when this bit transitions from 1 to 0, an interrupt request is generated. The state of this bit is also used for the NIOAIRQ[D:A]# pins; see Dev[B, A]:0x40[NIOAMODE].
15	TM: trigger mode. Read-write. IntrInfo[5] in the link interrupt request packet. 0=Edge sensitive. 1=Level sensitive. Normally, it is expected that this bit be programmed for level-sensitive interrupts. Note: this bit is ignored for delivery modes of SMI, NMI, Init, and ExtINT, which are always treated as edge sensitive.
14	IRR: interrupt request receipt. Read only. This bit is not defined for edge-triggered interrupts. For level-triggered interrupts, this bit is set by the hardware after an interrupt is detected. It is cleared by receipt of EOI as specified in section 4.5.2. .
13	POL: polarity. Read-write. 0=Active high for level-sensitive interrupts and rising edge for edge-sensitive interrupts. 1=Active low for level-sensitive interrupts and falling edge for edge-sensitive interrupts. This bit applies to the polarity of the [B, A]PIRQ[D:A]# pins as they enter the IC. Normally, it is expected that this bit be programmed for active low interrupts. This bit has no effect on the NIOAIRQ[D, C, B, A]# pins.
12	DS: delivery status. Read only. 0=Idle. 1=Interrupt message pending.
11	DM: destination mode. Read-write. IntrInfo[6] in the link interrupt request packet. 0=Physical mode. 1=Logical mode.

10:8	<p>MT: message type. Read-write. These bits are physically located in IDRDR[MT] (See Dev[B, A]:0x[BC, B8]). Accesses to this field result in translated accesses to the register bits in IDRDR[MT]. The value in IDRDR[MT] becomes the IntrInfo[4:2] field in link interrupt request packets. The translation is as follows:</p> <table border="1" style="margin-left: 40px;"> <thead> <tr> <th><u>Access to RDR[MT]</u></th> <th><u>Interrupt type</u></th> <th><u>Value in IDRDR[MT]</u></th> </tr> </thead> <tbody> <tr><td>000b</td><td>Fixed</td><td>000b</td></tr> <tr><td>001b</td><td>Lowest priority</td><td>001b</td></tr> <tr><td>010b</td><td>SMI</td><td>010b</td></tr> <tr><td>011b</td><td>Reserved</td><td>111b</td></tr> <tr><td>100b</td><td>NMI</td><td>011b</td></tr> <tr><td>101b</td><td>Init</td><td>100b</td></tr> <tr><td>110b</td><td>Reserved</td><td>101b</td></tr> <tr><td>111b</td><td>ExtINT</td><td>110b</td></tr> </tbody> </table> <p>So, for example, a write of 111b to RDR[MT] results in a write of 110b in IDRDR[MT]. Subsequent reads of RDR[MT] provide 111b. Subsequent reads of IDRDR[MT] provide 110b. The value placed in link interrupt request packets is as specified in IDRDR[MT] (110b). A write of 110b in IDRDR[MT] would be read as 111b through RDR[MT].</p>	<u>Access to RDR[MT]</u>	<u>Interrupt type</u>	<u>Value in IDRDR[MT]</u>	000b	Fixed	000b	001b	Lowest priority	001b	010b	SMI	010b	011b	Reserved	111b	100b	NMI	011b	101b	Init	100b	110b	Reserved	101b	111b	ExtINT	110b
<u>Access to RDR[MT]</u>	<u>Interrupt type</u>	<u>Value in IDRDR[MT]</u>																										
000b	Fixed	000b																										
001b	Lowest priority	001b																										
010b	SMI	010b																										
011b	Reserved	111b																										
100b	NMI	011b																										
101b	Init	100b																										
110b	Reserved	101b																										
111b	ExtINT	110b																										
7:0	<p>IV: interrupt vector. Read-write. IntrInfo[23:16] in the link interrupt request packet.</p>																											

5.5 SHPC Working Registers

These registers are accessed through either:

- Indexed configuration space (see Dev[B, A]:0x90[SELECT] and Dev[B, A]:0x94[DATA]), or
- Non-indexed memory space (see SHPC[B, A]:00).

See section 5.1.2 for a description of the register naming convention. If DevA:0x48[HPENA] = 0 then the SHPCA:XX registers are all reserved; if DevA:0x48[HPENB] = 0 then the SHPCB:XX registers are all reserved.

SHPC Base Offset Register

SHPC[B, A]:00

Default: 0000 0000h.

Attribute: Read only.

Bits	Description
31:0	BASE_OFFSET. This register is hard-wired low to indicate that the memory-space base address of the SHPC register set is specified only by Dev[B,A]:0x10[SHPCBAR].

SHPC Slots Available Register I**SHPC[B, A]:04**

Default: 0000 0000h

Attribute: Write once.

Bits	Description
31:29	Reserved.
28:24	N_133PCIX. Indicates maximum number of hot plug slots available to be enabled when the bus is running at 133 MHz in PCI-X® mode.
23:21	Reserved.
20:16	N_100PCIX. Indicates maximum number of hot plug slots available to be enabled when the bus is running at 100 MHz in PCI-X® mode.
15:13	Reserved.
12:8	N_66PCIX. Indicates maximum number of hot plug slots available to be enabled when the bus is running at 66 MHz in PCI-X® mode.
7:5	Reserved.
4:0	N_33CONV. Indicates maximum number of hot plug slots available to be enabled when the bus is running at 33 MHz in conventional PCI mode.

SHPC Slots Available Register II**SHPC[B, A]:08**

Default: 0000 0000h

Attribute: Write once.

Bits	Description
31:5	Reserved.
4:0	N_66CONV. Indicates maximum number of hot plug slots available to be enabled when the bus is running at 66 MHz in conventional PCI mode.

SHPC Slot Configuration Register**SHPC[B, A]:0C**

Default: 0000 0000h

Attribute: Write once.

Bits	Description
31	ABI: attention button implemented. 1=Hot plug slots implement the attention button. 0=Hot plug slots do not implement the attention button.
30	MRLSI: MRL sensor implemented. 1=Hot plug slots implement the MRL sensor. 0=Hot plug slots do not implement the MRL sensor.
29	PSN_UP: physical slot number up/down. 1=Each external slot label increments by 1 from the value in SHPC[B, A]:0C[PSN]. 0=Each external slot label decrements by 1 from the value in SHPC[B, A]:0C[PSN].
28:27	Reserved.
26:16	PSN: physical slot number. Specifies the physical slot number of the device specified by SHPC[B, A]:0C[FDN].
15:13	Reserved.
12:8	FDN: first device number. Specifies the device number assigned to the first hot plug slot on the secondary bridge bus.

7:5	Reserved.
4:0	NSI: number of slots implemented. Specifies the number of hot plug slots on the bridge.

SHPC Secondary Bus Configuration Register**SHPC[B, A]:10**

Default: 0100 0000h

Attribute: Read only.

Bits	Description
31:24	SHPC Programming Interface. Identifies the format of the SHPC Working Register set.
23:3	Reserved.
2:0	MODE. Indicates the current speed and mode at which the secondary bridge bus operates. 000b = 33 MHz conventional mode. 001b = 66 MHz conventional mode. 010b = 66 MHz PCI-X® mode. 011b = 100 MHz PCI-X® mode. 100b = 133 MHz PCI-X® mode. 101b, 110b, and 111b are reserved.

SHPC Command Register**SHPC[B, A]:14**

Writes to SHPC[B, A]:14 are ignored if SHPC[B, A]:16[BSY] = 1.

Default: 0000h

Attribute: Read-write.

Bits	Description
15:13	Reserved.
12:8	TGT: target slot. Specifies the slot to which SHPC[B, A]:14[CMD] is applied for the Slot Operation command.
7:0	CMD: SHPC command code. Specifies the SHPC command to be executed (see below).

Command Name	CMD[7:0]							
	0	0	Attention Indicator		Power Indicator		Slot State	
Slot Operation	0	0						
Set Bus Segment Speed/Mode	0	1	0	0	0	Bus Speed/Mode		
Power Only All Slots	0	1	0	0	1	0	0	0
Enable All Slots	0	1	0	0	1	0	0	1

Decodings for SHPC command code fields are:

- Attention Indicator and Power Indicator specify LED states. 00b=No change; 01b=On; 10b=Blink; 11b=Off.
- Slot State specifies the command to the slot. 00b=No Change; 01b=Power only; 10b=Enable slot; 11b=Disable slot.
- Bus Speed/Mode specifies the bridge speed and mode. See SHPC[B, A]:10[MODE] for the encoding.

SHPC Status Register**SHPC[B, A]:16**

The Controller Command Error Code field consists of SHPC[B, A]:16[INVS_M_ERR, INVCMD_ERR, MRLO_ERR]. No bits or one bit of the Controller Command Error Code field may be updated when SHPC[B, A]:16[BSY] transitions from 1 to 0, indicating command completion with an error. If a bit in the Controller Command Error Code field is set, then it remains set until the next 1 to 0 transition of BSY.

Default: 0000h

Attribute: Read only.

Bits	Description
15:4	Reserved.
3	<p>INVS_M_ERR: Invalid Speed/Mode. This is set high when one of the following errors occurs:</p> <ul style="list-style-type: none"> • The target slot specified by SHPC[B, A]:14[TGT] is not capable of running at the current speed or mode when the Slot Operation Command Enable command is issued. • A slot on the bus is not capable of running at the current bus speed or mode when the Enable All Slots Command is issued. • An enabled slot on the bus segment is not capable of running at the requested bus speed or mode when the Set Bus Segment Speed/Mode Command is issued. • The Set Bus Segment Speed/Mode Command is issued when the number of slots available at the requested bus speed or mode (specified by SHPC[B, A]:[08, 04]) is greater than zero and less than the number of slots enabled.
2	<p>INVCMD_ERR: invalid SHPC command. This is set high when one of the following errors occurs:</p> <ul style="list-style-type: none"> • A reserved command code is used. • The target slot specified by SHPC[B, A]:14[TGT] is zero or is greater than the SHPC[B, A]:0C[NSI] for any Slot Operation Command. • The target slot specified by SHPC[B, A]:14[TGT] is greater than the number of slots available at the current bus speed or mode (specified by SHPC[B, A]:[08, 04]) when Slot Operation Command Enable is issued. • The target slot specified by SHPC[B, A]:14[TGT] is enabled when Slot Operation Command Power Only is issued. • One or more slots on the bus segment are already enabled when Power Only All Slots Command or Enable All Slots Command is issued. • The Set Bus Segment Speed/Mode Command is issued when SHPC[B, A]:[08, 04] indicate no slots are available at the requested speed or mode.
1	<p>MRLO_ERR: MRL open. 1=The MRL of the target slot specified by SHPC[B, A]:14[TGT] was open when Slot Operation Command Power Only or Slot Operation Command Enable was issued.</p>
0	<p>BSY: Controller Busy. 1=An SHPC command (see SHPC[B, A]:14) is in progress.</p>

SHPC Interrupt Locator Register**SHPC[B, A]:18**

Default: 0000 0000h

Attribute: Read only.

Bits	Description
31:5	Reserved.
4:1	IP[4:1]: slot interrupt pending. Each bit <i>n</i> of this field corresponds to slot <i>n</i> . 1=A slot status bit capable of generating interrupts is set and the corresponding interrupt mask is 0. Slot status bits capable of generating interrupts are SHPC[B, A]:[30, 2C, 28, 24][CPC_STS, IPF_STS, ABP_STS, MRLSC_STS, CPF_STS]. The corresponding interrupt masks are SHPC[B, A]:[30, 2C, 28, 24][CP_IM, IPF_IM, AB_IM, MRLS_IM, CPF_IM].
0	CC_IP: command complete interrupt pending. 1=SHPC[B, A]:20[CC_STS] is 1 and SHPC[B, A]:20[CC_IM] is 0.

SHPC SERR Locator Register**SHPC[B, A]:1C**

Default: 0000 0000h

Attribute: Read only.

Bits	Description
31:5	Reserved.
4:1	SERRP[4:1]: slot SERR pending. Each bit <i>n</i> of this field corresponds to slot <i>n</i> . 1=A slot status bit capable of generating SERR is set and the corresponding SERR mask is 0. Slot status bits capable of generating SERR are SHPC[B, A]:[30, 2C, 28, 24][MRLSC_STS, CPF_STS]. The corresponding SERR masks are SHPC[B, A]:[30, 2C, 28, 24][MRLS_SERRM, CPF_SERRM].
0	A_SERRP: arbiter SERR pending. 1=SHPC[B, A]:20[ATOUT_STS] is 1 and SHPC[B, A]:20[A_SERRM] is 0.

SHPC SERR-INT Register**SHPC[B, A]:20**

The wakeup signal shown below sets Dev[B, A]:0x9C[PME_STS].

$$\text{SHPC_WAKEUP} = (\text{SHPC}[\text{B}, \text{A}]:18[\text{IP}] \neq 0000\text{b}) \\ | \sim\text{SHPC}[\text{B}, \text{A}]:20[\text{CC_IM}] \& \text{SHPC}[\text{B}, \text{A}]:20[\text{CC_STS}];$$

The SHPC interrupt shown below is routed to the [B, A]_PIRQA# pin.

$$\text{SHPC_INTR} = \sim\text{SHPC}[\text{B}, \text{A}]:20[\text{GIM}] \& \text{SHPC_WAKEUP};$$

The SHPC system error shown below sets Dev[B, A]:0x1C[RSE] (see also Dev[B, A]:0x3C[SERREN]).

$$\text{SHPC_SERR} = \sim\text{SHPC}[\text{B}, \text{A}]:20[\text{GSERRM}] \& \\ ((\text{SHPC}[\text{B}, \text{A}]:1\text{C}[\text{SERRP}] \neq 0000\text{b}) \\ | \sim\text{SHPC}[\text{B}, \text{A}]:20[\text{A_SERRM}] \& \text{SHPC}[\text{B}, \text{A}]:20[\text{ATOUT_STS}]);$$

Default: 0000 000Fh

Attribute: See below.

Bits	Description
31:18	Reserved.
17	ATOUT_STS: arbiter timeout status. Read; set by hardware; write 1 to clear. Set when an arbiter timeout is detected by the SHPC logic. The arbiter timeout occurs when the PCI bus is requested (from the internal arbiter) for a hot plug operation and it is not granted for 2^{23} [B, A]_PCLK cycles.
16	CC_STS: command completion status. Read; set by hardware; write 1 to clear. Set when an SHPC[B, A]:16[BSY] transition from 1 to 0 is detected.
15:4	Reserved.
3	A_SERRM: arbiter SERR mask. Read-write. 1=SERR indication for arbiter timeout is disabled.
2	CC_IM: command complete interrupt mask. Read-write. 1=SHPC interrupt generation for command completion is disabled.
1	GSERRM: global SERR mask. Read-write. 1=SERR indication is disabled.
0	GIM: global interrupt mask. Read-write. 1=SHPC interrupt generation is disabled.

SHPC Logical Slot Registers**SHPC[B, A]:[30, 2C, 28, 24]**

The offset for the SHPC Logical Slot Register (or LSR) for slot 1 is 24h, for slot 2 is 28h, for slot 3 is 2Ch, and for slot 4 is 30h. "LSR" is used instead of SHPC[B, A]:[30, 2C, 28, 24] in the description below. See SHPC[B, A]:20 for information about how these registers may affect interrupts, events, and system errors.

Default: 7F00 3F3Fh

Attribute: See below.

Bits	Description
31	Reserved.
30	CPF_SERRM: connected power fault SERR mask. Read-write. 1=SERR generation is disabled when LSR[CPF_STS] is set. 0=SERR generation is enabled when LSR[CPF_STS] is set.
29	MRLS_SERRM: MRL sensor SERR mask. Read-write. 1=SERR generation is disabled when LSR[MRLSC_STS] is set. 0=SERR generation is enabled when LSR[MRLSC_STS] is set.
28	CPF_IM: connected power fault interrupt mask. Read-write. 1=Interrupt generation is disabled when LSR[CPF_STS] is set. 0=Interrupt generation is enabled when LSR[CPF_STS] is set.
27	MRLS_IM: MRL sensor interrupt mask. Read-write. 1=Interrupt generation is disabled when LSR[MRLSC_STS] is set. 0=Interrupt generation is enabled when LSR[MRLSC_STS] is set.
26	AB_IM: attention button interrupt mask. Read-write. 1=Interrupt generation is disabled when LSR[ABP_STS] is set. 0=Interrupt generation is enabled when LSR[ABP_STS] is set.
25	IPF_IM: isolate power fault interrupt mask. Read-write. Read-write. 1=Interrupt generation is disabled when LSR[IPF_STS] is set. 0=Interrupt generation is enabled when LSR[IPF_STS] is set.
24	CP_IM: card presence interrupt mask. Read-write. 1=Interrupt generation is disabled when LSR[CPC_STS] is set. 0=Interrupt generation is enabled when LSR[CPC_STS] is set.
23:21	Reserved.
20	CPF_STS: connected power fault status. Read; set by hardware; write 1 to clear. Set when LSR[PF] changes from 0 to 1 while LSR[SS] = 10b (slot is enabled).

19	MRLSC_STS: MRL sensor change status. Read; set by hardware; write 1 to clear. Set when LSR[MRLS] changes its value.
18	ABP_STS: attention button press status. Read; set by hardware; write 1 to clear. Set when LSR[AB] transitions from 0 to 1.
17	IPF_STS: isolated power fault status. Read; set by hardware; write 1 to clear. Set when LSR[PF] changes from 0 to 1 while LSR[SS] != 10b (slot is not in the enabled state). .
16	CPC_STS: card presence change status. Read; set by hardware; write 1 to clear. Set when LSR[PRSNT1_2] field changes value.
15:14	Reserved.
13:12	PCI-X_CAP: PCI-X® capability. Read-only. Reflects the current PCI-X® capability of the add-in-card. These bits are not valid if the slot is empty. 00b = Conventional PCI. 01b = 66 MHz PCI-X® mode. 10b = Reserved. 11b = 133 MHz PCI-X® mode.
11:10	PRSNT1_2: PRSNT1#/PRSNT2#. Read-only. Reflects the current debounced state of the PRSNT1# and PRSNT2# pins on the slot. 00b = Card present; 7.5W. 01b = Card present; 15W. 10b = Card present; 25W. 11b = Slot Empty.
9	M66_CAP: 66 MHz capable. Read-only. This bit is valid only when the slot is occupied and powered. 1=Add-in card is capable of running at 66 MHz conventional mode. 0=Add-in-card is capable of running at 33 MHz conventional mode only.
8	MRLS: MRL sensor. Read-only. Reflects the current state of the debounced MRL sensor. 1=MRL sensor is open. 0=MRL sensor is closed.
7	AB: attention button. Read-only. Reflects the current state of the debounced attention button. 1=Attention button is being pressed. 0=Attention button is released.
6	PF: power fault. Read-only. Reflects the current state of the power fault latch in the slot power control circuitry. 1=Power fault (isolated or connected) is detected.
5:4	AIS: attention indicator state. Read-only. Reflects the current state of the attention indicator. 00b = reserved. 01b = on. 10b = blink. 11b = off.
3:2	PIS: power indicator state. Read-only. Reflects the current state of the power indicator. 00b = reserved. 01b = on. 10b = blink. 11b = off.
1:0	SS: slot state. Read-only. Reflects the current state of the slot. 00b = reserved. 01b = powered only. 10b = enabled. 11b = disabled.

6 Electrical Data

6.1 Absolute Ratings

The IC is not designed to operate beyond the parameters shown in the following table.

Parameter	Minimum	Maximum	Comments
VDD12[B, A]	-0.5 V	1.7 V	
VDD18, VDDA18	-0.5 V	2.0 V	
VDD33	-0.5 V	3.6 V	
T _{CASE} (Under Bias)		85 °C	
T _{STORAGE}	-65 °C	150 °C	

Table 11: Absolute maximum ratings.

6.2 Operating Ranges

The IC is designed to provide functional operation if the voltage and temperature parameters are within the limits defined in the following table.

Parameter	Minimum	Typical	Maximum	Units	Comments
VDD12[B, A]	1.14	1.2	1.26	V	
VDD18, VDDA18	1.71	1.8	1.89	V	
VDDA18 peak-to-peak noise			50	mV	Maximum sinusoidal amplitude at frequency range from 50 KHz to 20 MHz.
VDD33	3.135	3.3	3.465	V	
T _{CASE} (Under Bias)			85	deg C	

Table 12: Operating ranges.

6.3 DC Characteristics

See the HyperTransport™ Technology Electrical Specification for the DC characteristics of link signals.

The following table shows current consumption in amps and power in watts for each power plane. Unless otherwise noted, values assume that both bridges are operating at 133 MHz.

Parameter Description	Typical		Max	
	Current	Power	Current	Power
VDD12A current, power	0.13 A	0.16 W	0.19 A	0.24 W
VDD12B current, power	0.07 A	0.08 W	0.09 A	0.12 W
VDD18 current, power; operational	1.90 A	3.42 W	2.40 A	4.54 W
VDD18 current, power; internal clock gated mode (see section 4.3.3); both bridges operating at 66 MHz.	0.30 A	0.54 W	0.43 A	0.82 W
VDD18 current, power; internal clock gated mode (see section 4.3.3); both bridges operating at 133 MHz.	0.52 A	0.94 W	0.75 A	1.42 W
VDDA18 current, power	0.02 A	0.04 W	0.03 A	0.06 W
VDD33 current, power; assumes no current load on PCI bus signals	0.50 A	1.65 W	0.60 A	2.08 W
Total power; operational (no clock gating)		5.4 W		7.0 W

Table 13: Current and power consumption.

The following table shows DC characteristics for signals on the VDD33 power plane.

Symbol	Parameter Description	Min	Max	Units	Comments
V _{IL}	Input low voltage	-0.5	0.35 VDD33	V	
V _{IH}	Input high voltage	0.5 VDD33	0.5 + VDD33	V	
V _{OL}	Output low voltage; I _{OUT} = 1.5 mA		0.1 VDD33	V	
V _{OH}	Output high voltage; I _{OUT} = -0.5 mA	0.9 VDD33		V	
I _{LI}	Input leakage current		+/- 10	uA	
C _{IN}	Input capacitance		8	pF	
V _{XCL}	[B, A]_PCIXCAP voltage for low state		0.15 VDD33	V	
V _{XCM}	[B, A]_PCIXCAP voltage for mid state	0.25 VDD33	0.70 VDD33	V	
V _{XCH}	[B, A]_PCIXCAP voltage for high state	0.85 VDD33		V	

Table 14: DC characteristics for signals on the VDD33 power plane.

6.4 AC Characteristics

See the HyperTransport Technology Electrical Specification for the AC characteristics of link signals, PWROK, RESET#, and LDTSTOP#.

The following table shows AC requirements for REFCLK.

Symbol	Parameter Description	Min	Max	Units	Comments
t _{REF}	REFCLK cycle time	15	18	ns	
t _{REFCJ}	REFCLK cycle to cycle jitter; difference in the period of any two adjacent REFCLK cycles		0.250	ns	
t _{REFPJ}	REFCLK period jitter; difference between the nominal period and the period of any REFCLK cycle	-0.300	+0.300	ns	
t _{REFSW}	REFCLK slew rate	1	4	V/ns	
t _{REFSS}	REFCLK 33 KHz spread spectrum frequency change from t _{REF}	-0.5	0	%	

Table 15: AC requirements for REFCLK.

The following table shows AC specification data for PCI clocks.

Symbol	Parameter Description	133 MHz		100 MHz		66 MHz		33 MHz		Units	Comments
		Min	Max	Min	Max	Min	Max	Min	Max		
t _{CYC}	[B, A]_PCLK[4:0] cycle time	7.5		10		15		30		ns	
t _{HIGH}	[B, A]_PCLK[4:0] high time	3		4		6		11		ns	
t _{LOW}	[B, A]_PCLK[4:0] low time	3		4		6		11		ns	
t _{SLEW}	[B, A]_PCLK[4:0] slew rate	1.5	4	1.5	4	1.5	4	1.5	4	V/ns	

Table 16: AC data for PCI clocks.

The following table shows general AC specification data. "PCLK" refers to [B, A]_PCLK[4:0].

Symbol	Parameter Description	PCI-X® 133, 100, 66 MHz		Conven- tional PCI 66 MHz		Conven- tional PCI 33 MHz		Units	Comments
		Min	Max	Min	Max	Min	Max		
T _{val}	PCLK to signal valid delay	0.7	3.8	2	6	2	11	ns	
T _{on}	PCLK to signal active delay	0		2		2		ns	
T _{off}	PCLK to signal float delay		7		14		28	ns	
T _{su}	Input setup time to PCLK	1.2		3		7		ns	
T _h	Input hold time from PCLK	0.5		0		0		ns	

Table 17: AC data for PCI bus.

7 Ball Designations

A			LDT COMP1	VSS	VSS	LTACA D_P0	LTACA D_N0	LTACA D_P2	LTACA D_N2	LTACL K0_P	LTACL K0_N	LTACA D_P5	LTACA D_N5	LTACA D_P7	LTACA D_N7	LRACT L_N	LRACT L_P	LRACA D_N6	LRACA D_P6	LRACA D_N4	LRACA D_P4	LRACA D_N3	LRACA D_P3	LRACA D_N1	LRACA D_P1	VSS	CMP OVR							A	
B		LDT COMP2	VSS	LDT COMP0	VSS	VSS	LTACA D_P1	VDD18	LTACA D_P3	VSS	LTACA D_P4	VDD18	LTACA D_P6	VSS	LTACT L_P0	VDD18	LRACA D_N7	VSS	LRACA D_N5	VDD18	LRACL K0_N	VSS	LRACA D_N2	VDD18	LRACA D_P0	VSS	STRAP L2	RESET #							B
C	VSS	VSS	LDT COMP3	VSS	VSS	LTACA D_N8	LTACA D_N10	LTACA D_N3	LTACL K1_N	LTACA D_N4	LTACA D_N3	LTACA D_P13	LTACA D_N15	LTACT L_N0	RSVD3	LRACA D_P7	LRACA D_P14	LRACA D_P5	LRACA D_P12	LRACL K0_P	LRACA D_P11	LRACA D_P2	LRACA D_P9	LRACA D_P0	VSS	LDT STOP#	PWR OK	REF CLK					C		
D	FREE 12	FREE 10	VSS	FREE1	VSS	LTACA D_P8	VDD18	LTACA D_P10	VSS	LTACL K1_P	VDD18	LTACA D_P13	VSS	LTACA D_P15	VDD18	RSVD2	VSS	LRACA D_N14	VDD18	LRACA D_N12	VSS	LRACA D_N11	VDD18	LRACA D_N9	VSS	VSS	TEST	VSS	VSS					D	
E	FREE 18	FREE 19	FREE 13	FREE 14	VSS	LTACA D_P9	LTACA D_N9	LTACA D_P11	LTACA D_N11	LTACA D_P12	LTACA D_N12	LTACA D_P14	LTACA D_N14	RSVD0	RSVD1	LRACA D_N15	LRACA D_P15	LRACA D_N13	LRACA D_P13	LRACL K1_N	LRACL K1_P	LRACA D_N10	LRACA D_P10	LRACA D_N8	LRACA D_P8	VSS	VSS	VSS	LTBCA D_P0	LTBCA D_P0	LTBCA D_P0			E	
F	B_AD 35	B_AD 34	B_AD 33	B_AD 32	FREE5	VSS	VSS	VDD18	VSS	VDD18	VSS	VDD18	VSS	VDD18	VSS	VDD18	VSS	VDD18	VSS	VDD18	VSS	VDD18	VSS	VDD18	NIOA IROA#	VSS	LTBCA D_N1	LTBCA D_P1	LTBCA D_P0					F	
G	B_AD 38	VSS	B_AD 37	VDD33	B_AD 36	VSS	VDD18	VSS	VDD18	VSS	VDD18	VSS	VDD18	VSS	VDD18	VSS	VDD18	VSS	VDD18	VSS	VDD18	VSS	VDD18	NIOA IROB#	NIOA IROD#	NIOA IROE#	VSS	VDD18	LTBCA D_P2	LTBCA D_P2	LTBCA D_P2			G	
H	B_AD 44	B_AD 43	B_AD 42	B_AD 41	B_AD 40	B_AD 39	VSS	VDD18	VSS	VDD12 A	VSS	VDD18	VSS	VDD18	VSS	VDD18	VSS	VDD12 A	VSS	VDD18	VSS	VDD18	VSS	VDD18	VDD33	STRAP L3	VDD18	LTBCA D_N3	LTBCA D_P3	LTBCA D_N2				H	
J	B_AD 50	B_AD 49	B_AD 48	B_AD 47	B_AD 46	B_AD 45	VDD33	VSS	VDD33	VSS	VDD12 A	VSS	VDD18	VSS	VDD18	VSS	VDD12 A	VSS	VDD18	VSS	VDD18	VSS	VDD18	VSS	HPSOD	HPSIL#	RSVD4	VSS	LTBCL K0_P					J	
K	B_AD 53	VSS	B_AD 52	VDD33	B_AD 51	VDD33	VSS	VDD33	VSS	VDD12 A	VSS	VDD18	VSS	VDD18	VSS	VDD18	VSS	VDD12 A	VSS	VDD18	VSS	VDD18	VSS	VDD18	A_COM PAT	RSVD5	VSS	LTBCA D_N4	LTBCA D_P4	LTBCL K0_N				K	
L	B_AD 59	B_AD 58	B_AD 57	B_AD 56	B_AD 55	B_AD 54	VDD33	VSS	VDD33	VSS	VDD12 A	VSS	VDD18	VSS	VDD18	VSS	VDD12 A	VSS	VDD12 B	VSS	VDD12 B	VSS	VDD18	RSVD6	RSVD7	A_REO 4#	RSVD8	VDD18	LTBCA D_P5	LTBCA D_P5				L	
M	B_CBE _L4	B_PAR 64	B_AD 63	B_AD 62	B_AD 61	B_AD 60	VSS	VDD33	VSS	VDD33	VSS	VDD18	VSS	VDD18	VSS	VDD18	VSS	VDD12 B	VSS	VDD12 B	VSS	VDD18	VSS	A_PLL CLK0	A_PLL CLK1	VDD18	LTBCA D_N6	LTBCA D_P6	LTBCA D_N5					M	
N	B_CBE _L7	VSS	B_CBE _L6	VDD33	B_CBE _L5	VSS	VDD33	VSS	VDD33	VSS	VDD33	VSS	VDD18	VSS	VDD18	VSS	VDD18	VSS	VDD18	VSS	VDD18	VSS	VDD18	A_GNT 4#	A_PCLK4	RSVD9	RSVD 10	VSS	LTBCA D_P7	LTBCA D_P7				N	
P	B_AD3	B_AD2	B_AD1	B_AD0	B_ACK 64#	B_REO 64#	VSS	VDD33	VSS	VDD33	VSS	VDD18	VSS	VDD18	VSS	VDD18	VSS	VDD18	VSS	VDD18	VSS	VDD18	VSS	VDD18	A_PCLK3	A_GNT 3#	VSS	LTBCT L_N	LTBCT L_P	LTBCT D_N7				P	
R	B_AD4	B_AD5	B_AD6	B_AD7	B_CBE _L0	B_AD8	VDD33	VSS	VDD33	VSS	VDD33	VSS	VDD18	VSS	VDD18	VSS	VDD18	VSS	VDD18	VSS	VDD18	VSS	VDD18	A_PCLK2	A_GNT 2#	A_REO 3#	RSVD 11	VDD18	LRBCT L_N					R	
T	B_AD9	VSS	B_M66 EN	VDD33	B_AD 10	VDD33	VSS	VDD33	VSS	VDD33	VSS	VDD18	VSS	VDD18	VSS	VDD18	VSS	VDD18	VSS	VDD18	VSS	VDD18	VSS	VDD18	A_PCLK1	A_REO 2#	VDD18	LRBCA D_P7	LRBCA D_N7	LRBCT L_P				T	
U	B_AD1	B_AD 12	B_AD 13	B_AD 14	B_AD 15	B_CBE _L1	VDD33	VSS	VDD33	VSS	VDD33	VSS	VDD18	VSS	VDD18	VSS	VDD18	VSS	VDD18	VSS	VDD18	VSS	VDD18	VSS	VDD18	A_PIRO A#	A_GNT 1#	RSVD 12	RSVD 13	VSS	LRBCA D_N6	LRBCA D_P6			U
V	B_PAR SERR#	B_PERR#	B_STOP#	B_PCIX CAP	B_DEV SEL#	VSS	VDD33	VSS	VDD33	VSS	VDD18	VSS	VDD18	VSS	VDD18	VSS	VDD12 B	VSS	VDD12 B	VSS	VDD18	VSS	VDD18	A_REO 1#	A_PIRO B#	VSS	LRBCA D_N5	LRBCA D_P5	LRBCA D_P6					V	
W	B_TRDY#	VSS	B_IRDY #	VDD33	B_FRA ME#	VSS	VDD33	VSS	VDD33	VSS	VDD33	VSS	VDD33	VSS	VDD33	VSS	VDD12 B	VSS	VDD12 B	VSS	VDD18	VSS	VDD18	RSVD 14	RSVD 15	RSVD 16	RSVD 17	VDD18	LRBCA D_N4	LRBCA D_P4				W	
Y	B_CBE _L2	B_AD 16	B_AD 17	B_AD 18	B_AD 19	B_AD 20	VSS	VDD33	VSS	VDD33	VSS	VDD33	VSS	VDD33	VSS	VDD33	VSS	VDD33	VSS	VDD33	VSS	VDD33	VSS	VDD18	A_PIRO C#	RSVD 18	VDD18	LRBCL K0_P	LRBCL K0_N	LRBCA D_P4				Y	
AA	B_AD 21	B_AD 22	B_AD 23	B_CBE _L3	B_AD 24	B_AD 25	VDD33	VSS	VDD33	VSS	VDD33	VSS	VDD33	VSS	VDD33	VSS	VDD33	VSS	VDD33	VSS	VDD33	VSS	VDD33	VSS	VDD18	A_PIRO C#	RSVD 19	RSVD 20	RSVD 21	VSS	LRBCA D_N3	LRBCA D_P3			AA
AB	B_AD 26	VSS	B_AD 27	VDD33	B_AD 28	VDD33	VSS	VDD33	VSS	VDD33	VSS	VDD33	VSS	VDD33	VSS	VDD33	VSS	VDD33	VSS	VDD33	VSS	VDD33	VSS	VDD33	A_PCLK0	A_PIRO D#	VSS	LRBCA D_P2	LRBCA D_N2	LRBCA D_P3					AB
AC	B_AD 29	B_AD 30	B_AD 31	B_REO 0#	B_GNT 0#	B_PCLK0	VDD33	VSS	VDD33	VSS	VDD33	VSS	VDD33	VSS	VDD33	VSS	VDD33	VSS	VDD33	VSS	VDD33	VSS	VDD33	VSS	VDD33	A_REO 0#	A_GNT 0#	A_PRE SET#	RSVD 22	VDD18	LRBCA D_N0	LRBCA D_P1			AC
AD	B_PRE SET#	B_PIRO D#	B_PIRO C#	B_PIRO B#	FREE6 #	B_PME #	A_PME #	VSS	FREE 15	FREE8	A_AD 33	VDD33	A_AD 42	A_AD 48	VSS	A_AD6 2	A_CBE _L7	VDD33	A_AD6	A_AD 10	VSS	A_PERR#	A_FRA ME#	VDD33	A_AD 23	VDD18	LRBCA D_N0	LRBCA D_P1	LRBCA D_P1					AD	
AE	B_PIRO A#	VSS	P_CAL	VDD33	FREE 17	B_PLL CLK1	B_PLL CLK0	B_REO 4#	FREE 22	VSS	A_AD 34	A_AD 39	A_AD 43	A_AD 49	A_AD 54	A_AD 61	A_CBE _L6	A_AD0	A_AD5	A_M66 EN	A_AD 13	A_SERR#	A_IRDY #	A_AD 17	A_AD 22	VSS	VSS	VSS	VSS					AE	
AF	B_REO 1#	P_CAL #	B_GNT 1#	B_PCLK1	VDD33	B_REO 2#	B_REO 3#	B_PCLK4	VSS	A_AD 35	VDD33	A_AD 44	A_AD 50	VDD33	A_AD 60	A_CBE _L5	VDD33	A_AD4	A_AD9	VDD33	A_PAR	VDD33	A_TRDY#	VDD33	A_AD 21	A_AD 26	VDD33	A_AD 31	VDDA 18					AF	
AG	FREE 16	FREE 20	VDD33	NC0	NC1	B_GNT 2#	B_GNT 3#	B_GNT 4#	FREE3	VSS	A_AD 36	A_AD 40	A_AD 45	A_AD 51	A_AD 55	A_AD 59	A_CBE _L4	A_ACK 64#	A_AD3	A_AD8	A_AD 12	A_CBE _L1	A_DEV SEL#	A_AD 16	A_AD 20	A_AD 25	A_AD 28	A_AD 30	VDDA 18					AG	
AH		VSS	NC2	NC3	VSS	B_PCLK2	B_PCLK3	VSS	VSS	VSS	A_AD 37	A_AD 46	A_AD 52	A_AD 58	A_AD 64	A_PAR	VSS	A_AD2	A_CBE _L0	VSS	A_AD 15	A_PCIX CAP	VSS	A_AD 19	A_AD 24	VSS	A_AD 29							AH	
AJ			FREE9	FREE 21	FREE2	HPSIC	HPSOC	FREE4	FREE7	A_AD 32	A_AD 38	A_AD 41	A_AD 47	A_AD 53	A_AD 56	A_AD 57	A_AD 63	A_REO 64#	A_AD1	A_AD7	A_AD1 1	A_AD 14	A_STOP#	A_CBE _L2	A_CBE _L3	A_AD 27								AJ	
1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20	21	22	23	24	25	26	27	28	29							

Top side view

Figure 13: Ball designations.

Alphabetical listing of signals and corresponding BGA designators.

Signal name	Ball	Signal name	Ball	Signal name	Ball	Signal name	Ball
A_ACK64#	AG18	A_AD51	AG14	A_REQ4#	L26	B_AD46	J5
A_AD0	AE18	A_AD52	AH14	A_REQ64#	AJ18	B_AD47	J4
A_AD1	AJ19	A_AD53	AJ14	A_SERR#	AE22	B_AD48	J3
A_AD2	AH19	A_AD54	AE15	A_STOP#	AJ23	B_AD49	J2
A_AD3	AG19	A_AD55	AG15	A_TRDY#	AF23	B_AD50	J1
A_AD4	AF19	A_AD56	AJ15	B_ACK64#	P5	B_AD51	K5
A_AD5	AE19	A_AD57	AJ16	B_AD0	P4	B_AD52	K3
A_AD6	AD19	A_AD58	AH16	B_AD1	P3	B_AD53	K1
A_AD7	AJ20	A_AD59	AG16	B_AD2	P2	B_AD54	L6
A_AD8	AG20	A_AD60	AF16	B_AD3	P1	B_AD55	L5
A_AD9	AF20	A_AD61	AE16	B_AD4	R1	B_AD56	L4
A_AD10	AD20	A_AD62	AD16	B_AD5	R2	B_AD57	L3
A_AD11	AJ21	A_AD63	AJ17	B_AD6	R3	B_AD58	L2
A_AD12	AG21	A_CBE_L0	AH20	B_AD7	R4	B_AD59	L1
A_AD13	AE21	A_CBE_L1	AG22	B_AD8	R6	B_AD60	M6
A_AD14	AJ22	A_CBE_L2	AJ24	B_AD9	T1	B_AD61	M5
A_AD15	AH22	A_CBE_L3	AJ26	B_AD10	T5	B_AD62	M4
A_AD16	AG24	A_CBE_L4	AG17	B_AD11	U1	B_AD63	M3
A_AD17	AE24	A_CBE_L5	AF17	B_AD12	U2	B_CBE_L0	R5
A_AD18	AJ25	A_CBE_L6	AE17	B_AD13	U3	B_CBE_L1	U6
A_AD19	AH25	A_CBE_L7	AD17	B_AD14	U4	B_CBE_L2	Y1
A_AD20	AG25	A_COMPAT	K24	B_AD15	U5	B_CBE_L3	AA4
A_AD21	AF25	A_DEVSEL#	AG23	B_AD16	Y2	B_CBE_L4	M1
A_AD22	AE25	A_FRAME#	AD23	B_AD17	Y3	B_CBE_L5	N5
A_AD23	AD25	A_GNT0#	AC25	B_AD18	Y4	B_CBE_L6	N3
A_AD24	AH26	A_GNT1#	U25	B_AD19	Y5	B_CBE_L7	N1
A_AD25	AG26	A_GNT2#	R25	B_AD20	Y6	B_DEVSEL#	V6
A_AD26	AF26	A_GNT3#	P25	B_AD21	AA1	B_FRAME#	W5
A_AD27	AJ27	A_GNT4#	N24	B_AD22	AA2	B_GNT0#	AC5
A_AD28	AG27	A_IRDY#	AE23	B_AD23	AA3	B_GNT1#	AF3
A_AD29	AH28	A_M66EN	AE20	B_AD24	AA5	B_GNT2#	AG6
A_AD30	AG28	A_PAR	AF22	B_AD25	AA6	B_GNT3#	AG7
A_AD31	AF28	A_PAR64	AH17	B_AD26	AB1	B_GNT4#	AG8
A_AD32	AJ10	A_PCIXCAP	AH23	B_AD27	AB3	B_IRDY#	W3
A_AD33	AD11	A_PCLK0	AB24	B_AD28	AB5	B_M66EN	T3
A_AD34	AE11	A_PCLK1	T24	B_AD29	AC1	B_PAR	V1
A_AD35	AF11	A_PCLK2	R24	B_AD30	AC2	B_PAR64	M2
A_AD36	AG11	A_PCLK3	P24	B_AD31	AC3	B_PCIXCAP	V5
A_AD37	AH11	A_PCLK4	N25	B_AD32	F4	B_PCLK0	AC6
A_AD38	AJ11	A_PERR#	AD22	B_AD33	F3	B_PCLK1	AF4
A_AD39	AE12	A_PIRQA#	U24	B_AD34	F2	B_PCLK2	AH6
A_AD40	AG12	A_PIRQB#	V25	B_AD35	F1	B_PCLK3	AH7
A_AD41	AJ12	A_PIRQC#	AA24	B_AD36	G5	B_PCLK4	AF9
A_AD42	AD13	A_PIRQD#	AB25	B_AD37	G3	B_PERR#	V3
A_AD43	AE13	A_PLLCLKI	M25	B_AD38	G1	B_PIRQA#	AE1
A_AD44	AF13	A_PLLCLKO	M24	B_AD39	H6	B_PIRQB#	AD4
A_AD45	AG13	A_PME#	AD7	B_AD40	H5	B_PIRQC#	AD3
A_AD46	AH13	A_PRESET#	AC26	B_AD41	H4	B_PIRQD#	AD2
A_AD47	AJ13	A_REQ0#	AC24	B_AD42	H3	B_PLLCLKI	AE6
A_AD48	AD14	A_REQ1#	V24	B_AD43	H2	B_PLLCLKO	AE7
A_AD49	AE14	A_REQ2#	T25	B_AD44	H1	B_PME#	AD6
A_AD50	AF14	A_REQ3#	R26	B_AD45	J6	B_PRESET#	AD1

Table 18: Alphabetical listing of signals A_ACK64# to B_PRESET#.

Signal name	Ball	Signal name	Ball	Signal name	Ball	Signal name	Ball	Signal name	Ball
B_REQ0#	AC4	LRACAD_N12	D20	LTACAD_N6	C13	NC0	AG4	VDD12B	V20
B_REQ1#	AF1	LRACAD_N13	E18	LTACAD_N7	A15	NC1	AG5	VDD12B	W19
B_REQ2#	AF6	LRACAD_N14	D18	LTACAD_N8	C6	NC2	AH3	VDD12B	W21
B_REQ3#	AF7	LRACAD_N15	E16	LTACAD_N9	E7	NC3	AH4	VDD18	AA23
B_REQ4#	AE8	LRACAD_P0	C25	LTACAD_N10	C8	NIOAIRQA#	F25	VDD18	AC28
B_REQ64#	P6	LRACAD_P1	A25	LTACAD_N11	E9	NIOAIRQB#	G24	VDD18	AD26
B_SERR#	V2	LRACAD_P2	C23	LTACAD_N12	E11	NIOAIRQC#	G26	VDD18	B12
B_STOP#	V4	LRACAD_P3	A23	LTACAD_N13	C12	NIOAIRQD#	G25	VDD18	B16
B_TRDY#	W1	LRACAD_P4	A21	LTACAD_N14	E13	P_CAL	AE3	VDD18	B20
CMPOVR	A27	LRACAD_P5	C19	LTACAD_N15	C14	P_CAL#	AF2	VDD18	B24
FREE1	D4	LRACAD_P6	A19	LTACAD_P0	A6	PWROK	C28	VDD18	B8
FREE2	AJ5	LRACAD_P7	C17	LTACAD_P1	B7	REFCLK	C29	VDD18	D11
FREE3	AG9	LRACAD_P8	E25	LTACAD_P2	A8	RESET#	B28	VDD18	D15
FREE4	AJ8	LRACAD_P9	C24	LTACAD_P3	B9	RSVD0	E14	VDD18	D19
FREE5	F5	LRACAD_P10	E23	LTACAD_P4	B11	RSVD1	E15	VDD18	D23
FREE6	AD5	LRACAD_P11	C22	LTACAD_P5	A12	RSVD2	D16	VDD18	D7
FREE7	AJ9	LRACAD_P12	C20	LTACAD_P6	B13	RSVD3	C16	VDD18	F10
FREE8	AD10	LRACAD_P13	E19	LTACAD_P7	A14	RSVD4	J27	VDD18	F12
FREE9	AJ3	LRACAD_P14	C18	LTACAD_P8	D6	RSVD5	K25	VDD18	F14
FREE10	D2	LRACAD_P15	E17	LTACAD_P9	E6	RSVD6	L24	VDD18	F16
FREE12	D1	LRACLK0_N	B21	LTACAD_P10	D8	RSVD7	L25	VDD18	F18
FREE13	E3	LRACLK0_P	C21	LTACAD_P11	E8	RSVD8	L27	VDD18	F20
FREE14	E4	LRACLK1_N	E20	LTACAD_P12	E10	RSVD9	N26	VDD18	F22
FREE15	AD9	LRACLK1_P	E21	LTACAD_P13	D12	RSVD10	N27	VDD18	F24
FREE16	AG1	LRACLK_N	A16	LTACAD_P14	E12	RSVD11	R27	VDD18	F8
FREE17	AE5	LRACLK_P	A17	LTACAD_P15	D14	RSVD12	U26	VDD18	G11
FREE18	E1	LRBCAD_N0	AD28	LTACLK0_N	A11	RSVD13	U27	VDD18	G13
FREE19	E2	LRBCAD_N1	AC29	LTACLK0_P	A10	RSVD14	W24	VDD18	G15
FREE20	AG2	LRBCAD_N2	AB28	LTACLK1_N	C10	RSVD15	W25	VDD18	G17
FREE21	AJ4	LRBCAD_N3	AA29	LTACLK1_P	D10	RSVD16	W26	VDD18	G19
FREE22	AE9	LRBCAD_N4	W29	LTACTL_N0	C15	RSVD17	W27	VDD18	G21
HPSIC	AJ6	LRBCAD_N5	V28	LTACTL_P0	B15	RSVD18	Y25	VDD18	G23
HPSIL#	J26	LRBCAD_N6	U29	LTBCAD_N0	F29	RSVD19	AA25	VDD18	G28
HPSOC	AJ7	LRBCAD_N7	T28	LTBCAD_N1	F27	RSVD20	AA26	VDD18	G7
HPSOD	J25	LRBCAD_P0	AD27	LTBCAD_N2	H29	RSVD21	AA27	VDD18	G9
LDTCOMP0	B4	LRBCAD_P1	AD29	LTBCAD_N3	H27	RSVD22	AC27	VDD18	H12
LDTCOMP1	A3	LRBCAD_P2	AB27	LTBCAD_N4	K27	STRAPL2	B27	VDD18	H14
LDTCOMP2	B2	LRBCAD_P3	AB29	LTBCAD_N5	M29	STRAPL3	H25	VDD18	H16
LDTCOMP3	C3	LRBCAD_P4	Y29	LTBCAD_N6	M27	TEST	D27	VDD18	H20
LDTSTOP#	C27	LRBCAD_P5	V27	LTBCAD_N7	P29	VDD12A	H10	VDD18	H22
LRACAD_N0	B25	LRBCAD_P6	V29	LTBCAD_P0	E29	VDD12A	H18	VDD18	H26
LRACAD_N1	A24	LRBCAD_P7	T27	LTBCAD_P1	F28	VDD12A	J11	VDD18	H8
LRACAD_N2	B23	LRBCLK0_N	Y28	LTBCAD_P2	G29	VDD12A	J17	VDD18	J13
LRACAD_N3	A22	LRBCLK0_P	Y27	LTBCAD_P3	H28	VDD12A	K10	VDD18	J15
LRACAD_N4	A20	LRBCTL_N	R29	LTBCAD_P4	K28	VDD12A	K18	VDD18	J19
LRACAD_N5	B19	LRBCTL_P	T29	LTBCAD_P5	L29	VDD12A	L11	VDD18	J21
LRACAD_N6	A18	LTACAD_N0	A7	LTBCAD_P6	M28	VDD12A	L17	VDD18	J23

Table 19: Alphabetical listing of signals B_REQ0# to VDD18.

LRACAD_N7	B17	LTACAD_N1	C7	LTBCAD_P7	N29	VDD12B	L19	VDD18	K12
LRACAD_N8	E24	LTACAD_N2	A9	LTBCLK0_N	K29	VDD12B	L21	VDD18	K14
LRACAD_N9	D24	LTACAD_N3	C9	LTBCLK0_P	J29	VDD12B	M18	VDD18	K16
LRACAD_N10	E22	LTACAD_N4	C11	LTBCTL_N	P27	VDD12B	M20	VDD18	K20
LRACAD_N11	D22	LTACAD_N5	A13	LTBCTL_P	P28	VDD12B	V18	VDD18	K22

Table 19: Alphabetical listing of signals B_REQ0# to VDD18.


Signal name	Ball	Signal name	Ball	Signal name	Ball	Signal name	Ball	Signal name	Ball	Signal name	Ball	Signal name	Ball
VDD18	L13	VDD33	AA17	VDD33	P8	VSS	AB7	VSS	D3	VSS	K19	VSS	R8
VDD18	L15	VDD33	AA19	VDD33	R11	VSS	AB9	VSS	D5	VSS	K2	VSS	T11
VDD18	L23	VDD33	AA21	VDD33	R7	VSS	AC10	VSS	D9	VSS	K21	VSS	T13
VDD18	L28	VDD33	AA7	VDD33	R9	VSS	AC12	VSS	E26	VSS	K23	VSS	T15
VDD18	M12	VDD33	AA9	VDD33	T10	VSS	AC14	VSS	E27	VSS	K26	VSS	T17
VDD18	M14	VDD33	AB10	VDD33	T4	VSS	AC16	VSS	E28	VSS	K7	VSS	T19
VDD18	M16	VDD33	AB12	VDD33	T6	VSS	AC18	VSS	E5	VSS	K9	VSS	T2
VDD18	M22	VDD33	AB14	VDD33	T8	VSS	AC20	VSS	F11	VSS	L10	VSS	T21
VDD18	M26	VDD33	AB16	VDD33	U11	VSS	AC22	VSS	F13	VSS	L12	VSS	T23
VDD18	N13	VDD33	AB18	VDD33	U7	VSS	AC8	VSS	F15	VSS	L14	VSS	T7
VDD18	N15	VDD33	AB20	VDD33	U9	VSS	AD15	VSS	F17	VSS	L16	VSS	T9
VDD18	N17	VDD33	AB22	VDD33	V10	VSS	AD21	VSS	F19	VSS	L18	VSS	U10
VDD18	N19	VDD33	AB4	VDD33	V8	VSS	AD8	VSS	F21	VSS	L20	VSS	U12
VDD18	N21	VDD33	AB6	VDD33	W11	VSS	AE10	VSS	F23	VSS	L22	VSS	U14
VDD18	N23	VDD33	AB8	VDD33	W13	VSS	AE2	VSS	F26	VSS	L8	VSS	U16
VDD18	P12	VDD33	AC11	VDD33	W15	VSS	AE26	VSS	F6	VSS	M11	VSS	U18
VDD18	P14	VDD33	AC13	VDD33	W17	VSS	AE27	VSS	F7	VSS	M13	VSS	U20
VDD18	P16	VDD33	AC15	VDD33	W4	VSS	AE28	VSS	F9	VSS	M15	VSS	U22
VDD18	P18	VDD33	AC17	VDD33	W7	VSS	AE29	VSS	G10	VSS	M17	VSS	U28
VDD18	P20	VDD33	AC19	VDD33	W9	VSS	AF10	VSS	G12	VSS	M19	VSS	U8
VDD18	P22	VDD33	AC21	VDD33	Y10	VSS	AG10	VSS	G14	VSS	M21	VSS	V11
VDD18	R13	VDD33	AC23	VDD33	Y12	VSS	AH10	VSS	G16	VSS	M23	VSS	V13
VDD18	R15	VDD33	AC7	VDD33	Y14	VSS	AH12	VSS	G18	VSS	M7	VSS	V15
VDD18	R17	VDD33	AC9	VDD33	Y16	VSS	AH15	VSS	G2	VSS	M9	VSS	V17
VDD18	R19	VDD33	AD12	VDD33	Y18	VSS	AH18	VSS	G20	VSS	N10	VSS	V19
VDD18	R21	VDD33	AD18	VDD33	Y20	VSS	AH2	VSS	G22	VSS	N12	VSS	V21
VDD18	R23	VDD33	AD24	VDD33	Y22	VSS	AH21	VSS	G27	VSS	N14	VSS	V23
VDD18	R28	VDD33	AE4	VDD33	Y8	VSS	AH24	VSS	G6	VSS	N16	VSS	V26
VDD18	T12	VDD33	AF12	VDD33	H24	VSS	AH27	VSS	G8	VSS	N18	VSS	V7
VDD18	T14	VDD33	AF15	VDDA18	AF29	VSS	AH5	VSS	H11	VSS	N2	VSS	V9
VDD18	T16	VDD33	AF18	VDDA18	AG29	VSS	AH8	VSS	H13	VSS	N20	VSS	W10
VDD18	T18	VDD33	AF21	VSS	A26	VSS	AH9	VSS	H15	VSS	N22	VSS	W12
VDD18	T20	VDD33	AF24	VSS	A4	VSS	B10	VSS	H17	VSS	N28	VSS	W14
VDD18	T22	VDD33	AF27	VSS	A5	VSS	B14	VSS	H19	VSS	N6	VSS	W16
VDD18	T26	VDD33	AF5	VSS	AA10	VSS	B18	VSS	H21	VSS	N8	VSS	W18
VDD18	U13	VDD33	AF8	VSS	AA12	VSS	B22	VSS	H23	VSS	P11	VSS	W2
VDD18	U15	VDD33	AG3	VSS	AA14	VSS	B26	VSS	H7	VSS	P13	VSS	W20
VDD18	U17	VDD33	G4	VSS	AA16	VSS	B3	VSS	H9	VSS	P15	VSS	W22
VDD18	U19	VDD33	J7	VSS	AA18	VSS	B5	VSS	J10	VSS	P17	VSS	W6
VDD18	U21	VDD33	J9	VSS	AA20	VSS	B6	VSS	J12	VSS	P19	VSS	W8
VDD18	U23	VDD33	K4	VSS	AA22	VSS	C1	VSS	J14	VSS	P21	VSS	Y11
VDD18	V12	VDD33	K6	VSS	AA28	VSS	C2	VSS	J16	VSS	P23	VSS	Y13
VDD18	V14	VDD33	K8	VSS	AA8	VSS	C26	VSS	J18	VSS	P26	VSS	Y15
VDD18	V16	VDD33	L7	VSS	AB11	VSS	C4	VSS	J20	VSS	P7	VSS	Y17
VDD18	V22	VDD33	L9	VSS	AB13	VSS	C5	VSS	J22	VSS	P9	VSS	Y19

Table 20: Alphabetical listing of signals VDD18 to VSS.

VDD18	W23	VDD33	M10	VSS	AB15	VSS	D13	VSS	J24	VSS	R10	VSS	Y21
VDD18	W28	VDD33	M8	VSS	AB17	VSS	D17	VSS	J28	VSS	R12	VSS	Y23
VDD18	Y24	VDD33	N11	VSS	AB19	VSS	D21	VSS	J8	VSS	R14	VSS	Y7
VDD18	Y26	VDD33	N4	VSS	AB2	VSS	D25	VSS	K11	VSS	R16	VSS	Y9
VDD33	AA11	VDD33	N7	VSS	AB21	VSS	D26	VSS	K13	VSS	R18		
VDD33	AA13	VDD33	N9	VSS	AB23	VSS	D28	VSS	K15	VSS	R20		
VDD33	AA15	VDD33	P10	VSS	AB26	VSS	D29	VSS	K17	VSS	R22		

Table 20: Alphabetical listing of signals VDD18 to VSS.

8 Package Specification



GENERAL NOTES

- All dimensions are specified in millimeters (mm).
- Dimensioning and tolerancing per ASME-Y14.5M-1994.
- This corner which consists of a triangle on both sides of the package identifies ball A1 corner and can be used for handling and orientation purposes.
- Symbol "M" determines ball matrix size and "N" is number of balls.
- Dimension "b" is measured at maximum solder ball diameter on a plane parallel to datum C.
- "x" in front of package variation denotes non-qualified package per AMD 01-002.3.
- The following features are not shown on drawings:
 - Marking on die, label on package
 - Laser elements
 - Die and passive fudicials

Figure 14: Package mechanical drawing.

9 Test

The IC includes the following test modes.

Mode	TEST	A_REQ2#	A_REQ1#	A_REQ0#	Notes
Operational	0	X	X	X	
High impedance	1	0	0	0	
NAND tree	1	0	0	1	

Table 21: Test modes.

9.1 High Impedance Mode

In high-impedance mode, all the signals of the IC are placed into the high-impedance state.

9.2 NAND Tree Mode

There are several NAND trees in the IC. Some of the inputs are differential (e.g., LR[B, A] pins); for these, the _P and _N pairs of signals are converted into a single signal that is part of the NAND tree, as shown in Signal_3 in the following diagram.

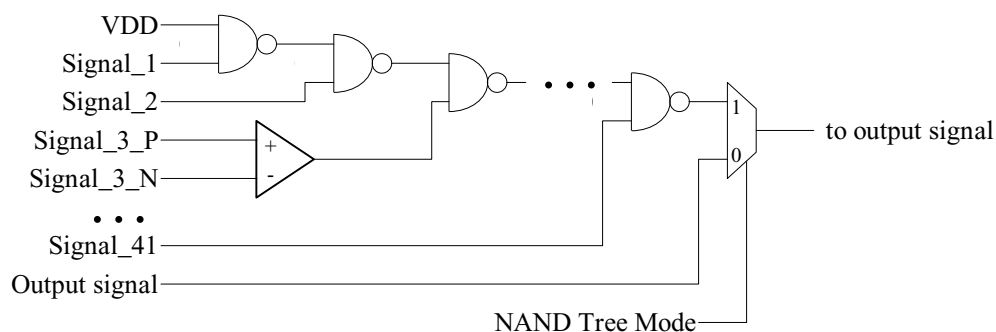


Figure 15: NAND tree.

NAND tree 1: output signal is B_REQ[3]#. However, the gate connected to the last signal in this NAND tree (LDTCOMP[3]) is an AND gate rather than a NAND gate; so the expected output of this NAND tree is inverted compared to the other NAND trees.

1	LRBCLK0_[P,N]	11	LTBCLK0_P	21	LTBCAD_P[4]	31	LDTCOMP[2]
2	LRBCAD_[P,N][0]	12	LTBCLK0_N	22	LTBCAD_N[4]	32	LDTCOMP[3]
3	LRBCAD_[P,N][1]	13	LTBCAD_P[0]	23	LTBCAD_P[5]		
4	LRBCAD_[P,N][2]	14	LTBCAD_N[0]	24	LTBCAD_N[5]		
5	LRBCAD_[P,N][3]	15	LTBCAD_P[1]	25	LTBCAD_P[6]		
6	LRBCAD_[P,N][4]	16	LTBCAD_N[1]	26	LTBCAD_N[6]		
7	LRBCAD_[P,N][5]	17	LTBCAD_P[2]	27	LTBCAD_P[7]		
8	LRBCAD_[P,N][6]	18	LTBCAD_N[2]	28	LTBCAD_N[7]		
9	LRBCAD_[P,N][7]	19	LTBCAD_P[3]	29	LTBCTL_P		
10	LRBCTL_[P,N]	20	LTBCAD_N[3]	30	LTBCTL_N		

NAND tree 2: output signal is B_REQ[2]#.

1	LRACLK0_[P,N]	21	LTACLK0_N	41	LTACAD_N[4]
2	LRACLK1_[P,N]	22	LTACLK1_P	42	LTACAD_P[12]
3	LRACAD_[P,N][0]	23	LTACLK1_N	43	LTACAD_N[12]
4	LRACAD_[P,N][8]	24	LTACAD_P[0]	44	LTACAD_P[5]
5	LRACAD_[P,N][1]	25	LTACAD_N[0]	45	LTACAD_N[5]
6	LRACAD_[P,N][9]	26	LTACAD_P[8]	46	LTACAD_P[13]
7	LRACAD_[P,N][2]	27	LTACAD_N[8]	47	LTACAD_N[13]
8	LRACAD_[P,N][10]	28	LTACAD_P[1]	48	LTACAD_P[6]
9	LRACAD_[P,N][3]	29	LTACAD_N[1]	49	LTACAD_N[6]
10	LRACAD_[P,N][11]	30	LTACAD_P[9]	50	LTACAD_P[14]
11	LRACAD_[P,N][4]	31	LTACAD_N[9]	51	LTACAD_N[14]
12	LRACAD_[P,N][12]	32	LTACAD_P[2]	52	LTACAD_P[7]
13	LRACAD_[P,N][5]	33	LTACAD_N[2]	53	LTACAD_N[7]
14	LRACAD_[P,N][13]	34	LTACAD_P[10]	54	LTACAD_P[15]
15	LRACAD_[P,N][6]	35	LTACAD_N[10]	55	LTACAD_N[15]
16	LRACAD_[P,N][14]	36	LTACAD_P[3]	56	LTACTL_P
17	LRACAD_[P,N][7]	37	LTACAD_N[3]	57	LTACTL_N
18	LRACAD_[P,N][15]	38	LTACAD_P[11]		
19	LRACL_[P,N]	39	LTACAD_N[11]		
20	LTACLK0_P	40	LTACAD_P[4]		

NAND tree 3: output signal is B_REQ[1]#.

1	B_AD[32]	22	B_AD[54]	43	B_AD[3]	64	B_DEVSEL#	85	B_PIRQD#
2	B_AD[36]	23	B_AD[55]	44	B_AD[8]	65	B_IRDY#	86	B_PIRQA#
3	B_AD[33]	24	B_AD[53]	45	B_CBE_L[0]	66	B_CBE_L[2]	87	B_AD[28]
4	B_AD[34]	25	B_AD[56]	46	B_AD[7]	67	B_AD[16]	88	B_PIRQC#
5	B_AD[39]	26	B_AD[57]	47	B_AD[6]	68	B_FRAME#	89	B_GNT[0]#
6	B_AD[35]	27	B_AD[58]	48	B_AD[5]	69	B_AD[17]	90	B_PIRQB#
7	B_AD[37]	28	B_AD[60]	49	B_AD[4]	70	B_AD[21]	91	B_PCLK[1]
8	B_AD[40]	29	B_AD[59]	50	B_AD[9]	71	B_AD[18]	92	B_PCLK[0]
9	B_AD[41]	30	B_AD[61]	51	B_M66EN	72	B_AD[22]	93	B_GNT[1]#
10	B_AD[38]	31	B_AD[62]	52	B_AD[10]	73	B_AD[19]	94	B_PLLCLKI
11	B_AD[45]	32	B_AD[63]	53	B_AD[11]	74	B_AD[23]	95	B_GNT[2]#
12	B_AD[42]	33	B_PAR64	54	B_AD[12]	75	B_AD[20]	96	B_PCLK[2]
13	B_AD[46]	34	B_CBE_L[4]	55	B_AD[13]	76	B_CBE_L[3]	97	HPSIC
14	B_AD[43]	35	B_CBE_L[5]	56	B_AD[14]	77	B_AD[26]	98	B_PLLCLKO
15	B_AD[44]	36	B_CBE_L[6]	57	B_AD[15]	78	B_AD[27]	99	B_REQ[4]#
16	B_AD[47]	37	B_CBE_L[7]	58	B_CBE_L[1]	79	B_AD[29]	100	B_GNT[3]#
17	B_AD[48]	38	B_REQ64#	59	B_PAR	80	B_AD[24]	101	B_PCLK[4]
18	B_AD[49]	39	B_ACK64#	60	B_SERR#	81	B_AD[30]	102	B_GNT[4]#
19	B_AD[50]	40	B_AD[0]	61	B_PERR#	82	B_PRESET#	103	B_PCLK[3]
20	B_AD[51]	41	B_AD[1]	62	B_STOP#	83	B_AD[25]	104	HPSOC
21	B_AD[52]	42	B_AD[2]	63	B_TRDY#	84	B_AD[31]		

NAND tree 4: output signal is A_REQ[3]#.

1	A_AD[32]	22	A_AD[53]	43	A_AD[5]	64	A_TRDY#	85	A_GNT[0]#
2	A_AD[33]	23	A_AD[54]	44	A_AD[6]	65	A_FRAME#	86	A_PRESET#
3	A_AD[34]	24	A_AD[55]	45	A_AD[2]	66	A_IRDY#	87	A_PCLK[0]
4	A_AD[35]	25	A_AD[56]	46	A_AD[1]	67	A_CBE_L[2]	88	A_PIRQD#
5	A_AD[36]	26	A_AD[62]	47	A_CBE_L[0]	68	A_AD[18]	89	A_PIRQC#
6	A_AD[37]	27	A_AD[61]	48	A_AD[8]	69	A_CBE_L[3]	90	A_PIRQB#
7	A_AD[38]	28	A_AD[60]	49	A_AD[9]	70	A_AD[27]	91	A_GNT[1]#
8	A_AD[39]	29	A_AD[59]	50	A_M66EN	71	A_AD[16]	92	A_PIRQA#
9	A_AD[40]	30	A_AD[58]	51	A_AD[7]	72	A_AD[19]	93	A_PCLK[1]
10	A_AD[41]	31	A_AD[57]	52	A_AD[11]	73	A_AD[24]	94	A_GNT[2]#
11	A_AD[42]	32	A_CBE_L[4]	53	A_AD[12]	74	A_AD[29]	95	A_PCLK[2]
12	A_AD[43]	33	A_CBE_L[5]	54	A_AD[13]	75	A_AD[20]	96	A_PCLK[3]
13	A_AD[44]	34	A_CBE_L[6]	55	A_AD[10]	76	A_AD[25]	97	A_GNT[3]#
14	A_AD[45]	35	A_CBE_L[7]	56	A_AD[14]	77	A_AD[28]	98	A_GNT[4]#
15	A_AD[46]	36	A_PAR64	57	A_AD[15]	78	A_AD[21]	99	A_PCLK[4]
16	A_AD[47]	37	A_AD[63]	58	A_CBE_L[1]	79	A_AD[26]	100	A_PLLCLKI
17	A_AD[48]	38	A_AD[0]	59	A_PAR	80	A_AD[22]	101	A_REQ[4]#
18	A_AD[49]	39	A_ACK64#	60	A_SERR#	81	A_AD[30]	102	A_PLLCLKO
19	A_AD[50]	40	A_REQ64#	61	A_STOP#	82	A_AD[31]	103	A_COMPAT
20	A_AD[51]	41	A_AD[4]	62	A_PERR#	83	A_AD[17]		
21	A_AD[52]	42	A_AD[3]	63	A_DEVSEL#	84	A_AD[23]		

NAND tree 5: output signal is B_REQ[0]#.

1	HPSOD	6	NIOAIRQB#	11	RESET#
2	HPSIL#	7	REFCLK	12	STRAPL[2]
3	STRAPL[3]	8	NIOAIRQA#	13	CMPOVR
4	NIOAIRQC#	9	PWROK		
5	NIOAIRQD#	10	LDTSTOP#		

Notes:

- [B, A]_PCIXCAP, A_REQ[2:0]#, TEST, LDTCOMP[1:0], P_CAL, P_CAL#, [B, A]_PME# are not included in the NAND trees.
- While in NAND-tree mode, the link and PCI-X compensation is placed at a “mid-band” value.
- Internal PLLs are disabled by placing them in bypass mode

10 Appendix

10.1 Revision History

Revision 3.01

- Initial Release.

Revision 3.02

- Removed Preliminary

- Added Error Conditions and Handling.