



Intel® 80333 I/O Processor

Specification Update

| *April 2006*

The Intel® 80333 I/O Processor 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.

Order Number: 305435-011US



INFORMATION IN THIS DOCUMENT IS PROVIDED IN CONNECTION WITH INTEL® PRODUCTS. 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 Corporation may have patents or pending patent applications, trademarks, copyrights, or other intellectual property rights that relate to the presented subject matter. The furnishing of documents and other materials and information does not provide any license, express or implied, by estoppel or otherwise, to any such patents, trademarks, copyrights, or other intellectual property rights.

Intel products are not intended for use in medical, life saving, life sustaining, critical control or safety systems, or in nuclear facility applications.

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

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

Copies of documents which have an ordering number and are referenced in this document, or other Intel literature may be obtained by calling 1-800-548-4725 or by visiting Intel's website at <http://www.intel.com>.

AnyPoint, AppChoice, BoardWatch, BunnyPeople, CablePort, Celeron, Chips, CT Media, Dialogic, DM3, EtherExpress, ETOX, FlashFile, i386, i486, i960, iCOMP, InstantIP, Intel, Intel Centrino, Intel logo, Intel386, Intel486, Intel740, IntelDX2, IntelDX4, IntelSX2, Intel Create & Share, Intel GigaBlade, Intel InBusiness, Intel Inside, Intel Inside logo, Intel NetBurst, Intel NetMerge, Intel NetStructure, Intel Play, Intel Play logo, Intel SingleDriver, Intel SpeedStep, Intel StrataFlash, Intel TeamStation, Intel Xeon, Intel XScale, IPLink, Itanium, MCS, MMX, MMX logo, Optimizer logo, OverDrive, Paragon, PC Dads, PC Parents, PDCharm, Pentium, Pentium II Xeon, Pentium III Xeon, Performance at Your Command, RemoteExpress, SmartDie, Solutions960, Sound Mark, StorageExpress, The Computer Inside., The Journey Inside, TokenExpress, VoiceBrick, VTune, and Xircom 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 © 2006, Intel Corporation. All rights reserved.

Contents

Revision History	5
Preface	6
Summary Table of Changes	7
Identification Information	13
Non-Core Errata	14
Core Errata	27
Specification Changes	31
Specification Clarifications	33
Documentation Changes	48

THIS PAGE INTENTIONALLY LEFT BLANK

Revision History

Date	Version	Description
April 2006	011	Added: - A-1 stepping information to all Summary tables and to Table 1, "80333 Die Details" on page 13 and Table 2, "80333 Device ID Registers" on page 13 - Status change of Erratum 35 - Status change of Specification Change 6 - Status change of Specification Clarification 52 - Documentation Change 4
February 2006	010	Added: - Non-Core Errata 36 - Specification Clarification 52
January 2006	009	Added: - Non-Core Errata 35 - Specification Change 6
December 2005	008	Added: - Revised Specification Clarification 9 - Specification Clarification 49, 50, and 51 - Documentation Change 3
October 2005	007	Added Specification Clarification 47 and 48
August 2005	006	Added: - Non-Core Errata 33 and 34 - Specification Clarification 45 and 46 - Documentation Change 2
June 2005	005	Added: - Non-Core Errata 32 - Specification Change 5
May 2005	004	Added: - Non-Core Errata 31 - Specification Change 3 and 4
April 2005	003	Added: - Non-Core Errata 30 - Specification Clarification 44
March 2005	002	Added: - Specification Change 1 and 2 - Documentation Change 1 - PB-free markings in Table 1
February 2005	001	Initial release

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.

Affected Documents/Related Documents

Title	Order
Intel® 80333 I/O Processor Developer's Manual	305432
Intel® 80333 I/O Processor Datasheet	305433
Intel® 80333 I/O Processor Design Guide	305434
Intel® 6700PXH 64-bit PCI Hub Specification Update	302706

Nomenclature

Errata are design defects or errors. These may cause the behavior of the Intel® 80333 I/O Processor¹ (80333) 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 on 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.

Note: Errata remain in the specification update throughout the product's life cycle, or until a particular stepping is no longer commercially available. Under these circumstances, errata removed from the specification update are archived and available upon request. Specification changes, specification clarifications and documentation changes are removed from the specification update when the appropriate changes are made to the appropriate product specification or user documentation (datasheets, manuals, etc.).

1. ARM architecture compliant.

Summary Table of Changes

The following table indicates the errata, specification changes, specification clarifications, or documentation changes which apply to the Intel® 80333 I/O Processor. 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.

Page

(Page):	Page location of item in this document.
---------	---

Status

Doc:	Document change or update will be implemented.
Fixed:	This erratum has been previously fixed.
No Fix:	There are no plans to fix this erratum.
Plan Fix:	This erratum may be fixed in a future stepping of the product.

Row

	Change bar to left of table row indicates this erratum is either new or modified from the previous version of the document.
--	---

Non-Core Errata (Sheet 1 of 2)

No.	Stepping		Page	Status	Errata
	A-0	A-1			
1	X	X	14	No Fix	CAS latency of three not supported for DDR-II On-Die Termination (ODT)
2	X	X	14	No Fix	Legacy power fail mechanism does not work
3	X	X	14	No Fix	A_REQ64# and B_REQ64# initialization pattern timing violation in PCI-33 mode
4	X	X	15	No Fix	Secondary Bus Number register (PEBSBBNR) provides incorrect bus number
5	X	X	15	No Fix	Boundary scan multi-chip module implementation
6	X	X	15	No Fix	PCI Express* traffic class (TC) bit[2] ignored for malformed packet checks
7	X	X	16	No Fix	Auto-Refresh command also generates a Precharge All command on DDR bus
8	X	X	16	No Fix	Coalesced writes to 32-bit memory can cause data corruption
9	X	X	17	No Fix	ATU passing rules operation in PCI mode
10	X	X	17	No Fix	Secondary bus PCI RST# pulse prior to the rising edge of PWRGD
11	X	X	18	No Fix	VPD Data Register bit[19] is not read/write
12	X	X	18	No Fix	PCI Express* Correctable Error Mask Bits
13	X	X	18	No Fix	DMA CRC result is byte-reversed
14	X	X	18	No Fix	CRC corruption on PCI-to-local DMA transfers
15	X	X	19	No Fix	IOAPIC End of Interrupt (EOI) Register is Read-Write, should be Write-Only
16	X	X	19	No Fix	Unreliable PCI Express* link operation when L0s active state power management is enabled
17	X	X	19	No Fix	SSE bit set for PERR# assertion when error reporting is masked
18	X	X	19	No Fix	Data Parity Error detected on PCI/X interface fails to propagate bad parity
19	X	X	20	No Fix	ATU claims PCI commands 8 and 9 when issued as Dual Address Cycle (DAC)
20	X	X	20	No Fix	Failure to train down in presence of degraded lane
21	X	X	21	No Fix	PCI Express* and PCI-X header logs and first-error pointers do not remain sticky through reset
22	X	X	21	No Fix	Incorrect default value for PCI Express* Flow Control Protocol Error Severity bit
23	X	X	21	No Fix	Power State bits in PCI Express* Power Management Status and Control Register mistakenly accept reserved values
24	X	X	22	No Fix	Performance across an upstream x1 PCI Express* link is less than expected
25	X	X	22	No Fix	PCI Express* ESD enhancement requires a change to register setting
26	X	X	22	No Fix	SKP ordered set might not be sent within required interval during link recovery when a packet is pending
27	X	X	23	No Fix	SERR fatal/non-fatal error message enabled with incorrect error message enabled bit
28	X	X	23	No Fix	Configuration write to offset 70h of A- and B-bridge (PM_CSR - PCI Express Power Management Control/Status Register) using non-continuous byte enables does not capture the data value
29	X	X	23	No Fix	The 80333 might become unresponsive when transitioning into the D3 power state
30	X	X	24	No Fix	Bus Interface Unit (BIU) claims DAC addresses in the range of the Memory Mapped Registers (MMR)

Non-Core Errata (Sheet 2 of 2)

No.	Stepping		Page	Status	Errata
	A-0	A-1			
31	X	X	24	No Fix	Tc1(min) of the PCI-X clock observed to be marginally less than the requirement specified for the PCI-X (Mode 1, class1) clock jitter.
32	X	X	25	No Fix	I2C Control Register reset bit does not function
33	X	X	25	No Fix	MSI Hot-Plug Interrupt issue
34	X	X	25	No Fix	Under certain conditions, inbound prefetched PCI read requests may return wrong data to the requestor
35	X		26	Fixed	Internal Clock Misalignment Can Cause Processor Hang.
36	X	X	26	No Fix	Spurious DMA0 End-Of-Transfer Interrupt

Core Errata

No.	Steppings		Page	Status	Errata
	A-0	A-1			
1	X	X	27	No Fix	Abort is missed when lock command is outstanding
2	X	X	27	No Fix	Aborted store that hits the data cache can mark write-back data as dirty
3	X	X	28	No Fix	Performance Monitor Unit event 0x1 can be incremented erroneously by unrelated events
4	X	X	28	No Fix	In Special Debug State, back-to-back memory operations—in which the first instruction aborts—can cause a hang
5	X	X	29	No Fix	Accesses to the CP15 ID Register with opcode2 > 001b returns unpredictable values
6	X	X	29	No Fix	Disabling and re-enabling the MMU can hang the core or cause it to execute the wrong code
7	X	X	30	No Fix	Updating the JTAG parallel registers requires an extra TCK rising edge
8	X	X	30	No Fix	Non-branch instruction in vector table might execute twice after a thumb mode exception

Specification Changes

No.	Steppings		Page	Status	Specification Changes
	A-0	A-1			
1	X	X	31	Doc	Recommended DLL register values
2	X	X	31	Doc	DDR-II JEDEC initialization sequence includes writes to EMRS2 and EMRS3
3	X	X	32	Doc	REFCLK relationship to voltage rails
4	X	X	32	Doc	Case temperature (Tcase) change
5	X	X	32	Doc	PWRGD and RSTIN# sequencing
6	X		32	Fixed	Internal Clock Misalignment.

Specification Clarifications (Sheet 1 of 2)

No.	Steppings		Page	Status	Specification Clarifications
	A-0	A-1			
1	X	X	33	No Fix	PCI Express* to PCI-X Bridge does not support Device ID Messaging (DIM)
2	X	X	33	No Fix	64 MB and 2 GB DDR333 capacities not tested in post-silicon validation
3	X	X	33	No Fix	DDR-II 400 unbuffered DIMMs are not supported
4	X	X	33	No Fix	Memory map for 2 GByte of DDR memory
5	X	X	33	No Fix	PCI configuration write anomaly when clearing BINIT[1]
6	X	X	33	No Fix	PWRGD and PERST# are the same signals
7	X	X	34	No Fix	Back-to-back MCU MMR reads
8	X	X	34	No Fix	Reserved IDSELs on A-segment
9	X	X	34	No Fix	Retry Disable Sequence
10	X	X	35	No Fix	Write requirements for the Peripheral Bus Interface
11	X	X	35	No Fix	PCI-X Status Register during PCI mode
12	X	X	35	No Fix	M_RST# driven to DDR-II or DDR-I voltage levels
13	X	X	35	No Fix	BIU master abort causes two interrupts on reads
14	X	X	36	No Fix	Potential race condition with Interrupt Controller Unit status bits
15	X	X	36	No Fix	Reset Internal Bus (PCSR[5]) usage
16	X	X	36	No Fix	PCI Express* Transaction Header Log register repeats on offset 124 and 128
17	X	X	37	No Fix	SHPC sequence
18	X	X	38	No Fix	Bus Interface Unit follows PCI ordering rules
19	X	X	39	No Fix	UART, I2C and GPIO memory mapped registers should be addressed with 32-bit accesses
20	X	X	39	No Fix	UART Interrupt Identification Register
21	X	X	39	No Fix	Reads on 16-bit PBI bus operate as 32-bit
22	X	X	39	No Fix	3.3 V to 1.5 V leakage
23	X	X	40	No Fix	Accessing extended bridge configuration space from the Intel XScale® processor
24	X	X	40	No Fix	Recommended B-segment termination when using the 80333 on PCI Express* adapter cards
25	X	X	40	No Fix	B-segment arbiter might not park on the last master when in PCI-X 133 MHz mode
26	X	X	40	No Fix	Configuring XINTx# signals as PCI interrupts
27	X	X	41	No Fix	Power plane isolation for Battery Back-Up (BBU) mode
28	X	X	41	No Fix	AAU result can be written directly to PCI host memory
29	X	X	41	No Fix	SMBus connection recommendations for PCI Express* adapter cards
30	X	X	42	No Fix	PBI lockout condition
31	X	X	42	No Fix	PFREQ functionality
32	X	X	42	No Fix	PWRDELAY functionality during power sequencing
33	X	X	42	No Fix	HPI# (High Priority Interrupt) is a maskable interrupt

Specification Clarifications (Sheet 2 of 2)

No.	Steppings		Page	Status	Specification Clarifications
	A-0	A-1			
34	X	X	42	No Fix	OCD and Receive Enable calibration de-featured
35	X	X	43	No Fix	PWRDELAY needs only a pull-up for battery back-up mode
36	X	X	43	No Fix	PCI Express* L0s functionality not supported in the 80333
37	X	X	43	No Fix	DDRRES2 can be pulled down to reduce current during self-refresh
38	X	X	43	No Fix	DDRSLWCRES resistor values
39	X	X	43	No Fix	B_PME# routing recommendation when using Parallel Hot Plug 1-slot, no-glue mode
40	X	X	44	No Fix	Multi-Transaction Timer grants fewer clocks in PCI mode than expected
41	X	X	44	No Fix	Byte Enables (BE) not included in PCI delayed reads can cause data corruption
42	X	X	44	No Fix	Interleaving AAU descriptors
43	X	X	45	No Fix	RCVDLY setting for DDR-I memory
44	X	X	45	No Fix	Embedded Usage Models
45	X	X	45	No Fix	ATUBAR3 Functionality
46	X	X	45	No Fix	VREF isolation for Battery Back-up (BBU) mode
47	X	X	45	No Fix	I2C Unit Enabling
48	X	X	46	No Fix	DMA transactions from local memory to a conventional PCI target can complete out of order
49	X	X	46	No Fix	SBR1 Programming with Bank 1 Unpopulated
50	X	X	46	No Fix	32-bit Writes to Unaligned 64-bit Addresses are Promoted to 64-bit Aligned Writes
51	X	X	46	No Fix	ATU Retry Response Through the Bridge
52	X		47	Fixed	Case Temperature Clarification.



Documentation Changes

No.	Document Revision	Page	Status	Documentation Changes
1	305433-001	48	Doc	PCI clock timings table missing note
2	305433-002	48	Doc	Wrong Voltage Values in Table 23
3	305432-001	48	Doc	SBR1 Programming When Bank 1 is Unpopulated
4	305433-003	48	Doc	PCI Express clock cycle time minimum

Identification Information

Die Details

Table 1. 80333 Die Details

Stepping	Part Number	QDF (Q)/ Specification Number (SL)	Processor Speed (MHz)	Notes
A-0	NQ80333M500 QG80333M500	Q010 Q108	500	Engineering Samples PB-free Eng Samples
A-0	NQ80333M667 QG80333M667	Q011 Q109	667	Engineering Samples PB-free Eng Samples
A-0	NQ80333M800 QG80333M800	Q012 Q110	800	Engineering Samples PB-free Eng Samples
A-0	NQ80333M500 QG80333M500	SL82B SL8CC	500	Production PB-free Production
A-0	NQ80333M667 QG80333M667	SL82C SL8CD	667	Production PB-free Production
A-0	NQ80333M800 QG80333M800	SL82D SL8CE	800	Production PB-free Production
A-1	NQ80333M500 QG80333M500	Q611 Q614	500	Engineering Samples PB-free Eng Samples
A-1	NQ80333M667 QG80333M667	Q612 Q615	667	Engineering Samples PB-free Eng Samples
A-1	NQ80333M800 QG80333M800	Q613 Q616	800	Engineering Samples PB-free Eng Samples
A-1	NQ80333M500 QG80333M500	SL9BA SL9BH	500	Production PB-free Production
A-1	NQ80333M667 QG80333M667	SL9BB SL9BJ	667	Production PB-free Production
A-1	NQ80333M800 QG80333M800	SL9BC SL9BK	800	Production PB-free Production

Table 2. 80333 Device ID Registers

Device and Stepping	Processor Device ID (CP15, Register 0 - opcode_2=0)	Revision ID	JTAG Device ID
A-0	0x69054210	0x00	0x09268013
A-1	0x69054210	0x00	0x09268013

NOTE: Processor core speed can be identified by reading CCLKCFG[3:0] (CP14, register 6).
500 MHz = 0x0, 667 MHz = 0x2, 800 MHz = 0x4

NOTE: A-Bridge = 0370h, A-IOAPIC = 0371h, B-Bridge = 0372h, B-IOAPIC = 0373h, ATU = 0374h

Non-Core Errata

1. CAS latency of three not supported for DDR-II On-Die Termination (ODT)

Problem: For DDR-II memory with a CAS Latency (CL) of three, the memory controller unit (MCU) does not provide the proper timing for the On-Die Termination signals (ODT[1:0]). The *JEDEC DDR-II SDRAM Specification*, September 2002 states that ODT must be driven one cycle prior to the write command, but the MCU does not meet this timing.

Implication: CAS latency of three is not supported in the 80333; therefore, there is minimal performance impact as compared to CAS latency of four.

Workaround: Use CAS latency = 4 or do not use the ODT feature.

Status: **No Fix.** Not to be fixed. See the [Table , “Summary Table of Changes” on page 7.](#)

2. Legacy power fail mechanism does not work

Problem: For previous I/O processor generations, an external clock was required to maintain the incoming PCI clock (P_CLK) long enough for the power-fail sequence to be sent to the memory. This is what is referred to as “legacy power fail.” The internal control circuit that enables the legacy power-fail method is broken.

Implication: Legacy power-fail cannot be used.

Workaround: For the 80333, a new feature was added which keeps internal clocks running on power fail.

Status: **No Fix.** Not to be fixed. See the [Table , “Summary Table of Changes” on page 7.](#)

3. A_REQ64# and B_REQ64# initialization pattern timing violation in PCI-33 mode

Problem: The *PCI Local Bus Specification*, Revision 2.3 states RST# to REQ64# hold time is 0 ns minimum and 50 ns maximum (Trrh). The 80333 drives the REQ64# signal for three cycles after RST# deasserts. Therefore, in PCI-33 mode, this is 90 ns, which violates the 50 ns maximum.

All other PCI and PCI-X modes are within the *PCI Local Bus Specification*, Revision 2.3 (that is, PCI-66 is 45 ns).

Implication: No negative impact expected.

Workaround: There is no overlap between the 80333 assertion of REQ64# for the initialization sequence, and the first assertion of FRAME#. As per the *PCI Local Bus Specification*, Revision 2.3, Trhff defines the minimum time from RST# high to first FRAME# assertion as 5 clocks. Since the 80333 deasserts REQ64# at three clocks following RST# high, there is no overlap with FRAME# being asserted as it is not allowed to be asserted any sooner than five clocks.

Status: **No Fix.** Not to be fixed. See the [Table , “Summary Table of Changes” on page 7.](#)

4. Secondary Bus Number register (PEBSBBNR) provides incorrect bus number

Problem: The PCI Express*-to-PCI Bridge Secondary B-Segment Bus Number Register (PEBSBBNR) does not provide the B-segment bus number. PEBSBBNR is to mirror the Secondary Bus Number (SCBN) field of the BNUM register (bits[15:8] at offset 18h) in the B-segment configuration header space. Instead, PEBSBBNR provides the Subordinate Bus Number (SBBN) of the BNUM register (bits[23:16] at offset 18h) from the A-segment configuration header space.

Implication: The I/O processor does not support transactions from the Intel XScale® processor on A-segment to the B-segment.

Workaround: To successfully read PCI B-segment devices, software can scan the PCI B-segment configuration space using configuration reads to determine the bus number.

Status: **No Fix.** Not to be fixed. See the [Table](#) , “Summary Table of Changes” on page 7.

5. Boundary scan multi-chip module implementation

Problem: The 80333 is not BSDL-compliant for the SAMPLE and BYPASS instructions specified by the JTAG specification 1149.1. It is compliant for most board-level testing. When boards are tested for opens and shorts, the 80333 BSDL can define the boundary scan length as +1 to encompass the BYPASS register in the Intel XScale® processor (therefore not visible).

Implication: When doing ID, SAMPLE, or BYPASS, the Intel XScale® processor BYPASS register makes the path from TDI to TDO one flop longer than the JTAG specification 1149.1 requires, which can cause canned software to error.

Note: When neither of these are used by the board vendors during manufacturing testing, there is no issue.

Workaround: Intel can provide two BSDL files which allow opens and shorts testing, as long as it does not test the ID and BYPASS instructions. One BSDL file covers the Intel XScale® processor unit, and the other BSDL file covers the I/O processor, with the exception that both instruction sets are reduced from 14 to 7, since they are operating independently.

Status: **No Fix.** Not to be fixed. See the [Table](#) , “Summary Table of Changes” on page 7.

6. PCI Express* traffic class (TC) bit[2] ignored for malformed packet checks

Problem: *PCI Express* Specification*, Revision 1.0a Packet formation rules state that Message (Unlock, Slot Power Limit), I/O, CFG TLP must use Traffic Class 0, such that TC[2:0] of the header field is set to 000. Receivers that implement the optional Malformed Packet Check must check for TC usage. The 80333 checks only TC[1:0] and ignores bit[2].

Implication: Packet formation violations are not detected and flagged as Malformed Packets.

Workaround: None

Status: **No Fix.** See the [Table](#) , “Summary Table of Changes” on page 7.

This was originally written with the understanding that malformed packet checks using TC[2:0] were required. It is now understood that the *PCI Express* to PCI/PCI-X bridge Specification 1.0* requires PCI Express* bridges to ignore TC and forward all requests regardless of TC labeling. In order to comply with the *PCI Express* to PCI/PCI-X bridge Specification 1.0*, TC checks are to be removed.

7. Auto-Refresh command also generates a Precharge All command on DDR bus

Problem: When an auto-refresh command is issued to the MCU SDRAM Initialization Register (write 0x6 to SDIR), the hardware state machine executes a Precharge All command and then an Auto-Refresh command.

Implication: Some DIMMs might fail to initialize.

Workaround: There is no way to decouple the Precharge All and Auto-Refresh commands in the MCU. However, the DCAL can issue an Auto-Refresh command, which can be used instead to issue the Auto Refresh for initialization.

For both DDRI and DDRII initialization sequences, there are two back-to-back Auto-Refresh (ARF) commands issued to the DIMM (accomplished by writing 0x6 to MCU_SDIR 0xFFFF_E500 twice). This sequence is replaced by writing to the DCAL Control and Status Register (DCALSR) 0xFFFF_F500 to issue an Auto-Refresh to each Chip Select; thus, there are four writes to the DCALSR as opposed to the previous two writes to MCU_SDIR. The exact address and data values are as follows:

ADDRESS	WRITE_DATA		NOTE
0xFFFF_F500	0x8000_0001	--	ARF to CS[0]
0xFFFF_F500	0x8000_0001	--	ARF to CS[0]
0xFFFF_F500	0x8010_0001	--	ARF to CS[1]
0xFFFF_F500	0x8010_0001	--	ARF to CS[1]

Status: **No Fix.** Not to be fixed. See the [Table , “Summary Table of Changes” on page 7.](#)

8. Coalesced writes to 32-bit memory can cause data corruption

Problem: The Bus Interface Unit (BIU) can cause data corruption within the memory controller unit when either an 8-byte write with a starting address offset of four (that is, Dword alignment) or a 12-byte write is done to 32-bit memory (or the 32-bit memory region).

Implication: This issue can cause data corruption. This is not an issue with 64-bit memory subsystems.

Workaround: When the use of 32-bit memory or the 32-bit memory region is required, write coalescing must be turned off.

Status: **No Fix.** Not to be fixed. See the [Table , “Summary Table of Changes” on page 7.](#)

9. ATU passing rules operation in PCI mode

Problem: When the A-segment bus is in PCI bus mode, the PCI passing rule enforcement logic within the ATU allows a read completion to pass write data, until at least four inbound delayed reads, inbound configuration writes, inbound configuration reads, or any combination of these have occurred from the PCI bus.

Implication: This issue causes a deadlock condition in legacy devices that contain shared read and write data queues, where the device allocates the data buffer for the requested delayed read data and is also being addressed by the outbound ATU write data.

This ATU functionality does not exist in PCI-X mode.

Workaround: Use the A-segment bus in PCI-X mode.

When the A-segment bus is in PCI mode, configuration retry is enabled, and the legacy buffer allocating device is present in the system, the ATU must not issue writes to that device until the configuration cycles or reads have completed. When this situation cannot be achieved by the host, the ATU can be momentarily programmed in loopback mode and issue delayed reads to itself.

This workaround is required only when the ATU is sending upstream traffic at the same time as the host is configuring it. This should not happen since the ATU needs to wait for the driver to be enabled after enumeration completes.

This erratum does not occur when configuration retry is deasserted at power-on and the host meets the above configuration cycles.

Status: **No Fix.** Not to be fixed. See the [Table , “Summary Table of Changes” on page 7.](#)

10. Secondary bus PCI RST# pulse prior to the rising edge of PWRGD

Problem: During system power-on and prior to the 80333 receiving the rising edge of PWRGD, a pulse is observed on the secondary bus PCI RST# signals.

This functionality is a result of the 3.3v to 1.5v leakage described in Specification Clarification [22](#). Other signals that may see a pulse during power-on include the following: all Peripheral Bus Interface (PBI), PCI, GPIO, UART, JTAG and PWRDELAY signals. Refer to Specification Clarification [32](#) for specifics on PWRDELAY. Signals not included are DDR, PCI-Express and I2C.

Implication: PCI/PCI-X controllers on the secondary bus segments might interpret this PCI RST# pulse as a true rising edge and initialize into an undetermined state. Pulses on PWE# and PCEx# may cause data corruption for memory devices connected to the PBI bus.

Workaround: A hardware workaround has been identified. Use the PWRGD signal that is received by the 80333 to gate the secondary bus PCI RST# signals. For example, use PWRGD and A_RST# as inputs to an AND gate and connect the output to the secondary device RST# pin. The gate delay must be kept down to ~2 ns, so as to not interfere with the PCI initialization pattern.

Another workaround is to bring up 3.3 V and 1.5 V power rails simultaneously, but continue to maintain the power-sequencing requirements (as specified in the *Intel® 80333 I/O Processor Design Guide*) for these two power rails.

Status: **No Fix.** Not to be fixed. See the [Table , “Summary Table of Changes” on page 7.](#)

11. VPD Data Register bit[19] is not read/write

- Problem:** VPD (Vital Product Data) is an extended capability of the ATU. Bit[19] of the VPD Data Register (FFFF_E1BCh) is implemented as RC (read clear) instead of RW (read/write).
- Implication:** This is application-dependent, as VPD provides the system with information that uniquely identifies hardware and, potentially, software elements of a system.
- Workaround:** When the VPD feature is needed, care must be taken to mask bit[19].
- Status:** **No Fix.** Not to be fixed. See the [Table , “Summary Table of Changes” on page 7.](#)

12. PCI Express* Correctable Error Mask Bits

- Problem:** PCI Express* Correctable Error Mask (offset 114) requires both PCI/PCI-X segment A and segment B register bits to be set to take effect.
- Implication:** When the mask bit is only set for one segment, the error is still reported to the host bridge.
- Workaround:** The 80333 requires mask bits in both A- and B-segments to be set to mask the PCI Express* correctable errors.
- Status:** **No Fix.** Not to be fixed. See the [Table , “Summary Table of Changes” on page 7.](#)

13. DMA CRC result is byte-reversed

- Problem:** The DMA CRC result value is byte-reversed.
- Example CRC operation with the DMA:
- CRC Seed: 0x0000 0000
Data Pattern (16 bytes): 0x1234 5678, 0x1457 9098, 0x1234 5678, 0x1457 9098
Expected CRC Value: 0x6791 25DA
Actual CRC Value: 0xDA25 9167
- CRC Seed: 0x1122 3344
Data Pattern(16 bytes): 0x1234 5678, 0x1457 9098, 0x1234 5678, 0x1457 9098
Expected CRC Value: 0x1D2C B941
Actual CRC Value: 0x41B9 2C1D
- Implication:** Data corruption occurs when software does not take care of the byte reversal.
- Workaround:** Software must byte reverse the CRC result after the DMA completes.
- Status:** **No Fix.** Not to be fixed. See the [Table , “Summary Table of Changes” on page 7.](#)

14. CRC corruption on PCI-to-local DMA transfers

- Problem:** CRC corruption can occur when polling DMA registers during PCI-to-local transfers. CRC corruption happens regardless of whether the transfer data is written to memory (DMA Transfer disable bit DCR.7, is set). Both DMA channels are affected.
- Implication:** CRC data corruption occurs. The DMA transfer itself is not affected by this erratum.
- Workaround:** Use local address sources only when calculating CRC.
- Status:** **No Fix.** Not to be fixed. See the [Table , “Summary Table of Changes” on page 7.](#)

15. IOAPIC End of Interrupt (EOI) Register is Read-Write, should be Write-Only

Problem: IOAPIC End of Interrupt register (APIC_EOI, offset 40h) should be write-only. The APIC specification specifies that this register be implemented as write-only. This register was inadvertently implemented as read-write.

Implication: When implemented as write-only, this register returns FFh when read. Since this register is implemented as read-write, the 80333 returns the “real” value of the register contents when read. No negative impact expected.

Workaround: None

Status: **No Fix.** Not to be fixed. See the [Table](#) , “Summary Table of Changes” on page 7.

16. 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 80333.

Implication: When L0s is enabled, the system hangs and other system instability occurs.

Workaround: None

Status: **No Fix.** Not to be fixed. Refer to Specification Clarification 36 (“PCI Express* L0s functionality not supported in the 80333” on page 43) and the latest version of the *Intel® 6700PXH 64-bit PCI Hub BIOS Specification Update* for details on how to disable L0s support. See the [Table](#) , “Summary Table of Changes” on page 7.

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

Problem: During a downstream memory write to the 80333, the following erroneous behavior is seen when PERR# is asserted on the secondary bus:

- Signaled System Error (SSE) in the PSTS register (D:0, F:0 and 2, offset 06h, bit[14]) is set when SERR# Enable (SEE) (D:0, F:0 and 2, offset 04h, bit[8]) and Parity Error Response Enable (PERE) (D:0, F:0 and 2, offset 04h, bit[6]) are set in the PCICMD register.
- The PERE bit in the BCTRL register is set (D:0, F:0 and 2, offset 3Eh, bit[0]).
- Error reporting is disabled in the PCI_XERRUNC_MSK register (D:0, F:0 and 2, offset 130h).

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

Workaround: None

Status: **No Fix.** Not to be fixed. See the [Table](#) , “Summary Table of Changes” on page 7.

18. 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 is 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 when error escalation is not enabled.

Workaround: Uncorrectable error escalation must be enabled in the MCH and the 80333 to contain this data parity escape. Therefore, a complete workaround for this erratum also includes MCH/root complex specific BIOS updates. Refer to the latest version of the *Intel® 6700PXH 64-bit PCI Hub BIOS Specification Update* for details on this workaround.

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

Status: **No Fix.** Not to be fixed. See the [Table , “Summary Table of Changes” on page 7.](#)

19. ATU claims PCI commands 8 and 9 when issued as Dual Address Cycle (DAC)

Problem: In PCI mode, commands 8 and 9 are reserved. The appropriate PCI response to these commands is to master abort. When these commands are issued as a Dual Address Cycle (DAC), the Address Translation Unit (ATU) claims them, and they are executed as Memory Read (command 8) and Memory Write (command 9) on the internal bus. The ATU properly master aborts SAC PCI commands 8 and 9.

This issue does not occur in PCI-X mode.

Implication: No negative impact is expected, since these PCI commands are “reserved” and are never normally issued to the ATU.

Workaround: None

Status: **No Fix.** Not to be fixed. See the [Table , “Summary Table of Changes” on page 7.](#)

20. Failure to train down in presence of degraded lane

Problem: During the PCI Express* training sequence, when a broken endpoint has correct receiver termination on a lane and transmits training sequences on the lane which are invalid, the 80333 fails to link train.

Implication: The PCI Express* specification intends that, when some lanes are transmitting invalid data instead of valid training sequences, those lanes must be treated as broken, and the link must fail down to an acceptable width, such as $\times 1$. When Lane 0 is failing in this manner, the PCI Express* specification anticipates that the link would fail to train. When a higher-numbered lane is failing in this manner, the PCI Express* specification requires that the link attempt to train as $\times 1$ on lane 0. In either case, the 80333 does 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. The 80333 trains properly for these failure modes, since either the receiver termination is missing, or the transmitted signals are not seen at the 80333. In order to see invalid transmitted data on lanes at the 80333, either a logic bug in the other PCI Express* endpoint is required, or a signal integrity issue so severe as to make operation impossible, such as a broken or intermittent connection.

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

Status: **No Fix.** Not to be fixed. See the [Table , “Summary Table of Changes” on page 7.](#)

21. 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 do not maintain their values after a warm/hot reset. These registers are supposed to be unaffected by a warm/hot reset, but instead, they are reset to default values. The following registers with “sticky” bits are affected:

- ADVERR_CTL (offset 118h)
- HDR_LOG (offset 11Ch)
- PCI-XERRUNC_PTR (offset 138h)
- PCI-XHDR_LOG (offset 13Ch)
- PCI-XD_LOG (offset 14Ch)
- PCI-X_ERRLOGCTL (offset 154h)

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

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

Workaround: No workaround

Status: **No Fix.** Not to be fixed. See the [Table](#) , “Summary Table of Changes” on page 7.

22. 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 is reported as non-fatal. This is in contradiction with the *PCI Express* Specification*, Revision 1.0a, which requires a default value of 1, indicating an uncorrectable flow control error is reported as fatal.

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

Workaround: This bit can be reprogrammed to match the specified default value when desired. BIOS or firmware must set register 10Ch bit[13] to 1 for both A- and B-segments. For accessing extended bridge configuration space from the Intel XScale® processor, see specification clarification [23](#), “Accessing extended bridge configuration space from the Intel XScale® processor” on page 40.

Status: **No Fix.** Not to be fixed. See the [Table](#) , “Summary Table of Changes” on page 7.

23. 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 PM_CSR at offset 70h), allow a reserved value of 01b or 10b to be written. This is contrary to the specification, which originally stated that when software attempts to write an unsupported reserved state to this field, the data must be discarded and a state change must not occur.

Implication: When a reserved state is written to this field, there is 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 can cause the 80333 to transition to the D0 power state, regardless of the previous power state.

Workaround: Do not write a reserved value to this bit field.

Status: **No Fix.** Not to be fixed. See the [Table](#) , “Summary Table of Changes” on page 7.

24. Performance across an upstream ×1 PCI Express* link is less than expected

Problem: When the 80333 is configured with an upstream ×1 PCI Express* link, the realized performance is significantly less than the predicted linear assumption that a x1 link provides 1/4 the performance of a x4 link. This is caused by circumstances where the 80333 must discard a large portion of the data it receives across the upstream link. Notably, anytime the 80333 services an incorrect prefetch, or anytime the 80333 services interleaved requests from multi-function devices, the 80333 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: When possible, system designers must reduce the amount of prefetching allowed to devices behind the 80333.

Status: **No Fix.** Not to be fixed. See the [Table , “Summary Table of Changes” on page 7.](#)

25. PCI Express* ESD enhancement requires a change to register setting

Problem: Validation has shown PCI Express* ESD enhancement, with changes to undisclosed registers in the 80333.

Implication: The workaround below increases the margin for the eye and therefore results in a healthier, improved PCI Express* link.

Workaround: BIOS or firmware must set F0/F2:R260h bit[15] to 1 and clear F0/F2:R270h bit[31] to 0. This workaround is required for both cold and warm reset.

Status: **No Fix.** Not to be fixed. The BIOS/firmware workaround must be left in place for all 80333 steppings. See the [Table , “Summary Table of Changes” on page 7.](#)

26. SKP ordered set might not be sent within required interval during link recovery when a packet is pending

Problem: During Link Recovery on the PCI Express* port, the 80333 might fail to transmit a SKP ordered set within the required time interval as defined in the *PCI Express* Specification*, Revision 1.0a when a TLP or DLLP was pending when the link entered Recovery.Idle state.

Implication: When the receiving device depends upon receipt of an SKP ordered set to progress through Link Recovery, a time-out occurs, 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 time-out and reinitialization, the PCI Express* link can be expected to resume normal operation unless the original Link Recovery condition was entered as a result of a hard failure mechanism.

Workaround: None

Status: **No Fix.** Not to be fixed. See the [Table , “Summary Table of Changes” on page 7.](#)

27. SERR fatal/non-fatal error message enabled with incorrect error message enabled bit

Problem: When the SERR# Enable (SEE) bit in the 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: The SERR fatal error message can be generated only when 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: **No Fix.** Not to be fixed. See the [Table](#) , “Summary Table of Changes” on page 7.

28. Configuration write to offset 70h of A- and B-bridge (PM_CSR - PCI Express Power Management Control/Status Register) using non-continuous byte enables does not capture the data value

Problem: A configuration write to single-byte offset 71h of the A-bridge and B-bridge configuration space does not capture the data value written to it. Specifically, performing this configuration write to offset 70h with byte enables of 2h, 6h, or Ah in order to write to offset 71h does not work.

Implication: Register offset 70h is the PCI Express* Power Management Control/Status Register (PM_CSR). The byte at location 71h contains the PME enable bit (bit[8] of 70h) and PME status bit (bit[15] of 70h). The above mechanism does not work to either set or clear the PME enable bit, or to clear the PME status bit.

Workaround: Write to offset 71h by performing a word (byte enable 3h) or a dword (byte enable Fh) to offset 70h.

Status: **No Fix.** Not to be fixed. See the [Table](#) , “Summary Table of Changes” on page 7.

29. The 80333 might become unresponsive when transitioning into the D3 power state

Problem: When the 80333 is transitioned to a power state of D3 or lower, the 80333 device might become unresponsive.

Implication: There have been no observed failures on systems with currently available software. Operating systems that independently manage the power state of the 80333, outside the scope of system-level power state transitions, might result in the loss of link communications to the MCH.

Workaround: Independent device power-state management of the 80333 must be avoided. If the 80333 becomes unresponsive, a fundamental device reset must be asserted to return the system to normal operation.

Status: **No Fix.** Not to be fixed. See the [Table](#) , “Summary Table of Changes” on page 7.

30. Bus Interface Unit (BIU) claims DAC addresses in the range of the Memory Mapped Registers (MMR)

Problem: The BIU incorrectly decodes and claims Dual Address Cycle (DAC) addresses in the xxxx_xxxx_FFFF_E000h to xxxx_xxxx_FFFF_FFFFh range (e.g. - 'x' represents any bit being set to '1'). The 32-bit address range of FFFF_E000h to FFFF_FFFFh on the internal bus, represents MMR and reserved space. When a 64-bit address in this range is presented on the internal bus, multiple internal bus units, one of them being the BIU, will claim the transaction.

Note: Only the DMA can generate a 64-bit address (DAC) on the internal bus.

Implication: Using DAC addresses in the xxxx_xxxx_FFFF_E000h to xxxx_xxxx_FFFF_FFFFh range on the internal bus will cause an internal bus conflict that may result in the reception of undesired data and setting of error flags.

Workaround: Avoid using the DMA with DAC addresses in the xxxx_xxxx_FFFF_E000h to xxxx_xxxx_FFFF_FFFFh range. This can be done by utilizing one of the two 64MB ATU outbound memory windows (8000 0000H or 8400 0000H) and its corresponding outbound translation registers (OMWTVR0/OUMWTR0 or OMWTVR1/OUMWTR1) in order to present a 32 bit address on the internal bus and generate a 64 bit address on the PCI bus.

The upper translate value register should be programmed with the upper 32 bits of the desired PCI address. The lower translate value register would then be configured to OR in the appropriate value such that the desired lower 32 bits appear on the PCI bus after translation. Refer to Section 3.2.2 "Outbound Transactions – Single Address Cycle (SAC) Internal Bus Transactions" in the *Intel® 80333 I/O Processor Developer's Manual* for more information on how the windowing and translation scheme works.

It is possible to now generate the 32 bit internal bus transaction either using the core or the DMA. Both options are viable and one may be preferable over the other depending on the application.

Note: Use of the DMA to generate the transaction would require not only the modification of the descriptors in question and setting of the memory-memory transfer enable bit (DCRx) for those descriptors, but care must also be taken not to allow any individual transfer to overrun the size of the outbound window (64MB).

Status: No Fix. Not to be fixed. See the [Table , "Summary Table of Changes" on page 7.](#)

31. Tc1(min) of the PCI-X clock observed to be marginally less than the requirement specified for the PCI-X (Mode 1, class1) clock jitter.

Problem: The 80333 generates the PCI-X clock with a nominal frequency of 133 MHz (Tc1 of 7.5 ns). After considering the clock jitter, the minimum clock period (Tc1(min)) observed at the pin may be less than 7.5 ns. The PCI-X class 1 clock jitter specification for mode 1 requires the Tcyc-min to be 7.5 ns with jitter consideration. The same applies for the PCI-X clock generated @ 66 MHz.

Implication: No negative impact is expected, when compliant with the routing guidelines.

Workaround: There is no workaround to adjust the minimum clock period (Tc1(min)) of the PCI-X clocks. However, the routing guidelines for the PCI-X clock signal take into consideration the effect of the jitter on the minimum clock period (Tc1(min)). Conforming to the routing guidelines in the 80333 design guide will offset the effect of the marginally reduced minimum clock period (Tc1(min)) towards the setup and hold times. Therefore, for system boards that are compliant with the routing guidelines, the risk of violating the setup and hold time requirements and any resulting functional impact, is low. Please refer to the 80333 I/O processor design guide (305434) for further information on the PCI-X clock routing guidelines.

Status: No Fix. Not to be fixed. See the [Table , "Summary Table of Changes" on page 7.](#)

32. I2C Control Register reset bit does not function

Problem: The I2C Control Register (ICR0 and ICR1) bit 14 is supposed to be used for resetting the I2C unit, but writing a '1' to this bit does not reset the I2C unit. Writing a '1' to bit 14 has no effect.

Implication: The I2C unit cannot be reset by using ICRx.14.

Workaround: Depends on what needs to be accomplished. Asserting P_RST# or setting BCR.6 will reset the I2C unit but will also reset the entire chip or the secondary bus/ATU. For an I2C bus lock condition, it may be cleared by software doing a toggle of the GPOD[11:10] to toggle SCL[1:0]. If SDA[1:0] need to be toggled, then an external device or unused GPIO will need to be used to control this sequence.

Status: **No Fix.** Not to be fixed. See the [Table , “Summary Table of Changes” on page 7.](#)

33. MSI Hot-Plug Interrupt issue

Problem: An MSI generated by the standard hot-plug controller may get corrupted in the 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 80333 SHPC MSI. By default MSI_MC (Offset 5Eh, function 2 only) bit 0 MSI Enable is clear. BIOS/FW must keep this default value.

Status: **No Fix.** Not to be fixed. See the [Table , “Summary Table of Changes” on page 7.](#)

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

Problem: With some prefetch policy settings, the 80333 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 exists on both 80333 PCI segments. 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 80333.

Workaround: BIOS or firmware must set D0:F0/F2 Offset 184h (Dword) bit [2] to 1. The workaround corrects the problem at all supported frequencies and all prefetch policy settings.

Status: **No Fix** Not to be fixed. The BIOS/firmware workaround must be left in place for all 80333 steppings. See the [Table , “Summary Table of Changes” on page 7.](#)

35. Internal Clock Misalignment Can Cause Processor Hang

- Problem:** After a reset, the 80333 can hang during initial accesses to SDRAM, due to a possible clock misalignment, aggravated by a race condition in the clock divider clear circuit.
- Implication:** A failure will manifest itself as a hang of the I/O processor after reset, during initial accesses to SDRAM. Subsequent warm or cold resets may clear the condition and allow the 80333 to continue operation.
- Workaround:** In most cases, doing a cold or warm reset will clear this condition. Increasing the 1.5v power supply will reduce the probability of a processor hang. Intel is screening parts to eliminate the probability of occurrence (refer to [Specification Change #6, “Internal Clock Misalignment”](#) on page 32).
- Status:** Fixed. This issue was fixed in the A-1 stepping of the product (this is also related to [Section 6, “Internal Clock Misalignment”](#) on page 32). See the [Table , “Summary Table of Changes”](#) on page 7.

36. Spurious DMA0 End-Of-Transfer Interrupt

- Problem:** When the interrupt controller goes from having no interrupts asserted to one or more asserted, there is a 1-clock cycle window in which the IINTVEC (IRQ Interrupt Vector register; FFFF_E7C8h or CP6, register 14) or FINTVEC (FIQ Interrupt Vector register; FFFF_E7CCh or CP6, register 15) may report the value of the INTBASE register (Interrupt Base register; FFFF_E7C0h or CP6, register 12), which is the vector address for interrupt 0, DMA0 End-of-Transfer. This condition can occur even if the DMA0 EOT interrupt is masked, INTCTL0.0 = 0.
- Implication:** No negative impact expected. If the routine that reads the IINTVEC/FINTVEC qualifies the return value against IINTSRC0.0/FINTSRC0.0, it will either see there is nothing to do or it will validly call the DMA0 End-of-Transfer handler.
- Workaround:** If IINTVEC/FINTVEC equals INTBASE, then re-read the IINTVEC/FINTVEC register.
- Status:** [No Fix](#). Not to be fixed. See the [Table , “Summary Table of Changes”](#) on page 7.

Core Errata

1. Abort is missed when lock command is outstanding

Problem: A bus abort occurs on a code fetch while an instruction TLB or I-Cache lock MCR, Move to Coprocessor from ARM Register, command is outstanding. The core fails to abort, instead executing the instruction returned on the aborting transaction. Parity errors are not affected. The bus abort can be due to either an abort pin assertion or a multi-bit ECC error.

Workaround: Branch flush after every I-TLB or I-Cache lock. For example, the following instruction does this: SUB PC, PC #4;flush the pipe.

Status: No Fix. See the [Table , “Summary Table of Changes” on page 7.](#)

2. Aborted store that hits the data cache can mark write-back data as dirty

Problem: When there is an aborted store that hits clean data in the data cache (data in an aligned four-word range that has not been modified from the core since it was last loaded in from memory or cleaned), the data in the array is not modified (the store is blocked), but the dirty bit is set. When the line is then aged out of the data cache or explicitly cleaned, the data in that four-word range is evicted to external memory, even though it has never been changed. In normal operation this is nothing more than an extra store on the bus that writes the same data to memory that is already there.

Here is the boundary condition where this might be visible:

1. A cache line is loaded into the cache at address A.
2. Another master externally modifies address A.
3. A core store instruction attempts to modify A, hits the cache, aborts because of MMU permissions, and is backed out of the cache. That line normally is not marked dirty, but because of this erratum, it is marked as dirty.
4. The cache line at A then ages out or is explicitly cleaned. The original data from location A is evicted to external memory, overwriting the data written by the external master. This happens only when software is allowing an external master to modify memory that is, write-back or write-allocate in the core page tables, and depending on the fact that the data is not “dirty” in the cache, to preclude the cached version from overwriting the external memory version. When there are any semaphores or any other handshaking to prevent collisions on shared memory, this is not a problem.

Workaround: For this shared memory region, mark it as write-through memory in the core page table. This prevents the data from ever being written out as dirty.

Status: No Fix. See the [Table , “Summary Table of Changes” on page 7.](#)

3. Performance Monitor Unit event 0x1 can be incremented erroneously by unrelated events

Event 0x1 in the Performance Monitor Unit (PMU) can be used to count cycles in which the instruction cache cannot deliver an instruction. The only cycles counted are supposed to be those due to an instruction cache miss or an instruction TLB miss. The following unrelated events in the core, also causes the corresponding count to increment when event number 0x1 is being monitored:

- Any architectural event (for example, IRQ, data abort)
- MSR instructions which alter the CPSR control bits
- Some branch instructions, including indirect branches and those mispredicted by the BTB
- CP15 MCR instructions to registers 7, 8, 9, or 10, which involve the instruction cache or the instruction TLB.

Each of the preceding items can cause the performance-monitoring count to increment several times. The resulting performance monitoring count can be higher than expected when the preceding items occur, but has not been observed to be lower than expected.

Workaround: There is no way to obtain the correct number of cycles stalled due to instruction cache misses and instruction TLB misses. Extra counts, due to branch instructions mispredicted by the BTB, might be one component of the unwanted count that can be filtered out.

The number of mispredicted branches also can be monitored using performance-monitoring event 0x6 during the same time period as event 0x1. The mispredicted branch number can then be subtracted from the instruction cache stall number generated by the performance monitor to get a value closer to the correct one. This workaround addresses only counts contributed by branches that the BTB is able to predict.

All the items in the preceding bulleted list still affect the count. Depending on the nature of the code being monitored, this workaround might have limited value.

Status: No Fix. See the Table , “Summary Table of Changes” on page 7.

4. In Special Debug State, back-to-back memory operations—in which the first instruction aborts—can cause a hang

Problem: When back-to-back memory operations occur in the Special Debug State (SDS, used by ICE and Debug vendors), and the first memory operation gets a precise data abort, the first memory operation is correctly cancelled and no abort occurs. Depending on the timing, however, the second memory operation might not work correctly. The data cache might internally cancel the second operation, but the register file may have scoreboarded registers for that second memory operation. The effect is that the core might hang (due to a permanently scoreboarded register) or that a store operation might be incorrectly cancelled.

Workaround: In Special Debug State, any memory operation that can cause a precise data abort must be followed by a write-buffer drain operation. This precludes further memory operations from being in the pipe when the abort occurs. Load Multiple/Store Multiple that might cause precise data aborts must not be used.

Status: No Fix. See the Table , “Summary Table of Changes” on page 7.

5. Accesses to the CP15 ID Register with opcode2 > 001b returns unpredictable values

Problem: The *ARM Architecture Reference Manual* (ARM DDI 0100E) states the following in chapter B-2, section 2.3:

When an <opcode2> value corresponding to an unimplemented or reserved ID register is encountered, the System Control processor returns the value of the main ID register. ID registers other than the main ID register are defined so that when implemented, their value cannot be equal to that of the main ID register. Software can therefore determine whether they exist by reading both the main ID register and the desired register and comparing their values. When the two values are not equal, the desired register exists.

The Intel XScale® processor does not implement any CP15 ID code registers other than the Main ID Register (opcode2 = 000b) and the Cache Type register (opcode2 = 001b). When any of the unimplemented registers are accessed by software (for example, mrc p15, 0, r3, c15, c15, 2), the value of the Main ID register is to be returned. Instead, an unpredictable value is returned.

Workaround: None

Status: No Fix. See the Table , “Summary Table of Changes” on page 7.

6. Disabling and re-enabling the MMU can hang the core or cause it to execute the wrong code

Problem: When the MMU is disabled via the CP15 control register (CP15, CR1, opcode_2 = 0, bit[0]) after being enabled, certain timing cases can cause the processor to hang. In addition to this, re-enabling the MMU after disabling it can cause the processor to fetch and execute code from the wrong physical address. To avoid these issues, the code sequence below must be used whenever disabling the MMU or re-enabling it afterwards.

Workaround: The following code sequence can be used to disable and/or re-enable the MMU safely. The alignment of the mcr instruction that disables or re-enables the MMU must be controlled carefully, so that it resides in the first word of an instruction cache line.

```
@ The following code sequence takes r0 as a parameter. The value of r0 will be
@ written to the CP15 control register to either enable or disable the MMU.
mcr p15, 0, r0, c10, c4, 1 @ unlock I-TLB
mcr p15, 0, r0, c8, c5, 0 @ invalidate I-TLB
mrc p15, 0, r0, c2, c0, 0 @ CPWAIT
mov r0, r0
sub pc, pc, #4
b 1f @ branch to aligned code
.align 5
1:
mcr p15, 0, r0, c1, c0, 0 @ enable/disable MMU, caches
mrc p15, 0, r0, c2, c0, 0 @ CPWAIT
mov r0, r0
sub pc, pc, #4
```

Status: No Fix. See the Table , “Summary Table of Changes” on page 7.

7. Updating the JTAG parallel registers requires an extra TCK rising edge

Problem: The IEEE 1149.1 specification states that the effects of updating all parallel JTAG registers are expected to be seen on the falling edge of TCK in the Update-DR state. The Intel XScale® processor parallel JTAG registers incorrectly require an extra TCK rising edge to make the update visible. Therefore, operations like hold-reset, JTAG break, and vector traps require either an extra TCK cycle by going to Run-Test-Idle or by cycling through the state machine again in order to trigger the expected hardware behavior.

Workaround: When the JTAG interface is polled continuously, this erratum has no effect. When not, an extra TCK cycle can be caused by going to Run-Test-Idle after writing a parallel JTAG register.

Status: No Fix. See the Table , “Summary Table of Changes” on page 7.

8. Non-branch instruction in vector table might execute twice after a thumb mode exception

Problem: When an exception occurs in thumb mode and a non-branch instruction is executed at the corresponding exception vector, that instruction might execute twice. Typically instructions located at exception vectors must be branch instructions which go to the appropriate handler, but the ARM architecture allows the FIQ handler to be placed directly at the FIQ vector (0x0000001c/0xffff001c) without requiring a branch. Because of this bug, the first instruction of such an FIQ handler might be executed twice when it is not a branch instruction.

Workaround: When a “NOP” is placed at the beginning of the FIQ handler, the “NOP” executes twice and no incorrect behavior results. When a branch instruction is placed at the beginning of the handler, it is not executed twice.

Status: No Fix. See the Table , “Summary Table of Changes” on page 7.

Specification Changes

1. Recommended DLL register values

Issue: Using the default DLL values in combination with low duty cycle DIMMs may result in low hold time margins on read transactions.

The default values for the DLL master and slave registers do not center the internal DQS with respect to the eye of the incoming data. Using the default values may result in a reduced hold time due to DQS being late in the data eye, which could lead to ECC errors. Errors have only been observed when using DIMMs that have a low DQS duty cycle.

Note: All of the Intel validation up to this point has been with the default, worst case DLL values and all DIMMs used in validation have passed.

Firmware should be updated and tested with these new DLL values, in order to add margin to the hold timing during memory reads.

DDR-II 400 settings

SLVLMIX0 - Address FFFF_F554h; Recommended value - 3333_3333h

SLVLMIX1 - Address FFFF_F558h; Recommended value - 0000_0003h

SLVHMIX0 - Address FFFF_F55Ch; Recommended value - 3333_3333h

SLVHMIX1 - Address FFFF_F560h; Recommended value - 0000_0003h

SLVLEN - Address FFFF_F564h; Recommended value - 0000_0003h

MASTMIX - Address FFFF_F568h; Recommended value - 0000_000Ah

MASTLEN - Address FFFF_F56Ch; Recommended value - 0000_0002h

DDR-I 333 settings

SLVLMIX0 - Address FFFF_F554h; Recommended value - 6666_6666h

SLVLMIX1 - Address FFFF_F558h; Recommended value - 0000_0006h

SLVHMIX0 - Address FFFF_F55Ch; Recommended value - 6666_6666h

SLVHMIX1 - Address FFFF_F560h; Recommended value - 0000_0006h

SLVLEN - Address FFFF_F564h; Recommended value - 0000_0003h

MASTMIX - Address FFFF_F568h; Recommended value - 0000_0000h

MASTLEN - Address FFFF_F56Ch; Recommended value - 0000_0002h

2. DDR-II JEDEC initialization sequence includes writes to EMRS2 and EMRS3

Issue: The JEDEC DDR-II specification includes a write to EMRS2 and EMRS3 (Extended Mode Register Set) during the initialization sequence. Step 5 is 'Issue EMRS2 command' and step 6 is 'Issue EMRS3 command'. In order to be JEDEC compliant, these steps should be added to the memory controller initialization sequence.

Note: Before implementing, check with your DIMM/memory manufacturer to determine if these steps are necessary. Software should always follow the initialization sequence provided by the DIMM/memory manufacturer guidelines.

The following pseudo code shows the EMRS initialization steps that are required to be compliant with the JEDEC DDR-II initialization sequence.

```
// Step 5 and 6 - EMRS(2) and EMRS(3) programming
```

```
if MemoryType is DDR-II
Write 0x2 to DCALADDR (Setting BA[1:0] for EMRS(2))
Write 0x81000003 to DCALCSR (issue EMRS command to CS0)
Wait for DCALCSR[31] to report 'Operation Completed'
Write 0x81100003 to DCALCSR (issue EMRS command to CS1)
Wait for DCALCSR[31] to report 'Operation Completed'
Write 0x3 to DCALADDR (Setting BA[1:0] for EMRS(3))
Write 0x81000003 to DCALCSR (issue EMRS command to CS0)
Wait for DCALCSR[31] to report 'Operation Completed'
Write 0x81100003 to DCALCSR (issue EMRS command to CS1)
Wait for DCALCSR[31] to report 'Operation Completed'
endif
```

3. REFCLK relationship to voltage rails

Issue: The current 80333 design guide (305434) states the following in section 10.1:

"Also, all 80333 voltage rails must be stable and within their operating ranges before the PCI Express differential clocks REFCLK+ and REFCLK- begin running. This is a requirement for all devices with PCI Express interfaces."

When the 80333 is on an add-in card, only 3.3V and 12V are provided to the slot, therefore, a local regulator is required for 1.5V and 2.5V generation. Due to the delay by the local regulators, REFCLK may already be provided before the power rails are stable. If this is the case, no device overstress will occur, provided that the REFCLK input current does not exceed 900mA and the input voltage does not exceed the PCI Express specification of 1.15V. REFCLK buffers on many Intel platforms show an input current of 15.6mA, well under the 900mA limit.

The requirement for "all 80333 voltage rails to be stable before the PCI Express differential clocks REFCLK+ and REFCLK- begin running" is no longer a requirement.

4. Case temperature (Tcase) change

Issue: To be consistent with the production test environment, the case temperature (Tcase) for the 80333 I/O processor has been changed from 105° c to 95° c.

5. PWRGD and RSTIN# sequencing

Issue: Strapping RSTIN# high (as stated in the 80333 design guide v001, table 3), bringing RSTIN# up prior to PWRGD rising edge or tying PWRGD and RSTIN# together, could result in platform failures.

The correct sequencing between PWRGD and RSTIN#, is for the rising edge of RSTIN# to follow the rising edge of PWRGD by at least 1ms.

6. Internal Clock Misalignment

Issue: Due to non-core [Erratum 35, Internal Clock Misalignment Can Cause Processor Hang, on page 26](#), Intel is screening parts to eliminate the probability of occurrence. Until this is fixed in a future stepping, a screen has been implemented which will screen out parts exhibiting this issue with a VCC15 greater than 1.46v.

With the screen at 1.46v, the on-board 1.5v power rail should not be allowed to go below 1.46v, as this would increase the risk of failure. The 1.5v rail minimum is currently specified in the datasheet as 1.425v, therefore for screened parts, the minimum is changed to 1.46v.

Status: Fixed. This issue was fixed in the A-1 stepping of the product.

Specification Clarifications

1. PCI Express* to PCI-X Bridge does not support Device ID Messaging (DIM)

Issue: The PCI Express* to PCI-X Bridge does not forward DIM transactions. This is an optional feature in the *PCI Express* to PCI/PCI-X Bridge Specification, Rev. 1.0*.

Status: No Fix. See the [Table](#) , “Summary Table of Changes” on page 7.

2. 64 MB and 2 GB DDR333 capacities not tested in post-silicon validation

Issue: Intel is not able to test 64 MB and 2 GB DDR333 DIMMs due to availability. Intel cannot guarantee proper functionality since validation cannot be completed.

Status: No Fix. See the [Table](#) , “Summary Table of Changes” on page 7.

3. DDR-II 400 unbuffered DIMMs are not supported

Issue: The 80333 supports DDR333 buffered (registered) and unbuffered DIMMs, but supports only DDR-II 400 buffered (registered) DIMMs. DDR-II 400 unbuffered DIMMs are not supported.

Status: No Fix. See the [Table](#) , “Summary Table of Changes” on page 7.

4. Memory map for 2 GByte of DDR memory

Issue: The 80333 can support up to 2 Gbytes of DDR SDRAM, but it cannot cross a 2 GB boundary. Therefore it must be mapped to either 0x00000000–0x7FFFFFFF or 0x80000000–0xFFFFFFFF. Either range conflicts with one or more of the statically assigned regions. The recommendation is to disable the direct outbound ATU window, in order to use the larger 2 GB memory, by clearing ATUCR[8] (default setting is 0–disabled).

Status: No Fix. See the [Table](#) , “Summary Table of Changes” on page 7.

5. PCI configuration write anomaly when clearing BINIT[1]

Issue: There is an anomaly when the ATU does a configuration write to BINIT[1] of the A-segment configuration header. When clearing bit[1] by writing 00000009h to the BINIT Configuration Register (offset FCh), the PCI bridge retries the write the first time. The PCI specification states that no data is to be transferred on retries. The configuration register is updated to 00000009h; therefore, data did get transferred on the retry. The PCI bridge ignores the second write, giving a master abort. BINIT[1] has a bit attribute of Read and Write Once (RWOS).

The worst-case implication is that a master abort error occurs when turning off inbound configuration accesses (BINIT[1]–A-segment).

This anomaly is not seen when the A-segment is running in PCI-X mode, but only when in PCI mode.

Status: No Fix. See the [Table](#) , “Summary Table of Changes” on page 7.

6. PWRGD and PERST# are the same signals

Issue: The 80333 uses PWRGD to identify the primary reset signal as described in the *PCI Express Card Electromechanical Specification*, Revision 1.0. This is the same signal as PERST#, which is described in the *PCI Express Card Electromechanical Specification*, Revision 1.0a.

Status: No Fix. See the [Table](#) , “Summary Table of Changes” on page 7.

7. Back-to-back MCU MMR reads

Issue: The memory controller unit (MCU) returns the wrong memory-mapped register (MMR) read data, when two MMR read transactions are enqueued into the transaction queues at the same time. This cannot happen from the BIU port as mapping the MMRs to this space is not allowed.

The only way this can occur is for two internal bus devices to request info from the MCU MMRs at the same time (with different addresses). For example, the BIU (via the Intel XScale® processor) and the ATU (via the host), which is a very unlikely usage model.

Status: No Fix. See the Table , “Summary Table of Changes” on page 7.

8. Reserved IDSELs on A-segment

Issue: The Intel® 80333 I/O Processor Developer’s Manual added Table 9, which shows AD24–AD27 as “reserved” IDSELs for the A-segment. These are reserved so that the Intel XScale® processor can access the extended configuration space of the bridge by using type-0 configuration cycles.

The A-side bridge can also be accessed by type-1 configuration cycles, but the primary bus number of the bridge must be tracked. On power-up, before the configuration retry bit is cleared, it is zero. When this technique is used, the AD24 and AD25 can be used for public or private devices based on BINIT[4], and AD26 and AD27 can be used for public devices.

Status: No Fix. See the Table , “Summary Table of Changes” on page 7.

9. Retry Disable Sequence

Issue: The 80333 has three integrated units (A-bridge, B-bridge, and ATU) that include retry bits, and these must be disabled in a specific sequence. When the retry bit is set, the PCI interface responds to all type-0 configuration cycles with a retry transaction. When the retry bit is cleared, type-0 transactions are completed normally. The default condition of the retry bit is determined by the RETRY reset strap muxed on AD[6].

Note: Since a PCIX reset independent of a PCIE reset will result in a reset of the Intel XScale® processor without resetting the bridge registers, the bridge retry release sequence MUST NOT be executed if the bridge retry bit is not set. Code can check the A segment BINIT register (bit 3) to determine if the bridge retry bits need to be cleared or not.

Proper retry disable sequence as initiated by the Intel XScale® processor:

1. Set the Bus Master Enable bit for the ATU in ATUCMD[2] (offset 04h) and enable the outbound ATU in the ATUCR.
2. Read A segment BINIT to determine if bridge is currently retrying type 0 config cycles.
3. If the bridge is retrying cycles (read A-bridge BINIT register and check if bit 3 is set):
 - a. Set the Bus Master Enable bit for the A-bridge in PCICMD[2] (offset 04h).
 - b. Read the primary bus number from the A-bridge for use in creating the type 1 transaction to the B side.
 - c. Enable the bridge (A side at a minimum, B side if application requires) to pass config retry cycles upstream by setting bit 15 in the PCI Express Device Control Register register at offset 0x4c.
 - d. Implement any errata workarounds that require bridge configuration cycle setup.
 - e. Clear the B-bridge retry bit in BINIT[3] (offset FCh using a type 1 transaction as defined in the PCI specification).

f. Restore the A-bridge command register and clear the A-bridge retry bit in BINIT[3] (offset FCh).

4. Clear (restore) the Bus Master Enable bit for the ATU in ATUCMD[2] (offset 04h).

5. Clear the ATU retry bit in PCSR[2] (offset 84h).

Note: Access to the bridge configuration header space is done by using a configuration cycle, which is generated by the ATU OCCAR and OCCDR registers.

Status: No Fix. See the [Table , “Summary Table of Changes” on page 7.](#)

10. Write requirements for the Peripheral Bus Interface

Issue: PBI write requirements are as follows:

- The Flash or memory must be set up to accept writes.
- Each write must be checked to make sure it has completed, before the next write can start.
- It is now allowed to burst writes to the PBI.

The data presented must not be larger than the PBI width.

Every PBI device must be mapped in the MMU the same way. Making the device address space cacheable can result in buffered/coalesced writes, which are burst to the PBI.

XCB = 000 and XCB = 101 are the only cache policies that can be used for PBI. All other cache policies result in multi-byte transactions.

Status: No Fix. See the [Table , “Summary Table of Changes” on page 7.](#)

11. PCI-X Status Register during PCI mode

Issue: The PCI-X Status Register (PX_SR, offset E4h-E5h) in the ATU has meaning only in PCI-X mode. The device number and bus number fields are always updated when a configuration write is detected. The ATU always grabs AD[15:11] for configuration writes, whether it is in PCI or PCI-X mode. When a PCI configuration transaction occurs, the Device Number bits[7:3] are updated to a value of 00000b from the default value of 11111b. The bus number bits[15:8] are grabbed only during the attribute phase (which does not exist for PCI).

Status: No Fix. See the [Table , “Summary Table of Changes” on page 7.](#)

12. M_RST# driven to DDR-II or DDR-I voltage levels

Issue: The de-asserted voltage level on M_RST# with DDR-II is 1.8 V and with DDR-I is 2.5 V. When M_RST# is needed for other devices (for example, Flash), make sure these voltage levels are appropriate for the target device.

Status: No Fix. See the [Table , “Summary Table of Changes” on page 7.](#)

13. BIU master abort causes two interrupts on reads

Issue: The Bus Interface Unit (BIU) will generate an interrupt when the BIU gets master aborted on an internal bus write. The functionality is different between a read and write case. For a write, the BIU master abort asserts IINTSRC[29] (when enabled) and does not assert an error to the core. For a read, the BIU master abort asserts both IINTSRC[29] (when enabled) and an error to the core. Therefore, on a read case two interrupts is generated. When the interrupt source is masked, then only the master abort on reads is detected, and this is from the direct core error.

Status: No Fix. See the [Table , “Summary Table of Changes” on page 7.](#)

14. Potential race condition with Interrupt Controller Unit status bits

Issue: There is a slight lag in the time it takes between clearing a status bit inside the unit and the corresponding bit in the Interrupt Controller Unit Status Register getting cleared. This has the potential of generating a false interrupt, meaning that the Intel XScale® processor is interrupted, but the handler is not able to find any source reported in the ICU registers. This condition can be avoided by adding a read from any ICU register, after the bit is cleared in the local unit before returning from the interrupt handler. The data from this read can be ignored, but the read itself creates enough latency to allow the updated status to propagate to the ICU.

Status: No Fix. See the Table , “Summary Table of Changes” on page 7.

15. Reset Internal Bus (PCSR[5]) usage

Issue: PCI bus data corruption can occur when the reset internal bus bit is not used correctly. PCSR[5] can be used to reset the internal bus. This resets all memory-mapped registers and the ATU configuration header space. This bit is read/write-capable from both the internal bus and PCI bus.

For both PCI and PCI-X modes, make sure the following steps occur when setting PCSR[5]:

1. Clear Bus Master (ATUCMD[2]) Enable and Memory Enable (ATUCMD[1]).
2. Wait for both the outbound (PCSR[15]) and inbound (PCSR[14]) read transaction queue busy bits to clear.
3. Make sure no PCI configuration read/write cycles targeting the ATU are in progress, except the reset configuration write, when applicable.
4. Set the Reset Internal Bus bit.
5. Wait for 40 PCI clocks.

Status: No Fix. See the Table , “Summary Table of Changes” on page 7.

16. PCI Express* Transaction Header Log register repeats on offset 124 and 128

Issue: The lower 32 bits of the address of the Single Address Cycle (SAC) are mirrored onto the upper 32 bits in the log register giving the appearance that a Dual Address Cycle (DAC, 64-bit address transaction) occurred instead of a SAC (Single Address Cycle, 32-bit address transaction).

The lower 32 bits of the address are on the third DWord and the upper are on the fourth DWord.

The fourth DWord is “reserved” for a SAC transaction. So, instead of the fourth DWord defaulting to 0 as expected, it actually mirrors the third DWord.

For example: For a 32-bit TLP (memory write with poisoned bit set) downstream transaction to the 80333, the header log (offset 11C–12B) could look like this:

Offset: 0x11C, Before: 0x00000000, After: 0x40004005
Offset: 0x120, Before: 0x00000000, After: 0x000000FF
Offset: 0x124, Before: 0x00000000, After: 0x06000000
Offset: 0x128, Before: 0x00000000, After: 0x06000000

Status: No Fix. See the Table , “Summary Table of Changes” on page 7.

17. SHPC sequence

Issue: When using parallel 1-slot, no-glue Hot-Plug mode, the frequency does not update correctly with a single reset; therefore, two B-segment resets are required. The first reset is needed to correctly latch in the B_PCIXCAP and B_M66EN, and the second reset is needed to execute the change frequency command. This sequence must be followed:

1. PWRGD 0 -> 1 (HW)
2. B_RST# 0 -> 1 (HW), at this time HPRST# is 0
3. Enable Slot Power (SW normal sequence)
4. Issue Hot-Plug change frequency with code 40 (PCI = 33M).
5. HPRST# 0 -> 1 (SW slot enable)
6. HPRST# = 1 (change frequency) [Note: this is a software step. Bus reset happens automatically due to the change frequency.]
...operating...
7. HPRST# 1 -> 0 (SW slot disable)
...off/removal...
...new card inserted...
8. Go to step 2

Status: [No Fix](#). See the [Table , “Summary Table of Changes”](#) on page 7.

18. Bus Interface Unit follows PCI ordering rules

Issue: The Core Bus Interface Unit orders transactions based on PCI rules. This allows outgoing writes to pass incoming reads. For most devices on the internal bus, this does not cause problems since the devices function asynchronously with respect to each other. For transactions between the Intel XScale® processor and memory via the internal bus, this can result in unexpected data. For example:

```
0: ldr r0, =0x40000
1: ldr r1, =0xaaaaaaaa
2: ldr r2, =0x55555555
3: str r1, [r0]
4: ... <time-delay to allow the previous transactions to complete>
5: ldr r3, [r0]
6: str r2, [r0]
```

This code can potentially load the register r3 with the value 0x55555555, since the store in line 6 might pass the load in line 5. This happens only with uncached transactions on the internal bus.

The MCU core port enforces strict ordering and does not exhibit this behavior. When the MCU core port is used, this issue does not occur.

When caching is enabled, the initial read initiates a cache-line fill. The subsequent write is pended in the Intel XScale® processor until the line fill and the ldr instruction complete. In this case, this issue does not occur.

When caching is disabled, and the caching policy is stall-until-complete ($X = 0, C = 0, B = 0$), this issue does not occur. For other MMU settings with caching disabled, the issue can occur. Specifically regions with data cache and write buffer policies of bufferable ($X = 0, C = 0, B = 1$) or coalescing-disabled-bufferable ($X = 1, C = 0, B = 1$) are vulnerable to this issue. In addition, regions configured as write through ($X = 0, C = 1, B = 0$) are also vulnerable to this issue.

In addition, when caching is enabled in the MMU page tables, but the DCache is disabled in the CP15 ARM Control Register, then the effective caching policy is bufferable and this reordering must be accounted for.

It is important to realize that any code that accesses memory spaces on the internal bus must account for this possibility. Code that dynamically disables cache (for example, Flash programming routines) must ensure that the caching policy for the appropriate memory region is set to “stall until complete” until the cache is re-enabled.

The simplest scenario to reproduce is a pair of back-to-back function calls, for example:

```
main:
    bl    fun1
    bl    fun2
    ...
    ...
    ...

fun1:
    stmfd sp!, {r4, r5, r6, r7, r8, r9, r10, lr}
    ldmfd sp!, {r4, r5, r6, r7, r8, r9, r10, pc}

fun2:
    stmfd sp!, {r4, fp, ip, lr}
```

```
ldmfd    sp!, {r4, fp, ip, pc}
```

Status: No Fix. See the [Table](#) , “Summary Table of Changes” on page 7.

19. **UART, I²C and GPIO memory mapped registers should be addressed with 32-bit accesses**

Issue: The UART, I²C and GPIO units sit on a dedicated low-speed internal bus that does not support byte enables. Due to this functionality, accessing any of these unit memory-mapped registers (MMR) with any accesses less than 32-bits can result in corruption of the other bits in the 32-bit MMR. For example, the GPOD register (located at FFFF_F788h) implements functionality for bit[10] and bit[11]. When software performs a byte access to GPOD, this could cause bit[10] and bit[11] to be written with incorrect data.

While most of these registers implement only the lower 8 bits (the upper three bytes are “reserved”), the recommendation is that all UART, I²C, and GPIO MMRs must be accessed only as 32-bit registers. While it is desired that 32-bit accesses be performed, it is acceptable to access with less than 32 bits, as long as all non-reserved bits are accessed. For purposes of future expansion, 32-bit accesses are preferred.

Status: No Fix. See the [Table](#) , “Summary Table of Changes” on page 7.

20. **UART Interrupt Identification Register**

Issue: The UART Interrupt Identification Register (UxIIR) is read by software to determine the type and source of UART interrupts. This register gathers and priority encodes the various sources of UART interrupts. The register is read after an interrupt occurs. Enabling and disabling of interrupts (via the Interrupt Enable Register [UxIER] or the Modem Control Register [UxMCR]) affects whether or not the interrupt to the processor occurs. This does not effect the logging of the status of what is happening in the UART. The UART operates in interrupt or polling mode. In polling mode, all interrupts to the processor are disabled.

Status: No Fix. See the [Table](#) , “Summary Table of Changes” on page 7.

21. **Reads on 16-bit PBI bus operate as 32-bit**

Issue: Two-byte and four-byte read transactions on the Peripheral Bus Interface (PBI) bus operate as burst reads (in other words, two 16-bit read cycles). All the read transactions from the Intel XScale[®] processor to PBI devices (such as SRAM, flash, and so on) are translated to burst reads with burst size of 2, even though there is no necessity to generate a burst transaction. Therefore, devices on the 16-bit PBI bus must be configured as pre-fetchable.

Note: Single-byte transactions on 16-bit PBI bus is not a supported case. Also, a PBI bus configured as 8-bit does not operate this way.

Status: No Fix. See the [Table](#) , “Summary Table of Changes” on page 7.

22. **3.3 V to 1.5 V leakage**

Issue: There is a leakage path from the 3.3 V rail to the 1.5 V rail. When the 3.3 V is powered-on and the 1.5 V is not, then ~500 mV is seen on the 1.5 V rail. This leakage is expected and does not cause any long-term reliability issues.

For related issues, see Non-Core Erratum 10 (“Secondary bus PCI RST# pulse prior to the rising edge of PWRGD” on page 17).

Status: No Fix. See the [Table](#) , “Summary Table of Changes” on page 7.

23. Accessing extended bridge configuration space from the Intel XScale® processor

Issue: In certain cases, the Intel XScale® processor might need to access the extended configuration space of the bridge headers (for example, see Non-Core Erratum 22, “[Incorrect default value for PCI Express* Flow Control Protocol Error Severity bit](#)” on page 21). When the Intel XScale® processor issues the configuration cycle address using OCCAR, bits[27:24] can be used for addressing the extended space. For example, to access offset 10Ch of A-bridge configuration space, use an address of 0x0101.000C for a type-zero configuration cycle.

Status: No Fix. See the [Table](#) , “Summary Table of Changes” on page 7.

24. Recommended B-segment termination when using the 80333 on PCI Express* adapter cards

Issue: For PCI Express* adapter cards, the B-segment is not used. Assuming PCIODTEN = 1 to enable internal pull-ups, the following external connections must be made for an unused B-segment:

1. Connect B_CLKOUT to B_CLKIN. The other clockouts can be turned off by software (PCI Clock Control register at offset 43 in the B-segment bridge configuration space).
2. For B_PME#, use an 8.2 K Ω pull-up.
3. For B_PCIXCAP, use a 3.3 K Ω pull-up.
4. For B_RCOMP, use 100 Ω to ground.

Status: No Fix. See the [Table](#) , “Summary Table of Changes” on page 7.

25. B-segment arbiter might not park on the last master when in PCI-X 133 MHz mode

Issue: When the B-segment is running in PCI-X 133 MHz mode, the arbiter might not park on the last master, when bit[8] of the Internal Arbiter Control register (ARB_CNTRL[8], offset 16Ah) is cleared.

This happens only when the B-segment is running at 133 MHz. When running at 66 MHz or 100 MHz, the arbiter parks the bus on the last master when ARB_CNTRL[8] is cleared. The A-segment does not exhibit this behavior with the same setup.

Status: No Fix. See the [Table](#) , “Summary Table of Changes” on page 7.

26. Configuring XINTx# signals as PCI interrupts

Issue: The *PCI Specification Rev. 2.3* states that PCI interrupts are defined as “level sensitive” and “active low”.

To configure the XINTx# signals of the 80333 as “level sensitive”, APIC_RDLxx[15] (Trigger Mode) must be changed to 1 from its default of 0.

To configure the XINTx# signals of the 80333 as “active low”, APIC_RDLxx[13] (Interrupt Input Pin Polarity) must be changed to 1 from its default of 0.

Status: No Fix. See the [Table](#) , “Summary Table of Changes” on page 7.

27. Power plane isolation for Battery Back-Up (BBU) mode

Issue: During battery back-up (BBU) mode, when the battery powers the DIMM and the VCC25/18 signals (1.8 V or 2.5 V, depending on the memory type being used) on a single power plane, the battery life will probably be reduced due to leakage.

To attain longer battery life, the DIMM and VCC25/18 power planes must be isolated. The power plane isolation can be accomplished by using a FET.

Status: **No Fix.** See the [Table , “Summary Table of Changes” on page 7.](#)

28. AAU result can be written directly to PCI host memory

Issue: The Application Accelerator Unit (AAU) can write results not only to local memory but also to the PCI bus host memory via the ATU.

This feature can be applied to degraded RAID-5 reads, where the AAU result is the reconstructed data for the host I/O read. The AAU can write its result to PCI; therefore, the degraded read XOR result can be written directly to host memory. This eliminates the need for a DMA operation to transfer the result from local memory to host memory via PCI.

Savings for the RAID application include the following:

- No DMA descriptor needs to be generated.
- No DMA interrupt needs to be serviced.
- Memory and internal bus bandwidth is saved (result write by AAU and read by DMA).
- Read I/O is serviced faster (eliminates latency of DMA operation).

Note: AAU source reads from PCI are not supported; only local memory can be used for this.

Status: **No Fix.** See the [Table , “Summary Table of Changes” on page 7.](#)

29. SMBus connection recommendations for PCI Express* adapter cards

Issue: PCI Express* cards based on the 80333 must implement the SMBus signals in one of the following ways:

1. The SMDAT and SMCLK signals from the PCI Express* connector must be left as “no connects”. The SCLK and SDTA signals on the 80333 must have pull-ups even when they are not used. The pull-ups prevent the inputs from oscillating and potentially causing other problems.
2. When the SMBus feature is required, an isolation device (for example, the LTC4301) must be placed between the SMBus signals on the PCI Express* connector and the 80333, so that the system has no connection to the 80333 on these two signals when power is off.

For motherboard designs, it is assumed that the SMBus is routed only to devices that are required and that remain powered.

Status: **No Fix.** See the [Table , “Summary Table of Changes” on page 7.](#)

30. PBI lockout condition

Issue: When the core is in a tight loop writing to the PBI bus, while the DMA is doing a large block transfer (for example, from SRAM, located on the PBI, to DDR memory), the DMA can be locked out of accessing the PBI and the transaction will never complete.

If this condition occurs, use one of these workarounds:

1. Change the MTTR1 from 98h (default) to a lower value (such as 01h). A lower value allows the DMA (or ATU which can also master a transaction to the PBI) to gain access to the PBI, because the BIU is given shorter access for back-to-back BIU internal bus transactions.
2. Add a core read along with the core write, causing it stall and preventing it from starving the DMA.
3. Add NOPs or dummy instructions to ensure the loop spans greater than two cache lines.
4. Modify the loop such that the write is not done on every iteration.

Status: No Fix. See the [Table](#) , “Summary Table of Changes” on page 7.

31. PFREQ functionality

Issue: The PFREQ bits (BCNF[10:9], offset 40h) cannot be used to change the actual PCI bus frequency. The only way to change the PCI bus frequency is to reset the bus or use the hot-plug controller.

Note that this is an artifact of specification clarification 17 (“SHPC sequence” on page 37). The same mechanism that requires two resets with the hot-plug controller, prevents the 33 MHz to/from 66 MHz PCI transition without hot plug. Other transitions (such as PCI-X) do work correctly; it is only the PCI33 that cannot change to or from another bus speed.

Status: No Fix. See the [Table](#) , “Summary Table of Changes” on page 7.

32. PWRDELAY functionality during power sequencing

Issue: When the 3.3 V rail is powered on and the 1.5 V rail is powered off, the PWRDELAY input signal drives out until the 1.5 V rail powers up. This is important to understand if still using the legacy power-fail circuit, because it might cause other circuitry to function incorrectly. The proper usage of PWRDELAY is described in Specification Clarification 35 (“PWRDELAY needs only a pull-up for battery back-up mode” on page 43), which recommends that PWRDELAY be isolated from all circuitry and tied to a 1.5 K Ω pull-up resistor.

Status: No Fix. See the [Table](#) , “Summary Table of Changes” on page 7.

33. HPI# (High Priority Interrupt) is a maskable interrupt

Issue: The HPI# interrupt input is both maskable and masked by default (as are all interrupts). It is controlled by INTCTL1[31]. HPI# operates the same as the other external interrupt inputs (XINT[7:0]#).

Status: No Fix. See the [Table](#) , “Summary Table of Changes” on page 7.

34. OCD and Receive Enable calibration de-featured

Issue: The ability to adjust the electrical interface to account for out-of-specification DDR-II DIMMs using OCD (off-chip driver) and receive enable calibration, is no longer a supported feature.

Status: No Fix. See the [Table](#) , “Summary Table of Changes” on page 7.

35. PWRDELAY needs only a pull-up for battery back-up mode

Issue: The I²C bus is no longer included in the processor reset equation. This allows the PWRDELAY circuit to be simplified to a single pull-up resistor, to enable battery back-up mode. PWRDELAY must be isolated from all other circuitry, and only a 1.5 K Ω pull-up to 3.3 V is required. When battery back-up is not required, PWRDELAY must have a 1.5 K Ω pull-down resistor.

Status: No Fix. See the Table , “Summary Table of Changes” on page 7.

36. PCI Express* L0s functionality not supported in the 80333

Issue: Due to the unreliable behavior described in Non-core Erratum 16 (“Unreliable PCI Express* link operation when L0s active state power management is enabled” on page 19), L0s active power management state is not supported. The BIOS must be updated to leave L0s disabled on the 80333 by default, and eliminate any setup option with which to enable it. Additionally, the BIOS must not allow L0s to be enabled on any PCI Express* root ports connected to the 80333 end-devices.

Status: No Fix. See the Table , “Summary Table of Changes” on page 7.

37. DDRRES2 can be pulled down to reduce current during self-refresh

Issue: DDRRES2 is used as compensation for DDR-II OCD. Since OCD is not supported in the 80333 (see Specification Clarification 34, “OCD and Receive Enable calibration de-featured” on page 42), DDRRES2 can be pulled down to reduce current draw during self-refresh mode.

Status: No Fix. See the Table , “Summary Table of Changes” on page 7.

38. DDRSLWCRES resistor values

Issue: The following resistors can be used on DDRSLWCRES:

- DDR-I memory: 845 Ω or 1.13 K Ω resistor, 1%
- DDR-II memory: 825 Ω or 976 Ω resistor, 1%

Status: No Fix. See the Table , “Summary Table of Changes” on page 7.

39. B_PME# routing recommendation when using Parallel Hot Plug 1-slot, no-glue mode

Issue: When the B-segment is in parallel hot plug, one-slot, no-glue mode, the B_PME# signal might be pulled to ground. Since B_PME# is active low and might be connected to the GIPx pin of ICHx, the ICHx interprets the low voltage of B_PME# as a power-management event from the PCI-X slot. This might cause the system not to remain in the appropriate power mode when no add-in card is plugged into the slot. B_PME# must be high when there is no PME request from the PCI-X slot. To work around this issue, add a diode between the buffer and the 80333, and add a 1 K Ω –10 K Ω pull-up resistor to the B_PME# signal. The diode can prevent ICHx from getting a B_PME# low logic level when there is no PME request from the PCI-X slot.

Status: No Fix. See the Table , “Summary Table of Changes” on page 7.

40. Multi-Transaction Timer grants fewer clocks in PCI mode than expected

Issue: The Multi-Transaction Timer (MTT, offset 42), in the bridge configuration header of the 80333, specifies the number of clocks that GNT# asserts to a master device on the PCI bus. When running in PCI-X mode, the number of clocks granted matches the value programmed in the MTT, but when in PCI mode, the master is granted 16 clocks fewer than the value programmed in the MTT. The number of GNT# cycles is measured from the assertion of FRAME# to the deassertion of GNT# with REQ# asserted during the whole time. This behavior occurs on both the A- and B-bus segments and at both PCI frequencies (33 MHz and 66 MHz). For example, when the MTT value is 0x18, it produces 0x8 clocks between the assertion of FRAME# and deassertion of GNT#. Another example: when the MTT value is 0x30, it produces 0x20 clocks between the assertion of FRAME# and deassertion of GNT#.

Status: No Fix. See the Table , “Summary Table of Changes” on page 7.

41. Byte Enables (BE) not included in PCI delayed reads can cause data corruption

Issue: A PCI device on one of the secondary busses that generates a zero-length read request can cause data corruption in platforms utilizing non-Intel® MCH components, as the byte enables (BE) are not included by the 80333 bridge in matching completions to PCI delayed read requests. All Intel® MCH devices return data consistent with the address of a zero-length read request. Thus, no corruption can occur when a subsequent non-zero-length read is inadvertently completed with data returned on behalf of the zero-length request, but this behavior is not required by specification.

The following is an example case:

1. A memory read request with zero BE is issued over PCI; a corresponding zero-length read results on PCI Express* to the host.
2. A PCI device on the same PCI segment issues a MR/MRL/MRM to the same address with valid BEs.
3. The 80333 bridge matches the completion for the Memory Read request on line 1 to the request on line 2 (in other words, BEs are ignored).
4. Unspecified data (returned for the zero-length read request) is driven to the PCI card, resulting in data corruption.

Note: This affects PCI mode exclusively, and is not an issue when the secondary busses are operating in PCI-X mode. Whether corruption can occur through this mechanism is dependent upon the behavior of the non-Intel® MCH component. If the MCH in use behaves similarly to Intel® MCH designs, there is no exposure to data corruption, and the incomplete completion match does not have any side effects.

Status: No Fix. See the Table , “Summary Table of Changes” on page 7.

42. Interleaving AAU descriptors

Issue: The P+Q capability is enabled in the AAU globally (ACR.3), not on a descriptor by descriptor basis. Therefore if enabled, all descriptors are processed as P+Q descriptor formats. If disabled, all descriptors are processed as prior AAU definitions (ie - straight XOR). In order to mix RAID-5 with P+Q RAID-6, enable P+Q RAID-6 GF Multiply for the AAU, and build all RAID-5 XOR descriptors as P+Q RAID-6 descriptors, where the GF Multiplier Byte values are all 0x01.

Status: No Fix. See the Table , “Summary Table of Changes” on page 7.

43. RCVDLY setting for DDR-I memory

Issue: The Receive Enable Delay Register (RCVDLY at FFFF_F550h) is used for DQS receive enable calibration. In other words, RCVDLY adjusts the memory controllers relationship of DQS to an internal M_CLK.

The RCVDLY value is highly dependant on the board layout and DIMM characteristics. Also, the memory controller only supports a non integer CAS latency (tCAS = 2.5, SDCR0.9:8) for DDR-I, which means that RCVDLY may need to be adjusted because DQS is no longer synchronized with M_CLK.

Therefore, when using DDR-I memory the RCVDLY default setting of 5, may need to be changed to 6 or 7 to operate correctly with a specific DIMM based on the board layout. For example, the Redboot reference code provided by Intel uses a value to 7 in order to allow for a wider compatibility with various DIMMs.

Status: **No Fix.** See the [Table , “Summary Table of Changes” on page 7.](#)

44. Embedded Usage Models

Issue: The 80333 I/O Processor was designed to be used as a PCI-Express endpoint. The PCI-Express interface on the 80333 does not have root complex support, therefore it cannot be used in an embedded application unless there is an upstream root complex that can control the 80333 PCI-Express port.

Also, data transfers between the A-segment PCI bus and B-segment PCI bus is not a supported configuration. These two PCI busses are independent and only bridge to the PCI-Express interface.

Status: **No Fix.** See the [Table , “Summary Table of Changes” on page 7.](#)

45. ATUBAR3 Functionality

Issue: The private memory window of the PCI A-segment defines an address range that the 80333 uses to map private devices to, and to locate local memory for private device access. This range is intended to be mapped to the ATUs private BAR window (ATUBAR3) and the private device BARs. Note that even when the private addressing is enabled, the normal 80333 behavior defined for BME, MSE, IOSE bits in the ATUCMD register are still true. Therefore, when the ATU Memory Space Enable bit is cleared, all ATU BARs including ATUBAR3 will be unable to claim any memory transactions. For example, this bit is typically cleared during a PCI bus scan / enumeration.

Status: **No Fix.** See the [Table , “Summary Table of Changes” on page 7.](#)

46. VREF isolation for Battery Back-up (BBU) mode

Issue: During battery back-up (BBU) mode, the DIMM power can be isolated from the 80333 IOP power. This isolation should also include the VREF signal for the DIMM interface. Due to leakage, the VREF signal for the DIMM should be isolated from the VREF signal for the IOP. This is to ensure that VREF for the DIMM is not disturbed as the IOP powers down when entering battery back-up (BBU) mode. The isolation can be provided by using separate voltage dividers or a FET.

For related issues, refer to Specification Clarification [27](#), “[Power plane isolation for Battery Back-Up \(BBU\) mode](#)” on page 41.

47. I2C Unit Enabling

Issue: Software must guarantee that the I2C bus is idle before enabling the I2C unit. Failure to do so could result in unstable behavior. The IBMR register can be used to monitor the state of the SCL and SDA pins in order to determine bus activity. The I2C Bus Busy bit in the I2C Status Register (ISR.3) can not be used for this purpose, as it is only valid when the I2C unit is enabled.

Status: **No Fix.** See the [Table , “Summary Table of Changes” on page 7](#)

48. DMA transactions from local memory to a conventional PCI target can complete out of order

Issue: It is possible for the Outbound ATU to get backed up and retry the DMA unit. When this occurs, the possibility exists for DMA transactions to complete out of order as each DMA has 3 separate data queues and there is no guarantee as to which one of the queues will drain on the retry.

Status: No Fix. See the Table , “Summary Table of Changes” on page 7

49. SBR1 Programming with Bank 1 Unpopulated

Issue: When using single-banked DDR-II DIMMs and the SDRAM Boundary Register 1 (SBR1; FFFF_E514h) is not programmed to match the SDRAM Boundary Register 0 (SBR0; FFFF_E510h), the wrong ODT signal will become activated. The memory controller provides two ODT (On Die Termination) signals (ODT[1:0]) which are used with DDR-II DIMMs to turn on termination during writes.

Section 8.7.6 in the *Intel® 80333 I/O Processor Developer’s Manual* (305432) states, “If bank 1 is unpopulated, SBR1[6:0] is programmed either with all zeroes or a value equal to SBR0[6:0].” To clarify this statement for single-banked DDR-II DIMMs, if bank 1 is unpopulated, then the entire SBR1 must be programmed the same as SBR0. This includes lower bits 06:00 and upper bits 31:30. Bits 30:07 are reserved and should not be written. The memory controller compares the entire range of the SBR0 and SBR1 to determine which ODT signal to enable. If the upper bits 31:30 in SBR0 and SBR1 do not match, then ODT1 will become active instead of ODT0.

In addition, when bank 1 is unpopulated, SBR1[6:0] should never be zero if SBR0[6:0] is non-zero.

Status: No Fix. See the Table , “Summary Table of Changes” on page 7

50. 32-bit Writes to Unaligned 64-bit Addresses are Promoted to 64-bit Aligned Writes

Issue: In applications that run the PCI bus segment in 32-bit PCI Mode or 64-bit PCI Mode with 32-bit targets, write transactions that are on unaligned 64-bit addresses are promoted to 64-bit aligned writes. The first half of the 64-bit write is on a 64-bit aligned address and has the BE# signals disabled. Therefore, the write is invalid. The second half on the 64-bit write is a valid write with the BE# enabled and the write is to the intended 32-bit address.

Per the PCI Local Bus Specification, Revision 2.2, PCI compliant devices should ignore the first half of the 64-bit write due to the BE# signals being disabled.

For devices that support using the I/O window, the 64-bit write does not occur when using the ATU I/O Window and only the expected 32-bit write occurs.

For memory mapped devices, the only option is to run in PCI-X mode, where the byte count and starting address are consistent with the actual number of bytes to be written (i.e., 4). When a 64-bit PCI-X request gets downshifted, the requester can use the starting address/byte count to recognize that the write request does not cross a DWORD address boundary and only perform a single 32-bit wide data cycle.

Status: No Fix. See the Table , “Summary Table of Changes” on page 7

51. ATU Retry Response Through the Bridge

Issue: The bridge will not pass ATU retry responses upstream through the bridge unless bit 15 in the PCI Express Device Control Register (EXP_DCTL) register, located at offset 0x4c, is set to '1'.

Status: No Fix. See the Table , “Summary Table of Changes” on page 7

52. Case Temperature Clarification

Issue: Internal models indicate that certain elevated case temperatures (Tcase) of the Intel® 80333 I/O processor may cause elevated field failures in later years of operation. This clarification is precautionary, as Intel has not seen any failures of the 80333 products in the field. Internal modeling data indicates that the degree of failure risk decreases with lower operating frequency (i.e. – 500, 667, or 800MHz) and temperature (i.e. – case temperature). If a failure should occur, it would manifest itself as a hard failure of the I/O processor which can cause a system failure.

A detailed analysis of the system operating conditions, specifically case temperature of the 80333 I/O processor, is required. The 80333 should be run with Tcase temperatures that, over the lifetime of the part, average less than the 95° C maximum Tcase that is currently specified in the *Intel® 80333 I/O Processor Datasheet* (305433).

Tcase recommendations for the 80333:

Frequency	5 year operating life	6 year operating life	7 year operating life
500MHz	95° C	93° C	91° C
667MHz	85° C	83° C	80° C
800MHz	78° C	76° C	74° C

Note: Tcase temperatures in the table reflect an average Tcase temperature over the lifetime of the product.

Status: Fixed. This issue was fixed in the A-1 stepping of the product. See the [Table , “Summary Table of Changes” on page 7](#).

Documentation Changes

1. PCI clock timings table missing note

Issue: Table 25 needs a note added referencing support of class 2 clock jitter.

Workaround: A fourth note has been added to Table 25 and appears as follows:

4 - Clock jitter class 2, per PCI-X Electrical and Mechanical Rev 2.0a specification

Affected Docs: Intel® 80333 I/O Processor Datasheet (305433-001)

2. Wrong Voltage Values in Table 23

Issue: Table 23 shows wrong voltage values.

Workaround: Replaced two rows in Table 23. The two rows now appear as follows:

V _{OL1}	Output Low Voltage (DDR SDRAM)		0.35	V	I _{OL} = 12.5 mA (1, 2)
V _{OH1}	Output High Voltage (DDR SDRAM)	1.95		V	I _{OH} = -12.5 mA (1, 2)

Affected Docs: Intel® 80333 I/O Processor Datasheet (305433-002).

3. SBR1 Programming When Bank 1 is Unpopulated

Issue: Section 8.7.6 incorrectly states: “If bank 1 is unpopulated, SBR1[6:0] is programmed either with all zeroes or a value equal to SBR0[6:0].”

Workaround: The sentence should be changed to “If bank 1 is unpopulated, SBR1[6:0] and SBR1[31:30] should be programmed with a value equal to SBR0[6:0] and SBR0[31:30].”

Affected Docs: Intel® 80333 I/O Processor Developer’s Manual (305432-001)

4. PCI Express clock cycle time minimum

Issue: Table 27 in the 80333 datasheet shows the PCI Express clock cycle time, Tc2, as 10ns minimum.

Workaround: Since the 80333 is compliant to PCI-Express 1.0a specification, this should be changed to 9.872ns minimum.

Affected Docs: Intel® 80333 I/O Processor Datasheet (305433-003)

