Intel® E8870 (870) Chipset

Specification Update

May 2004

Notice: The Intel® E8870 chipset 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.

Document Number: 251114-012

Contents

Revision History

Date	Version	Description
April 2004	012	Added SNC erratum #32.
November 2003	011	Updated workaround for SIOH erratum #6.
September 2003	010	Added SNS erratum #31.
August 2003	009	Added SPS erratum #7.
May 2003	008	Added SNC Errata #29-30; added DMH Documentation Change #1.
April 2003	007	Added SNC Erratum #28 and SIOH Erratum #16; updated workaround for SPS Erratum #6.
March 2003	006	Added SNC Erratum #27; updated workarounds for SNC Erratum #16 and 22; SIOH Errata #4, 6, 7, 13, and 15; SPS Erratum #5
February 2003	005	Added SNC Erratum #26, SIOH Errata # 14-15 and SPS Erratum #6; updated SNC Erratum #1; added SNC C1 and SIOH C2 steppings.
December 2002	004	Added SNC Errata #21-25, SIOH Errata #8-13 and SPS Errata #1-5; added SNC Specification Clarification #2 and SIOH Specification Clarification #1; added support for the E8870IO (SIOH) C1-step component and support for the E8870SP (SPS) component.
November 2002	003	Added SNC Errata #18-20; added SNC Specification Clarification #1.
October 2002	002	Added SIOH Errata #6-7; added DMH Specification Clarification #1.1
September 2002	001	Initial release of this document.

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.

Affected Documents /Related Documents

Title	Document #
Intel® E8870 Scalable Node Controller (SNC) Datasheet	251112
Intel® E8870IO Server I/O Hub (SIOH) Datasheet	251111
Intel® E8870DH DDR Memory Hub (DMH) Datasheet	251113
Intel® E8870SP Scalability Port Switch (SPS) Datasheet	252034

Nomenclature

S-Spec Number is used to identify products. Products are differentiated by their unique characteristics, e.g. core speed, L3 cache size, package types, etc. Care should be taken to read all notes associated with each S-Spec number.

Errata are design defects or errors. These may cause the Intel® E8870 (870) chipset 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 unless otherwise noted.

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

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 the next release of the specification.

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

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.).

Summary Table of Changes

The following tables indicate the errata, specification changes, specification clarifications, or documentation changes which apply to the individual E8870 component steppings. Intel may fix some of the errata in future steppings of the components, and account for the other outstanding issues through documentation or specification changes as noted. These tables use the notations indicated below.

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.
Plan Fix:	This erratum may be fixed in a future stepping of the component.
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 this document.

SNC Errata

No.	Steppings Affected		Pg	Status	Erratum Title
	C0	C1			
1			15	Fixed	FERRST[90] register bit set when a BINIT# is generated by the SNC
2	X	X	15	No Fix	Idle Flit acknowledge bit not cleared on a failed SP LLR
3	X	X	15	No Fix	Strobe Bus Busy (SBSY) asserted indefinitely
4	X	X	15	No Fix	SP performance monitor address compare mode not functioning
5	X	X	16	No Fix	Memory test feature not functioning as specified
6	X	X	16	No Fix	SNC SYRE[10] reset sequence issue
7	X	X	16	No Fix	Performance monitor interval timer cannot generate an interrupt when expired
8	X	X	16	No Fix	RESETO# assertion may not align with the 100 ms memory maintenance cycle
9	X	X	17	No Fix	Hot Page and max count compare status bits not functioning as specified
10	X	X	17	No Fix	Hot Page SRAM index doesn't update when in update mode
11	X	X	17	No Fix	SNC BNR# signal asserted indefinitely
12	X		17	Fixed	Memory initialization optimization not functioning as specified
13	X	X	18	No Fix	SNC.FSBC[10] register bit always returns '0'
14	X	X	18	No Fix	CVCR [2:0] always read as '0'
15	X		18	Fixed	REDMEM register overwritten prematurely
16	X	X	18	No Fix	STM register tRCD setting of 20 ns causes memory read data corruption
17	X		19	Fixed	Memory read queue overflow
18	X	X	19	No Fix	Performance monitors not counting events during response phase
19	X	X	19	No Fix	Performance monitors count only rising edges on monitored signals
20	X	X	19	No Fix	Performance monitor address bus utilization reported incorrectly
21	X	X	20	No Fix	Double deallocation if an invalid speculative memory read conflicts with a coherent FSB request
22	X	X	20	No Fix	False SP Data ECC error may be indicated when an LLR event occurs
23	X		20	Fixed	Inbound read transaction may get lost in the SNC
24	X	X	21	No Fix	Multiple errors may be logged in the FERRST register
25	X	X	21	No Fix	False partial write merge errors may be indicated under specific conditions
26	X	X	22	No Fix	SNC fails to freeze on ERR# assertion
27	X	X	22	No Fix	SPP Request Coherency transaction data not logged upon error
28	X	X	22	No Fix	False Scalability Port fatal errors logged in the SERRST registerr
29	X	X	22	No Fix	ERRMASK bits not set after system reset
30	X	X	23	No Fix	SNC debug feature is non-functional
31	X	X	23	No Fix	SNC may return incorrect data
32	X	X	23	No Fix	Recoverable LLR error logged in SERRST and not in the FERRST

SIOH Errata

No.	Steppings Affected			Pg	Status	Erratum Title
	C0	C1	C2			
1	X	X	X	24	No Fix	Idle Flit acknowledge bit not cleared on a failed SP LLR
2	X			24	Fixed	Incorrect RID value in C0-step components
3	X			24	Fixed	Hub interface failures due to RCOMP induced noise
4	X			24	Fixed	In-bound I/O reads may receive stale data
5	X	X	X	25	No Fix	CBC register allows only DWORD writes
6	X	X	X	25	No Fix	Back-to-back PRC and PRLC could cause loss of cache coherency
7	X	X		25	Fixed	Multi-node deadlock due to read cache invalidations blocking completions
8	X	X		26	Fixed	Possible deadlock situation when the write cache fills
9	X	X		26	Fixed	Starvation mechanism may corrupt outbound deferred transactions
10	X	X		26	Fixed	Inbound write merge stall and snoop hit may return invalid data
11	X	X		26	Fixed	Hub interface may issue an extra packet in the case of a transaction retry
12	X	X		27	Fixed	Write cache forward progress stall
13	X	X	X	27	No Fix	False SP Data ECC error may be indicated when an LLR event occurs
14	X	X	X	27	No Fix	Split read transaction never completes
15	X	X		28	Fixed	SIOH may return invalid data to PCI-X agents on the Hub Interface
16	X	X	X	28	No Fix	False Scalability Port fatal errors logged in the FERRST register

DMH Errata

No.	Steppings Affected	Pg.	Status	Erratum Title
	A1			
1	X	29	No Fix	DMH RAC power-up must be executed before SNC RAC initializes

SPS Errata

No.	Steppings Affected	Pg.	Status	Erratum Title
	B0			
1	X	30	No Fix	Hard reset clears the snoop filter content
2	X	30	No Fix	Idle Flit Acknowledge bit not cleared on a failed SP LLR
3	X	30	No Fix	No response when accessing a disconnected SP port
4	X	30	No Fix	SMBus hangs after receiving invalid address
5	X	30	No Fix	False SP Data ECC error may be indicated when an LLR event occurs
6	X	31	No Fix	False Scalability Port fatal errors logged in the FERRST register
7	X	31	No Fix	SPS starvation prevention logic may cause system live-lock

SNC Specification Changes

No.	Pg.	Title of Change
		No changes at this time.

SIOH Specification Changes

No.	Pg.	Title of Change
		No changes at this time.

DMH Specification Changes

No.	Pg.	Title of Change
		No changes at this time.

SPS Specification Changes

No.	Pg.	Title of Clarification
		No changes at this time.

SNC Specification Clarifications

No.	Pg.	Title of Clarification
1	33	Clarification to SNC SP{0/1}INCO register definitions
2	33	Scalability Port link error bits cannot be cleared

SIOH Specification Clarifications

No.	Pg.	Title of Clarification
1	33	Scalability Port link error bits cannot be cleared

DMH Specification Clarifications

No.	Pg.	Title of Clarification
1	33	DRAM considerations upon deassertion of DMH PWRGOOD

SPS Specification Clarifications

No.	Pg.	Title of Change
		No clarifications at this time.

SNC Documentation Changes

No.	Pg.	Title of Clarification
		No changes at this time.

SIOH Documentation Changes

No.	Pg.	Title of Change
		No changes at this time.

DMH Documentation Changes

No.	Pg.	Title of Change
1	33	DRAM considerations upon deassertion of DMH PWRGOOD

SPS Documentation Changes

No.	Pg.	Title of Change
		No changes at this time.

Identification Information

Markings

The following section details the product markings for the E8870 chipset components and is provided as an identification aid.

SNC, SIOH, SPS Component Markings

Figure 1-1 shows an example of the SNC and SIOH component markings which include the following information:

- INTEL Brand/ INTEL Product ID
- Legal Mark
- Serial Number
- S-Spec Number

Figure 1-1. SNC/SIOH/SPS Markings

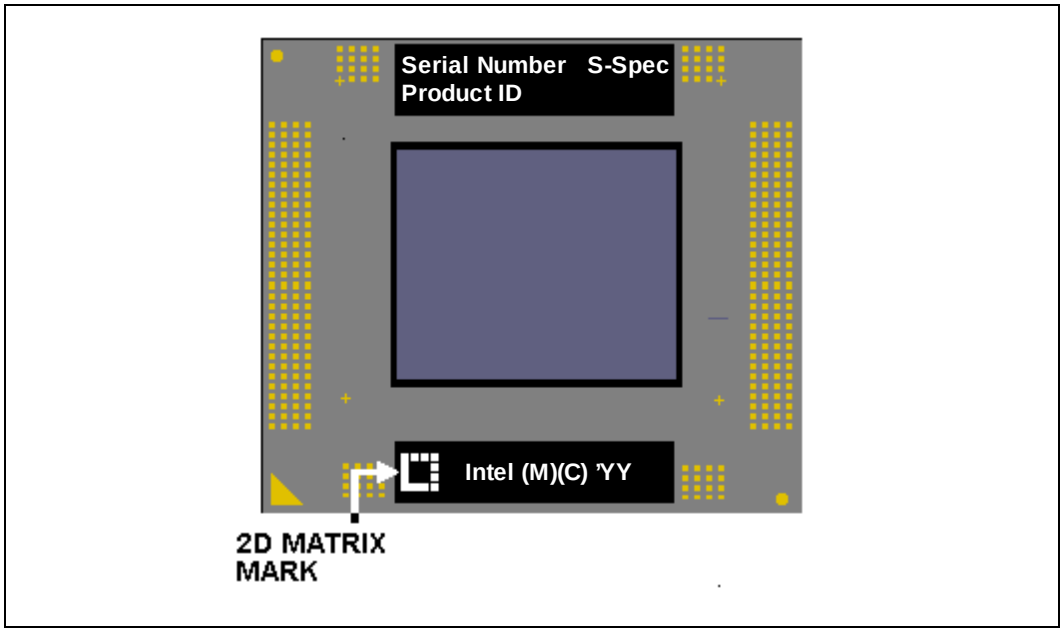

DMH Component Markings

Figure 1-2 shows an example of the DMH component markings which include the following information:

- INTEL Brand/ INTEL Product ID
- Legal Mark
- Serial Number
- S-Spec Number

Figure 1-2. DMH Markings

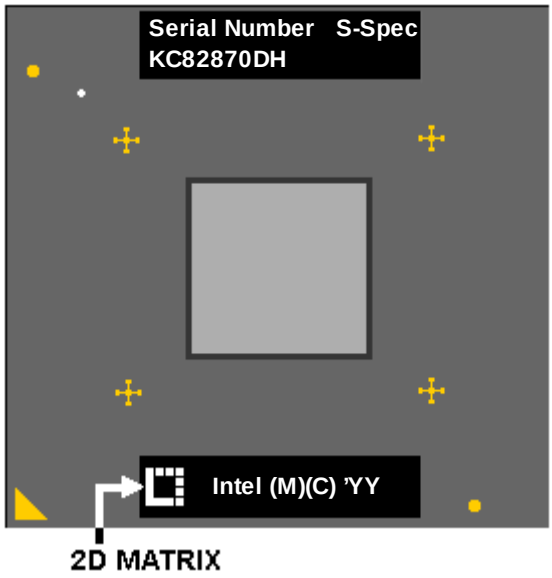

Intel® E8870 Identification and Package Information

Component	Product ID	S-Spec Number(s)	Stepping(s)	Notes
SNC	KW82870MC	SL5X5 SL6XR	C0 C1	
SIOH	KW82870SH	SL5X3 SL6Q4 SL6WX	C0 C1 C2	
DMH	KC82870DH	SL5X2	A1	
SPS	KW82870SP	SL5X6	B0	

Errata

SNC Errata

1. FERRST[90] register bit set when a BINIT# is generated by the SNC

Problem: The SNC erroneously asserts it's FERRST[90] bit even when BINIT# is generated by the SNC itself. The specification states that this bit should not be set if BINIT# is driven by the SNC.

Implication: When BINIT# is generated by the SNC the FERRST "BINIT observed" bit will be set (FERRST[90]). This may affect RAS handling routines.

Workaround: None

Status: Refer to the Summary Table of Changes for affected steppings and status.

2. Idle Flit acknowledge bit not cleared on a failed SP LLR

Problem: The "Idle Flit acknowledge" bit in the SP Interface Control (SP0INCO[4]/SP1INCO[4]) registers is not cleared on a failed SP Link Level Retry (SP LLR).

Implication: The "Idle Flit acknowledge" bit in the SPINCO register does not correctly report the status of the SP bus after a Link Level Retry (LLR) failure.

Workaround: None

Status: Refer to the Summary Table of Changes for affected steppings and status.

3. Strobe Bus Busy (SBSY) asserted indefinitely

Problem: When the SNC receives a RESETI# signal and the SYRE register SAVMEM bit is set to 1, the SNC blocks off all processor requests by asserting a Block Next Request (BNR#) signal and begins flushing any posted writes to preserve memory through reset. Once this is completed, the SNC releases BNR# and resets the processor bus agents. Between deassertion of BNR# and assertion of reset, it is possible for a processor to drive ADS. If this occurs, the SNC will not function correctly after coming out of reset.

Implication: Receipt of an ADS signal prior to the reset will place the SNC in an invalid state. The outcome is an SBSY signal hang on the processor bus.

Workaround: To ensure there is no processor-generated bus traffic prior to reset, all processors on the node will need to be placed in a software loop (running from the cache) prior to initiating the SAVMEM reset.

Status: Refer to the Summary Table of Changes for affected steppings and status.

4. SP performance monitor address compare mode not functioning

Problem: The address compare mode feature of the SNC scalability port (SP) performance monitors is not functioning as specified. When the SPPMR[1:0] register is set for address compare mode, address comparisons are not detected and reported as expected.

Implication: The SP Performance Monitor functions are not capable of triggering on an SP Address Compare.

Workaround: None

Status: Refer to the Summary Table of Changes for affected steppings and status.

5. Memory test feature not functioning as specified

Problem: The SNC includes a feature to autonomously initialize and test main system memory to facilitate fast boots. However, during the memory test function, writes do not always get written to the DRAM array but instead stay in the DMH write buffers. Subsequent reads will be provided from the buffers rather than from the DRAM array as intended.

Implication: SNC hardware memory testing is not possible.

Workaround: Memory testing must be facilitated using software routines.

Status: Refer to the Summary Table of Changes for affected steppings and status.

6. SNC SYRE[10] reset sequence issue

Problem: After initiating a system reset via the SYRE[10] register bit, the SNC internal reset deassertion occurs one clock cycle prior to the SNC driving RESET# to the processors. During this clock cycle, it is possible for a processor to drive ADS which the SNC will accept as the beginning of a new transaction. After the subsequent reset from the SNC, one of the processors will then assert ADS again to initiate its first processor bus transaction. However, at this point, the SNC will hang since it is not expecting a second ADS.

Implication: After a SYRE[10] initiated reset, the SNC will cause a system hang if an ADS was recognized between SNC internal reset deassertion and assertion of the processor reset signal.

Workaround: Modify system BIOS to place the processor(s) in a state that will not produce processor bus transactions while the SNC reset sequence is occuring.

Status: Refer to the Summary Table of Changes for affected steppings and status.

7. Performance monitor interval timer cannot generate an interrupt when expired

Problem: Setting the PTCTL[15] register bit (Timer Status Output Enable) inadvertently affects other SNC registers. Bit 31 of several other “Function 3” registers become set as a result. PTCTL[15] is used to enable the performance monitor timer completion status to be reflected on external event pins. Typically the event pin would be used to generate a system interrupt.

Implication: The unexpected modification of the Function 3 registers could have unpredictable effects on system operation.

Workaround: Do not set PTCTL[15]. This bit must be left at its default value (“0”) or unpredictable chipset operation may result. Timer completion status can be monitored via the PERFCON[6] or PTCTL[6] bits.

Status: Refer to the Summary Table of Changes for affected steppings and status.

8. RESETO# assertion may not align with the 100 ms memory maintenance cycle

Problem: For a deterministic software-initiated reset (with register bit SYRE[12] set), the RESETO# signal assertion should be aligned to the SNC 240 clock boundary and its 100 ms memory maintenance cycle. RESETO# assertion always aligns to the 240 clock boundary. However, if a SYRE reset does not occur within 480 clocks of the end of a memory maintenance cycle, RESETO# assertion may not be aligned with the 100 ms memory maintenance cycle boundary.

Implication: Assertion of RESETO# is not deterministic.

Workaround: For deterministic system implementations, do not rely solely on RESETO# assertion during a hard reset. The system must guarantee that the subsequent RESETI# assertion to other E8870 chipset components has a consistent phase relationship with the SNC memory maintenance cycle. This doesn't imply the need for synchronization with the beginning or end of the maintenance cycle, just a consistent timing relationship with it.

Status: Refer to the Summary Table of Changes for affected steppings and status.

9. Hot Page and max count compare status bits not functioning as specified

Problem: The Hot Page RAM Index bit and Max Count Compare Status bit in the HPPMR register don't behave as expected. The Max Count Compare Status bit (HPPMR[13]) does not get cleared upon initiating a new sample. In addition, this bit is not recognized by internal circuitry while operating in Immediate Sample Enable Mode. This results in the SRAM Index not being updated appropriately in this mode.

Implication: Once set, the Max Count Compare Status bit will remain set until cleared by software.

Workaround: Software should clear HPPMR[13] (Max Count Compare Status bit) prior to starting another sample period. Use the "Local Count Enable" bit in the PERFCON register to start and stop the sample interval instead of using the Immediate Sample Enable mode (HPPMR[4:2] = 001).

Status: Refer to the Summary Table of Changes for affected steppings and status.

10. Hot Page SRAM index doesn't update when in update mode

Problem: The Hot Page RAM Index (HPPMR[31:21]) does not function when operating in update mode (HPPMR[19:18]= 11). After performing a sample period in the update mode, the index to the address which updated the compare register (with the maximum page hit value) is not reflected in the RAM index as it should be.

Implication: The SRAM Index value will not be accurate after the completion of a sample period in update mode.

Workaround: For this mode of operation, all of the Hot Page SRAM entries need to be queried (for the maximum value) to determine which memory block updated the compare register.

Status: Refer to the Summary Table of Changes for affected steppings and status.

11. SNC BNR# signal asserted indefinitely

Problem: Under certain conditions, the SNC asserts the BNR# signal indefinitely. This situation occurs when the internal memory controller logic does not properly handle a conflict between a partial write request and a pending write flush operation.

Implication: Under certain conditions the SNC memory controller will hang causing a BNR# hang of the processor bus.

Workaround: Clear bit 7 of the Memory Control Settings Register (MC[7]). Normally, write buffer flushes occur only when the buffer is full. By clearing bit 7 of the control register, all writes will flush in the absence of pending reads.

Status: Refer to the Summary Table of Changes for affected steppings and status.

12. Memory initialization optimization not functioning as specified

Problem: The SNC Memory Initialization feature does not test all of the installed memory.

Implication: This feature was intended to provide a hardware-based memory initialization mechanism to quicken the initialization process and free up the processor to perform other tasks in parallel. This would potentially allow shorter system boot times. This optimization will be lost.

Workaround: Memory initialization must be performed by system BIOS software routines.

Status: Refer to the Summary Table of Changes for affected steppings and status.

13. SNC.FSBC[10] register bit always returns ‘0’

Problem: The FSBC[10] register bit always returns a ‘0’ when read. However, writing to this bit does cause the corresponding effects described in the SNC datasheet.

Implication: The FSBC[10] bit will not return the correct value when read. However, there are no functional impacts as a result of this erratum.

Workaround: Do not rely on the status read from FSBC[10].

Status: Refer to the Summary Table of Changes for affected steppings and status.

14. CVCR [2:0] always read as ‘0’

Problem: The CVCR register captures the state of specific bus pins during reset. Bits [2:0] are designated to capture frontside bus signals A[34:32]. However, these bits always read from the register as 0s regardless of their state at reset.

Implication: The use of the bus bits associated with CVCR[2:0] is “reserved”. There is no functional impact due to this erratum.

Implication: None

Status: Refer to the Summary Table of Changes for affected steppings and status.

15. REDMEM register overwritten prematurely

Problem: The SNC REDMEM register (Memory Read Data Error Log) latches data information for the first single-bit memory read error detected by the SNC. The contents of this register should not change until the associated bit is cleared in the FERRST register. However, the REDMEM register is being overwritten by the occurrence of subsequent single-bit or multi-bit memory errors.

Implication: The REDMEM register is used to help locate single-bit errors to a single DIMM. Since subsequent errors will overwrite this register, the contents may not be valid. This will limit isolation of the failure to a DIMM row (4 DIMMs sharing the same physical location across 4 DMHs) rather than a single DIMM.

Workaround: The REDMEM register information may not be reliable. Therefore, the RECMEM register contents must be used to determine the location of the memory failure. The RECMEM register can identify the logical DIMM location for a given DMH but is unable to identify to which of the four DMHs the failing DIMM is connected. Therefore, the error can only be narrowed down to a set of four DIMMs (one DIMM on each DMH).

Status: Refer to the Summary Table of Changes for affected steppings and status.

16. STM register tRCD setting of 20 ns causes memory read data corruption

Problem: The SNC DDR-SDRAM Timing Register (STM Register) defines timing parameters for all of the installed DDR SDRAMs. This register includes bits for setting the tRCD (SDRAM RAS to CAS minimum delay) timing parameter according to the capabilities of the installed DIMMs. However, setting tRCD to 20 ns causes a violation of the DMH write-to-read turnaround time specification when DIMMs are installed in both branch channels.

Implication: Setting tRCD to 20ns will cause the SNC to receive stale data when performing a read immediately following a write to the same address.

Workaround: Use the recommended tRCD setting of 30 ns.

Status: Refer to the Summary Table of Changes for affected steppings and status.

17. Memory read queue overflow

Problem: The SNC memory read queue can overflow when subjected to specific transaction patterns. Specific combinations of outbound back-to-back reads, inbound back-to-back “non-snooped” transactions, and simultaneous outbound partial writes can cause the read queue to overflow.

Implication: If the read queue overflows, a read transaction will be lost and will not complete. This situation will eventually result in a processor time-out condition and assertion of BINIT# on the system bus. This erratum has only been observed on system implementations in a synthetic, high stress validation environment.

Workaround: None

Status: Refer to the Summary Table of Changes for affected steppings and status.

18. Performance monitors not counting events during response phase

Problem: The SNC FSB performance monitors do not count system bus events while the bus is in the response phase.

Implication: SNC FSB events that occur while the system bus is in the response bus phase will not be recorded by the performance monitor counter.

Workaround: None

Status: Refer to the Summary Table of Changes for affected steppings and status.

19. Performance monitors count only rising edges on monitored signals

Problem: The SNC performance monitor counters only record the rising edge transitions of monitored event signals.

Implication: Both the SP and FSB performance monitors count only the rising edges of the monitored event signals. Therefore, monitoring signals that remain active for several clocks at a time may give erroneous results. An example would be monitoring the utilization of BNR# through the FSBPMEU register. The performance monitor would only report the transitions of BNR# rather than the number of clock cycles the signal remains active as intended. Monitoring utilitzation of any signal that remains asserted for multiple clock cycles may be affected by this erratum.

Workaround: None

Status: Refer to the Summary Table of Changes for affected steppings and status.

20. Performance monitor address bus utilization reported incorrectly

Problem: The performance monitor address bus utilization function selected through bit 6 of the FSBPMEU register does not correctly report the number of clocks. The specification states the performance monitor will report three clocks per ADS. However, only two clocks are counted per ADS.

Implication: The FSB address bus utilization function will not accurately report the number of clocks representing the bus utilitzation. Only two clocks will be counted for each instance of an ADS providing a bus utilization much lower than actual.

Workaround: To obtain the correct bus utilization value, multiply the value obtained from the perfmon counter by 1.5.

Status: Refer to the Summary Table of Changes for affected steppings and status.

21. Double deallocation if an invalid speculative memory read conflicts with a coherent FSB request

Problem: Normally, an invalid speculative memory read transaction (request to an invalid memory address) should be dropped by the SNC. However, if the SNC receives an invalid speculative read that conflicts with a coincident FSB coherent request, and the two transactions are accepted back-to-back into the conflict resolution logic, the resources associated with the speculative read are not deallocated appropriately, resulting in unpredictable SNC behavior.

Implication: If the above sequence of events occurs, unpredictable system behavior could result, including system hangs and possible data corruption.

Invalid speculative read transactions to the SNC are not expected under normal operating conditions. Therefore, this is not expected to impact normal system operation.

Workaround: An instance of an invalid speculative read address will be logged in the FERRST[3] error register bit. Therefore, this may be used as an indicator of a pending error.

Status: Refer to the Summary Table of Changes for affected steppings and status.

22. False SP Data ECC error may be indicated when an LLR event occurs

Problem: If an LLR occurs, the Scalability Port Link (SPL) layer may incorrectly evaluate an SP Flit header, which does not have valid ECC bits, and thus will flag an ECC error. If no LLR occurs there will be no false error reported.

Implication: Because all SP interfaces are an intermediate point and not an endpoint, a true ECC error would propagate through the chipset. Therefore, if there are no corresponding ECC errors reported elsewhere in the system (i.e. hub interface, system bus, memory), the ECC error is false.

Workaround: If the First Error Status register (FERRST) indicates an SP LLR (initiated by an Idle Flit Duplication error [bit 18] or Parity Error on the Link [bit 17]) and the Subsequent Error Status register (SERRST) indicates an SP single- or double-bit Data ECC error (bits 16 and 19 respectively), or if the SERRST register contains both an SP LLR error and an SP Data ECC error, then the ECC error may be false. Error handling routines should ignore these errors unless an ECC error is also reported at an endpoint.

Status: Refer to the Summary Table of Changes for affected steppings and status.

23. Inbound read transaction may get lost in the SNC

Problem: Under certain internal conditions, an inbound read transaction can get stuck in the SNC Remote Address Translation Table (RATT). This issue can arise during heavy, conflicting, and concurrent memory requests from local and remote sources. If the memory controller begins to exert back pressure to remote requests, subsequent invalidates, reads or snoop transactions (from the SP) that conflict with a pending write-back transaction in the Local Address Translation Table (LATT) may cause concurrent remote read transactions to get stuck in the RATT indefinitely.

Implication: An inbound read may be lost in the SNC RATT, resulting in a processor transaction timeout and subsequent system hang.

Workaround: For single-node implementations, set the SNC SP Control register bit 25 to 1 (SPC[25] = 1).

For multi-node implementations, the Scalability Port Switch (SPS) needs to be reconfigured. Use the algorithm below to program the appropriate SPS internal register values for interleaves 0 - 3 via PCI Functions 6 and 7 of each SPS component. Refer to Section 3.8 of the *E8870SP Scalability Port Switch (SPS) Datasheet* for information on the organization of the SPS interleaves.

1. Write 0x0900h into the PR_IDX register.
2. Write 0062000Ch into the PR_DAT register.

3. Write 0x01h into the PR_CMD register, which will store the contents of PR_DAT.
4. Repeat 1-3 with PR_IDX = 0x0901, PR_DAT = 03B063EC.
5. Repeat 1-3 with PR_IDX = 0x0902, PR_DAT = 63EFB063.
6. Repeat 1-3 with PR_IDX = 0x0903, PR_DAT = B063EC00.
7. Repeat 1-3 with PR_IDX = 0x0904, PR_DAT = 00006303.
8. Write 0x02h into the PR_CMD register.
9. Repeat 1-3 with PR_IDX = 0x0A00, PR_DAT = 0022000C.
10. Repeat 1-3 with PR_IDX = 0x0A01, PR_DAT = 039023E4.
11. Repeat 1-3 with PR_IDX = 0x0A02, PR_DAT = 23E79023.
12. Repeat 1-3 with PR_IDX = 0x0A03, PR_DAT = 9023E400.
13. Repeat 1-3 with PR_IDX = 0x0A04, PR_DAT = 00002303.
14. Write 0x02h into the PR_CMD register.

The following summarizes the needed register addresses within the SPS configuration space:

- PR_IDX[1:0] = Functions 6,7; Offsets [D8h:58h], Bits [15:0]
- PR_CMD[1:0] = Functions 6,7; Offsets [DAh:5Ah], Bits [7:0]
- PR_DAT[1:0] = Functions 6,7; Offsets [DCh:5Ch], Bits [31:0]

Status: Refer to the Summary Table of Changes for affected steppings and status.

24. Multiple errors may be logged in the FERRST register

Problem: The SNC FERRST may have multiple fatal and/or non-fatal errors logged simultaneously. Only one of each error type should be logged at a time. Additional errors should be logged in the SERRST.

Implication: As a result of this erratum, the FERRST register may have multiple fatal and/or non-fatal errors indicated. Although these errors may be logged incorrectly, any indicated error should be considered valid.

Workaround: Error handling routines must be aware that multiple errors of the same type may be simultaneously indicated in the FERRST register.

Status: Refer to the Summary Table of Changes for affected steppings and status.

25. False partial write merge errors may be indicated under specific conditions

Problem: Under certain rare conditions, heavy partial write transactions from remote processors may result in false errors indicated in the FERRST and Subsequent Error Status (SERRST) registers. The partial write merge multi- and single-bit Data ECC error bits (FERRST[5:4], SERRST[5:4]) may be incorrectly asserted.

Implication: Assertion of these error bits may indicate a false error.

Workaround: If an actual ECC error has occurred, it will propagate to a system endpoint (i.e. system bus, hub interface, memory). These register bits should be ignored unless additional ECC errors are indicated at a system endpoint.

Status: Refer to the Summary Table of Changes for affected steppings and status.

26. SNC fails to freeze on ERR# assertion

Problem: The System Level Debug feature, ‘Error Freeze’ does not function as intended. The SNC will only freeze on errors generated internally by the SNC. If an error is signalled on the ERR[2:0]# pins by another chipset component, the SNC will not freeze.

Implication: Setting the FTLFRZ, NCOFRZ and the CORFRZ bits in the ERRCOM register (ERRCOM [2:0]) may not have the intended results as only internally generated SNC errors will be acknowledged.

Workaround: None at this time.

Status: Refer to the Summary Table of Changes for affected steppings and status.

27. SPP Request Coherency transaction data not logged upon error

Problem: The SNC RECoverable Scalability Port Protocol (RECSPP) and Non-RECoverable Scalability Port Protocol (NRECSPP) registers do not correctly log the Coherent/Non-coherent transaction data (bit 62) on error. The SNC logs bit NRECSPP/RECSPP[63] as data/nodata(PktType[1]) and bit NRECSPP/RECSPP[62] as Req/Resp(PktType[0]).

Implication: As a result of this erratum the SPP Request Coherency transaction data is not saved when an error log is generated.

Workaround: Software should decode these bits accordingly for the SNC, SPP Request error logs.

Status: Refer to the Summary Table of Changes for affected steppings and status.

28. False Scalability Port fatal errors logged in the SERRST register

Problem: False Scalability Port strobe glitch errors may be reported causing the SERRST[20] register bit to be set to 1. This bit is used to indicate a fatal system error and when set, will also cause assertion of the SNC fatal error signaling pin(s) (ERR[2]#).

Implication: Depending on the system implementation, this may cause a system hang.

Workaround: Qualify the reported error prior to recognizing it. By setting the SNC ERRMASK[20] bit, the ERR[2] pin(s) will not be asserted when SERRST[20] is set. If the contents of this register are requested then SAL should clear this register to avoid causing false alarms.

Then the FERRST/SERRST registers should be polled for Link Parity, Idle Flit Duplication errors or by checking the other chipset components for any corresponding errors. If no corresponding errors are found, the error is false and the SNC SERRST[20] bit can be cleared since any fatal errors indicated by assertion of SERRST[20] would also cause a failure on the corresponding bus of other chipset components.

Status: Refer to the Summary Table of Changes for affected steppings and status.

29. ERRMASK bits not set after system reset

Problem: The SNC ERRMASK register bits [39:32] will not return to their default state after a system (warm) reset. These bits should assume their default state (set to 1) after a power-on or on a reset. However, these bits act “sticky” and will maintain their state through a reset.

Implication: The ERRMASK register is used to mask potential system errors from being signaled on the chipset ERR [2:0] pins. The ERRMASK [39:32] bits are associated with certain memory subsystem errors and should be set after a reset to mask errors resulting from un-initialized memory. With this erratum, these bits can remain cleared through the reset allowing errors to be signaled unexpectedly.

Workaround: The system BIOS routine should set these bits either before a reset or immediately after to prevent signaling of errors resulting from un-initialized memory

Status: Refer to the Summary Table of Changes for affected steppings and status.

30. SNC debug feature is non-functional

Problem: The SNC Interface Control register (SNCINCO) bit 5 will not disable the SNC system bus interface as was documented in the product datasheet. This feature was provided only as a debug feature.

Implication: There is no functional implication beyond the loss of this debug test feature.

Workaround: None.

Status: Refer to the Summary Table of Changes for affected steppings and status.

31. SNC may return incorrect data

Problem: The SNC may return incorrect data under the following conditions:

- A zero length memory write transaction is issued on the system bus.
- An Implicit Write Back (IWB) transaction, which was initiated by a modified cache line, has been deferred.
- In addition there must be contention for that same modified cache line from at least two other bus agents.

Depending on the order and timing of all the conditions noted above, the SNC may return incorrect data to one of the bus agents requesting ownership of the modified cache line.

Implication: As a result of this erratum, the system will perform in an inconsistent and unpredictable manner.

Workaround: For Itanium® 2-based systems the zero length write can be avoided by ensuring that software does not call PAL_CACHE_FLUSH with Type = 1, 2 or 3 and inv=0.

Status: Refer to the Summary Table of Changes for affected steppings and status.

32. Recoverable LLR error logged in SERRST and not in the FERRST

Problem: Due to certain internal conditions the first recoverable SP LLR error may initially be logged in the SERRST register.

Implication: SNC error logs may generate some confusion as the first recoverable LLR error is flagged in the SERRST register rather than the FERRST register as expected. Reading both error status registers is recommended in the *Intel® E8870/E9870 (870) Chipset Firmware Bootpath Guideline* (Ref No. 11186). In the case of multiple SP errors the logging behavior is correct. This erratum has only been observed with a focused synthetic stress test in a system validation environment.

Workaround: None at this time.

Status: Refer to the Summary Table of Changes for affected steppings and status.

SIOH Errata

1. Idle Flit acknowledge bit not cleared on a failed SP LLR

Problem: The “Idle Flit acknowledge” bit in the SP interface control (SP0INCO[4]/SP1INCO[4]) registers is not cleared on a failed SP LLR.

Implication: The “Idle Flit acknowledge” bit in the SPINCO register does not correctly report the status of the SP bus after a LLR failure.

Workaround: None

Status: Refer to the Summary Table of Changes for affected steppings and status.

2. Incorrect RID value in C0-step components

Problem: The C0-step SIOH RID (Revision ID) value is incorrect.

Implication: The system BIOS typically uses the revision ID to determine chipset configuration register programming. With the SIOH reporting the wrong RID value, there may be an impact to chipset functionality as a result. In addition, this may create other logistical issues associated with tracking component steppings.

Workaround: Modify the system BIOS to remove any dependencies on the SIOH RID. It may be practical to use only the SNC RID to determine the proper system configuration by assuming a consistent combination of SNC/SIOH component revisions.

Status: Refer to the Summary Table of Changes for affected steppings and status.

3. Hub interface failures due to RCOMP induced noise

Problem: Under certain conditions, the hub interface I/O buffer impedance compensation circuitry (RCOMP) will inject noise onto the hub interface I/O signals.

Implication: This event can cause unreliable hub interface electrical signaling, possibly resulting in data errors.

Workaround: The hub interface 1.5/2.0 RCOMP circuits must be disabled as follows:

- Read and record the RCOMP registers (Function 0,1,2,3,4; Offset 88h), bits [13:8] and bits [5:0].
- Write the RCOMP registers (Offset 88h) to disable the RCOMP updates:
  Set bit [15] = 1
  Set bits [13:8] with values read previously
  Set bits [5:0] with values read previously

Status: Refer to the Summary Table of Changes for affected steppings and status.

4. In-bound I/O reads may receive stale data

Problem: During heavy I/O traffic, an in-bound read immediately following a snoop invalidate to the same SIOH read cache line may receive stale data. This issue is applicable to all hub interface ports of the SIOH.

Implication: This issue can result in silent corruption of in-bound read data on any hub interface port.

Workaround: The SIOH read caches must be disabled for each hub interface. To accommodate this change and optimize performance, the following configuration register settings are recommended:

Disable the SIOH read cache and data prefetch feature:
SIOH HLCTL register (Function 0,1,2,3,4, Offset 40h)
Set HLCTL[1:0] = 11

Associated P64H2 settings:
CNF register (Function 0, Offset 40)
For SIOH C0 and C1 with PCI devices Set CNF[5:2] = 1000, for PCI-X devices Set CNF[5:2] = 0000.

Prefetch Control Registers Settings (Function 0):

Control Register	Subsequent Threshold (Ts)	Subsequent Request (Rs)	Initial Threshold (Ti)	Initial Request (Ri)	PCI Offset
PC33	3	3	2	3	F8h
PC66	7	7	3	7	FAh
PC100	B	F	F	F	FCh
PC133	B	F	F	F	FEh

Status: Refer to the Summary Table of Changes for affected steppings and status.

5. CBC register allows only DWORD writes

Problem: BYTE and WORD length writes to the SIOH CBC register do not function. The register will only accept DWORD length writes.

Implication: BYTE/WORD length writes to the CBC register will have no affect.

Workaround: Only perform DWORD writes to the CBC register.

Status: Refer to the Summary Table of Changes for affected steppings and status.

6. Back-to-back PRC and PRLC could cause loss of cache coherency

Problem: An inbound hub interface read, resulting in a port read current (PRC) transaction on the SP, followed immediately by an inbound hub interface read to the same cache line address, resulting in a port read line code (PRLC) command, could be reordered by the SIOH.

Implication: This event would cause a loss of coherency between the internal SIOH snoop filter and the read cache, possibly resulting in data corruption. This issue can only occur on transactions originating from PCI-X agents with differing hub interface IDs.

Workaround: Disable the SIOH read caches as detailed in SIOH erratum #4. Stepping C2 no longer requires the SIOH read cache to be disabled, instead the PRCs must be disabled by setting the IOCTL[13]=1.

Status: Refer to the Summary Table of Changes for affected steppings and status.

7. Multi-node deadlock due to read cache invalidations blocking completions

Problem: Under certain conditions, the SIOH will allow completions to be blocked by “back pressure” resulting from prior requests. If two SIOHs are interacting (for instance, while competing for a common memory resource) and the read/write and snoop transactions are sufficient to cause “back pressure,” completions may also become blocked.

Implication: This will result in neither SIOH being able to make forward progress, resulting in a deadlock situation. This issue impacts only multi-node implementations.

Workaround: Disable all read caches in all but one of the SIOHs. Refer to SIOH erratum #4 workaround section for additional details of applying this workaround; or

Prevent all peer-to-peer traffic between SIOHs. This would require that SIOHs not share any cache lines and not allow peer-to-peer MMIO writes.

Status: Refer to the Summary Table of Changes for affected steppings and status.

8. Possible deadlock situation when the write cache fills

Problem: It is possible for the Least Recently Used (LRU) mechanism in the SIOH write cache to inadvertently block incoming snoops from the SP.

Implication: In this situation, the SIOH would not be able to process the snoop result, eventually leading to deadlocks between incoming snoops and outgoing requests.

Workaround: Two workarounds are available:

1. Disable the SIOH LRB starvation counter by setting bit 30 of the SIOH register accessed through PCI Function 5, offset F4h. This will disable the arbitration fairness mechanism described in Section 5.5.3 of the *RS - Intel® E8870/E9870 (870 Chipset) Server I/O Hub External Design Specification* (Ref# 11002), or
2. Disable write combining by setting bit 12 of the SIOH Control register (IOCTL: PCI Function 5, offset 40h).

Status: Refer to the Summary Table of Changes for affected steppings and status.

9. Starvation mechanism may corrupt outbound deferred transactions

Problem: For multi-node system configurations, the outbound starvation mechanism can cause the SIOH to overrun its outbound deferred buffers and corrupt an outbound deferred transaction. The corrupted transaction will have the same transaction ID as the previous transaction. This results in multiple responses sharing common transaction IDs on the hub interface (HI) and SP interfaces possibly resulting in “stray transaction” (on SP) and “illegal response” (on HI) errors.

Implication: Multi-node system hangs may occur from multiple processors performing simultaneous MMIO accesses to a common HI resource. A typical failure signature is a BPRI# assertion on the system bus with a noncoherent read transaction (targeting the common HI resource) stuck in an SNC LATT, a strayed transaction error in another SNC, and possibly a “Recvd Illegal/Invalid Request/Response” error in the SIOH servicing the targeted HI resource.

Workaround: Disable the outbound anti-starvation logic by setting bit 29 of the SIOH register accessed through PCI Function 5, at offset F4h.

Status: Refer to the Summary Table of Changes for affected steppings and status.

10. Inbound write merge stall and snoop hit may return invalid data

Problem: On an inbound partial write to the SIOH, the entire cache line is modified by merging the partial data. When there are multiple inbound partial writes, merging of the later partial writes could be delayed until the previous writes have been merged. Heavy partial write traffic can cause a later write merge to stall long enough for a snoop to that line to bypass the merge and return the contents of the cache line before the merge has been executed.

Implication: This erratum could cause the SIOH to return the contents of a cache line before a merge is executed and may result in incorrect system or application behavior.

Workaround: Disable write combining in the SIOH by setting the write combining disable bit, IOCTL[12] = 1, in the IOCTL register. This significantly alters the partial write merge traffic within the SIOH.

Status: Refer to the Summary Table of Changes for affected steppings and status.

11. Hub interface may issue an extra packet in the case of a transaction retry

Problem: Under specific circumstances involving heavy I/O traffic on the hub interface, the SIOH can send an extra hub interface transaction following a retry issued to the SIOH from another hub interface agent.

Implication: If the extra packet is interpreted as a completion, the receiving hub interface device (P64H2 or ICH4) may report an error due to a transaction mismatch. However, if the packet is interpreted as a request, the completion may be reported as an “illegal request response” error by the SIOH. If the completion passes the SIOH, the SP interface device (SNC or SPS) will report a “stray transaction” error.

Workaround: This issue has only been encountered in a synthetic system stress test environment. No failures have been observed under normal, OS-controlled operating conditions. In addition, these failures have only been observed in multi-node system test configurations.

Due to these test results and the complexity of the failing sequence, the impact of this issue is considered to be insignificant to normal system operation. No workarounds are recommended.

Status: Refer to the Summary Table of Changes for affected steppings and status.

12. Write cache forward progress stall

Problem: When a partial write that is ‘completing’ a line enters the write cache concurrent with a cache line eviction, the partial write transaction may be delayed causing back pressure from the eviction mechanism. A subsequent sequence of specific back-to-back transactions to the eviction mechanism may then result in further delay to completing the partial write transaction.

Implication: If this sequence of events occurs such that the pending transactions in the write cache block the forward progress of each other, it may eventually result in a system hang.

Workaround: Disable SIOH write combining by setting bit 12 of the SIOH Control register (IOCTL[12] = 1).

Status: Refer to the Summary Table of Changes for affected steppings and status.

13. False SP Data ECC error may be indicated when an LLR event occurs

Problem: If an LLR occurs, the Scalability Port Link (SPL) layer may incorrectly evaluate an SP Flit header, which does not have valid ECC bits, and thus will flag an ECC error. If no LLR occurs there will be no false error reported.

Implication: Because all SP interfaces are an intermediate point and not an endpoint, a true ECC error would propagate through the chipset. Therefore, if there are no corresponding ECC errors reported elsewhere in the system (i.e. hub interface, system bus, memory), the ECC error is false.

Workaround: If the FERRST indicates an SP LLR (Initiated by an Idle Flit Duplication error [bit 9] or Parity Error on the Link [bit 8]) and the SERRST indicates an SP single- or double-bit Data ECC error (bits 7 and 10 respectively), or if the SERRST register contains both an SP LLR error and an SP Data ECC error, the ECC error may be false. Error handling routines should ignore these errors unless an ECC error is also reported at an endpoint.

Status: Refer to the Summary Table of Changes for affected steppings and status.

14. Split read transaction never completes

Problem: The SIOH may receive a split completion for an outbound read transaction from the ICH component. If two or more split completions are received and a read completion spans two flits or more, the Hub Link (HL) 0 block could incorrectly merge the split completions, resulting in the master of the read not receiving an expected completion or a subsequent split completion.

Implication: If a PCI agent on the ICH’s PCI bus disconnects, such that the ICH has at least 2 split completions and the first is aligned on an odd DW boundary and the second contains greater than 2 flits of data, then the possibility exists that the data structures may be misaligned by the SIOH on HL0. Disconnects occur if a device on the ICH’s PCI bus is designated as cacheable and accesses to that space are larger than the target device can accept. If these conditions occur the system may hang as a result of an outbound read transaction that never completes.

Workaround: Restrict all PCI memory address space for the ICH4 to uncacheable (UC/WC).

Status: Refer to the Summary Table of Changes for affected steppings and status.

15. SIOH may return invalid data to PCI-X agents on the Hub Interface

Problem: Under certain rare boundary conditions, the SIOH may return invalid data on the Hub Interface. The required sequence is as follows:

- A read/write is issued from any Hub Interface.
- This is followed by an un-cached read (“hard read”) to the same address.
- This is then followed by a cached read (“soft read”) to the same address from the same hub interface.

  After this sequence, an invalidate from the SP bus or a write from another Hub Interface to the affected address, will not invalidate the associated cache line. Any subsequent reads to the affected cache line prior to eviction, will return invalid data.

Implication: Any PCI-X agents initiating a hard read followed by a soft read to the same address over the same Hub Interface may cause the SIOH to return invalid data on subsequent reads. Any Hub Interfaces with only PCI agents are not susceptible to this issue.

Workaround: Disable read caches on the Hub Interfaces that are configured for PCI-X agents. Refer to SIOH erratum #4 workaround section for additional details of applying this workaround. BIOS can dynamically detect Hub Interfaces configured for PCI-X agents and disable the appropriate read caches.

Status: Refer to the Summary Table of Changes for affected steppings and status.

16. False Scalability Port fatal errors logged in the FERRST register

Problem: False Scalability Port strobe glitch errors may be reported causing the FERRST[11] register bit to be set to 1. This bit is used to indicate a fatal system error and when set, will also cause assertion of the SIOH fatal error signaling pin(s) (ERR[2]#).

Implication: Depending on the system implementation, this may cause a system hang.

Workaround: Qualify the reported error prior to recognizing it. By setting the SIOH ERRMASK[11] bit, the ERR[2] pin(s) will not be asserted when FERRST[11] is set. If the contents of this register are requested then SAL should clear this register to avoid causing false alarms.

Then the FERRST/SERRST registers should be polled for Link Parity, Idle Flit Duplication errors or by checking the other chipset components for any corresponding errors. If no corresponding errors are found, the error is false and the SIOH FERRST[11] bit can be cleared since any fatal errors indicated by assertion of FERRST[11] would also cause a failure on the corresponding bus of other chipset components.

Status: Refer to the Summary Table of Changes for affected steppings and status.

DMH Errata

1. DMH RAC power-up must be executed before SNC RAC initializes

Problem: The DMH has an uninitialized state in the RAC such that it may drive arbitrary values on the DQ pins until it receives the serial command to perform the RAC power-up sequence.

Implication: RAC operations such as temperature/current calibrations, as well as DMH time sync. packets might fail. A DMH sync. failure would cause a system hang.

Workaround: Perform a DMH RAC initialization immediately before the SNC RAC initialization.

Status: Refer to the Summary Table of Changes for affected steppings and status.

SPS Errata

1. Hard reset clears the snoop filter content

Problem: Contrary to what is described in the specification, a hard reset will cause a clearing of the snoop filter contents in the SPS.

Implication: There is no impact to normal operation. It may not be possible to recover the snoop filter contents after a system hang for debug purposes.

Workaround: None

Status: Refer to the Summary Table of Changes for affected steppings and status.

2. Idle Flit Acknowledge bit not cleared on a failed SP LLR

Problem: The “Idle Flit Acknowledge” bit in the SP Interface Control (SP0INCO[4]/SP1INCO[4]) registers is not cleared on a failed SP LLR.

Implication: The “Idle Flit Acknowledge” bit in the SPINCO register does not correctly report the status of the SP bus after a LLR failure.

Workaround: None

Status: Refer to the Summary Table of Changes for affected steppings and status.

3. No response when accessing a disconnected SP port

Problem: In the event of a programming error causing an access to a disconnected SP port, the SPS should respond with a master abort. However, if the node ID of the device that issued the access is 0x1F, the master abort will be lost.

Implication: A system hang could result when the requesting agent does not receive a response from the SPS. Accesses to a disconnected port are not expected under normal operating conditions. Therefore, this erratum is not expected to impact normal system operation.

Workaround: Don't assign the SNC/SIOH node IDs to 0x1F.

Status: Refer to the Summary Table of Changes for affected steppings and status.

4. SMBus hangs after receiving invalid address

Problem: When communicating with the SPS SMBus interface, if the transmitted BUS NUMBER is not equal to 0xFF or the Device portion of the Device/Function byte is not correct (should be the same as the SPS NODE ID), then the SMBus interface will hang until a subsequent power-on reset.

Implication: If a programming error occurs causing the SMBus interface to hang, SPS configuration register accesses from any interface will no longer be possible until after a subsequent power-on reset.

Workaround: SMBus communication software must ensure the BUS NUMBER and Device portion of the Device/Function byte are always correct for the given SPS being addressed. The BUS NUMBER should always be 0xFF. The Device portion [bits 7:3] of the Device/Function byte must be assigned the value “11xxx” (where xxx corresponds to the values of the SPS NODE ID pins).

Status: Refer to the Summary Table of Changes for affected steppings and status.

5. False SP Data ECC error may be indicated when an LLR event occurs

Problem: If an LLR occurs, the SPL layer may incorrectly evaluate an SP Flit header, which does not have valid ECC bits, and thus will flag an ECC error. If no LLR occurs there will be no false error reported.

Implication: Because all SP interfaces are an intermediate point and not an endpoint, a true ECC error would propagate through the chipset. Therefore, if there are no corresponding ECC errors reported elsewhere in the system (i.e. hub interface, system bus, memory), the ECC error is false.

Workaround: If the FERRST indicates an SP LLR (Initiated by an Idle Flit Duplication error [bit 6] or Parity Error on the Link [bit 5]) and the SERRST indicates an SP single- or double-bit Data ECC error (bits 4 and 7 respectively), or if the SERRST register contains both an SP LLR error and an SP Data ECC error, the ECC error may be false. Error handling routines should ignore these errors unless an ECC error is also reported at an endpoint.

Status: Refer to the Summary Table of Changes for affected steppings and status.

6. False Scalability Port fatal errors logged in the FERRST register

Problem: False SP strobe glitch errors may be reported causing the FERRST[8] register bit to be set to 1. This bit is used to indicate a fatal system error and when set, will also cause assertion of the SPS fatal error signaling pin(s) (ERRDx[2]#).

Implication: Depending on the system implementation, this may cause a system hang.

Workaround: Qualify the reported error prior to recognizing it. By setting the SPS ERRMASK[8] bit, the ERRDx[2] pin(s) will not be asserted when FERRST[8] is set. If the contents of this register are requested then SAL should clear this register to avoid causing false alarms.

Then the FERRST/SERRST registers should be polled for Link Parity, Idle Flit Duplication errors or by checking the other chipset components for any corresponding errors. If no corresponding errors are found, the error is false and the SPS FERRST[8] bit can be cleared since any fatal errors indicated by assertion of FERRST[8] would also cause a failure on the corresponding bus of other chipset components.

Status: Refer to the Summary Table of Changes for affected steppings and status.

7. SPS starvation prevention logic may cause system live-lock

Problem: Under the following conditions it is possible for the SPS to temporarily enter a live-lock state.

These conditions must occur simultaneously:

1. A processor owns a modified local cache line.
2. Two or more processors on a remote node attempt to read-modify the same cache line. These processors are correctly receiving the HIT_M response but the Starvation Prevention logic allows them to continuously try to get ownership of the cache line.
3. There is no other traffic on the SP bus.

Any incoming SP traffic will break the live-lock condition and the system will resume normal operation.

Implication: Under these rare conditions, a delay may be noted during system operation. If the live-lock condition occurs during the boot process it could cause the system to reset. Typically the system will continue to operate normally.

Workaround: None.

Status: Refer to the Summary Table of Changes for affected steppings and status.

Specification Changes

There are no Specification Changes for this document revision.

Specification Clarifications

SNC Specification Clarifications

1. Clarification to SNC SP{0/1}INCO register definitions

The *Intel® E8870 Scalable Node Controller (SNC) Datasheet* (Document Number 251112) states the SP Interface Control Registers (SP0INCO,SP1INCO) Response/ Request credits bits, “Must be set to a value of <=25 for reliable operation.” Instead it should state, “Must be set to a value greater than 8 and less than or equal to 25 for reliable operation.”

2. Scalability Port link error bits cannot be cleared

The SNC Link Error bits (FERRST[20], SERRST[20]) cannot be cleared while the associated SP is disabled. The *Intel® E8870 Scalable Node Controller (SNC) Datasheet* (Document Number 251112) will be updated accordingly upon any future revisions.

SIOH Specification Clarifications

1. Scalability Port link error bits cannot be cleared

The SIOH Link Error bits (FERRST[11], SERRST[11]) cannot be cleared while the associated SP is disabled. The *Intel® E8870IO Server I/O Hub (SIOH) Datasheet* (Document Number 251111) will be updated accordingly upon any future revisions.

DMH Specification Clarifications

1. DRAM considerations upon deassertion of DMH PWRGOOD

When the DMH PWRGOOD signal is deasserted, the DMH will stop the SDRAM clocks. If, at this time, there is traffic in flight in the memory subsystem, the loss of the clock may cause the DRAMs to hang or malfunction. Therefore, prior to PWRGOOD deassertion, the memory traffic should be halted or power should be removed from the DRAMs.

Documentation Changes

DMH Documentation Changes

1. The DDR bus periodic slew rate calibration should not be used

The E8870 chipset DDR Memory Controller Hub (DMH) includes a feature called periodic slew-rate calibration, intended to allow automatic optimization of the DDR output slew rate across changing environmental operating conditions. Extensive testing of the DMH determined this additional optimization feature is not needed to maintain signaling margin on the DMH DDR memory bus. As a result, the periodic slew rate calibration settings have not been optimized or validated during testing. Intel is recommending that this feature remain disabled in all E8870 chipset applications. Section 4.5.1 of the *Intel® E8870DH DDR Memory Hub (DMH) Datasheet* (Document Number 251113) will be updated accordingly upon any future revisions.