



Intel® 6702PXH 64-bit PCI Hub

Specification Update

August 2005

Notice: Intel® 6702PXH 64-bit PCI Hub product may contain design defects or errors known as errata which may cause the product to deviate from published specifications. Current characterized errata are documented in this specification update.

INFORMATION IN THIS DOCUMENT IS PROVIDED IN CONNECTION WITH INTEL® PRODUCTS. NO LICENSE, EXPRESS OR IMPLIED, BY ESTOPPEL OR OTHERWISE, TO ANY INTELLECTUAL PROPERTY RIGHTS IS GRANTED BY THIS DOCUMENT. EXCEPT AS PROVIDED IN INTEL'S TERMS AND CONDITIONS OF SALE FOR SUCH PRODUCTS, INTEL ASSUMES NO LIABILITY WHATSOEVER, AND INTEL DISCLAIMS ANY EXPRESS OR IMPLIED WARRANTY, RELATING TO SALE AND/OR USE OF INTEL PRODUCTS INCLUDING LIABILITY OR WARRANTIES RELATING TO FITNESS FOR A PARTICULAR PURPOSE, MERCHANTABILITY, OR INFRINGEMENT OF ANY PATENT, COPYRIGHT OR OTHER INTELLECTUAL PROPERTY RIGHT. Intel products are not intended for use in medical, life saving, or life sustaining applications.

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

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

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

Intel and the Intel logo are trademarks or registered trademarks of Intel Corporation or its subsidiaries in the United States and other countries.

*Other names and brands may be claimed as the property of others.

Copyright © 2004-2005, Intel Corporation



Contents

Revision History	4
Preface	5
Summary Table of Changes	6
Identification Information	9
Errata	10
Specification Changes	16
Specification Clarifications	17
Documentation Changes	18

Revision History

Revision	Description	Date
-001	<ul style="list-style-type: none">Initial Release	September 2004
-002	<ul style="list-style-type: none">Added erratum 13; added Documentation Changes 1-8.	November 2004
-003	<ul style="list-style-type: none">Added erratum 14; added Documentation Change 9.	January 2005
-004	<ul style="list-style-type: none">Added Documentation Change 10.	March 2005
-005	<ul style="list-style-type: none">Added Specification Clarification 2; updated Documentation Change 10; added Documentation Changes 11-16.	May 2005
-006	<ul style="list-style-type: none">Added errata 15-16.	June 2005
-007	<ul style="list-style-type: none">Added erratum 17.	August 2005

Preface

This document is an update to the specifications contained in the Affected Documents/Related Documents table below. This document is a compilation of device and documentation errata, specification clarifications and changes. It is intended for hardware system manufacturers and software developers of applications, operating systems, or tools.

Information types defined in Nomenclature are consolidated into the specification update and are no longer published in other documents.

This document may also contain information that was not previously published.

Note: This document contains information related to the Intel® 6702PXH 64-bit PCI Hub component and their respective Datasheet document. Please refer to the Affected Documents section of each item to determine which component or document is affected by that particular item.

Affected Documents/Related Documents

Title	Document #
Intel® 6702PXH 64-bit PCI Hub Datasheet	303633

Nomenclature

Errata are design defects or errors. Errata may cause the Intel 6702PXH 64-bit PCI Hub's behavior to deviate from published specifications. Hardware and software designed to be used with any given stepping must assume that all errata documented for that stepping are present in all devices.

Specification Changes are modifications to the current published specifications. These changes will be incorporated in any new release of the specification.

Specification Clarifications describe a specification in greater detail or further highlight a specification's impact to a complex design situation. These clarifications will be incorporated in any new release of the specification.

Documentation Changes include typos, errors, or omissions from the current published specifications. These will be incorporated in any new release of the specification.

Summary Table of Changes

The following table indicates the Specification Changes, Errata, or Documentation Changes which apply to the Intel 6702PXH 64-bit PCI Hub. Intel may fix some of the errata in a future stepping of the component, and account for the other outstanding issues through documentation or specification changes as noted. This table uses the following notations:

Codes Used in Summary Table

Stepping

X: Errata exists in the stepping indicated. Specification Change or Clarification that applies to this stepping.

(No mark) or (Blank Box): This erratum is fixed in listed stepping or specification change does not apply to listed stepping.

Status

Doc:	Document change or update that will be implemented.
Plan Fix:	This erratum may be fixed in a future stepping of the product.
Fixed:	This erratum has been previously fixed.
No Fix:	There are no plans to fix this erratum.
	Change bar to left of table row indicates this item is either new or modified from the previous version of this document.

Errata

No.	Stepping	Status	ERRATA/SIGHTINGS
	C1		
1	X	No Fix	Secondary bus PxPCIRST# pulse prior to the rising edge of PWROK
2	X	No Fix	Unreliable PCI Express* link operation when L0s Active State Power Management is enabled
3	X	No Fix	Signal integrity issues may occur when the Intel 6702PXH 64-bit PCI Hub is driving the secondary bus
4	X	No Fix	SSE bit set for PERR# assertion when error reporting is masked
5	X	No Fix	Data parity error detected on PCI-X interface fails to propagate bad parity
6	X	No Fix	Intel 6702PXH 64-bit PCI Hub fails to train down in presence of degraded lane
7	X	No Fix	PCI Express and PCI-X header logs and First Error Pointers do not remain sticky through reset
8	X	No Fix	Incorrect default value for PCI Express Flow Control Protocol Error Severity bit
9	X	No Fix	Power State bits in PCI Express Power Management Status and Control Register mistakenly accept reserved values
10	X	No Fix	Performance across an upstream x1 PCI Express link is less than expected
11	X	No Fix	SKP ordered set may not be sent within required interval during Link Recovery if a packet is pending
12	X	No Fix	SERR Fatal/Non-Fatal Error message enabled with incorrect Error Message Enabled bit
13	X	No Fix	Configuration write to offset x70 of D0:F0, F2 (EXP_PMSTSCNTL - PCI Express Power Management Status and Control Register) using non-continuous byte enables does not capture the data value
14	X	No Fix	PXH may become unresponsive when transitioning into the D3 power state
15	X	No Fix	MSI Hot-Plug Interrupt issue
16	X	No Fix	Tcyc(min) of the PXH output clocks deviates from the PCI-X (Mode 1, class1) jitter clock specification
17	X	No Fix	Under certain conditions, inbound prefetched PCI read requests may return wrong data to the requestor

Specification Changes

No.	Specification Changes
1	PCI Express L0s functionality not supported in Intel 6702PXH 64-bit PCI Hub
2	CMODE Power-on Strap for PCI Express 1.0a support

Specification Clarifications

No.	Specification Clarifications
1	The relationship between PWROK and the PCI Express clocks timing specification clarification
2	Table 4-20 will be updated to include the following information in columns PCI-X66, Conv. PCI 66 (ref), and Conv. PCI 33 (ref)

Documentation Changes

No.	Documentation Changes
1	Updated Section 2.7, Table 2-9
2	Updated Section 2.17.2.1
3	In Section 3.8.1.6, added C1 stepping, removed all other stepping and changed default value to I/OxAPIC REVID register (Offset 08h)
4	In Section 3.4.1.39, updated EXP_LCAP (Offset 50h) bit 14:12 description and register's default value
5	In Section 3.8.1.6, added C1 Stepping to I/OxAPIC REVID register (Offset 08h)
6	In Section 3.5.1.39, updated EXP_LCAP (Offset 50h) bit 14:12 description and register's default value
7	In Section 3.5.1.32, updated PCLKC register (Offset 43h) description
8	In Section 3.5.1.41, documentation error in EXP_LSTS (offset 56h) bit field description
9	Sections 2.19.2.2 (Error Logging) and 2.19.2.3 (Error Escalation) are replaced by a combined PCI error section (Section 2.19.2.2, "Error Logging and Escalation") with the following description
10	Section 1.2.1 (page 12), the documentation incorrectly includes "lane reversal" that PXH does not support
11	Figures 5-1 will be updated. The following pins in the figures need to match the pin list:
12	Table 2-2 (page 17) will be updated
13	Table 2-7 (page 23) will be updated
14	Table 2-7 (page 22) will be updated
15	Table 2-7 (page 21) will be updated
16	Section 2.13.6.1 (page 50-51) will be updated

Identification Information

Component Marking Information

The Intel 6702PXH 64-bit PCI Hub stepping can be identified by the following component markings:

Stepping	S-Spec	Top Marking	Notes
C1	SL7X4	865149	Production

Errata

1. Secondary bus PxPCIRST# pulse prior to the rising edge of PWROK

Problem: During system power on and prior to the Intel 6702PXH 64-bit PCI Hub receiving the rising edge of PWROK, a pulse is observed on the secondary bus PxPCIRST# signals.

Implication: PCI/PCI-X* controllers on the secondary bus segments could interpret this PxPCIRST# pulse as a true rising edge and initialize into an undetermined state.

Workaround: A hardware workaround has been identified. The PWROK signal that is received by the Intel 6702PXH 64-bit PCI Hub component should be used to gate the secondary bus PxPCIRST# signals.

Status: For the steppings affected, see the *Summary Table of Changes*.

2. Unreliable PCI Express* link operation when L0s Active State Power Management is enabled

Problem: PCI Express* link operation is unreliable after the L0s state is enabled in the Intel 6702PXH 64-bit PCI Hub.

Implication: System hangs and other system instability if L0s is enabled.

Workaround: None at this time.

Status: For the steppings affected, see the *Summary Table of Changes*.

3. Signal integrity issues may occur when the Intel 6702PXH 64-bit PCI Hub is driving the secondary bus

Problem: Signal integrity issues may occur at when Intel 6702PXH 64-bit PCI Hub is driving the PCI/PCI-X bus in Conventional PCI or PCI-X mode 1.

Implication: Parity errors and system hangs may occur.

Workaround: A BIOS workaround has been identified. Please refer to the latest version of the *Intel® 6702PXH 64-bit PCI Hub BIOS Specification Update* for details on the workaround. The workaround is required for reliable Conventional PCI and PCI-X mode 1 operation and must be left in place for all steppings of Intel® 6702PXH 64-bit PCI Hub.

Status: For the steppings affected, see the *Summary Table of Changes*.

4. SSE bit set for PERR# assertion when error reporting is masked

Problem: During a downstream memory write to Intel 6702PXH 64-bit PCI Hub, the following erroneous behavior is seen when PERR# is asserted on the secondary bus:

- Signaled System Error (SSE) in the STS_REG register (D:0, F:0&2, offset 06h, bit 14) is set when SERR# Enable (SEE) (D:0, F:0&2, offset 04h, bit 8) and Parity Error Response Enable (PERE) (D:0, F:0&2, offset 04h, bit 6) are set in the CMD register.
- The PERE bit in the BRDG_CNTL register is set (D:0, F:0&2, offset 3Eh, bit 0).
- Error reporting is disabled in the UNC_PXERRMSK register (D:0, F:0&2, offset 130h).

Implication: False indication of an error message escalated as recorded in SSE of the STS_REG register being set. This is considered low risk since the escalation of the message is functioning properly.

Workaround: None at this time.

Status: For the steppings affected, see the *Summary Table of Changes*.

5. Data parity error detected on PCI-X interface fails to propagate bad parity

Problem: In PCI and PCI-X mode using 32-bit data transfers, when a read request gets disconnected at an even dword boundary with data parity error, such that the subsequent request for partial data gets retried, the completion for this request is issued over PCI Express to the MCH (root complex) without the poisoned data EP field set in the PCI Express TLP header.

Implication: Corrupted Data forwarded without error indication if error escalation is not enabled.

Workaround: Uncorrectable error escalation must be enabled in the MCH and Intel 6702PXH 64-bit PCI Hub to contain this data parity escape. Therefore, a complete workaround for this Erratum will also include MCH/root complex specific BIOS updates. Please refer to the latest version of the *Intel® 6702PXH 64-bit PCI Hub BIOS Specification Update* for details on this workaround.

Note: If error escalation is not properly enabled in the MCH/root complex, there would be a potential race condition between the completion being returned to the CPU and the error escalation. Read completion may complete normally at the CPU followed by error escalation as an Interrupt/SERR.

Status: For the steppings affected, see the *Summary Table of Changes*.

6. Intel 6702PXH 64-bit PCI Hub fails to train down in presence of degraded lane

Problem: During the PCI Express training sequence, if a broken endpoint has correct receiver termination on a lane and transmits training sequences on the lane which are invalid, the Intel 6702PXH 64-bit PCI Hub will fail to link train.

Implication: The PCI Express specification intends that, if some lanes are transmitting invalid data instead of valid training sequences, those lanes should be treated as broken, and the link should fail down to an acceptable width, such as x1. If Lane 0 were failing in this manner, the PCI Express specification would anticipate that the link would fail to train. If a higher-numbered lane were failing in this manner, the PCI Express specification requires that the link attempt to train as a x1 on lane 0. In either case, Intel 6702PXH 64-bit PCI Hub will not train for the problem scenario.

On production material, failures are anticipated to be either a broken transmitter path or a broken receiver path, or a silent transmitter. Intel 6702PXH 64-bit PCI Hub will train properly for these failure modes, since either the receiver termination will be missing, or the transmitted signals will not be seen at the Intel 6702PXH 64-bit PCI Hub. In order to see invalid transmitted data on lanes at the Intel 6702PXH 64-bit PCI Hub, either a logic bug in the other PCI Express endpoint would be required, or a signal integrity issue so sever as to make operation impossible, such as a broken or intermittent connection.

Workaround: None; A non-compliant or broken device could exhibit this erratum.

Status: For the steppings affected, see the *Summary Table of Changes*.

7. PCI Express and PCI-X header logs and First Error Pointers do not remain sticky through reset

Problem: The PCI Express and PCI-X header logs and First Error pointers are not maintaining their values after a warm/hot reset. These registers should be unaffected by a warm/hot reset, but instead, they are reset to default values. The following registers with “sticky” bits are affected: ADVERR_CTRL (offset 118h), EXP_TXNHDLOG (offset 11Ch), UNC_PXERRPTR (offset 138h), PX_TXNHDLOG (offset 13Ch), PX_DERRLOG (offset 14Ch), and PX_MISCERRLOG (offset 154h).

Implication: Errors detected will be logged and escalated properly, but after a warm/hot reset, the header logs and first error pointers will reset to their default values.

Note: Error status registers are unaffected, and properly maintain their values through reset.

Workaround: None at this time.

Status: For the steppings affected, see the *Summary Table of Changes*.

8. Incorrect default value for PCI Express Flow Control Protocol Error Severity bit

Problem: The PCI Express Flow Control Error Severity bit, register offset 10C, bit 13, is programmed to a default value of 0, indicating an uncorrectable flow control error will be reported as non-fatal. This is in contradiction with the *PCI Express 1.0a Specification*, which requires a default value of 1, indicating an uncorrectable flow control error will be reported as fatal.

Implication: Implications for this erratum depend upon the error response strategy implemented in a specific system.

Workaround: This bit can be programmed to match the specified default value if desired. Please refer to the latest version of the *Intel® 6702PXH 64-bit PCI Hub BIOS Specification Update* for details on this workaround.

Status: For the steppings affected, see the *Summary Table of Changes*.

9. Power State bits in PCI Express Power Management Status and Control Register mistakenly accept reserved values

Problem: The Power State bits, bits 1:0 of EXP_PMSTSCNTL (offset 70h) will allow a reserved value of 01b or 10b to be written. This is contrary to the specification, which originally stated that if software attempted to write an unsupported reserved state to this field, the data would be discarded and no state change would occur.

Implication: If a reserved state is written to this field, there will be a mismatch between the actual power state of the part and the state reported in configuration space. In some cases, writing a reserved value to this field could cause the Intel 6702PXH 64-bit PCI Hub to transition to the D0 power state, regardless of the previous power state.

Workaround: Never write a reserved value to this bit field.

Status: For the steppings affected, see the *Summary Table of Changes*.

10. Performance across an upstream x1 PCI Express link is less than expected

Problem: When the PXH is configured with an upstream x1 PCI Express link, the realized performance is significantly less than the predicted linear assumption that a x1 link will provide 1/4 the performance of a x4 link. This is caused by circumstances where PXH must discard a large portion of the data it receives across the upstream link. Notably; anytime PXH services an incorrect prefetch, or anytime PXH services interleaved requests from multi-function devices, PXH must discard data.

Implication: Devices that rely heavily on prefetching, or multi-function devices that request data in an interleaved fashion are the most likely to experience degraded performance.

Workaround: System designers should reduce the amount of prefetching allowed to devices behind PXH if possible.

Status: For the steppings affected, see the *Summary Table of Changes*.

11. SKP ordered set may not be sent within required interval during Link Recovery if a packet is pending

Problem: During Link Recovery on the PCI Express port, the PXH may fail to transmit a SKP ordered set within the required time interval as defined in the *PCI Express 1.0a Specification* if a TLP or DLLP was pending when the link entered Recovery.Idle state.

Implication: If the receiving device depends upon receipt of a SKP ordered set to progress through Link Recovery, a timeout will occur, resulting in Link Down and automatic reinitialization of the PCI Express Link. A link transitions through Recovery only under exceptional operational conditions. Following the Link Recovery timeout and reinitialization, the PCI Express link should resume normal operation unless the original Link Recovery condition was entered as a result of a hard failure mechanism.

Workaround: None at this time.

Status: For the steppings affected, see the *Summary Table of Changes*.

12. SERR Fatal/Non-Fatal Error message enabled with incorrect Error Message Enabled bit

Problem: If SERR# Enable (SEE) bit in PCI Command Register (offset 0x4) is not set then In the Advanced Error Reporting scheme, when SERR is configured as a fatal error, the generation of a SERR fatal error message is mistakenly gated by the Non-Fatal Error Reporting Enabled bit (bit 1) instead of the Fatal Error Reporting enabled bit (bit 2) of the Device Control Register, offset 4Ch. Likewise, when SERR is configured as a non-fatal error, the generation of a SERR non-fatal error message is gated by the Fatal Error Reporting enabled bit.

Implication: SERR fatal error message may only be generated if non-fatal error messaging is enabled, and vice versa.

Workaround: Set both Fatal and Non-Fatal Error Reporting Enabled bits, and mask errors individually via Uncorrectable PCI-X Error Mask Register (offset 130h) when escalation of these errors is not desired.

Status: For the steppings affected, see the *Summary Table of Changes*.

13. Configuration write to offset x70 of D0:F0, F2 (EXP_PMSTSCNTL - PCI Express Power Management Status and Control Register) using non-continuous byte enables does not capture the data value

Problem: A configuration write to single byte offset x71 of D0:F0, F2 does not capture the data value written to it. Specifically, performing this configuration write to offset x70 with byte enables of x2, x6 or xA in order to write to offset x71 does not work.

Implication: Register offset x70 is the PCI Express Power Management Status and Control Register (EXP_PMSTSCNTL). The byte at location x71 contains the PME enable bit (bit 8 of x70) and PME status bit (bit 15 of x70). The above mechanism will not work to either set or clear the PME enable bit, or to clear the PME status bit.

Workaround: Write to offset x71 by performing a word (byte enable x3) or a dword (byte enable xF) to offset x70.

Status: For the steppings affected, see the *Summary Table of Changes*.

14. PXH may become unresponsive when transitioning into the D3 power state

Problem: When the PXH is transitioned to a power state of D3 or lower, the PXH device may become unresponsive.

Implication: There have been no observed failures on systems with currently available software. Operating systems which independently manage the power state of the PXH, outside the scope of system level power state transitions, may result in the loss of link communications to the MCH.

Workaround: Independent device power state management of the PXH should be avoided. Should the PXH become unresponsive, a fundamental device reset must be asserted to return the system to normal operation.

Status: For the steppings affected, see the *Summary Table of Changes*.

15. MSI Hot-Plug Interrupt issue

Problem: An MSI is generated by the standard hot-plug controller may get corrupted in presence of another ACPI hot-plug driver. The ACPI driver performs configuration reads of DWSEL/DWORD register in order to determine the hot-plug capability of all the ACPI devices. If the MSI is generated by the Standard Hot-Plug Controller (SHPC) in this time period, there is a possibility of the MSI getting corrupted. As a result the MSI may not get issued upstream to the MCH. The above is a result of interaction of separate events that are unpredictable.

Implication: With the above condition described, the hot-plug device may not get recognized by the OS. Currently this issue is susceptible to only MSI aware systems.

Workaround: Disable PXH SHPC MSI, refer to the *Intel® 6700PXH 64-bit PCI Hub and Intel® 6702PXH 64-bit PCI Hub BIOS Specification* for further description.

Status: For the steppings affected, see the *Summary Table of Changes*.

16. Tcyc(min) of the PXH output clocks deviates from the PCI-X (Mode 1, class1) jitter clock specification

Problem: PXH generates the PCI-X 133 MHz clock with a nominal frequency of 133 MHz (Tcyc of 7.5 ns). After considering the clock jitter, Tcyc(min) (minimum clock period) observed at the pin may be marginally less than 7.5 ns. The PCI-X class 1 jitter clock specification in Mode 1 requires the Tcyc(min) to be 7.5 ns with jitter consideration. Tcyc(min) deviation may also exists on PCI-X clock generated @ 100/66 MHz and PCI @ 66/33 MHz.

Implication: None at this time.

Workaround: There is no workaround to adjust the Tcyc(min) of the PCI-X clocks. However, the routing guidelines for the PXH output clock signals take into consideration the effect of the jitter on the Tcyc(min). Conforming to the guidelines published by Intel for the platforms that involve Intel 6700PXH/6702PXH 64-bit PCI Hub will offset the effect of the marginally reduced Tcyc(min) towards the set-up and hold times. Hence, in system boards that are compliant with the routing guidelines, the risk of deviating from the set-up and hold time requirements and any resulting functional impact, is low. Refer to the appropriate Platform Design Guides for further information on the PCI-X clock routing guidelines.

Status: For the steppings affected, see the *Summary Table of Changes*.

17. Under certain conditions, inbound prefetched PCI read requests may return wrong data to the requestor

Problem: With some prefetch policy settings, the PXH may over-aggressively prefetch data for PCI reads and subsequently return the wrong data to the requestor. This problem only exists when there is more than one active agent on the PCI bus. This problem exists for all supported frequencies. This problem does not affect PCI-X operation at any supported frequency.

Implication: Inbound read requests that are enabled for prefetching may return invalid data when multiple agents exist on the same PCI bus. No error is reported by the PXH.

Workaround: The problem may be corrected via a SW workaround. The workaround corrects the problem at all supported frequencies and all prefetch policy settings. Please refer to the latest version of the *Intel® 6700PXH 64-bit PCI Hub and Intel® 6702PXH 64-bit PCI Hub BIOS Specification* for details on the workaround.

Status: For the steppings affected, see the *Summary Table of Changes*.

Specification Changes

There are no new Specification Changes for this revision of the *Intel® 6702PXH 64-bit PCI Hub Specification Update*.

1. PCI Express L0s functionality not supported in Intel 6702PXH 64-bit PCI Hub

The L0s functionality will not achieve optimum performance due to limitations in the circuit design as described in erratum #2. Because of this, the L0s functionality will not be supported in the Intel 6702PXH 64-bit PCI Hub. The BIOS should be updated to leave L0s disabled on Intel 6702PXH 64-bit PCI Hub (by default), and eliminate any setup option with which to enable it.

2. CMODE Power-on Strap for PCI Express 1.0a support

This pin strap is used to configure PCI Express 1.0a support and is muxed with HAATNLED_1#. The CMODE/HAATNLED_1# pin does not have ODT (On-Die Termination), need an external resistor to pull up or pull down.

Specification Clarifications

There are no new Specification Clarifications for this revision of the *Intel® 6702PXH 64-bit PCI Hub Specification Update*.

1. The relationship between PWROK and the PCI Express clocks timing specification clarification

The PWROK signal is used to indicate when the power supply is within its specified voltage tolerance and is stable. It also initializes the Intel 6702PXH 64-bit PCI Hub's state machines and other logic once power supplies stabilize. On power up, the assertion of PWROK is delayed 100 ms (T_{PVPERL}) from the power rails achieving specified operating limits. Also, within this time, the PCI Express clocks should be stable at least 100 μ s before PWROK is asserted.

2. Table 4-20 will be updated to include the following information in columns PCI-X66, Conv. PCI 66 (ref), and Conv. PCI 33 (ref)

Sym	Parameter	PCI-X 133		PCI-X 66		Conv. PCI 66 (ref)		Conv. PCI 33 (ref)		Unit
		Min	Max	Min	Max	Min	Max	Min	Max	
Clock Jitter Class 1										
T_{cyc}	CLK Cycle Time	7.5	20	15	20	15	30	30		ns
T_{high}	CLK High Time	3		6		6		11		ns
T_{low}	CLK Low Time	3		6		6		11		ns
Clock Jitter Class 2										
T_{cyc}	CLK Cycle Time	Average	7.5	20	15	20	15	30	30	ns
		Absolute minimum	7.375		14.8		14.8		29.7	ns
T_{high}	CLK High Time	2.5		5.5		5.5		10		ns
T_{low}	CLK Low Time	2.5		5.5		5.5		10		ns
Slew Rate										
-	CLK Slew Rate	1.5	4	1.5	4	1.5	4	1	4	V/ns
Spread Spectrum Requirements										
f_{mod}	Modulation Frequency	30	33	30	33	30	33			kHz
f_{spread}	Frequency Spread	-1	0	-1	0	-1	0			%

Documentation Changes

There are no new Documentation Changes for this revision of the *Intel® 6702PXH 64-bit PCI Hub Specification Update*.

1. Updated Section 2.7, Table 2-9

Replaced “*This signal can be used to gate the starting of the PCI Express* clocks, since it is required that all voltages supplied to the Intel® 6700PXH 64-bit PCI Hub be valid and stable prior to starting the PCI Express* clocks.*” with “*Please refer to section 2.17.2.1 for details on the timing relationship between PWROK, the PCI Express* clocks, and all voltages supplied to the Intel® 6700PXH 64-bit PCI Hub.*”

2. Updated Section 2.17.2.1

At the beginning of the third paragraph, added “*The PWROK signal is used to indicate when the power supply is within its specified voltage tolerance and is stable. It also initializes the Intel® 6700PXH 64-bit PCI Hub’s state machines and other logic once power supplies stabilize. On power up, the assertion of PWROK is delayed 100ms (T_{PVPERL}) from the power rails achieving specified operating limits. Also, within this time, the reference clocks (PCI Express* clocks) become stable at least $T_{PWROK-CLK}$ (100 μ s) before PWROK is asserted.*”

3. In Section 3.8.1.6, added C1 stepping, removed all other stepping and changed default value to I/OxAPIC REVID register (Offset 08h)

3.8.1.6 Offset 08h: REVID—Revision ID Register (D0: F1, F3)

Offset:	08h	Attribute:	RO
Default Value:	09h	Size:	8 bits

Identifies the I/OxAPIC stepping of the Intel 6700PXH 64-bit PCI Hub.

Bits	Type	Reset	Description
7:0	RO	0	Revision ID (RID): Indicates the step of the I/OxAPIC of the Intel® 6700PXH 64-bit PCI Hub. <i>09h = C1 stepping.</i>

4. In Section 3.4.1.39, updated EXP_LCAP (Offset 50h) bit 14:12 description and register’s default value

Offset:	50 – 53h	Attribute:	RO
Default Value:	0003F481h	Size:	32 bits

Bits	Type	Reset	Description
14:12	RO	111b	L0s Exit Latency (L0EL): The value in these bits is influenced by bit 6 in the link control register. Note that software could write the bit 6 in link control register to either a 1 or 0 and these bits should change accordingly. The mapping is shown below: Bit 6 PCI Express* Link Control Link Capabilities Bits 14:12 0 111b = <i>more than 4 ns</i> 1 010b = 128 ns to less than 256 ns

5. In Section 3.8.1.6, added C1 Stepping to I/OxAPIC REVID register (Offset 08h)

Bits	Type	Reset	Description
7:0	RO	0	Revision ID (RID): Indicates the step of the I/OxAPIC of the Intel® 6700PXH 64-bit PCI Hub. 00h = A0 stepping. 04h = B0 stepping. 08h = C0 stepping. 09h = C1 stepping.

6. In Section 3.5.1.39, updated EXP_LCAP (Offset 50h) bit 14:12 description and register's default value

Offset: 50 – 53h Attribute: RO
 Default Value: **0003F481h** Size: 32 bits

Bits	Type	Reset	Description
14:12	RO	111b	L0s Exit Latency (L0EL): The value in these bits is influenced by bit 6 in the link control register. Note that software could write the bit 6 in link control register to either a 1 or 0 and these bits should change accordingly. The mapping is shown below: Bit 6 PCI Express* Link Control Link Capabilities Bits 14:12 0 111b = more than 4 us 1 010b = 128 ns to less than 256 ns

7. In Section 3.5.1.32, updated PCLKC register (Offset 43h) description

Offset: 43h Attribute: RW, RO
 Default Value: FFh Size: 8 bits

Bits	Type	Reset	Description
7	RO	1	Reserved
6:0	RW	1111111	PCI Clock Control (PCLKC): These bits enable the PCI clock output buffer to toggle, when 1. Otherwise the buffers are driven to a stable DC value. Corresponds to PCLKO[6], bit 5 for PCICLK[5] etc.

8. In Section 3.5.1.41, documentation error in EXP_LSTS (offset 56h) bit field description

In the Datasheet, the bits 15:10 have been indicated as:

Bits	Type	Reset	Description
15:10	RO	0	Reserved

This should change to:

Bits	Type	Reset	Description
15:13	RO	0	Reserved

Bits	Type	Reset	Description
12	RWO	0	Slot Clock Configuration – This bit indicates whether a PXH that is on a 3GIO connector uses the same reference clock as is provided at the connector or an independent reference clock. A value of 0 indicates independent reference clock and a value of 1 indicates same reference clock. This bit is initialized by BIOS.
11:10	RO	0	Reserved

9. **Sections 2.19.2.2 (Error Logging) and 2.19.2.3 (Error Escalation) are replaced by a combined PCI error section (Section 2.19.2.2, “Error Logging and Escalation”) with the following description**

2.19.2.2 Error Logging and Escalation

PXH supports the PCI Express-advanced error logging capability with bridge extensions and this capability is located in the enhanced PCI Express configuration space. All errors on PCI Express and PCI/X are classified as either correctable or uncorrectable fatal or uncorrectable non-fatal. Refer to PCI Express and bridge specs for details of the classification. Each of the uncorrectable errors is individually tracked with separate status bits in the capability and can be programmed to generate either a ERR_FATAL or ERR_NONFATAL message on PCI Express. Also the first uncorrectable error on each interface (which has its corresponding mask bit cleared) causes a header log and subsequent errors cause only a status bit to be set. Once a header is logged on a given interface, further header logging on that interface is reenabled only after the error status bit corresponding to the first error is cleared by software. On the PCI interface, PXH also logs the data on the transaction cycle on which a data parity/ECC error was detected.

2.19.2.2.1 Bridge and APIC Error Escalation Rules on PCI Express

All uncorrectable errors detected by PXH on PCI Express, PCI/X and internally are escalated to ERR_FATAL or ERR_NONFATAL message on PCI Express only if the corresponding mask bit for that error is cleared in the advanced error register and the ERR_FATAL or ERR_NONFATAL message is enabled. ERR_FATAL and ERR_NONFATAL messages are both enabled when the SERR enable bit in the PCICMD register is set. ERR_FATAL/NON_FATAL messages are also enabled when the corresponding enable bits in the PCI Express capability structure are set.

The following additional rules apply for all bridge error escalation:

1. SERR assertion forwarding from PCI/X to PCI Express is enabled either if the SERR enable bit in the bridge control register is set or the mask bit for the SERR detected, in the advanced error reporting register is cleared
2. Posted Write Master aborts on PCI/X can be escalated on to PCI Express either if the Master Abort Mode bit in the Bridge control register is set. (i.e. if the master abort mode bit is clear, posted write master aborts do not generate any PCI Express error message) or the “PCI/X detected master abort” mask bit in the advanced capability register is clear
3. Posted Write Target Aborts and Posted Write Data Parity errors on PCI/X when PXH is master on PCI/X bus, are escalated to PCI Express if the SERR enable bit is set, even if the advanced error mask bit is set
4. Secondary discard timer expiry will generate an error message on PCI Express either if the “SERR enable on discard timer expiry” bit in the bridge control register is set or the corresponding mask bit in the PCI Express advanced error capability register is clear

5. Uncorrectable address parity errors (PCI/X interface only and not PCI Express) generate error messages on PCI Express either if the parity error response enable bit in the BCTL register is set (or) the appropriate mask bit in the PCI Express advanced error capability register is clear
6. The signaled system error bit is set in the PCISTS register anytime an ERR_FATAL or ERR_NONFATAL message is signaled with the SERR enable bit is set, with the following exceptions:
7. Correctable ECC errors on PCIX Mode2 generate an ERR_COR message on PCI Express (if the ERR_CORR enable bit in the PCI Express device control register is set) and also set the ‘Detected Correctable Error’ Status bit in the PCI Express capability.
8. All uncorrectable data errors forwarded from PCI Express or from internal RAMs to PCIX would cause the appropriate status bits to be set in the Sec.STS (master data parity error bit for writes/delayed PCI reads) and PCIXERRUNC_STS (PCI/X PERR# detected bit) registers (when PERR# is observed on PCI bus), and also cause the appropriate header to be logged on the PCI/X side of the advanced error capability (if enabled per the rules for header logging), but would not cause any error messages to be sent on PCI Express. Error message is only generated on PCIE on observance of PERR# asserted on PCI, if it was a bus error rather than a forwarded error from PXH. Note that for error logging and escalation for errors from the PCIE interface follow the rules per the PCI Express specification 1.0a.

In all cases, PXH uses the severity bits in the advanced error capability to escalate between FATAL and NONFATAL message.

10. Section 1.2.1 (page 12), the documentation incorrectly includes “lane reversal” that PXH does not support

“The Intel® 6700PXH 64-bit PCI Hub also supports X8 lane reversal, plus X4 lane reversal on the lower 4 lanes only” will be removed.

11. Figures 5-1 will be updated. The following pins in the figures need to match the pin list:

PIN	Change from:	Change to:
C18	RESERVED[3]	PASTRAP0
A19	RESERVED[2]	PBSTRAP0
M20	PAVIOSEL	RESERVED
M8	PAVIOSEL	RESERVED
D19	RESERVED	RESERVED
D18	RESERVED[0]	RESERVED
C19	PASTRAP0	RESERVED
B19	RESERVED[1]	RESERVED
U11	PBSTRAP0	RESERVED
C6	TDIOANODE	RESERVED
B6	TDIOCATHODE	RESERVED

12. Table 2-2 (page 17) will be updated

PAPCIRST#/PBPCIRST# will be updated to add “Only used in non Hot Plug mode” at the end of the signal description.

13. Table 2-7 (page 23) will be updated

HPA_PRST#/HPA_RST1#/ HPB_PRST#/HPB_RST1# description will change from “one-slot-not-glue” to “one-slot-no-glue”.

14. Table 2-7 (page 22) will be updated

The last line in the pin PAIRQ_[9]#/HPACIXCAP2_1 and PBIRQ_[9]#/HPBCIXCAP2_1 description should read: *These pins are used only in single-slot parallel or dual-slot parallel hot plug mode (HPx_SLOT[3:0] =1001 or 1010).*

15. Table 2-7 (page 21) will be updated

The last line in the pin PAIRQ_[10]#/HPACIXCAP1_1 and PBIRQ_[10]#/HPBCIXCAP1_1 description should read: *These pins are used only in single-slot parallel or dual-slot parallel hot plug mode (HPx_SLOT[3:0] =1001 or 1010).*

16. Section 2.13.6.1 (page 50-51) will be updated

References to PxPCIRST# and PCIRST# will be replaced with HxPRST#.