



Intel[®] 955X Express Chipset

Specification Update

For the Intel[®] 82955X Memory Controller Hub (MCH)

November 2005

Notice: The Intel[®] 82955X MCH may contain design defects or errors known as errata which may cause the product to deviate from published specifications. Current characterized errata are documented in this Specification Update.



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

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

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

The Intel® 82955X 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.

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

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

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

Copyright © 2005, Intel Corporation



Contents

Revision History	4
Preface	5
Errata.....	9
Specification Change	15
Specification Clarifications	16
Documentation Changes	17

Revision History

Version	Description	Date
-001	<ul style="list-style-type: none">Initial Release	April 2005
-002	<ul style="list-style-type: none">Added errata #7 - #14Added document changes #1 - #11	June 2005
-003	<ul style="list-style-type: none">Added document change #12	July 2005
-004	<ul style="list-style-type: none">Added document change #13 - #14Added specification change #1	September 2005
-005	<ul style="list-style-type: none">Added document change #15	October 2005
-006	<ul style="list-style-type: none">Added document change #16	November 2005

§

Preface

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

Information types defined in the Nomenclature section of this document are consolidated into this update document and are no longer published in other documents. This document may also contain information that has not been previously published.

Affected Documents/Related Documents

Document Title	Document Number
<i>Intel® 955X Express Chipset Datasheet</i>	306828

Nomenclature

Errata are design defects or errors. Errata may cause the behavior of the 82955X MCH to deviate from published specifications. Hardware and software designed to be used with any given stepping must assume that all errata documented for that stepping are present on all devices.

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

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



Component Identification via Programming Interface

The 82955X MCH may be identified by the following register contents:

Stepping	Vendor ID ¹	Device ID ²	Revision Number ³
A1	8086h	2774h	81h

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.

Component Marking Information

The 82955X MCH may be identified by the following component markings:

Stepping	Product	S-Spec	Top Marking	Notes
A1	MCH	SL8FW	QG82955X	Production – Pb Free

Summary Table of Changes

The following table indicates the Errata, Specification Changes, Specification Clarifications or Documentation Changes which apply to the listed 82955X MCH steppings. Intel intends to fix some of the errata in a future stepping of the component and to account for the other outstanding issues through documentation or Specification Changes as noted. This table uses the following notations:

Codes Used in Summary Table

Doc:	Document change or update that will be implemented.
PlanFix:	This erratum is intended to be fixed in a future stepping of the component.
Fixed:	This erratum has been previously fixed.
NoFix	There are no plans to fix this erratum.
(No mark) or (Blank Box):	This erratum is fixed in listed stepping or specification change does not apply to listed stepping.
Shaded:	This item is either new or modified from the previous version of the document

NO.	A1	PLANS	ERRATA
1	X	NoFix	PCI Express* x16 Port Related Status Register Bits that Drive SERR Generation Logic are Never Automatically Cleared.
2	X	NoFix	PCIEXBAR Decode Fails When Using Size=64MB or 128 MB and MCHBAR is Not Aligned to 256MB.
3	X	NoFix	PCI Express SKP/InitFCx Contention.
4	X	NoFix	Corrupted Packets Incorrectly Detected as Fatal Errors on PCI Express* Graphics Port
5	X	NoFix	Packet Dropped When Replay Timer Expires and Replay is in Progress
6	X	NoFix	LOCK to non-DRAM Memory Flag (Register C8, Bit 9) is Getting Asserted
7	X	NoFix	The PCI Express Port Does Not Send The Correct TLP Type Downstream When There is a Memory Read Request-Locked TLP
8	X	NoFix	PCI Express Port Skip Sequence is Not Transmitted When Entering Recovery State
9	X	NoFix	STPCLK# Throttling May Cause System to Hang
10	X	NoFix	The Transaction Layer Resets the Completion Timer Counter before Receiving a Passing CRC from the Link Layer on the PCI Express* Port
11	X	NoFix	Malformed Upstream IO or Configuration Write Cycles Are Not Being Detected As Malformed on PCI Express* Port
12	X	NoFix	Excessive Clock Jitter Observed on Intel® 955X Reference Design Platforms

NO.	A1	PLANS	ERRATA
13	X	NoFix	PCI Express* Port is Recognizing an Invalid Transaction with a CRC Error from an Agent as Completed with CRS Status
14	X	NoFix	PCI Express Port Flow Control Updates Being Sent in During PM_REQ_ACK Stream

NO.	SPECIFICATION CHANGES
1	There are no specification changes

NO.	SPECIFICATION CLARIFICATIONS
1	There are no specification clarifications.

NO.	DOCUMENTATION CHANGES
1	Correct DMIUEMSK Register Name
2	Remove Reference to DMI Uncorrectable Status Register
3	Remove DMI Uncorrectable Status Register
4	Remove Reference to DMI Correctable Status Register
5	Remove DMI Correctable Status Register
6	Remove Reference to Correctable Error Mask Register
7	Remove Correctable Error Mask Register
8	Modify Bit Definitions in Device Status (DSTS) Register
9	Modify Bit Definitions in Uncorrectable Error Mask (UEMSK) Register
10	Modify Bit Definitions in DMI Uncorrectable Error Mask (DMIUEMSK) Register
11	Table 10-1 Ball M17 Signal Name Correction
12	Modify Bit Default Values in Prefetchable Memory Limit Address (PMLIMIT1) Register
13	Modify Activate to Precharge Delay Bit Assignment in Channel A and Channel B DRAM Timing (C0DRT1 and C1DRT1) Registers
14	Modify the Register Property for SMRAM—System Management RAM Control
15	Addition of the PEGCNF – PCI Express-G Configuration Register
16	Correction to the maximum memory limit in MCH feature summary

§

Errata

1. PCI Express* x16 Port Related Status Register Bits that Drive SERR Generation Logic are Never Automatically Cleared

Problem: PCI Express* x16 Port related status register bits that drive the SERR generation logic are never automatically cleared. This includes being sticky through warm reset.

Implication: Follow-on (any after first occurrence) errors of same type are ignored because associated status bit is not cleared.

Workaround: BIOS workaround available. Contact your Intel Field Representative for the latest BIOS information.

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

2. PCIEXBAR Decode Fails When Using Size=64MB or 128 MB and MCHBAR is Not Aligned to 256 MB

Problem: When accesses are to Device 0 and 1 on the configuration bus using the enhanced configuration mechanism and the size is set to 64 MB and the address is aligned to a 128 M or 64 M boundary, the transaction gets decoded as a type 1 transaction on the backbone instead of the configuration bus.

Implication: May cause failure to boot or may lead to a system hang.

Workaround: Set the length to 64 MB or 128 MB and align it to a 256 MB boundary.

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

3. PCI Express* SKP/InitFCx Contention

Problem: During 955X PCI Express initialization, if a SKP is being transmitted immediately before an InitFCx DLLP, then a partial InitFCx may be transmitted.

Implication: A slight delay (less than 100 ns) may occur during link initialization. Device may report correctable error. InitFCx will automatically be repeated.

Workaround: None

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

4. **Corrupted Packets Incorrectly Detected as Fatal Errors on PCI Express* Graphics Port**

Problem: When the MCH receives a corrupted packet it may incorrectly detect the corrupted packet as a malformed TLP.

Implication: This issue has only been observed in a synthetic stress test environment and requires the packet corruption to occur in combination with very specific boundary conditions. The MCH may end up detecting a malformed TLP (fatal error) instead of just identifying it as a correctable error on the next valid packet.

Workaround: None.

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

5. **Packet Dropped When Replay Timer Expires and Replay is in Progress**

Problem: When a packet replay is in progress on the PCI Express* Port and the replay timer expires, the next packet in the replay buffer may be sent with an old sequence number. That packet is seen by receiver side as a duplicate and subsequently dropped.

Note: This has only been reproduced in a synthetic test environment.

Implication: A fatal error may be registered by the MCH and the system may hang.

Workaround: None.

Status: No Plans to Fix. For affected steppings, see the *Summary Table of Changes*.

6. **LOCK to non-DRAM Memory Flag (Register C8, Bit 9) is Getting Asserted**

Problem: A processor lock cycle request is unintentionally being recognized as request to a non-system memory destination.

Implication: The MCH may incorrectly flag an error for a valid lock cycle that targets DRAM. A System Error (SERR) may be generated if enabled by System BIOS.

Note: The default setting for ERRCMD[9] Bus 0 Device 0 Offset CAh is to disable this reporting.

Workaround: Do not enable or change default setting of ERRCMD[9] Bus 0 Device 0 Offset CAh (SERR reporting for Lock cycles to non-DRAM Memory).

Status: No Plans to Fix. For affected steppings, see the *Summary Table of Changes*.



7. The PCI Express* Port Does Not Send The Correct TLP Type Downstream When There is a Memory Read Request-Locked TLP

Problem: Upstream Transaction Layer Packet (TLP) Type, Memory Read Request-Locked (MRdLk), is a unsupported request for the MCH on the PCI Express* Port. The Transaction Layer receives the MRdLk and sends downstream a Completion without Data (Cpl) TLP type with an unsupported request status. The correct behavior should be to send a Completion for Locked Memory Read without Data (CplLk) TLP type with an unsupported request status.

Implication: None. PCI Express* 1.0a compliant devices are not allowed to send locked requests upstream.

Workaround: None.

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

8. PCI Express* Port Skip Sequence is Not Transmitted When Entering Recovery State

Problem: PCI Express* Port Skip Sequence in a non-common clock configuration is not transmitted when the skip latency counter expires exactly at the same time the MCH is entering the recovery state. The MCH sends the COM symbol (K28.5) followed by idles instead of skip sequence symbol (K28.0).

Note: This has only been reproduced in a synthetic test environment and only applies to systems that use a non-common clock configuration.

Implication: None. Skip Sequence Symbol generation is not a requirement for proper operation in systems that implement common clock configurations.

Workaround: None.

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

9. STPCLK# Throttling May Cause System to Hang

Problem: In platforms that use STPCLK# throttling in conjunction with devices that invoke the PHOLD mechanism in the ICH (e.g. floppy drives), a boundary condition can occur in the system resulting in the number of STPCLK# acknowledges to be out of synchronization. The failure occurs if a STPCLK# acknowledge cycle is retried on the front side bus at the same time as an internal MCH throttling counter is incremented.

Note: This has only been reproduced in a synthetic test environment under extreme thermal throttling conditions.

Implication: The system may hang.

Workaround: STPCLK# throttling is not necessary in desktop/workstation systems that meet Intel's thermal guidelines and therefore should be disabled by the BIOS. Please reference latest BIOS Specification Update.

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

10. The Transaction Layer Resets the Completion Timer Counter before Receiving a Passing CRC from the Link Layer on the PCI Express* Port

Problem: The PCI Express Port is resetting the completion timer before receiving a Transaction Layer Packet (TLP) with a passing Cyclic Redundancy Check (CRC) indicator from the Link Layer. The completion timer should only be resetting when there is a passing CRC indicator from the Link Layer.

Note: This has only been reproduced in a synthetic test environment.

Implication: None known.

Workaround: None.

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

11. Malformed Upstream IO or Configuration Write Cycles Are Not Being Detected As Malformed on PCI Express* Port

Problem: Malformed upstream IO or configuration write cycles are not being properly detected. The IO or configuration write cycles are put in the upstream non-posted queue as an invalid cycle and an unsupported request completion is returned instead of a fatal error.

Note: This has only been reproduced in a synthetic test environment.

Implication: None. PCI Express 1.0a compliant devices are not allowed to send I/O or Configuration cycles upstream.

Workaround: None.

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



12. Excessive Clock Jitter Observed on Intel® 955X Reference Design Platforms

Problem: DDR2-667 system memory clocks outperform the tCL/tCH spec of 48/52 by reaching 49/51, but do not meet the below listed JEDEC balloted DDR2-667 DRAM Device jitter values at all times. The jitter limits were measured at about the 9-sigma level.

Parameter	JEDEC Value	MCH Value
tJIT(per)	125	290
tJIT(cc)	250	470
tJIT(duty)	125	150
tERR(2per)	175	350
tERR(3per)	225	450
tERR(4per)	250	545
tERR(5per)	250	600

Implication: None. Intel has characterized the system memory clocks and system timing margins and shared the data with the major DRAM suppliers. Intel has determined and the major DRAM suppliers agree that this system clock errata should not cause memory-clock functionality or timing related issues providing all other DRAM related interface timing specifications are fulfilled according to DDR2 Intel specification addendum and JEDEC DDR2 DRAM specification, and the Intel® 955 Platform Design Guide.

Workaround: None.

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

13. PCI Express* Port is Recognizing an Invalid Transaction with a CRC Error from an Agent as Completed with CRS Status

If the MCH has a downstream I/O cycle or Memory Read outstanding, and receives for that cycle a completion TLP that has been corrupted in such a way that the status is “Configuration Retry” (which is illegal), and another corruption within the same TLP appears as a premature “END” symbol, then the MCH may violate system ordering rules.

Note: This has only been reproduced in a synthetic simulation test environment with heavy error injection.

Implication: Anomalous system behavior could result if the exact scenario described above occurs.

Workaround: None.

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

14. PCI Express* Port Flow Control Updates Being Sent in During PM_REQ_ACK Stream

A flow control update DLLP may be sent in the middle of continuous PM_REQ_ACK packets while entering L2/L3 Ready state. This link state is only used when entering the S3/S4/S5 system power management states.

Implication: None known. No system failures have been observed. System will still enter S3/S4/S5 power management states.

Workaround: None.

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

§

Specification Change

1. DDR2-667 1Gbit Memory Technology is supported.

DDR2-667 1Gbit memory technology validation is complete. This memory technology is now supported on the Intel® 955X Express Chipset Family.

§



Specification Clarifications

There are no specification clarifications in this Specification Update revision.

§

Documentation Changes

1. Correct DMIUEMSK Register Name.

Change “DIM Uncorrectable Error Mask” to “DMI Uncorrectable Error Mask” in Table 6-1, page 158.

1C8–1CBh	DMIUEMSK	DMI Uncorrectable Error Mask	DMIBAR	RO, R/W/S
----------	----------	------------------------------	--------	--------------

2. Remove Reference to DMI Uncorrectable Status Register

Remove the following reference from the Table 6-1, page 158.

1C4–1C7h	DMIUESTS	DMI Uncorrectable Error Status	DMIBAR	RO, R/WC/S
----------	----------	--------------------------------	--------	---------------

3. Remove DMI Uncorrectable Status Register

Remove DMIUESTS—DMI Uncorrectable Error Status register on page 166, section 6.1.14.

4. Remove Reference to DMI Correctable Status Register

Remove the following reference from the Table 6-1, page 158.

1D0–1D3h	DMICESTS	DMI Correctable Error Status	DMIBAR	RO, R/WC/S
----------	----------	------------------------------	--------	---------------

5. Remove DMI Correctable Status Register

Remove DMICESTS—DMI Correctable Error Status register on page 168, section 6.1.16.

6. Remove Reference to Correctable Error Mask Register

Remove the following reference from the Table 5-1, page 106.

1D4–1D7h	CEMSK	Correctable Error Mask	00000000h	RO, R/W/S
----------	-------	------------------------	-----------	--------------

7. Remove Correctable Error Mask Register

Remove CEMSK—Correctable Error Mask (D1:F0) on page 155, section 5.1.64.

8. Modify Bit Definitions in Device Status (DSTS) Register

The following are the new bit descriptions that apply to the DSTS Register, on page 134, section 5.1.37.

2	R/WC 0b	Fatal Error Detected: 0 = Fatal Error Not detected. 1 = Indicates that fatal error(s) were detected. Errors are logged in this register regardless of whether error reporting is enabled or not in the Device Control Register.
1	R/WC 0b	Non-Fatal Error Detected: 0 = Non-fatal error Not detected. 1 = Indicates that non-fatal error(s) were detected. Errors are logged in this register regardless of whether error reporting is enabled or not in the Device Control Register.
0	R/WC 0b	Correctable Error Detected: 0 = Correctable error Not detected. 1 = Indicates that correctable error(s) were detected. Errors are logged in this register regardless of whether error reporting is enabled or not in the Device Control Register.

9. Modify Bit Definitions in Uncorrectable Error Mask (UEMSK) Register

The following are the new bit descriptions that apply to the UEMSK Register, on page 153, section 5.1.62.

Bit	Access & Default	Description
31:21		Reserved
20	R/W/S 0b	Unsupported Request Error Mask: 0 = Not Masked 1 = Masked
19:0		Reserved

10. Modify Bit Definitions in DMI Uncorrectable Error Mask (DMIUEMSK) Register

The following are the new bit descriptions that apply to the DMIUEMSK Register, on page 167, section 6.1.15.

Bit	Access & Default	Description
31:21		Reserved
20	R/W/S 0b	Unsupported Request Error Mask: 0 = Not Masked 1 = Masked
19:0		Reserved

11. Table 10-1 Ball M17 Signal Name Correction

In table 10-1 in section 10.1 of the Intel® 955X Express Chipset Datasheet on page 221 the signal name for M17 is incorrect and should be changed to the following:
M17 = VCC

12. Modify Bit Default Values in Prefetchable Memory Limit Address (PMLIMIT1) Register

The following are correct default values for the PMLIMIT1 register on page 119 section 5.1.18

PCI Device: 1
 Address Offset: 26–27h
 Default Value: 0001h
 Access: RO, R/W
 Size: 16 bits

Bit	Access & Default	Description
15:4	R/W 000h	Prefetchable Memory Address Limit (PMLIMIT): This field corresponds to A[31:20] of the upper limit of the address range passed to PCI Express*.
3:0	RO 1h	64-bit Address Support: This field indicates that the bridge has 32-bit address support only.

13. Modify Activate to Precharge Delay Bit Assignment in Channel A and Channel B DRAM Timing (C0DRT1 and C1DRT1) Registers

The Activate to Precharge Delay (t_{RAS}) bit assignment has been corrected from bits 23:20 to bits 22:19 in section 4.2.9.

Bit	Access & Default	Description
31:23		Reserved
22:19	R/W 9h	Activate to Precharge delay (t_{RAS}): This bit controls the number of DRAM clocks for t_{RAS} . The minimum recommendations are beside their corresponding encodings. 0h – 3h = Reserved 4h – Fh = Four to fifteen clocks respectively.
18:10		Reserved
9:8	R/W 01b	CAS# Latency (t_{CL}): This value is programmable on DDR2 DIMMs. The value programmed here must match the CAS Latency of every DDR2 DIMM in the system. 00 = 5 01 = 4 10 = 3 11 = 6
7		Reserved
6:4	R/W 010b	DRAM RAS to CAS Delay (t_{RCB}): This bit controls the number of clocks inserted between a row activate command and a read or write command to that row. 000 = 2 DRAM clocks 001 = 3 DRAM clocks 010 = 4 DRAM clocks 011 = 5 DRAM clocks 100 = 6 DRAM clocks 101–111 = Reserved
3		Reserved
2:0	R/W 010b	DRAM RAS Precharge (t_{RP}): This bit controls the number of clocks that are inserted between a row precharge command and an activate command to the same rank. 000 = 2 DRAM clocks 001 = 3 DRAM clocks 010 = 4 DRAM clocks 011 = 5 DRAM clocks 100 = 6 DRAM clocks 101–111 = Reserved

14. Modify the Register Property for SMRAM—System Management RAM Control

The register property for SMRAM—System Management RAM Control (D0:F0), offset 9Dh, bit 5 (SMM Space Closed) is described as R/W/L. However, it is should be documented as R/W.

15. Addition of the PEGCNF – PCI Express*-G Configuration Register

The datasheet will be updated to include the PEGCNF – PCI Express-G Configuration Register supporting the optional peer-to-peer capability

PCI Device: 1
 Address Offset: 200h-203h
 Default Value: 00201F6Eh
 Access: RO, R/W
 Size: 32 bits

Bit	Access & Default	Description
31:22		Reserved
21	R/W 01b	Peer Write Traffic Disable (PWTB): 0: Peer-to-Peer write cycles will be forwarded 1: Peer-to-Peer write cycles will be blocked and treated as unsupported cycles.
20:0		Reserved

16. Correction to the maximum memory limit in MCH feature summary

In the datasheet on page 13, the system memory section contained a sub bullet incorrectly stating that 8MB maximum memory was supported. The maximum memory limit is 8GB.

§