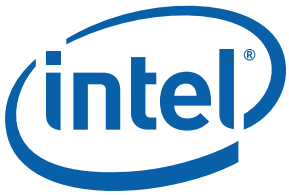

Intel® 5400 Chipset Memory Controller Hub (MCH)

Specification Update

October 2009

Notice: The Intel® 5400 chipset MCH may contain design defects or errors known as errata which may cause the product to deviate from published specifications. Current characterized errata are available on request.

Order Number: 318638-006

Contents

§

1 Revision History

Date	Revision	Description
November 2007	-001	Initial Release
November 2007	-002	Added Erratum upto #25
November 2007	-003	Added Erratum #26 and #27
October 2009	-004	Added Erratum #28
March 2009	-005	Added Documentation Change #1 and #2
October 2009	-006	Added Errata #29 and #30

§

2 Preface

This document is an update to the Intel® 5400 chipset specification 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.

2.1 Affected Documents/Related Documents

Document Title	Reference Number
Intel® Xeon® 5400 Chipset Memory Controller Hub (MCH) Datasheet	318610
Intel® 5400 Chipset Memory Controller Hub (MCH) Thermal/Mechanical Design Guidelines	318639

2.2 Nomenclature

Errata are design defects or errors. These may cause the MCH behavior to deviate from published specifications. Hardware and software designed to be used with any given stepping must assume that all errata documented for that stepping are present on all devices.

S-Spec Number is a five-digit code used to identify production products. Products are differentiated by their unique characteristics,e.g., stepping, speed, package type, etc. as described in the identification information table. Read all notes associated with each S-Spec number.

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 lifecycle, 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.).

§

3 Summary Table of Changes

The following table indicates the errata, specification changes, specification clarifications, or documentation changes which apply to the MCH. 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 Tables

Stepping

X:	Erratum exists in the stepping indicated. Specification Change or Clarification that applies to this stepping.
(No mark) or (Blank box):	This erratum is fixed in listed stepping or specification change does not apply to listed stepping.

Status

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

Row

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

3.1 Errata

Table 3-1. Errata

Number	Stepping	Status	Erratum
	C0		
1	X	No Fix	Error logging logic (x16 mode) may not indicate the first uncorrectable error in the First Error register
2	X	No Fix	MCH CPU reset may cause spurious errors in CPU to be logged
3	X	No Fix	FBD IBIST pattbuf2enable register does not work for normal operation
4	X	No Fix	FBD IBIST operation does not stop when IBSTR is cleared in Continuous Mode
5	X	No Fix	Simultaneous errors of the same type in both FBD branches result in only Branch 0 FERR being updated
6	X	No Fix	MCH will incorrectly direct all MMCFG accesses to the PCI-Express MMCFG space if HECBASE is enabled with lengths not aligned on 256MB boundaries
7	X	No Fix	Clearing Bus Master Enable (BME) Can Generate Address Errors
8	X	No Fix	When VT-d is enabled and the Guest Physical Address (GPA) increments beyond FFFF_FFFF_FFFF, an address error is not flagged in the CHANERR register
9	X	No Fix	FBD Chan_Indx set incorrectly for single channel mode
10	X	No Fix	Incorrect JTAG scan chain error log
11	X	No Fix	PCI Express: MCH not meeting Receiver high impedance Spec
12	X	No Fix	MCH does not support PCI Express 2.0 compliance mode features
13	X	No Fix	Store to Write-Through (WT) Memory Data may be seen in wrong order by two subsequent loads when snoop filter is enabled
14	X	No Fix	Link Bandwidth Management Status set for autonomous bandwidth change initiated by downstream device
15	X	No Fix	PCI Express: PCI Express LTSSM does not go REC.SPEED from REC.RCVRCFG on receiving valid TS sequences only on some lanes with speed change bit set
16	X	No Fix	PCI Express data corruption in presence of bad and good completion in the same clock
17	X	No Fix	Peer-to-Peer traffic may cause completion timeouts, primarily in conjunction with PCI Express, FBD, or FSB throttling
18	X	No Fix	PCI Express: Electrical Idle inference priority mode causes Polling.Compliance load board failure
19	X	No Fix	MCH IOxAPIC PCISTS.DPE not cleared correctly
20	X	No Fix	CRC errors logged by AMB during C2 or S1 transition
21	X	No Fix	All OB Cfg type transactions are allowed during D3-hot
22	X	No Fix	System hang with large number of transaction retries
23	X	No Fix	Low and High Protected Memory Regions in Intel (R) Virtualization Technology for Directed I/O (VT-d) are not supported
24	X	No Fix	PCI Express Differential Tx Voltage swing exceeds spec limits at -6.0dB
25	X	No Fix	MCH FBD TX compliance Eye Height in violation at 4.8Gb/s
26	X	No Fix	MCH logs correctable errors when L0s is enabled and when the end point has short EI times (<224ns for x16)
27	X	No Fix	MCH can timeout in Recovery.Rcvrcfg on L1 exit with PCI Express Gen2 endpoint device that has short L1.idle times
28	X	No Fix	Read Transactions may be delayed
29	X	No Fix	PCI express TLP packets not flagged “Malformed” under certain conditions.
30	X	No Fix	Error Source Identification (ID) is not properly reporting the Requestor ID when the uncorrectable (Non-fatal/fatal) error is detected in the PCI express Root Port

3.2 Specification Changes

Number	SPECIFICATION CHANGES
N/A	None

3.3 Specification Clarifications

Number	SPECIFICATION CLARIFICATIONS
N/A	None

3.4 Documentation Changes

Number	DOCUMENTATION CHANGES
1	Device ID (DID) Correction
2	Intel 5400 Chipset inadvertently written as Intel 54000 Chipset

§

4 Identification Information

4.1 Component Identification via Programming Interface

The Intel® 5400 Chipset can be identified by the following register contents:

MCH Version	Stepping	Vendor ID[1]	Device ID[2]	Revision Number[3]
5400A (1333 FSB)	C-0	8086h	4001h	20h
5400B (1600 FSB)	C-0	8086h	4003h	20h

Notes:
1. The Vendor ID corresponds to bits 15:0 of the Vendor ID Register located at offset 00 - 01h in the PCI function 0 configuration space.
2. The Device ID corresponds to bits 15:0 of the Device ID Register located at offset 02 - 03h in the PCI function 0 configuration space.
3. The Revision Number corresponds to bits 7:0 of the Revision ID Register located at offset 08h in the PCI function 0 configuration space.

4.2 Component Marking Information

The Intel® 5400 Chipset can be identified by the following component markings:

MCH	Stepping	S-Spec/QDF	Top Marking	Notes
5400A Pb-free (1333 FSB)	C-0	QP35	NU82014MCH	Product Sample
5400B Pb-free (1600 FSB)	C-0	QP36	NU82014MCH	Product Sample
5400A Pb-free (1333 FSB)	C-0	SLARD	82014MCH	Production
5400B Pb-free (1600 FSB)	C-0	SLARE	82014MCH	Production

The Intel® 5400 Chipset stepping can be identified by the following component markings:

Figure 4-1. Top-Side Marking Example

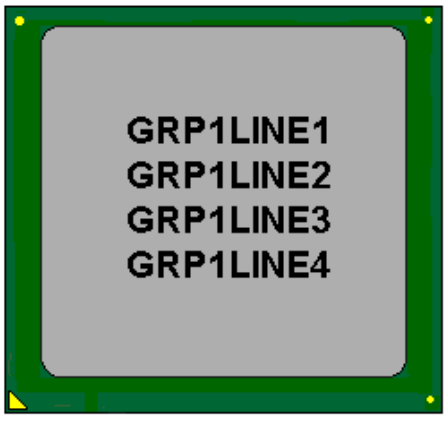

§

5 Errata

1. Error logging logic (x16 mode) may not indicate the first uncorrectable error in the First Error register

Problem: If there are two back to back TLP with uncorrectable errors, then the pointer will point to the second uncorrectable error instead of the first one.

Implication: MCH PCI Express ports 1 and 5, when operating as Gen2 x16, may indicate the wrong first uncorrectable error in the First Error Pointer Field of the AERRCAPCTRL register.

Workaround: Ignore the first error pointer in AERRCAPCTRL. Read the header logging in HDRLOG register for the correct error information.

Status: No Fix

2. MCH CPU reset may cause spurious errors in CPU to be logged

Problem: While the MCH is in reset and has asserted FSB_CPURST#, the CPU may log spurious errors that are detected by the machine check architecture.

Implication: CPU spurious error log.

Workaround: Error handler should ignore CPU errors when doing CPURST during traffic.

Status: No Fix

3. FBD IBIST pattbuf2enable register does not work for normal operation

Problem: Running FBD IBIST with pattbuf2 enabled (Device 21,22, Function 0, offset 1ACh, 2ACh, bit [13:0]), will result in IBIST failures if SB and NB is not mapped correctly.

Implication: IBIST failure

Workaround: Program the pattbuf2 enable registers with valid SBNB mapping. This workaround is similar to what is done for loopback with inversion.

Status: No Fix

4. FBD IBIST operation does not stop when IBSTR is cleared in Continuous Mode

Problem: When IBIST operation is stopped while operating in Continuous Mode (by clearing the IBISTR register), IBIST operation is not halted as intended.

Implication: IBIST will continue to make TX/RX comparisons based on “junk” data, and will report errors.

Workaround: Before stopping IBIST in Continuous Mode, should first clear the continuous bit in the appropriate IBIST-capable device control registers.

Status: No Fix

5. Simultaneous errors of the same type in both FBD branches result in only Branch 0 FERR being updated

Problem: If two errors of the same type occur in the same cycle on both MCH FBD Branch0 and Branch1, only a single error is logged.

Implication: Simultaneous branch errors of the same type will be reflected in the FERR always as a Branch 0 event.

Workaround: None

Status: No Fix

6. MCH will incorrectly direct all MMCFG accesses to the PCI-Express MMCFG space if HECBASE is enabled with lengths not aligned on 256MB boundaries

Problem: If HECBASE.HECBASEEN is 1 and HECBASE.LENGTH is not aligned on a 256 MB boundary, all MMCFG accesses will be directed towards PCI-E MMCFG rather than the MCH MMCFG space.

Implication: MMCFG transactions intended for MCH will be directed to the PCI-E MMCFG space leaving the system in an unknown state.

Workaround: Force BIOS to always make HECBASE.LENGTH 256 MB aligned when HECBASE is enabled. TOLM is already required to be 256 MB aligned so HECBASE should be set just after TOLM to maximize memory use (placing HECBASE below TOLM would waste memory).

Status: No Fix

7. Clearing Bus Master Enable (BME) Can Generate Address Errors

Problem: The DMA Engine can log Address Errors in the CHANERR register if MCH PCICMD.BME = 0 and the channels are active. This can be due to MCH PCICMD.BME being forcibly cleared while channels are active, or the DMA Engine being started while MCH PCICMD.BME is clear.

Implication: Address errors are flagged in the CHANERR register when they should not be.

Workaround: Clearing BME during traffic is illegal and would cause loss of synchronization and unexpected behavior. Software must quiesce (halt) all DMA traffic before disabling bus master by clearing the BME bit.

Status: No Fix

8. When VT-d is enabled and the Guest Physical Address (GPA) increments beyond FFFF_FFFF_FFFF, an address error is not flagged in the CHANERR register

Problem: In VT-d mode, MCH does not flag an address error in the CHANERR register if an GPA address and transfer length cause an address to increment to a value that is above "Top of Platform Address Space".

Implication: An address error is not flagged in the CHANERR register.

Workaround: VMM driver should not map its last memory page near the 48-bit boundary.

Status: No Fix

9. FBD Chan_Indx set incorrectly for single channel mode

Problem: In single channel mode, VALIDLOG reports the incorrect channel that triggered the recovered memory error

Implication: SW that relies on the FBD channel index field would incorrectly identify the source of failure.

Workaround: SW must ignore the Channel ID when in single channel mode

Status: No Fix

10. Incorrect JTAG scan chain error log

Problem: When accessing a non-existent device incorrectly, chipset logs B17 error when scanning for non-existent device/function.

Implication: B17 is logged

Workaround: None

Status: No Fix

11. PCI Express: MCH not meeting Receiver high impedance Spec

Problem: When the Receiver is in high impedance the measured impedance is lower than required impedance.

Implication: When MCH is powered off or in reset, some adapters may go into compliance mode.

Workaround: None

Status: No Fix

12. MCH does not support PCI Express 2.0 compliance mode features

Problem: The MCH does not fully support all the compliance mode features as defined in the PCI Express base specification version 2.0. Only PCI Express base specification version 1.1 compliance mode is supported fully. PCI Express 2.0 compliance mode features are not supported for following reasons. 1) The concurrent development of the MCH while the PCI Express 2.0 specification was being defined. 2) Some late changes to the compliance mode did not make it into the design. 3) The design/implementation problems due to being a first generation 2.0 product. Since compliance mode is only used for debug and compliance test, there should not be any interoperability problems under normal operation. Below are the list of MCH compliance mode issues by category.

SOS insertion during compliance mode:

- MCH does not support insertion of SOS (Skip Ordered Sets) when transmitting (modified) compliance pattern as required in bit 11 of the Link Control 2 register (PXPLNKCTRL2). MCH has no capability to insert SOS when sending compliance patterns. Compliance mode testing in distinct clocking mode is restricted.

Software exposed registers to control entrance into compliance mode:

- MCH root ports will not set the compliance bit in the TS (Training Set) ordered sets in polling Active when directed by bit 4 (Enter Compliance) of the Link Control 2 Register if bit 10 (Modified Compliance) is set.
- MCH will transition the link to the compliance mode with the link sending modified compliance pattern if bit 10 (Modified Compliance) is set without bit 4 (Enter Compliance) of the link control 2 register being set when entering the Polling LTSSM state. So setting bit 10 without bit 4 will cause the root port to initiate the link to enter compliance mode and send modified compliance patterns. It is recommended that SW never set bit 10 of the link control 2 register.
- MCH will transition the link to compliance mode using the standard compliance pattern when entering the Polling LTSSM state even when bit 10 (Modified Compliance) and bit 4 (Enter Compliance) of Link Control 2 register are set in a MCH root port. Setting bit 4 and bit 10 should have caused the root port to transition the link to compliance mode using the modified compliance pattern.
- Bit 12 (Compliance De-emphasis Level) of the Link Control 2 Register has no affect on MCH root ports. It does not control the compliance mode de-emphasis level.

Returning from Compliance Mode:

- MCH will not send any of the required 8 EIOS (Electrical Idle Ordered Set) when exiting compliance mode (Polling.Compliance) when the link speed was at 5 Gb/s before entering Electrical Idle.
- MCH root port does not properly interpret and lock on to the Modified compliance pattern received from the end point on lanes where the polarity is inverted.

Supporting Compliance Mode request from End Point:

- MCH ignores the compliance bit in the received TS (Training Set) ordered sets for lanes with inverted polarity. If all the lanes of a port have an inverted polarity, the link will not enter compliance mode (Polling.Compliance).

Compliance Mode when Root Port is in Loopback:

- MCH can not change the link speed in Loopback.LTSSM state when the compliance mode bit in the training sets is asserted. If an end point initiates loopback and sets the compliance bit, MCH will ignore the required speed change at 5 Gb/s speed. Speed change should be independent of compliance bit.

Implication: Violation of PCI Express Specification revision 2.0.

Workaround: None

Status: No Fix

13. Store to Write-Through (WT) Memory Data may be seen in wrong order by two subsequent loads when snoop filter is enabled

Problem: If the data of a store transaction to WT memory is used by two subsequent loads of one thread and another thread performs a store to the same address, and the MCH filters out the snoop request triggered by the second store, then the first load may get the data from external memory or the L2 cache as written by another core while the second load gets the data straight from the WT store transaction.

Implication: Software that uses WT memory with shared data may violate proper store ordering. Intel has not observed this erratum with any commercially available software.

Workaround: After the write operation to shared data area of WT memory type, software may use the SFENCE instruction before accessing this data.

Status: No Fix

14. Link Bandwidth Management Status set for autonomous bandwidth change initiated by downstream device

Problem: When a device is “up plugged” into a MCH x8 or x16 slot where the lanes are reversed, any autonomous speed change event by the downstream device will be erroneously logged in the Link Bandwidth Management status instead of the Link Autonomous Bandwidth status in root port’s Link Status Register. This occurs in a MCH x16 port if a x8, x4, x2, or x1 device is inserted into a x16 slot where the lanes are reversed between slot and device. (i.e. x8 device in x16 slot lanes reversed: Device.lane 0 -> RP.lane 15, Device.lane 1 -> RP.lane 14, … Device.lane 7, RP.lane 8). Again this occurs in a MCH x8 port if a x4, x2, or x1 device is inserted into a x8 slot where the lanes are reserved between slot and device. (i.e. x4 device in x8 slot lanes reversed Device.lane 0 -> RP.lane 7 … Device.lane 3 -> RP.lane 4). The autonomous bandwidth event is also erroneously logged if the PCI Express link degrades to the upper 4 lanes or upper 8 lanes of a x8 and x16 RP respectively where the active lane match the up plugged scenarios.

Implication: The bandwidth management status bit is incorrectly logged by a autonomous speed change event from the downstream device.

Workaround: None

Status: No Fix

15. PCI Express: PCI Express LTSSM does not go REC.SPEED from REC.RCVRCFG on receiving valid TS sequences only on some lanes with speed change bit set

Problem: When in the Recovery Substate, the MCH PCI-Express training sequence may not proceed from REC.RCVRCFG to REC.SPEED.

Implication: When successfully trained at 5.0GT/s and one or more MCH PCI-Express lanes encounters an electrical failure of semi-permanent nature (at least 24ms), the link will enter the Recovery Substate per PCI-Express Spec, but MCH will not be able to negotiate a degraded speed and will instead drop the link. Because there may be traffic on the link when it is dropped, this may result in a system hang.

Workaround: None

Status: No Fix

16. PCI Express data corruption in presence of bad and good completion in the same clock

Problem: When a x16 PCI-Express port's inbound data queue for a particular flow class is almost full and a bad packet in that flow class is received that would have exactly filled the data queue is immediately followed by a good packet of that same flow class in the same clock, MCH discards the bad packet and stores the header of the good packet but doesn't push the data for the good packet creating data corruption. For this case to occur, the good packet can have at most 2 DWORDs of data and the start of the packet must immediately follow the end of the bad packet in the same symbol time.

Implication: Data corruption. The state required to expose this issue is extremely unlikely during normal operation.

Workaround: BIOS workaround for posted and non-posted flow classes. No workaround for the completion flow class.

Status: Posted and non-posted classes -- No Fix, Permanent Workaround. Completion class -- Fixed in C0 stepping

17. Peer-to-Peer traffic may cause completion timeouts, primarily in conjunction with PCI Express, FBD, or FSB throttling

Problem: Inbound non-posted requests get blocked behind inbound writes on the source bus due to ordering rules, and inbound completions get blocked behind inbound writes on the target bus which can not be passed due to ordering rules.

Implication: Completion timeouts could occur at the endpoint.

Workaround: Based on the expected level of peer to peer traffic adjust PCI Express, FBD, or FSB throttling registers such that completion timeouts do not occur. The completion timeout value at the endpoint may also be adjustable.

Status: No Fix

18. PCI Express: Electrical Idle inference priority mode causes Polling.Compliance load board failure

Problem: When the load board card initiates polling compliance, the compliance mode can not toggle between different modes when sending the 100MHz burst signal to RX lane.

Implication: PCI Express link get stuck in compliance Gen1 mode

Workaround: In order to do load board compliance testing, set bit 9 of offset 4Ch in Device 9-0, Function 0 to “0” if toggling through the 3 (Gen1, Gen2-3.5dB and Gen2-6dB) transmitter modes is required, for those ports that have load boards connected/ installed.

Status: No Fix

19. MCH IOxAPIC PCISTS.DPE not cleared correctly

Problem: The IOxAPIC data parity error bit (PCISTS[18].DPE, bit 15) does not clear on a write “1” after parity error detection.

Implication: Results in the parity error signal remaining asserted, even after an attempt to directly clear the bit.

Workaround: In order to clear PCISTS[18].DPE after a parity error on RTE access, all RTE locations must be written once (since the RTE location which contained a parity error is unknown). One additional RTE location must then be written. PCISTS[18].DPE will then be cleared by MCH.

Status: No Fix

20. CRC errors logged by AMB during C2 or S1 transition

Problem: AMB occasionally logs CRC errors while the SB lanes are transitioning into Electrical Idle for C2 or S1 power state (MCH FBD fails to set IgnoreERR in the SYNC packet to the AMB preceding Electrical Idle entry).

Implication: Results in bogus errors being set in the AMB error log

Workaround: The following AMB register fields must be cleared upon C2 or S1 exit: Function 1, Offset 90h and 94h, bit 0.

Status: No Fix

21. All OB Cfg type transactions are allowed during D3-hot

Problem: Type 1 Configuration accesses are not blocked from being issued when the RP device state is non-D0 state.

Implication: Type 1 Configuration accesses are allowed to proceed to devices even when the RP is in a non-D0 state. The configuration access should complete normally because devices in D3hot are required to respond to configuration accesses. It is possible that software may be confused if it was expecting a UR response for a type 1 configuration access. However properly behaving software should not be accessing configuration space for devices below a Root Port in a non-D0 state. Software should return the root port to a D0 state before accessing device below the root port.

Workaround: None

Status: No Fix

22. System hang with large number of transaction retries

Problem: A system hang can occur in the presence of a large number of transaction retries. When anti-starvation logic is activated on the starved bus, it is waiting for an inbound request to complete which in turn cannot complete because a prior ordered inbound request in the retry queue cannot complete as the other bus is repeatedly reading that same line.

Implication: Possible system hang.

Workaround: Use MCH MRC 1.0 or later. MRC 1.0 enables memory patrol scrub which prevents the hang condition.

Status: No Fix

23. Low and High Protected Memory Regions in Intel (R) Virtualization Technology for Directed I/O (VT-d) are not supported

Problem: MCH incorrectly reports support for Low and High Protected Memory Regions through bits 5 (PLMR) and 6 (PHMR) of the “VT-d Chipset Capabilities Register” (VTD_CAP).

Implication: Protected Low/High memory region base/limit and enable registers are not supported by MCH.

Workaround: Any software that supports use of Protected Memory Regions must identify this hardware implementation, and not enable protected memory routines.

Status: No Fix

24. PCI Express Differential Tx Voltage swing exceeds spec limits at -6.0dB

Problem: Differential p-p Tx voltage swing measured with compliance load board for 5.0 GT/s data rate at -6.0dB de-emphasis yield worst case eye margins below the required 800mV voltage swing.

Implication: De-emphasis of -6.0dB should not be used with MCH. This mode is intended for trace lengths greater than 12 inches.

Workaround: None

Status: No Fix

25. MCH FBD TX compliance Eye Height in violation at 4.8Gb/s

Problem: MCH FBD Differential p-p Tx voltage Swing measured with compliance load board for 4.8GT/s data rate violates the compliant spec at the pin. The Vswing voltage at the Tx pin does not meet the 800mV minimum voltage requirement for the regular voltage swing.

Implication: Even though the MCH FBD Tx Differential p-p does not meet the compliant spec, lab measurements and simulation data showed that the receiver Vswing meets/exceeds the specs with the design guideline released by Intel.

Workaround: None

Status: No Fix

26. MCH logs correctable errors when L0s is enabled and when the end point has short EI times (<224ns for x16)

Problem: MCH can fail to properly exit L0s when operating at Gen2 data rate if the endpoint device has a series of short L0s.idle times (<=224ns for x16). MCH will timeout from L0s.FTS to Recovery.

Implication: Correctable errors are logged by MCH. No industry adapters are currently known to have this behavior.

Workaround: Disable endpoint L0s if correctable errors are unacceptable

Status: No Fix

27. MCH can timeout in Recovery.Rcvrcfg on L1 exit with PCI Express Gen2 endpoint device that has short L1.idle times

Problem: MCH can fail to properly exit L1 to L0 when operating at Gen2 data rate if endpoint device has a series of short L1.idle times (<=224ns for x16). MCH will timeout in RECOVERY.RCVRCFG and drop to detect which will reset the endpoint.

Implication: System hang

Workaround: It is possible for the BIOS to contain a workaround for this erratum. Contact your OEM for the workaround.

Status: No Fix

28. Read Transactions may be delayed

Problem: Under a certain sequence of read and write transactions issued from processors or bus mastering I/O devices, the read transaction may be delayed.

Implication: A read transaction may be delayed. Intel has not observed this behavior with any commercially available software.

Workaround: None

Status: No Fix

29. PCI express TLP packets not flagged “Malformed” under certain conditions.

Problem: PCI express Transaction Layer Protocol (TLP) packets that are of exactly 256B in length will not be flagged with Malformed TLP when the Maximum Payload Size (MPS) is set to 128B in the PEXDEVCTRL register. All packet sizes greater than 128B except for exactly 256B are correctly reported as a Malformed TLP in the error reporting registers. This does not affect systems with the MPS set to 256B, in this case all packets greater than 256B are correctly reported as Malformed TLP in the error reporting registers.

Implication: When the MPS is set to 128B, if the PCI express end point incorrectly transmits a packet that is exactly 256B the Intel® 5100 MCH will process the packet and will not report a malformed TLP error.

Workaround: None

Status: No Fix

30. Error Source Identification (ID) is not properly reporting the Requestor ID when the uncorrectable (Non-fatal/fatal) error is detected in the PCI express Root Port

Problem: The event collector for uncorrectable error source ID in the Root Complex of the PCI express ports reported in RPERRSID[7:2,0] register under bits[31:16] ERR_FAT_NOFAT_SID field is not capturing the Requestor ID of the source when a Fatal or Non Fatal error is received by the Root Port.

Implication: The value reported in the RPERRSID[7:2,0][ERR_FAT_NOFAT_SID] does not represent the source of the uncorrectable (Non-fatal/fatal) error detected by the root port.

Workaround: Do not use RPERRSID[7:2,0][ERR_FAT_NOFAT_SID] information when a uncorrectable error is detected by the PCI express Root Port.

Status: No Fix

§

6 Specification Changes

This document revision contains no Specification Changes.

§

7 Specification Clarifications

This document version contains no Specification Clarifications.

§

8 Documentation Changes

1. Device ID (DID) Correction

Problem: In version 001 of the Intel® 5400 Chipset Memory Controller Hub (MCH) Datasheet the Device ID is incorrectly listed with the value of 4000h in Table 3-2 on page 53. That table will read as follows in future revisions of the document:

Table 3-2.) Memory Control Hub ESI Device Identification

Component	Register Group	DID	Device	Function	Comment
Intel® 5400 Chipset Memory Controller Hub (MCH)	Enterprise Southbridge Interface	4001h 4003h	0	0	4001h = 1333FSB 4003h = 1600FSB

2. Intel 5400 Chipset inadvertently written as Intel 54000 Chipset

Problem: On page 318 in section 5.3.11.2 and again on page 326 in Table 5-9 there was a typo and where the chipset is listed as the Intel 54000 Chipset it should read Intel 5400 Chipset.

§