Intel® E7525 Memory Controller Hub (MCH) Chipset

Datasheet

June 2004

Document Number: 302405-001

Contents

Figures

Tables

Revision History

Doc. No.	Rev. No.	Description	Date
302405	-001	Initial public release.	June 2004

Intel® E7525 MCH Chipset Feature Overview

- Supports Intel® Xeon™ processors with 800 MHz system bus
  - —800 MHz system bus (2X address, 4X data)
  - —Symmetric Multiprocessing Protocol (SMP) for up to two processors at 800 MHz
  - —Parity protection on address, data, request, and response signals
  - —Supports Hyper-Threading technology
  - —Dynamic Bus Inversion (DBI)
  - —36-bit host interface addressing support
  - —12-deep in-order queue
  - —AGTL+ technology with on-die termination
- Memory System
  - —Support for 128 Mb, 256 Mb, 512 Mb, and 1 Gb DRAM densities
  - —Up to two registered memory DDR333 or DDR2-400 channels operating in lock-step
  - —Data bandwidth per channel of 2.67 GB/s (DDR333) or 3.2 GB/s (DDR2-400)
  - —Maximum memory size 16 GB
  - —Hardware memory initialization
- High-Speed Serial PCI Express* Interface
  - —One x16 PCI Express* interface.
  - —One x8 PCI Express* interface; can be configured as two independent x4 interfaces
  - —32-bit CRC and hardware link-level retry
  - —Compatible with *PCI Express Interface Specification*, Rev 1.0a
  - —High bandwidth connection of 4 GB/s per x8 port to I/O processor, PCI-X, Ethernet, or Infiniband* Technology bridge devices
  - —Supports 36-bit addressing using 64-bit semantics
  - —Support for peer segment destination write traffic between PCI Express* ports
  - —Support for non-snooped traffic to memory
  - —Support for remote boot
  - —Support for link active-state and ACPI power management
- Hub Interface 1.5
  - —Connects to either Intel® 82801ER ICH5R or Intel® 6300ESB ICH
  - —266 MB/s bandwidth
  - —Parity protected
  - —Support for differentiated, high priority requests
  - —32-bit downstream addressing
  - —64-bit upstream addressing (full DAC support) truncated to 36 bits internally
  - —Power management messaging
- RASUM
  - —Support for memory sparing
  - —Support for automatic read retry on uncorrectable errors
  - —Support for RAS fail-over to an on-line spare DIMM device
  - —Hardware periodic memory scrubbing, including demand scrub support
  - —Full access to configuration registers via SMBus and IEEE 1149.1 JTAG ports
  - —Support for Intel® x4 Single Device Data Correction (x4 SDDC)
  - —Support for standard SEC-DED (72, 64) ECC on each channel when x4 SDDC technology is disabled
- Package
  - —1077-ball, 42.5 mm, FC-BGA package

Introduction 1

This document details the system architecture supported by the Intel® E7525 MCH, its external interfaces, and other features visible to hardware and software designers implementing a E7525 platform. Included in this specification are descriptions and pin listings for all external electrical interfaces, and descriptions of supported chipset components.

1.1 Terminology

Term	Description
Agent	A logical device connected to a bus or shared interconnect that can either initiate and/or be the target of accesses.
Asserted	Signal is set to a level that represents logical true.
Asynchronous	1. An event that causes a change in state with no relationship to a clock signal. 2. When applied to transactions or a stream of transactions, a classification for those that do not require service within a fixed time interval.
Buffer	1. A random access memory structure. 2. The term I/O buffer is also used to describe a low level input receiver and output driver combination.
Cache Line	The unit of memory that is copied to and individually tracked in a cache. Specifically, 64 bytes of data or instructions aligned on a 64-byte physical address boundary.
CAM	Content Addressable Memory
Cfg	Used as a qualifier for transactions that target PCI configuration address space.
Coherent	Transactions that ensure that the processor's view of memory through the cache is consistent with that obtained through the I/O subsystem.
Command	The distinct phases, cycles, or packets that make up a transaction. Requests and Completions are referred to generically as Commands.
Completion	A packet, phase, or cycle used to terminate a Transaction on an interface, or within a component. A Completion will always refer to a preceding Request and may or may not include data and/or other information.
Core	The internal, base logic of a component.
CRC	Cyclic Redundancy Check. A number derived from, and stored or transmitted with, a block of data in order to detect corruption. By recalculating the CRC and comparing it to the value originally transmitted, the receiver can detect some types of transmission errors.
Deasserted	Signal is set to a level that represents logical false.
DED	Double-bit Error Detect
Deferred Transaction	A processor bus Split Transaction. The requesting agent receives a Deferred Response which allows other transactions to occur on the bus. Later, the response agent completes the original request with a separate Deferred Reply transaction.
Delayed Transaction	A transaction where the target retries an initial request, but unknown to the initiator, forwards or services the request on behalf of the initiator and stores the completion or the result of the request. The original initiator subsequently re-issues the request and receives the stored completion
DMA	Direct Memory Access. Method of accessing memory on a system without interrupting the processors on that system

Term	Description
Downstream	Describes commands or data flowing away from the processor-memory complex and toward I/O. The terms Upstream and Downstream are never used to describe transactions as a whole. For example, downstream data may be the result of an Outbound Write, or an Inbound Read. The Completion to an Inbound Read travels Downstream.
DP	Dual Processor
Dual Address Cycle (DAC)	As described in the *PCI Local Bus Specification, Rev 2.3*, Dual Address Cycles are used by PCI devices that utilize 64-bit addressing for memory transactions.
DW	Double Word. A reference to 32 bits of data on a naturally aligned four-byte boundary (i.e., the least significant two bits of the address are 00b).
ECC	Error Correcting Code
Gb/s	Gigabits per second (10^9 bits per second)
GB/s	Gigabytes per second (10^9 bytes per second)
Host	This term is used synonymously with “processor.”
Hub Interface (HI)	Proprietary interface that connects the MCH to the ICH.
I/O	1. Input/Output. 2. When used as a qualifier to a transaction type, specifies that transaction targets Intel® Architecture-specific I/O space. (e.g., I/O read)
ICH	Intel® I/O Controller Hub (ICH) device. This refers to the Intel® 82801ER I/O Controller Hub 5-R (ICH5R)or Intel® 6300ESB ICH.
Implicit Writeback	A snoop initiated data transfer from the bus agent with the modified Cache Line to the memory controller due to an access to that line.
Inbound	A transaction where the request destination is the processor-memory complex and is sourced from I/O. The terms Inbound and Outbound refer to transactions as a whole and never to Requests or Completions in isolation. For example, an Inbound Read generates Downstream data, whereas an Inbound Write has Upstream data. The Completion to an Inbound Read travels Downstream.
Industry Standard Architecture (ISA)	A 16-bit bus architecture associated with the IBM AT motherboard designed to connect motherboard circuitry to expansion card devices that is now considered Legacy.
Initiator	The source of requests. For example, an agent sending a request packet on a PCI Express* interface is referred to as the Initiator for that Transaction. The Initiator may receive a completion for the Request.
ISA Regime	A special Legacy mode to support ISA-based devices which have been integrated into the chipset that opens a dedicated channel from the peripheral device to the processor bus. While in this mode, the Legacy device is granted exclusive accesses to memory and the ability to use Tenured Transactions.
Isochronous	A classification of transactions or a stream of transactions that require service within a fixed time interval.
Layer	A level of abstraction commonly used in interface specifications as a tool to group elements related to a basic function of the interface within a layer and to identify key interactions between layers.
Legacy	Functional requirements handed down from previous chipsets or PC compatibility requirements from the past.
Link	A full duplex transmission path between any two PCI Express* devices.
LSb	Least Significant Bit
LSB	Least Significant Byte
Master	A device or logical entity that is capable of initiating transactions. A Master is any potential Initiator.
MB/s	Megabytes per second (10^6 bytes per second)
MCH	Memory Controller Hub
Mem	Used as a qualifier for transactions that target memory space. (e.g., a Mem read to I/O)
MSb	Most Significant Bit

Term	Description
MSB	Most Significant Byte
MSI	Message Signaled Interrupt. MSIs allow a PCI device to request interrupt service via a Memory Write transaction rather than a hardware signal.
MTBF	Mean Time Between Failure
Non-coherent	Transactions that may cause the processor's view of memory through the cache to be different with that obtained through the I/O subsystem.
Outbound	A transaction where the request destination is I/O and is sourced from the processor-memory complex. The terms Inbound and Outbound refer to transactions as a whole and never to Requests or Completions in isolation. For example, an Outbound Read generates Upstream data, whereas an Outbound Write has Downstream data. The Completion to an Outbound Read travels Upstream.
Packet	The indivisible unit of data transfer and routing, consisting of a header, data, and CRC.
PCI	Peripheral Component Interconnect Local Bus. A 32- or 64-bit bus with multiplexed address and data lines that is primarily intended for use as an interconnect mechanism within a system between processor/memory and peripheral components or add-in cards.
Plesiochronous	Each end of a link uses an independent clock reference. Support of this operational mode places restrictions on the absolute frequency difference, as specified by PCI Express*, which can be tolerated between the two independent clock references.
Posted	A Transaction that is considered complete by the initiating agent or source before it actually completes at the Target of the Request or destination. All agents or devices handling the Request on behalf of the original Initiator must then treat the Transaction as being system visible from the initiating interface all the way to the final destination. Commonly refers to memory writes.
Push Model	Method of messaging or data transfer that predominately uses writes instead of reads.
Intel® 6700PXH	PCI Express* to PCI-X bridge device.
Queue	A first-in first-out (FIFO) structure.
Receiver	1. The Agent that receives a Packet across an interface regardless of whether it is the ultimate destination of the packet. 2. More narrowly, the circuitry required to convert incoming signals from the physical medium to more perceptible forms.
Request	A packet, phase, or cycle used to initiate a Transaction on a interface, or within a component.
Reserved	The contents or undefined states or information are not defined at this time. Using any reserved area is not permitted. Reserved register bits must be set to '0'.
RMW	Read-Modify-Write operation
S4EC/D4ED	A specific data protection algorithm that distributes data and ECC across 144 bits. Enables correction of all error combination within a single nibble, and therefore provides the capability to map out a faulty x4 DRAM device. Allows detection of all error combinations contained within two nibbles, and greater than 90% of error combinations affecting more than two nibbles.
SEC	Single-bit Error Correct
Snooping	A means of ensuring cache coherency by monitoring all memory accesses on a common multi-drop bus to determine if an access is to information resident within the cache of another bus agent.
SSTL	Stub-Series Terminated Logic
Symbol	An expanded and encoded representation of a data Byte in an encoded system (e.g., the 10-bit value in an 8-bit/10-bit encoding scheme). This is the value that is transmitted over the physical medium.
Symbol Time	The amount of time required to transmit a symbol.
System Bus	The processor-to-MCH interface, which includes control signals and source-synchronous address and data signals.
Target	A device that responds to bus Transactions. On a PCI Express* interface, the agent receiving a request packet is referred to as the Target for that Transaction.

Term	Description
Tenured Transaction	A transaction that holds the bus, or interconnect, until complete, effectively blocking all other transactions while the Target is servicing the Request.
Transaction	An operation between two or more agents that can be comprised of multiple phases, cycles, or packets.
Transmitter	1. The Agent that sends a Packet across an interface regardless of whether it was the original generator of the packet. 2. More narrowly, the circuitry required to drive signals onto the physical medium.
UP	Uniprocessor
Upstream	Describes commands or data flowing toward the processor-memory complex and away from I/O. The terms Upstream and Downstream are never used to describe transactions as a whole (e.g., Upstream data may be the result of an Inbound Write, or an Outbound Read. The Completion to an Outbound Read travels Upstream).

1.2 Reference Documentation

Document	Document Number / Sources
Double Data Rate (DDR) SDRAM Specification	http://www.jedec.org/
Intel® 6700PXH PCI-X Hub Datasheet	http://developer.intel.com/
Intel® 82801ER I/O Controller Hub 5-R (ICH5R) Datasheet	http://developer.intel.com/
Intel® 6300 ESB I/O Controller Datasheet	http://developer.intel.com/
Intel® Xeon™ Processor with 800 MHz System Bus Datasheet	http://developer.intel.com/
PCI Local Bus Specification, Rev 2.3	http://www.pcisig.com/home
PCI-X Addendum to the PCI Local Bus Specification, Rev 1.0b	http://www.pcisig.com/home
PCI Express Interface Specification, Rev 1.0a	http://www.pcisig.com/home
System Management Bus (SMBus) Specification, Rev 2.0	http://www.smbus.org
PCI Bus Power Management Interface Specification, Rev 1.1	http://www.pcisig.com/home
Advanced Configuration and Power Interface Specification (ACPI)	http://www.acpi.info

1.3 Intel® E7525 MCH System Architecture

The architecture of the chipset provides the performance and feature set required for dual-processor based workstations. To accomplish this, the MCH has numerous reliability and manageability features on multiple interfaces. Detailed descriptions of the interfaces and feature set are provided in Chapter 5, “Functional Description.”

The chipset consists of the following components: Memory Controller Hub (MCH), ICH and the Intel® 6700PXH PCI-X Hub. Although a brief overview is provided here, detailed component information can be found in each device's respective documentation.

1.3.1 Intel® Xeon™ Processor with 800 MHz System Bus

The MCH supports either single or dual population of the Xeon processors. The front side bus supports a base system bus frequency of 200 MHz. The address and request interface is double pumped to 400 MHz while the 64-bit data interface (+ parity) is quad pumped to 800 MHz. This provides a matched system bus address and data bandwidths of 6.4 GB/s.

1.3.2 Memory Subsystem

The MCH provides an integrated memory controller for direct connection to two channels of registered DDR333 or DDR2-400 memory. Peak theoretical memory data bandwidth using DDR333 technology is 5.33 GB/S. For DDR2-400 technology, this increases to 6.4 GB/s.

When both DDR channels are populated and operating, they function in lock-step mode. The maximum supported DDR333 or DDR2-400 memory configuration is 16 GB. There are several high-end features for the MCH memory interface:

- DIMM sparing allows for one DIMM per channel to be held in reserve and brought on-line if another DIMM in the channel becomes defective.
- Hardware periodic memory scrubbing, including demand scrub support
- Retry on uncorrectable memory errors
- x4 SDDC for memory error detection and correction of any number of bit failures in a single x4 memory device

1.3.3 PCI Express*

The MCH supports the new *PCI Express Interface Specification, Rev 1.0a.* The MCH provides one x16 and one configurable x8 PCI Express* interface. The x8 PCI Express* interface may alternatively be configured as two independent x4 PCI Express* interfaces. The max theoretical bandwidth of a x8 PCI Express* port is 4 GB/s. This value scales linearly with port size.

The MCH is a root class component as defined in the *PCI Express Interface Specification, Rev 1.0a*. In addition to the 6700PXH detailed below, the PCI Express* interfaces support connection of the MCH to a variety of other bridges compliant with *PCI Express Interface Specification, Rev 1.0a*. Other compatible PCI Express* devices, available from Intel and/or third-party vendors, implement functionality such as graphics, H/W RAID controllers and TCP/IP offload engines.

1.3.4 Hub Interface 1.5

The MCH interfaces with the Intel® 82801ER I/O Controller Hub 5-R or the Intel® 6300ESB ICH via a dedicated Hub Interface 1.5 supporting a peak bandwidth of 266 MB/s using a x4 base clock of 66 MHz. Throughout this document, they both will routinely be referred to simply as the “ICH”.

1.3.4.1 Intel® 82801ER I/O Controller Hub 5-R (ICH5R)

The ICH5R provides legacy function support similar to that of previous ICH-family devices, but with extensions in RAID 0 support, Serial-ATA technology, and an integrated ASF controller. The ICH5R also includes integrated USB 2.0 and USB Classic support, ATA-100 IDE support, integrated 10/100 Ethernet, an LPC firmware hub (FWH) and SuperIO interface, a system

management interface, a power management interface, integrated IOxAPIC and 8259 interrupt controllers, and a 32-bit/33 MHz PCI bus interface. The ICH5R also includes desktop features such as integrated 20-bit audio.

1.3.4.2 Intel® 6300ESB ICH

The 6300ESB ICH provides legacy function support similar to that of previous ICH-family devices, but with extensions in Serial-ATA technology and 64-bit/66MHz PCI-X support. The 6300ESB ICH also includes integrated USB 2.0 and USB Classic support, Ultra ATA/100 IDE support, an LPC firmware hub (FWH) and SuperIO interface, a system management interface, a power management interface, integrated IOxAPIC and 8259 interrupt controllers, and an integrated DMA controller. The 6300ESB ICH also includes desktop features such as integrated 20-bit audio.

1.3.5 Intel® 6700PXH

The PXH provides connection between a PCI Express* interface and two independent PCI bus interfaces configurable for standard PCI 2.3 protocol, as well as the enhanced high-frequency PCI-X 1.0b protocol. The PXH provides configurable support for 32- or 64-bit PCI devices, as well as support for interface operation at 33 MHz, 66 MHz, 100 MHz, or 133 MHz.

1.3.6 Platform Summary

Figure 1-1 summarizes one possible system configuration using the MCH and associated peripheral components.

Figure 1-1. System Block Diagram

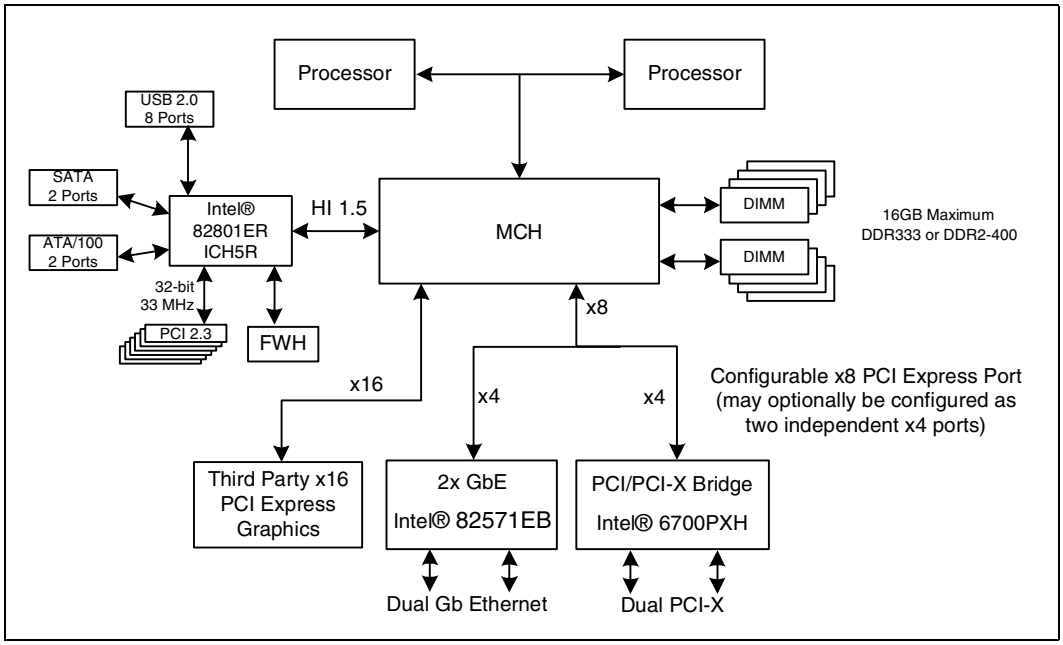

Signal Description 2

The following notations are used to describe the signal type:

I	Input pin
O	Output pin
I/O	Bidirectional input/output pin
s/t/s	Sustained tri-state. This pin is driven to its inactive state prior to tri-stating.
as/t/s	Active sustained tristate. This applies to some of the HI signals. This pin is weakly driven to its last driven value.
2x	Double-pump clocking. Addressing at 2x of HCLK.
4x	Quad-pump clocking. Data transfer at 4x of HCLK. The signal description also includes the type of buffer used for the particular signal.
AGTL+	Open drain AGTL+ interface signal. Refer to the AGTL+ I/O Specification for complete details. The MCH integrates AGTL+ termination resistors.
CMOS	CMOS buffers
SCHMITT	CMOS buffers with integrated Schmitt trigger circuit to compensate for slow edges
SSTL-2	Stub series terminated logic for 2.5V. Used for the source-synchronous DDR interface signals. Refer to the JEDEC Specification D8-9A for complete details.
OD	Open drain output. The MCH can drive a logic 0 on the pin, but relies on an external pull-up resistor to attain a logic 1 level. The voltage associated with logic 1 is specified for each OD type signal. The platform designer must ensure that these pins are not pulled to an unsupported voltage level.
DIFF	Differential high-speed serial signals (PCI Express*)

2.1 System Bus Interface Signals

Table 2-1. System Bus Signal Description (Sheet 1 of 3)

Signal Name	Type	Description
ADS#	I/O AGTL+	**Address Strobe:** The system bus owner asserts **ADS#** to indicate the first of two cycles of a request phase.
AP[1:0]#	I/O AGTL+	**Address Parity:** The **AP[1:0]#** lines are driven by the request initiator and provide parity protection for the Request Phase signals. **AP[1:0]#** are common clock signals and are driven one common clock after the Request Phase. Address parity is correct if there are an even number of electrically low signals (low voltage) in the set consisting of the covered signals plus the parity signal. Note that the MCH only connects to **HA[35:3]#**. The MCH may be configured to send an error message to the ICH when it detects an address/parity error on the system bus.
BINIT#	I AGTL+	**Bus Initialize:** This signal indicates an unrecoverable error and can be driven by the processor. It is latched by the MCH.
BNR#	I/O AGTL+	**Block Next Request:** Used to block the current request bus owner from issuing new requests. This signal is used to dynamically control the system bus pipeline depth.
BPRI#	O AGTL+	**Priority Agent Bus Request:** The MCH is the only Priority Agent on the system bus. It asserts this signal to obtain the ownership of the address bus. The MCH has priority over symmetric bus requests and will cause the current symmetric owner to stop issuing new transactions unless the **HLOCK#** signal is asserted.
BREQ0#	I/O AGTL+	**Bus Request 0#:** The MCH pulls the processor bus **BREQ0#** signal low during **CPURST**#. The signal is sampled by the processors on the active-to-inactive transition of **CPURST**#. The minimum setup time for this signal is 4 HCLKs. The minimum hold time is 2 HCLKs and the maximum hold time is 20 HCLKs. **BREQ0#** should be tristate after the hold time requirement has been satisfied.
BREQ1#	I AGTL+	**Bus Request 1#:** The MCH does not drive this pin. It is used by the processors for symmetric bus arbitration. The MCH samples this signal during run time only.
CPURST#	O AGTL+	**Processor Reset:** The MCH asserts **CPURST#** while **RSTIN#** (PCIRST# from ICH) is asserted and for approximately 1 ms after **RSTIN#** is deasserted. The **CPURST#** allows the processors to begin execution in a known state.
DBSY#	I/O AGTL+	**Data Bus Busy:** Used by the data bus owner to hold the data bus for transfers requiring more than one cycle.
DEFER#	O	**Defer:** Signals that the MCH will terminate the transaction currently being snooped with either a deferred response or with a retry response.
DEP[3:0]#	I/O AGTL+	**Host Data Parity:** The **DP[3:0]#** signals provide parity protection for **HD[63:0**]#. The **DP[3:0]#** signals are common clock signals and are driven one common clock after the data phases they cover. **DP[3:0]#** are driven by the same agent driving **HD[63:0]**#. Data parity is correct if there are an even number of electrically low signals (low voltage) in the set consisting of the covered signals plus the parity signal.
DBI[3:0]#	I/O AGTL+	**Dynamic Bus Inversion:** Driven along with the **HD[63:0]#** signals. Indicates when the associated signals are inverted. **DBI[3:0]#** are asserted such that the number of data bits driven electrically low (low voltage) within the corresponding 16-bit group never exceeds eight.

Table 2-1. System Bus Signal Description (Sheet 2 of 3)

Signal Name	Type	Description
DRDY#	I/O AGTL+	**Data Ready:** Asserted for each cycle that data is transferred.
HA[35:3]#	I/O AGTL+ 2x	**Host Address Bus: HA[35:3]#** connect to the system address bus. During processor cycles, **HA[35:3]#** are inputs. The MCH drives **HA[35:3]#** during snoop cycles on behalf of Hub Interface (HI) initiators.
HACVREF	I	**Host Address/Control and Common Clock Bus Reference Voltage:** Reference voltage input for the Address and Control signals of the Host AGTL+ interface.
HADSTB[1:0]#	I/O AGTL+ 2x	**Host Address Strobe:** The source synchronous strobes are used to transfer **HA[35:3]#** and **HREQ[4:0]#** at the 2x transfer rate.
HCLKINN HCLKINP	I CMOS	**Differential Host Clock In:** These pins receive a differential host clock from the external clock synthesizer. This clock is used by all the MCH logic in the host clock domain.
HCRES0	I/O Analog	**Host Compensation Resistor Return:** Common return for host bus compensation resistors on HODTCRES and HSLWCRES.
HD[63:0]#	I/O AGTL+ 4x	**Host Data:** These signals are connected to the system data bus.
HDSTBN[3:0]# HDSTBP[3:0]#	I/O AGTL+ 4x	**Differential Host Data Strobes:** The differential source synchronous strobes are used to transfer **HD[63:0]#** and **DBI[3:0]#** at the 4x transfer rate. Strobe Data Bits HDSTBP3#, HDSTBN3# HD[63:48]#, DBI3# HDSTBP2#, HDSTBN2# HD[47:32]#, DBI2# HDSTBP1#, HDSTBN1# HD[31:16]#, DBI1# HDSTBP0#, HDSTBN0# HD[15:0]#, DBI0#
HDVREF[1:0]	I	**Host Data Reference Voltage:** Reference voltage input for the 4x Data Signals of the Host AGTL+ interface.
HIT#	I/O AGTL+	**Hit:** Indicates that a caching agent holds an unmodified version of the requested line. Also, driven in conjunction with **HITM#** by the target to extend the snoop window.
HITM#	I/O AGTL+	**Hit Modified:** Indicates that a caching agent holds a modified version of the requested line and that this agent assumes responsibility for providing the line. **HITM#** is driven in conjunction with **HIT#** to extend the snoop window.
HLOCK#	I AGTL+	**Host Lock:** All system bus cycles that are sampled with the assertion of **HLOCK#** and **ADS#**, until the negation of **HLOCK#** must be atomic.
HODTCRES	I Analog	**Compensation Resistor:** Compensation for system bus on-die termination.
HREQ[4:0]#	I/O AGTL+ 2x	**Host Request Command:** Defines the attributes of the request. **HREQ[4:0]#** are transferred at the 2x rate. Asserted by the requesting agent during both halves of a Request Phase. In the first half the signals define the transaction type to a level of detail that is sufficient to begin a snoop request. In the second half the signals carry additional information to define the complete transaction type.
HSLWCRES	I Analog	**Compensation Resistor:** Compensation for system bus slew rate.
HTRDY#	O AGTL+	**Host Target Ready:** Indicates that the target of the processor transaction is able to enter the data transfer phase.

Table 2-1. System Bus Signal Description (Sheet 3 of 3)

Signal Name	Type	Description
MCERR#	I/O AGTL+	**Machine Check Error:** May be asserted by the MCH or the processors to indicate an unrecoverable error without a bus protocol violation. The following assertion options are configurable at a system level. • Enable/Disable • Asserted, if configured, along with IERR# for processor internal errors • Asserted, if configured, by the request initiator of a bus transaction after it observes and error • Asserted by any bus agent when it observes an error in a bus transaction
RS[2:0]#	O AGTL+	**Response Status: RS[2:0]#** indicate the type of response according to the following table: **RS[2:0] Response Type** 000 Idle state 001 Retry response 010 Deferred response 011 Reserved (not driven by MCH) 100 Hard Failure (not driven by MCH) 101 No data response 110 Implicit Writeback 111 Normal data response
RSP#	O AGTL+	**Processor Bus Response Parity**: Provides parity protection for **RS[2:0]** signals. **RSP#** is always driven by the MCH, and is valid on all clocks. Parity is correct if there are an even number of electrically low signals (low voltage) in the set consisting of the covered signals plus the parity signal.

2.2 DDR Interface A Signals

Table 2-2. DDR Channel_A Interface Signals (Sheet 1 of 2)

Signal Name	Type	Description
DDRA_DQ[63:0]	I/O SSTL-2	**DDR Channel A Data Bus:** The DDR data bus provides the data from the DRAM devices.
DDRA_CB[7:0]	I/O SSTL-2	**DDR Channel A Check Bits:** These check bits are required to provide ECC Support
DDRA_DQSP[17:0] DDRA_DQSN[17:0] (DDR2 Only)	I/O SSTL-2	**DDR Channel A Data Strobes:** The DDR data strobes. Each data strobe is used to strobe a set of four or eight data signals (depending on whether x4 or x8 DRAM devices are used). These signals function as differential pairs for DDR2 technology only.
DDRA_CMDCLKP[3:0]/ DDRA_CMDCLKN[3:0]	O CMOS	**DDR Channel A Command Clock:** The DDR command clocks used by the DDR DRAMs to latch **MA[13:0], BA[2:0], RAS#, CAS#, WE#, CKE#,** and **CS#** signals.
DDRA_MA[13:0]	O CMOS	**DDR Channel A Memory Address:** The DDR memory address signals.
DDRA_BA[2:0]	O CMOS	**DDR Channel A Bank Address:** The DDR bank address signals. These signals are outputs of the MCH and select which bank within a row is selected.
DDRA_RAS#	O CMOS	**DDR Channel A Row Address Strobe:** Used to indicate a valid row address and open a row.
DDRA_CAS#	O CMOS	**DDR Channel A Column Address Strobe:** Used to indicate a valid column address and initiate a transaction.

Table 2-2. DDR Channel_A Interface Signals (Sheet 2 of 2)

Signal Name	Type	Description
DDRA_WE#	O CMOS	**DDR Channel A Write Enable:** Used to indicate a write cycle.
DDRA_CS[7:0]#	O CMOS	**DDR Channel A Chip Select:** Used to indicate to which DRAM device cycles are targeted.
DDRA_VREF	I Analog	**DDR Channel A Voltage Reference:** DDR Reference voltage input.

2.3 DDR Interface B Signals

Table 2-3. DDR Channel_B Interface Signals

Signal Name	Type	Description
DDRB_DQ[63:0]	I/O SSTL-2	**DDR Channel B Data Bus:** The DDR data bus provides the data from the DRAM devices.
DDRB_CB[7:0]	I/O SSTL-2	**DDR Channel B Check Bits:** These check bits are required to provide ECC Support
DDRB_DQSP[17:0] DDRB_DQSN[17:0] (DDR2 Only)	I/O SSTL-2	**DDR Channel B Data Strobes:** The DDR data strobes. Each data strobe is used to strobe a set of four data signals. These signals function as differential pairs for DDR2 technology only.
DDRB_CMDCLKP[3:0]/ DDRB_CMDCLKN[3:0]	O CMOS	**DDR Channel B Command CLOCK:** The DDR command clocks used by the DDR DRAMs to latch **MA[13:0], BA[2:0], RAS#, CAS#, WE#, CKE#,** and **CS#** signals.
DDRB_RAS#	O CMOS	**DDR Channel B Row Address Strobe:** Used to indicate a valid row address and open a row.
DDRB_CAS#	O CMOS	**DDR Channel B Column Address Strobe:** Used to indicate a valid column address and initiate a transaction.
DDRB_BA[2:0]	O CMOS	**DDR Channel B Bank Address:** The DDR bank address signals. These signals are outputs of the MCH and select which bank within a row is selected.
DDRB_MA[13:0]	O CMOS	**DDR Channel B Memory Address:** The DDR memory address signals.
DDRB_WE#	O CMOS	**DDR Channel B Write Enable:** Used to indicate a write cycle.
DDRB_CS[7:0]#	O CMOS	**DDR Channel B Chip Select:** Used to indicate to which DRAM device cycles are targeted.
DDRB_VREF	I Analog	**DDR Channel B Voltage Reference:** DDR Reference voltage input.

2.4 DDR Interface Shared Signals

Table 2-4. DDR Interface Shared Signals (Sheet 1 of 2)

Signal Name	Type	Description
DDR_CRES0	I/O Analog	**DDR Compensation Resistor Return:** Common return for DDR interface compensation resistors on **DDRSLWCRES** and **DDRIMPCRES**.
DDR_SLWCRES	I/O Analog	**Compensation Resistor:** Slew rate compensation for DDR interface

Table 2-4. DDR Interface Shared Signals (Sheet 2 of 2)

Signal Name	Type	Description
DDR_IMPCRES	I/O Analog	**Compensation Resistor:** Impedance compensation for DDR interface
DDRCKE[7:0]	O CMOS	**DDR Clock Enable:** Two operating modes exist.: • Shared across channels in lock-step, where [7:0] correspond to the rows as defined in the DRAM Row Boundary registers (two CKE signals for each DIMM slot) • Independent per channel, with one CKE for each DIMM slot.
DDRRES[2:1]	I/O Analog	**DDR Resistor**: Additional compensation resistors for DDR interface

2.5 PCI Express* Interface Port A Signals

Note: For product scheduling reasons, the MCH was unable to adopt the PCI Express* signal naming convention ratified by the PCI SIG in Q3 of 2002. Refer to the following conversion table to reconcile between the MCH and PCI SIG naming conventions.

Table 2-5. PCI Express* Naming Convention Definitions

MCH Convention:	EXP_<port>_<T/R>X<P/N>[lane #]
PCI SIG Convention:	PE<port><T/R><p/n><lane #>

Where:

- T/R defines a transmit (T) or receive (R) signal
- P/N or p/n defines the polarity of the differential signal –p (+), n (–)

For example, Table 2-6 lists equivalent signal names across the two conventions:

Table 2-6. Example Naming Convention Conversions

MCH Convention	PCI SIG Convention
EXP_A_TXN[4]	PEATn4
EXP_A_RXp[2]	PEARp2
EXP_B_RXN[7]	PEBRn7

Table 2-7. PCI Express* Interface Port A Signals

Signal Name	Type	Description
EXP_A_TXP[7:0] EXP_A_TXN[7:0]	O DIFF	**PCI Express* Interface Port A Output (transmit) Data Pair** (differential). These signals are output (TX) from the MCH's perspective, and should be connected to the input (receive or RX) signals of the other PCI Express* device. Note that the signals for this x8 interface (EXP_A_TX*[7:0]) can be split into two x4 interfaces: • EXP_A_TX*[3:0] -> EXP_A0_TX*[3:0] • EXP_A_TX*[7:4] -> EXP_A1_TX*[3:0]
EXP_A_RXP[7:0] EXP_A_RXN[7:0]	I DIFF	**PCI Express* Interface Port A Input (receive) Data Pair** (differential). These signals are input (RX) from the MCH's perspective, and should be connected to the output (transmit or TX) signals of the other PCI Express* device. Note that the signals for this x8 interface (EXP_A_RX*[7:0]) can be split into two x4 interfaces: • EXP_A_RX*[3:0] -> EXP_A0_RX*[3:0] • EXP_A_RX*[7:4] -> EXP_A1_RX*[3:0]

2.6 PCI Express* Interface Port B Signals

Table 2-8. PCI Express* Interface Port B Signals

Signal Name	Type	Description
EXP_B_TXP[15:0] EXP_B_TXN[15:0]	O DIFF	**PCI Express* Interface Port B Output (transmit) Data Pair** (differential). These signals are output (TX) from the MCH's perspective, and should be connected to the input (receive or RX) signals of the other PCI Express* device.
EXP_B_RXP[15:0] EXP_B_RXN[15:0]	I DIFF	**PCI Express* Interface Port B Input (receive) Data Pair** (differential). These signals are input (RX) from the MCH's perspective, and should be connected to the output (transmit or TX) signals of the other PCI Express* device.

2.7 PCI Express* Interface Shared Signals

Table 2-9. PCI Express* Interface Shared Signals

Signal Name	Type	Description
EXP_COMP[1:0]	I/O CMOS	**PCI Express* Compensation:** Used to calibrate the PCI Express* **high-speed serial input/output** buffers
EXP_CLKP EXP_CLKN	I DIFF	**PCI Express* Clock Reference Input** (differential)

2.8 Hub Interface Signals

Table 2-10. Hub Interface Signals

Signal Name	Type	Description
HI[11:0]	I/O CMOS	**Hub Interface Signals**: Interface between the MCH and the ICH.
HI_STBF	I/O CMOS	**Hub Interface Strobe First**: First of two strobe signals used to transmit and receive data through the Hub Interface.
HI_STBS	I/O CMOS	**Hub Interface Strobe Second:** Second of two strobe signals used to transmit and receive data through the Hub Interface.
HICOMP	I Analog	**Hub Interface Compensation:** Used for Hub Interface buffer compensation
HI_VSWING	I Analog	**Hub Interface Voltage Swing:** Reference voltage input for the Hub Interface
HIVREF	I Analog	**HI Reference:** Reference voltage input for the Hub Interface
HICLK	I CMOS	**66 MHz Clock In:** This pin receives a 66 MHz clock from the clock synthesizer.

2.9 Reset, Power, and Miscellaneous Signals

The voltage reference pins are described in the signal description sections for the associated interface.

Table 2-11. Reset, Power, and Miscellaneous Signals (Sheet 1 of 2)

Signal Name	Type	Description
SMBus Interface		
SMBSCL	I/O OD (3V) SCHMITT	**SMBus Clock:** Clock signal for the SMBus interface
SMBSDA	I/O OD (3V) SCHMITT	**SMBus Data:** Data signal for the SMBus interface (3.3V)
Power Pins		
VCC	I	**Power**: 1.5V power supply pins for the Hub interface and the MCH core.
VCCDDR	I	**Power**: Power supply pins for the DDR interfaces. Voltage is 2.5V for DDR(266 or 333), or 1.8V for DDR2-400 technology.
VCCEXP	I	**Power**: 1.5V power supply pins for the PCI Express* interface.
VTT	I	**Power**: Power supply pins for the system bus interface. Voltage is the same as processor VTT.
VCCA_CORE	I	**Power**: 1.5V analog VCC supply
VCCA_DDR	I	**Power**: 1.5V analog VCC supply
VCCA_EXP	I	**Power**: 1.5V analog VCC supply
VCCA_HI	I	**Power**: 1.5V analog VCC supply
VSSA_CORE	I	**Ground**: Analog ground pin
VSSA_EXP	I	**Ground**: Analog ground pin
VSSA_HI	I	**Ground**: Analog ground pin

Table 2-11. Reset, Power, and Miscellaneous Signals (Sheet 2 of 2)

Signal Name	Type	Description
VSS	I	**Ground:** Digital ground pins
V3REF	I	Signal-routed 3.3V supply.
EXP_VCCBG	I	**Power**: 2.5V PCI Express* analog voltage supply
EXP_VSSBG	I	**Ground**: PCI Express* analog ground
JTAG Interface		
TRST#	I SCHMITT	JTAG Tap Reset
TMS	I SCHMITT	JTAG Tap Mode Select
TDI	I SCHMITT	JTAG Serial Data Input
TCK	I SCHMITT	JTAG Clock
TDO	OD (1.2 – 1.35V)	JTAG Serial Data Output
Miscellaneous Signals		
PLLSEL[1:0]#	I CMOS	**PLL Select**: Used by the MCH at power-on to configure the internal timing relationship between the DRAM interface and the system bus. All other encodings are reserved by Intel. PLLSEL1# / PLLSEL0# DDR266: Open (High) / Open (High) DDR333: Ground / Ground DDR2-400: Ground / Ground Though the **PLLSEL[1:0]#** pins are active low, the table encodings reflect actual voltage values measured at the pins.
RSTIN#	I SCHMITT	**Reset Input**: **RSTIN#** assertion asynchronously resets the MCH logic, except for certain “sticky” bits. Typically connected to the ICH **PCIRST#** output.
PWRGD	I SCHMITT	**Power Good**: Asynchronously resets the entire MCH component, including “sticky” bits. Driven by system logic to indicate all board power supplies are valid.
MCHPME#	O OD (3V)	**MCH Power Management Event:** Allows devices connected to the MCH’s PCI Express* ports to request power state changes via the ICH **PME#** signal.
MCHGPE#	O OD (3V)	**MCH General Purpose Event:** Typically connects to an ICH input capable of generating an ACPI GPE event.
TEST#	I	Internal Test Mode Select Pin
DEBUG[7:0]	O	XDP (Extended Debug Port) Interface Pins

Register Descriptions 3

The MCH contains two sets of software accessible registers, accessed via the Host processor I/O address space:

- Control registers – These registers are mapped into the processor I/O space, and control access to PCI configuration space (see Section 3.3, "I/O Mapped Registers" on page 3-36).
- Internal configuration registers – These registers, which reside within the MCH, are partitioned into multiple logical device register sets ("logical" since they reside within a single physical device). One register set is dedicated to Host-Hub Interface (HI) Bridge functionality (controls PCI_A, DRAM configuration, other chipset operating parameters and optional features). Additional sets of registers map to the virtual PCI-to-PCI bridges for the PCI Express* ports.

The MCH internal registers (I/O Mapped and Configuration registers) are accessible by the host processor. The registers can be accessed as Byte (8-bit), Word (16-bit), or Dword (32-bit) quantities, with the exception of the CONFIG_ADDRESS register, which can only be accessed as a Dword. All multi-byte numeric fields use "little-endian" ordering (i.e., lower address bits contain the least significant parts of the field).

3.1 Register Terminology

Type	Description
RO	**Read Only** – Software/BIOS can only read registers/bits with this attribute. Contents either hardwired or set by hardware. Attempted Writes to RO locations have no effect.
WO	**Write Only** – Attempted Reads of registers/bits with this attribute do not return valid data. Writes to these locations cause some hardware event to take place.
R/W	**Read / Write** – Software/BIOS can read and write to registers/bits with this attribute.
R/WC	**Read / Write Clear** – Software/BIOS can read and write to registers/bits with this attribute. However, writing a value of 0 to a bit with this attribute has no effect. A R/WC bit can only be set to a value of 1 by a hardware event. To clear a R/WC bit (i.e., change the value to '0'), a value of 1 must be written to the bit location.
R/WS	**Read / Write Set** – Software/BIOS can read and write to registers/bits with this attribute. However, writing a value of 0 to a bit with this attribute has no effect. Software can set a R/WS bit by writing a value of 1 to the location, but the bit can only be cleared to '0' by hardware.
R/W/L	**Read / Write / Lock** – Software/BIOS can read and write to registers/bits with this attribute. Hardware or some other configuration bit can lock a R/W/L bit and prevent it from being updated.
R/WO	**Read / Write Once** – Software/BIOS can read, but only write to a register/bit with this attribute once after reset. It is a special form of RWL. Once a R/WO bit has been written, it is locked and subsequent attempts to Write this bit will have no effect unless the system is reset.
Reserved Bits	Some of the MCH registers described in this section contain reserved bits. These bits are labeled "Reserved". Software must deal correctly with fields that are reserved. On reads, software must use appropriate masks to extract the defined bits and not rely on reserved bits being any particular value. On writes, software must ensure that the values of reserved bit positions are preserved. That is, the values of reserved bit positions must first be read, merged with the new values for other bit positions and then written back. The software does not need to perform read, merge, write operation for the configuration address register.

Type	Description
Reserved Registers	Registers that are marked as “Reserved” must not be modified by system software. Reads to “Reserved” registers may return non-zero values. Writes to “reserved” registers may cause system failure.
Default Value upon Reset	Upon a Reset, the MCH sets its entire internal configuration registers to predetermined default states. Some register values at reset are determined by external strapping options. A register's default value represents the minimum functionality feature set required to successfully bring up the system. Hence, it does not represent the optimal system configuration. It is the responsibility of the system initialization software (usually BIOS) to properly determine the DRAM configurations, operating parameters and optional system features that are applicable, and to program the MCH registers accordingly.
Sticky Registers	Certain registers in the MCH are sticky through a hard reset. They will only be reset on a Power Good reset. These registers, in general, are the error logging registers and a few special cases. The error command registers are not sticky, in order to avoid incorrect error reporting through a mechanism that has not yet been set up in code. Only those registers that are explicitly marked as such are sticky. The Device 0, Function 1 error logging registers are sticky. The command registers are not.

3.1.1 Platform Configuration

The MCH and the ICH are physically connected by the legacy Hub Interface (HI). From a configuration standpoint, the HI is logically PCI bus 0. As a result, all devices internal to the MCH and ICH appear to be on PCI bus 0. The system’s primary PCI expansion bus is physically attached to the ICH and, from a configuration perspective, appears to be a hierarchical PCI bus behind a PCI-to-PCI bridge and therefore has a programmable bus number.

The primary PCI bus is referred to in this document as PCI_A, and is not PCI bus 0 from a configuration standpoint.

The PCI Express* ports appear to system software as PCI buses behind PCI-to-PCI bridges that reside as devices on PCI bus 0.

The MCH decodes multiple PCI device numbers. The configuration registers for the devices are mapped as devices residing on PCI bus 0. Each device number may contain multiple functions.

The PCI predefined header has five fields that deal with device identification. All devices are required to implement these fields. Generic configuration software will easily be able to identify devices that are available for use. These registers are read-only. The five fields are Vendor ID, Device ID, Revision ID, Header Type, and Class Code.

The 16-bit Vendor ID value assigned to Intel is 8086h.

The following table shows the device # assignment for the various internal MCH devices:

Table 3-1. PCI Device Number Assignment (Sheet 1 of 2)

MCH Function	Device #, Function #
Memory Controller Hub (Hub Interface) MCH Error Reporting (Hub Interface)	Device 0, Function 0 Device 0, Function 1
Host-to-PCI Express* A Bridge (x8 or x4)	Device 2
Host-to-PCI Express* A1 Bridge (x4 only)	Device 3

Table 3-1. PCI Device Number Assignment (Sheet 2 of 2)

MCH Function	Device #, Function #
Host-to-PCI Express B Bridge (x16, x8, x4)	Device 4
Extended Configuration Registers	Device 8

A disabled or non-existent device's configuration register space is hidden. A disabled or non-existent device will return all ones for reads and will drop writes just as if the cycle terminated with a Master Abort on PCI.

When a PCI Express* interface is unpopulated or fails to train, the associated configuration register space is hidden, returning all ones for all registers just as if the cycle terminated with a Master Abort on PCI. Also, if PCI Express* port A is configured for x8 operation rather than x4, PCI Express* port A1 configuration space will be hidden.

Logically, the ICH appears as multiple PCI devices within a single physical component also residing on PCI bus 0. One of the ICH devices is a PCI-to-PCI bridge. Logically, the primary side of the bridge resides on PCI bus 0 while the secondary is a standard PCI expansion bus.

Note: A physical PCI bus 0 does not exist and that the HI, the internal devices in the MCH and the ICH logically constitute PCI Bus 0 to configuration software.

3.2 General Routing Configuration Accesses

The MCH is responsible for routing PCI configuration cycles to the proper interface. PCI configuration cycles to the ICH internal devices and Primary PCI (including downstream devices) are routed to the ICH via the HI (Hub Interface). Routing of configuration accesses to any of the PCI Express* ports is controlled via the standard PCI-to-PCI bridge mechanism using information contained within the PRIMARY BUS NUMBER, the SECONDARY BUS NUMBER, and the SUBORDINATE BUS NUMBER registers of the corresponding PCI-to-PCI bridge device.

A detailed description of the mechanism for translating processor I/O bus cycles to configuration cycles on one of the buses is described below.

Note: The MCH supports a variety of connectivity options. When any of the MCH's interfaces are disabled, the associated interface's device registers are not visible. Attempted configuration cycles to these registers will return all ones for Reads, and Master Abort for writes.

Figure 3-1. PCI Device Map

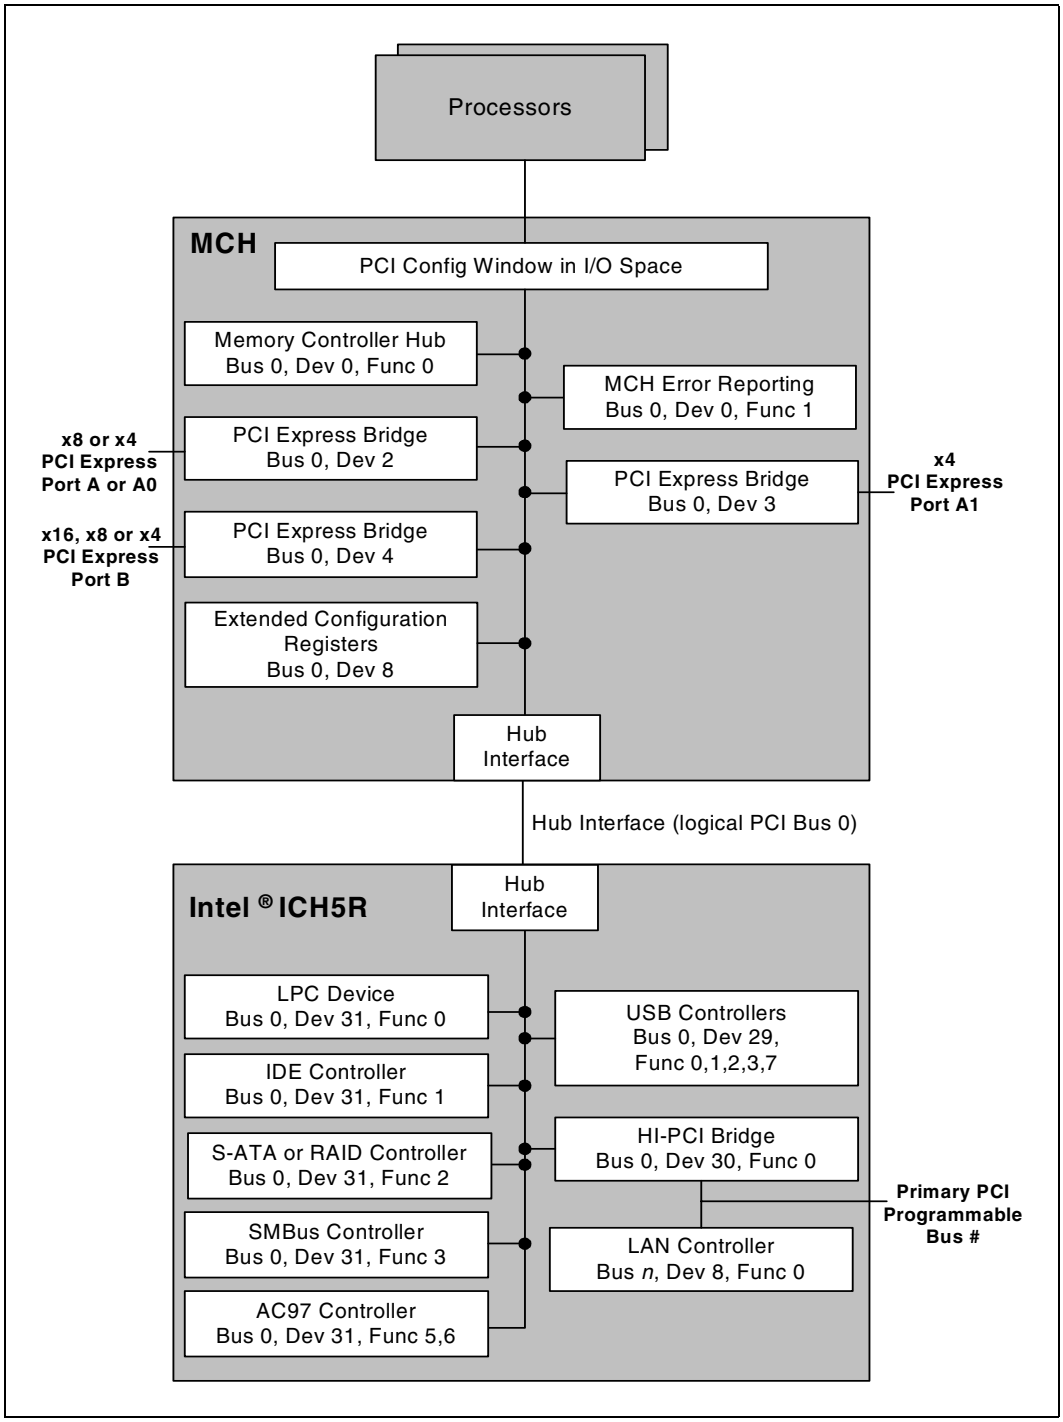

3.2.1 Standard PCI Bus Configuration Mechanism

The PCI Bus defines a slot-based configuration space that allows each device to contain up to eight functions; each function containing up to 256, 8-bit configuration registers. The *PCI Local Bus Specification, Rev 2.3* defines two bus cycles to access the PCI configuration space: Configuration Read and Configuration Write. Memory and I/O spaces are supported directly by the processor. Configuration space is supported by a mapping mechanism implemented within the MCH. Previous versions of the *PCI Local Bus Specification* define two mechanisms to access configuration space, Mechanism 1 and Mechanism 2. **The MCH supports only Mechanism 1.**

The configuration access mechanism makes use of the CONFIG_ADDRESS register and CONFIG_DATA register. To reference a configuration register, a Dword (32-bit) I/O write cycle is used to place a value into CONFIG_ADDRESS that specifies the PCI bus, the device on that bus, the function within the device, and a specific configuration register of the device function being accessed. CONFIG_ADDRESS[31] must be 1 to enable a configuration cycle. CONFIG_DATA then becomes a window into the four bytes of configuration space specified by the contents of CONFIG_ADDRESS. Any read or write to CONFIG_DATA will result in the MCH translating the CONFIG_ADDRESS into the appropriate configuration cycle.

The MCH is responsible for translating and routing the processor's I/O accesses to the CONFIG_ADDRESS and CONFIG_DATA registers to internal MCH configuration registers for the HI and PCI Express*.

3.2.2 Logical PCI Bus 0 Configuration Mechanism

The MCH decodes the Bus Number (bits 23:16) and the Device Number fields of the CONFIG_ADDRESS register. If the Bus Number field of CONFIG_ADDRESS is 0 the configuration cycle is targeting a PCI Bus 0 device.

Configuration cycles to any of the MCH's enabled internal devices are confined to the MCH and not sent over the HI. Accesses to disabled or non-existent internal devices are forwarded over the HI as Type 0 Configuration Cycles. The ICH decodes the Type 0 access and generates a configuration access to the selected internal device.

3.2.3 Primary PCI and Downstream Configuration Mechanism

If the Bus Number in the CONFIG_ADDRESS is non-zero, and does not lie between the SECONDARY BUS NUMBER register and the SUBORDINATE BUS NUMBER register for one of the PCI Express* bridges, the MCH will generate a Type 1 HI Configuration Cycle.

When the cycle is forwarded to the ICH via the HI, the ICH compares the non-zero Bus Number with the SECONDARY BUS NUMBER and SUBORDINATE BUS NUMBER registers of its PCI-to-PCI bridges to determine if the configuration cycle is meant for the Primary PCI, or a downstream PCI bus.

3.2.4 PCI Express* Bus Configuration Mechanism

From the chipset configuration perspective, the PCI Express* ports are seen as PCI bus interfaces residing on a Secondary Bus side of the "virtual" PCI-to-PCI bridges referred to as the MCH Host-to-PCI Express* bridges. On the Primary bus side, the "virtual" PCI-to-PCI bridge is attached to PCI Bus 0. Therefore the PRIMARY BUS NUMBER register is hardwired to '0'. The "virtual" PCI-to-PCI bridge entity converts Type 1 PCI Bus Configuration cycles on PCI Bus 0 into Type 0

or Type 1 configuration cycles on the PCI Express* interfaces. Type 1 configuration cycles on PCI Bus 0 that have a BUS NUMBER that matches the SECONDARY BUS NUMBER of one of the MCH's "virtual" PCI-to-PCI bridges will be translated into Type 0 configuration cycles on the appropriate PCI Express* interface.

Note: There is no requirement that the secondary and sub-ordinate bus number values from one PCI Express* port be contiguous with another PCI Express* port. For example, it is possible that PCI Express_A will decode buses 2 through 5, and PCI Express_B will decode buses 8 through 12. In this case there is a gap where buses 6 and 7 are subtractively decoded to the HI.

If the Bus Number is non-zero, greater than the value programmed into the SECONDARY BUS NUMBER register, and less than or equal to the value programmed into the corresponding SUBORDINATE BUS NUMBER register the configuration cycle is targeting a PCI bus downstream of the targeted PCI Express* interface. The MCH will generate a Type 1 configuration cycle on the appropriate PCI Express* interface. The address bits will be mapped as described in Figure 3-2 "Type 1 Configuration Address to PCI Address Mapping" on page 3-36.

Figure 3-2. Type 1 Configuration Address to PCI Address Mapping

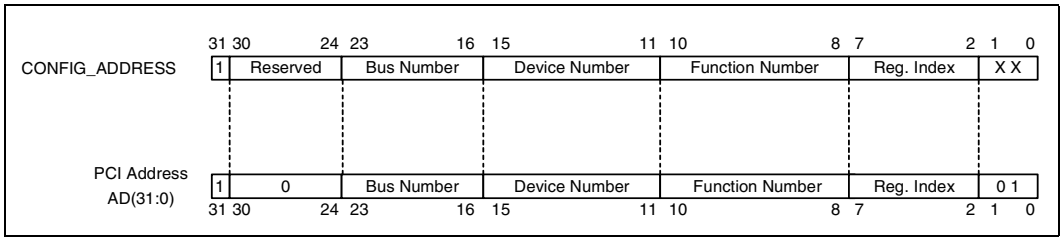

To prepare for mapping of the configuration cycles on PCI Express* the initialization software will go through the following sequence:

1. Scan all devices residing on the PCI Bus 0 using Type 0 configuration accesses.
2. For every device residing at bus 0 which implements PCI-to-PCI bridge functionality, it will configure the secondary bus of the bridge with the appropriate number and scan further down the hierarchy. This process will include the configuration of the "virtual" PCI-to-PCI bridges within the MCH used to map the PCI Express* devices' address spaces in a software specific manner.

3.3 I/O Mapped Registers

The MCH contains two registers that reside in the processor I/O address space – the Configuration Address (CONFIG_ADDRESS) register and the Configuration Data (CONFIG_DATA) register. The Configuration Address register enables/disables the configuration space and determines what portion of configuration space is visible through the Configuration Data window.

3.3.1 CONFIG_ADDRESS – Configuration Address Register

I/O Address: 0CF8h Accessed as a DWord
Access: R/W
Size: 32 Bits
Default: 00000000h

CONFIG_ADDRESS is a 32-bit register that can be accessed only as a DWord. A Byte or Word reference will “pass through” the Configuration Address register and HI onto the PCI_A bus as an I/O cycle. The CONFIG_ADDRESS register contains the Bus Number, Device Number, Function Number, and Register Number for which a subsequent configuration access is intended.

Bit	Descriptions
31	**Configuration Enable (CFGE)**. 0 = Accesses to PCI configuration space are disabled. 1 = Accesses to PCI configuration space are enabled.
30:24	Reserved (These bits are read only and have a value of 0).
23:16	**Bus Number.** Contains the bus number being targeted by the config cycle.
15:11	**Device Number.** Selects one of the 32 possible devices per bus.
10:8	**Function Number.** Selects one of eight possible functions within a device.
7:2	**Register Number.** This field selects one register within the particular Bus, Device, and Function as specified by the other fields in the Configuration Address register. This field is mapped to A[7:2] during HI or PCI Express* configuration cycles.
1:0	Reserved

3.3.2 CONFIG_DATA – Configuration Data Register

I/O Address: 0CFCh
Access: R/W
Size: 32 Bits
Default: 00000000h

CONFIG_DATA is a 32-bit read/write window into configuration space. The portion of configuration space that is referenced by CONFIG_DATA is determined by the contents of CONFIG_ADDRESS.

Bit	Descriptions
31:0	**Configuration Data Window (CDW).** If bit 31 of CONFIG_ADDRESS is set to ‘1’, any I/O access to the CONFIG_DATA register will be mapped to configuration space using the contents of CONFIG_ADDRESS.

3.4 PCI Express* Enhanced Configuration Mechanisms

PCI Express* extends the configuration space to 4096 bytes per device/function as compared to 256 bytes allowed by PCI 2.3 configuration space. PCI Express* configuration space is divided into a PCI 2.3 compatible region, which consists of the first 256 bytes of a logical device’s configuration space and an extended PCI Express* region which consists of the remaining configuration space. The PCI 2.3 compatible region can be accessed using either the mechanisms defined in the *PCI Local Bus Specification, Rev 2.3* or using the enhanced PCI Express* configuration access mechanism. All changes made using either access mechanism are equivalent; however, software is not allowed to interleave PCI Express* and PCI access mechanisms to access the configuration registers of devices. The extended PCI Express* region can only be accessed using the enhanced PCI Express* configuration access mechanism.

3.4.1 PCI Express* Configuration Transaction Header

The PCI Express* Configuration Transaction Header includes an additional four bits for the Register Number field (ExtendedRegisterAddress[3:0]) to provide additional configuration space.

Figure 3-3. PCI Express* Configuration Transaction Header

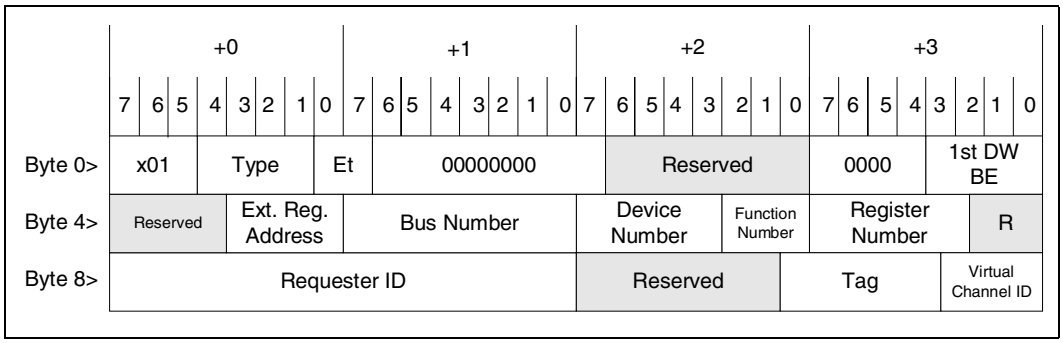

The PCI 2.3 compatible configuration access mechanism uses the same Request format as the enhanced PCI Express* mechanism. For PCI-compatible Configuration Requests, the Extended Register Address field must be all zeros.

To maintain compatibility with PCI configuration addressing mechanisms, system software must access the enhanced configuration space using DWORD operations (DWORD-aligned) only.

3.4.2 Enhanced Configuration Memory Address Map

The Enhanced Configuration Memory Address Map is positioned into memory space by use of the PCI Express* Enhanced Configuration Base register known as EXPECBASE. This register contains the address that corresponds to bits 31 to 28 of the base address for PCI Express* enhanced configuration space below 4 GB. Configuration software will read this register to determine where the 256 MB range of memory addresses resides for enhanced configuration. This register defaults to a value of Eh, which corresponds to E000 0000h. It is not intended that this value is ever changed by BIOS.

Figure 3-4. Enhanced Configuration Memory Address Map

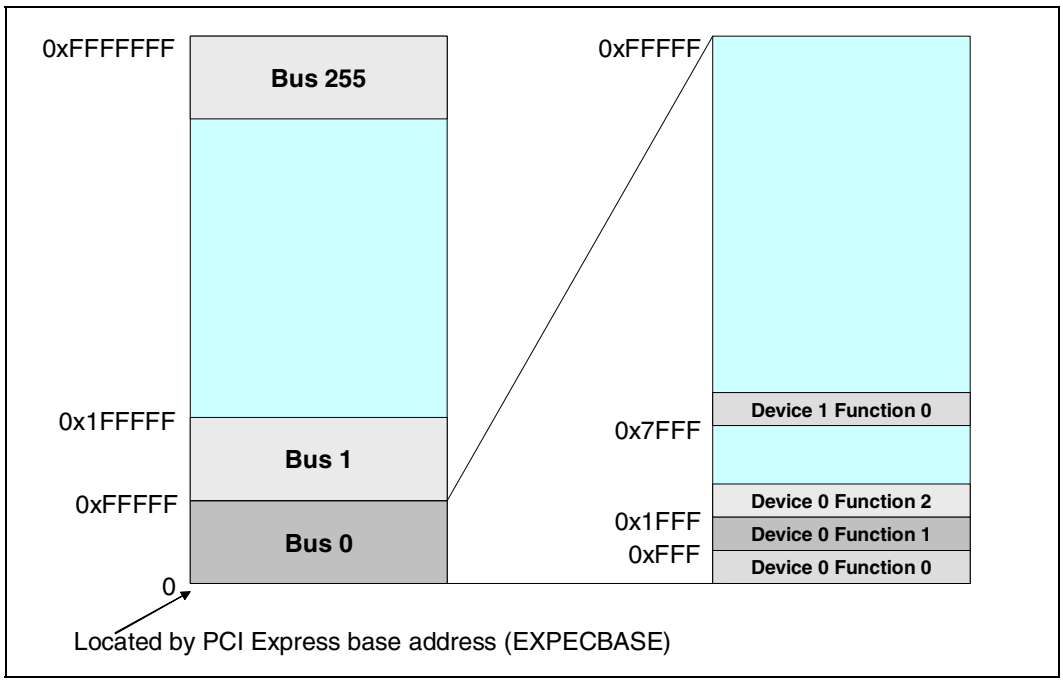

3.5 MCH Control Registers (D0:F0)

The MCH Control registers are in Device 0 (D0), Function 0 (F0). Table 3-2 provides the register address map for this device and function.

Warning: Address locations that are not listed are considered Reserved register locations. Reads to Reserved registers may return non-zero values. Writes to reserved locations may cause system failure.

Table 3-2. MCH Control PCI Configuration Register Map (D0:F0) (Sheet 1 of 2)

Offset	Mnemonic	Register Name	Access	Default
00 – 01h	VID	Vendor Identification	RO	8086h
02 – 03h	DID	Device Identification	RO	359Eh
04 – 05h	PCICMD	PCI Command Register	RO, R/W	0006h
06 – 07h	PCISTS	PCI Status Register	R/WC, RO	0090h
08h	RID	Revision Identification	RO	09h
0Ah	SUBC	Sub-Class Code	RO	00h
0Bh	BCC	Base Class Code	RO	06h
0Dh	MLT	Master Latency Timer	RO	00h
0Eh	HDR	Header Type	RO	80h
2C – 2Dh	SVID	Subsystem Vendor Identification	R/WO	0000h
2E – 2Fh	SID	Subsystem Identification	R/WO	0000h
34h	CAPPTR	Capabilities Pointer	RO	40h
50h	MCHCFG0	MCH Configuration 0	RO	0Ch

Table 3-2. MCH Control PCI Configuration Register Map (D0:F0) (Sheet 2 of 2)

Offset	Mnemonic	Register Name	Access	Default
52 – 53h	MCHSCRB	MCH Memory Scrub and Init Configuration	RO, R/W	0000h
58h	FDHC	Fixed DRAM Hole Control	R/W, RO	00h
59 – 5Fh	PAM 0:6	Programmable Attribute Map 0 – 6	R/W, RO	00h
60 – 67h	DRB 0:7	DRAM Row Boundary Register 0 – 7	R/W	00h
70 – 73h	DRA 0:3	DRAM Row Attribute Register 0 – 3	R/W	00h
78 – 7Bh	DRT	DRAM Timing Register	R/W	9599_9604h
7C – 7Fh	DRC	DRAM Controller Mode Register	RO, R/W, R/WO	0000_0008h
80 – 81	DRM	DRAM Mapping Register	R/W	8421h
82	DRORC	Opportunistic Refresh Control Register	R/W	71h
84 – 87h	ECCDIAG	ECC Detection/Correction Diagnostic Register	R/W, RO, R/WS	0000_0000h
88 – 8Bh	SDRC	DDR SDRAM Secondary Control Register	R/W	0000_0000h
8Ch	CKDIS	CK/CK# Disable	R/W	FFh
8Dh	CKEDIS	CKE/CKE# Disable	R/W	00h
9A – 9Bh	DDRCSR	DDR Channel Configuration Control/Status Register	RO, R/WS, R/W	0000h
9Ch	DEVPRES	Device Present	R/WO,RO	01h
9Dh	ESMRC	Extended System Management RAM Control	R/WL, R/W	00h
9Eh	SMRC	System Management RAM Control	R/W, RO, R/WL, R/WS	02h
9Fh	EXSMRC	Expansion System Management RAM Control	R/WC, RO	07h
B0 – B3h	DDR2ODTC	DDR2 ODT Control Register	R/W	0000_0000h
C4 – C5h	TOLM	Top of Low Memory Register	R/W	0800h
C6 – C7h	REMAPBASE	Remap Base Address Register	R/W	03FFh
C8 – C9h	REMAPLIMIT	Remap Limit Address Register	R/W	0000h
CA – CBh	REMAPOFFSET	Remap Offset	R/W	0000h
CC – CDh	TOM	Top of Memory Register	R/W	0000h
CE – CFh	EXPECBASE	PCI Express* Enhanced Config Base Address	R/WO	E000h
DE – DFh	SKPD	Scratchpad Data	R/W	0000h
F4h	DEVPRES1	Device Present 1 Register	RO, R/W	18h
F5	MCHTST	MCH Test Register	RO, R/W	01h

3.5.1 VID – Vendor Identification (D0:F0)

Address Offset: 00 – 01h
Access: RO
Size: 16 Bits
Default: 8086h

The VID register contains the vendor identification number. This 16-bit register combined with the Device Identification register uniquely identifies any PCI device. Writes to this register have no effect.

Bit Field	Default & Access	Description
15:0	8086h RO	**Vendor Identification (VID)**. This register field contains the PCI standard identification for Intel, 8086h.

3.5.2 DID – Device Identification (D0:F0)

Address Offset: 02 – 03h
Access: RO
Size: 16 Bits
Default: 359Eh

This 16-bit register combined with the Vendor Identification register uniquely identifies any PCI device. Writes to this register have no effect.

Bit Field	Default & Access	Description
15:0	359Eh RO	**Device Identification Number (DID)**. This is a 16-bit value assigned to the MCH Host-HI Bridge Function #0.

3.5.3 PCICMD – PCI Command Register (D0:F0)

Address Offset: 04 – 05h
Access: RO, R/W
Size: 16 Bits
Default: 0006h

Since MCH Device 0 does not physically reside on a real PCI bus, portions of this register are not implemented.

Bit Field	Default & Access	Description
15:10	00h	Reserved
9	0b RO	**Fast Back-to-Back Enable (FB2B).**This bit controls whether or not the master can do a fast back-to-back write. Since Device 0 is strictly a target, this bit is hardwired to '0'.
8	0b R/W	**SERR Enable (SERRE).** This is a global enable bit for Device 0 SERR messaging. The MCH does not have an SERR signal. The MCH communicates the SERR condition by sending an SERR message over the HI to the ICH. 0 = Disable. The SERR message is not generated by the MCH for Device 0. 1 = Enable. The MCH is enabled to generate SERR messages over HI for specific Device 0 error conditions that are individually enabled in the HI_SERRCMD register (D0:F1:5Ch). **NOTE:** This bit only controls SERR messaging for the following Device 0 errors: Hub Interface Data Parity Errors and Hub Interface Address/Command Parity Errors. These errors are reported via the Hub Interface SERR Command Registers (HI_SERRCMD, Bus 0, Device 0, Function 1, Offset 5Ch, bit 4 and 0 respectively). Devices 1-7 have their own SERR bits to control error reporting for error conditions occurring on their respective devices. The control bits are used in a logical OR manner to enable the SERR HI message mechanism.
7	0b RO	**Address/Data Stepping Enable (ADSTEP)**. Address/data stepping is not implemented in the MCH.

Bit Field	Default & Access	Description
6	0b R/W	**Parity Error Enable (PERRE)**. 0 = Disable. The MCH does not take any action when it detects a parity error on HI. 1 = Enable. The MCH generates an SERR message over the HI to the ICH when an address or data parity error is detected by the MCH on the HI (DPE set in PCISTS) and SERRE is set to '1'.
5	0b RO	**VGA Palette Snoop Enable (VGASNOOP).** The MCH does not implement this bit.
4	0b RO	**Memory Write and Invalidate Enable (MWIE).** The MCH will never issue memory write and invalidate commands.
3	0b RO	**Special Cycle Enable (SCE)**. The MCH does not implement this bit.
2	1b RO	**Bus Master Enable (BME)**. The MCH is always enabled as a master, so this bit is hardwired to '1'.
1	1b RO	**Memory Access Enable (MAE)**. This bit is hardwired to '1'.
0	0b RO	**I/O Access Enable (IOAE)**. This bit is not implemented in the MCH.

3.5.4 PCISTS – PCI Status Register (D0:F0)

Address Offset: 06 – 07h
Access: R/WC, RO
Size: 16 Bits
Default: 0090h

PCISTS is a 16-bit status register that reports the occurrence of error events on Device 0's PCI interface. Since MCH Device 0 does not physically reside on a real PCI bus, many of the bits are not implemented.

Bit Field	Default & Access	Description
15	0b R/WC	**Detected Parity Error (DPE)**. 0 = No parity error detected. 1 = MCH detected an address or data parity error on the HI interface.
14	0b R/WC	**Signaled System Error (SSE)**. 0 = Software clears this bit by writing a '1' to the bit location. 1 = MCH Device 0 generates an SERR message over HI for any enabled Device 0 error condition. Device 0 error conditions are enabled in the PCICMD and ERRCMD registers. Device 0 error flags are read/reset from the PCISTS or Error registers.
13	0b RO	**Received Master Abort Status (RMAS)**. The ICH will never send up a Master Abort completion.
12	0b R/WC	**Received Target Abort Status (RTAS)**. 0 = Software clears this bit by writing a '1' to the bit location. 1 = MCH generated a HI request that received a Target Abort.
11	0b RO	**Signaled Target Abort Status (STAS)**. The MCH will not generate a Target Abort on the HI.

Bit Field	Default & Access	Description
10:9	00b RO	**DEVSEL Timing (DEVT).** Device 0 does not physically connect to PCI_A. These bits are set to 00 (fast decode) so that optimum DEVSEL timing for PCI_A is not limited by the MCH.
8	0b RO	**Master Data Parity Error Detected (DPD).** PERR signaling and messaging are not implemented by the MCH.
7	1b RO	**Fast Back-to-Back (FB2B).** Device 0 does not physically connect to PCI_A. This bit is set to '1' (indicating fast back-to-back capability) so that the optimum setting for PCI_A is not limited by the MCH.
6:5	00b	Reserved
4	1b RO	**Capability List (CLIST)**. Indicates to the configuration software that this device/function implements a list of new capabilities. A list of new capabilities is accessed via register CAPPTR at configuration address offset 34h.
3:0	0h	Reserved

3.5.5 RID – Revision Identification (D0:F0)

Address Offset: 08h
Access: RO
Size: 8 Bit
Default: 09h

This register contains the revision number of the MCH Device 0.

Bit Field	Default & Access	Description
7:0	00h R/WO	**Revision Identification Number (RID).** This value indicates the revision identification number for the MCH Device 0. 09h= C1 stepping. 0Ah= C2 stepping.

3.5.6 SUBC – Sub-Class Code (D0:F0)

Address Offset: 0Ah
Access: RO
Size: 8 Bits
Default: 00h

Bit Field	Default & Access	Description
7:0	00h RO	**Sub-Class Code (SUBC)**. This value indicates the Sub Class Code into which the MCH Device 0 falls. 00h = Host Bridge.

3.5.7 BCC – Base Class Code (D0:F0)

Address Offset: 0Bh
Access: RO
Size: 8 Bits
Default: 06h

Bit Field	Default & Access	Description
7:0	06h RO	**Base Class Code (BASEC)**. This value indicates the Base Class Code for the MCH Device 0. 06h = Bridge device.

3.5.8 MLT – Master Latency Timer (D0:F0)

Address Offset: 0Dh
Access: RO
Size: 8 Bits
Default: 00h

Device 0 in the MCH is not a PCI master. Therefore this register is not implemented.

Bit Field	Default & Access	Description
7:0	00h	Reserved

3.5.9 HDR – Header Type (D0:F0)

Address Offset: 0Eh
Access: RO
Size: 8 Bits
Default: 80h

Bit Field	Default & Access	Description
7:0	80h RO	**PCI Header (HDR)**. The header type of the MCH Device 0. 80h = multi-function device with standard header layout. 00h = single function device when function 1 is disabled.

3.5.10 SVID – Subsystem Vendor Identification (D0:F0)

Address Offset: 2C – 2Dh
Access: R/WO
Size: 16 Bits
Default: 0000h

This value is used to identify the vendor of the subsystem.

The MCH treats the SVID and SID as a single 32 bit register with regard to R/WO functionality. Any time a write access to the address 2C – 2Fh occurs, regardless of byte enables, entire 32 bit register comprised of SVID and SID locks.

Bit Field	Default & Access	Description
15:0	0000h R/WO	**Subsystem Vendor ID (SUBVID)**. This field should be programmed during boot-up to indicate the vendor of the system board.

3.5.11 SID – Subsystem Identification (D0:F0)

Address Offset: 2E – 2Fh
Access: R/WO
Size: 16 Bits
Default: 0000h

This value is used to identify a particular subsystem.

The MCH treats the SVID and SID as a single 32 bit register with regard to R/WO functionality. Any time a write access to the address 2C – 2Fh occurs, regardless of byte enables, entire 32 bit register comprised of SVID and SID locks.

Bit Field	Default & Access	Description
15:0	0000h R/WO	**Subsystem ID (SUBID)**. This field should be programmed during BIOS initialization.

3.5.12 CAPPTR – Capabilities Pointer (D0:F0)

Address Offset: 34h
Access: RO
Size: 8 Bits
Default: 40h

The CAPPTR provides the offset that is the pointer to the location where the first set of capabilities registers is located.

Bit Field	Default & Access	Description
7:0	40h RO	**Capabilities Pointer (CAP_PTR)**. Indicates the offset in configuration space for the location of the Capability Identification register block.

3.5.13 MCHCFG0 – MCH Configuration 0 (D0:F0)

Address Offset: 50h
Access: RO
Size: 8 Bits
Default: 0Ch

Bit Field	Default & Access	Description
7:3	01h	Reserved
2	1b RO	**In-Order Queue Depth (IOQD)**. This bit reflects the value sampled on HA[7]# on the deassertion of the CPURST#. It indicates the depth of the system bus in-order queue. 0 = HA[7]# has been sampled asserted (i.e., logic 1, or electrical low). The depth of the IOQ is set to '1' (i.e., no pipelining on the processor bus). HA[7]# may be driven low during CPURST# by an external source. 1 = HA[7]# was sampled as deasserted (i.e., logic 0 or electrical high). The depth of the processor bus in-order queue is configured to the maximum (i.e., 12).
1:0	00b	Reserved

3.5.14 MCHSCRB – MCH Memory Scrub and Initialization Configuration Register (D0:F0)

Address Offset: 52 – 53h
Access: RO, R/W
Size: 16 Bits
Default: 0000h

This register contains control and status bit for the DRAM memory scrubber, and status bits for the DRAM thermal management feature.

Bit Field	Default & Access	Description
15:10	000000b	Reserved
9	0b R/W	**Thermally Managed-Write Occurred**. 0 = Writing a zero clears this bit. 1 = This bit is set by hardware when a write is thermally managed. This happens when the maximum allowed number of writes has been reached during a time-slice and there is at least one more write to be completed.
8	0b R/W	**Thermally Managed-Read Occurred**. 0 = Writing a zero clears this bit. 1 = This bit is set by hardware when a read is thermally managed. This happens when the maximum allowed number of reads has been reached during a time-slice and there is at least one more read to be done.
7	0b	Reserved
6:5	00b RO	**Scrub Completion Counter**. This field reflects scrub iterations. Upon the first scrub completion, this would increment to 01, subsequent scrub rollovers would increment this field through values 10,11,00,01 in sequence.

Bit Field	Default & Access	Description
4:2	0h	Reserved
1:0	0b R/W	**Scrub Mode Enable**. 00 = Scrub Engine turned off 01 = Reserved 10 = Scrub Engine turned on 11 = Reserved

3.5.15 FDHC – Fixed DRAM Hole Control (D0:F0)

Address Offset: 58h
Access: RO, R/W
Size: 8 Bits
Default: 00h

This eight-bit register controls a fixed DRAM hole from 15 MB–16 MB.

Bit Field	Default & Access	Description
7	0b R/W	**Hole Enable (HEN)**. This field enables a memory hole in DRAM space. The DRAM that lies "behind" this space is not remapped. 0 = No memory hole. 1 = Memory hole from 15 MB to 16 MB. Accesses in this range will be sent to the HI
6:0	00h	Reserved

3.5.16 PAM 0:6 – Programmable Attribute Map Registers 0 – 6 (D0:F0)

Address Offset: 59 – 5Fh (PAM0 – PAM6)
Access: R/W,RO
Size: 8 Bits Each
Default: 00h

The MCH allows programmable memory attributes on 13 legacy memory segments of various sizes in the 768 KB to 1 MB (C0000h – FFFFFh) address range. Seven Programmable Attribute Map (PAM) registers are used to support these features. Not all seven of these registers are identical. PAM0 controls only one segment (high), while PAM[1:6] each control two segments (high and low). Cacheability of these areas is controlled via the MTRR registers in the processor. The following two bits are used to specify the memory attributes for each memory segment:

RE Read Enable. When RE = 1, the processor read accesses to the corresponding memory segment are claimed by the MCH and directed to main memory. Conversely, when RE = 0, the host read accesses are directed to PCI_A.

WE Write Enable. When WE = 1, the host write accesses to the corresponding memory segment are claimed by the MCH and directed to main memory. Conversely, when WE = 0, the host write accesses are directed to PCI_A.

Together, these two bits specify memory attributes (Read Only, Write Only, Read/Write and Disabled) for each memory segment. These bits only apply to host-initiated access to the PAM areas. The MCH forwards to main memory any PCI Express* initiated accesses to the PAM areas. At the time such PCI Express* accesses to the PAM region may occur, the targeted PAM segment must be programmed to Read/Write. It is illegal to issue a PCI Express* initiated transaction to a PAM region with the associated PAM register not set to Read/Write.

As an example, consider a BIOS that is implemented on the expansion bus. During the initialization process, BIOS can be shadowed to main memory to increase system performance. When BIOS is shadowed to main memory it should be copied to the same address location. To shadow the BIOS, the attributes for that address range should be set to Write Only. The BIOS is shadowed by first doing a read of that address, which is forwarded to the expansion bus. The host then does a write of the same address, which is directed to main memory. After BIOS is completely

shadowed, the attributes for that memory area are changed to Read Only so that all writes are forwarded to the expansion bus. Figure 3-5 and Table 3-3 show the PAM registers and the associated attribute bits.

Bit Field	Default & Access	Description
7:6	00b	Reserved
5:4	00b R/W	**High Attribute Register (HIENABLE)**. This field controls the steering of read and write cycles that address the BIOS 0 0= DRAM Disabled – All accesses are directed to the HI 0 1= Read Only – All reads are serviced by DRAM. Writes are forwarded to HI 1 0= Write Only – All writes are sent to DRAM. Reads are serviced by the HI 1 1= Normal DRAM Operation – All reads and writes are serviced by DRAM
3:2	00b	Reserved
1:0	00b R/W	**Low Attribute Register (LOENABLE)**. This field controls the steering of read and write cycles that address the BIOS 0 0= DRAM Disabled – All accesses are directed to the HI 0 1= Read Only – All reads are serviced by DRAM. Writes are forwarded to HI 1 0= Write Only – All writes are sent to DRAM. Reads are serviced by the HI 1 1= Normal DRAM Operation – All reads and writes are serviced by DRAM **NOTE:** The Low segment for PAM0 is Reserved, as shown in Figure 3-5

Figure 3-5. PAM Associated Attribute Bits

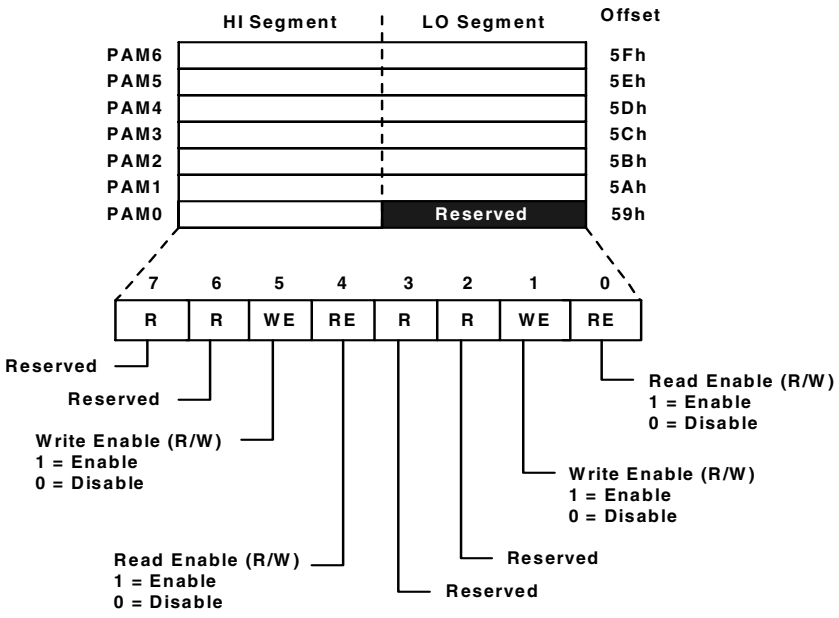

Table 3-3. PAM Associated Attribute Bits

PAM Reg	Attribute Bits		Memory Segment	Comments	Offset
PAM0 3:0, 7:6	Reserved		—	Reserved	59h
PAM0 5:4	WE	RE	0F0000h–0FFFFFh (64k)	BIOS Area	59h
PAM1 3:2, 7:6	Reserved		—	Reserved	5Ah
PAM1 1:0	WE	RE	0C0000h–0C3FFFh (16k)	BIOS Area	5Ah
PAM1 5:4	WE	RE	0C4000h–0C7FFFh (16k)	BIOS Area	5Ah
PAM2 3:2, 7:6	Reserved		—	Reserved	5Bh
PAM2 1:0	WE	RE	0C8000h–0CBFFFh (16k)	BIOS Area	5Bh
PAM2 5:4	WE	RE	0CC000h–0CFFFFh (16k)	BIOS Area	5Bh
PAM3 3:2, 7:6	Reserved		—	Reserved	5Ch
PAM3 1:0	WE	RE	0D0000h–0D3FFFh (16k)	BIOS Area	5Ch
PAM3 5:4	WE	RE	0D4000h–0D7FFFh (16k)	BIOS Area	5Ch
PAM4 3:2, 7:6	Reserved		—	Reserved	5Dh
PAM4 1:0	WE	RE	0D8000h–0DBFFFh (16k)	BIOS Area	5Dh
PAM4 5:4	WE	RE	0DC000h–0DFFFFh (16k)	BIOS Area	5Dh
PAM5 3:2, 7:6	Reserved		—	Reserved	5Eh
PAM5 1:0	WE	RE	0E0000h–0E3FFFh (16k)	BIOS Extension	5Eh
PAM5 5:4	WE	RE	0E4000h–0E7FFFh (16k)	BIOS Extension	5Eh
PAM6 3:2, 7:6	Reserved		—	Reserved	5Fh
PAM6 1:0	WE	RE	0E8000h–0EBFFFh (16k)	BIOS Extension	5Fh
PAM6 5:4	WE	RE	0EC000h–0EFFFFh (16k)	BIOS Extension	5Fh

3.5.17 DRB 0:7 – DRAM Row Boundary Register 0 - 7 (D0:F0)

Address Offset: 60 – 67h
Access: R/W
Size: 8 Bits Each
Default: 00h

The DRAM Row Boundary register defines the upper boundary address of each DRAM row with a granularity of 64 MB in single-channel mode or 128 MB in dual-channel mode. Each row has its own, single-byte DRB. The value in a given DRB corresponds to the cumulative memory size in the rows up to and including the corresponding row. For example, in dual-channel mode, a value of 1 (0000 0001) in DRB0 indicates that 128 Mbyte of DRAM has been populated in the first row.

Bit Field	Default & Access	Description
7:0	00b R/W	**DRAM Row Boundary Address.** This value defines the upper addresses for each row of DRAM. This value is compared against a set of address lines to determine the upper address limit of a particular row. This field corresponds to bits 34:27 of the address in dual channel mode and bits 33:26 in single channel mode.

DIMM #	Even Row (or single-sided)		Odd Row (present if double-sided)	
	Row Number	Address of DRB	Row Number	Address of DRB
1	Row 0	60h	Row 1	61h
2	Row 2	62h	Row 3	63h
3	Row 4	64h	Row 5	65h
4	Row 6	66h	Row 7	67h

DRB0 = Total memory in row0 (in 64-MB or 128-MB increments)
DRB1 = Total memory in row0 + row1 (in 64-MB or 128-MB increments)
...
DRB7 = Total memory in row0 + row1 + row2 + row3 + row4 + row5 + row6 + row7
(in 64-MB or 128-MB increments)

The row referred to by this register is defined by the DIMM chip select used during non-symmetric mode. Double-rank DIMMs use both Row 0 and Row 1, even though there is one physical slot for the two rows. Single rank DIMMs use only the even row number, since single sided DIMMs only support CS0#. For single-rank DIMMs the value BIOS places in the odd row should equal the same value as what was placed in the even row field. A row in single-channel mode is the 64-bit (72-bit with ECC) wide interface consisting of one rank of one DIMM. A row in dual channel mode is the 128-bit (144-bit with ECC) wide interface consisting one rank of two identical DIMMs.

Unpopulated rows must be programmed with the value of the last populated row.

The functionality of DRB7 is somewhat different from that of DRB[6:0]. If DRB(6:0) are non-zero and DRB7 is set to a value of 00h after memory configuration, the value is interpreted as 100h. This functionality avoids a 64MB/128-MB “hole” at the top of memory. This behavior is unique to the DRB7 register.

Programming Example:

Configuration is dual channel (128-MB granularity), with DIMMs populated as in the table below:

DIMM pair 1	256 MB in even row, none in odd row (single-rank DIMM)
DIMM pair 2	512 MB in even row, 512 MB in odd row (double-rank DIMM)
DIMM pair 3	128 MB in even row, 128 MB in odd row (double-rank DIMM)
DIMM pair 4	256 MB in even row, 256 MB in odd row (double-rank DIMM)

Address	Row	Size of Row	Accumulative Size	Register Value
60h	Row 0 (DIMM 1, even)	256 MB	256 MB	02h
61h	Row 1 (DIMM 1, odd)	empty	256 MB	02h
62h	Row 2 (DIMM 2, even)	512 MB	768 MB	06h
63h	Row 3 (DIMM 2, odd)	512 MB	1280 MB	0Ah
64h	Row 4 (DIMM 3, even)	128 MB	1408 MB	0Bh
65h	Row 5 (DIMM 3, odd)	128 MB	1536 MB	0Ch

Address	Row	Size of Row	Accumulative Size	Register Value
66h	Row 6 (DIMM 4, even)	256 MB	1792 MB	0Eh
67h	Row 7 (DIMM 4, odd)	256 MB	2048 MB	10h

3.5.18 DRA 0:3 – DRAM Row Attribute Register 0 - 3 (D0:F0)

Address Offset: 70 – 73h
Access: R/W
Size: 8 Bits Each
Default: 00h

The DRAM Row Attribute register defines the DRAM technology and the DQ/DQS signal mapping to be used for each row of memory. Each nibble of information in the DRA registers describes the page size of a row. For this register, a row is defined by the chip select used by the DIMM, so that a double-sided DIMM would have both an even and an odd entry. For single-sided DIMMs, only the even side is used. See Table 3-4 on page 3-54.

Bit Field	Default & Access	Description
7:6	00b R/W	**Device Width for Odd-numbered Row.** BIOS sets this bit according to the width of the DDR-SDRAM devices populated in this row. This is used to determine the page size, as well as the DQS to DQ signal mapping. 00 = Reserved 01 = x8 DDR-SDRAM 10 = x4 DDR-SDRAM 11 = x8 DDR2 (one strobe pair per nibble like x4 devices)
5:4	00b R/W	**DRAM Technology for Odd-numbered Row**. BIOS sets this bit according to the density of the DDR-SDRAM devices populated in this row. This is used along with the Device Width to determine the page size. 00 = 128-Mb DDR-SDRAM 01 = 256-Mb DDR-SDRAM 10 = 512-Mb DDR-SDRAM 11 = 1-Gb DDR-SDRAM
3:2	00b R/W	**Device Width for Even-Numbered Row.** BIOS sets this bit according to the width of the DDR-SDRAM devices populated in this row. This is used to determine the page size, as well as the DQS to DQ signal mapping. 00 = Reserved 01 = x8 DDR-SDRAM 10 = x4 DDR-SDRAM 11 = x8 DDR2 (one strobe pair per nibble like x4 devices)
1:0	00b R/W	**DRAM Technology for Even-numbered Row**. BIOS sets this bit according to the density of the DDR-SDRAM devices populated in this row. This is used along with the Device Width to determine the page size, as well as the DQS to DQ signal mapping. 00 = 128-Mb DDR-SDRAM 01 = 256-Mb DDR-SDRAM 10 = 512-Mb DDR-SDRAM 11 = 1-Gb DDR-SDRAM

Table 3-4. DIMM to DRA Register Mapping

DIMM #	Even Row (or Single Rank)		Odd Row (present if double rank)	
	Row Number	Address of DRA	Row Number	Address of DRA
DIMM 1	Row 0	70h bits 3:0	Row 1	70h bits 7:4
DIMM 2	Row 2	71h bits 3:0	Row 3	71h bits 7:4
DIMM 3	Row 4	72h bits 3:0	Row 5	72h bits 7:4
DIMM 4	Row 6	73h bits 3:0	Row 7	73h bits 7:4

3.5.19 DRT – DRAM Timing Register (D0:F0)

Address Offset: 78-7Bh
Access: R/W
Size: 32 Bits
Default: 9599_9604h

This register controls the timing of the DRAM Interface.

Bit Field	Default & Access	Description
31:30	10b R/W	**Programmable Read Pointer Delay**. This bit determines the read pointer delay, which is based on both DIMM topology and technology. The encodings in this table refer to additional delays beyond one command clock. Encoding Number of CMDCLK at Specified Frequency 167 MHz200 MHz 00 0 (0 ns)0 (0 ns) 01 1 (6 ns)1 (5 ns) 10 2 (12 ns)2 (10 ns) 11 ReservedReserved
29:28	01b R/W	**Back-To-Back Write-Read Turn Around.** This field determines the data bubble duration in CMDCLKs between Write-Read commands. It applies to WR-RD pairs to any destinations (in same or different rows). The purpose of this bit is to control the turnaround time on the DQ bus. Encoding Number of CMDCLK at Specified Frequency 167 MHz200 MHz 00 ReservedReserved 01 1 (6 ns)1 (5 ns) 10 2 (12 ns)2 (10 ns) 11 Reserved3 (15 ns)
27:26	01b R/W	**Back-To-Back Read-Write Turn Around.** This field determines the minimum number of CMDCLK between Read-Write commands. It applies to RD-WR pairs to any destinations (in same or different rows). The purpose of this bit is to control the turnaround time on the DQ bus. Encoding Number of CMDCLK at Specified Frequency 167 MHz200 MHz 00 1 (6 ns)1 (5 ns) 01 2 (12 ns)2 (10 ns) 10 3 (18 ns)3 (15 ns) 11 4 (24 ns)4 (20 ns)

Bit Field	Default & Access	Description
25:24	01b R/W	**Back To Back Read Turn Around.** This field determines the minimum number of CMDCLK between two Reads destined to different rows. The purpose of these bits is to control the turnaround time on the DQ bus. Encoding — Number of CMDCLK at Specified Frequency 167 MHz / 200 MHz 00 — ReservedReserved 01 — 1 (6 ns) / 1 (5 ns) 10 — 2 (12 ns) / 2 (10 ns) 11 — Reserved / 3 (15 ns)
23:22	10b R/W	**Auto refresh Cycle Time (Trfc)**. The required time between/after auto refresh cycles to any particular DIMM. Encoding — Number of CMDCLK at Specified Frequency 167 MHz / 200 MHz 00 — 12 (72 ns) / 15 (75 ns) 01 — Reserved / 21 (105 ns) 10 — 20 (120 ns) / 26 (130 ns) 11 — Reserved / Reserved
21:20	01b R/W	**Row Delay (Trrd)**. The required row delay period between two activate commands accessing the same CS# of a DIMM. Encoding — Number of CMDCLK at Specified Frequency 167 MHz / 200 MHz 00 — 1 (6 ns) / 1 (5 ns) 01 — 2 (12 ns) / 2 (10 ns) 10 — 3 (18 ns) / 3 (15 ns) 11 — Reserved / 4 (20 ns)
19:18	10b R/W	**Trasmax**. Indicates allowable number of clocks to allow sequential page hits prior to forcing a precharge to close the page. Encoding — Number of CMDCLK at Specified Frequency 167 MHz / 200 MHz 00 — 32 / 32 01 — 64 / 64 10 — 128 / 128 11 — 512 / 512
17:16	01b R/W	**Write Recovery Delay (Twr)**. The required write recovery delay before being able to issue a precharge command to the same page accessing the same CS/bank of a DIMM. Encoding — Number of CMDCLK at Specified Frequency 167 MHz / 200 MHz 00 — Reserved / Reserved 01 — 2 (12 ns) / Reserved 10 — 3 (18 ns) / 3 (15 ns) 11 — Reserved / 4 (20 ns)
15:14	10b R/W	**RAS Cycle Time (Trc)**. These bits control the required delay from an activate command before issuing another activate or refresh command to the same CS/bank of a DIMM. Encoding — Number of CMDCLK at Specified Frequency 167MHz / 200 MHz 00 — Reserved / 11 (55 ns) 01 — Reserved / 12 (60 ns) 10 — 10 (60 ns) / 13 (65 ns) 11 — Reserved / Reserved

Bit Field	Default & Access	Description
13:12	01b R/W	**Write with Auto precharge Recovery Delay (Tdal)**. The time required before being able to issue an activate command to the same page accessing the same CS / Bank of a DIMM. The value must be the sum of Trp and Twr (Tdal = Trp + Twr). Encoding Number of CMDCLK at Specified Frequency 167 MHz200 MHz 00 Reserved6 (30 ns) 0 5 (30 ns)7 (35 ns) 10 6 (36 ns)Reserved 11 ReservedReserved
11:10	01b R/W	**Write RAS# to CAS# Delay (Trcd)**. Controls the number of clocks inserted between a row activate command and a read or write command to that row Encoding Number of CMDCLK at Specified Frequency 167 MHz200 MHz 00 2 (12 ns)3 (15 ns) 01 3 (18 ns)4 (20 ns) 10 ReservedReserved 11 ReservedReserved
9:8	10b R/W	**DRAM RAS# Precharge (Trp)**. This bit controls the number of clocks that are inserted between a row precharge command and an activate command to the same row. Encoding Number of CMDCLK at Specified Frequency 167 MHz200 MHz 00 ReservedReserved 01 2 (12 ns)3 (15 ns) 10 3 (18 ns)4 (20 ns) 11 ReservedReserved
7:6	00b R/W	**Back-To-Back Write Turn Around**. This field determines the data bubble duration between Write data bursts. It applies to WR-WR pairs to different ranks, and is only expected to be used in DDR2 mode with ODT enabled in the event that ODT selections must change between ranks. The purpose of this field is to control the data burst spacing on the DQ bus. Encoding Number of CMDCLK at Specified Frequency 167 MHz200 MHz 00 0 (0 ns)0 (0 ns) 01 1 (6 ns)1 (5 ns) 10 2 (12 ns)2 (10 ns) 11 Reserved3 (15 ns)
5	0b R/W	**Turn Around Cycle Add**. Setting this bit to a '1' adds an extra turn around cycle between a read to DIMM4 (furthest DIMM from MCH) and a read to DIMM1 (nearest DIMM to MCH). It is only intended for use in a 4 DIMM configuration where the farthest logical DIMM is also the farthest physical DIMM (DRM setting of 1248h).
4	0b R/W	**CKE Guard Band**. 0 = CKE is driven high for one CMDCLK prior to a new command 1 = CKE is driven high for two CMDCLKs prior to a new command

Bit Field	Default & Access	Description
3:2	01 R/W	**CAS# Latency (Tcl).** The number of clocks between the rising edge used by DRAM to sample the Read Command and the rising edge that is used by the DRAM to drive read data. Encoding Number of CMDCLK at Specified Frequency 167 MHz200 MHz 00 2 or 2.5Reserved 01 Reserved3 10 3 4 11 ReservedReserved
1:0	00b R/W	**CKE Idle Selection.** Specifies the number of CMDCLKs with no command activity to a row before deasserting CKE, putting the row into low power mode. Encoding Number of CMDCLK at Specified Frequency 167 MHz200 MHz 00 32 32 01 128128 10 512512 11 2K 2K

3.5.20 DRC – DRAM Controller Mode Register (D0:F0)

Address Offset: 7C – 7Fh
Access: RO, R/W, R/WO
Size: 32 Bits
Default: 0000_0008h

This register controls the mode of the DRAM controller.

Bit Field	Default & Access	Description
31:30	00b RO	**Revision Number (REV)**. The DDR register definition format revision number. This field is reserved and will always return zeroes when read.
29	0b R/W	**Initialization Complete (IC).** This bit is used for communication of software state between the memory controller and the BIOS. 0 = The DRAM interface has NOT been initialized. 1 = The DRAM interface has been initialized.
28	0b R/W	**Dynamic Power-Down Mode Enable.** When set, the DRAM controller will put pair of rows into power down mode when all banks are precharged (closed). Once a bank is accessed, the relevant pair of rows is taken out of Power Down mode. DIMM power down mode is controlled by the CKE signal. 0 = DRAM Power-down disabled (DDR, DDR2 with ODT) 1 = DRAM Power-down enabled (DDR)
27	0b RW	**DED Retry Enable**: 0 = Disabled. No retries occur on double-bit errors. 1 = Enable a single retry of read accesses on detection of a double-bit error.
26	0b R/W	**Overlap Enable**. 0 = Inhibits overlap scheduling of row/col tenures. 1 = Allows overlapped scheduling of activates prior to completing the outstanding column command.

Bit Field	Default & Access	Description
25:24	00b R/W	**Auto-Precharge Mode for Writes**. 00 = Intelligent 01 = Always Auto-precharge 10 = Never Auto-precharge 11 = Reserved
23:22	00b R/W	**Auto-Precharge Mode for Reads**. 00 = Intelligent 01 = Always Auto-precharge 10 = Never Auto-precharge 11 = Reserved
21:20	00b R/W	**DRAM Data Integrity Mode (DDIM):** These bits select one of four DRAM data integrity modes. When in non-ECC mode, no ECC correction is done and no ECC errors are logged in the FERR/NERR registers. **NOTE 1:** Non-ECC mode should be used for debug purposes only, and non-ECC DIMMs are not supported. **NOTE 2:** If mode 10b (x4 SDDC) is selected in single channel mode, the MCH will effectively change this to 01 and enforce 72-bit ECC. 00 = Non-ECC mode. 01 = 72-bit ECC 10 = x4 Chip-Fail ECC 11 = Reserved
19:11	000h	Reserved
10:8	000b R/W	**Refresh Mode Select (RMS)**: This field determines whether refresh is enabled and, if so, at what rate refreshes will be executed. 000 = Refresh disabled 001 = Refresh enabled. Refresh interval 15.6 µs 010 = Refresh enabled. Refresh interval 7.8 µs 011 = Refresh enabled. Refresh interval 64 µs 100 = Refresh enabled. Refresh interval 3.9 µs 101 = Reserved 110 = Reserved 111 = Refresh enabled. Refresh interval 64 clocks (fast refresh mode)
7	0b R/W	**CMD Disable Channel B.** Used during fail-down to disable the DDR/DDR2 command bus. This bit is sticky through reset. 0 = DDR channel B command bus is enabled 1 = DDR channel B command bus is disabled
6	0b R/W	**CMD Disable Channel A.** Used during fail-down to disable the DDR/DDR2 command bus. This bit is sticky through reset. 0 = DDR channel A command bus is enabled 1 = DDR channel A command bus is disabled
5	0b R/W	**DRAM ODT Disable**. Used to disable ODT when running in DDR2-400 mode. If running in DDR this bit has no effect. This bit is sticky through reset. 0 = DRAM ODT enabled in DDR2-400 mode 1 = DRAM ODT disabled in DDR2-400 mode
4	0b R/W	**CKE pin mode.** DDR Clock Enables have two operating modes. This bit is sticky through reset. 0 = All eight CKEs are shared across both lock-stepped channels, where 0..7 correspond to the rows indicated by DRB0..7. 1 = Independent CKEs per channel with one CKE per DIMM slot 0, 2, 4, 6 for channel A, and 1, 3, 5, 7 for channel B (This mode needed to support single channel faildown state).

Bit Field	Default & Access	Description
3:2	10b R/WO	**FSBFREQSEL.** Front Side Bus Frequency Select in MHz. **The PLL only supports one update of ratio (the lower nibble of this register).** The lower nibble of this register becomes locked down only when the lower byte of the register is written for the first time, and then cannot be further updated. This bit is sticky through reset. 00 = Reserved 01 = Reserved 10 = 166 MHz 11 = 200 MHz **Note:** Failure to properly program these bits can result in the inability to access configuration space.
1:0	00b R/W	**DRAM Type (DT).** Used to select between supported SDRAM types. The bit setting reflects whether or not it is DDR2, and hence differential strobes. The PLL only supports one update of ratio (the lower nibble of this register). This bit is sticky through reset. To meet the JEDEC spec, the hardware assumes that these ODT/CS# pins are ODT pins and drives them low until the DRAM type is established. This bit field must be written to establish whether or not the DRAMs are in fact DDR or DDR2 before accessing the DIMMs. 00 = Reserved 01 = DDR333 10 = DDR2-400 11 = Reserved **Note:** Failure to properly program these bits can result in the inability to access configuration space.

3.5.21 DRM – DRAM Mapping Register (D0:F0)

Address Offset: 80 – 81h
Access: R/W
Size: 16 Bits
Default: 8421h

This register is used for mapping CS logical to CS physical. Encodings not listed are not supported. This register is sticky through reset.

Bit Field	Default & Access	Description
15:12	8h R/W	**Logical CS to physical CS[7:6] mapping**. 0001 Logical CS(1:0) maps to physical CS(7:6) 0010 Logical CS(3:2) maps to physical CS(7:6) 0100 Logical CS(5:4) maps to physical CS(7:6) 1000 Logical CS(7:6) maps to physical CS(7:6)
11:8	4h R/W	**Logical CS to physical CS[5:4] mapping**. 0001 Logical CS(1:0) maps to physical CS(5:4) 0010 Logical CS(3:2) maps to physical CS(5:4) 0100 Logical CS(5:4) maps to physical CS(5:4) 1000 Logical CS(7:6) maps to physical CS(5:4)

Bit Field	Default & Access	Description
7:4	2h R/W	**Logical CS to physical CS[3:2] mapping**. 0001 Logical CS(1:0) maps to physical CS(3:2) 0010 Logical CS(3:2) maps to physical CS(3:2) 0100 Logical CS(5:4) maps to physical CS(3:2) 1000 Logical CS(7:6) maps to physical CS(3:2)
3:0	1h R/W	**Logical CS to physical CS[1:0] mapping**. 0001 Logical CS(1:0) maps to physical CS(1:0) 0010 Logical CS(3:2) maps to physical CS(1:0) 0100 Logical CS(5:4) maps to physical CS(1:0) 1000 Logical CS(7:6) maps to physical CS(1:0)

3.5.22 DRORC – Opportunistic Refresh Control Register (D0:F0)

Address Offset: 82h
Access: R/W
Size: 8 Bits
Default: 71h

The MCH contains a 4b refresh counter that allows the counting of up to 16 refreshes. Using the counter, refresh requests can be queued when the DRAM interface is busy performing cycles. Ideally, refreshes are performed when the DRAM interface is idle. This opportunistic refresh scheme utilizes two watermarks, which the following register is used to control.

Bit Field	Default & Access	Description
7:4	0111b R/W	**High Watermark**. When the refresh-counter reaches or exceeds the value in the high watermark field, the DRAM controller performs a refresh in the highest priority mode. In such a case, refresh will be processed as soon as the currently pending DRAM cycle is completed. Once a high priority refresh is internally launched (through the command queue), the DRAM controller may schedule an additional refresh immediately if the refresh counter high watermark condition remains “true”. Note that current DDR components require DLL refresh every 9 refresh periods. As a result, this register must be set at 7 or lower. Bit field encoding: 0000 Illegal value 0001 One Refresh is the watermark …… 1111 15 Refreshes is the watermark
3:0	0001b R/W	**Low Watermark**. When the refresh-counter reaches or exceeds the value in the low watermark field, the DRAM controller performs a refresh if there is no other request pending to DRAM. It means that low watermark refresh is performed as the lowest priority request, opportunistically. Once a low priority refresh is internally launched (through the command queue), the DRAM controller will not schedule an additional low priority refresh until the already launched refresh operation is completed (low watermark refresh is blocked when the command queue contains a low watermark refresh request). Once low watermark refresh counter is reached or exceeded, the DRAM controller will opportunistically perform low priority refreshes until the refresh counter is down to 0. Bit field encoding: 0000 Illegal value 0001 One Refresh is the watermark …… 1111 15 Refreshes is the watermark

3.5.23 ECCDIAG – ECC Detection /Correction Diagnostic Register (D0:F0)

Address Offset: 84 – 87h
Access: R/W, RO, R/WS
Size: 32 Bits
Default: 0000_0000h

This register is used for diagnostic testing of ECC from the DRAM. This feature is presented for validation purposes only. Functionality is not guaranteed and may not be supported.

Bit Field	Default & Access	Description
31:18	000h RO	Reserved
18	0b R/W	**Memory Poison Enable.** Allows for propagation of data errors not initiated by this register to DRAM. Error Injection is possible regardless of this bit setting. The setting of this bit has no effect on the reporting or logging of data errors. 0 = Error poisoning is disabled. Bad ECC arriving at the Memory Interface (Inbound, during memory writes), will be recalculated based on the data, and then, forwarded to the memory. 1 = Error poisoning enabled. Bad ECC arriving at the Memory Interface (Inbound, during memory writes), will be forwarded as bad ECC to the Memory.
17:0	00000h	Reserved

3.5.24 SDRC – DDR SDRAM Secondary Control Register (D0:F0)

Address Offset: 88 - 8Bh
Access: RO, R/W
Size: 32 bits
Default Value: 0000_0000h

Bit Field	Default & Access	Description
31:30	00b RW	**Channel B On Die Termination Enable.** Enables the MCH on-die termination for the DDR2 channel B data signals. The MCH terminates the DQ/DQS signals for a read. The MCH has ODT regardless of its operating in either DDR or DDR2 mode even though of the two types of memory devices, only DDR2 has ODT. The termination values are dependent on the DDRIMPCRES external resistor (Rodt). **Encoding** / **DDR MCH ODT (Rodt=383)** / **DDR2 MCH ODT (Rodt = 287)** 00 / Off / Off 01 (Rodt/2) / ~200 / 150 10 (Rodt/4) / ~200 / 150 11 / ~100 / ~75
29:28	00b RW	**Channel A On Die Termination Enable.** These bits enable the MCH on-die termination for the DDR2 channel A data signals. Encodings match those provided above for channel B.
27:9	0_0000h RO	Reserved

Bit Field	Default & Access	Description
8	0 RW	**DQS Half Gain (DQSHALFGAIN).** Select for the DQS differential amplifier gain. This bit is '0' for DDR, but is set to '1'. 0 = Gain of 1 (DDR setting) 1 = Gain of 1/2 (cuts the gain roughly in half for differential strobe mode associated with DDR2)
7	0b RW	**Differential DQS Enable (DIFFDQSEN**). 0 = Disable (DDR) 1 = Enables differential data strobes for DDR2
6:0	0000h RW	Reserved

3.5.25 CKDIS – CK/CK# Disable Register (D0:F0)

Address Offset: 8Ch
Access: R/W
Size: 8 Bits
Default: FFh

The MCH uses one of the 4 clock pairs on each DDR channel to properly calibrate the voltage crossing. The clocks used to determine VOX crossing are as follows: DDRA_CMDCLK1/1# for channel A and DDRB_CMDCLK0/0# for channel B. The MCH requires that these clocks be enabled to perform VOX calibration.

Bit Field	Default & Access	Description
7:0	FFh R/W	**CK/CK# Disable.** Each bit corresponds to a CMDCLK/CMDCLK# pair of pins on one of the DIMMs, as defined in the table below. 0 = Enable CMDCLK signals to the corresponding DIMM (or DIMM pair) 1 = Disable CMDCLK signals to the corresponding DIMM (or DIMM pair). When disabled, the corresponding CMDCLK/CMDCLK# pair are tri-stated. Bit DIMM 7 DDRB_CMDCLK3 / DDRB_CMDCLK3# 6 DDRB_CMDCLK2 / DDRB_CMDCLK2# 5 DDRB_CMDCLK1 / DDRB_CMDCLK1# 4 DDRB_CMDCLK0 / DDRB_CMDCLK0# (Calibration Clock) 3 DDRA_CMDCLK3 / DDRA_CMDCLK3# 2 DDRA_CMDCLK2 / DDRA_CMDCLK2# 1 DDRA_CMDCLK1 / DDRA_CMDCLK1# (Calibration Clock) 0 DDRA_CMDCLK0 / DDRA_CMDCLK0#

3.5.26 CKEDIS – CKE/CKE# Disable Register (D0:F0)

Address Offset: 8Dh
Access: R/W
Size: 8 Bits
Default: 00h

Bit Field	Default & Access	Description
7:0	00h R/W	**CKE disable. Each bit corresponds to a CKE pin as defined in the table below.** 0 = Enable CKE signal (Default) 1 = Disable CKE signal Bit — Device Pin 7 — DDRCKE7 6 — DDRCKE6 5 — DDRCKE5 4 — DDRCKE4 3 — DDRCKE3 2 — DDRCKE2 1 — DDRCKE1 0 — DDRCKE0

3.5.27 DDRCSR – DDR Channel Configuration Control/Status Register (D0:F0)

Address Offset: 9A – 9Bh
Access: RO, R/W, R/WS
Size: 16 Bits
Default: 0000h

Bit Field	Default & Access	Description
15	0b R/WS	**Transition Enable**. Written with a 1 to set and enable DDR/DDR2 Channel Configuration FSM transitions, a write of 0 does nothing. This bit is cleared by the MCH when the Channel Configuration FSM state change has occurred. This bit is used to take the DDRCSR FSM out of idle and also to initiate data migration for sparing. 0 = No change 1 = Set and enable DDR Channel Configuration FSM transitions.
14:10	00h	Reserved
9	0b R/W	**DIMM Sparing Enable**. To enter any of the sparing states we must not be in symmetric decoding. This bit written by BIOS informs the memory subsystem whether or not on-line DIMM sparing is enabled, and thereby also disables symmetric decode when set. (DEFAULT is disabled.)
8:7	00b R/W	**Failing DIMM**. Used to specify the copy source for DIMM sparing. The encoding for this two-bit field shown is for the logical CS pair. Note that setting these bits by themselves does not initiate data migration. For the copy to commence, the Transition Enable bit must be set. 00 = CS pair 0 and 1 01 = CS pair 2 and 3 10 = CS pair 4 and 5 11 = CS pair 6 and 7
6:5	00b	Reserved

Bit Field	Default & Access	Description
4	0b R/O	**Symmetric Mode**. The DRAM logic when it detects that it is not in sparing mode, and it has four equal ranks, puts the DIMMs into interleaved mode for improved performance. This state is reflected for use by BIOS.
3:0	0h RO	**Channel Configuration FSM Current State.** This field details the current state of the DDRCSR_FSM. 0000 – Idle 0100 – Single channel A normal 0101 – Single channel A sparing copy in progress 0111 – Single channel A sparing complete 1000 – Single channel B normal 1001 – Single channel B sparing copy in progress 1011 – Single channel B sparing complete 1100 – Dual channel normal 1101 – Dual channel sparing copy in progress 1111 – Dual channel sparing complete 1010 – Reserved 1110 – Reserved

3.5.28 DEVPRES – Device Present (D0:F0)

Address Offset: 9C
Access: R/WO, RO
Size: 8 bits
Default: 01h

The Device Present bits can be used to disable the various devices within the MCH and make their PCI configuration space invisible to software. Once software has disabled devices, they can only be re-enabled via reset. The MCH does not support turning off a device, and then turning it back on. This register should only be written to at boot time when there is no traffic to or from the PCI Express* ports.

Bit Field	Default & Access	Description
7:5	000b	Reserved
4	0b R/WO	**Device 4 Present**. 0 = PCI Express* port B is disabled. 1 = PCI Express* port B is enabled.
3	0b R/WO	**Device 3 Present**. 0 = PCI Express* port A1 (x4) is disabled. In this state, port A (Device 2) can operate with a maximum x8 link width. 1 = PCI Express* port A1 (x4) is enabled. In this state, port A (Device 2) can operate with a maximum x4 link width.
2	0b R/WO	**Device 2 Present**. 0 = PCI Express* port A is disabled. 1 = PCI Express* port A is enabled.
1	0b	Reserved
0	1b RO	**Device 0 Present.** The MCH Device 0 cannot be disabled.

3.5.29 ESMRC – Extended System Management RAM Control (D0:F0)

Address Offset: 9Dh
Access: R/W/L, R/W
Size: 8 Bits
Default: 00h

The Extended SMRAM register controls the configuration of Extended SMRAM space. The Extended SMRAM (E_SMRAM) memory provides a write-back cacheable SMRAM memory space that is above 1 Mbyte.

Bit Field	Default & Access	Description
7	0b R/W/L	**Enable High SMRAM (H_SMRAME)**. Controls the SMM memory space location (i.e. above 1 Mbyte or below 1 Mbyte). When G_SMRAME is '1' and H_SMRAME is set to '1', the high SMRAM memory space is enabled. SMRAM accesses within the range 0FEDA_0000h to 0FEDA_FFFFh are remapped to DRAM addresses within the range 000A0000h to 000BFFFFh. Once D_LCK has been set, this bit becomes read only.
6	0b R/W	**MDA Present (MDAP)**. This bit works with the VGA Enable bits in the BCTRL registers of Devices 2–57 to control the routing of processor-initiated transactions targeting MDA compatible I/O and memory address ranges. This bit should not be set if none of the VGA Enable bits are set. If none of the VGA enable bits are set, then accesses to I/O address range x3BCh – x3BFh are forwarded to the HI. If the VGA enable bit is not set then accesses to I/O address range x3BCh – x3BFh are treated just like any other I/O accesses, i.e., the cycles are forwarded to PCI Express* if the address is within the corresponding IOBASE and IOLIMIT and ISA enable bit is not set, otherwise they are forwarded to HI. **NOTE:** Since the logic performs address decoding on a DW boundary, the DW that includes address 3BFh also includes addresses 3BCh, 3BDh and 3BEh, and accesses to any of these byte addresses are handled as MDA references. MDA resources are defined as the following: Memory: 0B0000h – 0B7FFFh I/O: 3B4h, 3B5h, 3B8h, 3B9h, 3BAh, 3BFh, (including ISA address aliases, A[15:10] are not used in decode) **NOTE:** The VGA region includes I/O space ranges 3B0 – 3BBh and 3C0 – 3DF, so there is an overlap between these two I/O regions. Any I/O reference that includes the I/O locations listed above, or their aliases, will be forwarded to Hub Interface even if the reference includes I/O locations not listed above. The following table shows the behavior for all combinations of MDA and VGA: **VGA / MDA / Behavior** 0 / 0 / All References to MDA and VGA go to HI 0 / 1 / Illegal Combination (DO NOT USE) 1 / 0 / All References to VGA go to device with VGA enable set. MDA-only references (I/O address 3BF and aliases) will go to HI. 1 / 1 / VGA-only references go to the PCI Express* port which has its VGA Enable bit set. MDA References go to the HI
5	0b R/W	**APIC Memory Range Disable (APICDIS)**. When set to '1', the MCH forwards accesses to the IOAPIC regions to the appropriate interface, as specified by the memory and PCI configuration registers. When this bit is clear, the MCH will send cycles between 0_FEC0_0000 and 0_FEC7_FFFF to the HI. Accesses between 0_FEC8_0000 and 0_FEC8_0FFF will be sent to PCI Express A, between 0_FEC8_1000 and 0_FEC8_1FFF will be sent to PCI Express A1, between 0_FEC8_2000 and 0_FEC8_2FFF will be sent to PCI Express B.

Bit Field	Default & Access	Description
4	0b R/W	**Hub Interface RCOMP Disable**. 0 = Enable the RCOMP operation for the Hub Interface. 1 = Disable the RCOMP operation for the Hub Interface.
3	0b R/W/L	**Global SMRAM Enable (G_SMRAME)**. 0 = Disable 1 = Enable compatible and extended SMRAM functions. Provides 128 KB of DRAM accessible at the A0000h address while in SMM (ADS# with SMM decode). Refer to the section on SMM for more details. Once D_LCK is set, this bit becomes read only.
2:1	00b R/W/L	**TSEG Size (TSEG_SZ)**. Selects the size of the TSEG memory block if enabled. Memory from the top of DRAM space (TOLM – TSEG_SZ) to TOLM is partitioned away so that it may only be accessed by the processor interface and only then when the SMM bit is set in the request packet. Non-SMM accesses to this memory region are sent to HI when the TSEG memory block is enabled. Note that once D_LCK is set, these bits become read only. 0 0= (TOLM – 128 k) to TOLM 0 1= (TOLM – 256 k) to TOLM 1 0= (TOLM – 512 k) to TOLM 1 1= (TOLM – 1 M) to TOLM
0	0b R/W/L	**TSEG Enable (T_EN)**. Enabling of SMRAM memory for Extended SMRAM space only. When G_SMRAME = 1 and TSEG_EN = 1, the TSEG is enabled to appear in the appropriate physical address space. Note that once D_LCK is set, this bit becomes read only.

3.5.30 SMRC – System Management RAM Control Register (D0:F0)

Address Offset: 9Eh
Access: RO, R/W, R/WL, R/WS
Size: 8 Bits
Default: 02h

This register controls how accesses to Compatible and Extended SMRAM spaces are treated. The open, close, and lock bits function only when G_SMRAME bit is set to ‘1’. The open bit must be reset before the lock bit is set.

Bits	Default, Access	Description
7	0b	Reserved
6	0b R/WL	**SMM Space Open (D_OPEN).** When D_OPEN=1 and D_LCK=0, the SMM space DRAM is made visible even when SMM decode is not active. This is intended to help BIOS initialize SMM space. Software should ensure that D_OPEN=1 and D_CLS=1 are not set at the same time. This bit becomes RO when D_LCK is set to ‘1’.
5	0b R/W	**SMM Space Closed (D_CLS).** When D_CLS = 1 SMM space DRAM is not accessible to data references, even if SMM decode is active. Code references may still access SMM space DRAM. This will allow SMM software to reference through SMM space to update the display even when SMM is mapped over the VGA range. Software should ensure that D_OPEN=1 and D_CLS=1 are not set at the same time. Note that the D_CLS bit only applies to Compatible SMM space.

Bits	Default, Access	Description
4	0b R/WS	**SMM Space Locked (D_LCK).** When D_LCK is set to '1', D_OPEN is reset to '0' and D_LCK, D_OPEN, H_SMRAM_EN, TSEG_SZ and TSEG_EN become read only. D_LCK can be set to '1' via a normal configuration space write but can only be cleared by a Full Reset. The combination of D_LCK and D_OPEN provide convenience and security. The BIOS can use the D_OPEN function to initialize SMM space and then use D_LCK to "lock down" SMM space in the future so that no application software (or the BIOS itself) can violate the integrity of SMM space, even if the program has knowledge of the D_OPEN function.
3	0b	Reserved
2:0	010b RO	**Compatible SMM Space Base Segment (C_BASE_SEG).** This field indicates the location of SMM space. SMM DRAM is not remapped. It is simply made visible if the conditions are right to access SMM space, otherwise the access is treated as a VGA access. The MCH supports only the SMM space between A0000h and BFFFFh.

3.5.31 EXSMRC – Expansion System Management RAM Control (D0:F0)

Address Offset: 9Fh
Access: R/WC, RO
Size: 8 Bits
Default: 07h

Controls the configuration of Extended SMRAM space. The Extended SMRAM (E_SMRAM) memory provides a write-back cacheable SMRAM memory space that is above 1 Mbyte.

Bit Field	Default & Access	Description
7	0b R/WC	**Invalid SMRAM Access (E_SMERR).** It is software's responsibility to clear this bit. This bit is cleared by writing a '1' to the bit location. 0 = No Invalid SMRAM accesses detected. 1 = Processor has accessed the defined memory ranges in the Extended SMRAM (High Memory and T-segment) while not in SMM space, and with the D_OPEN bit = 0.
6:3	0h	Reserved
2	1b RO	**SMRAM Cacheable (SM_CACHE)**.
1	1b RO	**L1 Cache Enable for SMRAM (SM_L1)**.
0	1b RO	**L2 Cache Enable for SMRAM (SM_L2)**.

3.5.32 DDR2ODTC – DDR2 ODT Control Register (D0:F0)

Address Offset: B0 – B3h
Access: R/W
Size: 32 Bits
Default: 0000_0000h

This register consists of 4 bytes of ODT enable vectors, one per DDR2 rank, each of which contains a nibble for read accesses and a nibble for write accesses. The selected ODT vector for a given access type is sent to both channels when in dual channel mode. The encodings for each nibble are a direct reflection of the four ODT outputs driven on the odd CS lines of each channel during an access to the corresponding rank. Refer to the lower nibble bit field description for clarification.

Bit Field	Default & Access	Description
31:28	0h R/W	**CS6ODTWR**. Logical ODT vector for write access to logical CS 6 routed rank.
27:24	0h R/W	**CS6ODTRD**. Logical ODT vector for write access to logical CS 6 routed rank.
23:20	0h R/W	**CS4ODTWR**. Logical ODT vector for write access to logical CS 4 routed rank.
19:16	0h R/W	**CS4ODTRD**. Logical ODT vector for read access to logical CS 4 routed rank.

Bit Field	Default & Access	Description
15:12	0h R/W	**CS2ODTWR**. Logical ODT vector for write access to logical CS 2 routed rank.
11:8	0h R/W	**CS2ODTRD**. Logical ODT vector for read access to logical CS 2 routed rank.
7:4	0h R/W	**CS0ODTWR**. Logical ODT vector for write access to logical CS 0 routed rank. The four bits correspond directly to CS[7,5,3,1] values driven on the multi-function CS lines used for ODT in DDR2 mode. For each asserted bit in the nibble, the corresponding ODT output will be driven during read accesses. So a value of 5h would cause CS5 and CS1 to be driven asserted, and CS7 and CS3 to remain deasserted during writes.
3:0	0h R/W	**CS0ODTRD**. Logical ODT vector for read access to logical CS 0 routed rank. The four bits correspond directly to CS[7,5,3,1] values driven on the multi-function CS lines used for ODT in DDR2 mode when the DRM register is set to 0x8421 which maps logical CS directly to physical CS. For each asserted bit in the nibble, the corresponding ODT output will be driven during read accesses to rank 1. So a value of 5h would cause logical CS5 and logical CS1 to be driven asserted, and logical CS7 and logical CS3 to remain deasserted during reads.

3.5.33 TOLM – Top of Low Memory Register (D0:F0)

Address Offset: C4 – C5h
Access: R/W
Size: 16 Bits
Default: 0800h

This register contains the maximum address below 4 GB that should be treated as a memory access, and is defined on a 128-MB boundary. Usually it will sit below the areas configured for PCI Express*, HI and PCI memory. Note that the memory address found in DRB7 reflects the total amount of memory populated. In the event that the combined DRAM and PCI memory space in the system is less than 4 GB, these two registers will be identical.

Bit Field	Default & Access	Description

15:11	00001b R/W	**Top of Low Memory (TOLM)**. This register corresponds to bits 31 to 27 of the system address which is 1 greater than the maximum DRAM location below 4GB. Configuration software should set this value to either the maximum amount of memory in the system or to the minimum address allocated for PCI memory or the graphics aperture, whichever is smaller. Address bits 26:0 are assumed to be 0 for the purposes of address comparison. Addresses equal to or greater than the TOLM, and less than 4G, are treated as non-memory accesses. All accesses less than the TOLM are treated as DRAM accesses (except for the 15-16MB or PAM gaps). This register must be set to at least 0800h, for a minimum of 128MB of DRAM. There is also a minimum of 128MB of PCI space, since this register is on a 128MB boundary. Configuration software should set this value to either the maximum amount of memory in the system (same as DRB7), or to the lower 128MB boundary of the Memory Mapped IO range, whichever is smaller. Programming example: 1100_0b = 3GB (assuming that DBR7 is set > 4GB): An access to 0_C000_0000h or above (but <4GB) will be considered above the TOLM and therefore not to DRAM. It may go to PCI Express* or the HI or be subtractively decoded to HI. An access to 0_BFFF_FFFFh and below will be considered below the TOLM and go to DRAM.
10:0	000h	Reserved

3.5.34 REMAPBASE – Remap Base Address Register (D0:F0)

Address Offset: C6 – C7h
Access: R/W
Size: 16 Bits
Default: 03FFh

Bit Field	Default & Access	Description
15:10	0	Reserved
9:0	3FFh R/W	**Remap Base Address [35:26]**. The value in this register defines the lower boundary of the remap window. The remap window is inclusive of this address. In the decoder A[25:0] of the remap Base Address are assumed to be zeros. Thus the bottom of the defined memory range will be aligned to a 64-MB boundary. When the value in this register is greater than the value programmed into the Remap Limit register, the Remap window is disabled.

3.5.35 REMAPLIMIT – Remap Limit Address Register (D0:F0)

Address Offset: C8 – C9h
Access: R/W
Size: 16 Bits
Default: 0000h

Bit Field	Default & Access	Description
15:10	0	Reserved
9:0	00h R/W	**Remap Limit Address [35:26].** The value in this register defines the upper boundary of the remap window. The remap window is inclusive of this address. In the decoder A[25:0] of the Remap Limit Address are assumed to be F's. Thus the top of the defined range will be one less than a 64-MB boundary. When the value in this register is less than the value programmed into the Remap Base register, the remap window is disabled.

3.5.36 REMAPOFFSET – Remap Offset (D0:F0)

Address Offset: CA – CBh
Access: R/W
Size: 16 Bits
Default: 0000h

This register contains the difference between the REMAPBASE and TOLM.

Bit Field	Default & Access	Description
15:10	0	Reserved
9:0	000h R/W	**Remap Offset**. This register contains the difference between the REMAPBASE and TOLM. This register value corresponds to address bits 35:26. It is used to translate the physical host-bus address to the system memory address for accesses to the remap region.

3.5.37 TOM – Top of Memory Register (D0:F0)

Address Offset: CC – CDh
Access: R/W
Size: 16 Bits
Default: 0000h

BIOS uses this register to determine the memory size reported to the OS. The value in this register hides any DIMMs that can't be directly addressed due to sparing

Bit Field	Default & Access	Description
15:9	0	Reserved
8:0	00h R/W	**Top of Memory (TOM)**. This field reflects the effective size of memory, taking into account sparing. These bits correspond to address bits 35:27 (128-MB granularity). Bits 26:0 are assumed to be 0.

3.5.38 EXPECBASE – PCI Express* Enhanced Configuration Base Address Register (D0:F0)

Address Offset: CE – CFh
Access: R/WO
Size: 16 Bits
Default: E000h

Configuration software will read this register to determine where the 256-MB range of addresses resides for this particular host bridge.

Bit Field	Default & Access	Description
15:12	1110b R/WO	**PCI Express* Enhanced Configuration Base (EXPECBASE)**. This field contains the address that corresponds to bit 31:28 of the base address for the PCI Express* enhanced configuration space below 4 GB. Configuration software will read this register to determine where the 256 MB range of addresses resides for this particular host bridge. BIOS needs to write this register at boot time. Settings 0 and F are not valid. When any byte or combination of bytes of this register is written, the register value locks down and cannot be further updated.
11:0	000h	Reserved

3.5.39 SKPD – Scratchpad Data (D0:F0)

Address Offset: DE – DFh
Access: R/W
Size: 16 Bits
Default: 0000h

Bit Field	Default & Access	Description
15:0	0000h R/W	**Scratchpad (SCRTCH)**. These bits are R/W storage bits that have no effect on the MCH functionality. BIOS typically programs this register to the revision ID of the Memory Reference Code.

3.5.40 DEVPRES1 – Device Present 1 Register (D0:F0)

Address Offset: F4h
Access: RO, R/W
Size: 8 bits
Default Value: 18h

Bit Field	Default & Access	Description
7:6	00b	Reserved
5	0b RW	**Device 0 Function 1 Enable**. 0 = Disable, Device #0 Function #1 is invisible in PCI Configuration space. 1 = Enable, Device #0 Function #1 is visible in PCI Configuration space.

Bit Field	Default & Access	Description
4:2	100b	Reserved
1	0b R/W	**Device #8 Enable.** 0 = Disable, Device #8 is invisible in PCI Configuration space. 1 = Enable, Device #8 is visible in PCI Configuration space.
0	0b RO	Reserved

3.5.41 MCHTST - MCH Test Register (D0:F0)

Address Offset: F5h
Access: RO, R/W
Size: 8 bits
Default Value: 01h

Bit Field	Default & Access	Description
7:1	00h RO	Reserved.
0	1b R/W	**PCI Express* Compliance Mode Disable.** This bit allows or prevents entry by the MCH into Compliance Mode. 0 = Compliance Mode entry is enabled on all PCI Express* ports. 1 = Compliance Mode entry is disabled on all PCI Express* ports.

3.6 MCH Error Reporting Registers (D0:F1)

Table 3-5. Error Reporting PCI Configuration Register Map (D0:F1) (Sheet 1 of 3)

Address Offset	Mnemonic	Register Name	Access	Default
00 – 01h	VID	Vendor Identification	RO	8086h
02 – 03h	DID	Device Identification	RO	3593h
04 – 05h	PCICMD	PCI Command Register	R/W	0000h
06 – 07h	PCISTS	PCI Status Register	R/WC	0000h
08h	RID	Revision Identification	RO	09h
0Ah	SUBC	Sub-Class Code	RO	00h
0Bh	BCC	Base Class Code	RO	FFh
0Dh	MLT	Master Latency Timer	RO	00h
0Eh	HDR	Header Type	RO	00h
2C – 2Dh	SVID	Subsystem Vendor Identification	R/WO	0000h
2E – 2Fh	SID	Subsystem Identification	R/WO	0000h
40 – 43h	FERR_GLOBAL	Global First Error	R/WC	0000_0000h
44 – 47h	NERR_GLOBAL	Global Next Error	R/WC	0000_0000h
50h	HI_FERR	Hub Interface First Error	R/WC	00h

Table 3-5. Error Reporting PCI Configuration Register Map (D0:F1) (Sheet 2 of 3)

Address Offset	Mnemonic	Register Name	Access	Default
52h	HI_NERR	Hub Interface Next Error	R/WC	00h
54h	HI_ERRMASK	Hub Interface Error Mask	R/W	00h
58h	HI_SCICMD	Hub Interface SCI Command	R/W	00h
5Ah	HI_SMICMD	Hub Interface SMI Command	R/W	00h
5Ch	HI_SERRCMD	Hub Interface SERR Command	R/W	00h
5Eh	HI_MCERR	Hub Interface MCERR	R/W	00h
60 – 61h	SYSBUS_FERR	System Bus First Error	R/WC	0000h
62 – 63h	SYSBUS_NERR	System Bus Next Error	R/WC	0000h
64 – 65h	SYSBUS_ERRMASK	System Bus Error Mask	R/W	0009h
68 – 69h	SYSBUS_SCICMD	System Bus SCI Command	R/W	0000h
6A – 6Bh	SYSBUS_SMICMD	System Bus SMI Command	R/W	0000h
6C – 6Dh	SYSBUS_SERRCMD	System Bus SERR Command	R/W	0000h
6E – 6Fh	SYSBUS_MCERR	System Bus MCERR#	R/W	0000h
70h	BUF_FERR	Memory Buffer First Error	R/WC	00h
72h	BUF_NERR	Memory Buffer Second Error	R/WC	00h
74h	BUF_ERRMASK	Memory Buffer Error Mask	R/W	00h
78h	BUF_SCICMD	Memory Buffer SCI Command	R/W	00h
7Ah	BUF_SMICMD	Memory Buffer SMI Command	R/W	00h
7Ch	BUF_SERRCMD	Memory Buffer SERR Command	R/W	00h
7Eh	BUF_MCERRCMD	Memory Buffer MCERR# Command	R/W	00h
80 – 81h	DRAM_FERR	DRAM First Error	R/WC	0000h
82 – 83h	DRAM_NERR	DRAM Next Error	R/WC	0000h
84h	DRAM_ERRMASK	DRAM Error Mask	R/W	00h
88h	DRAM_SCICMD	DRAM SCI Command	R/W	00h
8Ah	DRAM_SMICMD	DRAM SMI Command	R/W	00h
8Ch	DRAM_SERRCMD	DRAM SERR Command	R/W	00h
8Eh	DRAM_MCERR	DRAM MCERR#	R/W	00h
98 – 99h	THRESH_SEC0	DIMM0 SEC Threshold Register	R/W	0000h
9A – 9Bh	THRESH_SEC1	DIMM1 SEC Threshold Register	R/W	0000h
9C – 9Dh	THRESH_SEC2	DIMM2 SEC Threshold Register	R/W	0000h
9E – 9Fh	THRESH_SEC3	DIMM3 SEC Threshold Register	R/W	0000h
A0 – A3h	DRAM_SEC1_ADD	DRAM First SEC Address	RO	0000_0000h
A4 – A7h	DRAM_DED_ADD	DRAM DED Error Address	RO	0000_0000h
A8 – ABh	DRAM_SCRB_ADD	DRAM Scrub Error Address	RO	0000_0000h
AC – AFh	DRAM_RETR_ADD	DRAM DED Retry Address	RO	0000_0000h
B0 – B1h	DRAM_SEC_D0A	DIMM 0 Channel A SEC Counter	R/W	0000h
B2 – B3h	DRAM_DED_D0A	DIMM 0 Channel A DED Counter	R/W	0000h

Table 3-5. Error Reporting PCI Configuration Register Map (D0:F1) (Sheet 3 of 3)

Address Offset	Mnemonic	Register Name	Access	Default
B4 – B5h	DRAM_SEC_D1A	DIMM 1 Channel A SEC Counter	R/W	0000h
B6 – B7h	DRAM_DED_D1A	DIMM 1 Channel A DED Counter	R/W	0000h
B8 – B9h	DRAM_SEC_D2A	DIMM 2 Channel A SEC Counter	R/W	0000h
BA – BBh	DRAM_DED_D2A	DIMM 2 Channel A DED Counter	R/W	0000h
BC – BDh	DRAM_SEC_D3A	DIMM 3 Channel A SEC Counter	R/W	0000h
BE – BFh	DRAM_DED_D3A	DIMM 3 Channel A DED Counter	R/W	0000h
C2 – C3h	THRESH_DED	DED Threshold	R/W	0000h
C4 – C5h	DRAM_SEC1_SYNDROME	First SEC Syndrome	RO	0000h
C6 – C7h	DRAM_SEC2_SYNDROME	Second SEC Syndrome	RO	0000h
C8 – CBh	DRAM_SEC2_ADD	DRAM Next SEC Address	RO	0000_0000h
CC – CDh	DRAM_SEC_D0B	DIMM 0 Channel B SEC Counter	R/W	0000h
CE – CFh	DRAM_DED_D0B	DIMM 0 Channel B DED Counter	R/W	0000h
D0 – D1h	DRAM_SEC_D1B	DIMM 1 Channel B SEC Counter	R/W	0000h
D2 – D3h	DRAM_DED_D1B	DIMM 1 Channel B DED Counter	R/W	0000h
D4 – D5h	DRAM_SEC_D2B	DIMM 2 Channel B SEC Counter	R/W	0000h
D6 – D7h	DRAM_DED_D2B	DIMM 2 Channel B DED Counter	R/W	0000h
D8 – D9h	DRAM_SEC_D3B	DIMM 3 Channel B SEC Counter	R/W	0000h
DA – DBh	DRAM_DED_D3B	DIMM 3 Channel B DED Counter	R/W	0000h
DC – DDh	DIMM_THR_EX	DIMM Threshold Exceeded	R/WC	0000h
E0 – E3h	SYSBUS_ERR_CTL	Host Error Control	R/W	0020_0000h
E4 – E7h	HI_ERR_CTL	Legacy Error Control	R/W	0004_0000h
EC – EFh	DRAM_ERR_CTL	DRAM Error Control	R/W	0000_0000h

3.6.1 VID – Vendor Identification (D0:F1)

Address Offset: 00 – 01h
Access RO
Size 16 Bits
Default 8086h

The VID register contains the vendor identification number. This 16-bit register combined with the Device Identification register uniquely identify any PCI device.

Bit Field	Default & Access	Description
15:0	8086h RO	**Vendor Identification (VID)**. This register field contains the PCI standard identification for Intel, 8086h.

3.6.2 DID – Device Identification (D0:F1)

Address Offset: 02 – 03h
Access RO
Size 16 Bits
Default 3593h
This 16-bit register combined with the Vendor Identification register uniquely identifies any PCI device.

Bit Field	Default & Access	Description
15:0	3593h RO	**Device Identification Number (DID)**. This is a 16-bit value assigned to the MCH Host-HI Bridge Function 1.

3.6.3 PCICMD – PCI Command Register (D0:F1)

Address Offset: 04 – 05h
Access: R/W
Size: 16 Bits
Default: 0000h

Since MCH Device 0 does not physically reside on PCI_A many of the bits are not implemented.

Bit Field	Default & Access	Description
15:9	00h	Reserved
8	0b R/W	**SERR Enable (SERRE).** This is a global enable bit for Device 0 SERR messaging. The MCH does not have an SERR signal. The MCH communicates the SERR condition by sending an SERR message over the HI to the ICH. 0 = Disable. The SERR message is not generated by the MCH for Device 0. 1 = Enable. The MCH is enabled to generate SERR messages over HI for specific Device 0 error conditions that are individually enabled in the Hub Interface SERR Command register (D0:F1:5Ch), System Bus SERR Command Register (D0:F1:6Ch), Memory Buffer SERR Command Register(D0:F1:7Ch), and the DRAM SERR Command Register(D0:F1:8Ch). **NOTE:** This bit only controls SERR messaging for the Device 0 Function 1 error conditions not handled by the Device 0 Function 0 SERR enable bit. Devices 1-7 have their own SERR bits to control error reporting for error conditions occurring on their respective devices. The control bits are used in a logical OR manner to enable the SERR HI message mechanism.
7:0	0b	Reserved

3.6.4 PCISTS – PCI Status Register (D0:F1)

Address Offset: 06 – 07h
Access R/WC
Size 16 Bits
Default 0000h

PCISTS is a 16-bit status register that reports the occurrence of error events on Device 0's PCI interface. Since MCH Device 0 does not physically reside on PCI_A many of the bits are not implemented.

Bit Field	Default & Access	Description
15	0b	Reserved
14	0b R/WC	**Signaled System Error (SSE)**. 0 = SERR not generated by MCH Device 0 1 = MCH Device 0 generated a SERR Software clears this bit by writing a '1' to the bit location.
13:0	000h	Reserved

3.6.5 RID – Revision Identification (D0:F1)

Address Offset: 08h
Access RO
Size 8 Bits
Default 09h

This register contains the revision number of the MCH Device 0.

Bit Field	Default & Access	Description
7:0	00h RO	**Revision Identification Number (RID)**. This value indicates the revision identification number for the MCH Device 0. This number should always be the same as the RID for Function 0. 09h = C1 stepping. 0Ah = C2 stepping.

3.6.6 SUBC – Sub-Class Code (D0:F1)

Address Offset: 0Ah
Access RO
Size 8 Bits
Default 00h

Bit Field	Default & Access	Description
7:0	00h RO	**Sub-Class Code (SUBC)**. This value indicates the Sub Class Code into which the MCH Device 0 Function 1 falls. 00h = Bridge

3.6.7 BCC – Base Class Code (D0:F1)

Address Offset: 0Bh
Access RO
Size 8 Bits
Default FFh

Bit Field	Default & Access	Description
7:0	FFh RO	**Base Class Code (BASEC)**. This value indicates the Base Class Code for the MCH Device 0 Function 1. FFh = A 'non-defined' device. Since this function is used for error conditions, it does not fall into any other class.

3.6.8 MLT – Master Latency Timer (D0:F1)

Address Offset: 0Dh
Access RO
Size 8 Bits
Default 00h

Device 0 in the MCH is not a PCI master, therefore this register is not implemented.

Bit Field	Default & Access	Description
7:0	00h	Reserved

3.6.9 HDR – Header Type (D0:F1)

Address Offset: 0Eh
Access RO
Size 8 Bits
Default 00h

Bit Field	Default & Access	Description
7:0	00h RO	**PCI Header (HDR)**. This value indicates the Header Type for the MCH Device 0. 00h = MCH is a multi-function device with a standard header layout.

3.6.10 SVID – Subsystem Vendor Identification (D0:F1)

Address Offset: 2C – 2D
Access R/WO
Size 16 Bits
Default 0000h

The MCH treats the SVID and SID as a single 32 bit register with regard to R/WO functionality. Any time a write access to the address 2C – 2Fh occurs, regardless of byte enables, entire 32 bit register comprised of SVID and SID locks.

Bit Field	Default & Access	Description
15:0	0000h R/WO	**Subsystem Vendor ID (SUBVID)**. This field should be programmed during boot-up to indicate the vendor of the system board.

3.6.11 SID – Subsystem Identification (D0:F1)

Address Offset: 2E – 2F
Access R/WO
Size 16 Bits
Default 0000h

The MCH treats the SVID and SID as a single 32 bit register with regard to R/WO functionality. Any time a write access to the address 2C – 2Fh occurs, regardless of byte enables, entire 32 bit register comprised of SVID and SID locks.

Bit Field	Default & Access	Description
15:0	0000h R/WO	**Subsystem ID (SUBID)**. This field should be programmed during BIOS initialization.

3.6.12 FERR_GLOBAL – Global First Error Register (D0:F1)

Address Offset: 40 – 43h
Access R/WC
Size 32 Bits
Default 0000_0000h

This register is used to log various error conditions at the "unit" level. These bits are "sticky" through reset, and are set regardless of whether or not any error messages (SCI, SMI, SERR#, MCERR#) are enabled and generated at the unit level. Specific error conditions within the various functional units are logged in the unit-specific error registers that follow.

Errors are reported in the FERR/NERR registers indicating the detection of either a fatal or non-fatal error. For these global error registers, a non-fatal error can be either an uncorrectable error which is non-fatal, or a correctable error.

This register captures the FIRST global Fatal and the FIRST global Non-Fatal errors. For these global error registers, a non-fatal error can be either an uncorrectable error which is non-fatal, or a correctable error. Any future errors (NEXT errors) will be captured in the NERR_Global register. No further error bits in this register will be set until the existing error bit is cleared.

Note: If multiple errors are reported in the same clock as the first error, all errors are latched.

Bit Field	Default & Access	Description
31:28	000b	Reserved
27	0b R/WC	**DRAM Controller Channel Fatal Error.** This bit is sticky through reset. System software clears this bit by writing a 1 to the location. 0 = No fatal DRAM I/F error. 1 = The MCH detected a fatal DRAM I/F error.
26	0b R/WC	**System Bus Fatal Error.** This bit is sticky through reset. System software clears this bit by writing a 1 to the location. 0 = No fatal system bus error. 1 = The MCH detected a fatal system bus error.
25	0b R/WC	**Hub Interface Fatal Error.** This bit is sticky through reset. System software clears this bit by writing a 1 to the location. 0 = No fatal HI error. 1 = The MCH detected a fatal HI error.
24	0b	Reserved
23	0b R/WC	**PCI Express* Port A Fatal Error.** This bit is sticky through reset. System software clears this bit by writing a 1 to the location. 0 = No fatal PCI Express* Port A error. 1 = The MCH detected a fatal PCI Express* Port A error.
22	0b R/WC	**PCI Express* Port A1 Fatal Error.** This bit is sticky through reset. System software clears this bit by writing a 1 to the location. 0 = No fatal PCI Express* Port A1 error. 1 = The MCH detected a fatal PCI Express* Port A1 error.
21	0b R/WC	**PCI Express* Port B Fatal Error.** This bit is sticky through reset. System software clears this bit by writing a 1 to the location. 0 = No fatal PCI Express* Port B error. 1 = The MCH detected a fatal PCI Express* Port B error.
20:15	00h	Reserved

Bit Field	Default & Access	Description
14	0b R/WC	**Internal Buffer Non-Fatal Error.** This bit is sticky through reset. System software clears this bit by writing a 1 to the location 0 = No non-fatal internal buffer error. 1 = The MCH detected a non-fatal internal buffer error.
13	0b R/WC	**DRAM Controller Non-Fatal Error.** This bit is sticky through reset. System software clears this bit by writing a 1 to the location. 0 = No non-fatal DRAM Controller error. 1 = The MCH detected a non-fatal DRAM Controller error.
12	0b R/WC	**System Bus Non-Fatal Error.** This bit is sticky through reset. System software clears this bit by writing a 1 to the location. 0 = No non-fatal system bus error. 1 = The MCH detected a non-fatal system bus error.
11	0b R/WC	**Hub Interface Non-Fatal Error.** This bit is sticky through reset. System software clears this bit by writing a 1 to the location. 0 = No non-fatal HI error. 1 = The MCH detected a non-fatal HI error.
10	0b	Reserved
9	0b R/WC	**PCI Express* Port A Non-Fatal Error.** This bit is sticky through reset. System software clears this bit by writing a 1 to the location. 0 = No non-fatal PCI Express* Port A error. 1 = The MCH detected a non-fatal PCI Express* Port A error.
8	0b R/WC	**PCI Express* Port A1 Non-Fatal Error.** This bit is sticky through reset. System software clears this bit by writing a 1 to the location. 0 = No non-fatal PCI Express* Port A1 error. 1 = The MCH detected a non-fatal PCI Express* Port A1 error.
7	0b R/WC	**PCI Express* Port B Non-Fatal Error.** This bit is sticky through reset. System software clears this bit by writing a 1 to the location. 0 = No non-fatal PCI Express* Port B error. 1 = The MCH detected a non-fatal PCI Express* Port B error.
6:0	000b	Reserved

3.6.13 NERR_GLOBAL – Global Next Error Register (D0:F1)

Address Offset: 44 – 47h
Access R/WC
Size 32 Bits
Default 0000_0000h

The bit definitions are the same as defined for FERR_GLOBAL, and are sticky through reset.

3.6.14 HI_FERR – Hub Interface First Error Register (D0:F1)

Address Offset: 50h
Access R/WC
Size 8 Bits
Default 00h

This register stores the first error related to the Hub Interface. Only one error bit will be set in this register. Any future errors (NEXT errors) will be net in the HI_NERR register. No further error bits in the HI_FERR register will be set until the existing error bit is cleared. These bits are sticky through reset. Software clears these bits by writing a '1' to the bit location.

Note: If multiple errors are reported in the same clock as the first error, all errors are latched.

Bit Field	Default & Access	Description
7	0b	Reserved
6	0b R/WC	**Hub Interface Target Abort.** This bit is sticky through reset. System software clears this bit by writing a '1' to the location. 0 = No Target Abort on HI 1 = An MCH-initiated HI cycle terminated with a Target Abort. Non-fatal
5	0b R/WC	**Enhanced Configuration Access Error.** This bit is sticky through reset. System software clears this bit by writing a '1' to the location. 0 = No Enhanced Configuration Access error 1 = A PCI Express* Enhanced Configuration access was mistakenly targeting the legacy interface. Fatal
4	0b R/WC	**HI Data Parity Error Detected.** This bit is sticky through reset. System software clears this bit by writing a '1' to the location. 0 = No HI data parity error. 1 = MCH has detected a parity error on the data phase of a HI transaction. Non-fatal.
3	0b R/WC	**Out of Range Address Error Detected.** This bit is sticky through reset. System software clears this bit by writing a '1' to the location. 0 = No Out of Range Address error detected. 1 = MCH has detected an attempted HI access to an Out of Range Address (includes addresses above 4G sent to the HI). Fatal.
2	0b R/WC	**HI Internal Parity Error Detected.** This bit is sticky through reset. System software clears this bit by writing a '1' to the location. 0 = No Internal Parity error detected. 1 = MCH HI bridge has detected an Internal Parity error. Non-fatal.
1	0b R/WC	**Illegal Access from HI Detected.** This bit is sticky through reset. System software clears this bit by writing a '1' to the location. 0 = No Illegal Access error detected. 1 = MCH has detected an attempted illegal access from the HI. Fatal.
0	0b R/WC	**HI Address/Command Parity Error Detected**. This bit is sticky through reset. System software clears this bit by writing a '1' to the location. 0 = No HI address or command parity error detected. 1 = MCH has detected a parity error on a HI address or command. Fatal.

3.6.15 HI_NERR – Hub Interface Next Error Register (D0:F1)

Address Offset: 52h
Access R/WC
Size 8 Bits
Default 00h

As mentioned in Section 3.6.14, the first HI error will be stored in the HI_FERR register. This register stores all subsequent HI errors. Multiple bits in this register may be set. The bit definitions in this register are identical to those in Section 3.6.14, "HI_FERR – Hub Interface First Error Register (D0:F1)" on page 3-81.

3.6.16 HI_ERRMASK – Hub Interface Error Mask Register (D0:F1)

Address Offset: 54h
Access R/W
Size 8 Bits
Default 00h

This register masks HI unit errors from being recognized, preventing them from being logged at the unit or global level, and no interrupt/messages are generated. These bits are sticky through reset.

Bit Field	Default & Access	Description
7	0b	Reserved
6	0b R/W	**Hub Interface Target Abort Mask**. This bit is sticky through reset. 0 = Enable HI Target Abort detection and reporting 1 = Mask HI Target Abort detection and reporting
5	0b R/W	**Enhanced Configuration Access Error Mask**. This bit is sticky through reset. 0 = Enable HI Enhanced Config Access error detection and reporting 1 = Mask HI Enhanced Config Access error detection and reporting
4	0b R/WC	**HI Data Parity Error Mask**. This bit is sticky through reset. 0 = Enable HI Data Parity error detection and reporting 1 = Mask HI Data Parity error detection and reporting
3	0b R/WC	**Out of Range Address Error Mask**. This bit is sticky through reset. 0 = Enable HI Out of Range Address error detection and reporting 1 = Mask HI Out of Range Address error detection and reporting
2	0b R/WC	**HI Internal Parity Error Mask**. This bit is sticky through reset. 0 = Enable HI Internal Parity error detection and reporting 1 = Mask HI Internal Parity error detection and reporting
1	0b RO	**Illegal Access from HI Mask**. This bit is sticky through reset. 0 = Enable HI Illegal Access error detection and reporting 1 = Mask HI Illegal Access error detection and reporting
0	0b R/WC	**HI Address/Command Parity Error Mask**. This bit is sticky through reset. 0 = Enable HI Address/Command Parity error detection and reporting 1 = Mask HI Address/Command Parity error detection and reporting

3.6.17 HI_SCICMD – Hub Interface SCI Command Register (D0:F1)

Address Offset: 58h
Access R/W
Size 8 Bits
Default 00h

This register enables various errors to generate an SCI HI special cycle. When an error flag is set in the FERR or NERR registers, it can generate an SCI HI special cycle when enabled in the SCICMD registers. Note that one and only one message type can be enabled.

Bit Field	Default & Access	Description
7	0b	Reserved
6	0b R/W	**Hub Interface Target Abort SCI Enable**. Controls whether or not an SCI is generated when bit 6 of either the HI_FERR or HI_NERR register is set. 0 = No SCI generated on HI Target Abort detection 1 = Enable SCI generation on HI Target Abort detection
5	0b R/W	**Enhanced Configuration Access Error SCI Enable**. Controls whether or not an SCI is generated when bit 5 of either the HI_FERR or HI_NERR register is set. 0 = No SCI generated on HI Enhanced Config Access error detection. 1 = Enable SCI generation on HI Enhanced Config Access error detection
4	0b R/W	**HI Data Parity Error SCI Enable**. Controls whether or not an SCI is generated when bit 4 of either the HI_FERR or HI_NERR register is set. 0 = No SCI generated on HI Data Parity error detection 1 = Enable SCI generation on HI Data Parity error detection
3	0b R/W	**Out of Range Address Error SCI Enable**. Controls whether or not an SCI is generated when bit 3 of either the HI_FERR or HI_NERR register is set. 0 = No SCI generated on HI Out of Range Address error detection 1 = Enable SCI generation on HI Out of Range Address error detection
2	0b R/W	**HI Internal Parity Error SCI Enable**. Controls whether or not an SCI is generated when bit 2 of either the HI_FERR or HI_NERR register is set. 0 = No SCI generated on HI Internal Parity error detection 1 = Enable SCI generation on HI Internal Parity error detection
1	0b R/W	**Illegal Access from HI SCI Enable**. Controls whether or not an SCI is generated when bit 1 of either the HI_FERR or HI_NERR register is set. 0 = No SCI generated on HI Illegal Access error detection 1 = Enable SCI generation on HI Illegal Access error detection
0	0b R/W	**HI Address/Command Parity Error SCI Enable**. Controls whether or not an SCI is generated when bit 0 of either the HI_FERR or HI_NERR register is set. 0 = No SCI generated on HI Address/Command Parity error detection 1 = Enable SCI generation on HI Address/Command Parity error detection

3.6.18 HI_SMICMD – Hub Interface SMI Command Register (D0:F1)

Address Offset: 5Ah
Access R/W
Size 8 Bits
Default 00h

This register enables various errors to generate an SMI HI special cycle. When an error flag is set in the FERR or NERR registers, it can generate an SMI HI special cycle when enabled in the SMICMD register. Note that one and only one message type can be enabled.

Bit Field	Default & Access	Description
7	0b	Reserved
6	0b R/W	**Hub Interface Target Abort SMI Enable.** Controls whether or not an SMI is generated when bit 6 of either the HI_FERR or HI_NERR register is set. 0 = No SMI generated on HI Target Abort detection 1 = Enable SMI generation on HI Target Abort detection
5	0b R/W	**Enhanced Configuration Access Error SMI Enable**. Controls whether or not an SMI is generated when bit 5 of either the HI_FERR or HI_NERR register is set. 0 = No SMI generated on HI Enhanced Config Access error detection. 1 = Enable SMI generation on HI Enhanced Config Access error detection
4	0b R/W	**HI Data Parity Error SMI Enable**. Controls whether or not an SMI is generated when bit 4 of either the HI_FERR or HI_NERR register is set. 0 = No SMI generated on HI Data Parity error detection 1 = Enable SMI generation on HI Data Parity error detection
3	0b R/W	**Out of Range Address Error SMI Enable**. Controls whether or not an SMI is generated when bit 3 of either the HI_FERR or HI_NERR register is set. 0 = No SMI generated on HI Out of Range Address error detection 1 = Enable SMI generation on HI Out of Range Address error detection
2	0b R/W	**HI Internal Parity Error SMI Enable**. Controls whether or not an SMI is generated when bit 2 of either the HI_FERR or HI_NERR register is set. 0 = No SMI generated on HI Internal Parity error detection 1 = Enable SMI generation on HI Internal Parity error detection
1	0b R/W	**Illegal Access from HI SMI Enable**. Controls whether or not an SMI is generated when bit 1 of either the HI_FERR or HI_NERR register is set. 0 = No SMI generated on HI Illegal Access error detection 1 = Enable SMI generation on HI Illegal Access error detection
0	0b R/W	**HI Address/Command Parity Error SMI Enable**. Controls whether or not an SMI is generated when bit 0 of either the HI_FERR or HI_NERR register is set. 0 = No SMI generated on HI Address/Command Parity error detection 1 = Enable SMI generation on HI Address/Command Parity error detection

3.6.19 HI_SERRCMD – Hub Interface SERR Command Register (D0:F1)

Address Offset: 5Ch
Access R/W
Size 8 Bits
Default 00h

This register enables various errors to generate an SERR HI special cycle. When an error flag is set in the FERR or NERR registers, it can generate an SERR HI special cycle when enabled in the SERRCMD register. Note that one and only one message type can be enabled.

Bit Field	Default & Access	Description
7	0b	Reserved
6	0b R/W	**Hub Interface Target Abort SERR Enable**. Controls whether or not an SERR is generated when bit 6 of either the HI_FERR or HI_NERR register is set. 0 = No SERR generated on HI Target Abort detection 1 = Enable SERR generation on HI Target Abort detection
5	0b R/W	**Enhanced Configuration Access Error SERR Enable**. Controls whether or not an SERR is generated when bit 5 of either the HI_FERR or HI_NERR register is set. 0 = No SERR generated on HI Enhanced Config Access error detection. 1 = Enable SERR generation on HI Enhanced Config Access error detection
4	0b R/W	**HI Data Parity Error SERR Enable**. Controls whether or not an SERR is generated when bit 4 of either the HI_FERR or HI_NERR register is set. 0 = No SERR generated on HI Data Parity error detection 1 = Enable SERR generation on HI Data Parity error detection
3	0b R/W	**Out of Range Address Error SERR Enable**. Controls whether or not an SERR is generated when bit 3 of either the HI_FERR or HI_NERR register is set. 0 = No SERR generated on HI Out of Range Address error detection 1 = Enable SERR generation on HI Out of Range Address error detection
2	0b R/W	**HI Internal Parity Error SERR Enable**. Controls whether or not an SERR is generated when bit 2 of either the HI_FERR or HI_NERR register is set. 0 = No SERR generated on HI Internal Parity error detection 1 = Enable SERR generation on HI Internal Parity error detection
1	0b R/W	**Illegal Access from HI SERR Enable**. Controls whether or not an SERR is generated when bit 1 of either the HI_FERR or HI_NERR register is set. 0 = No SERR generated on HI Illegal Access error detection 1 = Enable SERR generation on HI Illegal Access error detection
0	0b R/W	**HI Address/Command Parity Error SERR Enable**. Controls whether or not an SERR is generated when bit 0 of either the HI_FERR or HI_NERR register is set. 0 = No SERR generated on HI Address/Command Parity error detection 1 = Enable SERR generation on HI Address/Command Parity error detection

3.6.20 HI_MCERRCMD – Hub Interface MCERR# Register (D0:F1)

Address Offset: 5Eh
Access R/W
Size 8 Bits
Default 00h

This register enables various errors to assert the MCERR# signal on the system bus. When an error flag is set in the FERR or NERR registers, it can generate a MCERR# when enabled in the MCERRCMD.

Bit Field	Default & Access	Description
7	0b	Reserved
6	0b R/W	**Hub Interface Target Abort MCERR# Enable**. Controls whether or not an MCERR# is generated when bit 6 of either the HI_FERR or HI_NERR register is set. 0 = No MCERR# generated on HI Target Abort detection 1 = Enable MCERR# generation on HI Target Abort detection
5	0b R/W	**Enhanced Configuration Access Error MCERR# Enable**. Controls whether or not an MCERR# is generated when bit 5 of either the HI_FERR or HI_NERR register is set. 0 = No MCERR# generated on HI Enhanced Config Access error detection. 1 = Enable MCERR# generation on HI Enhanced Config Access error detection
4	0b R/W	**HI Data Parity Error MCERR# Enable**. Controls whether or not an MCERR is generated when bit 4 of either the HI_FERR or HI_NERR register is set. 0 = No MCERR# generated on HI Data Parity error detection 1 = Enable MCERR# generation on HI Data Parity error detection
3	0b R/W	**Out of Range Address Error MCERR# Enable**. Controls whether or not an MCERR# is generated when bit 3 of either the HI_FERR or HI_NERR register is set. 0 = No MCERR# generated on HI Out of Range Address error detection 1 = Enable MCERR# generation on HI Out of Range Address error detection
2	0b R/W	**HI Internal Parity Error MCERR# Enable**. Controls whether or not an MCERR# is generated when bit 2 of either the HI_FERR or HI_NERR register is set. 0 = No MCERR# generated on HI Internal Parity error detection 1 = Enable MCERR# generation on HI Internal Parity error detection
1	0b R/W	**Illegal Access from HI MCERR# Enable**. Controls whether or not an MCERR# is generated when bit 1 of either the HI_FERR or HI_NERR register is set. 0 = No MCERR# generated on HI Illegal Access error detection 1 = Enable MCERR# generation on HI Illegal Access error detection
0	0b R/W	**HI Address/Command Parity Error MCERR# Enable**. Controls whether or not an MCERR# is generated when bit 0 of either the HI_FERR or HI_NERR register is set. 0 = No MCERR# generated on HI Address/Command Parity error detection 1 = Enable MCERR# generation on HI Address/Command Parity error detection

3.6.21 SYSBUS_FERR – System Bus First Error Register (D0:F1)

Address Offset: 60 – 61h
Access R/WC
Size 16 Bits
Default 0000h

This register stores the first error related to the system bus. Only one error bit will be set in this register. Any future errors (NEXT errors) will be net in the SYSBUS_NERR register. No further error bits in the SYSBUS_FERR register will be set until the existing error bit is cleared. These bits are sticky through reset. Software clears these bits by writing a '1' to the bit location.

Note: If multiple errors are reported in the same clock as the first error, all errors are latched.

Bit Field	Default & Access	Description
15:10	00h	Reserved
9	0b R/WC	**Parity error from I/O subsystem.** This bit is sticky through reset. System software clears this bit by writing a '1' to the location. 0 = No parity error detected. 1 = Parity error detected on data from I/O subsystem heading for FSB. Non-fatal
8	0b R/WC	**Parity error from memory.** This bit is sticky through reset. System software clears this bit by writing a '1' to the location. 0 = No parity error detected. 1 = Parity error detected on data from memory heading for FSB. Non-fatal
7	0b R/WC	**System Bus BINIT# detected**. This bit is sticky through reset. System software clears this bit by writing a '1' to the location. 0 = No system bus BINIT# detected 1 = This bit is set on an electrical high-to-low transition (0 to 1) of BINIT#. Fatal
6	0b R/WC	**System Bus MCERR# detected**. This bit is sticky through reset. System software clears this bit by writing a '1' to the location. 0 = No system bus MCERR# detected 1 = This bit is set on an electrical high-to-low transition (0 to 1) of MCERR# when the chipset is not driving. Non-fatal
5	0b R/WC	**Non-DRAM Lock Error (NDLOCK)**. This bit is sticky through reset. System software clears this bit by writing a '1' to the location. 0 = No DRAM Lock Error detected 1 = MCH detected a lock operation to memory space that did not map into DRAM. Non-fatal
4	0b R/WC	**System Bus Address Above TOM (SBATOM)**. This bit is sticky through reset. System software clears this bit by writing a '1' to the location. 0 = No system bus address above TOM detected 1 = MCH has detected an address above DRB[7], which is the Top of Memory and above 4 GB. If the system has less than 4 GB of DRAM, then addresses between DRB[7] and 4 GB are sent to HI. Non-fatal
3	0b R/WC	**System Bus Data Parity Error (SBDPAR)**. This bit is sticky through reset. System software clears this bit by writing a '1' to the location. 0 = No system bus parity error detected. 1 = MCH has detected a data parity error on the system bus. Non-fatal
2	0b R/WC	**System Bus Address Strobe Glitch Detected (SBAGL).** This bit is sticky through reset. System software clears this bit by writing a '1' to the location. 0 = No system bus address strobe glitch detected. 1 = MCH has detected a glitch one of the System Bus address strobes. Fatal

Bit Field	Default & Access	Description
1	0b R/WC	**System Bus Data Strobe Glitch Detected (SBDGL)**. This bit is sticky through reset. System software clears this bit by writing a '1' to the location. 0 = No system bus data strobe glitch detected. 1 = MCH has detected a glitch one of the System Bus data strobes. Fatal
0	0b R/WC	**System Bus Request/Address Parity Error Detected (SBRPR)**. This bit is sticky through reset. System software clears this bit by writing a '1' to the location. 0 = No system bus request/address parity error detected. 1 = MCH has detected a parity error on either the address or request signals of the System Bus. Fatal

3.6.22 SYSBUS_NERR – System Bus Next Error Register (D0:F1)

Address Offset: 62 – 63h
Access R/WC
Size 16 Bits
Default 0000h

For bit definitions, see SYSBUS_FERR.

3.6.23 SYSBUS_ERRMASK – System Bus Error Mask Register (D0:F1)

Address Offset: 64 – 65h
Access R/W
Size 16 Bits
Default 0009h

This register masks the system bus unit errors from being recognized, preventing them from being logged at the unit or global level, and no interrupt/messages are generated. These bits are sticky through reset.

Bit Field	Default & Access	Description
15:10	00h	Reserved
9	0b R/W	**Parity error from I/O subsystem masked.** This bit is sticky through reset. 0 = Enable parity error detection and reporting. 1 = Mask parity error detection and reporting.
8	0b R/W	**Parity error from memory masked.** This bit is sticky through reset. 0 = Enable parity error detection and reporting. 1 = Mask parity error detection and reporting.
7	0b R/W	**System Bus BINIT# detected Mask.** This bit is sticky through reset. 0 = Enable System Bus BINIT# detection and reporting 1 = Mask System Bus BINIT# detection and reporting
6	0b R/W	**System Bus MCERR# detected Mask.** This bit is sticky through reset. 0 = Enable System Bus MCERR# detection and reporting 1 = Mask System Bus MCERR# detection and reporting
5	0b R/W	**Non-DRAM Lock Error Mask.** This bit is sticky through reset. 0 = Enable Non-DRAM Lock Error detection and reporting 1 = Mask Non-DRAM Lock Error detection and reporting

Bit Field	Default & Access	Description
4	0b R/W	**System Bus Address Above TOM Mask.** This bit is sticky through reset. 0 = Enable System Bus address above TOM detection and reporting 1 = Mask System Bus address above TOM detection and reporting
3	1b R/W	**System Bus Data Parity Error Mask.** This bit is sticky through reset. Note that system Bus Data Parity Errors are masked at boot. 0 = Enable System Bus Data Parity Error detection and reporting 1 = Mask System Bus Data Parity Error detection and reporting
2	0b R/W	**System Bus Address Strobe Glitch Detected Mask.** This bit is sticky through reset. 0 = Enable System Bus address strobe glitch detection and reporting 1 = Mask System Bus address strobe glitch detection and reporting
1	0b R/W	**System Bus Data Strobe Glitch Detected Mask.** This bit is sticky through reset. 0 = Enable System Bus data strobe glitch detection and reporting 1 = Mask System Bus data strobe glitch detection and reporting
0	1b R/W	**System Bus Request/Address Parity Error Detected Mask.** This bit is sticky through reset. Note that system Bus Request/Address Parity Errors are masked at boot. 0 = Enable System Bus request/address parity error detection and reporting 1 = Mask System Bus request/address detection and reporting

3.6.24 SYSBUS_SCICMD – System Bus SCI Command Register (D0:F1)

Address Offset: 68 – 69h
Access R/W
Size 16 Bits
Default 0000h

This register enables various errors to generate an SCI HI special cycle. When an error flag is set in the FERR or NERR registers, it can generate an SCI HI special cycle when enabled in the SCICMD register. Note that one and only one message type can be enabled.

Bit Field	Default & Access	Description
15:10	00h	Reserved
9	0b R/W	**Parity error from I/O subsystem SCI Enable.** Controls whether or not an SCI is generated when bit 9 of the SYSBUS_FERR or SYSBUS_NERR is set. 0 = No SCI generated on parity error detection. 1 = Enable SCI generation on parity error detection.
8	0b R/W	**Parity error from memory SCI Enable.** Controls whether or not an SCI is generated when bit 8 of the SYSBUS_FERR or SYSBUS_NERR is set. 0 = No SCI generated on parity error detection. 1 = Enable SCI generation on parity error detection.
7	0b R/W	**System Bus BINIT# detected SCI Enable**. Controls whether or not an SCI is generated when bit 7 of the SYSBUS_FERR or SYSBUS_NERR register is set. 0 = No SCI generated on System Bus BINIT# detection 1 = Enable SCI generation on System Bus BINIT# detection

Bit Field	Default & Access	Description
6	0b R/W	**System Bus MCERR# detected SCI Enable**. Controls whether or not an SCI is generated when bit 6 of the SYSBUS_FERR or SYSBUS_NERR register is set. 0 = No SCI generated on System Bus MCERR# detection 1 = Enable SCI generation on System Bus MCERR# detection
5	0b R/W	**Non-DRAM Lock Error SCI Enable**. Controls whether or not an SCI is generated when bit 5 of the SYSBUS_FERR or SYSBUS_NERR register is set. 0 = No SCI generated on Non-DRAM Lock Error detection 1 = Enable SCI generation on Non-DRAM Lock Error detection
4	0b R/W	**System Bus Address Above TOM SCI Enable**. Controls whether or not an SCI is generated when bit 4 of the SYSBUS_FERR or SYSBUS_NERR register is set. 0 = No SCI generated on System Bus address above TOM detection 1 = Enable SCI generation on System Bus address above TOM detection
3	0b R/W	**System Bus Data Parity Error SCI Enable**. Controls whether or not an SCI is generated when bit 3 of the SYSBUS_FERR or SYSBUS_NERRR register is set. 0 = No SCI generated on System Bus Data Parity Error detection 1 = Enable SCI generation on System Bus Data Parity Error detection
2	0b R/W	**System Bus Address Strobe Glitch Detected SCI Enable**. Controls whether or not an SCI is generated when bit 2 of the SYSBUS_FERR or SYSBUS_NERR register is set. 0 = No SCI generated on System Bus address strobe glitch detection 1 = Enable SCI generation on System Bus address strobe glitch detection
1	0b R/W	**System Bus Data Strobe Glitch Detected SCI Enable**. Controls whether or not an SCI is generated when bit 1 of the SYSBUS_FERR or SYSBUS_NERR register is set. 0 = No SCI generated on System Bus data strobe glitch detection 1 = Enable SCI generation on System Bus data strobe glitch detection
0	0b R/W	**System Bus Request/Address Parity Error Detected SCI Enable**. Controls whether or not an SCI is generated when bit 0 of the SYSBUS_FERR or SYSBUS_NERR register is set. 0 = No SCI generated on System Bus request/address parity error detection 1 = Enable SCI generation on System Bus request/address detection

3.6.25 SYSBUS_SMICMD – System Bus SMI Command Register (D0:F1)

Address Offset: 6A – 6Bh
Access R/W
Size 16 Bits
Default 0000h

This register enables various errors to generate an SMI HI special cycle. When an error flag is set in the SYSBUS_FERR or SYSBUS_NERR register, it can generate an SMI HI special cycle when enabled in the SMICMD register. Note that one and only one message type can be enabled.

Bit Field	Default & Access	Description
15:10	00h	Reserved
9	0b R/W	**Parity error from I/O subsystem SMI Enable.** Controls whether or not an SMI is generated when bit 9 of the SYSBUS_FERR or SYSBUS_NERR is set. 0 = No SMI generated on parity error detection. 1 = Enable SMI generation on parity error detection.
8	0b R/W	**Parity error from memory SMI Enable.** Controls whether or not an SMI is generated when bit 8 of the SYSBUS_FERR or SYSBUS_NERR is set. 0 = No SMI generated on parity error detection. 1 = Enable SMI generation on parity error detection.
7	0b R/W	**System Bus BINIT# detected SMI Enable**. Controls whether or not an SMI is generated when bit 7 of the SYSBUS_FERR or SYSBUS_NERR register is set. 0 = No SMI generated on System Bus BINIT# detection 1 = Enable SMI generation on System Bus BINIT# detection
6	0b R/W	**System Bus MCERR# detected SMI Enable**. Controls whether or not an SMI is generated when bit 6 of the SYSBUS_FERR or SYSBUS_NERR register is set. 0 = No SMI generated on System Bus MCERR# detection 1 = Enable SMI generation on System Bus MCERR# detection
5	0b R/W	**Non-DRAM Lock Error SMI Enable**. Controls whether or not an SMI is generated when bit 5 of the SYSBUS_FERR or SYSBUS_NERR register is set. 0 = No SMI generated on Non-DRAM Lock Error detection 1 = Enable SMI generation on Non-DRAM Lock Error detection
4	0b R/W	**System Bus Address Above TOM SMI Enable**. Controls whether or not an SMI is generated when bit 4 of the SYSBUS_FERR or SYSBUS_NERR register is set. 0 = No SMI generated on System Bus address above TOM detection 1 = Enable SMI generation on System Bus address above TOM detection
3	0b R/W	**System Bus Data Parity Error SMI Enable**. Controls whether or not an SMI is generated when bit 3 of the SYSBUS_FERR or SYSBUS_NERRR register is set. 0 = No SMI generated on System Bus Data Parity Error detection 1 = Enable SMI generation on System Bus Data Parity Error detection
2	0b R/W	**System Bus Address Strobe Glitch Detected SMI Enable**. Controls whether or not an SMI is generated when bit 2 of the SYSBUS_FERR or SYSBUS_NERR register is set. 0 = No SMI generated on System Bus address strobe glitch detection 1 = Enable SMI generation on System Bus address strobe glitch detection
1	0b R/W	**System Bus Data Strobe Glitch Detected SMI Enable**. Controls whether or not an SMI is generated when bit 1 of the SYSBUS_FERR or SYSBUS_NERR register is set. 0 = No SMI generated on System Bus data strobe glitch detection 1 = Enable SMI generation on System Bus data strobe glitch detection
0	0b R/W	**System Bus Request/Address Parity Error Detected SMI Enable**. Controls whether or not an SMI is generated when bit 0 of the SYSBUS_FERR or SYSBUS_NERR register is set. 0 = No SMI generated on System Bus request/address parity error detection 1 = Enable SMI generation on System Bus request/address detection

3.6.26 SYSBUS_SERRCMD – System Bus SERR Command Register (D0:F1)

Address Offset: 6C – 6Dh
Access R/W
Size 16 Bits
Default 0000h

This register enables various errors to generate an SERR HI special cycle. When an error flag is set in the SYSBUS_FERR or SYSBUS_NERR register, it can generate an SERR HI special cycle when enabled in the SERRCMD register. Note that one and only one message type can be enabled.

Bit Field	Default & Access	Description
15:10	00h	Reserved
9	0b R/W	**Parity error from I/O subsystem SERR Enable.** Controls whether or not an SERR is generated when bit 9 of the SYSBUS_FERR or SYSBUS_NERR is set. 0 = No SERR generated on parity error detection. 1 = Enable SERR generation on parity error detection.
8	0b R/W	**Parity error from memory SERR Enable.** Controls whether or not an SERR is generated when bit 8 of the SYSBUS_FERR or SYSBUS_NERR is set. 0 = No SERR generated on parity error detection. 1 = Enable SERR generation on parity error detection.
7	0b R/W	**System Bus BINIT# detected SERR Enable**. Controls whether or not an SERR is generated when bit 7 of the SYSBUS_FERR or SYSBUS_NERR register is set. 0 = No SERR generated on System Bus BINIT# detection 1 = Enable SERR generation on System Bus BINIT# detection
6	0b R/W	**System Bus MCERR# detected SERR Enable**. Controls whether or not an SERR is generated when bit 6 of the SYSBUS_FERR or SYSBUS_NERR register is set. 0 = No SERR generated on System Bus MCERR# detection 1 = Enable SERR generation on System Bus MCERR# detection
5	0b R/W	**Non-DRAM Lock Error SERR Enable**. Controls whether or not an SERR is generated when bit 5 of the SYSBUS_FERR or SYSBUS_NERR register is set. 0 = No SERR generated on Non-DRAM Lock Error detection 1 = Enable SERR generation on Non-DRAM Lock Error detection
4	0b R/W	**System Bus Address Above TOM SERR Enable**. Controls whether or not an SERR is generated when bit 4 of the SYSBUS_FERR or SYSBUS_NERR register is set. 0 = No SERR generated on System Bus address above TOM detection 1 = Enable SERR generation on System Bus address above TOM detection
3	0b R/W	**System Bus Data Parity Error SERR Enable**. Controls whether or not an SERR is generated when bit 3 of the SYSBUS_FERR or SYSBUS_NERRR register is set. 0 = No SERR generated on System Bus Data Parity Error detection 1 = Enable SERR generation on System Bus Data Parity Error detection
2	0b R/W	**System Bus Address Strobe Glitch Detected SERR Enable**. Controls whether or not an SERR is generated when bit 2 of the SYSBUS_FERR or SYSBUS_NERR register is set. 0 = No SERR generated on System Bus address strobe glitch detection 1 = Enable SERR generation on System Bus address strobe glitch detection

Bit Field	Default & Access	Description
1	0b R/W	**System Bus Data Strobe Glitch Detected SERR Enable**. Controls whether or not an SERR is generated when bit 1 of the SYSBUS_FERR or SYSBUS_NERR register is set. 0 = No SERR generated on System Bus data strobe glitch detection 1 = Enable SERR generation on System Bus data strobe glitch detection
0	0b R/W	**System Bus Request/Address Parity Error Detected SERR Enable**. Controls whether or not an SERR is generated when bit 0 of the SYSBUS_FERR or SYSBUS_NERR register is set. 0 = No SERR generated on System Bus request/address parity error detection 1 = Enable SERR generation on System Bus request/address detection

3.6.27 SYSBUS_MCERRCMD – System Bus MCERR# Command Register (D0:F1)

Address Offset: 6E – 6Fh
Access R/W
Size 16 Bits
Default 0000h

This register enables various errors to assert the MCERR# signal on the system bus. When an error flag is set in the SYSBUS_FERR or SYSBUS_NERR register, it can generate a MCERR# when enabled in the MCERRCMD.

Bit Field	Default & Access	Description
15:10	00h	Reserved
9	0b R/W	**Parity error from I/O subsystem MCERR# Enable.** Controls whether or not an MCERR is generated when bit 9 of the SYSBUS_FERR or SYSBUS_NERR is set. 0 = No MCERR generated on parity error detection. 1 = Enable MCERR generation on parity error detection.
8	0b R/W	**Parity error from memory MCERR# Enable.** Controls whether or not an MCERR is generated when bit 8 of the SYSBUS_FERR or SYSBUS_NERR is set. 0 = No MCERR generated on parity error detection. 1 = Enable MCERR generation on parity error detection.
7	0b R/W	**System Bus BINIT# detected MCERR# Enable**. Controls whether or not an MCERR# is generated when bit 7 of the SYSBUS_FERR or SYSBUS_NERR register is set. 0 = No MCERR# generated on System Bus BINIT# detection 1 = Enable MCERR# generation on System Bus BINIT# detection
6	0b R/W	**System Bus MCERR# detected MCERR# Enable**. Controls whether or not an MCERR# is generated when bit 6 of the SYSBUS_FERR or SYSBUS_NERR register is set. 0 = No MCERR# generated on System Bus MCERR# detection 1 = Enable MCERR# generation on System Bus MCERR# detection
5	0b R/W	**Non-DRAM Lock Error MCERR# Enable**. Controls whether or not an MCERR# is generated when bit 5 of the SYSBUS_FERR or SYSBUS_NERR register is set. 0 = No MCERR# generated on Non-DRAM Lock Error detection 1 = Enable MCERR# generation on Non-DRAM Lock Error detection

Bit Field	Default & Access	Description
4	0b R/W	**System Bus Address Above TOM MCERR# Enable**. Controls whether or not an MCERR# is generated when bit 4 of the SYSBUS_FERR or SYSBUS_NERR register is set. 0 = No MCERR# generated on System Bus address above TOM detection 1 = Enable MCERR# generation on System Bus address above TOM detection
3	0b R/W	**System Bus Data Parity Error MCERR# Enable**. Controls whether or not an MCERR# is generated when bit 3 of the SYSBUS_FERR or SYSBUS_NERRR register is set. 0 = No MCERR# generated on System Bus Data Parity Error detection 1 = Enable MCERR# generation on System Bus Data Parity Error detection
2	0b R/W	**System Bus Address Strobe Glitch Detected MCERR# Enable**. Controls whether or not an MCERR# is generated when bit 2 of the SYSBUS_FERR or SYSBUS_NERR register is set. 0 = No MCERR# generated on System Bus address strobe glitch detection 1 = Enable MCERR# generation on System Bus address strobe glitch detection
1	0b R/W	**System Bus Data Strobe Glitch Detected MCERR# Enable**. Controls whether or not an MCERR# is generated when bit 1 of the SYSBUS_FERR or SYSBUS_NERR register is set. 0 = No MCERR# generated on System Bus data strobe glitch detection 1 = Enable MCERR# generation on System Bus data strobe glitch detection
0	0b R/W	**System Bus Request/Address Parity Error Detected MCERR# Enable**. Controls whether or not an MCERR# is generated when bit 0 of the SYSBUS_FERR or SYSBUS_NERR register is set. 0 = No MCERR# generated on System Bus request/address parity error detection 1 = Enable MCERR# generation on System Bus request/address detection

3.6.28 BUF_FERR – Memory Buffer First Error Register (D0:F1)

Address Offset: 70h
Access: R/WC
Size: 8 bits
Default Value: 00h

This register stores the first error related to the coherent Posted Memory Write Buffer (PMWB). Only one error bit will be set in this register. Any future errors (NEXT errors) will be set in the BUF_NERR register. No further error bits in the BUF_FERR register will be set until the existing error bit is cleared. These bits are sticky through reset. Software clears these bits by writing a 1 to the bit location.

Note: If multiple errors are reported in the same clock as the first error, all errors are latched.

Bit Field	Default & Access	Description
7:4	0h	Reserved
3	0b R/WC	**Internal DRAM to PMWB Parity Error Detected**. Error detected when a cache line read from DRAM was written to the PMWB as part of a Read/Modify/Write operation (partial write). This bit is sticky through reset. System software clears this bit by writing a 1 to the location. 0 = No internal DRAM I/F to PMWB parity error detected 1 = Internal DRAM I/F to PMWB parity error detected. Non fatal

Bit Field	Default & Access	Description
2	0b R/WC	**Internal System Bus or I/O to PMWB Parity Error Detected**. Error detected on a System Bus or I/O write of a line to the PMWB. This bit is sticky through reset. System software clears this bit by writing a 1 to the location. 0 = No internal data parity error detected on line write 1 = Internal data parity error detected on line write to PMWB. Non fatal
1	0b R/WC	**Internal PMWB to System Bus Parity Error Detected.** This bit is sticky through reset. System software clears this bit by writing a 1 to the location. 0 = No internal PMWB to System Bus parity error detected 1 = Internal PMWB to System Bus parity error detected. Non-fatal
0	0b R/WC	**Internal PMWB to DRAM Parity Error Detected.** Error detected when PMWB is flushed to DRAM. This bit is sticky through reset. System software clears this bit by writing a 1 to the location. 0 = No internal PMWB to DRAM I/F parity error detected 1 = Internal PMWB to DRAM I/F parity error detected. Non-fatal

3.6.29 BUF_NERR – Memory Buffer Next Error Register (D0:F1)

Address Offset: 72h
Access: R/W
Size: 8 bits
Default Value: 00h

This register stores all errors after the first error related to the coherent Posted Memory Write Buffer (PMWB). These bits are sticky through reset. Software clears these bits by writing a 1 to the bit location.

Bit Field	Default & Access	Description
7:4	0h	Reserved
3	0b R/WC	**Internal DRAM to PMWB Parity Error Detected**. Error detected when a cache line read from DRAM was written to the PMWB as part of a Read/Modify/Write operation (partial write). This bit is sticky through reset. System software clears this bit by writing a 1 to the location. 0 = No internal DRAM I/F to PMWB parity error detected 1 = Internal DRAM I/F to PMWB parity error detected. Non fatal
2	0b R/WC	**Internal System Bus or I/O to PMWB Parity Error Detected**. Error detected on a System Bus or I/O write of a line to the PMWB. This bit is sticky through reset. System software clears this bit by writing a 1 to the location. 0 = No internal data parity error detected on line write 1 = Internal data parity error detected on line write to PMWB. Non fatal
1	0b R/WC	**Internal PMWB to System Bus Parity Error Detected.** This bit is sticky through reset. System software clears this bit by writing a 1 to the location. 0 = No internal PMWB to System Bus parity error detected 1 = Internal PMWB to System Bus parity error detected. Non-fatal
0	0b R/WC	**Internal PMWB to DRAM Parity Error Detected.** Error detected when PMWB is flushed to DRAM. This bit is sticky through reset. System software clears this bit by writing a 1 to the location. 0 = No internal PMWB to DRAM I/F parity error detected 1 = Internal PMWB to DRAM I/F parity error detected. Non-fatal

3.6.30 BUF_ERRMASK – Memory Buffer Error Mask Register (D0:F1)

Address Offset: 74h
Access: R/W
Size: 8 bits
Default Value: 00h

This register masks PMWB errors from being recognized, preventing them from being logged at the unit or global level, and no interrupt messages are generated. These bits are sticky through reset.

Bit Field	Default & Access	Description
7:4	00h	Reserved
3	0b R/W	**Internal DRAM to PMWB Parity Error Mask.** This bit is sticky through reset. 0 = Enable Internal DRAM to PMWB Parity Error detection and reporting 1 = Mask Internal DRAM to PMWB Parity Error detection and reporting
2	0b R/W	**Internal System Bus or I/O to PMWB Parity Error Mask.** This bit is sticky through reset. 0 = Enable internal System Bus or I/O to PMWB Parity Error detection and reporting 1 = Mask internal System Bus or I/O to PMWB Parity Error detection and reporting
1	0b R/W	**Internal PMWB to System Bus Parity Error Mask.** This bit is sticky through reset. 0 = Enable Internal PMWB to System Bus Parity Error detection and reporting 1 = Mask Internal PMWB to System Bus Parity Error detection and reporting
0	0b R/W	**Internal PMWB to DRAM Parity Error Mask.** This bit is sticky through reset. 0 = Enable Internal PMWB to DRAM Parity Error detection and reporting 1 = Mask Internal PMWB to DRAM Parity Error detection and reporting

3.6.31 BUF_SCICMD – Memory Buffer SCI Command Register (D0:F1)

Address Offset: 78h
Access: R/W
Size: 8 bits
Default Value: 00h

This register enables various errors to generate an SCI Hub Interface (HI) special cycle. When an error flag is set in the FERR or NERR registers, it can generate an SCI HI special cycle when enabled in the SCICMD registers. Note that only one message type can be enabled.

Bit Field	Default & Access	Description
7:4	0h	Reserved
3	0b R/W	**Internal DRAM I/F to PMWB Parity Error SCI Enable**. Controls whether or not an SCI is generated when bit 3 of the BUF_FERR or BUF_NERR register is set. 0 = No SCI on internal DRAM I/F to PMWB parity error detection 1 = Enable SCI generation on internal DRAM I/F to PMWB parity error detection

Bit Field	Default & Access	Description
2	0b R/W	**Internal System Bus or I/O to PMWB Parity Error SCI Enable**. Controls whether or not an SCI is generated when bit 2 of the BUF_FERR or BUF_NERR register is set. 0 = No SCI on internal System Bus or I/O to PMWB parity error detection 1 = Enable SCI generation on internal System Bus or I/O to PMWB data parity error detection
1	0b R/W	**Internal PMWB to System Bus Parity Error SCI Enable.** Controls whether or not an SCI is generated when bit 0 of the BUF_FERR or BUF_NERR register is set. 0 = No SCI on internal PMWB to System Bus parity error detection 1 = Enable SCI generation on internal PMWB to System Bus parity error detection
0	0b R/W	**Internal PMWB to DRAM I/F Parity Error SCI Enable.** Controls whether or not an SCI is generated when bit 1of the BUF_FERR or BUF_NERR register is set. 0 = No SCI on internal PMWB to DRAM I/F parity error detection 1 = Enable SCI generation on internal PMWB to DRAM I/F parity error detection

3.6.32 BUF_SMICMD – Memory Buffer SMI Command Register (D0:F1)

Address Offset: 7Ah
Access: R/W
Size: 8 bits
Default Value: 00h

This register enables various errors to generate an SMI HI special cycle. When an error flag is set in the FERR or NERR registers, it can generate an SMI HI special cycle when enabled in the SMICMD register. Note that one and only one message type can be enabled.

Bit Field	Default & Access	Description
7:4	0h	Reserved
3	0b R/W	**Internal DRAM I/F to PMWB Parity Error SMI Enable**. Controls whether or not an SMI is generated when bit 3 of the BUF_FERR or BUF_NERR register is set. 0 = No SMI on internal DRAM I/F to PMWB parity error detection 1 = Enable SMI generation on internal DRAM I/F to PMWB parity error detection
2	0b R/W	**Internal System Bus or I/O to PMWB Parity Error SMI Enable**. Controls whether or not an SMI is generated when bit 2 of the BUF_FERR or BUF_NERR register is set. 0 = No SMI on internal System Bus or I/O to PMWB parity error detection 1 = Enable SMI generation on internal System Bus or I/O to PMWB data parity error detection
1	0b R/W	**Internal PMWB to System Bus Parity Error SMI Enable.** Controls whether or not an SMI is generated when bit 0 of the BUF_FERR or BUF_NERR register is set. 0 = No SMI on internal PMWB to System Bus parity error detection 1 = Enable SMI generation on internal PMWB to System Bus parity error detection
0	0b R/W	**Internal PMWB to DRAM I/F Parity Error SMI Enable.** Controls whether or not an SMI is generated when bit 1 of the BUF_FERR or BUF_NERR register is set. 0 = No SMI on internal PMWB to DRAM I/F parity error detection 1 = Enable SMI generation on internal PMWB to DRAM I/F parity error detection

3.6.33 BUF_SERRCMD – Memory Buffer SERR Command Register (D0:F1)

Address Offset: 7Ch
Access: R/W
Size: 8 bits
Default Value: 00h

This register enables various errors to generate an SERR HI special cycle. When an error flag is set in the FERR or NERR registers, it can generate an SERR HI special cycle when enabled in the SERRCMD register. Note that only one message type can be enabled.

Bit Field	Default & Access	Description
7:4	0h	Reserved
3	0b R/W	**Internal DRAM I/F to PMWB Parity Error SERR Enable**. Controls whether or not an SERR is generated when bit 3 of the BUF_FERR or BUF_NERR register is set. 0 = No SERR on internal DRAM I/F to PMWB parity error detection 1 = Enable SERR generation on internal DRAM I/F to PMWB parity error detection
2	0b R/W	**Internal System Bus or I/O to PMWB Parity Error SERR Enable**. Controls whether or not an SERR is generated when bit 2 of the BUF_FERR or BUF_NERR register is set. 0 = No SERR on internal System Bus or I/O to PMWB parity error detection 1 = Enable SERR generation on internal System Bus or I/O to PMWB data parity error detection
1	0b R/W	**Internal PMWB to System Bus Parity Error SERR Enable.** Controls whether or not an SERR is generated when bit 0 of the BUF_FERR or BUF_NERR register is set. 0 = No SERR on internal PMWB to System Bus parity error detection 1 = Enable SERR generation on internal PMWB to System Bus parity error detection
0	0b R/W	**Internal PMWB to DRAM I/F Parity Error SERR Enable.** Controls whether or not an SERR is generated when bit 1 of the BUF_FERR or BUF_NERR register is set. 0 = No SERR on internal PMWB to DRAM I/F parity error detection 1 = Enable SERR generation on internal PMWB to DRAM I/F parity error detection

3.6.34 BUF_MCERRCMD – Memory Buffer MCERR# Command Register (D0:F1)

Address Offset: 7Eh
Access: R/W
Size: 8 bits
Default Value: 00h

This register enables various errors to assert a MCERR# signal on the system bus. When an error flag is set in the FERR or NERR registers, it can generate a MCERR# when enabled in the MCERRCMD register.

Bit Field	Default & Access	Description
7:4	0h	Reserved
3	0b R/W	**Internal DRAM I/F to PMWB Parity Error MCERR# Enable**. Controls whether or not an MCERR# is generated when bit 3 of the BUF_FERR or BUF_NERR register is set. 0 = No MCERR# on internal DRAM I/F to PMWB parity error detection 1 = Enable MCERR# generation on internal DRAM I/F to PMWB parity error detection
2	0b R/W	**Internal System Bus or I/O to PMWB Parity Error MCERR# Enable**. Controls whether or not an MCERR# is generated when bit 2 of the BUF_FERR or BUF_NERR register is set. 0 = No MCERR# on internal System Bus or I/O to PMWB parity error detection 1 = Enable MCERR# generation on internal System Bus or I/O to PMWB data parity error detection
1	0b R/W	**Internal PMWB to System Bus Parity Error MCERR# Enable.** Controls whether or not an MCERR# is generated when bit 0 of the BUF_FERR or BUF_NERR register is set. 0 = No MCERR# on internal PMWB to System Bus parity error detection 1 = Enable MCERR# generation on internal PMWB to System Bus parity error detection
0	0b R/W	**Internal PMWB to DRAM I/F Parity Error MCERR# Enable.** Controls whether or not an MCERR# is generated when bit 1 of the BUF_FERR or BUF_NERR register is set. 0 = No MCERR# on internal PMWB to DRAM I/F parity error detection 1 = Enable MCERR# generation on internal PMWB to DRAM I/F parity error detection

3.6.35 DRAM_FERR – DRAM First Error Register (D0:F1)

Address Offset: 80 – 81h
Access R/WC
Size 16 Bits
Default 0000h

This register stores the first error related to the DRAM Controller. Only one error bit will be set in this register. Any future errors (NEXT errors) will be net in the DRAM_NERR register. No further error bits in the DRAM_FERR register will be set until the existing error bit is cleared. These bits are sticky through reset. Software clears these bits by writing a '1' to the bit location.

Note: If multiple errors are reported in the same clock as the first error, all errors are latched.

Bit Field	Default & Access	Description
15	0b R/WC	**Memory Test Complete for Channel B**. Not an error condition. This bit is set by hardware to signal BIOS that HW testing of the channel is complete. This bit is sticky through reset. 0 = System software clears this bit by writing a '1' to the location. 1 = Hardware-based test of DRAM channel B is complete.
14	0b R/WC	**Uncorrectable Error on Write to Channel B**. This bit will be set on a detected error regardless of ECC mode, even if ECC is disabled. This bit is sticky through reset. System software clears this bit by writing a '1' to the location. 0 = No poisoned write to DRAM Channel B detected. 1 = Poisoned write to DRAM Channel B detected. Non-fatal.
13	0b R/WC	**DED Retry Initiated for Channel B**. (Uncorrectable) System software clears this bit by writing a '1' to the location. This bit is sticky through reset. This can be set for a normal demand data read, or for a scrub that is retried. 0 = No DED Retry Initiated for Channel B 1 = DED Retry Initiated for Channel B. Non-fatal.
12	0b R/WC	**Data Copy Complete for Channel B**. This bit does not indicate an error condition. This bit is set by hardware to signal BIOS that data copy for DIMM sparing is complete. This bit is sticky through reset. 0 = System software clears this bit by writing a '1' to the location. 1 = Data Copy Complete for Channel B.
11	0b R/WC	**Correctable Error Threshold Detect Channel B**. System software clears this bit by writing a '1' to the location. This bit is sticky through reset. 0 = No Error Threshold detected for Channel B 1 = Error Threshold detected for Channel B. Non-fatal.
10	0b R/WC	**Uncorrectable Scrubber Data Error Channel B.** System software clears this bit by writing a '1' to the location. This bit is sticky through reset. 0 = No Scrubber Error Detected for Channel B 1 = Scrubber Error Detected for Channel B, either for PMS (periodic memory scrubbing) or Sparing (which uses the scrubber to perform the data copy). Non-fatal.
9	0b R/WC	**Uncorrectable Read Memory Error Channel B**. Applies to non-scrub (normal demand fetch) reads and also indicated and unsuccessful retry if retry is enabled. System software clears this bit by writing a '1' to the location. This bit is sticky through reset. 0 = No Uncorrectable Non-Scrub Demand Read Memory Error Channel B 1 = Uncorrectable Non-Scrub Demand Read Memory Error Channel B. Non-fatal.

Bit Field	Default & Access	Description
8	0b R/WC	**Correctable Read Memory Error Channel B.** SEC (Single Bit Error Correction) detected by normal demand requests or scrubs are counted, and logged in FERR/NERR. This bit is sticky through reset. System software clears this bit by writing a '1' to the location. 0 = No Correctable Read Memory Error Channel B 1 = Correctable Read Memory Error Channel B. Non-fatal.
7	0b R/WC	**Memory Test Complete for Channel A**. Not an error condition. This bit is set by hardware to signal BIOS that HW testing of the channel is complete. This bit is sticky through reset. 0 = System software clears this bit by writing a '1' to the location. 1 = Hardware-based test of DRAM channel A is complete.
6	0b R/WC	**Uncorrectable Error on Write to Channel A**. This bit will be set on a detected error regardless of ECC mode, even if ECC is disabled. This bit is sticky through reset. System software clears this bit by writing a '1' to the location. 0 = No poisoned write to DRAM Channel A detected. 1 = Poisoned write to DRAM Channel A detected. Non-fatal.
5	0b R/WC	**DED Retry Initiated for Channel A**. This bit is sticky through reset. System software clears this bit by writing a '1' to the location. This can be set for a normal demand data read, or for a scrub that is retried. 0 = No DED Retry Initiated for Channel A 1 = DED Retry Initiated for Channel A. Non-fatal.
4	0b R/WC	**Data Copy Complete for Channel A**. This bit does not indicate an error condition. This bit is set by hardware to signal BIOS that data copy for DIMM sparing is complete. This bit sticky through reset. 0 = System software clears this bit by writing a '1' to the location. 1 = Data Copy Complete for Channel A. Non-fatal.
3	0b R/WC	**Correctable Error Threshold Detect Channel A**. This bit is sticky through reset. System software clears this bit by writing a '1' to the location. 0 = No Error Threshold detected for Channel B 1 = Error Threshold detected for Channel B. Non-fatal.
2	0b R/WC	**Uncorrectable Scrubber Data Error Channel A.** This bit is sticky through reset. System software clears this bit by writing a '1' to the location. 0 = No Scrubber Error Detected for Channel A 1 = Scrubber Error Detected for Channel A, either for PMS (periodic memory scrubbing) or Sparing. Non-fatal.
1	0b R/WC	**Uncorrectable Read Memory Error Channel A**. Applies to non-scrub (normal demand fetch) reads and also indicated and unsuccessful retry if retry is enabled. This bit is sticky through reset. System software clears this bit by writing a '1' to the location. 0 = No Uncorrectable Non-Scrub Demand Read Memory Error Channel A 1 = Uncorrectable Non-Scrub Demand Read Memory Error Channel A. Non-fatal.
0	0b R/WC	**Correctable Read Memory Error Channel A.** SEC (Single Bit Error Correction) detected by normal demand requests or scrubs are counted, and logged in FERR/NERR. This bit is sticky through reset. System software clears this bit by writing a '1' to the location. 0 = No Correctable Read Memory Error Channel A 1 = Correctable Read Memory Error Channel A. Non-fatal.

3.6.36 DRAM_NERR – DRAM Next Error Register (D0:F1)

Address Offset: 82 – 83h
Access R/WC
Size 16 Bits
Default 0000h

Signal errors occurring in the memory system. See DRAM_FERR for bit definitions

3.6.37 DRAM_ERRMASK – DRAM Error Mask Register (D0:F1)

Address Offset: 84h
Access R/W
Size 8 Bits
Default 00h

This register masks the DRAM Controller errors and events from being recognized, preventing them from being logged at the unit or global level, and no interrupt/messages are generated. Note that these bits apply to both channels A and B. These bits are sticky through reset.

Bit Field	Default & Access	Description
7	0b R/W	**Memory Test Complete Mask**. This bit is sticky through reset. 0 = Allow Memory Test Complete logging and signaling. 1 = Mask Memory Test Complete logging and signaling.
6	0b R/W	**Uncorrectable Error Detected on Write to DRAM Mask**. This bit is sticky through reset. 0 = Allow Poisoned Write to DRAM detection and signaling. 1 = Mask Poisoned Write to DRAM detection and signaling.
5	0b R/W	**DED Retry Initiated Mask**. This bit is sticky through reset. 0 = Allow DED Retry Initiated detection and signaling. 1 = Mask DED Retry Initiated detection and signaling.
4	0b R/W	**Data Copy Complete Mask**. This bit is sticky through reset. 0 = Allow Data Copy Complete logging and signaling. 1 = Mask Data Copy Complete logging and signaling.
3	0b R/W	**Error Threshold Detect Mask**. This bit is sticky through reset. 0 = Allow Error Threshold detection and signaling. 1 = Mask Error Threshold detection and signaling.
2	0b R/W	**Scrubber Data Error Mask.** This bit is sticky through reset. 0 = Allow Scrubber Data Error detection and signaling. 1 = Mask Scrubber Data Error detection and signaling.
1	0b R/W	**Uncorrectable Read Memory Error Mask**. This bit is sticky through reset. 0 = Allow Uncorrectable Memory Read Error detection and signaling. 1 = Mask Uncorrectable Memory Read Error detection and signaling.
0	0b R/W	**Correctable Read Memory Error Mask.** This bit is sticky through reset. 0 = Allow Correctable Memory Read Error detection and signaling. 1 = Mask Correctable Memory Read Error detection and signaling.

3.6.38 DRAM_SCICMD – DRAM SCI Command Register (D0:F1)

Address Offset: 88h
Access R/W
Size 8 Bits
Default 00h

This register enables various errors to generate an SCI HI special cycle. When an error flag is set in the FERR or NERR registers, it can generate an SERR, SMI, or SCI HI special cycle when enabled in the ERRCMD, SMICMD, or SCICMD registers, respectively. Note that one and only one message type can be enabled. These bits apply to both channel A and channel B.

Bit Field	Default & Access	Description
7	0b R/W	**Memory Test Complete SCI Enable.** Generate SCI when Bit 15 or 7 of DRAM_FERR or DRAM_NERR is set. 0 = Disable 1 = Enable
6	0b R/W	**Poisoned Write to DRAM SCI Enable.** Generate SCI when Bit 14 or 6 of DRAM_FERR or DRAM_NERR is set. 0 = Disable 1 = Enable
5	0b R/W	**DED Retry Initiated SCI Enable.** Generate SCI when Bit 13 or 5 of DRAM_FERR or DRAM_NERR is set. 0 = Disable 1 = Enable
4	0b R/W	**Data Copy Complete SCI Enable.** Generate SCI when Bit 12 or 4 of DRAM_FERR or DRAM_NERR is set. 0 = Disable 1 = Enable
3	0b R/W	**Error Threshold Detect SCI Enable.** Generate SCI when Bit 11 or 3 of DRAM_FERR or DRAM_NERR is set. 0 = Disable 1 = Enable
2	0b R/W	**Scrubber Data Error SCI Enable.** Generate SCI when Bit 10 or 2 of DRAM_FERR or DRAM_NERR is set. 0 = Disable 1 = Enable
1	0b R/W	**Uncorrectable Read Memory Error SCI Enable.** Generate SCI when Bit 9 or 1 of DRAM_FERR or DRAM_NERR is set. 0 = Disable 1 = Enable
0	0b R/W	**Correctable Read Memory Error SCI Enable.** Generate SCI when Bit 8 or 0 of DRAM_FERR or DRAM_NERR is set. 0 = Disable 1 = Enable

3.6.39 DRAM_SMICMD – DRAM SMI Command Register (D0:F1)

Address Offset: 8Ah
Access R/W
Size 8 Bits
Default 00h

This register enables various errors to generate an SMI HI special cycle. When an error flag is set in the FERR or NERR registers, it can generate an SERR, SMI, or SCI HI special cycle when enabled in the ERRCMD, SMICMD, or SCICMD registers, respectively. Note that one and only one message type can be enabled. These bits apply to both channel A and channel B.

Bit Field	Default & Access	Description
7	0b R/W	**Memory Test Complete SMI Enable.** Generate SMI when Bit 15 or 7 of DRAM_FERR or DRAM_NERR is set. 0 = Disable 1 = Enable
6	0b R/W	**Poisoned Write to DRAM SMI Enable.** Generate SMI when Bit 14 or 6 of DRAM_FERR or DRAM_NERR is set. 0 = Disable 1 = Enable
5	0b R/W	**DED Retry Initiated SMI Enable.** Generate SMI when Bit 13 or 5 of DRAM_FERR or DRAM_NERR is set. 0 = Disable 1 = Enable
4	0b R/W	**Data Copy Complete SMI Enable.** Generate SMI when Bit 12 or 4 of DRAM_FERR or DRAM_NERR is set. 0 = Disable 1 = Enable
3	0b R/W	**Error Threshold Detect SMI Enable.** Generate SMI when Bit 11 or 3 of DRAM_FERR or DRAM_NERR is set. 0 = Disable 1 = Enable
2	0b R/W	**Scrubber Data Error SMI Enable.** Generate SMI when Bit 10 or 2 of DRAM_FERR or DRAM_NERR is set. 0 = Disable 1 = Enable
1	0b R/W	**Uncorrectable Read Memory Error SMI Enable.** Generate SMI when Bit 9 or 1 of DRAM_FERR or DRAM_NERR is set. 0 = Disable 1 = Enable
0	0b R/W	**Correctable Read Memory Error SMI Enable.** Generate SMI when Bit 8 or 0 of DRAM_FERR or DRAM_NERR is set. 0 = Disable 1 = Enable

3.6.40 DRAM_SERRCMD – DRAM SERR Command Register (D0:F1)

Address Offset: 8Ch
Access R/W
Size 8 Bits
Default 00h

This register enables various errors to generate an SERR HI special cycle. When an error flag is set in the FERR or NERR registers, it can generate an SERR, SMI, or SCI HI special cycle when enabled in the SERRCMD, SMICMD, or SCICMD registers, respectively. Note that one and only one message type can be enabled. These bits apply to both channel A and channel B.

Bit Field	Default & Access	Description
7	0b R/W	**Memory Test Complete SERR Enable.** Generate SERR when Bit 15 or 7 of DRAM_FERR or DRAM_NERR is set. 0 = Disable 1 = Enable
6	0b R/W	**Poisoned Write to DRAM SERR Enable.** Generate SERR when Bit 14 or 6 of DRAM_FERR or DRAM_NERR is set. 0 = Disable 1 = Enable
5	0b R/W	**DED Retry Initiated SERR Enable.** Generate SERR when Bit 13 or 5 of DRAM_FERR or DRAM_NERR is set. 0 = Disable 1 = Enable
4	0b R/W	**Data Copy Complete SERR Enable.** Generate SERR when Bit 12 or 4 of DRAM_FERR or DRAM_NERR is set. 0 = Disable 1 = Enable
3	0b R/W	**Error Threshold Detect SERR Enable.** Generate SERR when Bit 11 or 3 of DRAM_FERR or DRAM_NERR is set. 0 = Disable 1 = Enable
2	0b R/W	**Scrubber Data Error SERR Enable.** Generate SERR when Bit 10 or 2 of DRAM_FERR or DRAM_NERR is set. 0 = Disable 1 = Enable
1	0b R/W	**Uncorrectable Read Memory Error SERR Enable.** Generate SERR when Bit 9 or 1 of DRAM_FERR or DRAM_NERR is set. 0 = Disable 1 = Enable
0	0b R/W	**Correctable Read Memory Error SERR Enable.** Generate SERR when Bit 8 or 0 of DRAM_FERR or DRAM_NERR is set. 0 = Disable 1 = Enable

3.6.41 DRAM_MCERRCMD – DRAM MCERR# Command Register (D0:F1)

Address Offset: 8Eh
Access R/W
Size 8 Bits
Default 00h

This register enables various errors to generate a MCERR# signal on the system bus. When an error flag is set in the FERR or NERR registers, it can generate a MCERR#, SMI, or SCI HI special cycle when enabled in the MCERRCMD register. These bits apply to both channel A and channel B.

Bit Field	Default & Access	Description
7	0b R/W	**Memory Test Complete MCERR# Enable.** Generate MCERR# when Bit 15 or 7 of DRAM_FERR or DRAM_NERR is set. 0 = Disable 1 = Enable
6	0b R/W	**Poisoned Write to DRAM MCERR# Enable.** Generate MCERR# when Bit 14 or 6 of DRAM_FERR or DRAM_NERR is set. 0 = Disable 1 = Enable
5	0b R/W	**DED Retry Initiated MCERR# Enable.** Generate MCERR# when Bit 13 or 5 of DRAM_FERR or DRAM_NERR is set. 0 = Disable 1 = Enable
4	0b R/W	**Data Copy Complete MCERR# Enable.** Generate MCERR# when Bit 12 or 4 of DRAM_FERR or DRAM_NERR is set. 0 = Disable 1 = Enable
3	0b R/W	**Error Threshold Detect MCERR# Enable.** Generate MCERR# when Bit 11 or 3 of DRAM_FERR or DRAM_NERR is set. 0 = Disable 1 = Enable
2	0b R/W	**Scrubber Data Error MCERR# Enable.** Generate MCERR# when Bit 10 or 2 of DRAM_FERR or DRAM_NERR is set. 0 = Disable 1 = Enable
1	0b R/W	**Uncorrectable Read Memory Error MCERR# Enable.** Generate MCERR# when Bit 9 or 1 of DRAM_FERR or DRAM_NERR is set. 0 = Disable 1 = Enable
0	0b R/W	**Correctable Read Memory Error MCERR# Enable.** Generate MCERR# when Bit 8 or 0 of DRAM_FERR or DRAM_NERR is set. 0 = Disable 1 = Enable

3.6.42 THRESH_SEC0 – DIMM0 SEC Threshold Register (D0:F1)

Address Offset: 98 – 99h
Access: R/W
Size: 16 Bits
Default: 0000h

Threshold compare value for SEC errors. An Error Threshold Detect is signaled if the logical DIMM0 SEC counter (Section 3.6.50 on page 3-111) exceeds the value programmed into this register. The bits in this register are sticky through reset.

Bit Field	Default & Access	Description
15:0	0000h R/W	**Threshold compare value for logical DIMM0 SEC errors**.

3.6.43 THRESH_SEC1 – DIMM1 SEC Threshold Register (D0:F1)

Address Offset: 9A – 9Bh
Access: R/W
Size: 16 Bits
Default: 0000h

Threshold compare value for SEC errors. An Error Threshold Detect is signaled if the logical DIMM1 SEC counter (Section 3.6.52 on page 3-112) exceeds the value programmed into this register. The bits in this register are sticky through reset.

Bit Field	Default & Access	Description
15:0	0000h R/W	**Threshold compare value for logical DIMM1 SEC errors**.

3.6.44 THRESH_SEC2 – DIMM2 SEC Threshold Register (D0:F1)

Address Offset: 9C – 9Dh
Access: R/W
Size: 16 Bits
Default: 0000h

Threshold compare value for SEC errors. An Error Threshold Detect is signaled if the logical DIMM2 SEC counter (Section 3.6.54 on page 3-113) exceeds the value programmed into this register. The bits in this register are sticky through reset.

Bit Field	Default & Access	Description
15:0	0000h R/W	**Threshold compare value for logical DIMM2 SEC errors**.

3.6.45 THRESH_SEC3 – DIMM3 SEC Threshold Register (D0:F1)

Address Offset: 9E – 9Fh
Access: R/W
Size: 16 Bits
Default: 0000h

Threshold compare value for SEC errors. An Error Threshold Detect is signaled if the logical DIMM3 SEC counter (Section 3.6.56 on page 3-113) exceeds the value programmed into this register. The bits in this register are sticky through reset.

Bit Field	Default & Access	Description
15:0	0000h R/W	**Threshold compare value for logical DIMM3 SEC errors**.

3.6.46 DRAM_SEC1_ADD – DRAM First Single-Bit Error Correct Address Register (D0:F1)

Address Offset: A0 – A3h
Access RO
Size 32 Bits
Default 0000_0000h

Captures the address of the first SEC error occurring in the memory system corresponding to the bit in the FERR register. This register saves the first address of a correctable read data error on a read, including scrubs. A correctable error on a sparing read would also load this register. The contents of this register correspond to a correctable error being set in the DRAM_FERR register (either bit 0 or 8 depending on channel).

Bit Field	Default & Access	Description
31	0b	Reserved
30:2	000_0000h RO	**First Correctable Error Address**. This field contains address bits 34:6 for the first correctable error. This field is set by HW, and represents a system address. This field can only be reset by a PWRGD reset.
1:0	00b	Reserved

3.6.47 DRAM_DED_ADD – DRAM DED Error Address (D0:F1)

Address Offset: A4 – A7h
Access RO
Size 32 Bits
Default 0000_0000h

Captures the address of the first DED error (uncorrectable, non-scrub engine) occurring in the memory system. If DED Retry is enabled, this register will capture the address of the first failed retry. The value in this register is only valid if the Uncorrectable Read Memory Error bit in either the DRAM_FERR or DRAM_NERR register has been set. The bits in this register are sticky through reset.

Bit Field	Default & Access	Description
31	0b	Reserved
30:2	000_0000h RO	**First Uncorrectable Error Address.** This field contains address bits 34:6 for the first uncorrectable error. If DED Retry is enabled, the address of the first failed retry will be captured. This field is set by HW, and represents a system address. This field can only be reset by a PWRGD reset.
1:0	00b	Reserved

3.6.48 DRAM_SCRB_ADD – DRAM Scrub Error Address Register (D0:F1)

Address Offset: A8 – ABh
Access RO
Size 32 Bits
Default 0000_0000h

Captures the address of the first uncorrectable error encountered by either the periodic memory scrubber or the sparing engine. The contents of this register correspond to the errors bits 2 and 10 of the DRAM_FERR/NERR registers (depending on channel). The bits in this register are sticky through reset.

Bit Field	Default & Access	Description
31	0b	Reserved
30:2	000_0000h RO	**Scrub Error Address**. This field contains address bits 34:6 for the first uncorrectable error encountered by the periodic memory scrubber or the sparing engine. This field is set by HW, and represents a system address. This field can only be reset by a PWRGD reset.
1:0	00b	Reserved

3.6.49 DRAM_RETR_ADD – DRAM DED Retry Address (D0:F1)

Address Offset: AC – AFh
Access RO
Size 32 Bits
Default 0000_0000h

Captures the address of the first DED retry occurring in the memory system. This retry address can be for a normal demand data read, or for a scrub that is retried. The bits in this register are sticky through reset. The contents of this register correspond to the errors bits 5 and 13 of the DRAM_FERR/NERR registers (depending on channel).

Bit Field	Default & Access	Description
31	0b	Reserved
30:2	0 RO	**DED Retry Address**. This field contains address bits 34:6 for the first DED Retry occurring in the memory system. This field is set by HW, and represents a system address. This field can only be reset by a PWRGD reset.
1:0	00b	Reserved

3.6.50 DRAM_SEC_D0A – DRAM DIMM0 Channel A SEC Counter Register (D0:F1)

Address Offset: B0 – B1h
Access R/W
Size 16 Bits
Default 0000h

Counter for SEC errors occurring for logical DIMM0 of channel A. The DIMM counters for SEC and DED are implemented using a leaky-bucket concept. The error count returned when this register is read is not an absolute count over time, but the sum of errors during a current specified time period (specified in SPRCTL – D0:F0:90-93h) plus half of the accumulated errors from past time periods. When a time period expires, the sum of the current time period accumulated errors and a value equal to half of the past accumulated errors is retained. Half of this registered error value will be added to the errors accumulated during the next time period. This method is employed because it is not the absolute number of errors that is most interesting, but the rate that errors occur. For more details on this feature, please refer to Section 5.3.2.6. The bits in this register are sticky through reset.

Bit Field	Default & Access	Description
15:0	0000h R/W	**DIMM0 Channel A SEC Count**

3.6.51 DRAM_DED_D0A – DRAM DIMM0 Channel A DED Counter Register (D0:F1)

Address Offset: B2 – B3h
Access R/W
Size 16 Bits
Default 0000h

Counter for DED errors occurring for logical DIMM0 of channel A. The functionality of this counter is described in Section 3.6.50 on page 3-111. The bits in this register are sticky through reset.

When DED retry is enabled, and the second read is successful, no DED is counted. It is only when the second read is unsuccessful, will a DED be counted for a DED retry. Data errors that occur for a DIMM spare will be counted by the same counter that logged for the errors for the victim.

Bit Field	Default & Access	Description
15:0	0000h R/W	**DIMM0 Channel A DED Count**.

3.6.52 DRAM_SEC_D1A – DRAM DIMM1 Channel A SEC Counter Register (D0:F1)

Address Offset: B4 – B5h
Access R/W
Size 16 Bits
Default 0000h

Counter for SEC errors occurring for logical DIMM1 of channel A. The functionality of this counter is described in Section 3.6.50 on page 3-111. The bits in this register are sticky through reset.

Bit Field	Default & Access	Description
15:0	0000h R/W	**DIMM1 Channel A SEC Count**

3.6.53 DRAM_DED_D1A – DRAM DIMM1 Channel A DED Counter Register (D0:F1)

Address Offset: B6 – B7h
Access R/W
Size: 16 Bits
Default: 0000h

Counter for DED errors occurring for logical DIMM1 of channel A. The functionality of this counter is described in Section 3.6.50 on page 3-111. The bits in this register are sticky through reset.

Bit Field	Default & Access	Description
15:0	0000h R/W	**DIMM1 Channel A DED Count**

3.6.54 DRAM_SEC_D2A – DRAM DIMM2 Channel A SEC Counter Register (D0:F1)

Address Offset: B8 – B9h
Access: R/W
Size: 16 Bits
Default: 0000h

Counter for SEC errors occurring for logical DIMM2 of channel A. The functionality of this counter is described in Section 3.6.50 on page 3-111. The bits in this register are sticky through reset.

Bit Field	Default & Access	Description
15:0	0000h R/W	**DIMM2 Channel A SEC Count**

3.6.55 DRAM_DED_D2A – DRAM DIMM2 Channel A DED Counter Register (D0:F1)

Address Offset: BA – BBh
Access: R/W
Size: 16 Bits
Default: 0000h

Counter for DED errors occurring for logical DIMM2 of channel A. The functionality of this counter is described in Section 3.6.50 on page 3-111. The bits in this register are sticky through reset.

Bit Field	Default & Access	Description
15:0	0000h R/W	**DIMM2 Channel A DED Count**

3.6.56 DRAM_SEC_D3A – DRAM DIMM3 Channel A SEC Counter Register (D0:F1)

Address Offset: BC – BDh
Access: R/W
Size: 16 Bits
Default: 0000h

Counter for SEC errors occurring for logical DIMM3 of channel A. The functionality of this counter is described in Section 3.6.50 on page 3-111. The bits in this register are sticky through reset.

Bit Field	Default & Access	Description
15:0	0000h R/W	**DIMM3 Channel A SEC Count**

3.6.57 DRAM_DED_D3A – DRAM DIMM3 Channel A DED Counter Register (D0:F1)

Address Offset: BE – BFh
Access: R/W
Size: 16 Bits
Default: 0000h

Counter for DED errors occurring for logical DIMM3 of channel A. The functionality of this counter is described in Section 3.6.50 on page 3-111. The bits in this register are sticky through reset.

Bit Field	Default & Access	Description
15:0	0000h R/W	**DIMM3 Channel A DED Count**

3.6.58 THRESH_DED – DED Threshold Register (D0:F1)

Address Offset: C2 – C3h
Access: R/W
Size: 16 Bits
Default: 0000h

Threshold compare value for DED errors. An Error Threshold Detect is signaled if any of the DED counters exceeds the value programmed into this register. The bits in this register are sticky through reset.

Bit Field	Default & Access	Description
15:0	0000h R/W	**Threshold compare value for DED errors**.

3.6.59 DRAM_SEC1_SYNDROME – First Single-Bit Error Correct Syndrome (D0:F1)

Address Offset: C4 – C5h
Access: RO
Size: 16 Bits
Default: 0000h

Syndrome for correctable errors occurring in the memory system. The contents of this register are set when a correctable error bit is being set in the DRAM_FERR register (either bit 0 or 8 or both). Syndrome is always logged for QW0/2 or QW1/3 pairs irrespective of single/dual channel and ECC modes, if transferring the lower half of the cache line, and logged for QW4/6 or QW5/7 if transferring the upper half of the cache line. ECC is checked half cacheline at a time. The syndrome logged in this register is for the lowest ordered QW pair. For example: If both QW0/2 and QW1/3 have correctable errors, the syndrome stored will be for QW0/2. If the correctable errors occurred on the same channel within this QW pair, correctable errors for both channels would be flagged in the DRAM_FERR register. When in S4EC (chip fail) mode, the syndrome stored is a 16 bit quantity across both channels. When in SEC mode, the syndrome is stored as two

separate quantities, bit 15:8 are for channel A while bits 7:0 are for channel B. A syndrome indicates error when it is a non-zero value. Syndrome logged for correctable data errors that occur when mirroring mode is enabled may not be recorded with the proper channel designation.

The bits in this register are sticky through reset.

Bit Field	Default & Access	Description
15:0	0000h RO	**ECC Syndrome for first correctable error**. Because only hardware writes to this register, it is read only. **SEC Mode (72-bit ECC)** Bits 15:8 Channel A Bits 7:0 Channel B **x4 SDDC Mode** Bits 15:0 Channels A & B

3.6.60 DRAM_SEC2_SYNDROME – Second Single-Bit Error Correct Syndrome (D0:F1)

Address Offset: C6 – C7h
Access: RO
Size: 16 Bits
Default: 0000h

Syndrome for next correctable error occurring in the memory system. The contents of this register correspond to the correctable error bit being set in the DRAM_NERR register (either bit 0 or 8 depending on channel). The bits in this register are sticky through reset.

Bit Field	Default & Access	Description
15:0	0000h RO	**ECC Syndrome for second correctable error**. Because only hardware writes to this register, it is read only. **SEC Mode (72-bit ECC)** Bits 15:8 Channel A Bits 7:0 Channel B **x4 SDDC Mode** Bits 15:0 Channels A & B

3.6.61 DRAM_SEC2_ADD – DRAM Next Single-Bit Error Correct Address Register (D0:F1)

Address Offset: C8 – CBh
Access: RO
Size: 32 Bits
Default: 0000_0000h

Captures the address of the next SEC error (either normal or scrub read) occurring in the memory system corresponding to the bit set in the NERR register. The value in this register correspond to a correctable error being set in the DRAM_NERR register (either bit 0 or 8 depending on channel). The bits in this register are sticky through reset.

Bit Field	Default & Access	Description
31	0b	Reserved
30:2	000_0000h RO	**Next Correctable Error Address**. This field contains address bits 35:12 for the next correctable error. This field is set by HW when the Correctable Read Memory Error bit in the DRAM_SERR register is set. This value represents a system address. This field can only be reset by a PWRGD reset.
1:0	00b	Reserved

3.6.62 DRAM_SEC_D0B – DRAM DIMM0 Channel B SEC Counter Register (D0:F1)

Address Offset: CC – CDh
Access: R/W
Size: 16 Bits
Default: 0000h

Counter for SEC errors occurring for logical DIMM0 of channel B. The functionality of this counter is described in Section 3.6.50 on page 3-111. The bits in this register are sticky through reset.

Bit Field	Default & Access	Description
15:0	0000h R/W	**DIMM0 Channel B SEC Count**

3.6.63 DRAM_DED_D0B – DRAM DIMM0 Channel B DED Counter Register (D0:F1)

Address Offset: CE – CFh
Access: R/W
Size: 16 Bits
Default: 0000h

Counter for DED errors occurring for logical DIMM0 of channel B. The functionality of this counter is described in Section 3.6.50 on page 3-111. The bits in this register are sticky through reset.

Bit Field	Default & Access	Description
15:0	0000h R/W	**DIMM0 Channel B DED Count**

3.6.64 DRAM_SEC_D1B – DRAM DIMM1 Channel B SEC Counter Register (D0:F1)

Address Offset: D0 – D1h
Access: R/W
Size: 16 Bits
Default: 0000h

Counter for SEC errors occurring for logical DIMM1 of channel B. The functionality of this counter is described in Section 3.6.50 on page 3-111. The bits in this register are sticky through reset.

Bit Field	Default & Access	Description
15:0	0000h R/W	**DIMM1 Channel B SEC Count**

3.6.65 DRAM_DED_D1B – DRAM DIMM1 Channel B DED Counter Register (D0:F1)

Address Offset: D2 – D3h
Access: R/W
Size: 16 Bits
Default: 0000h

Counter for DED errors occurring for logical DIMM1 of channel B. The functionality of this counter is described in Section 3.6.50 on page 3-111. The bits in this register are sticky through reset.

Bit Field	Default & Access	Description
15:0	0000h R/W	**DIMM1 Channel B DED Count**

3.6.66 DRAM_SEC_D2B – DRAM DIMM2 Channel B SEC Counter Register (D0:F1)

Address Offset: D4 – D5h
Access: R/W
Size: 16 Bits
Default: 0000h

Counter for SEC errors occurring for logical DIMM2 of channel B. The functionality of this counter is described in Section 3.6.50 on page 3-111. The bits in this register are sticky through reset.

Bit Field	Default & Access	Description
15:0	0000h R/W	**DIMM2 Channel B SEC Count**

3.6.67 DRAM_DED_D2B – DRAM DIMM2 Channel B DED Counter Register (D0:F1)

Address Offset: D6 – D7h
Access: R/W
Size: 16 Bits
Default: 0000h

Counter for DED errors occurring for logical DIMM2 of channel B. The functionality of this counter is described in Section 3.6.50 on page 3-111. The bits in this register are sticky through reset.

Bit Field	Default & Access	Description
15:0	0000h R/W	**DIMM2 Channel B DED Count**

3.6.68 DRAM_SEC_D3B – DRAM DIMM3 Channel B SEC Counter Register (D0:F1)

Address Offset: D8 – D9h
Access: R/W
Size: 16 Bits
Default: 0000h

Counter for SEC errors occurring for logical DIMM3 of channel B. The functionality of this counter is described in Section 3.6.50 on page 3-111. The bits in this register are sticky through reset.

Bit Field	Default & Access	Description
15:0	0000h R/W	**DIMM3 Channel B SEC Count**

3.6.69 DRAM_DED_D3B – DRAM DIMM3 Channel B DED Counter Register (D0:F1)

Address Offset: DA – DBh
Access: R/W
Size: 16 Bits
Default: 0000h

Counter for DED errors occurring for logical DIMM3 of channel B. The functionality of this counter is described in Section 3.6.50 on page 3-111. The bits in this register are sticky through reset.

Bit Field	Default & Access	Description
15:0	0000h RO	**DIMM3 Channel B DED Count**

3.6.70 DIMM_THR_EX – DIMM Threshold Exceeded Register (D0:F1)

Address Offset: DC – DDh
Access: R/WC
Size: 16 Bits
Default: 0000h

Preserves knowledge of DIMM error thresholds exceeded. The bits in this register are sticky through reset.

Bit Field	Default & Access	Description
15	0b R/WC	**Channel B logical DIMM 3 DED Threshold Status**. This bit is sticky through reset. Software can clear this bit by writing a '1' to the bit location. 0 = Threshold not exceeded 1 = Threshold exceeded.
14	0b R/WC	**Channel B logical DIMM 2 DED Threshold Status**. This bit is sticky through reset. Software can clear this bit by writing a '1' to the bit location. 0 = Threshold not exceeded 1 = Threshold exceeded
13	0b R/WC	**Channel B logical DIMM 1 DED Threshold Status**. This bit is sticky through reset. Software can clear this bit by writing a '1' to the bit location. 0 = Threshold not exceeded 1 = Threshold exceeded
12	0b R/WC	**Channel B logical DIMM 0 DED Threshold Status**. This bit is sticky through reset. Software can clear this bit by writing a '1' to the bit location. 0 = Threshold not exceeded 1 = Threshold exceeded
11	0b R/WC	**Channel B logical DIMM 3 SEC Threshold Status**. This bit is sticky through reset. Software can clear this bit by writing a '1' to the bit location. 0 = Threshold not exceeded 1 = Threshold exceeded
10	0b R/WC	**Channel B logical DIMM 2 SEC Threshold Status**. This bit is sticky through reset. Software can clear this bit by writing a '1' to the bit location. 0 = Threshold not exceeded 1 = Threshold exceeded
9	0b R/WC	**Channel B logical DIMM 1 SEC Threshold Status**. This bit is sticky through reset. Software can clear this bit by writing a '1' to the bit location. 0 = Threshold not exceeded 1 = Threshold exceeded
8	0b R/WC	**Channel B logical DIMM 0 SEC Threshold Status**. This bit is sticky through reset. Software can clear this bit by writing a '1' to the bit location. 0 = Threshold not exceeded 1 = Threshold exceeded
7	0b R/WC	**Channel A logical DIMM 3 DED Threshold Status**. This bit is sticky through reset. Software can clear this bit by writing a '1' to the bit location. 0 = Threshold not exceeded 1 = Threshold exceeded

Bit Field	Default & Access	Description
6	0b R/WC	**Channel A logical DIMM 2 DED Threshold Status**. This bit is sticky through reset. Software can clear this bit by writing a '1' to the bit location. 0 = Threshold not exceeded 1 = Threshold exceeded
5	0b R/WC	**Channel A logical DIMM 1 DED Threshold Status**. This bit is sticky through reset. Software can clear this bit by writing a '1' to the bit location. 0 = Threshold not exceeded 1 = Threshold exceeded
4	0b R/WC	**Channel A logical DIMM 0 DED Threshold Status**. This bit is sticky through reset. Software can clear this bit by writing a '1' to the bit location. 0 = Threshold not exceeded 1 = Threshold exceeded
3	0b R/WC	**Channel A logical DIMM 3 SEC Threshold Status**. This bit is sticky through reset. Software can clear this bit by writing a '1' to the bit location. 0 = Threshold not exceeded 1 = Threshold exceeded
2	0b R/WC	**Channel A logical DIMM 2 SEC Threshold Status**. This bit is sticky through reset. Software can clear this bit by writing a '1' to the bit location. 0 = Threshold not exceeded 1 = Threshold exceeded
1	0b R/WC	**Channel A logical DIMM 1 SEC Threshold Status**. This bit is sticky through reset. Software can clear this bit by writing a '1' to the bit location. 0 = Threshold not exceeded 1 = Threshold exceeded
0	0b R/WC	**Channel A logical DIMM 0 SEC Threshold Status.** This bit is sticky through reset. Software can clear this bit by writing a '1' to the bit location. 0 = Threshold not exceeded 1 = Threshold exceeded

3.6.71 SYSBUS_ERR_CTL – System Bus Error Control Register (D0:F1)

Address Offset: E0 – E3h
Access: R/W
Size: 32 Bits
Default: 0020_0000h

This register controls the way in which the MCH handles parity errors detected on incoming data streams into the MCH core from the system bus.

Bit Field	Default & Access	Description
31:19	0004h	Reserved
18	0b R/W	**Data Poisoning Enable**. This bit controls whether or not the MCH marks data as "poisoned" when a parity error is detected on incoming data from the system bus. 0 = Error checking disabled. 1 = Error checking enabled. Incoming data with parity errors will be marked as "poisoned" before being sent on towards its destination.
17:0	0	Reserved

3.6.72 HI_ERR_CTL – Hub Interface Error Control Register (D0:F1)

Address Offset: E4 – E7h
Access: R/W
Size: 32 Bits
Default: 0004_0000h

This register controls the way in which the MCH handles parity errors on the Hub Interface.

Bit Field	Default & Access	Description
31:20	000h	Reserved
19	0b R/W	**Stop and Scream**. This is a special control for errors going to the Hub Interface, outgoing from the MCH core. 0 = Outgoing data errors are propagated. 1 = Outgoing data errors are reported, but not propagated.
18	1b R/W	**Data Poisoning Enable**. This bit controls whether or not the MCH marks data as "poisoned" when a parity error is detected from the HI. 0 = Error checking disabled. 1 = Error Checking Enabled. Incoming data with parity errors will be marked as "poisoned" before being sent on towards its destination.
17:0	0	Reserved

3.6.73 BUF_ERR_CTL – Buffer Error Control Register (D0:F1)

Address Offset: E8 – EBh
Access: R/W
Size: 32 Bits
Default: 0000_0000h

This register controls the MCH handling of errors on incoming data streams into the MCH core from the posted write buffer.

Bit Field	Default & Access	Description
31:19	0000h	Reserved
18	0b R/W	**Data Poisoning Enable**. This bit controls whether or not the MCH marks data as "poisoned" when a parity error is detected on incoming data from the DRAM I/F. 0 = Errors will not be propagated, only good internal parity generated. 1 = Error Poisoning Enabled. Incoming data with parity errors will be marked as "poisoned" before being sent on towards its destination when in either 72-bit or x4 SDDC ECC mode via the DRC register.
17:0	0	Reserved

3.6.74 DRAM_ERR_CTL – DRAM Error Control Register (D0:F1)

Address Offset: EC – EFh
Access: R/W
Size: 32 Bits
Default: 0000_0000h

This register controls the MCH handling of errors on incoming data streams into the MCH core from the DRAM interface.

Bit Field	Default & Access	Description
31:19	0000h	Reserved
18	0b R/W	**Data Poisoning Enable**. This bit controls whether or not the MCH marks data as "poisoned" when a parity error is detected on incoming data from the DRAM I/F. 0 = Errors will not be propagated, only good internal parity generated. 1 = Error Poisoning Enabled. Incoming data with parity errors will be marked as "poisoned" before being sent on towards its destination when in either 72-bit or x4 SDDC ECC mode via the DRC register.
17:0	0	Reserved

3.7 PCI Express* Port A Registers (D2:F0)

Device 2 is the PCI Express* port A (in x8 mode) or port A0 (in x4 mode) virtual PCI-to-PCI bridge. The registers described here include both the standard configuration space as well as the enhanced configuration space (starting at offset 100h).

Table 3-6. PCI Express* Port A PCI Configuration Register Map (D2:F0) (Sheet 1 of 3)

Address Offset	Mnemonic	Register Name	Access	Default
00 – 01h	VID	Vendor Identification	RO	8086h
02 – 03h	DID	Device Identification	RO	3595h
04 – 05h	PCICMD	PCI Command Register	RO, R/W	0000h
06 – 07h	PCISTS	PCI Status Register	RO, R/WC	0010h
08h	RID	Revision Identification	RO	09h
0Ah	SUBC	Sub-Class Code	RO	04h
0Bh	BCC	Base Class Code	RO	06h
0Ch	CLS	Cache Line Size	R/W	00h
0Eh	HDR	Header Type	RO	01h
18h	PBUSN	Primary Bus Number	RO	00h
19h	SBUSN	Secondary Bus Number	R/W	00h
1Ah	SUBUSN	Subordinate Bus Number	R/W	00h
1Ch	IOBASE	I/O Base Address Register	R/W, RO	F0h
1Dh	IOLIMIT	I/O Limit Address Register	R/W	00h
1E – 1Fh	SEC_STS	Secondary Status Register	R/WC, RO	0000h
20 – 21h	MBASE	Memory Base Address Register	R/W	FFF0h
22 – 23h	MLIMIT	Memory Limit Address Register	R/W	0000h
24 – 25h	PMBASE	Prefetchable Memory Base Address Reg.	R/W, RO	FFF1h
26 – 27h	PMLIMIT	Prefetchable Memory Limit Address Reg.	R/W, RO	0001h
28h	PMBASU	Prefetchable Mem Base Upper Addr. Reg.	RO, R/W	0Fh
2Ch	PMLMTU	Prefetchable Memory Limit Upper Address Register	RO, R/W	00h

Table 3-6. PCI Express* Port A PCI Configuration Register Map (D2:F0) (Sheet 2 of 3)

Address Offset	Mnemonic	Register Name	Access	Default
34h	CAPPTR	Capabilities Pointer	RO	50h
3Ch	INTRLINE	Interrupt Line Register	R/W	00h
3Dh	INTRPIN	Interrupt Pin Register	RWO	01h
3Eh	BCTRL	Bridge Control Register	RO, R/W	00h
44h	VS_CMD0	Vendor-Specific Command Register 0	RO	00h
45h	VS_CMD1	Vendor-Specific Command Register 1	RO, R/W, R/WS	00h
46h	VS_STS0	Vendor-Specific Status Register 0	RO	00h
47h	VS_STS1	Vendor-Specific Status Register 1	RO, R/WC	00h
50h	PMCAPID	Power Management Capabilities Structure	RO	01h
51h	PMNPTR	Power Management Next Capabilities Pointer	RO	58h
52 – 53h	PMCAPA	Power Management Capabilities	RO	C822h
54 – 55h	PMCSR	Power Management Status and Control	RO,R/W	0000h
56h	PMCSRBSE	Power Management Status and Control Bridge Extensions	RO	00h
57h	PMDATA	Power Management Data	RO	00h
58h	MSICAPID	MSI Capabilities Structure	RO	05h
59h	MSINPTR	MSI Next Capabilities Pointer	RO	64h
5A – 5Bh	MSICAPA	MSI Capabilities	RO, R/W	0002h
5C – 5Fh	MSIAR	MSI Address Register for PCI Express*	R/W	FEE0_0000h
60h	MSIDR	MSI Data Register	R/W	0000h
64h	EXP_CAPID	PCI Express Features Capabilities Structure	RO	10h
65h	EXP_NPTR	PCI Express Next Capabilities Pointer	RO	00h
66 – 67h	EXP_CAPA	PCI Express Features Capabilities	RO, R/WO	0041h
68 – 6Bh	EXP_DEVCAP	PCI Express Device Capabilities	RO	0002_8001h
6C – 6Dh	EXP_DEVCTL	PCI Express Device Control	R/W. RO	0000h
6E – 6Fh	EXP_DEVSTS	PCI Express Device Status	RO, R/WC	0000h
70 – 73h	EXP_LNKCAP	PCI Express Link Capabilities	RO	0203_E481h
74 – 75h	EXP_LNKCTL	PCI Express Link Control	RO, R/W, WO	0000h
76 – 77h	EXP_LNKSTS	PCI Express Link Status	RO	1001h
78 – 7Bh	EXP_SLTCAP	PCI Express Slot Capabilities	RO, R/WO	0000_0000h
7C – 7Dh	EXP_SLTCTL	PCI Express Slot Control	R/W	03C0h
7E – 7Fh	EXP_SLTSTS	PCI Express Slot Status	RO, R/WC	0040h
80 – 83h	EXP_RPCTL	PCI Express Root Port Control	R/W	0000_0000h
84 – 87h	EXP_RPSTS	PCI Express Root Port Status	RO, R/WC	0000_0000h
C4 – C7h	EXP_PFCCA	PCI Express Posted Flow Control Credits Allocated	RO	000C_0030h
C8 – CBh	EXP_NPFCCA	PCI Express Non Posted Flow Control Credits Allocated	RO	0008_0001h
100 – 103h	EXP_ENHCAPST	PCI Express Enhanced Capability Structure	RO	0001_0001h
104 – 107h	EXP_UNCERRSTS	PCI Express Uncorrectable Error Status	RO, R/WC	0000_0000h

Table 3-6. PCI Express* Port A PCI Configuration Register Map (D2:F0) (Sheet 3 of 3)

Address Offset	Mnemonic	Register Name	Access	Default
108 – 10Bh	EXP_UNCERRMSK	PCI Express Uncorrectable Error Mask	RO, R/W	0000_0000h
10C – 10Fh	EXP_UNCERRSEV	PCI Express Uncorrectable Error Severity	RO, R/W	0006_0011h
110 – 113h	EXP_CORERRSTS	PCI Express Correctable Error Status	R/WC	0000_0000h
114 – 117h	EXP_CORERRMSK	PCI Express Correctable Error Mask	R/W	0000_0000h
118 – 11Bh	EXP_AERCACR	PCI Express Advanced Error Capabilities and Control	RO, R/W	0000_0000h
11C – 11Fh	EXP_HDRLOG0	PCI Express Header Log DW0	RO	0000_0000h
120 – 123h	EXP_HDRLOG1	PCI Express Header Log DW1	RO	0000_0000h
124 – 127h	EXP_HDRLOG2	PCI Express Header Log DW2	RO	0000_0000h
128 – 12Bh	EXP_HDRLOG3	PCI Express Header Log DW3	RO	0000_0000h
12C – 12Fh	EXP_RPERRCMD	PCI Express Root Port Error Command	R/W	0000_0000h
130 – 133h	EXP_RPERRMSTS	PCI Express Root Port Error Message Status	RO, R/WC	0000_0000h
134 – 137h	EXP_ERRSID	PCI Express Error Source ID	RO	0000_0000h
140 – 143h	EXP_UNITERR	PCI Express Unit Error Status	RO, R/WC	0000_0000h
144 – 147h	EXP_MASKERR	PCI Express Mask Error	RO, R/W	0000_E000h
148 – 14Bh	EXP_ERRDOCMD	PCI Express Error Do Command Register	RO, R/W	0000_0000h
14C – 14Fh	EXP_UNCERRDMSK	PCI Express Uncorrectable Error Detect Mask	RO, R/W	0000_0000h
150 – 153h	EXP_CORERRDMSK	PCI Express Correctable Error Detect Mask	R/W	0000_0000h
158 – 15Bh	EXP_UNITERRDMSK	PCI Express Unit Error Detect Mask	R/W	0000_0000h
160 – 163h	EXP_FERR	PCI Express First Error	R/WC	0000_0000h
164 – 167h	EXP_NERR	PCI Express Next Error	R/WC	0000_0000h
168 – 16Bh	EXP_ERR_CTL	PCI Express Error Control	R/W	0000_0000h

3.7.1 VID – Vendor Identification (D2:F0)

Address Offset: 00 – 01h
Access: RO
Size: 16 Bits
Default: 8086h

The VID register contains the vendor identification number. This 16-bit register combined with the Device Identification register uniquely identify any PCI device.

Bit Field	Default & Access	Description
15:0	8086h RO	**Vendor Identification Device (VID)**. This is a 16-bit value assigned to Intel.

3.7.2 DID – Device Identification (D2:F0)

Address Offset: 02 – 03h
Access: RO
Size: 16 Bits
Default: 3595h

Bit Field	Default & Access	Description
15:0	3595h RO	**Device Identification Number (DID)**. This is a 16 bit value assigned to the MCH Device 2, Function 0.

3.7.3 PCICMD – PCI Command Register (D2:F0)

Address Offset: 04 – 05h
Access: RO, R/W
Size: 16 Bits
Default: 0000h

Many of these bits are not applicable since the primary side of this device is not an actual PCI bus.

Bit Field	Default & Access	Description
15:11	00h	Reserved
10	0b R/W	**INTx Disable**. Controls the ability of the PCI Express* device to assert INTx interrupts. When set, devices are prevented from asserting INTx. This bit only applies to legacy interrupts, and not MSIs. Also, this bit has no effect on PCI Express* messages that are converted to legacy interrupts... only internal, device generated interrupts. 0 = Enable INTx assertion 1 = Disable INTx assertion
9	0b RO	**Fast Back-to-Back Enable (FB2B)**. Not Applicable.
8	0b R/W	**SERR Enable (SERRE)**. This bit is a global enable bit for Device SERR messaging. The MCH does not have an SERR# signal. The MCH communicates the SERR# condition by sending an SERR message to the ICH via the Hub I/F. 0 = No SERR message is generated by the MCH for device (unless enabled through enhanced configuration registers). 1 = Enable SERR, SCI, or SMI messages over HI or asserting MCERR# for specific device error conditions.
7	0b RO	**Address/Data Stepping (ADSTEP)**. Not applicable.
6	0b R/W	**Parity Error Enable (PERRE)**. This bit determines the device behavior on detection of a parity error. See the *PCI Express Interface Specification, Rev 1.0a*, for details. 0 = Parity Errors are logged in the status register, but no other action is taken. 1 = Normal action is taken upon detection of Parity Error, as well as logging.
5	0b RO	**VGA Palette Snoop**. Not Applicable.
4	0b RO	**Memory Write and Invalidate**. Not applicable.

Bit Field	Default & Access	Description
3	0b RO	**Special Cycle Enable.** Not applicable.
2	0b R/W	**Bus Master Enable (BME).** This bit controls the PCI Express* port's ability to issue memory and I/O read/write requests on behalf of subordinate devices. Note that MSI interrupt messages are in-band memory writes, and clearing this bit disables MSI interrupt messages. 0 = Disable. The port will not respond to any I/O or memory transaction originating on the secondary interface. 1 = Enable.
1	0b R/W	**Memory Access Enable (MAE)**. Controls access to the Memory and Prefetchable memory address ranges defined in the MBASE2, MLIMIT2, PMBASE2, and PMLIMIT2 registers. 0 = Disable all of device memory space. 1 = Enable
0	0b R/W	**I/O Access Enable (IOAE)**. Controls access to the I/O address range defined in the IOBASE2 and IOLIMIT2 registers. 0 = Disable device I/O space. 1 = Enable

3.7.4 PCISTS – PCI Status Register (D2:F0)

Address Offset: 06 – 07h
Access: RO, R/WC
Size: 16 Bits
Default: 0010h

PCISTS2 is a 16-bit status register that reports the occurrence of error conditions associated with the primary side of the "virtual" PCI-to-PCI bridge embedded within the MCH.

Bit Field	Default & Access	Description
15	0b RO	**Detected Parity Error (DPE)**. Parity is supported on the primary side of this device. This bit is set by the MCH's PCI Express* Port logic when it receives a poisoned TLP, regardless of the state of the Parity Error Enable bit. Since the parity is not checked on the downstream side from the core, this bit can never be set. 0 = No parity Error detected
14	0b R/WC	**Signaled System Error (SSE)**. Indicates whether or not a Hub I/F SERR message was generated by this device. For the root port the fatal and non-fatal messages can be either received or virtual messages that are forwarded for reporting. Software clears this bit by writing a '1' to the bit location. 0 = SERR message not generated by this device. 1 = This device was the source of fatal or non-fatal error that has been enabled for generation of a System Error.
13	0b R/WC	**Received Master Abort Status (RMAS)**. Indicates whether or not this PCI Express* device received a completion with Unsupported Request Completion status. 0 = No Master Abort received. Software clears this bit by writing a '1' to the bit location. 1 = Set when this PCI Express* device receives a completion with Unsupported Request Completion Status.

Bit Field	Default & Access	Description
12	0b R/WC	**Received Target Abort Status (RTAS)**. Indicates whether or not this PCI Express* device received a completion with Completer Abort Completion Status. 0 = No Target Abort received. Software clears this bit by writing a '1' to the bit location. 1 = Set when this PCI Express* device receives a completion with Completer Abort Completion Status
11	0b RO	**Signaled Target Abort Status (STAS)**. Indicates whether or not this PCI Express* device completed a request using Completer Abort Completion Status. Not applicable to the primary side.
10:9	00b RO	**DEVSEL# Timing (DEVT)**. Not Applicable.
8	0b R/WC	**Master Data Parity Error Detected (DPD)**. Parity is supported on the primary side of this device. 0 = No Master Parity Error detected. Software clears this bit by writing a '1' to the bit location. 1 = Set when this PCI Express* device receives a completion marked poisoned, or when this device poisons a write Request. This bit can only be set if the Parity Error Enable bit is set.
7	0b RO	**Fast Back-to-Back (FB2B)**. Not Applicable.
6	0b	Reserved
5	0b RO	**66 MHz Capable**. Not Applicable.
4	1b RO	**Capabilities List**. Hardwired to 1 to indicate the presence of an Extended Capability List item
3	0b RO	**INTx Status**. Indicates that an interrupt is pending internal to this device. This bit does not get set for interrupts forwarded up from downstream devices, or for messages converted to interrupts by the root port. The INTx Assertion Disable bit has no effect on the setting of this bit. This bit is not set for an MSI.
2:0	00h	Reserved

3.7.5 RID – Revision Identification (D2:F0)

Address Offset: 08h
Access: RO
Size: 8 Bits
Default: 09h

This register contains the revision number of the device.

Bit Field	Default & Access	Description
7:0	00h RO	**Revision Identification Number (RID)**. This value indicates the revision identification number for the device. It is always the same as the value in Device 0 RID. 09h = C1 stepping. 0Ah = C2 stepping.

3.7.6 SUBC – Sub-Class Code (D2:F0)

Address Offset: 0Ah
Access: RO
Size: 8 Bits
Default: 04h

This register contains the Sub-Class Code for the device.

Bit Field	Default & Access	Description
7:0	04h RO	**Sub-Class Code (SUBC)**. This value indicates the category of Bridge into which device falls. 04h = PCI to PCI Bridge.

3.7.7 BCC – Base Class Code (D2:F0)

Address Offset: 0Bh
Access: RO
Size: 8 Bits
Default: 06h

This register contains the Base Class Code of the device.

Bit Field	Default & Access	Description
7:0	06h RO	**Base Class Code (BASEC)**. This value indicates the Base Class Code for the device. 06h = Bridge device

3.7.8 CLS – Cache Line Size (D2:F0)

Address Offset: 0Ch
Access: RW
Size: 8 Bits
Default: 00h

This register is normally set by system firmware and OS to the system cache line size. Legacy PCI 2.3 software may not always be able to program this field correctly. It is implemented as a read-write field for legacy compatibility purposes, but has no effect on this device's functionality.

Bit Field	Default & Access	Description
7:0	00h R/W	**Cache Line Size (CLS)** – This register is set by BIOS or OS to the system cache line size. Implemented as read-write field only for compatibility reasons. It has no effect on the device's functionality.

3.7.9 HDR – Header Type (D2:F0)

Address Offset: 0Eh
Access: RO
Size: 8 Bits
Default: 01h

This register identifies the header layout of the configuration space.

Bit Field	Default & Access	Description
7:0	01h RO	**Header Type Register (HDR)**. This value indicates the Header Type of the device. 01h = single-function device with Bridge layout.

3.7.10 PBUSN – Primary Bus Number (D2:F0)

Address Offset: 18h
Access: RO
Size: 8 Bits
Default: 00h

This register identifies that "virtual" PCI-to-PCI bridge is connected to bus #0.

Bit Field	Default & Access	Description
7:0	00h RO	**Primary Bus Number (BUSN)**. Configuration software typically programs this field with the number of the bus on the primary side of the bridge. Since device is an internal device and its primary bus is always 0, these bits are hardwired to '0'.

3.7.11 SBUSN – Secondary Bus Number (D2:F0)

Address Offset: 19h
Access: R/W
Size: 8 Bits
Default: 00h

This register identifies the bus number assigned to the second bus side of the "virtual" PCI-to-PCI bridge (the PCI Express* connection). This number is programmed by the PCI configuration software to allow mapping of configuration cycles to a subordinate device connected to PCI Express*.

Bit Field	Default & Access	Description
7:0	00h R/W	**Secondary Bus Number (BUSN)**. This field is programmed by configuration software with the bus number of the PCI Express* port.

3.7.12 SUBUSN – Subordinate Bus Number (D2:F0)

Address Offset: 1Ah
Access: R/W
Size: 8 Bits
Default: 00h

This register is programmed by PCI configuration software to the highest numbered subordinate bus (if any) that resides below another bridge device below the secondary PCI Express* interface.

Bit Field	Default & Access	Description
7:0	00h R/W	**Subordinate Bus Number (BUSN)**. This register is programmed by configuration software with the number of the highest subordinate bus that lies behind the device bridge.

3.7.13 IOBASE – I/O Base Address Register (D2:F0)

Address Offset: 1Ch
Access: RO, R/W
Size: 8 Bits
Default: F0h

The IOBASE and IOLIMIT registers control the processor-to-PCI Express* I/O access routing based on the following formula:

$$IOBASE \leq Address \leq IOLIMIT$$

Only the upper four bits are programmable. For the purpose of address decode address bits A[11:0] are treated as 0. Thus the bottom of the defined I/O address range will be aligned to a 4-KB boundary.

Bit Field	Default & Access	Description
7:4	Fh R/W	**I/O Address Base (IOBASE)**. Corresponds to A[15:12] of the I/O addresses passed by the device bridge to PCI Express*
3:0	0h RO	**I/O Addressing Capability**. Only 16-bit I/O addressing is supported, so these bits are hardwired to '0'.

3.7.14 IOLIMIT – I/O Limit Address Register (D2:F0)

Address Offset: 1Dh
Access: R/W
Size: 8 Bits
Default: 00h

This register controls the processor to PCI Express* I/O access routing based on the following formula:

$$IOBASE \leq Address \leq IOLIMIT$$

Only the upper four bits are programmable. For the purpose of address decode address bits A[11:0] are assumed to be FFFh. Thus, the top of the defined I/O address range will be at the top of a 4KB aligned address block.

Bit Field	Default & Access	Description
7:4	0h R/W	**I/O Address Limit (IOLIMIT)**. Corresponds to A[15:12] of the I/O address limit of device. Devices between this upper limit and IOBASE2 will be passed to PCI Express*.
3:0	0h RO	**I/O Addressing Capability**. Only 16-bit I/O addressing is supported, so these bits are hardwired to '0'.

3.7.15 SEC_STS – Secondary Status Register (D2:F0)

Address Offset: 1E – 1Fh
Access: R/WC, RO
Size: 16 Bits
Default: 0000h

SEC_STS is a 16-bit status register that reports the occurrence of error conditions associated with secondary side (i.e., PCI Express* side) of the "virtual" PCI-to-PCI bridge embedded within MCH.

Bit Field	Default & Access	Description
15	0b R/WC	**Detected Parity Error (2DPE)**. This bit is set by the MCH's PCI Express* Port logic when the secondary side receives a poisoned TLP, regardless of the state of the Parity Error Enable bit. Software clears this bit by writing a '1' to the bit location. See the *PCI Express Interface Specification, Rev 1.0a* for details. 0 = No parity Error detected. 1 = Parity Error Detected (poisoned TLP received).
14	0b R/WC	**Received System Error (2RSE)**. Indicates whether or not an ERR_FATAL or ERR_NONFATAL message was received via PCI Express*. The SERR Enable bit in the Bridge Control Register does not gate the setting of this bit. This bit is not set for virtual messages. 0 = Error message not received by this device. 1 = This device received fatal or non-fatal error message via PCI Express*.
13	0b R/WC	**Received Master Abort Status (2RMAS)**. Indicates whether or not this PCI Express* device received a completion with Unsupported Request Completion status. 0 = No Master Abort received. Software clears this bit by writing a '1' to the bit location. 1 = This PCI Express* device received a completion with Unsupported Request Completion Status.
12	0b R/WC	**Received Target Abort Status (2RTAS)**. Indicates whether or not this PCI Express* device received a completion with Completer Abort Completion Status. 0 = No Target Abort received. Software clears this bit by writing a '1' to the bit location. 1 = This PCI Express* device received a completion with Completer Abort Completion Status

Bit Field	Default & Access	Description
11	0b R/WC	**Signaled Target Abort Status (2STAS)**. Indicates whether or not this PCI Express* device completed a request using Completer Abort Completion Status. 0 = No Target Abort signaled. Software clears this bit by writing a '1' to the bit location. 1 = This PCI Express* device completed a request using Completer Abort Completion Status
10:9	00b RO	**DEVSEL# Timing (DEVT)**. Not Applicable.
8	0b R/WC	**Master Data Parity Error Detected (DPD)**. Parity is supported on the secondary side of this device. 0 = No Master Parity Error detected. Software clears this bit by writing a '1' to the bit location. 1 = This PCI Express* device received a completion marked poisoned, or when this device poisoned a write Request. This bit can only be set if the Parity Error Enable bit is set.
7	0b RO	**Fast Back-to-Back (FB2B)**. Not Applicable.
6	0b	Reserved
5	0b RO	**66 MHz Capable**. Not Applicable.
4:0	00h	Reserved

3.7.16 MBASE – Memory Base Address Register (D2:F0)

Address Offset: 20 – 21h
Access: R/W
Size: 16 Bits
Default: FFF0h

The MBASE and MLIMIT registers control the processor to PCI Express* non-prefetchable memory access routing based on the following formula:

$$MBASE \leq Address \leq MLIMIT$$

The upper 12 bits of both registers are read/write and correspond to the upper 12 address bits A[31:20] of the 32-bit address. The bottom four bits of these registers are read-only and return zeroes when read. These registers must be initialized by configuration software. For the purpose of address decode address bits A[19:0] of the Memory Base Address are assumed to be 0. Similarly, the bridge assumes that the lower 20 bits of the Memory Limit Address (A[19:0]) are F_FFFFh. Thus, the bottom of the defined memory address range will be aligned to a 1-Mbyte boundary, and the top of the defined memory range will be at the top of a 1-MB memory block.

Note: Memory range covered by MBASE and MLIMIT registers are used to map non-prefetchable PCI Express* address ranges (typically where control/status memory-mapped I/O data structures of the graphics controller will reside) and PMBASE and PMLIMIT are used to map prefetchable address ranges (typically graphics local memory). This segregation allows for the application of the write combine space attribute to be performed in a true plug-and-play manner to the prefetchable address range for improved PCI Express* memory access performance.

Note also that configuration software is responsible for programming all address range registers (prefetchable, non-prefetchable) with the values that provide exclusive address ranges (i.e. prevent overlap with each other and/or with the ranges covered with the main memory). There is no provision in the MCH hardware to enforce prevention of overlap and operations of the system in the case of overlap are not guaranteed.

Bit Field	Default & Access	Description
15:4	FFFh R/W	**Memory Address Base (MBASE)**. Corresponds to A[31:20] of the lower limit of the memory range that will be passed by the device bridge to PCI Express*.
3:0	0h	Reserved

3.7.17 MLIMIT – Memory Limit Address Register (D2:F0)

Address Offset: 22 – 23h
Access: R/W
Size: 16 Bits
Default: 0000h

Bit Field	Default & Access	Description
15:4	000h R/W	**Memory Address Limit (MLIMIT).** Corresponds to A[31:20] of the memory address that corresponds to the upper limit of the range of memory accesses that will be passed by the device bridge to PCI Express*.
3:0	0h	Reserved

3.7.18 PMBASE – Prefetchable Memory Base Address Register (D2:F0)

Address Offset: 24 – 25h
Access: RO, R/W
Size: 16 Bits
Default: FFF1h

The PMBASE and PMLIMIT registers control the processor-to-PCI Express* prefetchable memory accesses. The upper 12 bits of both registers are read/write and correspond to the upper 12 address bits A[31:20] of the 32-bit address. For the purpose of address decode, address bits A[19:0] of the Prefetchable Memory Base Address are assumed to be 0. Similarly, the bridge assumes that the lower 20 bits of the Prefetchable Memory Limit Address (A[19:0]) are F_FFFFh. Thus, the bottom of the defined memory address range will be aligned to a 1-Mbyte boundary, and the top of the defined memory range will be at the top of a 1-MB memory block.

Bit Field	Default & Access	Description
15:4	FFFh R/W	**Prefetchable Memory Address Base (PMBASE)**. Corresponds to A[31:20] of the lower limit of the address range passed by bridge device across PCI Express*.

Bit Field	Default & Access	Description
3:1	000b RO	**Memory addressing mode**.
0	1b RO	**Memory Base Upper Address Enabled**. This bit when set indicates that the base address is further defined by the upper address bits of the memory base upper address register.

3.7.19 PMLIMIT – Prefetchable Memory Limit Address Register (D2:F0)

Address Offset: 26 – 27h
Access: RO, R/W
Size: 16 Bits
Default: 0001h

This register controls the processor to PCI Express* prefetchable memory accesses. The upper 12 bits of the register are read/write and correspond to the upper 12 address bits A[31:20] of the 32 bit address. For the purpose of address decode, bits A[19:0] are assumed to be F_FFFFh. Thus, the top of the defined memory address range will be at the top of a 1MB aligned memory block.

Bit Field	Default & Access	Description
15:4	000h R/W	**Prefetchable Memory Address Limit (PMLIMIT)**. Corresponds to A[31:20] of the upper limit of the address range passed by bridge device across PCI Express*.
3:1	000b RO	**Memory addressing mode**.
0	1b RO	**Memory Limit Upper Address Enabled**. This bit when set indicates that the limit address is further defined by the upper address bits of the memory limit upper address register.

3.7.20 PMBASU – Prefetchable Memory Base Upper Address Register (D2:F0)

Address Offset: 28h
Access: RO, R/W
Size: 8 Bits
Default: 0Fh

These register expands the prefetchable memory base address by four bits. All other bits are reserved.

Bit Field	Default & Access	Description
7:4	0h	Reserved
3:0	Fh R/W	**Base Upper Address bits.** These four bits expand the prefetchable address base to 36 bits. Corresponds to A[35:32] of the lower limit of the address range passed by bridge device across PCI Express* interface.

3.7.21 PMLMTU – Prefetchable Memory Limit Upper Address Register (D2:F0)

Address Offset: 2Ch
Access: RO, R/W
Size: 8 Bits
Default: 00h

Bit Field	Default & Access	Description
7:4	0h	Reserved
3:0	0h R/W	**Limit Upper Address bits.** These four bits expand the prefetchable address limit to 36 bits. Corresponds to A[35:32] of the upper limit of the address range passed by bridge device across PCI Express* interface.

3.7.22 CAPPTR – Capabilities Pointer (D2:F0)

Address Offset: 34h
Access: RO
Size: 8 Bits
Default: 50h

The CAPPTR provides the offset that is the pointer to the location where the first set of capabilities registers is located.

Bit Field	Default & Access	Description
7:0	50h RO	**Capabilities Pointer (CAP_PTR)**. Pointer to first PCI Express* Capabilities Structure register block, which is the first of the chain of capabilities.

3.7.23 INTRLINE – Interrupt Line Register (D2:F0)

Address Offset: 3Ch
Access: R/W
Size: 8 Bits
Default: 00h

Bit Field	Default & Access	Description
7:0	00h R/W	**Interrupt Connection** – BIOS writes the interrupt routing information to this register to indicate which input of the interrupt controller this device is connected to.

3.7.24 INTRPIN – Interrupt Pin Register (D2:F0)

Address Offset: 3Dh
Access: R/WO
Size: 8 Bits
Default: 02h

Bit Field	Default & Access	Description
7:0	01h R/WO	**Interrupt Pin** – Designates the interrupt pin mapping for this device. 00h = Reserved 01h = INTA# 02h = INTB# 03h = INTC# 04h = INTD# 05 – FFh = Reserved

3.7.25 BCTRL – Bridge Control Register (D2:F0)

Address Offset: 3Eh
Access: RO, R/W
Size: 8 Bits
Default: 00h

This register provides extensions to the PCICMD register that are specific to PCI-to-PCI bridges. The BCTRL provides additional control for the secondary interface (i.e., PCI Express*) as well as some bits that affect the overall behavior of the "virtual" PCI-to-PCI bridge embedded within MCH (e.g., VGA compatible address range mapping).

Bit Field	Default & Access	Description
7	0b RO	**Fast Back-to-Back Enable (FB2BEN)**. Not Applicable.
6	0b R/W	**Secondary Bus Reset (SRESET)**. Setting this bit triggers a hot reset on the link for the corresponding PCI Express* port and the PCI Express* hierarchy domain subordinate to the port. This sends the LTSSM into the Training (or Link) Control Reset state, which necessarily implies a reset to the downstream device and all subordinate devices. Once this bit has been cleared, and the minimum transmission requirement has been met, the detect state will be entered by both ends of the link. Note also that a secondary bus reset will not reset the primary side configuration registers of the targeted PCI Express* port. This is necessary to allow software to specify special training configuration, such as entry into rollback mode.
5	0b RO	**Master Abort Mode (MAMODE)**. Not Applicable.
4	0b	Reserved
3	0b R/W	**VGA Enable (VGAEN)**. Controls the routing of processor-initiated transactions targeting VGA compatible I/O and memory address ranges. **NOTE:** Only one of Device 2-4's VGAEN bits are allowed to be set. This must be enforced via software.
2	0b R/W	**ISA Enable (ISAEN)**. Modifies the response by the MCH to an I/O access issued by the processor that target ISA I/O addresses. This applies only to I/O addresses that are enabled by the IOBASE and IOLIMIT registers. 0 = All addresses defined by the IOBASE and IOLIMIT for processor I/O transactions will be mapped to PCI Express*. 1 = MCH will not forward to PCI Express* any I/O transactions addressing the last 768 bytes in each 1-KB block even if the addresses are within the range defined by the IOBASE and IOLIMIT registers. Instead, these cycles will be forwarded to HI.

Bit Field	Default & Access	Description
1	0b R/W	**SERR Enable (2SERRE)**. This bit enables or disables forwarding of SERR messages from PCI Express* to the HI, where they can be converted into interrupts that are eventually delivered to the processor.
0	0b R/W	**Parity Error Response Enable (2PERRE)**. Controls MCH's response to poisoned TLPs on PCI Express*. Other types of error conditions can still be signaled via SERR messaging independent of this bit's state. 0 = Poisoned TLPs are not reported via the MCH HI_A SERR messaging mechanism. 1 = Poisoned TLPs are reported via the Hub Interface SERR messaging mechanism, if further enabled by SERRE2.

3.7.26 VS_CMD0 – Vendor Specific Command Register 0 (D2:F0)

Address Offset: 44h
Access: RO
Size: 8 Bits
Default: 00h

This register is for vendor specific Hot-Plug commands. As Hot-Plug is not implemented on this device, the contents of this register are reserved.

Bit Field	Default & Access	Description
7:0	00h	Reserved

3.7.27 VS_CMD1 – Vendor Specific Command Register 1 (D2:F0)

Address Offset: 45h
Access: RO, R/WC
Size: 8 Bits
Default: 00h

This register is for vendor specific commands.

Bit Field	Default & Access	Description
7:4	0h	Reserved
3	0b R/W	**Completion TO Timer Disable**. 0 = Enabled. 1 = Disabled.
2	0b	Reserved
1	0b RW	**Training Control Loopback Enable.** 0 = Disabled. 1 = When the TS1/TS2 ordered-sets are transmitted, the "Enable Loopback" bit will be set in the training control symbol.
0	0b R/WS	**PMETOR**. PME Turn Off Request, set by software, cleared by hardware when the acknowledge is returned from the link. The bit will also be cleared when the link layer is in the DL_down state.

3.7.28 VS_STS0 – Vendor Specific Status Register 0 (D2:F0)

Address Offset: 46h
Access: RO
Size: 8 Bits
Default: 00h

This register is for vendor specific Hot-Plug status. As Hot-Plug is not implemented on this device, the contents of this register are reserved.

Bit Field	Default & Access	Description
7:0	00h	Reserved

3.7.29 VS_STS1 – Vendor Specific Status Register 1 (D2:F0)

Address Offset: 47h
Access: RO, R/WC
Size: 8 Bits
Default: 00h

This register is for vendor specific status

Bit Field	Default & Access	Description
7:2	0b	Reserved
1	0b RO	**Link Active.** Bit will report whether transactions are being sent or aborted by the downstream transaction control, which is determined by the “link_active” status from the link layer reflected in this status bit. 1 = link up. 0 = link down.
0	0b R/WC	**PMETOA.** PME Turn Off Acknowledge is set by hardware when PMETOR is ON and the acknowledge is returned from the link. When this bit is set, the Turn Off request bit is cleared. Software writes a ‘1’ to this bit to clear it. The bit will also be cleared when the link layer is in the DL_down state.

3.7.30 PMCAPID – Power Management Capabilities Structure (D2:F0)

Address Offset: 50h
Access: RO
Size: 8 Bits
Default: 01h

This register identifies the capability structure and points to the next structure.

Bit Field	Default & Access	Description
7:0	01h RO	**CAP_ID**. This field has the value 01h to identify the CAP_ID assigned by the PCI SIG for vendor dependent capability pointers.

3.7.31 PMNPTR – Power Management Next Capabilities Pointer (D2:F0)

Address Offset: 51h
Access: RO
Size: 8 Bits
Default: 58h

This register identifies the capability structure and points to the next structure.

Bit Field	Default & Access	Description
7:0	58h RO	**Next Capability Pointer**. This field points to the next Capability ID in this device, which is the MSI

3.7.32 PMCAPA – Power Management Capabilities (D2:F0)

Address Offset: 52 – 53h
Access: RO
Size: 16 Bits
Default: C822h

This register identifies the capabilities for PM.

Bit Field	Default & Access	Description
15:11	11001b RO	**PME Support**. Identifies power states which assert PME. Bits 15, 14 and 11 must be set to '1' for PCI-PCI bridge structures representing ports on root complexes. The definition of these bits is taken from the *PCI Bus Power Management Interface Specification, Rev 1.1*. bit(11) XXXX1b - PME# can be asserted from D0 bit(12) XXX1Xb - PME# can be asserted from D1 — (not supported) bit(13) XX1XXb - PME# can be asserted from D2 — (not supported) bit(14) X1XXXb - PME# can be asserted from D3 hot bit(15) 1XXXXb - PME# can be asserted from D3 cold
10	0b RO	**D2 Support**. Hardwired to '0' to indicate that this function does not support the D2 state.
9	0b RO	**D1 Support**. Hardwired to '0' to indicate that this function does not support the D1 state.
8:6	000b RO	**AUX Current**. Hardwired to 000b to indicate a self-powered device.
5	1b RO	**DSI – Device Specific Initialization**. Device-specific initialization is required.
4	0b RO	Reserved
3	0b RO	**PME Clock**. No PCI clock is required for PME.
2:0	010b RO	**Version**. Hardwired to 010b to indicate compliance with *PCI Bus Power Management Interface Specification, Rev 1.1*.

3.7.33 PMCSR – Power Management Status and Control (D2:F0)

Address Offset: 54 – 55h
Access: RO, R/W
Size: 16 Bits
Default: 0000h

This register identifies the capabilities for PM.

Bit Field	Default & Access	Description
15	0b RO	**PME Status**. Indicates/clears PME# assertion. Only for generated PME, not forwarded PME. Field not supported by MCH. This bit is sticky.
14:13	00b RO	**Data Scale**. Not Applicable.
12:9	0h RO	**Data Select**. Not Applicable.
8	0b R/W	**PME Enable**. Controls PME# assertion. This bit is sticky through reset. Writes to this field have no effect. This bit is sticky. 0 = This device will not assert PME# 1 = Enables this device to assert PME#
7:2	00h RO	Reserved
1:0	00b RO	**Power State**. Since the PCI Express* bridge device supports only the D0 state, writes to this field have no effect.

3.7.34 PMCSRBSE – Power Management Status and Control Bridge Extensions (D2:F0)

Address Offset: 56h
Access: RO
Size: 8 Bits
Default: 00h

This register identifies the capabilities for PM.

Bit Field	Default & Access	Description
7	0b RO	**Bus Power/Clock Control Enable**. Not applicable.
6	0b RO	**B2/B3 Support**. Not applicable.
5:0	00h	Reserved

3.7.35 PMDATA – Power Management Data (D2:F0)

Address Offset: 56h
Access: RO
Size: 8 Bits
Default: 00h

This register identifies the data read based on the data select. This register is not supported in the MCH and is reserved.

Bit Field	Default & Access	Description
7:0	00h	Reserved

3.7.36 MSICAPID – MSI Capabilities Structure (D2:F0)

Address Offset: 58h
Access: RO
Size: 8 Bits
Default: 05h

This register identifies the MSI capability structure.

Bit Field	Default & Access	Description
7:0	05h RO	**CAP_ID**. This field has the value 05h to identify the CAP_ID assigned by the PCI SIG for a Message Signaled Interrupts capability list.

3.7.37 MSINPTR – MSI Next Capabilities Pointer (D2:F0)

Address Offset: 59h
Access: RO
Size: 8 Bits
Default: 64h

This register points to the next structure.

Bit Field	Default & Access	Description
7:0	64h RO	**Next Capability Pointer**. This field points to the next Capability ID in this device.

3.7.38 MSICAPA – MSI Capabilities (D2:F0)

Address Offset: 5A – 5Bh
Access: RO, R/W
Size: 16 Bits
Default: 0002h

The PCI Express* controller generates upstream interrupt message using MSI to the processor, bypassing IOxAPIC. The MSI is generated by a Memory Write to address 0FEEx_xxxxh. Three 32-bit registers exist in the PCI Express* controller to support this mechanism. The default values of these registers will be compatible to the default value of IOxAPIC. The software can reprogram these registers to required values. The three registers are MSI Control register (MSICR), MSI Address register (MSIAR) and MSI Data register (MSIDR). Depending on system requirement each PCI Express* channel can have a MSI block (provides better flexibility) or the PCI Express* controller as a whole will have one MSI block and all channels raise hardware interrupts to this block.

The MSI Control register (MSICR) contains all the information related to the capability of PCI Express* MSI interrupts. Note that the MSICR register has been separated into its components, MSIAPID, MSINPTR and MSICAPA for purposes of separate register definitions.

Bit Field	Default & Access	Description
15:8	00h	Reserved
7	0b RO	**64-bit Address Capable**. Hardwired to '0' to indicate that the PCI Express* bridge is capable of 32-bit MSI addressing.
6:4	0h R/W	**Multiple Message Enable**. The software writes this field to indicate the number of allocated messages, which is aligned to a power of two. When MSI is enabled, the software will allocate at least one message to the device.
3:1	001b RO	**Multiple Message Capable**. The PCI Express* requests a capability for two messages by initializing this field to a value of 001b.
0	0b R/W	**MSI Enable**. Selects the method of interrupt delivery. Interrupts are generated for one of five conditions as described in the descriptor control register for each channel. If none of these five conditions are selected, software must poll for status as no interrupts of either type will be generated. 0 = Legacy interrupts will be generated. 1 = Message Signaled Interrupts (MSI) will be generated.

3.7.39 MSIAR – MSI Address Register for PCI Express* (D2:F0)

Address Offset: 5C – 5Fh
Access: R/W
Size: 32 Bits
Default: FEE0_0000h

The MSI Address register (MSIAR) contains all the address related information to route MSI interrupts.

Bit Field	Default & Access	Description
31:20	FEEh R/W	**Address**. Most significant 12-bits of 32-bit address.
19:12	00h R/W	**Destination ID**. Should reflect the 63:56 bits of IOxAPIC redirection table entry. The MCH may substitute other values in this field when redirecting to the System Bus.
11:4	00h R/W	**Extended Destination ID**. Should reflect the 55:48 bits of IOxAPIC redirection table entry.
3	0b R/W	**Redirection Hint**. Used by the MCH to allow the interrupt message to be redirected. 0 = Direct. Message will be delivered to the agent listed in bits 19:12 1 = Redirect. Message will be delivered to an agent with a lower interrupt priority. This can be derived from bits 10:8 in the Data Field.
2	0b R/W	**Destination Mode**. Used only if Redirection Hint is set to '1'. 0 = Physical 1 = Logical
1:0	00b	Reserved

3.7.40 MSIDR – MSI Data Register (D2:F0)

Address Offset: 60h
Access: R/W
Size: 16 Bits
Default: 0000h

The MSI Data register (MSIDR) contains all the data related information to route MSI interrupts.

Bit Field	Default & Access	Description
15	0b R/W	**Trigger Mode.** Same as the corresponding bit in the I/O Redirection Table for that interrupt. 0 = Edge 1 = Level
14	0b R/W	**Delivery Status**. If using edge-triggered interrupts, this will always be 1, since only assertion is sent. If using level-triggered interrupts, then this bit indicates the state of the interrupt input.
13:12	0b	Reserved
11	0b R/W	**Destination Mode**. Same as bit 2 of MSIAR. 0 = Physical 1 = Logical

Bit Field	Default & Access	Description
10:8	0h R/W	**Delivery Mode**.Same as the corresponding bits in the I/O Redirection Table for that interrupt. 000 = Fixed 100 = NMI 001 = Lowest Priority 101 = INIT 010 = SMI/PMI 110 = Reserved 011 = Reserved 111 = ExtINT
7:0	00h R/W	**Interrupt Vector**. Same as the corresponding bits in the I/O Redirection Table for that interrupt.

3.7.41 EXP_CAPID – PCI Express* Features Capabilities Structure (D2:F0)

Address Offset: 64h
Access: RO
Size: 8 Bits
Default: 10h

This register identifies the PCI Express* features capability structure.

Bit Field	Default & Access	Description
7:0	10h RO	**CAP_ID**. This field has the value 10h to identify the CAP_ID assigned by the PCI SIG for the PCI Express* capability structure.

3.7.42 EXP_NPTR – PCI Express* Next Capabilities Pointer (D2:F0)

Address Offset: 65h
Access: RO
Size: 8 Bits
Default: 00h
This register identifies the next PCI Express* capability structure.

Bit Field	Default & Access	Description
7:0	00h RO	**Next Capability Pointer**. Indicates there are no additional capability structures.

3.7.43 EXP_CAPA – PCI Express* Features Capabilities (D2:F0)

Address Offset: 66 – 67h
Access: RO, R/WO
Size: 16 Bits
Default: 0041h

This register identifies PCI Express* device type and associated capabilities.

Bit Field	Default & Access	Description
15:14	00b	Reserved
13:9	00000b RO	**Interrupt Message Number**. If the function is allocated more than one MSI interrupt number, this field will contain the offset between the base Message Data and the MSI Message that is generated when any of the status bits in either the Slot Status or Root Port Status registers of this capability structure are set. Hardware will update this field so that it is correct if the number of MSI Messages assigned to the device (based on the setting of the Multiple Message Enable bits in the MSI Capabilities register).
8	0b R/WO	**Slot Implemented**. BIOS must set this bit at boot time if the PCI Express* link associated with this port is connected to a slot (as compared to being connected to a motherboard component, or being disabled). 0 = Slot not implemented. 1 = Slot implemented.
7:4	4h RO	**Device/Port Type**. Hardwired to a value of 4h to indicate a Root Port.
3:0	1h RO	**Capability Version**. Hardwired to 1h to indicate compliance with the *PCI Express Interface Specification, Rev 1.0a*.

3.7.44 EXP_DEVCAP – PCI Express* Device Capabilities (D2:F0)

Address Offset: 68 – 6Bh
Access: RO
Size: 32 Bits
Default: 0002_8001h

This register identifies the device capabilities for PCI Express*.

Bit Field	Default & Access	Description
31:28	0bS	Reserved
27:26	00b RO	**Slot Power Limit Scale**. Specifies the scale used for the Slot Power Limit value. This value is set by the Set_Slot_Power_Limit message. This is not applicable to root ports. 00b = 1.0x (25.5 – 255) 01b = 0.1x (2.55 – 25.5) 10b = 0.01x (0.255 – 2.55) 11b = 0.001x (0.0 – 0.255)
25:18	00h RO	**Slot Power Limit Value.** In combination with the Slot Power Limit Scale value, specifies the upper limit on power supplied by the slot. Power Limit (in watts) is calculated by multiplying the value in this field by the value in the Slot Power Limit Scale field. This is not applicable to root ports.
17:15	000b	Reserved

Bit Field	Default & Access	Description
14	0b RO	**Power Indicator Present.** Not Applicable to a Root port.
13	0b RO	**Attention Indicator Present.** Not Applicable to a Root port.
12	0b RO	**Attention Button Present.** Not Applicable to a Root port.
11:9	000b RO	**Endpoint L1 Acceptable Latency**. Not Applicable to a Root port.
8:6	000b RO	**Endpoint L0s Acceptable Latency**. Not Applicable to a Root port.
5	0b RO	**Extended Tag Field Supported**. Hardwired to 0b, indicating 5 bits, as required for a Root port.
4:3	00b RO	**Phantom Functions Supported**. Hardwired to 00b as required for Root ports, indicating that devices may implement all function numbers.
2:0	001b RO	**Max Payload Size Supported**. Hardwired to 001b to indicate a maximum 256B payload size. Note that this refers to an inbound payload size, since the outbound payload size is restricted to a cacheline size to a value of 64B.

3.7.45 EXP_DEVCTL – PCI Express* Device Control (D2:F0)

Address Offset: 6C – 6Dh
Access: RO, RW
Size: 16 Bits
Default: 0000h

This register details PCI Express* device-specific parameters.

Bit Field	Default & Access	Description
15	0b	Reserved
14:12	000b R/W	**Max Read Request Size**. This field sets the maximum Read Request size for the device as a requester. The MCH will not generate read requests with size exceeding the set value. 000b = 128B 100b = 2KB 001b = 256B 101b = 4KB 010b = 512B 110b = Reserved 011b = 1KB 111b = Reserved The MCH will break cache-line (64B) reads destined for the I/O subsystem into pairs of aligned sequential 32B reads. Because of this, the read request size is not limited by the value in this register.
11	0b RO	**Enable No Snoop.** Software override on usage of the “No Snoop” attribute.The MCH never issues transactions with this attribute set.
10	0b RO	**AUX Power PM Enable**. Not Applicable.
9	0b RO	**Phantom Functions Enable**. Not Applicable to a Root port.
8	0b RO	**Extended Tag Field Enable**. Not Applicable to a Root port.

Bit Field	Default & Access	Description
7:5	000b R/W	**Max Payload Size**. This field sets the maximum TLP payload size for the device. As a receiver, the device must handle TLPs as large as the set value; as transmitter, the device must not generate TLPs exceeding the set value. Permissible values that can be programmed are indicated by the max_payload_size supported in the device capabilities register. Defined encodings for this field are: 000b = 128B 100b = 2KB 001b = 256B 101b = 4KB 010b = 512B 110b = Reserved 011b = 1KB 111b = Reserved
4	0b RO	**Enable Relaxed Ordering.** Permits the device to set the Relaxed Ordering bit in the attributes field of transactions it issues that do not require strong write ordering. Writes cannot be performed on this bit field. MCH sends TLPs with Enable Relaxed Ordering (ERO) attribute set as follows: For Messages, I/O space, memory space & CFG space transactions, ERO = 0 For Peer-to-Peer transactions, ERO received from the transaction requester will be transmitted to the target as is. For Completions packets, ERO from the original request will be copied into completion packets.
3	0b R/W	**Unsupported Request Reporting Enable**. Enables/Disables reporting of Unsupported Request errors. Note that the reporting of error messages (ERR_CORR, ERR_NONFATAL, ERR_FATAL) is controlled exclusively by the Root Port Command register. 0 = Disable reporting of Unsupported Request errors 1 = Enable reporting of Unsupported Request errors
2	0b R/W	**Fatal Error Reporting Enable**. This bit controls the reporting of fatal errors. Note that the reporting of fatal errors is internal to the root. No external ERR_FATAL message is generated. PCICMD[SERRE] when set can also enable reporting of both internal and external errors to be reported. 0 = Disable fatal error reporting 1 = Enable fatal error reporting
1	0b R/W	**Non-Fatal Error Reporting Enable**. This bit controls the reporting of nonfatal errors. Note that the reporting of nonfatal errors is internal to the root. No external ERR_NONFATAL message is generated. PCICMD[SERRE] when set can also enable reporting of both internal and external errors to be reported. 0 = Disable nonfatal error reporting 1 = Enable nonfatal error reporting
0	0b R/W	**Correctable Error Reporting Enable**. This bit controls the reporting of correctable errors. Note that the reporting of correctable errors is internal to the root. No external ERR_CORR message is generated. 0 = Disable correctable error reporting 1 = Enable correctable error reporting

3.7.46 EXP_DEVSTS – PCI Express* Device Status (D2:F0)

Address Offset: 6E – 6Fh
Access: RO, R/WC
Size: 16 Bits
Default: 0000h

This register provides information about PCI Express* device-specific parameters.

Bit Field	Default & Access	Description
15:6	0	Reserved
5	0b RO	**Transactions Pending**. Indicates that the device has transactions pending. 0 = Cleared by hardware only when all pending transactions (including completions for any outstanding non-posted requests on any used virtual channel) have been completed. 1 = Set by hardware to indicate that transactions are pending (including completions for any outstanding non-posted requests for all used Traffic Classes).
4	0b RO	**AUX Power Detected**. Not Applicable.
3	0b R/WC	**Unsupported Request Detected**. Indicates that an Unsupported Request has been detected. This bit is set upon Unsupported Request detection regardless of whether or not error reporting is enabled in the Device Control register. Software clears this bit by writing a '1' to the bit location. 0 = No Unsupported Request detected 1 = Unsupported Request detected
2	0b R/WC	**Fatal Error Detected**. Indicates that a fatal error has been detected. This bit is set upon fatal error detection regardless of whether or not error reporting is enabled in the Device Control register. Software clears this bit by writing a '1' to the bit location. 0 = No fatal error detected 1 = Fatal error detected
1	0b R/WC	**Non-Fatal Error Detected**. Indicates that a nonfatal error has been detected. This bit is set upon nonfatal error detection regardless of whether or not error reporting is enabled in the Device Control register. Software clears this bit by writing a '1' to the bit location. 0 = No nonfatal error detected 1 = Nonfatal error detected
0	0b R/WC	**Correctable Error Detected**. Indicates that a correctable error has been detected. This bit is set upon correctable error detection regardless of whether or not error reporting is enabled in the Device Control register. Software clears this bit by writing a '1' to the bit location. 0 = No correctable error detected 1 = Correctable error detected

3.7.47 EXP_LNKCAP – PCI Express* Link Capabilities (D2:F0)

Address Offset: 70 – 73h
Access: RO
Size: 32 Bits
Default: 0203_E481h

This register identifies PCI Express* link-specific capabilities.

Bit Field	Default & Access	Description
31:24	02h RO	**Port Number**. This field indicates the PCI Express* port number for the associated PCI Express* link.
23:18	00h	Reserved
17:15	111b RO	**L1 Exit Latency**. Field is ignored if L1 ASPM is unsupported. 000b Less than 1 µs 001b 1 µs - 2 µs 010b 2 µs - 4 µs 011b 4 µs - 8 µs 100b 8 µs - 16 µs 101b 16 µs - 32 µs 110b 32 µs - 64 µs 111b More than 64 µs.
14:12	110b RO	**L0s Exit Latency**. Field is ignored as L0s is unsupported. Field reflects the required latency to exit from the L0s ASPM link state, and must be updated automatically by hardware when the common clock configuration bit is set in the Link Control register (bit 6, offset 74h). The MCH advertises the maximum exit latency (110) by default, and will switch to the encoding for 150 ns (010) if software sets the common clock configuration bit in Link Control. 000 Less than 64 ns 001 64 ns - 128 ns 010 128 ns - 256 ns (Reported value when common clock config bit is set.) 011 256 ns - 512 ns 100 512 ns - 1 µs 101 1 µs - 2 µs 110 2 µs - 4 µs (Default while common clock config bit is clear.) 111 Reserved
11:10	01b RO	**Active State PM**. MCH does not support ASPM L0s or L1. 00 = Reserved 01 = L0s Entry Supported 10 = Reserved 11 = L0s and L1 Entry Supported
9:4	001000b R/WO	**Maximum Link Width**. This field indicates the maximum width of the PCI Express* link. Device 2 will report a value of 001000b, indicating a max link width of x8. However, if two separate devices are connected to port A (Device 2) and port A1 (Device 3), the max link width for both ports will be x4.
3:0	1h R/WO	**Maximum Link Speed**. Hardwired to a value of 1h, to indicate a max link speed of 2.5Gb/s.

3.7.48 EXP_LNKCTL – PCI Express* Link Control (D2:F0)

Address Offset: 74 – 75h
Access: RO, R/W, WO
Size: 16 Bits
Default: 0000h

This register controls PCI Express* link-specific parameters.

Bit Field	Default & Access	Description
15:8	0	Reserved
7	0b R/W	**Extended Synch**. Provides external devices monitoring the link with additional time for to achieve bit and symbol lock before the link enters L0 state and resumes communication. 0 = Normal 1 = Force extended transmission of FTS ordered sets in FTS, and extra TS2 at exit of L1 prior to entering L0.
6	0b R/W	**Common Clock Configuration**. L1 and L0s are unsupported. Components use this common clock configuration information to report the correct L0s and L1 Exit Latencies. Components also communicate the “NFTS” parameter required to reacquire lock on L0s exit in the TS1/TS2 training sequence. Thus, it will be necessary to force a retrain to re-issue the NFTS parameter if this bit is updated, and consequently modifies the L0s exit latency value. 0 = This component and the component at the opposite end of the link are operating with asynchronous reference clocks. 1 = This component and the component at the opposite end of the link are operating with a distributed common reference clock.
5	0b WO	**Retrain Link.** Writing a ‘1’ to this bit initiates link retraining. This bit always returns a value of 0 on reads.
4	0b R/W	**Link Disable**. Disables/Enables the associated PCI Express* link. 0 = Enable 1 = Disable Does not affect address decode, so a transaction to a space allocated to the “present” but “disabled” port will get positively decoded and routed to that port. Once it gets there, the transaction layer will manufacture an abortive completion because the link layer is in the “down” state. Posted writes will be dropped, deferred I/O writes will be flagged complete, and any form of read will receive an all “F” (master abort) completion.
3	0b RO	**Read Completion Boundary.** Hardwired to ‘0’, indicating “RCB” capability of 64B.
2	0b	Reserved
1:0	00b R/W	**Active State Link PM Control**. Controls the level of active state power management supported on the associated PCI Express* link. 00b= Disabled 01b= Reserved 10b= Reserved 11b= Reserved

3.7.49 EXP_LNKSTS – PCI Express* Link Status (D2:F0)

Address Offset: 76 – 77h
Access: RO
Size: 16 Bits
Default: 1001h

This register provides information about PCI Express* link-specific parameters.

Bit Field	Default & Access	Description
15:13	000b	Reserved
12	1b RO	**Slot Clock Configuration**. Indicates whether or not the component uses the same physical reference clock that the platform provides on the connector. 0 = The component in the slot uses an independent reference clock, irrespective of the presence of a reference on the connector. 1 = The component in the slot uses the same physical reference clock provided on the connector.
11	0b RO	**Link Training.** Indicates link training in progress (Physical Layer LTSSM in Configuration or Recovery state); hardware clears this bit once Link Training is successfully trained to the L0 state 0 = Cleared by hardware once link training is complete. 1 = Set by hardware when link training is in progress.
10	0b RO	**Link Training Error**. 0 = No Link Training Error 1 = Link Training Error occurred during transition from config or recovery to detect
9:4	000000b RO	**Negotiated Width.** All other values are Reserved. 000001b = x1 000100b = x4 001000b = x8
3:0	1h RO	**Link Speed**. Value of 1h indicates 2.5 GB/s link.

3.7.50 EXP_SLTCAP – PCI Express* Slot Capabilities (D2:F0)

Address Offset: 78 – 7Bh
Access: RO, R/WO
Size: 32 Bits
Default: 0000_0000h

This register identifies PCI Express* slot-specific capabilities.

Bit Field	Default & Access	Description
31:19	000h R/WO	**Physical Slot Number –** This field indicates the physical slot number attached to this Port. This field must be initialized to a value that assigns a slot number that is globally unique within the chassis. This field should be initialized to '0' for ports connected to devices that are integrated on the motherboard.
18:17	0	Reserved

Bit Field	Default & Access	Description
16:15	00b R/WO	**Slot Power Limit Scale –** Specifies the scale used for the Slot Power Limit Value. 00b = 1.0x (25.5 – 255) 01b = 0.1x (2.55 – 25.5) 10b = 0.01x (0.255 – 2.55) 11b = 0.001x (0.0 – 0.255)
14:7	00h R/WO	**Slot Power Limit Value –** In combination with the Slot Power Limit Scale value, specifies the upper limit on power supplied by slot. Power limit (in watts) is calculated by multiplying the value in this field by the value in the Slot Power Limit Scale field. This field should be programmed at boot; writing to this field triggers a Set Slot Power Limit inband PCI Express* message.
6	0b RO	**Hot-Plug Capable**. Hardwired to '0'. Device does not support Hot-Plug.
5	0b RO	**Hot-Plug Surprise.** Hardwired to '0'. Device does not support Hot-Plug.
4	0b RO	**Power Indicator Present**. Hardwired to '0'. Device does not support Hot-Plug.
3	0b RO	**Attention Indicator Present**. Hardwired to '0'. Device does not support Hot-Plug.
2	0b RO	**MRL Sensor Present**. Hardwired to '0'. Device does not support Hot-Plug.
1	0b RO	**Power Controller Present**. Hardwired to '0'. Device does not support Hot-Plug.
0	0b RO	**Attention Button Present**. Hardwired to '0'. Device does not support Hot-Plug.

3.7.51 EXP_SLTCTL – PCI Express* Slot Control (D2:F0)

Address Offset: 7C – 7Dh
Access: R/W
Size: 16 Bits
Default: 03C0h

This register controls PCI Express* slot-specific parameters.

Bit Field	Default & Access	Description
15:11	0	Reserved
10	0b R/W	**Power Controller Control**. Not Applicable. Device does not support Hot-Plug.
9:8	11b R/W	**Power Indicator Control.** Not Applicable. Device does not support Hot-Plug.
7:6	11b R/W	**Attention Indicator Control.** Not Applicable. Device does not support Hot-Plug.
5	0b R/W	**Hot-Plug Interrupt Enable.** Not Applicable. Device does not support Hot-Plug.
4	0b R/W	**Command Complete Interrupt Enable.** Not Applicable. Device does not support Hot-Plug.

Bit Field	Default & Access	Description
3	0b R/W	**Presence Detect Changed Enable.** Not Applicable. Device does not support Hot-Plug.
2	0b R/W	**MRL Sensor Changed Enable.** Not Applicable. Device does not support Hot-Plug.
1	0b R/W	**Power Fault Detected Enable.** Not Applicable. Device does not support Hot-Plug.
0	0b R/W	**Attention Button Pressed Enable.** Not Applicable. Device does not support Hot-Plug.

3.7.52 EXP_SLTSTS – PCI Express* Slot Status (D2:F0)

Address Offset: 7E – 7Fh
Access: RO, R/WC
Size: 16 Bits
Default: 0040h

This register provides information about PCI Express* slot-specific parameters.

Bit Field	Default & Access	Description
15:7	0	Reserved
6	1b RO	**Presence Detect State.** This bit indicates the presence of a card in the slot. The bit reflects the Presence Detect status determined via in-band mechanism or via Present Detect pins on the slot itself. This register is required if a slot is implemented. 0 = Slot Empty 1 = Card present in slot
5	0b RO	**MRL Sensor State.** Not Applicable. Device does not support Hot-Plug.
4	0b R/WC	**Command Completed.** Not Applicable. Device does not support Hot-Plug.
3	0b R/WC	**Presence Detect Changed.** Not Applicable. Device does not support Hot-Plug.
2	0b R/WC	**MRL Sensor Changed.** Not Applicable. Device does not support Hot-Plug.
1	0b R/WC	**Power Fault Detected.** Not Applicable. Device does not support Hot-Plug.
0	0b R/WC	**Attention Button Pressed.** Not Applicable. Device does not support Hot-Plug.

3.7.53 EXP_RPCTL – PCI Express* Root Port Control (D2:F0)

Address Offset: 80 – 83h
Access: R/W
Size: 32 Bits
Default: 0000_0000h

This register enables the forwarding of error messages based on messages received.

Bit Field	Default & Access	Description
31:4	0	Reserved
3	0b R/W	**PME Interrupt Enable.** Enables/disables interrupt generation upon receipt of a PME message as reflected in the PME Status register bit. A PME interrupt is also generated if the PME Status register bit is already set when this bit is set from a cleared state. 0 = Disable PME interrupt generation 1 = Enable PME interrupt generation
2	0b R/W	**System Error on Fatal Error Enable.** Controls the Root Complex's response to fatal errors reported by any of the devices in the hierarchy associated with this Root Port. System error generation based on fatal errors also enabled by PCICMD[SERRE]. 0 = Disable System Error generation in response to fatal errors reported on this port. 1 = Enable System Error generation in response to fatal errors reported on this port.
1	0b R/W	**System Error on Non-Fatal Error Enable.** Controls the Root Complex's response to nonfatal errors reported by any of the devices in the hierarchy associated with this Root Port. System error generation based on non-fatal errors also enabled by PCICMD[SERRE]. 0 = Disable System Error generation in response to nonfatal errors reported on this port. 1 = Enable System Error generation in response to nonfatal errors reported on this port.
0	0b R/W	**System Error on Correctable Error Enable.** Controls the Root Complex's response to correctable errors reported by any of the devices in the hierarchy associated with this Root Port. 0 = Disable System Error generation in response to correctable errors reported on this port. 1 = Enable System Error generation in response to correctable errors reported on this port.

3.7.54 EXP_RPSTS – PCI Express* Root Port Status (D2:F0)

Address Offset: 84 – 87h
Access: RO, R/WC
Size: 32 Bits
Default: 0000_0000h

This register supports power management events.

Bit Field	Default & Access	Description
31:18	0	Reserved
17	0b RO	**PME Pending**. 0 = Cleared by hardware when no more PMEs are pending. 1 = Indicates that another PME is pending when the PME Status bit is set.

Bit Field	Default & Access	Description
16	0b R/WC	**PME Status.** 0 = Cleared by software writing a '1' to the bit location. 1 = PME has been asserted by the requestor indicated in the PME Requestor ID field.
15:0	0000h RO	**PME Requestor ID.** Indicates the PCI Requestor ID of the last PME requestor.

3.7.55 EXP_PFCCA – PCI Express* Posted Flow Control Credits Allocated (D2:F0)

Address Offset: C4-C7h
Access: RO
Size 32 bit
Default Value: 000C_0030h

Bit Field	Default & Access	Description
31:24	00h RO	Reserved
23:16	0Ch RO	**PRH flow control credits allocated.**
15:12	0h RO	Reserved
11:0	030h RO	**PRD flow control credits allocated.**

3.7.56 EXP_NPFCCA – PCI Express* Non Posted Flow Control Credits Allocated (D2:F0)

Address Offset: C8-CBh
Access: RO
Size: 32 bit
Default Value: 0008_0001h

Bit Field	Default & Access	Description
31:24	00h RO	Reserved
23:16	08h RO	**NPRH flow control credits allocated.**
15:12	0h RO	Reserved
11:0	001h RO	**NPRD flow control credits allocated.**

3.7.57 EXP_ENHCAPST – PCI Express* Enhanced Capability Structure (D2:F0)

Address Offset: 100 – 103h
Access: RO
Size: 32 Bits
Default: 0001_0001h

This register identifies the capability structure and points to the next structure. This enhanced configuration structure by definition starts at configuration offset 100h.

Bit Field	Default & Access	Description
31:20	000h RO	**Next Capability Pointer**. Hardwired to 000h to indicate that there are no other items in the capability list.
19:16	1h RO	**Capability Version.** Hardwired to 1h, to indicate *PCI Express Interface Specification, Rev 1.0a*.
15:0	0001h RO	**Extended Capability ID.** Hardwired to 0001h, to indicate advanced error reporting capability.

3.7.58 EXP_UNCERRSTS – PCI Express* Uncorrectable Error Status (D2:F0)

Address Offset: 104 – 107h
Access: RO, R/WC
Size: 32 Bits
Default: 0000_0000h

The Uncorrectable Error Status register reports the status of individual error sources on the PCI Express* device. An individual error status bit that is set indicates that a particular error occurred. Software may clear an error status bit by writing a '1' to the bit location. These bits are sticky through reset.

Bit Field	Default & Access	Description
31:21	0	Reserved
20	0b R/WC	**Unsupported Request Error Status.** This error, if the first uncorrectable error, will load the header log. This bit is sticky through reset. 0 = Cleared by writing a '1' to the bit location. 1 = Unsupported Request detected.
19	0b RO	**ECRC Error Status.** OPTIONAL. Not implemented.
18	0b R/WC	**Malformed TLP Status.** This error, if the first uncorrectable error, will load the header log. Malformed TLP errors include: data payload length issues, byte enable rule violations, and various other illegal field settings. This bit is sticky through reset. 0 = Cleared by writing a '1' to the bit location. 1 = Malformed TLP detected

Bit Field	Default & Access	Description
17	0b R/WC	**Receiver Overflow Status.** OPTIONAL. This error, if the first uncorrectable error, will load the header log. The MCH checks for overflows on the following upstream queues: posted, non-posted, and completion. This bit is sticky through reset. 0 = Cleared by writing a '1' to the bit location. 1 = Receiver Overflow detected
16	0b R/WC	**Unexpected Completion Status.** This bit will be set when the device receives a completion which does not correspond to any of the outstanding requests issued by that device. This error, if the first uncorrectable error, will load the header log. This bit is sticky through reset. 0 = Cleared by writing a '1' to the bit location. 1 = Unexpected Completion detected.
15	0b R/WC	**Completer Abort Status.** OPTIONAL If a request received violates the specific programming model of this device, but is otherwise legal, this bit will be set. This error, if the first uncorrectable error, will load the header log. This bit is sticky through reset. 0 = Cleared by writing a '1' to the bit location. 1 = Completer Abort detected.
14	0b R/WC	**Completion Timeout Status.** The Completion Timeout timer will expire if a Request is not completed within 16 ms, but must not expire earlier than 50 us. When the timer expires, this bit will be set. This bit is sticky through reset. 0 = Cleared by writing a '1' to the bit location. 1 = Completion timeout detected.
13	0b R/WC	**Flow Control Protocol Error Status.** OPTIONAL The MCH asserts this bit for one of two conditions: 1. An FC update has been received which describes header or data credits for P, NP, or CPL which were originally advertised as infinite during initialization but are now advertised with non-zero or non-infinite values. 2. The minimum number of credits is not being advertised. The MCH implements this error bit, but only implements checking for conditions 1 and 3. This bit is sticky through reset. 0 = Cleared by writing a '1' to the bit location. 1 = Flow Control Protocol Error detected.
12	0b R/WC	**Poisoned TLP Status.** This bit when set indicates that some portion of the TLP data payload was corrupt. This error, if the first uncorrectable error, will load the header log. This bit is sticky through reset. 0 = Cleared by writing a '1' to the bit location. 1 = Poisoned TLP detected.
11:5	0	Reserved
4	0b R/WC	**Data Link Protocol Error Status.** This bit is set when an ACK/NAK received does not specify the sequence number of an unacknowledged TLP, or of the most recently acknowledged TLP. This bit is sticky through reset. 0 = Cleared by writing a '1' to the bit location. 1 = Data Link Protocol Error detected.
3:1	0	Reserved
0	0b RO	**Training Error Status.** OPTIONAL. Not implemented.

3.7.59 EXP_UNCERRMSK – PCI Express* Uncorrectable Error Mask (D2:F0)

Address Offset: 108 – 10Bh
Access: RO, R/W
Size: 32 Bits
Default: 0000_0000h

The Uncorrectable Error Mask register controls reporting of individual errors by device to the PCI Express* Root Complex via a PCI Express* error message. A masked error (respective bit set in mask register) is not reported to the PCI Express* Root Complex by an individual device. However, masked errors are still logged in the Uncorrectable Error Status register. There is one mask bit corresponding to every implemented bit in the Uncorrectable Error Status register. These bits are sticky through reset.

Bit Field	Default & Access	Description
31:21	0	Reserved
20	0b R/W	**Unsupported Request Error Mask.** 0 = Report Unsupported Request Error 1 = Mask Unsupported Request Error
19	0b RO	**ECRC Error Mask.** OPTIONAL. Not implemented.
18	0b R/W	**Malformed TLP Mask.** 0 = Report Malformed TLP Error 1 = Mask Malformed TLP Error
17	0b R/W	**Receiver Overflow Mask.** OPTIONAL 0 = Report Receiver Overflow Error 1 = Mask Receiver Overflow Error
16	0b R/W	**Unexpected Completion Mask.** 0 = Report Unexpected Completion Error 1 = Mask Unexpected Completion Error
15	0b R/W	**Completer Abort Mask.** OPTIONAL 0 = Report Completer Abort Error 1 = Mask Completer Abort Error
14	0b R/W	**Completion Timeout Mask.** 0 = Report Completion Timeout Error 1 = Mask Completion Timeout Error
13	0b R/W	**Flow Control Protocol Error Mask.** OPTIONAL 0 = Report Flow Control Protocol Error 1 = Mask Flow Control Protocol Error
12	0b R/W	**Poisoned TLP Mask.** 0 = Report Poisoned TLP Error 1 = Mask Poisoned TLP Error
11:5	0	Reserved
4	0b R/W	**Data Link Protocol Error Mask.** 0 = Report Data Link Protocol Error 1 = Mask Data Link Protocol Error

Bit Field	Default & Access	Description
3:1	0	Reserved
0	0b R/O	**Training Error Mask.** OPTIONAL. Not implemented.

3.7.60 EXP_UNCERRSEV – PCI Express* Uncorrectable Error Severity (D2:F0)

Address Offset: 10C – 10Fh
Access: RO, R/W
Size: 32 Bits
Default: 0006_2011h

The Uncorrectable Error Severity register controls whether an individual error is reported as a nonfatal or fatal error. An error is reported as fatal when the corresponding error bit in the severity register is set. If the bit is cleared, the corresponding error is considered nonfatal. These bits are sticky through reset.

Bit Field	Default & Access	Description
31:21	0	Reserved
20	0b R/W	**Unsupported Request Error Severity.** 0 = Nonfatal 1 = Fatal
19	0b RO	**ECRC Error Severity.** OPTIONAL. Not implemented.
18	1b R/W	**Malformed TLP Severity.** 0 = Nonfatal 1 = Fatal
17	1b R/W	**Receiver Overflow Severity.** OPTIONAL 0 = Nonfatal 1 = Fatal
16	0b R/W	**Unexpected Completion Severity.** 0 = Nonfatal 1 = Fatal
15	0b R/W	**Completer Abort Severity.** OPTIONAL 0 = Nonfatal 1 = Fatal
14	0b R/W	**Completion Timeout Severity.** 0 = Nonfatal 1 = Fatal
13	1b R/W	**Flow Control Protocol Error Severity.** OPTIONAL 0 = Nonfatal 1 = Fatal
12	0b R/W	**Poisoned TLP Severity.** 0 = Nonfatal 1 = Fatal
11:5	0	Reserved

Bit Field	Default & Access	Description
4	1b R/W	**Data Link Protocol Error Severity.** 0 = Nonfatal 1 = Fatal
3:1	0	Reserved
0	0b R/O	**Training Error Severity.** OPTIONAL 0 = Nonfatal 1 = Fatal

3.7.61 EXP_CORERRSTS – PCI Express* Correctable Error Status (D2:F0)

Address Offset: 110 – 113h
Access: R/WC
Size: 32 Bits
Default: 0000_0000h

The Correctable Error Status register reports the status of individual error sources on the PCI Express* device. An individual error status bit that is set indicates that a particular error occurred. Software may clear an error status bit by writing a '1' to the bit location. These bits are sticky through reset.

Bit Field	Default & Access	Description
31:13	0	Reserved
12	0b R/WC	**Replay Timer Timeout Status**. The replay timer counts time since the last ACK or NAK DLLP was received. When the timer expires, this bit is set. This bit is sticky through system reset. 0 = Cleared by writing a '1' to the bit location. 1 = Replay Timer timeout detected.
11:9	0	Reserved
8	0b R/WC	**REPLAY_NUM Rollover Status.** A 2-bit counter counts the number of times the retry buffer has been retransmitted. When this counter rolls over, this bit is set. This bit is sticky through system reset. 0 = Cleared by writing a '1' to the bit location. 1 = REPLAY_NUM rollover detected.
7	0b R/WC	**Bad DLLP Status.** This bit is set when the calculated DLLP CRC is not equal to the received value. Also included are 8b/10b errors within the TLP including wrong disparity. An invalid sequence number also sets this bit. This bit is sticky through system reset. 0 = Cleared by writing a '1' to the bit location. 1 = Bad DLLP detected.
6	0b R/WC	**Bad TLP Status.** OPTIONAL. This bit is set when the calculated TLP CRC is not equal to the received value. Also included are 8b/10b errors within the TLP including wrong disparity. An invalid sequence number also sets this bit. This bit is sticky through system reset. 0 = Cleared by writing a '1' to the bit location. 1 = Bad TLP detected.

Bit Field	Default & Access	Description
5:1	0	Reserved
0	0b R/WC	**Receiver Error Status.** OPTIONAL. Receiver Error Status register is set for 8b/10b errors received, framing errors received irrespective of the packet boundaries. 0 = Cleared by writing a '1' to the bit location. 1 = Receiver Error detected.

3.7.62 EXP_CORERRMSK – PCI Express* Correctable Error Mask (D2:F0)

Address Offset: 114 – 117h
Access: R/W
Size: 32 Bits
Default: 0000_0000h

The Correctable Error Mask register controls reporting of individual errors by device to the PCI Express* Root Complex via a PCI Express* error message. A masked error (respective bit set in mask register) is not reported to the PCI Express* Root Complex by an individual device. However, masked errors are still logged in the Correctable Error Status register. There is one mask bit corresponding to every implemented bit in the Correctable Error Status register. These bits are sticky through reset.

Bit Field	Default & Access	Description
31:13	0	Reserved
12	0b R/W	**Replay Timer Timeout Mask**. This bit is sticky through system reset. 0 = Report Replay Timer Timeout error. 1 = Mask Replay Timer timeout error.
11:9	0	Reserved
8	0b R/W	**REPLAY_NUM Rollover Error Mask.** This bit is sticky through system reset. 0 = Report REPLAY_NUM rollover 1 = Mask REPLAY_NUM rollover.
7	0b R/W	**Bad DLLP Error Mask.** This bit is sticky through system reset. 0 = Report Bad DLLP error. 1 = Mask Bad DLLP error.
6	0b R/W	**Bad TLP Error Mask.** OPTIONAL. This bit is sticky through system reset. 0 = Report Bad TLP error. 1 = Mask Bad TLP error.
5:1	0	Reserved
0	0b R/W	**Receiver Error Mask.** OPTIONAL. This bit is sticky through system reset. 0 = Report Receiver error. 1 = Mask Receiver Error error.

3.7.63 EXP_AERCACR – PCI Express* Advanced Error Capabilities and Control (D2:F0)

Address Offset: 118 – 11Bh
Access: RO, R/W
Size: 32 Bits
Default: 0000_0000h

This register identifies the capability structure and points to the next structure. The first error pointer rearms after the unmasked errors have been cleared. Software, after clearing the errors, must read the register again to ensure that it is indeed cleared. If it finds that another error occurred, it cannot rely on the pointer or header, unless it detects that the error pointer changed from the last time it was read for the previous error. These bits are sticky through reset.

Bit Field	Default & Access	Description
31:9	0	Reserved
8	0b RO	**ECRC Check Enable**. Not supported on the MCH.
7	0b RO	**ECRC Check Capable**. Not supported on the MCH.
6	0b RO	**ECRC Generation Enable**. Not supported on the MCH.
5	0b RO	**ECRC Generation Capable**. Not supported on the MCH.
4:0	00h RO	**First Error Pointer.** Identifies the bit position of the first error reported in the Uncorrectable Error Status register. However, if a subsequent Uncorrectable Error occurs with a higher severity, this field will be over-written with the bit position of the subsequent error status bit. In the event of simultaneous errors, the pointer indicates the least significant bit of the group. This bit is sticky through system reset.

3.7.64 EXP_HDRLOG0 – PCI Express* Header Log DW0 (D2:F0)

Address Offset: 11C – 11Fh
Access: RO
Size: 32 Bits
Default: 0000_0000h

This register contains the first 32 bits of the header log locked down when the first uncorrectable error occurs that saves the header. To rearm this register all reported uncorrectable errors must be cleared from the register. Software after clearing the errors must read the register again to ensure that it is indeed cleared. If it finds that another error occurred, it can not rely on the pointer or header, unless it detects that the error pointer changed from the last time it was read for the previous error. Byte 0 of the header is located in byte 3 of the Header Log Register 0, byte 1 of the header is in byte 2 of the Header Log Register 0 and so forth. For 12 byte headers, only the first three of the four Header Log Registers are used, and values in HDRLOG3 are undefined. These bits are sticky through reset.

Bit Field	Default & Access	Description
31:0	0000_0000h RO	**Header Log 0**

3.7.65 EXP_HDRLOG1 – PCI Express* Header Log DW1 (D2:F0)

Address Offset: 120 – 123h
Access: RO
Size: 32 Bits
Default: 0000_0000h

The function of the Header Log registers is described in Section 3.7.64 on page 3-162. Header Log DW1 contains the second 32 bits of the header. Byte 4 of the header is located in byte 3 of the Header Log Register 1, byte 5 of the header is in byte 2 of the Header Log Register 1 and so forth. These bits are sticky through reset.

Bit Field	Default & Access	Description
31:0	0000_0000h RO	**Header Log 1**

3.7.66 EXP_HDRLOG2 – PCI Express* Header Log DW2 (D2:F0)

Address Offset: 124 – 127h
Access: RO
Size: 32 Bits
Default: 0000_0000h

The function of the Header Log registers is described in Section 3.7.64 on page 3-162. Header Log DW2 contains the third 32 bits of the header. Byte 8 of the header is located in byte 3 of the Header Log Register 2, byte 9 of the header is in byte 2 of the Header Log Register 2 and so forth. These bits are sticky through reset.

Bit Field	Default & Access	Description
31:0	0000_0000h RO	**Header Log 2**

3.7.67 EXP_HDRLOG3 – PCI Express* Header Log DW3 (D2:F0)

Address Offset: 128 – 12Bh
Access: RO
Size: 32 Bits
Default: 0000_0000h

The function of the Header Log registers is described in Section 3.7.64 on page 3-162. Header Log DW3 contains the fourth 32 bits of the header. For 16-byte headers, byte 12 of the header is located in byte 3 of the Header Log Register 3, byte 13 of the header is in byte 2 of the Header Log Register 3 and so forth. For 12 byte headers, values in this register are undefined. These bits are sticky through reset.

Bit Field	Default & Access	Description
31:0	0000_0000h RO	**Header Log 3**

3.7.68 EXP_RPERRCMD – PCI Express* Root Port Error Command (D2:F0)

Address Offset: 12C – 12Fh
Access: RW
Size: 32 Bits
Default: 0000_0000h

This register controls the generation of interrupts (beyond the basic root complex capability to generate system errors) upon detection of errors. System error generation in response to PCI Express* error messages may be turned off by system software using the PCI Express* Capability Structure when advanced error reporting via interrupts is enabled.

Bit Field	Default & Access	Description
31:3	0	Reserved
2	0b R/W	**Fatal Error Interrupt Enable**. Enables the generation of an interrupt when a fatal error is reported by any of the devices in the hierarchy associated with this Root Port. This bit is sticky through reset. 0 = Disable interrupt generation on fatal error. 1 = Enable interrupt generation on fatal error.
1	0b R/W	**Nonfatal Error Interrupt Enable**. Enables the generation of an interrupt when an nonfatal error is reported by any of the devices in the hierarchy associated with this Root Port. This bit is sticky through reset. 0 = Disable interrupt generation on nonfatal error. 1 = Enable interrupt generation on nonfatal error.
0	0b	**Correctable Error Interrupt Enable**. Enables the generation of an interrupt when a correctable error is reported by any of the devices in the hierarchy associated with this Root Port. This bit is sticky through reset. 0 = Disable interrupt generation on correctable error. 1 = Enable interrupt generation on correctable error.

3.7.69 EXP_RPERRMSTS – PCI Express* Root Port Error Message Status (D2:F0)

Address Offset: 130 – 133h
Access: RO, R/WC
Size: 32 Bits
Default: 0000_0000h

This register reports the status of errors received by the root complex. Each correctable and uncorrectable (nonfatal and fatal) error source has a First Error bit and a Next Error bit. When an error is received by the root complex, the associated First Error bit is set and the Requestor ID is logged in the Error Source Identification register. Software may clear an error status bit by writing a '1' to the bit location. If software does not clear the first reported error before another error is received, the Next Error status bit will be set, but the Requestor ID of the subsequent error message is discarded.

Bit Field	Default & Access	Description
31:27	0h RO	**Advanced Error Interrupt Message Number**. If this function has been allocated more than one MSI interrupt number, this field will reflect the offset between the base Message Data and the MSI Message that is generated when any of the status bits of this capability are set.
26:7	0_0000h	Reserved
6	0b R/WC	**Fatal Error Messages Detected**. This bit is used by error handling software to determine whether fatal errors are outstanding in the hierarchy. In hardware, this bit along with bits 4 and 2 is used to clear fatal error escalation. This bit is sticky through reset. 0 = Software clears this bit by writing a '1' to the bit location. 1 = Fatal error message detected.
5	0b R/WC	**Non-Fatal Error Messages Detected**. This bit is used by error handling software to determine whether non-fatal errors are outstanding in the hierarchy. In hardware, this bit along with bits 4 and 2 is used to clear non-fatal error escalation. This bit is sticky through reset. 0 = Software clears this bit by writing a '1' to the bit location. 1 = Non-fatal error message detected.
4	0b R/WC	**First Uncorrectable Fatal Flag.** This bit is sticky through reset. 0 = First uncorrectable error is non-fatal. 1 = First uncorrectable error is fatal.
3	0b R/WC	**Multiple Uncorrectable Errors Detected**. In the unlikely event of two first errors occurring during the same clock period, only the first uncorrectable error message bit will be set. It will take an error to occur in a subsequent clock to set this bit. This bit is sticky through reset. 0 = Software clears this bit by writing a '1' to the bit location. 1 = Set when either a fatal or nonfatal error is received, and the First Uncorrectable Error Detected bit is already set. This indicates that one or more message Requestor IDs were lost.
2	0b R/WC	**First Uncorrectable Error Detected**. This bit is sticky through reset. 0 = Software clears this bit by writing a '1' to the bit location. 1 = Set when the first fatal or nonfatal error is received.

Bit Field	Default & Access	Description
1	0b R/WC	**Multiple Correctable Errors Detected**. In the unlikely event of two first errors occurring during the same clock period, only the first correctable error message bit will be set. It will take an error to occur in a subsequent clock to set this bit. This bit is sticky through reset. 0 = Software clears this bit by writing a '1' to the bit location. 1 = Set when a correctable error is received, and the First Correctable Error Detected bit is already set. This indicates that one or more message Requestor IDs were lost.
0	0b R/WC	**First Correctable Error Detected**. This bit is sticky through reset. 0 = Software clears this bit by writing a '1' to the bit location. 1 = Set when the first correctable error is received.

3.7.70 EXP_ERRSID – PCI Express* Error Source ID (D2:F0)

Address Offset: 134 – 137h
Access: RO
Size: 32 Bits
Default: 0000_0000h

This register reports the source (Requestor ID) of the first correctable and uncorrectable (fatal or nonfatal) errors reported in the Root Error Status register. These bits are sticky through reset.

Bit Field	Default & Access	Description
31:16	0000h RO	**Uncorrectable Error Source ID**. Requestor ID of the source when an uncorrectable error (fatal or nonfatal) is received, and the First Uncorrectable Error Detected bit is not already set. Since this ID could be for an internally detected error or from a message received from the other end of the link, in the event of errors detected in the same clock, priority will be given to the error received from the link, and that ID is what will be logged. These bits are sticky through reset.
15:0	0000h RO	**Correctable Error Source ID**. Requestor ID of the source when an correctable error is received, and the First Correctable Error Detected bit is not already set. Since this ID could be for an internally detected error or from a message received from the other end of the link, in the event of errors detected in the same clock, priority will be given to the error received from the link, and that ID is what will be logged. These bits are sticky through reset.

3.7.71 EXP_UNITERR – PCI Express* Unit Error Status (D2:F0)

Address Offset: 140 – 143h
Access: RO, R/WC
Size: 32 Bits
Default: 0000_0000h

This register is specific to the MCH. It captures the non-PCI Express* unit errors (those beyond the scope of the bus specification). The unit error mechanism is parallel to that used by "compatible" error registers and masks, but cannot feed back into standard registers because that would confuse standardized error handling software (which would not understand the extracurricular error bits). Escalation is controlled via the EXP_ERRDOCMD register (D2:F0:140-143h) for both standard and MCH-specific error types. Uncorrectable fatal errors feed into the fatal reporting select,

uncorrectable non-fatal errors feed into the non-fatal reporting select, and correctable errors feed into the correctable reporting select. The lower nibble is for HPC related errors. These bits are sticky through reset.

Bit Field	Default & Access	Description
31:16	0000h	Reserved
15	0b R/WC	**Upstream Posted Queue Overflow Status.** This bit is one of the components of the Receiver Overflow Status bit in the UNCERRSTS register. Even though this bit can be set, it is only reported through the receiver overflow bit in the UNCERRSTS register. The setting of this bit will never be logged in the local FERR/NERR registers or subsequently the global FERR/NERR registers, nor will it cause a SCI/SMI/SERR or MCERR message. At most, when the report mask is disabled, it could affect the unit error pointer. This functionality is provided as an aid to debug. This bit is sticky through reset. 0 = Software clears this bit by writing a '1' to the bit position. 1 = Overflow occurred for one of the posted header or data queues.
14	0b R/WC	**Upstream Non-Posted Queue Overflow Status**. This bit is one of the components of the Receiver Overflow Status bit in the UNCERRSTS register. Even though this bit can be set, it is only reported through the receiver overflow bit in the UNCERRSTS register. The setting of this bit will never be logged in the local FERR/NERR registers or subsequently the global FERR/NERR registers, nor will it cause a SCI/SMI/SERR or MCERR message. At most, when the report mask is disabled, it could affect the unit error pointer. This functionality is provided as an aid to debug. This bit is sticky through reset. 0 = Software clears this bit by writing a '1' to the bit position. 1 = Overflow occurred for the non-posted header queues.
13	0b R/WC	**Upstream Completion Queue Overflow Status.** Even though this bit can be set, it is only reported through the receiver overflow bit in the UNCERRSTS register. The setting of this bit will never be logged in the local FERR/NERR registers or subsequently the global FERR/NERR registers, nor will it cause a SCI/SMI/SERR or MCERR message. This functionality is provided as an aid to debug. This bit is sticky through reset. 0 = Software clears this bit by writing a '1' to the bit position. 1 = Overflow occurred for one of the completion header or data queues.
12	0b R/WC	**LLE Protocol Error.** This bit is set when the transaction layer detects a protocol error on the receiver interface from the LLE. Such an event should cause retraining eventually, but not necessarily immediately. The transaction with a problem will be dropped. This bit is sticky through reset. 0 = Software clears this bit by writing a '1' to the bit position. 1 = Transaction layer detected a protocol error on the receiver interface from the LLE
11	0b R/WC	**Link Down Error.** This bit is set when the link transition from DL_UP to DL_DOWN. This error will not be set if any of the BCTRL[6] (Secondary Bus Reset, HSILNKCTL[5] (Link Disabled), or VSCMD1[1] (Loopback Enable) bits are set. All of these would be SW initiated Link Down events which should not be considered errors. 0 = Software clears this bit by writing a '1' to the bit position. 1 = Set when the link transitions from DL_UP to DL_DOWN.
10	0b R/WC	**Downstream Data Queue Parity Error**. This bit is sticky through reset. 0 = Software clears this bit by writing a '1' to the bit position. 1 = Parity error occurred in the downstream data queue.

Bit Field	Default & Access	Description
9	0b R/WC	**LUT Parity Error**. Setting this bit indicates the logic is now lost and no transfers can be processed. The link is forced into retraining. When this error occurs, it will most likely be reported as multiple errors, since the error will be reported as long as the row with the error is accessed. This bit is sticky through reset. 0 = Software clears this bit by writing a '1' to the bit position. 1 = Parity error occurred in the Lookup Table.
8	0b R/WC	**LLRB Data Parity Error**. This error can only occur during a retry. This bit is sticky through reset. 0 = Software clears this bit by writing a '1' to the bit position. 1 = Data Parity error occurred in the Link Level Retry Buffer.
7	0b R/WC	**LLRB Header Parity Error.** The current transaction will be terminated as marked as bad. No further transfers will be completed. This error can only occur during a retry. This bit is sticky through reset. 0 = Software clears this bit by writing a '1' to the bit position. 1 = LLRB header parity error detected. Fatal.
6	0b R/WC	**LLRB Control Parity Error**. Since the pointer information is somehow lost or corrupted, no further data transfers can be completed. This error can only occur during a retry. This bit is sticky through reset. 0 = Software clears this bit by writing a '1' to the bit position. 1 = LLRB control parity error detected. Fatal.
5	0b R/WC	**DLLP Timeout Error.** DLLP flow control traffic is not recieved within the expected time. This bit is not set fo rtimeouts related to ACK or NAK. This bit is sticky through reset. 0 = Software clears this bit by writing a '1' to the bit position. 1 = DLLP timeout error occured..
4:3	00b	Reserved
2	0b R/WC	**SMB Clock Low Timeout (SMBCLTO)**. This bit is sticky through reset. 0 = Software clears this bit by writing a '1' to the bit position. 1 = SMB CLK low greater than 25ms.
1	0b R/WC	**Unexpected NAK on SMB (UESMBN)**. This bit is sticky through reset. 0 = Software clears this bit by writing a '1' to the bit position. 1 = Unexpected NAK on SMB detected.
0	0b R/WC	**SMB Lost Bus Arbitration (SMBLA)**. (Correctable) This bit is sticky through reset. 0 = Software clears this bit by writing a '1' to the bit position. 1 = SMB lost bus arbitration.

3.7.72 EXP_MASKERR – PCI Express* Mask Error (D2:F0)

Address Offset: 144 – 147h
Access: RO, R/W
Size: 32 Bits
Default: 0000_E000h

This register masks individual non-PCI Express* unit errors from being reported. They are still logged when masked, but only in the PCI Express* Unit Error Register. They will not be logged in either the local (EXP_FERR/EXP_NERR) or global FERR/NERR registers. The lowest nibble of this register is for the Hot-Plug Controller, and thus not applicable to this device.

Bit Field	Default & Access	Description
31:16	0000h	Reserved
15	1b R/W	**Upstream Posted Queue Overflow Mask**. Defaults to masked, normally reported through PCI Express* receive overflow status bit. This bit is sticky through reset. 0 = Enable upstream posted queue overflow reporting. 1 = Disable upstream posted queue overflow reporting.
14	1b R/W	**Upstream Non-Posted Queue Overflow Mask.** Defaults to masked, normally reported through PCI Express* receive overflow status bit. This bit is sticky through reset. 0 = Enable upstream non-posted queue overflow reporting. 1 = Disable upstream non-posted queue overflow reporting.
13	1b R/W	**Upstream Completion Queue Overflow Mask.** Defaults to masked, normally reported through PCI Express* receive overflow status bit. This bit is sticky through reset. 0 = Enable upstream completion queue overflow reporting. 1 = Disable upstream completion queue overflow reporting.
12	0b R/W	**LLE Protocol Error Mask.** This bit is sticky through reset. 0 = Enable LLE protocol error reporting. 1 = Disable LLE protocol error reporting.
11	0b R/W	**Link Down Error Mask.** Mask reporting of detected link transitions from DL_UP to DL_DOWN. This bit is sticky through reset. 0 = Enable link down error mask reporting. 1 = Disable link down error mask reporting.
10	0b R/W	**Downstream Data Queue Parity Error Reporting Mask.** This bit is sticky through reset. 0 = Enable 1 = Disable
9	0b R/W	**LUT Parity Error Reporting Mask.** This bit is sticky through reset. 0 = Enable LUT parity error reporting. 1 = Disable LUT parity error reporting.
8	0b R/W	**LLRB Data Parity Error Reporting Mask.** This bit is sticky through reset. 0 = Enable LLRB data parity error reporting. 1 = Disable LLRB data parity error reporting.
7	0b R/W	**LLRB Header Parity Error Reporting Mask.** This bit is sticky through reset. 0 = Enable LLRB header parity error reporting. 1 = Disable LLRB header parity error reporting.

Bit Field	Default & Access	Description
6	0b R/W	**LLRB Control Parity Error Reporting Mask.** This bit is sticky through reset. 0 = Enable LLRB control parity error reporting. 1 = Disable LLRB control parity error reporting.
5	0b R/W	**DLLP Timeout Error Reporting Mask.** This bit is sticky through reset. 0 = Enable **DLLP timeout** error reporting. 1 = Disable **DLLP timeout** error reporting.
4:3	00b	Reserved
2	0b R/W	**SMBCLTO Reporting Mask.** This bit is sticky through reset. 0 = Enable SMBCLTO reporting. 1 = Disable SMBCLTO reporting.
1	0b R/W	**UESMBN Reporting Mask.** This bit is sticky through reset. 0 = Enable UESMBN reporting. 1 = Disable UESMBN reporting.
0	0b R/W	**SMBLA Reporting Mask.** This bit is sticky through reset. 0 = Enable SMBLA reporting. 1 = Disable SMBLA reporting.

3.7.73 EXP_ERRDOCMD – PCI Express* Error Do Command Register (D2:F0)

Address Offset: 148 – 14Bh
Access: RO, R/W
Size: 32 Bits
Default: 0000_0000h

This register supports PCI Express* error signaling commands. DO_SCI, DO_SMI, and DO_MCERR, DO_SERR must further be enabled by the PCI Express* Host Do Command register.

Bit Field	Default & Access	Description
31:29	000b	Reserved
28:24	0_0000b RO	**First Error Pointer for Unmasked PCI Express* Correctable Errors.** This pointer is rearmed when all unmasked errors have been cleared. In the event of simultaneous errors, the pointer indicates the least significant bit of the group. These bits are sticky.
23:21	000b	Reserved
20:16	0_0000b R/W	**First Error Pointer for PCI Express* unit errors**. This pointer is locked once any units errors are logged in the EXP_FERR. It is rearmed when all EXP_unit errors have been cleared. In the event of simultaneous errors, the pointer indicates the least significant bit of the group. This pointer is only valid for an error that is enabled for reporting. These bits are sticky.
15	0b R/W	**Enable Header Log Use for LLE Protocol Error.** The header log is used by PCI Express* uncorrectable errors. This feature is used to capture the header log for the LLE protocol error in the unit error register during debug.

Bit Field	Default & Access	Description
14	0b R/W	**PCI Express* Unit Report Enable.** This bit enables reporting of fatal, non-fatal and correctable unit errors. 0 = Disable 1 = Enable
13:12	00b R/W	**Fatal Unit Error Report Steering**. Selects the method of reporting fatal errors if unit error reporting is enabled (bit 14). 00b = SCI 10b = SERR 01b = SMI 11b = MCERR
11:10	00b R/W	**Non-Fatal Unit Error Report Steering**. Selects the method of reporting non-fatal errors if unit error reporting is enabled (bit 14). 00b = SCI 10b = SERR 01b = SMI 11b = MCERR
9:8	00b R/W	**Correctable Unit Error Report Steering**. Selects the method of reporting correctable errors if unit error reporting is enabled (bit 14). 00b = SCI 10b = SERR 01b = SMI 11b = MCERR
7:6	00b	Reserved
5:4	00b R/W	**Fatal Root Port Error Report Steering**. Selects the method of reporting a fatal error in any of the devices in the hierarchy associated with this root port (internally detected or via a message received on the link). This functionality is enabled via EXP_RPCTL[2]. 00b = SCI 10b = SERR 01b = SMI 11b = MCERR
3:2	00b R/W	**Nonfatal Root Port Error Report Steering**. Selects the method of reporting a non-fatal error in any of the devices in the hierarchy associated with this root port (internally detected or via a message received on the link). This functionality is enabled via EXP_RPCTL[1]. 00b = SCI 10b = SERR 01b = SMI 11b = MCERR
1:0	00b R/W	**Correctable Root Port Error Report Steering**. Selects the method of reporting a correctable error in any of the devices in the hierarchy associated with this root port (internally detected or via a message received on the link). This functionality is enabled via EXP_RPCTL[0]. 00b = SCI 10b = SERR 01b = SMI 11b = MCERR

3.7.74 EXP_UNCERRDMSK – PCI Express* Uncorrectable Error Detect Mask (D2:F0)

Address Offset: 14C – 14Fh
Access: RO, R/W
Size: 32 Bits
Default: 0000_0000h

The Uncorrectable Error Detect Mask register controls detection of individual errors. An error event that is masked in this register will not be logged in the Uncorrectable Error Status register, and will never be reported. There is one mask bit corresponding to every implemented bit in the Uncorrectable Error Status register. This register is specific to the MCH. These bits are sticky through reset.

Bit Field	Default & Access	Description
31:21	000h	Reserved
20	0b R/W	**Unsupported Request Error Detect Mask**. This bit is sticky through reset. 0 = Detect Unsupported Request Error 1 = Disable Unsupported Request Error detection
19	0b RO	**ECRC Error Detect Mask.** OPTIONAL. Not implemented.
18	0b R/W	**Malformed TLP Detect Mask**. This bit is sticky through reset. 0 = Detect Malformed TLP Error 1 = Disable Malformed TLP Error detection
17	0b R/W	**Receiver Overflow Detect Mask.** This bit is sticky through reset. OPTIONAL 0 = Detect Receiver Overflow Error 1 = Disable Receiver Overflow Error detection
16	0b R/W	**Unexpected Completion Detect Mask**. This bit is sticky through reset. 0 = Detect Unexpected Completion Error 1 = Disable Unexpected Completion Error detection
15	0b R/W	**Completer Abort Detect Mask.** This bit is sticky through reset. OPTIONAL 0 = Detect Completer Abort Error 1 = Disable Completer Abort Error detection
14	0b R/W	**Completion Timeout Detect Mask**. This bit is sticky through reset. 0 = Detect Completion Timeout Error 1 = Disable Completion Timeout Error detection
13	0b R/W	**Flow Control Protocol Error Detect Mask.** This bit is sticky through reset. OPTIONAL 0 = Detect Flow Control Protocol Error 1 = Disable Flow Control Protocol Error detection
12	0b R/W	**Poisoned TLP Detect Mask.** This bit is sticky through reset. 0 = Detect Poisoned TLP Error 1 = Disable Poisoned TLP Error detection
11:5	0	Reserved
4	0b R/W	**Data Link Protocol Error Detect Mask.** This bit is sticky through reset. 0 = Detect Data Link Protocol Error 1 = Disable Data Link Protocol Error detection

Bit Field	Default & Access	Description
3:1	0	Reserved
0	0b R/O	**Training Error Detect Mask.** Not Implemented.

3.7.75 EXP_CORERRDMSK – PCI Express* Correctable Error Detect Mask (D2:F0)

Address Offset: 150 – 153h
Access: R/W
Size: 32 Bits
Default: 0000_0000h

The Correctable Error Detect Mask register controls detection of individual errors. An error event that is masked in this register will not be logged in the Correctable Error Status register, and will never be reported. There is one mask bit corresponding to every implemented bit in the Correctable Error Status register. These bits are sticky through reset. This register is specific to the MCH.

Bit Field	Default & Access	Description
31:13	0	Reserved
12	0b R/W	**Replay Timer Timeout Detect Mask**. This bit is sticky through system reset. 0 = Detect Replay Timer Timeout error. 1 = Disable Replay Timer timeout error detection.
11:9	0	Reserved
8	0b R/W	**REPLAY_NUM Rollover Error Detect Mask.** This bit is sticky through system reset. 0 = Detect REPLAY_NUM rollover 1 = Disable REPLAY_NUM rollover detection.
7	0b R/W	**Bad DLLP Error Detect Mask.** This bit is sticky through system reset. 0 = Detect Bad DLLP error. 1 = Disable Bad DLLP error detection.
6	0b R/W	**Bad TLP Error Detect Mask.** OPTIONAL. This bit is sticky through system reset. 0 = Detect Bad TLP error. 1 = Disable Bad TLP error detection.
5:1	0	Reserved
0	0b R/W	**Receiver Error Detect Mask.** OPTIONAL. This bit is sticky through system reset. 0 = Detect Receiver error. 1 = Disable Receiver Error error detection.

3.7.76 EXP_UNITERRDMSK – PCI Express* Unit Error Detect Mask (D2:F0)

Address Offset: 158 – 15Bh
Access: R/W
Size: 32 Bits
Default: 0000_0000h

This register is specific to the MCH, and controls detection of the PCI Express* functional unit error conditions. These bits are sticky through reset.

Bit Field	Default & Access	Description
31:16	0000h	Reserved
15	0b R/W	**Upstream Posted Queue Overflow Detect Mask.** This bit is sticky through reset. 0 = Enable upstream posted queue overflow detection. 1 = Disable upstream posted queue overflow detection
14	0b R/W	**Upstream Non-Posted Queue Overflow Detect Mask.** This bit is sticky through reset. 0 = Enable upstream non-posted queue overflow detection. 1 = Disable upstream non-posted queue overflow detection
13	0b R/W	**Upstream Completion Queue Overflow Detect Mask.** This bit is sticky through reset. 0 = Enable completion queue overflow detection. 1 = Disable completion queue overflow detection
12	0b R/W	**LLE Protocol Error Detect Mask.** This bit is sticky through reset. 0 = Enable LLE protocol error detection. 1 = Disable LLE protocol error detection
11	0b R/W	**Link Down Error Detect Mask**. Mask detection of link transitions from DL_UP to DL_DOWN. This bit is sticky through reset. 0 = Enable link down error detection. 1 = Disable link down error detection
10	0b R/W	**Downstream Data Queue Parity Error Detect Mask.** This bit is sticky through reset. 0 = Enable downstream data queue parity error detection. 1 = Disable downstream data queue parity error detection
9	0b R/W	**LUT Parity Error Detect Mask.** This bit is sticky through reset. 0 = Enable LUT parity error detection. 1 = Disable LUT parity error detection
8	0b R/W	**LLRB Data Parity Error Detect Mask.** This bit is sticky through reset. 0 = Enable LLRB data parity error detection. 1 = Disable LLRB data parity error detection
7	0b R/W	**LLRB Header Parity Error Detect Mask.** This bit is sticky through reset. 0 = Enable LLRB header parity error detection. 1 = Disable LLRB header parity error detection
6	0b R/W	**LLRB Control Parity Error Detect Mask.** This bit is sticky through reset. 0 = Enable LLRB control parity error detection. 1 = Disable LLRB control parity error detection.

Bit Field	Default & Access	Description
5	0b R/W	DLLP Timeout Error Detect Mask. 0 = Enable DLLP Timeout error detection. 1 = Disable DLLP Timeout error detection
4:3	00b	Reserved
2	0b R/W	**SMB Clock Low Timeout Detect Mask**. This bit is sticky through reset. 0 = Enable SMB Clock Low Timeout error detection. 1 = Disable SMB Clock Low Timeout error detection.
1	0b R/W	**Unexpected NAK On SMB Detect Mask**. This bit is sticky through reset. 0 = Enable Unexpected NAK on SMB error detection. 1 = Disable Unexpected NAK on SMB error detection.
0	0b R/W	**SMB Lost Bus Arbitration Detect Mask**. This bit is sticky through reset. 0 = Enable SMB arbitration loss detection. 1 = Disable SMB arbitration loss detection.

3.7.77 EXP_FERR – PCI Express* First Error Register (D2:F0)

Address Offset: 160 – 163h
Access: R/WC
Size: 32 Bits
Default: 0000_0000h

Captures the first error in each category type. These bits are sticky through reset. This register is specific to the MCH.

Note: If multiple errors are reported in the same clock as the first error, all errors are latched.

Bit Field	Default & Access	Description
31:9	00_0000h	Reserved
8	0b R/WC	**Device Fatal Error Detected**. This bit is sticky through reset. 0 = Software clears this bit by writing a '1' to the bit position. 1 = Internal Fatal Error Detected
7	0b R/WC	**Device Nonfatal Error Detected**. This bit is sticky through reset. 0 = Software clears this bit by writing a '1' to the bit position. 1 = Device Nonfatal Error Detected
6	0b R/WC	**Device Correctable Error Detected**. This bit is sticky through reset. 0 = Software clears this bit by writing a '1' to the bit position. 1 = Device Correctable Error Detected
5	0b R/WC	**Unit Specific Fatal Error Detected**. This bit is for fatal errors not in the *PCI Express Interface Specification, Rev 1.0a* as logged by the EXP_UNITERR register. The EXP_MASKERR register only prevents reporting of the unit errors, but does not prevent the logging of errors in this register. This bit is sticky through reset. 0 = Software clears this bit by writing a '1' to the bit position. 1 = Unit Fatal Error Detected

Bit Field	Default & Access	Description
4	0b R/WC	**Unit Specific Non-fatal Error Detected**. This bit is for non-fatal errors not in the *PCI Express Specification* as logged by the EXP_UNITERR register. The EXP_MASKERR register only prevents reporting of the unit errors, but does not prevent the logging of errors in this register. This bit is sticky through reset. 0 = Software clears this bit by writing a '1' to the bit position. 1 = Unit Nonfatal Error Detected
3	0b R/WC	**Unit Specific Correctable Error Detected**. This bit is for correctable errors not in the *PCI Express Specification* as logged by the EXP_UNITERR register. The EXP_MASKERR register only prevents reporting of the unit errors, but does not prevent the logging of errors in this register. This bit is sticky through reset. 0 = Software clears this bit by writing a '1' to the bit position. 1 = Unit Correctable Error Detected
2	0b R/WC	**Fatal Error Message received**. This bit is set when an ERR_FATAL message is received. These received fatal error messages can be masked by the SERR enable bit in the Bridge Control register, if the SERR enable bit is a '0'. This bit is sticky through reset. 0 = Software clears this bit by writing a '1' to the bit position. 1 = Fatal Error Detected
1	0b R/WC	**Non-fatal Error Message received**. This bit is set when an ERR_NONFATAL message is received. These received non-fatal error messages can be masked by the SERR enable bit in the Bridge Control register, if the SERR enable bit is a '0'. This bit is sticky through reset. 0 = Software clears this bit by writing a '1' to the bit position. 1 = Non-fatal Error Detected
0	0b R/WC	**Correctable Error Message received**. This bit is set when an ERR_COR message is received. These received correctable error messages can be masked by the SERR enable bit in the Bridge Control register, if the SERR enable bit is a '0'. This bit is sticky through reset. 0 = Software clears this bit by writing a '1' to the bit position. 1 = Correctable Error Detected

3.7.78 EXP_NERR PCI Express* Next Error Register (D2:F0)

Address Offset: 164 – 167h
Access: R/WC
Size: 32 Bits
Default: 0000_0000h

This register logs errors after the corresponding field in EXP_FERR register is locked down. For bit definitions see Section 3.7.77 on page 3-175.

3.7.79 EXP_ERR_CTL – PCI Express* Error Control Register (D2:F0)

Address Offset: 168 – 16Bh
Access: R/W
Size: 32 Bits
Default: 0000_0000h

This register controls the MCH handling of errors on incoming data streams to the core.

Bit Field	Default & Access	Description
31:19	0000h	Reserved
18	0b R/W	**Data Poisoning Enable**. This bit controls whether or not the MCH marks data as "poisoned" when a parity error is detected on incoming data from PCI Express*. 0 = Error checking disabled. 1 = Error Checking Enabled. Incoming data with parity errors will be marked as "poisoned" before being sent on towards its destination.
17:0	0	Reserved

3.8 PCI Express* Port A1 Registers (D3:F0)

Device 3 is the PCI Express* port A1 virtual PCI-to-PCI bridge. The associated PCI Express* link has a maximum lane width of x4. When Device 2 is configured as a x8 PCI Express* link, this device is not available.

Table 3-7. PCI Express* Port A1 PCI Configuration Register Map (D3:F0) (Sheet 1 of 3)

Address Offset	Mnemonic	Register Name	Access	Default
00 – 01h	VID	Vendor Identification	RO	8086h
02 – 03h	DID	Device Identification	RO	3596h
04 – 05h	PCICMD	PCI Command Register	RO, R/W	0000h
06 – 07h	PCISTS	PCI Status Register	RO, R/WC	0010h
08h	RID	Revision Identification	RO	09h
0Ah	SUBC	Sub-Class Code	RO	04h
0Bh	BCC	Base Class Code	RO	06h
0Ch	CLS	Cache Line Size	R/W	00h
0Eh	HDR	Header Type	RO	01h
18h	PBUSN	Primary Bus Number	RO	00h
19h	SBUSN	Secondary Bus Number	R/W	00h
1Ah	SUBUSN	Subordinate Bus Number	R/W	00h
1Ch	IOBASE	I/O Base Address Register	R/W, RO	F0h
1Dh	IOLIMIT	I/O Limit Address Register	R/W	00h
1E – 1Fh	SEC_STS	Secondary Status Register	R/WC, RO	0000h
20 – 21h	MBASE	Memory Base Address Register	R/W	FFF0h

Table 3-7. PCI Express* Port A1 PCI Configuration Register Map (D3:F0) (Sheet 2 of 3)

Address Offset	Mnemonic	Register Name	Access	Default
22 – 23h	MLIMIT	Memory Limit Address Register	R/W	0000h
24 – 25h	PMBASE	Prefetchable Memory Base Address Reg.	R/W, RO	FFF1h
26 – 27h	PMLIMIT	Prefetchable Memory Limit Address Reg.	R/W, RO	0001h
28h	PMBASU	Prefetchable Mem Base Upper Addr. Reg.	RO, R/W	0Fh
2Ch	PMLMTU	Prefetchable Memory Limit Upper Address Register	RO, R/W	00h
34h	CAPPTR	Capabilities Pointer	RO	50h
3Ch	INTRLINE	Interrupt Line Register	R/W	00h
3Dh	INTRPIN	Interrupt Pin Register	RWO	01h
3Eh	BCTRL	Bridge Control Register	RO, R/W	00h
44h	VS_CMD0	Vendor-Specific Command Register 0	RO	00h
45h	VS_CMD1	Vendor-Specific Command Register 1	RO, R/W, R/WS	00h
46h	VS_STS0	Vendor-Specific Status Register 0	RO	00h
47h	VS_STS1	Vendor-Specific Status Register 1	RO, R/WC	00h
50h	PMCAPID	Power Management Capabilities Structure	RO	01h
51h	PMNPTR	Power Management Next Capabilities Pointer	RO	58h
52 – 53h	PMCAPA	Power Management Capabilities	RO	C822h
54 – 55h	PMCSR	Power Management Status and Control	RO,R/W	0000h
56h	PMCSRBSE	Power Management Status and Control Bridge Extensions	RO	00h
57h	PMDATA	Power Management Data	RO	00h
58h	MSICAPID	MSI Capabilities Structure	RO	05h
59h	MSINPTR	MSI Next Capabilities Pointer	RO	64h
5A – 5Bh	MSICAPA	MSI Capabilities	RO, R/W	0002h
5C – 5Fh	MSIAR	MSI Address Register for PCI Express*	R/W	FEE0_0000h
60h	MSIDR	MSI Data Register	R/W	0000h
64h	EXP_CAPID	PCI Express Features Capabilities Structure	RO	10h
65h	EXP_NPTR	PCI Express Next Capabilities Pointer	RO	00h
66 – 67h	EXP_CAPA	PCI Express Features Capabilities	RO, R/WO	0041h
68 – 6Bh	EXP_DEVCAP	PCI Express Device Capabilities	RO	0002_8001h
6C – 6Dh	EXP_DEVCTL	PCI Express Device Control	R/W. RO	0000h
6E – 6Fh	EXP_DEVSTS	PCI Express Device Status	RO, R/WC	0000h
70 – 73h	EXP_LNKCAP	PCI Express Link Capabilities	RO	0303_E441h
74 – 75h	EXP_LNKCTL	PCI Express Link Control	RO, R/W, WO	0000h
76 – 77h	EXP_LNKSTS	PCI Express Link Status	RO	1001h
78 – 7Bh	EXP_SLTCAP	PCI Express Slot Capabilities	RO, R/WO	0000_0000h

Table 3-7. PCI Express* Port A1 PCI Configuration Register Map (D3:F0) (Sheet 3 of 3)

Address Offset	Mnemonic	Register Name	Access	Default
7C – 7Dh	EXP_SLTCTL	PCI Express Slot Control	R/W	03C0h
7E – 7Fh	EXP_SLTSTS	PCI Express Slot Status	RO, R/WC	0040h
80 – 83h	EXP_RPCTL	PCI Express Root Port Control	R/W	0000_0000h
84 – 87h	EXP_RPSTS	PCI Express Root Port Status	RO, R/WC	0000_0000h
C4 – C7h	EXP_PFCCA	PCI Express Posted Flow Control Credits Allocated	RO	000C_0030h
C8 – CBh	EXP_NPFCCA	PCI Express Non Posted Flow Control Credits Allocated	RO	0008_0001h
100 – 103h	EXP_ENHCAPST	PCI Express Enhanced Capability Structure	RO	0001_0001h
104 – 107h	EXP_UNCERRSTS	PCI Express Uncorrectable Error Status	RO, R/WC	0000_0000h
108 – 10Bh	EXP_UNCERRMSK	PCI Express Uncorrectable Error Mask	RO, R/W	0000_0000h
10C – 10Fh	EXP_UNCERRSEV	PCI Express Uncorrectable Error Severity	RO, R/W	0006_0011h
110 – 113h	EXP_CORERRSTS	PCI Express Correctable Error Status	R/WC	0000_0000h
114 – 117h	EXP_CORERRMSK	PCI Express Correctable Error Mask	R/W	0000_0000h
118 – 11Bh	EXP_AERCACR	PCI Express Advanced Error Capabilities and Control	RO, R/W	0000_0000h
11C – 11Fh	EXP_HDRLOG0	PCI Express Header Log DW0	RO	0000_0000h
120 – 123h	EXP_HDRLOG1	PCI Express Header Log DW1	RO	0000_0000h
124 – 127h	EXP_HDRLOG2	PCI Express Header Log DW2	RO	0000_0000h
128 – 12Bh	EXP_HDRLOG3	PCI Express Header Log DW3	RO	0000_0000h
12C – 12Fh	EXP_RPERRCMD	PCI Express Root Port Error Command	R/W	0000_0000h
130 – 133h	EXP_RPERRMSTS	PCI Express Root Port Error Message Status	RO, R/WC	0000_0000h
134 – 137h	EXP_ERRSID	PCI Express Error Source ID	RO	0000_0000h
140 – 143h	EXP_UNITERR	PCI Express Unit Error Status	RO, R/WC	0000_0000h
144 – 147h	EXP_MASKERR	PCI Express Mask Error	RO, R/W	0000_E000h
148 – 14Bh	EXP_ERRDOCMD	PCI Express Error Do Command Register	RO, R/W	0000_0000h
14C – 14Fh	EXP_UNCERRDMSK	PCI Express Uncorrectable Error Detect Mask	RO, R/W	0000_0000h
150 – 153h	EXP_CORERRDMSK	PCI Express Correctable Error Detect Mask	R/W	0000_0000h
158 – 15Bh	EXP_UNITERRDMSK	PCI Express Unit Error Detect Mask	R/W	0000_0000h
160 – 163h	EXP_FERR	PCI Express First Error	R/WC	0000_0000h
164 – 167h	EXP_NERR	PCI Express Next Error	R/WC	0000_0000h
168 – 16Bh	EXP_ERR_CTL	PCI Express Error Control	R/W	0000_0000h

Except for the registers listed below, all registers for Device 3 are exactly as for Device 2. See Section 3.7 for the remaining register definitions.

3.8.1 DID – Device Identification (D3:F0)

Address Offset: 02 – 03h
Access: RO
Size: 16 Bits
Default: 3596h

Bit Field	Default & Access	Description
15:0	3596h RO	**Device Identification Number (DID)**. This is a 16 bit value assigned to the MCH Device 3, Function 0.

3.8.2 EXP_LNKCAP – PCI Express* Link Capabilities (D3:F0)

Address Offset: 70 – 73h
Access: RO
Size: 32 Bits
Default: 0303_E441h

Bit Field	Default & Access	Description
31:24	03h RO	**Port Number**. This field indicates the PCI Express* port number for the associated PCI Express* link.
23:18	00h	Reserved
17:15	111b RO	**L1 Exit Latency**. Field is ignored if L1 ASPM is unsupported. 000b Less than 1 µs 001b 1 µs - 2 µs 010b 2 µs - 4 µs 011b 4 µs - 8 µs 100b 8 µs - 16 µs 101b 16 µs - 32 µs 110b 32 µs - 64 µs 111b More than 64 µs
14:12	110b RO	**L0s Exit Latency**. Field is ignored as L0s is unsupported. Field reflects the required latency to exit from the L0s ASPM link state, and must be updated automatically by hardware when the common clock configuration bit is set in the Link Control register (bit 6, offset 74h). The MCH advertises the maximum exit latency (110) by default, and will switch to the encoding for 150 ns (010) if software sets the common clock configuration bit in Link Control. 000 Less than 64 ns 001 64 ns - 128 ns 010 128 ns - 256 ns (Reported value when common clock config bit is set.) 011 256 ns - 512 ns 100 512 ns - 1 µs 101 1 µs - 2 µs 110 2 µs - 4 µs (Default while common clock config bit is clear.) 111 Reserved

Bit Field	Default & Access	Description
11:10	01b RO	**Active State PM**. MCH does not support L0s or L1. 00 = Reserved 01 = L0s Entry Supported 10 = Reserved 11 = L0s and L1 Entry Supported
9:4	000100b RO	**Maximum Link Width**. This field indicates the maximum width of the PCI Express* link. The value reported in this field by Device 3 is 000100b, indicating a maximum link width of x4.
3:0	1h RO	**Maximum Link Speed**. Hardwired to a value of 1h, to indicate a max link speed of 2.5 Gb/s.

Note: All other bits is this register are the same as described in Section 3.7.47 on page 3-149

3.9 PCI Express* Port B Registers (D4:F0)

Device 4 is the PCI Express* port B x16 virtual PCI-to-PCI bridge.

Table 3-8. PCI Express* Port B PCI Configuration Register Map (D4:F0) (Sheet 1 of 3)

Address Offset	Mnemonic	Register Name	Access	Default
00 – 01h	VID	Vendor Identification	RO	8086h
02 – 03h	DID	Device Identification	RO	3597h
04 – 05h	PCICMD	PCI Command Register	RO, R/W	0000h
06 – 07h	PCISTS	PCI Status Register	RO, R/WC	0010h
08h	RID	Revision Identification	RO	09h
0Ah	SUBC	Sub-Class Code	RO	04h
0Bh	BCC	Base Class Code	RO	06h
0Ch	CLS	Cache Line Size	R/W	00h
0Eh	HDR	Header Type	RO	01h
18h	PBUSN	Primary Bus Number	RO	00h
19h	SBUSN	Secondary Bus Number	R/W	00h
1Ah	SUBUSN	Subordinate Bus Number	R/W	00h
1Ch	IOBASE	I/O Base Address Register	R/W, RO	F0h
1Dh	IOLIMIT	I/O Limit Address Register	R/W	00h
1E – 1Fh	SEC_STS	Secondary Status Register	R/WC, RO	0000h
20 – 21h	MBASE	Memory Base Address Register	R/W	FFF0h
22 – 23h	MLIMIT	Memory Limit Address Register	R/W	0000h
24 – 25h	PMBASE	Prefetchable Memory Base Address Reg.	R/W, RO	FFF1h
26 – 27h	PMLIMIT	Prefetchable Memory Limit Address Reg.	R/W, RO	0001h
28h	PMBASU	Prefetchable Mem Base Upper Addr. Reg.	RO, R/W	0Fh
2Ch	PMLMTU	Prefetchable Memory Limit Upper Address Register	RO, R/W	00h

Table 3-8. PCI Express* Port B PCI Configuration Register Map (D4:F0) (Sheet 2 of 3)

Address Offset	Mnemonic	Register Name	Access	Default
34h	CAPPTR	Capabilities Pointer	RO	50h
3Ch	INTRLINE	Interrupt Line Register	R/W	00h
3Dh	INTRPIN	Interrupt Pin Register	RWO	01h
3Eh	BCTRL	Bridge Control Register	RO, R/W	00h
44h	VS_CMD0	Vendor-Specific Command Register 0	RW	00h
45h	VS_CMD1	Vendor-Specific Command Register 1	RO, R/W, R/WS	00h
46h	VS_STS0	Vendor-Specific Status Register 0	RO	00h
47h	VS_STS1	Vendor-Specific Status Register 1	RO, R/WC	00h
50h	PMCAPID	Power Management Capabilities Structure	RO	01h
51h	PMNPTR	Power Management Next Capabilities Pointer	RO	58h
52 – 53h	PMCAPA	Power Management Capabilities	RO	C822h
54 – 55h	PMCSR	Power Management Status and Control	RO,R/W	0000h
56h	PMCSRBSE	Power Management Status and Control Bridge Extensions	RO	00h
57h	PMDATA	Power Management Data	RO	00h
58h	MSICAPID	MSI Capabilities Structure	RO	05h
59h	MSINPTR	MSI Next Capabilities Pointer	RO	64h
5A – 5Bh	MSICAPA	MSI Capabilities	RO, R/W	0002h
5C – 5Fh	MSIAR	MSI Address Register for PCI Express*	R/W	FEE0_0000h
60h	MSIDR	MSI Data Register	R/W	0000h
64h	EXP_CAPID	PCI Express Features Capabilities Structure	RO	10h
65h	EXP_NPTR	PCI Express Next Capabilities Pointer	RO	00h
66 – 67h	EXP_CAPA	PCI Express Features Capabilities	RO, R/WO	0041h
68 – 6Bh	EXP_DEVCAP	PCI Express Device Capabilities	RO	0002_8001h
6C – 6Dh	EXP_DEVCTL	PCI Express Device Control	R/W. RO	0000h
6E – 6Fh	EXP_DEVSTS	PCI Express Device Status	RO, R/WC	0000h
70 – 73h	EXP_LNKCAP	PCI Express Link Capabilities	RO	0403_E501h
74 – 75h	EXP_LNKCTL	PCI Express Link Control	RO, R/W, WO	0000h
76 – 77h	EXP_LNKSTS	PCI Express Link Status	RO	1001h
78 – 7Bh	EXP_SLTCAP	PCI Express Slot Capabilities	RO, R/WO	0000_0000h
7C – 7Dh	EXP_SLTCTL	PCI Express Slot Control	R/W	03C0h
7E – 7Fh	EXP_SLTSTS	PCI Express Slot Status	RO, R/WC	0040h
80 – 83h	EXP_RPCTL	PCI Express Root Port Control	R/W	0000_0000h
84 – 87h	EXP_RPSTS	PCI Express Root Port Status	RO, R/WC	0000_0000h
C4 – C7h	EXP_PFCCA	PCI Express Posted Flow Control Credits Allocated	RO	0010_0040h
C8 – CBh	EXP_NPFCCA	PCI Express Non Posted Flow Control Credits Allocated	RO	0008_0001h
100 – 103h	EXP_ENHCAPST	PCI Express Enhanced Capability Structure	RO	0001_0001h
104 – 107h	EXP_UNCERRSTS	PCI Express Uncorrectable Error Status	RO, R/WC	0000_0000h

Table 3-8. PCI Express* Port B PCI Configuration Register Map (D4:F0) (Sheet 3 of 3)

Address Offset	Mnemonic	Register Name	Access	Default
108 – 10Bh	EXP_UNCERRMSK	PCI Express Uncorrectable Error Mask	RO, R/W	0000_0000h
10C – 10Fh	EXP_UNCERRSEV	PCI Express Uncorrectable Error Severity	RO, R/W	0006_0011h
110 – 113h	EXP_CORERRSTS	PCI Express Correctable Error Status	R/WC	0000_0000h
114 – 117h	EXP_CORERRMSK	PCI Express Correctable Error Mask	R/W	0000_0000h
118 – 11Bh	EXP_AERCACR	PCI Express Advanced Error Capabilities and Control	RO, R/W	0000_0000h
11C – 11Fh	EXP_HDRLOG0	PCI Express Header Log DW0	RO	0000_0000h
120 – 123h	EXP_HDRLOG1	PCI Express Header Log DW1	RO	0000_0000h
124 – 127h	EXP_HDRLOG2	PCI Express Header Log DW2	RO	0000_0000h
128 – 12Bh	EXP_HDRLOG3	PCI Express Header Log DW3	RO	0000_0000h
12C – 12Fh	EXP_RPERRCMD	PCI Express Root Port Error Command	R/W	0000_0000h
130 – 133h	EXP_RPERRMSTS	PCI Express Root Port Error Message Status	RO, R/WC	0000_0000h
134 – 137h	EXP_ERRSID	PCI Express Error Source ID	RO	0000_0000h
140 – 143h	EXP_UNITERR	PCI Express Unit Error Status	RO, R/WC	0000_0000h
144 – 147h	EXP_MASKERR	PCI Express Mask Error	RO, R/W	0000_E000h
148 – 14Bh	EXP_ERRDOCMD	PCI Express Error Do Command Register	RO, R/W	0000_0000h
14C – 14Fh	EXP_UNCERRDMSK	PCI Express Uncorrectable Error Detect Mask	RO, R/W	0000_0000h
150 – 153h	EXP_CORERRDMSK	PCI Express Correctable Error Detect Mask	R/W	0000_0000h
158 – 15Bh	EXP_UNITERRDMSK	PCI Express Unit Error Detect Mask	R/W	0000_0000h
160 – 163h	EXP_FERR	PCI Express First Error	R/WC	0000_0000h
164 – 167h	EXP_NERR	PCI Express Next Error	R/WC	0000_0000h
168 – 16Bh	EXP_ERR_CTL	PCI Express Error Control	R/W	0000_0000h

Except for the registers listed below, all registers for Device 4 are exactly as for Device 2. See Section 3.7 on page 3-122 for the remaining register definitions.

3.9.1 DID – Device Identification (D4:F0)

Address Offset: 02 – 03h
Access: RO
Size: 16 Bits
Default: 3597h

Bit Field	Default & Access	Description
15:0	3597h RO	**Device Identification Number (DID)**. This is a 16 bit value assigned to the MCH Device 4, Function 0.

3.9.2 EXP_LNKCAP – PCI Express* Link Capabilities (D4:F0)

Address Offset: 70 – 73h
Access: RO
Size: 32 Bits
Default: 0403_E501h

Bit Field	Default & Access	Description
31:24	04h RO	**Port Number**. This field indicates the PCI Express* port number for the associated PCI Express* link.
23:18	00h	Reserved
17:15	111b RO	**L1 Exit Latency**. Field is ignored if L1 ASPM is unsupported. 000b Less than 1 µs 001b 1 µs - 2 µs 010b 2 µs - 4 µs 011b 4 µs - 8 µs 100b 8 µs - 16 µs 101b 16 µs - 32 µs 110b 32 µs - 64 µs 111b More than 64 µs
14:12	110b RO	**L0s Exit Latency**. Field is ignored as L0s is unsupported. Field reflects the required latency to exit from the L0s ASPM link state, and must be updated automatically by hardware when the common clock configuration bit is set in the Link Control register (bit 6, offset 74h). The MCH advertises the maximum exit latency (110) by default, and will switch to the encoding for 150 ns (010) if software sets the common clock configuration bit in Link Control. 000 Less than 64 ns 001 64 ns - 128 ns 010 128 ns - 256 ns (Reported value when common clock config bit is set.) 011 256 ns - 512 ns 100 512 ns - 1 µs 101 1 µs - 2 µs 110 2 µs - 4 µs (Default while common clock config bit is clear.) 111 Reserved
11:10	01b RO	**Active State PM**. MCH does not support L0s or L1. 00 Reserved 01 L0s Entry Supported 10 Reserved 11 L0s and L1 Entry Supported
9:4	010000b RO	**Maximum Link Width**. This field indicates the maximum width of the PCI Express* link. Device 4 will report a value of 010000b, indicating a maximum link width of x16.
3:0	1h RO	**Maximum Link Speed**. Hardwired to a value of 1h, to indicate a max link speed of 2.5 Gb/s.

3.9.3 EXP_LNKSTS – PCI Express* Link Status (D4:F0)

Address Offset: 76 – 77h
Access: RO
Size: 16 Bits
Default: 1001h

This register provides information about PCI Express* link-specific parameters.

Bit Field	Default & Access	Description
15:13	000b	Reserved
12	1b RO	**Slot Clock Configuration**. Indicates whether or not the component uses the same physical reference clock that the platform provides on the connector. 0 = The component in the slot uses an independent reference clock, irrespective of the presence of a reference on the connector. 1 = The component in the slot uses the same physical reference clock provided on the connector.
11	0b RO	**Link Training.** Indicates link training in progress (Physical Layer LTSSM in Configuration or Recovery state); hardware clears this bit once Link Training is successfully trained to the L0 state 0 = Cleared by hardware once link training is complete. 1 = Set by hardware when link training is in progress.
10	0b RO	**Link Training Error**. The Link Training machine must make it out of detect, before this bit can set to a '1'. If errors occur when in training, then this bit will be set. If the link successfully trains, this bit will be a '0'. 0 = No Link Training Error 1 = Link Training Error occurred
9:4	000000b RO	**Negotiated Width.** All other values are Reserved. 000001b = x1 000100b = x4 (Not Validated with Graphics Devices) 001000b = x8 (Not Validated with Graphics Devices) 010000b = x16
3:0	1h RO	**Link Speed**. Value of 1h indicates 2.5 GB/s link.

3.9.4 EXP_SLTCAP – PCI Express* Slot Capabilities (D4:F0)

Address Offset: 78 – 7Bh
Access: RO, R/WO
Size: 32 Bits
Default: 0000_0000h

Bit Field	Default & Access	Description
31:19	000h R/WO	**Physical Slot Number.** This field indicates the physical slot number attached to this Port. This field must be initialized to a value that assigns a slot number that is globally unique within the chassis. This field should be initialized to '0' for ports connected to devices that are integrated on the motherboard.
18:17	0	Reserved

Bit Field	Default & Access	Description
16:15	00b R/WO	**Slot Power Limit Scale.** Specifies the scale used for the Slot Power Limit Value. 00b = 1.0x (25.5 – 255) 01b = 0.1x (2.55 – 25.5) 10b = 0.01x (0.255 – 2.55) 11b = 0.001x (0.0 – 0.255)
14:7	00h R/WO	**Slot Power Limit Value.** In combination with the Slot Power Limit Scale value, specifies the upper limit on power supplied by slot. Power limit (in watts) calculated by multiplying the value in this field by the value in the Slot Power Limit Scale field. This field should be programmed at boot and writing to this field triggers a Set_Slot_Power_Limit in-band PCI Express* message.
6	0b RO	**Hot-Plug Capable**. Hardwired to '0'. Device 4 does not support Hot-Plug.
5	0b RO	**Hot-Plug Surprise.** Hardwired to '0'. Device 4 does not support Hot-Plug.
4	0b RO	**Power Indicator Present**. Hardwired to '0'. Device 4 does not support Hot-Plug.
3	0b RO	**Attention Indicator Present**. Hardwired to '0'. Device 4 does not support Hot-Plug.
2	0b RO	**MRL Sensor Present**. Hardwired to '0'. Device 4 does not support Hot-Plug.
1	0b RO	**Power Controller Present**. Hardwired to '0'. Device 4 does not support Hot-Plug.
0	0b RO	**Attention Button Present**. Hardwired to '0'. Device 4 does not support Hot-Plug.

3.9.5 EXP_SLTCTL – PCI Express* Slot Control (D4:F0)

Address Offset: 7C – 7Dh
Access: R/W
Size: 16 Bits
Default: 03C0h

Bit Field	Default & Access	Description
15:11	0	Reserved
10	0b	**Power Controller Control**. Not Applicable. Device 4 does not support Hot-Plug.
9:8	11b	**Power Indicator Control.** Not Applicable. Device 4 does not support Hot-Plug.
7:6	11b	**Attention Indicator Control.** Not Applicable. Device 4 does not support Hot-Plug.
5	0b	**Hot-Plug Interrupt Enable.** Not Applicable. Device 4 does not support Hot-Plug.
4	0b	**Command Complete Interrupt Enable.** Not Applicable. Device 4 does not support Hot-Plug.
3	0b	**Presence Detect Changed Enable.** Not Applicable. Device 4 does not support Hot-Plug.
2	0b	**MRL Sensor Changed Enable.** Not Applicable. Device 4 does not support Hot-Plug.

Bit Field	Default & Access	Description
1	0b	**Power Fault Detected Enable.** Not Applicable. Device 4 does not support Hot-Plug.
0	0b	**Attention Button Pressed Enable.** Not Applicable. Device 4 does not support Hot-Plug.

3.9.6 EXP_SLTSTS – PCI Express* Slot Status (D4:F0)

Address Offset: 7E – 7Fh
Access: RO, R/WC
Size: 16 bit
Default: 0040h

Bit Field	Default & Access	Description
15:7	0	Reserved
6	1b RO	**Presence Detect State.** Reflects the status of the Presence Detect pin, to indicate the presence of a card in the slot. 0 = Slot Empty 1 = Card present in slot
5	0b	**MRL Sensor State.** Not Applicable. Device 4 does not support Hot-Plug.
4	0b	**Command Completed.** Not Applicable. Device 4 does not support Hot-Plug.
3	0b	**Presence Detect Changed.** Not Applicable. Device 4 does not support Hot-Plug.
2	0b	**MRL Sensor Changed.** Not Applicable. Device 4 does not support Hot-Plug.
1	0b	**Power Fault Detected.** Not Applicable. Device 4 does not support Hot-Plug.
0	0b	**Attention Button Pressed.** Not Applicable. Device 4 does not support Hot-Plug.

3.9.7 EXP_PFCCA – PCI Express* Flow Control Credits Allocated (D4:F0)

Address Offset: C4-C7h
Access: RO
Size 32 bit
Default Value: 0010_0040h

Bit Field	Default & Access	Description
31:24	00h RO	Reserved
23:16	10h RO	**PRH flow control credits allocated.**
15:12	0h RO	Reserved
11:0	040h RO	**PRD flow control credits allocated.**

3.9.8 EXP_NPFCCA – PCI Express* Non Posted Flow Control Credits Allocated (D4:F0)

Address Offset: C8-CBh
Access: RO
Size: 32 bit
Default Value: 000C_0001h

Bit Field	Default & Access	Description
31:24	00h RO	Reserved
23:16	0Ch RO	**NPRH flow control credits allocated.**
15:12	0h RO	Reserved
11:0	001h RO	**NPRD flow control credits allocated.**

3.10 Extended Configuration Registers (D8:F0)

The Extended Configuration Registers comprise Device 8 (D8), Function 0 (F0). Table 3-9 provides the register address map for this device and function.

Table 3-9. Extended Configuration Registers PCI Configuration Register Map (D8:F0)

Offset	Mnemonic	Register Name	Access	Default
00 – 01h	VID	Vendor Identification	RO	8086h
02 – 03h	DID	Device Identification	RO	359Bh
04 – 05h	PCICMD	PCI Command Register	RO	0000h
06 – 07h	PCISTS	PCI Status Register	RO	0080h
08h	RID	Revision Identification	RO	09h
0Ah	SUBC	Sub-Class Code	RO	80h
0Bh	BCC	Base Class Code	RO	08h
0Eh	HDR	Header Type	RO	00h
2C – 2Dh	SVID	Subsystem Vendor Identification	R/WO	0000h
2E – 2Fh	SID	Subsystem Identification	R/WO	0000h
C8 – CBh	SCRUBLIM	Scrub Limit and Control Register	R0, R/W, R/WS	0000_0000h
CC – CFh	SCRBADD	Scrub Address	RO, R/W	0000_0000h
D0 – D3h	DTCL	DRAM Power Management Control Lower Register	RO, R/W	2000_0000h
D4 – D7h	DTCU	DRAM Power Management Control Upper Register	RO, R/W	0000_0000h
F3h	CORR	Processor Only Reset Register	RO, R/W, R/WS	00h

3.10.1 VID – Vendor Identification (D8:F0)

Address Offset: 00 – 01h
Access: RO
Size: 16 Bits
Default: 8086h

The VID register contains the vendor identification number. This 16-bit register combined with the Device Identification register uniquely identifies any PCI device. Writes to this register have no effect.

Bit Field	Default & Access	Description
15:0	8086h RO	**Vendor Identification (VID)**. This register field contains the PCI standard identification for Intel, 8086h.

3.10.2 DID – Device Identification (D8:F0)

Address Offset: 02 – 03h
Access: RO
Size: 16 Bits
Default: 359Bh

This 16-bit register combined with the Vendor Identification register uniquely identifies any PCI device. Writes to this register have no effect.

Bit Field	Default & Access	Description
15:0	359Bh	**Device Identification Number (DID)**. This is a 16-bit value assigned to the MCH Extended Configuration Registers registers, Device 8.

3.10.3 PCICMD – PCI Command Register (D8:F0)

Address Offset: 04 – 05h
Access: RO
Size: 16 Bits
Default: 0000h

Since MCH Device 8 does not physically reside on a real PCI bus, portions of this register are not implemented.

Bit Field	Default & Access	Description
15:10	00h	Reserved
9	0b RO	**Fast Back-to-Back Enable (FB2B).**This function is not enabled for Device 8.
8	0b RO	**SERR Enable (SERRE)**. This function is not enabled for Device 8.
7	0b RO	**Address/Data Stepping Enable (ADSTEP)**. This function is not enabled for Device 8.

Bit Field	Default & Access	Description
6	0b RO	**Parity Error Enable (PERRE)**. This function is not enabled for Device 8.
5	0b RO	**VGA Palette Snoop Enable (VGASNOOP).** This function is not enabled for Device 8.
4	0b RO	**Memory Write and Invalidate Enable (MWIE).** This function is not enabled for Device 8.
3	0b RO	**Special Cycle Enable (SCE)**. This function is not enabled for Device 8.
2	0b RO	**Bus Master Enable (BME)**. This function is not enabled for Device 8.
1	0b RO	**Memory Access Enable (MAE)**. This function is not enabled for Device 8.
0	0b RO	**I/O Access Enable (IOAE)**. This function is not enabled for Device 8.

3.10.4 PCISTS – PCI Status Register (D8:F0)

Address Offset: 06 – 07h
Access: RO
Size: 16 Bits
Default: 0080h

PCISTS is a 16-bit status register that reports the occurrence of error events on Device 8's PCI interface. Since MCH Device 8 does not physically reside on a real PCI bus, many of the bits are not implemented.

Bit Field	Default & Access	Description
15	0b RO	**Detected Parity Error (DPE)**. Not implemented for Device 8.
14	0b RO	**Signaled System Error (SSE)**. Not implemented for Device 8.
13	0b RO	**Received Master Abort Status (RMAS)**. Not implemented for Device 8.
12	0b RO	**Received Target Abort Status (RTAS)**. Not implemented for Device 8.
11	0b RO	**Signaled Target Abort Status (STAS)**. Not implemented for Device 8.
10:9	00b RO	**DEVSEL Timing (DEVT).** Device 8 does not physically connect to PCI_A. These bits are set to 00 (fast decode) so that optimum DEVSEL timing for PCI_A is not limited by the MCH.
8	0b RO	**Master Data Parity Error Detected (DPD).** Not implemented for Device 8.
7	1b RO	**Fast Back-to-Back (FB2B).** Device 8 does not physically connect to PCI_A. This bit is set to '1' (indicating fast back-to-back capability) so that the optimum setting for PCI_A is not limited by the MCH.
6:0	00h	Reserved

3.10.5 RID – Revision Identification (D8:F0)

Address Offset: 08h
Access: RO
Size: 8 Bit
Default: 09h

This register contains the revision number of the MCH Device 8.

Bit Field	Default & Access	Description
7:0	00h RO	**Revision Identification Number (RID).** This value indicates the revision identification number for the MCH Device 8. 09h = C1 stepping. 0Ah = C2 stepping.

3.10.6 SUBC – Sub-Class Code (D8:F0)

Address Offset: 0Ah
Access: RO
Size: 8 Bits
Default: 80h

Bit Field	Default & Access	Description
7:0	80h RO	**Sub-Class Code (SUBC)**. This value indicates the Sub Class Code into which the MCH Device 8 falls. 80h = Other system peripheral device

3.10.7 BCC – Base Class Code (D8:F0)

Address Offset: 0Bh
Access: RO
Size: 8 Bits
Default: 08h

Bit Field	Default & Access	Description
7:0	08h RO	**Base Class Code (BASEC)**. This value indicates the Base Class Code for the MCH Device 8. 08h = Other system peripheral device

3.10.8 HDR – Header Type (D8:F0)

Address Offset: 0Eh
Access: RO
Size: 8 Bits
Default: 00h

Bit Field	Default & Access	Description
7:0	00h RO	**PCI Header (HDR)**. The header type of the MCH Device 8. 00h = single-function device with standard header layout.

3.10.9 SVID – Subsystem Vendor Identification (D8:F0)

Address Offset: 2C – 2Dh
Access: R/WO
Size: 16 Bits
Default: 0000h

This value is used to identify the vendor of the subsystem.

The MCH treats the SVID and SID as a single 32 bit register with regard to R/WO functionality. Any time a write access to the address 2C – 2Fh occurs, regardless of byte enables, entire 32 bit register comprised of SVID and SID locks.

Bit Field	Default & Access	Description
15:0	0000h R/WO	**Subsystem Vendor ID (SUBVID)**. This field should be programmed during boot-up to indicate the vendor of the system board.

3.10.10 SID – Subsystem Identification (D8:F0)

Address Offset: 2E – 2Fh
Access: R/WO
Size: 16 Bits
Default: 0000h

This value is used to identify a particular subsystem.

The MCH treats the SVID and SID as a single 32 bit register with regard to R/WO functionality. Any time a write access to the address 2C – 2Fh occurs, regardless of byte enables, entire 32 bit register comprised of SVID and SID locks.

Bit Field	Default & Access	Description
15:0	0000h R/WO	**Subsystem ID (SUBID)**. This field should be programmed during BIOS initialization.

3.10.11 SCRUBLIM – Scrub Limit and Control Register (D8:F0)

Address Offset: C8 – CBh
Access: RO, R/W, R/WS
Size: 32 Bits
Default: 0000_0000h

This register is used to load the scrub engine with a particular limit address and other control information, such as the initialization data pattern.

Bit Field	Default & Access	Description
31	0b R/WS	**Scrublim Valid**. When this bit is written to '1' by software, the contents of this register will take affect on the next scrub address update. The hardware will clear this bit back to '0' as it loads these values, therefore software should never expect to read '1' on this bit.
30:29	00b	Reserved
28	0b R/W	**Mask periodic scrubbing**. 0 = Writes for the periodic scrubs are performed 1 = Writes for the periodic scrubs are not performed
27	0B R/W	**Mask demand scrubbing**.This bit when changed takes affect immediately and does not require the scrublim valid bit to be written to a 1. 0 = Writes for the demand scrubs are performed 1 = Writes for the demand scrubs are not performed
26:15	000h	Reserved
14:0	000h R/W	**Limit Address**. Defines address bits [34:20] to limit the top of the address range used for scrubs. These address bits require the scrublim valid bit to be written to a 1 in order to take affect on the next scrub address update.

3.10.12 SCRBADD – Scrub Address (D8:F0)

Address Offset: CC – CFh
Access: RO, R/W
Size: 32 Bits
Default: 0000_0000h

This register is used to load the scrub engine with a particular starting address. The scrub will be performed between this address and the address in the scrub limit register.

Bit Field	Default & Access	Description
31	0b R/W	**Scrbadd Valid**. When this bit is written to '1' by software the value of the address in bits 27:0 will be loaded into the scrub unit on the next scrub address update. The hardware will clear this bit back to '0' as it loads the scrub address from bits 27:0, therefore software should never expect to read '1' on this bit. **NOTE:** The address is loaded at the counter rollover time, which isn't externally visible. Therefore it is possible that after this register is written, there will be one more scrub using the internal counter address, instead of the address placed in this register. There will never be more than one 'extra' scrub before the value of this register is used.
30:29	00b	Reserved
28:0	000_000h R/W	**Starting Scrub Address**. This corresponds to address bits 34:6 of the scrub address, which is always 64B line based.

3.10.13 DTCL – DRAM Power Management Control Lower Register (D8:F0)

Address Offset: D0 – D3h
Access: RO, R/W
Size: 32 Bits
Default: 2000_0000h

DTCL and DTCU make up a conceptual 64-bit register. This thermal management mechanism is used for activates, reads, and writes. The memory subsystem uses this register to apply on a per DIMM basis. Once thermal management is invoked, no transactions are issued to the entire memory subsystem.

Bit Field	Default & Access	Description
31:30	00b	Reserved
29:28	10b R/W	**Transaction Weighting**. These two bits select the weighting between activate and read/write commands. A read or write will increment the throttle counter by 2. The activate command will increment the counter by the amount specified by this two-bit field. In a given cycle, a read or write can occur along with an activate to a different DIMM. Two different DIMM slot counters would be incremented for this cycle. Activate: read/write ratio. 00 = 2:2 (Increment by 2 for an activate, increment by 2 for a read or write) 01 = 3:2 (Increment by 3 for an activate, increment by 2 for a read or write) 10 = 4:2 (Increment by 4 for an activate, increment by 2 for a read or write) 11 = 5:2 (Increment by 5 for an activate, increment by 2 for a read or write)
27:22	00h R/W	**Thermal Management Time (TT)**. This value provides a multiplier between 0 and 63, which specifies how long thermal management remains in effect as a number of Global DRAM Sampling Windows. For example, if GDSW is programmed to 1000_0000b and TT is set to 01_0000b, then thermal management will be performed for ~8 seconds once invoked (128 • 4ms • 16).
21:15	00h R/W	**Thermal Management Monitoring Window (TMW)**. The value in this register is shifted left by 4 to specify a window of 0-2047 host clocks with a 16 clock granularity. While the thermal management mechanism is invoked, DRAM activate, reads, and writes are monitored during this window. If the weighted activity count during the window reaches the Thermal Management Activity Maximum for any DIMM slot, then requests are blocked for the remainder of the window for all DIMM slots.
14:3	000h R/W	**Thermal Management Activity Maximum (TAM)**. This value defines the maximum weighted activity count, between 0-4095, which is permitted to occur on a DIMM slot within one TMW.
2:1	00b R/W	**Thermal Management Mode (RM)**. 00 = Thermal management via counters and hardware throttle_on signal mechanisms disabled. 01 = Reserved 10 = Counter mechanism controlled through GDSW and GAT is enabled. When the threshold set in GDSW and GAT is reached, thermal management start/stop cycles occur based on the setting in TT, TMW and TAM. 11 = Reserved

Bit Field	Default & Access	Description
0	0b R/W	**Start Thermal Management (ST)**. Software writes to this bit to start and stop the write to thermal management. 0 = Write thermal management stops and the counters associated with TMW and TAM are reset. 1 = Write thermal management begins based on the settings in TMW and TAM; and remains in effect until this bit is reset to '0'.

3.10.14 DTCU – DRAM Power Management Control Upper Register (D8:F0)

Address Offset: D4 – D7h
Access: RO, R/W
Size: 32 Bits
Default: 0000_0000h

DTCL and DTCU make up a conceptual 64-bit register.

Bit Field	Default & Access	Description
31:30	00b R/W	**Thermal Management Lock (TLOCK)**. These bits secure the DRAM thermal management control registers. The bits default to '0'. Once a '1' is written to either bit, the configuration register bits in DTC become read-only, and no later writes can change the register until reset. See exception to this statement for ST below. **NOTE:** This register is not sticky. 00 = Not locked. All of the bits in DTC can be written. 01 = START Mode bits not locked. All bits in DTC(U&L) except for the ST (DTCL bit 0) is locked and cannot be written to (including bits 31:30 of this register). 10 = All bits locked. DTC is fully locked and cannot be changed, including bits 31:30 of this register. 11 = Reserved
29	0b R/W	**Thermal Management Test Mode Enable (TME)**. This bit is used to shorten test time. 0 = Normal operation 1 = Global DRAM Sampling Window, and the Global Activate Threshold are scaled down by ~1000. GDSW becomes a specification of microseconds rather than milliseconds and GAT is multiplied by 32 rather than 32k.
28:21	00h	Reserved
20:13	00h R/W	**Global DRAM Sampling Window (GDSW)**. This 8b value is multiplied by 4 to define the length of time in milliseconds (0-1020). If the weighted activity count during this window exceeds the Global Activity Threshold defined below, then the thermal management mechanism will be invoked to limit DRAM requests to a lower bandwidth checked over smaller time windows across all DIMM slots.
12:0	000h R/W	**Global Activity Threshold (GAT)**. This 13b value is multiplied by 2^{15} to arrive at the weighted activity count that must occur on a DIMM slot within the Global DRAM Sampling Window in order to cause the thermal management mechanism to be invoked.

System Address Map 4

4.1 Overview

A system based on the MCH supports up to 64 GB of host-addressable memory space and 16 KB+3 of host-addressable I/O space. The MCH supports full 36-bit addressing for both processor and I/O subsystem initiated memory space accesses, providing 64 GB of addressable space despite a main memory subsystem physically limited to 16 GB in size. The I/O and memory spaces are divided by system configuration software into non-overlapping regions. The memory ranges are useful either as system memory or as specialized memory, while the I/O regions are used solely to control the operation of devices in the System Memory Map

There are five basic regions of memory in the system: memory below 1 MB, memory between 1 MB and the Top of Low Memory (TOLM) register (see Section 3.5.32 on page 3-68), memory between the TOLM register and 4 GB, memory above 4 GB, and the high PCI memory range between the top of main memory and 64 GB. The high PCI memory range is added with the MCH, and was not available in previous generations of Intel Architecture 32-bit MCH components.

Note that the DRAM that physically overlaps the low PCI Memory Address Range (between TOLM and the 4-GB boundary) may be recovered for use by the system. For example, if there is 4 GB of physical DRAM and 1 GB of PCI space, then the system can address a total of 5 GB. In this instance the top GB of physical DRAM physically located from 3 GB to 4 GB is addressed between 4 GB and 5 GB by the system.

Figure 4-1. Basic Memory Regions

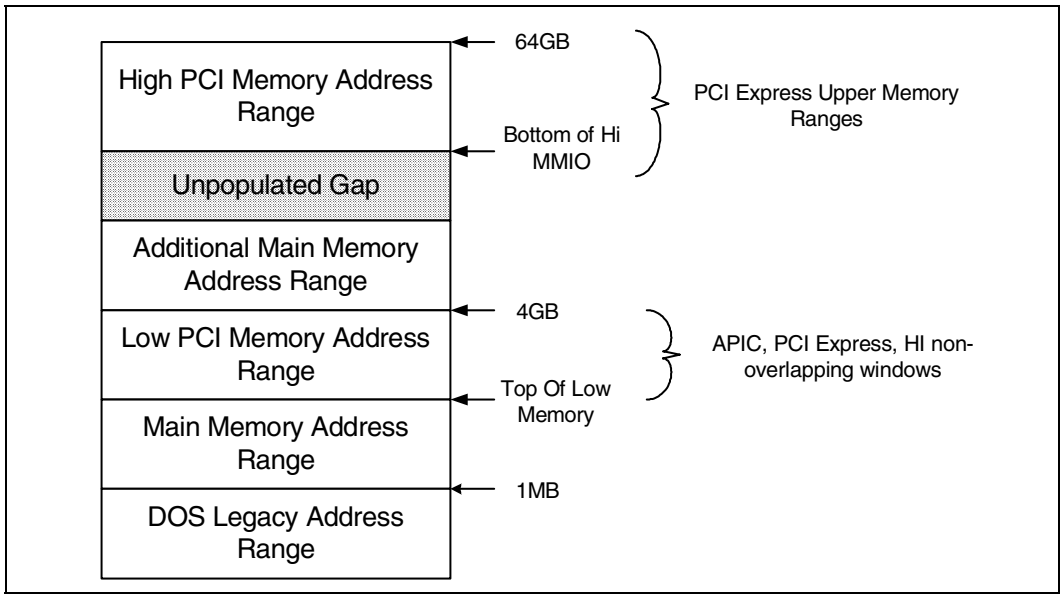

4.1.1 System Memory Spaces

The address ranges in this space are:

- DOSMEM 0_0000_0000 to 0_0009_FFFF
- MEM1_15 0_0010_0000 to 0_00EF_FFFF
- MAINMEM 0_0100_0000 to TOLM
- HIGHMEM 1_0000_0000 to 7_FFFF_FFFF

Note: Although the MCH can physically address up to 16 GB of DRAM memory, the addition of the high PCI memory range may create an unusable range if the allocated range extends below 16 GB. The MCH does not offer reclaim of this range when such overlap occurs, as supported processors are limited to 36 address bits regardless.

Note: Inbound accesses are now supported for the Extended BIOS region from 0xE0000 through 0xEFFFF. The rest of the 640KB--1MB range remains unsupported for inbound DMA -- the MCH will abort attempted accesses irrespective of PAM address access control settings.

These address ranges are always mapped to system memory, regardless of the system configuration. The Top of Low Memory (TOLM) register provides a mechanism to carve memory out of the MAINMEM segment for use by System Management Mode (SMM) hardware and software, PCI add-in devices, and other functions. The address of the highest 128 MB quantity of populated DRAM in the system is placed into the DRB7 register, which will match the Top of Memory (TOM) register when both DIMM sparing and memory mirroring are disabled. For systems with a total DRAM space and low PCI memory-mapped space of less than 4 GB, the value in TOM will match that of the TOLM register. For other memory configurations, the two are unlikely to be the same, since the PCI configuration portion of the BIOS software will program the TOLM register to the maximum value that is less than 4 GB and also allows enough room for the total memory space below 4 GB (LoPCI) allocated to populated PCI devices.

Figure 4-2. DOS Legacy Region

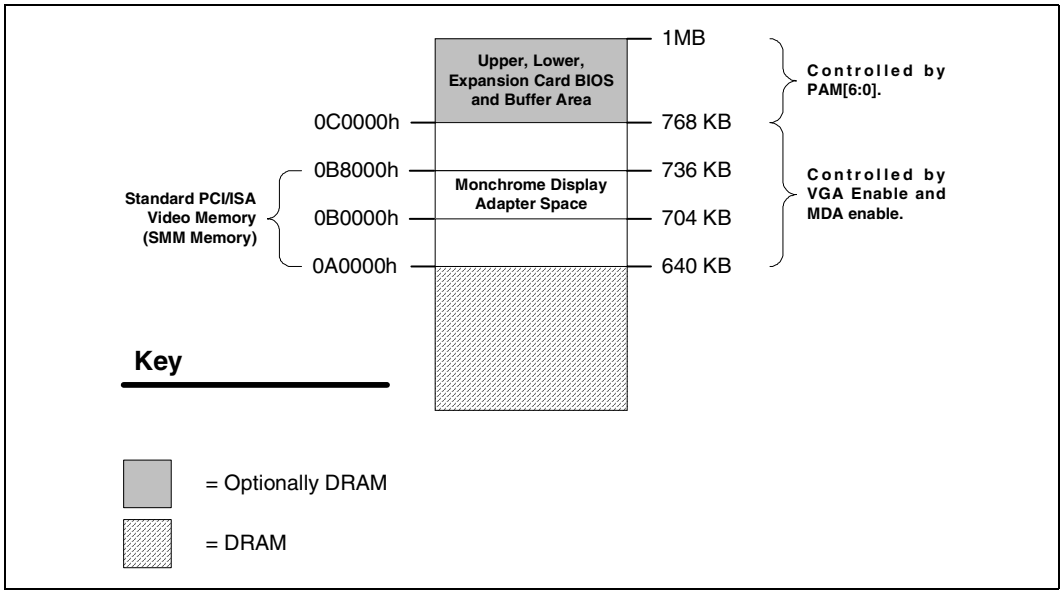

4.1.2 VGA and MDA Memory Spaces

- VGAA 0_000A_0000 to 0_000A_FFFF
- MDA 0_000B_0000 to 0_000B_7FFF
- VGAB 0_000B_8000 to 0_000B_FFFF

These legacy address ranges are used on behalf of video cards to map a frame buffer or a character-based video buffer into a dedicated location. By default, accesses to these ranges are forwarded to HI. However, if the VGA_EN bit is set in one of the BCTRL configuration registers, then transactions within the VGA and MDA spaces are sent to one of the PCI Express* interfaces.

The VGA_EN bit may be set in one and only one of the BCTRL registers. Software must not set more than one of the VGA_EN bits.

If the configuration bit EXSMRC.MDAP (see Section 3.5.29 on page 3-65) is set, then accesses that fall within the MDA range will be sent to HI without regard for the VGAEN bits. Legacy support requires the ability to have a second graphics controller (monochrome display adapter) in the system. In a MCH system with PCI graphics installed, accesses in the standard VGA range may be forwarded to any of the logical PCI Express* ports (depending on configuration bits). Since the monochrome adapter may be on the HI/PCI (or logical ISA) bus, the MCH must decode cycles in the MDA range and forward them to HI. This capability is controlled via the MDAP configuration bit. In addition to the memory range B0000h to B7FFFh, the MCH decodes I/O cycles at 3B4h, 3B5h, 3B8h, 3B9h, 3BAh and 3BFh and forwards them to HI.

An optimization allows the system to reclaim the memory displaced by these regions. If SMM memory space is enabled by SMRAM.G_SMRAME and either the SMRAM.D_OPEN bit (see Section 3.5.29 on page 3-65 and Section 3.5.30 on page 3-66) is set or the system bus receives an SMM-encoded request for code (not data), then the transaction is steered to system memory rather than HI. Under these conditions, both the VGAEN bits and the MDAP bit are ignored.

4.1.3 PAM Memory Spaces

The address ranges in this space are:

- PAMC0 0_000C_0000 to 0_000C_3FFF
- PAMC4 0_000C_4000 to 0_000C_7FFF
- PAMC8 0_000C_8000 to 0_000C_BFFF
- PAMCC 0_000C_C000 to 0_000C_FFFF
- PAMD0 0_000D_0000 to 0_000D_3FFF
- PAMD4 0_000D_4000 to 0_000D_7FFF
- PAMD8 0_000D_8000 to 0_000D_BFFF
- PAMDC 0_000D_C000 to 0_000D_FFFF
- PAME0 0_000E_0000 to 0_000E_3FFF
- PAME4 0_000E_4000 to 0_000E_7FFF
- PAME8 0_000E_8000 to 0_000E_BFFF
- PAMEC 0_000E_C000 to 0_000E_FFFF
- PAMF0 0_000F_0000 to 0_000F_FFFF

The 256-KB PAM region is divided into three parts:

- ISA expansion region, a 128-KB area between 0_000C_0000h – 0_000D_FFFFh
- Extended BIOS region, a 64-KB area between 0_000E_0000h – 0_000E_FFFFh
- System BIOS region, a 64-KB area between 0_000F_0000h – 0_000F_FFFFh.

The ISA expansion region is divided into eight, 16-KB segments. Each segment can be assigned one of four Read/Write states: read-only, write-only, read/write, or disabled. These segments are typically set to disabled for memory access, which leaves them routed to HI for ISA space.

The extended system BIOS region is divided into four, 16-KB segments. Each segment can be assigned independent read and write attributes so it can be mapped either to main DRAM or to HI. Typically, this area is used for RAM or ROM.

The system BIOS region is a single 64 KB segment. This segment can be assigned independent read and write attributes. It is by default (after reset) Read/Write disabled, and cycles are forwarded to HI. By manipulating the Read/Write attributes, the MCH can “shadow” BIOS into the main DRAM.

Note that the PAM regions are accessible from the logical PCI Express* ports. All inbound writes from any port that hit the PAM area are sent to HI, which prevents the corruption of non-volatile data shadowed in main memory. All inbound reads from any port that hit the PAM regions are harmlessly terminated internally; data are returned, but not necessarily from the requested address. Transaction routing is not hardware enforced based on the settings in the PAM registers.

4.1.4 ISA Hole Memory Space

BIOS software may optionally open a “window” between 15 MB and 16 MB (00F0_0000 to 0100_0000) that relays transactions to HI instead of completing them with a system memory access. This window is opened with the FDHC.HEN configuration field.

4.1.5 TSEG SMM Memory Space

The TSEG SMM space (TOLM to TSEG – TOLM) allows system management software to partition a region of main memory just below the top of low memory (TOLM) that is accessible only by system management software. This region may be 128 KB, 256 KB, 512 KB, or 1 MB in size, depending upon the EXSMRC.TSEG_SZ field (see Section 3.5.29 on page 3-65). This space must be below 4 GB, so is specified relative to TOLM and not relative to the top of physical memory. SMM memory is globally enabled by SMRAM.G_SMRARE. Requests may access SMM system memory when either SMM space is open (see SMRAM.D_OPEN in Section 3.5.30 on page 3-66) or the MCH receives an SMM code request on its system bus. In order to access the TSEG SMM space, the TSEG must be enabled by EXSMRC.T_EN (Section 3.5.29 on page 3-65). When all of these conditions are met, then a system bus access to the TSEG space is sent to system memory. If the high SMRAM is not enabled or if the TSEG is not enabled, then all memory requests from all interfaces are forwarded to system memory. If the TSEG SMM space is enabled, and an agent attempts a non-SMM access to TSEG space, then the transaction is specially terminated.

Hub Interface or PCI Express* accesses are not allowed to SMM space.

Figure 4-3. 1 MB through 4 GB Memory Regions

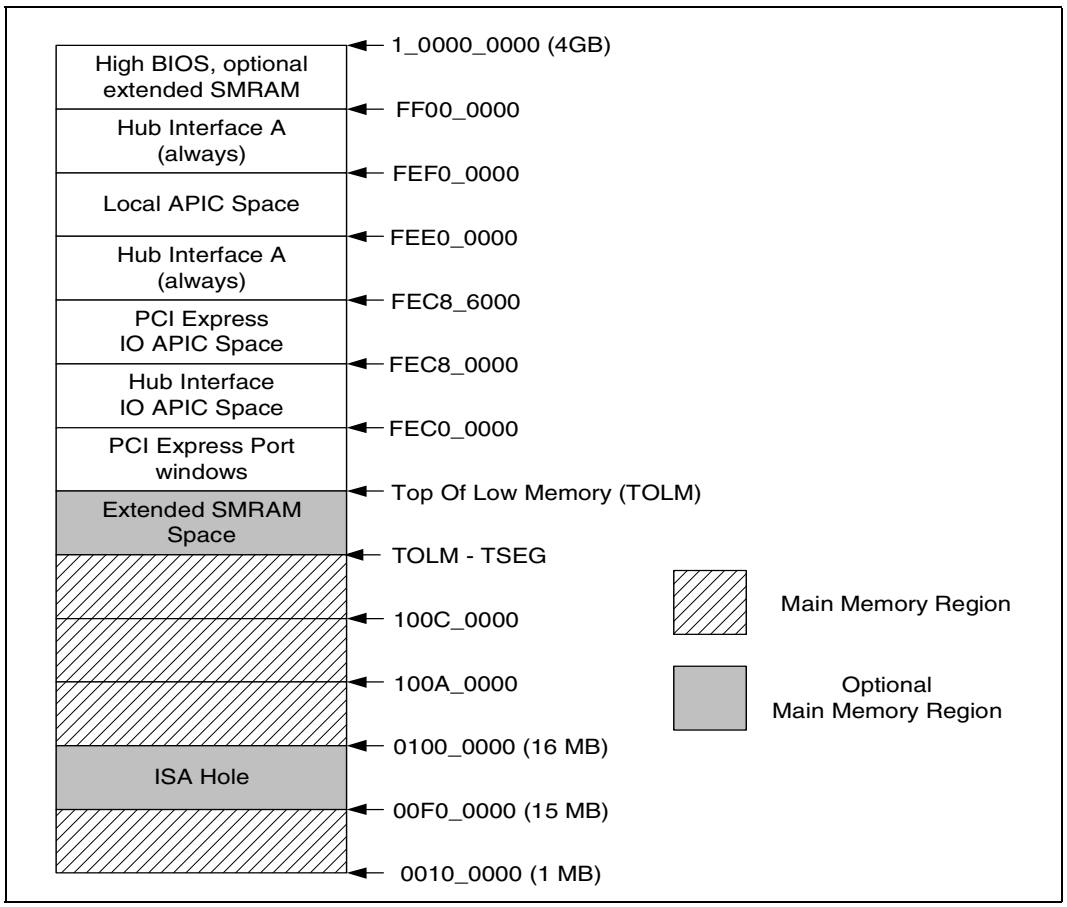

4.1.6 PCI Express* Enhanced Configuration Aperture

The address ranges in this space are:

- EXPREGION 0_E000_0000 to 0_EFFF_FFFF

PCI Express* defines a memory-mapped aperture mechanism through which to access 4KB of PCI configuration register space for each possible bus, device, and function number. This 4KB space includes the compatible 256B of register offsets that are traditionally accessed via the legacy CF8/CFC configuration aperture mechanism in I/O address space, making the enhanced configuration mechanism a full superset of the legacy mechanism. The enhanced mechanism has the advantage that full destination and type of access is specified in a single memory-mapped uncacheable transaction on the FSB, which is both faster and more robust than the historical I/O-mapped address and data register access pair.

The MCH places the enhanced configuration aperture at E000_0000h by default, as this is the first contiguous 256MB location below the 4GB boundary available for such usage.

4.1.7 I/O APIC Memory Space

The address ranges in this space are:

- IOAPIC0 (HI) 0_FEC0_0000 to 0_FEC7_FFFF
- IOAPIC2 (PCI Express A) 0_FEC8_0000 to 0_FEC8_0FFF
- IOAPIC3 (PCI Express A1) 0_FEC8_1000 to 0_FEC8_1FFF
- IOAPIC4 (PCI Express B) 0_FEC8_2000 to 0_FEC8_2FFF

The IOAPIC spaces are used to communicate with IOAPIC interrupt controllers that may be populated on the HI or PCI Express* interfaces. Since it is difficult to relocate an interrupt controller using plug-and-play software, fixed address decode regions have been allocated for them. Processor accesses to the IOAPIC0 region are always sent to HI. Processor accesses to the IOAPIC2 region are always sent to PCI Express A and so on. These regions are subject to the APIC Disable, which will be cleared by BIOS after the allocated regions have been reflected down to the base registers of APIC controllers discovered during normal enumeration. Until this step of the initialization sequence has been performed, accesses to these regions are treated as subtractive decode and routed to HI.

4.1.8 System Bus Interrupt Memory Space

The system bus interrupt space (0_FEE0_0000 to 0_FEEF_FFFF) is the address range used to deliver interrupts to the system bus. Any device below HI or a PCI Express* port may issue a Memory Write to 0FEEx_xxxxh. The MCH will forward this Memory Write along with its associated data to the system bus as a Message Signaled Interrupt (MSI) transaction. The MCH terminates the system bus transaction by asserting TRDY# and providing the response. This Memory Write cycle does not go to DRAM.

The processors may also use this region to send interprocessor interrupts (IPI) from one processor to another. MCH support for this feature includes the ability to handle redirectable MSI transactions according to values programmed into integrated task priority registers. Refer to the section on interrupt delivery in the “Architectural Features” chapter of this specification.

Reads to this address range are aborted by the MCH.

4.1.9 High SMM Memory Space

The HIGHSMM space (0_FEDA_0000 to 0_FEDB_FFFF) allows cacheable access to the compatible SMM space by remapping valid SMM accesses between 0_FEDA_0000 and 0_FEDB_FFFF to physical accesses between 0_000A_0000 and 0_000B_FFFF. The accesses are remapped when SMRAM space is enabled, an appropriate access is detected on the system bus, and when ESMRAMC.H_SMRAME allows access to high SMRAM space. Inbound SMM memory accesses from any port are specially terminated; reads are provided with data retrieved from address 0, while writes are ignored entirely (all byte enables deasserted).

4.1.10 PCI Device Memory (MMIO)

The MCH provides two distinct regions of memory that may be mapped to populated PCI devices. The first is the traditional (non-prefetchable) MMIO range, which must lie below the 4GB boundary. The registers associated with non-prefetchable MMIO (MBASE/MLIMIT) are unchanged from historical 32-bit architecture MCH implementations. The second is the

prefetchable MMIO range, which has been extended in a MCH such that it may lie on either side of the 4GB boundary. The registers associated with prefetchable MMIO (PMBASE/PMLIMIT) have been augmented by the PCI defined upper 32-bit base/limit register pair (PMBASU/PMLMTU), although only the first nibble of each register is implemented in the MCH (physical addressing is limited to 36 bits total).

The MBASE/MLIMIT pair must be programmed to lie between TOLM and 4GB. The PMBASE/PMLIMIT and PMBASU/PMLMTU registers must be programmed to lie either between TOLM and 4GB, or between the top of main memory and 64GB.

Because these registers define a PCI memory space, they are subject to the memory access enable (MAEN) control bit in the standard PCI command register.

4.1.10.1 Device 2 Memory and Prefetchable Memory

Plug-and-play software configures the PCI Express A memory window in order to provide enough memory space for the devices behind this virtual PCI-to-PCI bridge. Accesses whose addresses fall within these windows are decoded and forwarded to PCI Express A for completion. The address ranges in this space are:

- M2 MBASE2 to MLIMIT2
- PM2 PMBASE2/PMBASU2 to PMLIMIT2/PMLMTU2

Note that neither region should overlap with any other fixed or relocatable area of memory.

4.1.10.2 Device 3 Memory and Prefetchable Memory

Plug-and-play software configures the PCI Express A1 memory window in order to provide enough memory space for the devices behind this virtual PCI-to-PCI bridge. Accesses whose addresses fall within this window are decoded and forwarded to PCI Express A1 for completion. The address ranges in this space are:

- M3 MBASE3 to MLIMIT3
- PM3 PMBASE3/PMBASU3 to PMLIMIT3/PMLMTU3

Note that neither region should overlap with any other fixed or relocatable area of memory.

Note: If PCI Express A is configured to operate in x8 mode, all functional space for PCI Express A1 disappears; effectively collapsing M3/PM3 to match the limit addresses of M2/PM2.

4.1.10.3 Device 4 Memory and Prefetchable Memory

Plug-and-play software configures the PCI Express B memory window in order to provide enough memory space for the devices behind this virtual PCI-to-PCI bridge. Accesses whose addresses fall within this window are decoded and forwarded to PCI Express B for completion. The address ranges in this space are:

- M4 MBASE4 to MLIMIT4
- PM4 PMBASE4/PMBASU4 to PMLIMIT4/PMLMTU4

Note that neither region should overlap with any other fixed or relocatable area of memory.

4.1.10.4 HI Subtractive Decode

All accesses that fall between the address programmed into the TOLM register and 4 GB are subtractively decoded and forwarded to HI; that is, they will be mapped to HI if they do not positively decode to a space that corresponds to another port/device. Any gaps in the map between APIC, processor-specific, platform-specific, and compatibility PCI device memory regions fall into this category. Note that the MCH does not support subtractive decode for transactions initiated by the I/O subsystem; thus accesses which fall into this category received inbound will be specially terminated, rather than forwarded to the HI.

4.2 I/O Address Space

The MCH does not support the existence of any other I/O devices on the system bus. The MCH generates outbound transactions for all processor I/O accesses. The MCH contains two internal registers in the processor I/O space, Configuration Address register (CONF_ADDR) and the Configuration Data register (CONF_DATA). These locations are used to implement the configuration space access mechanism and are described in the Device Configuration registers section.

The processor allows 64 KB + 3 bytes to be addressed within the I/O space. The MCH propagates the processor I/O address without any translation to the targeted destination bus. Note that the upper three locations can be accessed only during I/O address wraparound when signal A16# is asserted on the system bus. A16# is asserted on the system bus whenever a DWord I/O access is made from address 0FFFDh, 0FFFEh, or 0FFFFh. In addition, A16# is asserted when software attempts a two-byte I/O access from address 0FFFFh.

All I/O accesses (read or write) that do not target the MCH's Configuration Address or Data registers will receive a Defer Response on the system bus, and be forwarded to the appropriate outbound port. The MCH never posts an I/O write.

The MCH never responds to inbound transactions to I/O or configuration space initiated on any port. Inbound I/O or configuration transactions requiring completion are terminated with "master abort" completion packets on the originating port interface. Outbound I/O or configuration write transactions not requiring completion are dropped.

4.3 System Management Mode (SMM) Space

The MCH supports the use of main memory as System Management RAM (SMM RAM) enabling the use of System Management Mode. The MCH supports three SMM options:

Compatible SMRAM (C_SMRAM)

- High Segment (HSEG)
- Top of Memory Segment (TSEG).
- System Management RAM space provides an access-protected memory area that is available for SMI handler code and data storage. This memory resource is normally hidden from the Operating System so that the processor has immediate access to this memory space upon entry to SMM (cannot be swapped out).

4.3.1 SMM Addressing Ranges

MCH provides three SMRAM options:

1. Below 1 MB option that supports compatible SMI handlers.
2. Above 1 MB option that allows new SMI handlers to execute with writeback cacheable SMRAM.
3. Optional larger writeback cacheable T_SEG area from 128 KB to 1 MB in size. The above 1-MB solutions require changes to compatible SMRAM handler code to properly execute above 1 MB.

Note that the first two options both map legal accesses to the same physical range of memory, while the third defines an independent region of addresses.

4.3.1.1 SMM Space Restrictions

If any of the following conditions are violated, the results of SMM accesses are unpredictable and may cause the system to hang:

The Compatible SMM space must not be set up as cacheable.

- Both D_OPEN and D_CLOSE must not be set to ‘1’ at the same time.
- When TSEG SMM space is enabled, the TSEG space must not be reported to the Operating System as available DRAM. This is a BIOS responsibility.

BIOS and SMM code must cooperate to properly configure the MCH in order to ensure reliable operation of the SMM function.

4.3.1.2 SMM Space Definition

SMM space is defined by both its addressed SMM space and its DRAM SMM space. The addressed SMM space is defined as the range of system bus addresses used by the processor to access SMM space. DRAM SMM space is defined as the range of physical DRAM memory locations containing SMM information.

SMM space can be accessed at one of three transaction address ranges: Compatible, High and TSEG. The Compatible and TSEG SMM space not remapped; therefore the addressed and DRAM SMM physical addresses are identical. The High SMM space is remapped; thus the addressed and DRAM SMM locations are different. Note that the High DRAM space is the same as the Compatible Transaction Address space.

Table 4-1 describes all three unique addressing combinations:

- Compatible Transaction Address
- High Transaction Address
- TSEG Transaction Address

Table 4-1. Supported SMM Ranges

SMM Space Enabled	Transaction Address Space	DRAM Space
Compatible (C)	A0000h to BFFFFh	A0000h to BFFFFh
High (H)	0FEDA0000h to 0FEDBFFFFh	A0000h to BFFFFh
TSEG (T)	(TOLM-TSEG_SZ) to TOLM	(TOLM-TSEG_SZ) to TOLM

NOTES:
1. High SMM: This implementation is consistent with the Intel® E7500 and Intel® E7501 MCH designs. Prior to these designs, the High segment was the 384 KB region from A_0000h to F_FFFFh. However, the C_0000h to F_FFFFh was not practically useful, so it has been deleted from the MCH design.
2. TSEG SMM: This implementation is consistent with the Intel E7500 and Intel E7501 MCH designs. Prior to these designs, the TSEG address space was offset by 256 MB to allow for simpler decoding, and the TSEG was remapped to just under the TOLM. In the MCH the TSEG region is not offset by 256 MB, and it is not remapped.

4.4 Memory Reclaim Background

The following Memory Mapped I/O devices are typically located below 4 GB:

- High BIOS
- H-Seg
- XAPIC
- Local APIC
- System Bus Interrupts
- PCI Express* M, PM and BAR Regions

In previous generation MCH architectures, the physical DRAM memory overlapped by the logical address space allocated to these Memory Mapped I/O devices was unusable. In server systems the memory allocated to memory mapped I/O devices could easily exceed 1 GB. The result is that a large amount of physical memory populated in the system would not be usable.

The MCH provides the capability to reclaim the physical memory overlapped by the Memory Mapped I/O logical address space. The MCH remaps physical memory from the Top of Low Memory (TOLM) boundary up the 4-GB boundary (or DRB7 if less than 4 GB) to an equivalent sized logical address range located just above the top of physical memory.

4.4.1 Memory Remapping

An incoming address (referred to as a logical address) is checked to see if it falls in the memory remap window. The bottom of the remap window is defined by the value in the REMAPBASE register. The top of the remap window is defined by the value in the REMAPLIMIT register. An address that falls within this window is remapped to the physical memory starting at the address defined by the TOLM register.

Functional Description 5

This chapter provides a functional description of the MCH chipset architecture. Coverage includes the MCH interface units (system bus, system memory, PCI Express*, Hub Interface (HI), SMBus, power management, MCH clocking, MCH system reset and power sequencing) as well as reliability and manageability features.

5.1 Internal Feature Set

5.1.1 Coherent Memory Write Buffer

The MCH integrates a coherent write buffer sized for 16, 64-byte cache lines (a total of 1 KB of storage). This feature enables the MCH to optimize memory read latency, allowing reads to pass less critical writes en route to the main memory store. The write buffer includes a CAM structure to enforce ordering among conflicting accesses to the same cache line, as well as to provide for read service from the write cache. In the latter case, the access to the main memory store never occurs, which both improves latency and conserves bandwidth on the memory interface.

Note that the write buffer is capable of servicing processor read requests directly via a “hit” to the internal location containing the data without initiation of any DDR subsystem accesses. Inbound read requests which “hit” the write buffer result in a flush of the target data, followed by retrieval via an external read request. The complexity of direct service for inbound read requests is not warranted given the extremely modest hit rate expected for a 1-KB structure running workstation I/O workloads.

5.1.2 Internal Data Protection

Due to the nature of having various data protection schemes on the different interfaces (ECC, parity, and CRC) it is necessary to be able to convert between them when transferring data internally. To accomplish this, protection of internal data is done with parity.

5.2 Front Side Bus (FSB)

The MCH supports either single or dual processor population. System bus implementation of arbitration and transaction identification allows for processors to operate in either single or dual-threaded modes (Hyper-Threading technology) for a maximum of four logical processors.

The MCH supports a base system bus frequency of 200 MHz. The address and request interface is double pumped to 400 MHz while the 64-bit data interface (+ parity) is quad pumped to 800 MHz. This provides a matched system bus address and data bandwidth of 6.4 GB/s.

5.2.1 In-Order Queue (IOQ)

The MCH has a 12 deep IOQ. This means that it does not need to limit the number of simultaneous outstanding transactions by asserting BNR#.

5.2.2 System Bus Interrupts

Intel Xeon processors support FSB interrupt delivery. They do not support the APIC serial bus interrupt delivery mechanism. Interrupt related messages are encoded on the FSB as “Interrupt Message Transactions”. In the MCH, platform FSB interrupts may originate from one of the processors on the FSB (IPIs – inter-processor interrupts), or from a downstream device on either the Hub Interface (HI) or one of the PCI Express* ports. In the latter case the MCH drives the “Interrupt Message Transaction” onto the FSB.

In the IOxAPIC environment an interrupt is generated from the IOxAPIC to a processor in the form of an upstream Memory Write. In the MCH environment, the ICH and PXH contain IOxAPICs, and their interrupts are generated as upstream HI/PCI Express* Memory Writes. Furthermore, PCI 2.3 defines MSIs (Message Signaled Interrupts) that are also in the form of Memory Writes. A PCI 2.3 device may generate an interrupt as an MSI cycle on its PCI bus instead of asserting a hardware signal to the IOxAPIC. The MSI may be directed to the IOxAPIC, which in turn generates an interrupt as an upstream HI/PCI Express* Memory Write. Alternatively the MSI may be directed directly to the FSB. The target of an MSI is dependent on the address of the interrupt Memory Write. The MCH forwards inbound HI/PCI Express* Memory Writes to address 0FEEx_xxxxh to the FSB as “Interrupt Message Transactions”.

The MCH supports redirecting Lowest Priority delivery mode interrupts to the processor that is executing the lowest priority task thread. The MCH redirects interrupts based on the task priority status of each processor thread. The task priority of each processor thread is periodically downloaded to the MCH via the xTPR (Task Priority Register) Special Transaction. The MCH redirects HI/PCI Express* and PCI originated interrupts as well as IPIs.

The MCH also broadcasts EOI cycles generated by a processor downstream to the HI/PCI Express* interfaces.

5.2.3 System Bus Dynamic Inversion

The MCH supports Dynamic Bus Inversion (DBI) both when driving and when receiving data from the system bus. DBI limits the number of data signals that are driven to a low voltage on each quad pumped data phase. This decreases the power consumption of the MCH. DBI[3:0]# indicate if the corresponding 16 bits of data are inverted on the bus for each quad pumped data phase. See Table 5-1 for DBI to data bit mapping.

Table 5-1. DBI Signals to Data Bit Mapping

DBI[3:0]#	Data Bits
DBI0#	HD[15:0]
DBI1#	HD[31:16]
DBI2#	HD[47:32]
DBI3#	HD[63:48]

When the MCH drives data, each 16-bit segment is analyzed. The corresponding DBI# signal asserts depending on the number of active data lines and the data path through the MCH.

- When data is driven by the MCH from memory to the FSB, the corresponding DBI# signal asserts and data inverts if ***eight or more*** of the sixteen signals would normally be driven low. Otherwise, the group is not inverted.
- For all other data driven by the MCH to the FSB, the corresponding DBI# signal asserts and data inverts if ***more than eight*** of the sixteen signals would normally be driven low. Otherwise, the group is not inverted.

This behavior ensures that the MCH drives low no more than eight data bits within a sixteen bit group at any given time.

When the processor drives data, the MCH monitors DBI[3:0]# to determine if the corresponding data segment requires inversion.

5.2.4 Front Side Bus Parity

Address/request, response, and data bus signals are protected by parity. The address/request and data busses can be configured to perform no error checking.

The FSB data parity scheme is not straightforward parity, so a description is warranted. Note that the data in a given clock cycle is quad-pumped while the parity that corresponds to this same data is common-clocked driven in the clock following the presentation of the data. The four subphases correspond to the data and data inversion bits that will be driven out during a single clock cycle. The corresponding parity bits are calculated by XORing the four table components. For example; DP3# which will be driven out in the clock following the data is an XOR of DP3a, DP3b, DP3c, and DP3d. As the Table 5-2 shows these four parity components are on different rows of the table. Ultimately an approximate 32 bytes of data is protected by four parity bits. This particular rotating parity scheme is able to detect stuck at faults more easily.

Table 5-2. FSB Parity Matrix

Data Signals	Subphase			
	1	2	3	4
D[15:0]#, DBI0#	DP3a	DP2b	DP1c	DP0d
D[31:16]#, DBI1#	DP0a	DP3b	DP2c	DP1d
D[47:32]#, DBI2#	DP1a	DP0b	DP3c	DP2d
D[63:48]#, DBI3#	DP2a	DP1b	DP0c	DP3d

5.3 Memory Interface

The MCH provides an integrated memory controller for direct connection to two channels of registered DDR333 or DDR2-400 memory (single or dual ranked). Peak theoretical memory data bandwidth using DDR333 technology is 5.33 GB/S. For DDR2-400 technology, this increases to 6.4 GB/s.

The MCH supports single channel operation using either of its memory channels. This yields a minimum supported memory size for DDR of 128 MB by using a single, 128 Mb technology DIMM (nine x8 devices). For DDR2, the minimum supported memory size is 256 MB, achieved by using as single, 256 Mb technology DIMM (nine x8 devices). The MCH DDR2 memory interface does not support 128 Mb technology.

The MCH memory interface implements a 14-bit address bus. ECC support allows for standard SEC-DED ECC as well as x4 SDDC. Although both ECC modes may be simultaneously disabled via the DRAM Controller Mode register, the MCH only supports ECC DIMMs. Additionally, the MCH supports 4 KB, 8 KB, 16 KB, 32 KB, and 64 KB page sizes.

When both DDR channels are populated and operating, they function in lock-step mode. For the MCH, the maximum supported DDR333 memory configuration is 16 GB obtained from different combinations of single and dual ranked x4, 1GB technology DIMMs (limit of up to three DIMMs and four ranks per channel). Note that only BGA DRAMs are supported at DDR333. The maximum supported DDR2-400 memory configuration is 16 GB using different combinations of single and dual ranked, x4, 1GB technology DIMMs (limit of up to four ranks per channel).

The operational frequency of the memory interface is related to that of the system bus by a gearing ratio, thus it is not limited to lock-step operation at the same frequency. This interface supports DDR333 or DDR2-400 data rates, yielding data bandwidths of 2.67 GB/s and 3.2 GB/s per channel, respectively.

The MCH supports a burst length of 4 whether in single or dual channel mode. In dual channel mode this results in eight 64-bit chunks (64-byte cache line) from a single read or write. In single channel mode two reads or writes are required to access a cache line of data.

Table 5-3. Memory Interface Capacities

Memory Interface Capacities (MB's)		128 Mb		256 Mb		512 Mb		1 Gb	
		Min	Max	Min	Max	Min	Max	Min	Max
DDR333	1 channel	128	1024	256	2048	512	4096	1024	8192
	2 channels	256	2048	512	4096	1024	8192	2048	16384
DDR2-400	1 channel	N/A	N/A	256	2048	512	4096	1024	8192
	2 channels	N/A	N/A	512	4096	1024	8192	2048	16384

5.3.1 Memory Interface Performance Optimizations

5.3.1.1 DDR Overlapped Command Scheduling

The MCH memory controller command scheduler overlaps scheduling of activate and precharge commands for successive accesses, thus hiding the latency of these events. An activate command may be issued prior to the current command completion depending upon the bank addressed. It may be issued during the Trcd or Tcl wait periods. Overlap Command Scheduling is enabled in the DRC register.

5.3.1.2 Aggressive Page-Closed Policy with Look-Ahead

The MCH uses an aggressive page closed policy with a single entry look-ahead. All commands are issued with auto-precharge unless a look-ahead indicates that a page-hit is already scheduled. The normal operating case is "page-empty". The "page-hit" case will occur as often as the

inbound/outbound and memory control arbiters are able to forward sequential requests from the same source back to back out to the DDR interface, making this the next most frequent occurrence. A “page-miss” case is rare in the MCH. Page closing policies are adjusted in the DRC register.

5.3.1.3 Symmetric Addressing Mode

The MCH automatically enables symmetric address bit permuting when precisely four identical ranks of memory per channel are available. When other than four identical ranks are available, the MCH operates in “non symmetric” addressing mode. The MCH always utilizes system address bits 8 and 9 to select among the first four banks on a DIMM. When symmetric bit permuting is enabled, the MCH also maps system address bits 13 and 14 to the chip-selects of the four symmetric ranks. This has the effect of spreading traffic with spatial locality across more of the populated DRAM, minimizing the occurrence of page misses, and allowing better access pipelining.

5.3.2 Memory Interface Feature Set

5.3.2.1 DRAM ECC – Intel® x4 Single Device Data Correction (x4 SDDC)

The DRAM interface uses two different ECC algorithms. The first is a standard SEC/DED ECC across a 64-bit data quantity. The second ECC method is a distributed, 144-bit S4EC-D4ED mechanism, which provides x4 SDDC protection for DIMMs that utilize x4 devices. Bits from x4 parts are presented in an interleaved fashion such that each bit from a particular part is represented in a different ECC word. DIMMs that use x8 devices, can use the same algorithm but will not have x4 SDDC protection, since at most only four bits can be corrected with this method. The algorithm does provide enhanced protection for the x8 parts over a standard SEC-DED implementation. With two memory channels, either ECC method can be utilized with equal performance, although single-channel mode only supports standard SEC/DED.

The x4 SDDC ECC feature is not supported in single-channel operation as it has significant performance impacts in that environment.

5.3.2.2 Integrated Memory Scrub Engine

The MCH includes an integrated engine to walk the populated memory space proactively seeking out soft errors in the memory subsystem. In the case of a single bit correctable error, this hardware detects, logs, and corrects the data except when an incoming write to the same memory address is detected. For any uncorrectable errors detected, the scrub engine logs the failure. Both types of errors may be reported via multiple alternate mechanisms under configuration control. The scrub hardware will also execute “demand scrub” writes when correctable errors are encountered during normal operation (on demand reads, rather than scrub-initiated reads). This functionality provides incremental protection against time-based deterioration of soft memory errors from correctable to uncorrectable.

Using this method, a 16GB system can be completely scrubbed in less than one day. (The effect of these scrub writes do not cause any noticeable degradation to memory bandwidth, although they will cause a greater latency for that one very infrequent read that is delayed due to the scrub write cycle.)

Note that an uncorrectable error encountered by the memory scrub engine is a “speculative error.” This designation is applied because no system agent has specifically requested use of the corrupt data, and no real error condition exists in the system until that occurs. It is possible that the error

resides in an unmodified page of memory that will be simply dropped on a swap back to disk. Were that to occur, the speculative error would simply “vanish” from the system undetected without adverse consequences.

5.3.2.3 Retry on Uncorrectable Error

The MCH includes specialized hardware to resubmit a memory read request upon detection of an uncorrectable error. When a demand fetch (as opposed to a scrub) of memory encounters an uncorrectable error as determined by the enabled ECC algorithm, the memory control hardware will cause a (single) full resubmission of the entire cache line request from memory to verify the existence of corrupt data. This feature is expected to greatly reduce or eliminate the reporting of false or transient uncorrectable errors in the DRAM array.

Note that any given read request will only be retried a single time on behalf of this error detection mechanism. Any uncorrectable error encountered on the retry will be logged and escalated as directed by device configuration. This feature may be enabled and disabled via configuration.

5.3.2.4 Integrated Memory Initialization Engine

The MCH provides hardware managed ECC auto-initialization of all populated DRAM space under software control. Once internal configuration has been updated to reflect the types and sizes of populated DIMM devices, the MCH will traverse the populated address space initializing all locations with good ECC. This not only speeds up the mandatory memory initialization step, but also frees the processor to pursue other machine initialization and configuration tasks.

Additional features have been added to the initialization engine to support high-speed population and verification of a programmable memory range with one of four known data patterns (0/F, A/5, 3/C, and 6/9). This function facilitates a limited, very high-speed memory test, as well as provides a BIOS accessible memory zeroing capability for use by the operating system.

5.3.2.5 DIMM Sparing Function

To provide a more fault tolerant system, the MCH includes specialized hardware to support fail-over to a spare DIMM device in the event that a primary DIMM in use exceeds a specified threshold of runtime errors. One of the DIMMs installed per channel will not be used, but kept in reserve. In the event of significant failures in a particular DIMM, it and its corresponding partner in the other channel (if applicable), will, over time, have its data copied over to the spare DIMM(s) held in reserve. When all the data has been copied, the reserve DIMM(s) will be put into service, and the failing DIMM(s) will be removed from service. Only one sparing cycle is supported. If this feature is not enabled, then all DIMMs will be visible in normal address space.

The rationale behind this feature is that a failing DIMM will experience an increasing error frequency prior to catastrophic failure, and a controller capable of detecting the increasing failure rate could take action to circumvent downtime.

Hardware additions for this feature include the implementation of tracking register per DIMM to maintain a history of error occurrence, and a programmable register to hold the fail-over error threshold level. See Section 5.3.2.6 for more details on this feature.

The fail-over mechanism is slightly more complex. Once fail-over has been initiated the MCH must execute every write twice; once to the primary DIMM, and once to the spare. (This requires that the spare DIMM be at least the size of the largest primary DIMM in use.) The MCH will also begin tracking the progress of its built-in memory scrub engine. Once the scrub engine has covered

every location in the primary DIMM, the duplicate write function will have copied every data location to the spare. At that point, the MCH can switch the spare into primary use, and take the failing DIMM off-line.

Hardware will detect the threshold initiating fail-over, and escalate the occurrence of that event as directed (signal an SMI, generate an interrupt, or wait to be discovered via polling). Whatever software routine responds to the threshold detection must select a victim DIMM (in case multiple DIMMs have crossed the threshold prior to sparing invocation) and initiate the memory copy. Hardware will automatically isolate the "failed" DIMM once the copy has completed. The data copy is accomplished by address aliasing within the DDR control interface, thus it does not require reprogramming of the DRAM row boundary (DRB) registers, nor does it require notification to the operating system that anything untoward has occurred in memory.

DDR:
The spare DIMM should have greater than or equal to the number of ranks as any other DIMM. In other words, the spare DIMM may be a single rank DIMM only if all other DIMMs are single rank. In addition, the rank size of the spare DIMM must be greater than or equal to the largest rank size of all other DIMMs.

DDR2:
The number of ranks/DIMM is irrelevant for DDR2 DIMM sparing. However, the size of the spare rank must be greater than or equal to the largest rank size of all other ranks.

5.3.2.6 DIMM Error Rate Threshold Counters

The error rate threshold counters of the memory subsystem utilize a simple leaky bucket counter in order to flag excessive error rates rather than a total error count. The operational model is straightforward: set the fail-over threshold register to a non-zero value in the THRESH_DED or THRESH_SEC0-3 registers to enable the feature. If the count of errors on any DIMM exceeds the threshold after a programmable time period, an error will occur and sparing fail-over will commence. The counter registers themselves are implemented as "leaky buckets," such that they do not contain an absolute cumulative count of all errors since power-on; rather, they contain an aggregate count of the number of errors received over a running time period. The "drip rate" of the bucket is selectable by software via the SPARECTL register, so it is possible to set the threshold to a value that will never be reached by a "healthy" memory subsystem experiencing the rate of errors expected for the size and type of memory devices in use.

5.3.2.6.1 Error Threshold Methodology

Figure 5-1 illustrates the Error Threshold methodology with the assumption that errors are accumulated at the Expected Average Rate (EAR). The EAR is defined as the rate at which a properly functioning DIMM accumulates errors. At the end of the first time period, the sum of the current errors plus the residue from past errors, zero in this case, is then divided by 2 to become the resultant value used for the comparison. The resultant value becomes the residue for the next time period (second in this example). At the end of the second time period, the counter contains the sum of new errors and the non-zero error residue. This, in turn, is divided by 2 to become the new resultant value for comparison and so on.

Figure 5-2 shows the comparison error count at the end of each time period. Assuming a counter accumulates errors at the average rate, the counter will reach a steady state after approximately five time periods. Given this, setting the threshold value at or just above the EAR will result in a threshold event if a significant surge in errors occurs.

Figure 5-1. Error Ramp Rate

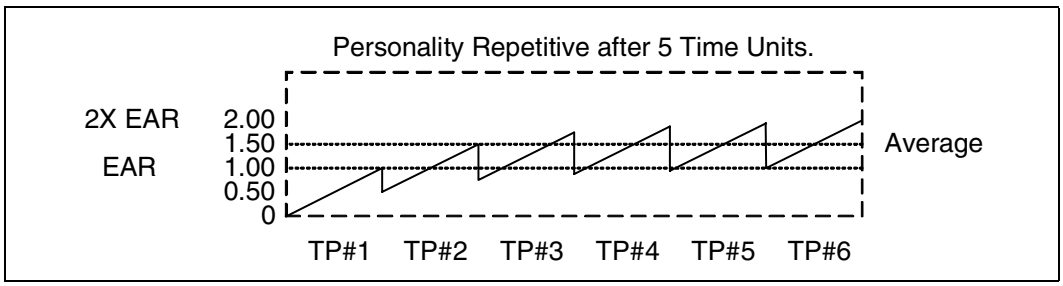

Figure 5-2. Error Count For Comparison

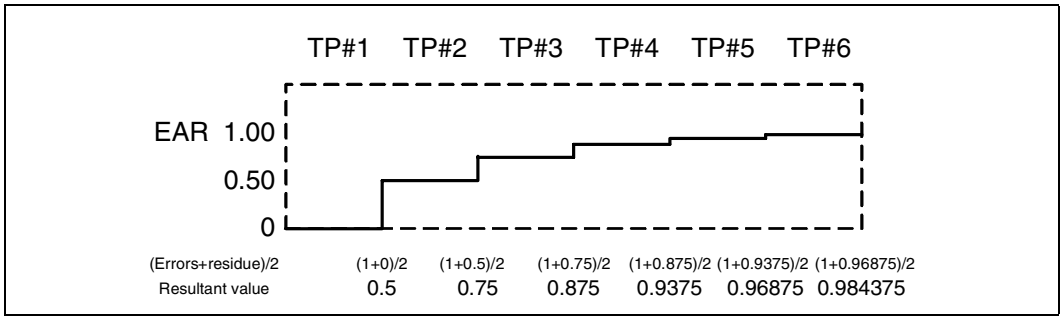

This methodology implies that the first non-zero comparison will occur only after the first time period has expired. Thus, an infinite time unit (bucket set to “never leak”) cannot generate a threshold event regardless of how many errors have been counted.

5.3.2.6.2 Error Reporting

At the expiration of a time period, all counters, SEC and DED, are halved and compared against their respective SEC and DED threshold registers. When an error counter exceeds its threshold, a per DIMM flag bit is set in the DIMM_THR_EX register. This one register can be read to determine if any DIMMs had an unexpected number of errors. Once software clears a flag bit in the DIMM_THR_EX register, the threshold detect hardware is not rearmed for that DIMM until such time that the error count decays to a value that matches or is less than the threshold value.

5.3.2.6.3 Sparing Implications

After a sparing operation has been performed, the spare DIMM will inherit the same SEC and DED counters utilized by the failing DIMM. In order to avoid an immediate threshold exceeded error on the new DIMM due to the error residue remaining in the counter from the failing DIMM, it is necessary to clear the counter by adjusting the time period via the SPARECTL register to the smallest non-zero value (1uS) and running for at least 10 time periods. The time period can then be adjusted back to its desired value. The side affect of this is that the error accumulations for all other DIMMs will essentially be cleared out as well, since they all use the same time period mechanism.

5.3.3 Memory Address Translation Tables

The following tables define the address bit translation from the system address to the DRAM row/column/bank address. Note that the count of DRAM devices per DIMM listed in the following tables excludes the devices utilized for ECC information, thus the actual DRAM counts in MCH platforms are expected to be 9 devices per rank of x8 technology, and 18 devices per rank of x4 technology.

Table 5-4. Dual Channel Non-Symmetric Address Map

Tech	Configuration (No. of RAMS / DIMM)	Row Size Page Size	R/C/B	ADDR		BA2	BA1	BA0	A13	A12	A11	A10	A9	A8	A7	A6	A5	A4	A3	A2	A1	A0
128Mb DDR	4Meg x 8 x 4 bks	256MB	12 x 10 x 2	Row	27	"0-"	9	8			26	25	24	23	22	21	20	19	18	17	16	15
	8	16KB		Col								AP	27	14	13	12	11	10	7	6	5	'0'
128Mb DDR	8Meg x 4 x 4 bks	512MB	12 x 11 x 2	Row	28	'0'	9	8			26	25	24	23	22	21	20	19	18	17	16	15
	16	32KB		Col							28	AP	27	14	13	12	11	10	7	6	5	'0'
256Mb DDR/ DDR2	8Meg x 8 x 4 bks	512MB	13 x 10 x 2	Row	28	'0'	9	8		27	26	25	24	23	22	21	20	19	18	17	16	15
	8	16KB		Col								AP	28	14	13	12	11	10	7	6	5	'0'
256Mb DDR/ DDR2	16Meg x 4 x 4 bks	1024MB	13 x 11 x 2	Row	29	'0'	9	8		27	26	25	24	23	22	21	20	19	18	17	16	15
	16	32KB		Col							28	AP	29	14	13	12	11	10	7	6	5	'0'
512Mb DDR	16Meg x 8 x 4 bks	1024MB	13 x 11 x 2	Row	29	'0'	9	8		27	26	25	24	23	22	21	20	19	18	17	16	15
	8	32KB		Col							28	AP	29	14	13	12	11	10	7	6	5	'0'
512Mb DDR	32Meg x 4 x 4 bks	2048MB	13 x 12 x 2	Row	30	'0'	9	8		27	26	25	24	23	22	21	20	19	18	17	16	15
	16	64KB		Col						30	28	AP	29	14	13	12	11	10	7	6	5	'0'
512Mb DDR2	16Meg x 8 x 4 bks	1024MB	14 x 10 x 2	Row	29	'0'	9	8	28	27	26	25	24	23	22	21	20	19	18	17	16	15
	8	16KB		Col								AP	29	14	13	12	11	10	7	6	5	'0'
512Mb DDR2	32Meg x 4 x 4 bks	2048MB	14 x 11 x 2	Row	30	'0'	9	8	30	27	26	25	24	23	22	21	20	19	18	17	16	15
	16	32KB		Col							28	AP	29	14	13	12	11	10	7	6	5	'0'
1Gb DDR	32Meg x 8 x 4 bks	2048MB	14 x 11 x 2	Row	30	'0'	9	8	28	27	26	25	24	23	22	21	20	19	18	17	16	15
	8	32KB		Col							30	AP	29	14	13	12	11	10	7	6	5	'0'
1Gb DDR	64Meg x 4 x 4 bks	4096MB	14 x 12 x 2	Row	31	'0'	9	8	28	27	26	25	24	23	22	21	20	19	18	17	16	15
	16	64KB		Col						30	31	AP	29	14	13	12	11	10	7	6	5	'0'
1Gb DDR2	16Meg x 8 x 8 bks	2048MB	14 x 10 x 3	Row	30	30	9	8	28	27	26	25	24	23	22	21	20	19	18	17	16	15
	8	16KB		Col								AP	29	14	13	12	11	10	7	6	5	'0'
1Gb DDR2	32Meg x 4 x 8 bks	4096MB	14 x 11 x 3	Row	31	30	9	8	28	27	26	25	24	23	22	21	20	19	18	17	16	15
	16	32KB		Col							31	AP	29	14	13	12	11	10	7	6	5	'0'

Table 5-5. Single Channel Non-Symmetric Address Map

Tech	Configuration (No. of RAMS / DIMM)	Row Size Page Size	R/C/B	ADDR		BA2	BA1	BA0	A13	A12	A11	A10	A9	A8	A7	A6	A5	A4	A3	A2	A1	A0
128Mb DDR	4Meg x 8 x 4 bks	128MB	12 x 10 x 2	Row	26	"0-"	9	8			25	24	23	22	21	20	19	18	17	16	15	14
	8	8KB		Col								AP	26	13	12	11	10	7	6	5	'0'	'0'
128Mb DDR	8Meg x 4 x 4 bks	2562MB	12 x 11 x 2	Row	27	'0'	9	8			25	24	23	22	21	20	19	18	17	16	15	14
	16	16KB		Col							27	AP	26	13	12	11	10	7	6	5	'0'	'0'
256Mb DDR/ DDR2	8Meg x 8 x 4 bks	256MB	13 x 10 x 2	Row	27	'0'	9	8		26	25	24	23	22	21	20	19	18	17	16	15	14
	8	8KB		Col								AP	27	13	12	11	10	7	6	5	'0'	'0'
256Mb DDR/ DDR2	16Meg x 4 x 4 bks	512MB	13 x 11 x 2	Row	28	'0'	9	8		26	25	24	23	22	21	20	19	18	17	16	15	14
	16	16KB		Col							27	AP	28	13	12	11	10	7	6	5	'0'	'0'
512Mb DDR	16Meg x 8 x 4 bks	512MB	13 x 11 x 2	Row	28	'0'	9	8		26	25	24	23	22	21	20	19	18	17	16	15	14
	8	16KB		Col							27	AP	28	13	12	11	10	7	6	5	'0'	'0'
512Mb DDR	32Meg x 4 x 4 bks	1024MB	13 x 12 x 2	Row	29	'0'	9	8		26	25	24	23	22	21	20	19	18	17	16	15	14
	16	32KB		Col						29	27	AP	28	13	12	11	10	7	6	5	'0'	'0'
512Mb DDR2	16Meg x 8 x 4 bks	512MB	14 x 11 x 2	Row	28	'0'	9	8	27	26	25	24	23	22	21	20	19	18	17	16	15	14
	8	8KB		Col								AP	28	13	12	11	10	7	6	5	'0'	'0'
512Mb DDR2	32Meg x 4 x 4 bks	1024MB	14 x 12 x 2	Row	29	'0'	9	8	29	26	25	24	23	22	21	20	19	18	17	16	15	14
	16	16KB		Col							27	AP	28	13	12	11	10	7	6	5	'0'	'0'
1Gb DDR	32Meg x 8 x 4 bks	1024MB	14 x 11 x 2	Row	29	'0'	9	8	27	26	25	24	23	22	21	20	19	18	17	16	15	14
	8	16KB		Col							29	AP	28	13	12	11	10	7	6	5	'0'	'0'
1Gb DDR	64Meg x 4 x 4 bks	2048MB	14 x 12 x 2	Row	30	'0'	9	8	27	26	25	24	23	22	21	20	19	18	17	16	15	14
	16	32KB		Col						29	30	AP	28	13	12	11	10	7	6	5	'0'	'0'
1Gb DDR2	16Meg x 8 x 4 bks	1024MB	14 x 10 x 3	Row	29	29	9	8	27	26	25	24	23	22	21	20	19	18	17	16	15	14
	8	8KB		Col								AP	28	13	12	11	10	7	6	5	'0'	'0'
1Gb DDR2	32Meg x 4 x 8 bks	2048MB	14 x 11 x 3	Row	30	29	9	8	27	26	25	24	23	22	21	20	19	18	17	16	15	14
	16	16KB		Col							30	AP	28	13	12	11	10	7	6	5	'0'	'0'

Table 5-6. Dual Channel Symmetric Address Map

Tech	Configuration (No. of RAMS / DIMM)	Row Size Page Size	R/C/B	ADDR		BA2	BA1	BA0	A13	A12	A11	A10	A9	A8	A7	A6	A5	A4	A3	A2	A1	A0
128Mb DDR	4Meg x 8 x 4 bks	256MB	12 x 10 x 2	Row	29	"0-"	9	8			26	25	24	23	22	21	20	19	18	17	16	15
	8	16KB		Col								AP	29	28	27	12	11	10	7	6	5	'0'
128Mb DDR	8Meg x 4 x 4 bks	512MB	12 x 11 x 2	Row	30	'0'	9	8			26	25	24	23	22	21	20	19	18	17	16	15
	16	32KB		Col							30	AP	29	28	27	12	11	10	7	6	5	'0'
256Mb DDR/ DDR2	8Meg x 8 x 4 bks	512MB	13 x 10 x 2	Row	30	'0'	9	8		29	26	25	24	23	22	21	20	19	18	17	16	15
	8	16KB		Col								AP	30	28	27	12	11	10	7	6	5	'0'
256Mb DDR/ DDR2	16Meg x 4 x 4 bks	1024MB	13 x 11 x 2	Row	31	'0'	9	8		29	26	25	24	23	22	21	20	19	18	17	16	15
	16	32KB		Col							30	AP	31	28	27	12	11	10	7	6	5	'0'
512Mb DDR	16Meg x 8 x 4 bks	1024MB	13 x 11 x 2	Row	31	'0'	9	8		29	26	25	24	23	22	21	20	19	18	17	16	15
	8	32KB		Col							30	AP	31	28	27	12	11	10	7	6	5	'0'
512Mb DDR	32Meg x 4 x 4 bks	2048MB	13 x 12 x 2	Row	32	'0'	9	8		29	26	25	24	23	22	21	20	19	18	17	16	15
	16	64KB		Col						32	30	AP	31	28	27	12	11	10	7	6	5	'0'
512Mb DDR2	16Meg x 8 x 4 bks	1024MB	14 x 10 x 2	Row	31	'0'	9	8	30	29	26	25	24	23	22	21	20	19	18	17	16	15
	8	16KB		Col								AP	31	28	27	12	11	10	7	6	5	'0'
512Mb DDR2	32Meg x 4 x 4 bks	2048MB	14 x 11 x 2	Row	32	'0'	9	8	32	29	26	25	24	23	22	21	20	19	18	17	16	15
	16	32KB		Col							30	AP	31	28	27	12	11	10	7	6	5	'0'
1Gb DDR	32Meg x 8 x 4 bks	2048MB	14 x 11 x 2	Row	32	'0'	9	8	30	29	26	25	24	23	22	21	20	19	18	17	16	15
	8	32KB		Col							32	AP	31	28	27	12	11	10	7	6	5	'0'
1Gb DDR	64Meg x 4 x 4 bks	4096MB	14 x 12 x 2	Row	33	'0'	9	8	30	29	26	25	24	23	22	21	20	19	18	17	16	15
	16	64KB		Col						32	33	AP	31	28	27	12	11	10	7	6	5	'0'
1Gb DDR2	16Meg x 8 x 8 bks	1024MB	14 x 10 x 3	Row	32	32	9	8	30	29	26	25	24	23	22	21	20	19	18	17	16	15
	8	16KB		Col								AP	31	28	27	12	11	10	7	6	5	'0'
1Gb DDR2	32Meg x 4 x 8 bks	2048MB	14 x 11 x 3	Row	33	32	9	8	30	29	26	25	24	23	22	21	20	19	18	17	16	15
	16	32KB		Col							33	AP	31	28	27	12	11	10	7	6	5	'0'

5.3.4 Quad Word Ordering

The ordering of quad words to memory is dependant upon the mode of memory. The figures below itemize the ordering.

Figure 5-3. Dual Channel Memory Read

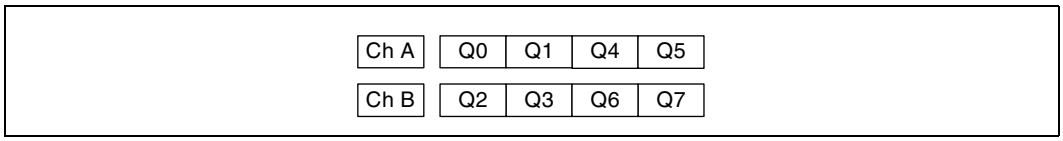

Figure 5-4. Single Channel Memory Read

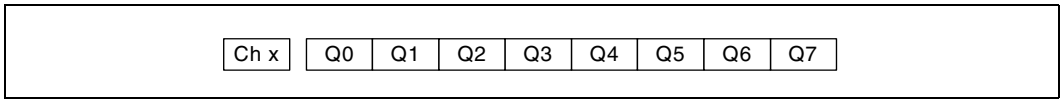

5.3.5 VOX Calibration

The MCH uses one of the 4 clock pairs on each DDR channel to properly calibrate the voltage crossing. The clocks used to determine VOX crossing are as follows: DDRA_CMDCLK1/1# for channel A and DDRB_CMDCLK0/0# for channel B. The MCH requires that these clocks be enabled to perform VOX calibration.

Figure 5-7 and Figure 5-8 enumerate the clock to DIMM assignments for channel A and B respectively for a 4 DIMM per channel implementation. For a three or two DIMM per channel, use the same assignments, but leave the DIMMs closest to the MCH a no connect.

Table 5-7. DDR Channel A Clock to DIMM assignment

Clock Signal	Associated DIMM
DDRA_CMDCLK0	DIMM1 (closest to MCH)
DDRA_CMDCLK3	DIMM2
DDRA_CMDCLK2	DIMM3
DDRA_CMDCLK1	DIMM4 (farthest from MCH)

Table 5-8. DDR Channel B Clock to DIMM assignment

Clock Signal	Associated DIMM
DDRB_CMDCLK3	DIMM1 (closest to MCH)
DDRB_CMDCLK1	DIMM2
DDRB_CMDCLK2	DIMM3
DDRB_CMDCLK0	DIMM4 (farthest from MCH)

5.3.6 Thermal Management

The MCH provides a DDR thermal management method that selectively reduces accesses to system memory when the access rate crosses the predetermined programmable threshold. In effect, this mechanism enforces a “traffic limit” to any given DIMM slot; thereby providing a

hardware-enforced ceiling on power dissipation. While available, DDR2 is not expected to require activation of the thermal management mechanism under any environmental conditions, as dissipation is significantly lower at 1.8V than at 2.5V.

5.3.6.1 Thermal Management Algorithmic Description

The goal of thermal management is to keep DRAM devices below their maximum specified case temperatures. To this end, the MCH algorithm is fairly straightforward. The amount of activity per DIMM is tracked per fixed window of time, and compared against a programmed threshold value set by the most restrictive DIMM. If any DIMM goes over the budget, the thermal management mechanism engages, limiting traffic to the memory array at large for fixed period of time to avoid a thermal spike. Once the thermal management period has expired, the mechanism is rearmed to look for further threshold crossings. Figure 5-5 provides an overview of how thermal management is implemented in the MCH. A single threshold register is implemented to reduce overhead, and to align with the "big hammer" nature of thermal management enforcement. Given that the design cannot freely reorder traffic to avoid the "hot" DIMM, thermal management simply limits all memory traffic while active.

Figure 5-5. Memory Thermal Management Operation

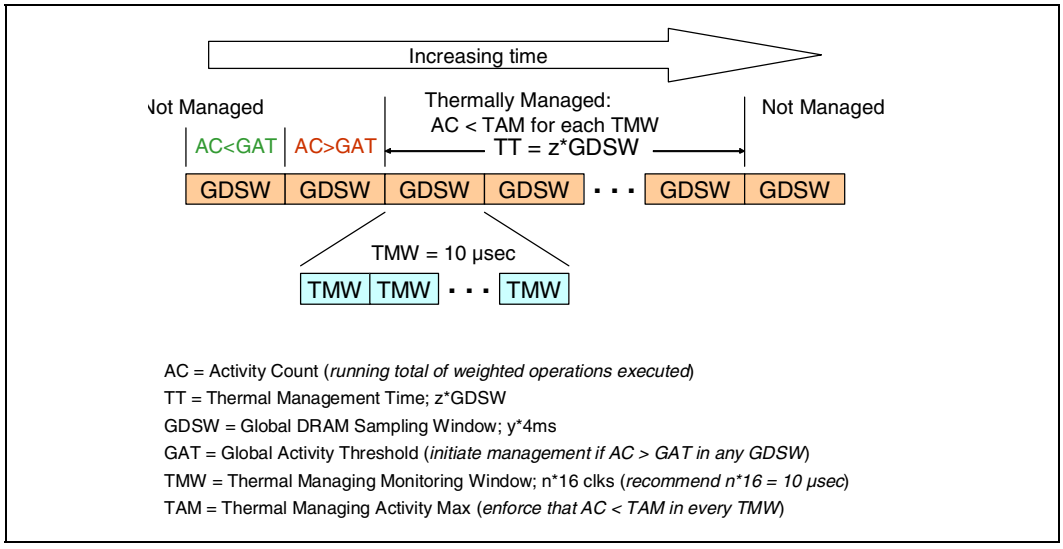

Referring to Figure 5-5, the various sampling and monitoring windows and thresholds are programmable settings in the DRAM thermal management control registers of the MCH (DTCL, DTCU). Fields in these registers define the "z, y, and n" parameters, as well as the activity count thresholds, and are set by firmware prior to enabling the thermal management feature.

5.3.6.2 Thermal Management Threshold Calculation

The dissipation on a given DIMM is a combination of endemic dissipation when powered-on, and dissipation due to access activity (refresh, activates, and read/write commands). The thermal management threshold is a measure of the amount of activity supported over a fixed period of time in addition to the "idle" dissipation for that DIMM. To calculate that budget, assumptions must be made about the maximum allowable case temperature, the inlet air temperature, and the effective "theta-CA" coefficient for the DRAM devices in the target chassis. (The theta-CA value denotes degrees at the case per Watt of dissipated power – better cooling yields a lower coefficient, and a higher budget for activity.)

The idle dissipation for a given DIMM depends upon the number and type of DRAM devices it carries; i.e: technology, x8 vs. x4 devices, single vs. dual rank, operating frequency, and device density. The total current drawn by the device can be calculated from DRAM datasheets. Typical numbers were used to generate the values used in these calculations.

The activity-based dissipation is the total of activates and commands issued over the monitored period, weighted for their relative power consumption. For most devices, activates burn more power than read and write commands, and the relative dissipation changes with device density. As a result, the MCH supports a programmable weight factor for activates vs. commands (DTCL register, bits 29:28).

So the threshold is just a measure of the number of "active" clocks per monitor period supported by the power budget after idle dissipation is taken into account. The weighting factor for activates vs. commands allows us to create a fixed "active current" for each non-idle clock, which in turn reduces the threshold value to the weighted number of active clocks allowed during the thermal management monitoring period. The equation to calculate the threshold value then looks like this:

Threshold = (current_allowed – idle_current) • (clocks_per_period) / active_current

The current_allowed parameter is set by max case temperature (85°C), less the inlet temperature (Ti), divided by the product of theta-CA and the operating voltage. ((85 – Ti)/(theta-CA • 2.5) for DDR). As an example, if the inlet temperature (at the DIMM, not at the chassis grille) is 40°C, and the effective theta-CA is 5°/W, the resulting current_allowed per DIMM is 3.6A before case temperature hits max.

The idle_current and active_current are derived from DRAM device data-sheet values. Active current is a composite of activate current and read command current, normalized for the weighting factor used for activates. (i.e: Calculated value of datasheet read plus activate current, divided by weight total to get a generic current value per active clock. The hardware weighting algorithm will convert the generic active clock value back to the data sheet values on the fly. For a 3:2 ratio, the normalized active current is 1/5 of the sum of activate current and read command current.)

Clocks per period is the number of DDR clocks in the global monitoring window (GDSW). The GDSW is programmable in 4ms increments, with a range from 0 (disabled) up to 1.02 seconds. Assuming the minimum 4ms GDSW, that would be 667K clocks per window for DDR333, and 533K for DDR266. The trade-off for GDSW size has two dimensions. Under normal operation the GDSW period is non-managed, so its duration limits how long DRAM is run "too hot" before thermal management commences, which would lead to a desire to keep GDSW on the smaller end of the scale. But thermal effects have a fair amount of "inertia" because it takes time for a DRAM package to change temperature, which leads to a desire to operate in fairly large time increments so that thermal management is effective. The thermal management time (TT) field provides some capability to play both ends against the middle, because the duration of thermal management is a 6-bit multiple of the GDSW. So it is possible to monitor GDSW at 16 ms granularity, and still use thermal management for an entire second once the threshold has been crossed.

Once thermal management has been invoked, MCH hardware will enforce that no DIMM receives more activity than that allowed by the value programmed into the "Thermal Management Activity Maximum" (TAM) field in any given "Thermal Management Monitor Window" (TMW). The threshold is again tracked per DIMM, so the total amount of memory activity may exceed the TAM value, but once any DIMM exceeds TAM, no traffic will be sent to any DIMM until the current TMW expires.

Putting all of that together, a sample GAT calculation for single-rank, x4, 512 Mb DDR333, assuming a 16 ms GDSW, an inlet temperature of 45°C, and a theta-CA of 4.0, would yield the following:

Threshold(16ms) = ((85 – 45)/(4.0 • 2.5)) – 1.890) • (2666,667) / 2.185 = 2575134 = 04E hex

A little more math: assuming that each cache line access took one activate and one read (in dual-channel mode) each cache line then costs 5 off the threshold), and translates this threshold value into an available bandwidth of around 2 GB/sec. Thermal Management would be rare under these circumstances as only a very well behaved traffic pattern could pull that kind of bandwidth out of a single DIMM pair.

5.4 PCI Express* Interface

The MCH utilizes the PCI Express* high-speed serial interface to allow for superior I/O bandwidth. The MCH is compatible with the *PCI Express Interface Specification, Rev 1.0a.* A PCI Express* port is defined as a collection of lanes. Each lane consists of two striped differential pairs in each direction (transmit and receive). Thus, a x8 port would contain eight transmit signal pairs and eight receive signal pairs.

The raw bit-rate on the data pins of 2.5 Gbit/s, results in a real bandwidth per pair of 250 MB/s given the 8/10 encoding used to transmit data across this interface. The result is a maximum theoretical realized bandwidth on a x8 PCI Express* port of 2 GB/s in each direction or an aggregate of 4 GB/s. On a x16 port, the maximum theoretical realized bandwidth is 4 GB/s in each direction for an aggregate 8 GB/s.

The MCH provides one x16 and one configurable x8 PCI Express* port. The x8 PCI Express* port may alternatively be configured as two independent x4 PCI Express* ports. An interface routed for x16 operation may train with a link width of either x16 or x1. An interface routed for x8 operation may train with a link width of either a x8, x4 or x1. A x4 routed interface may train as either x4 or x1 Note that this does not in any way imply a capability for the MCH to support more than three independent PCI Express* ports of any width simultaneously, nor does it imply that the remaining lane of a potential x4 port is useful once the associated link has been established for x1operation.

The MCH supports PCI Express* lane reversal at all native widths, and reversed x4 training on any x8 port.

5.4.1 PCI Express* Training

To establish a connection between PCI Express* endpoints, they both participate in a sequence of steps known as training. This sequence will establish the operational width of the link as well as adjust skews of the various lanes within a link so that the data sample points can correctly take a data sample off of the link. In the case of a x8 port, the x4 link pairs will first attempt to train independently, and will collapse to a single link at the x8 width upon detection of a single device returning link ID information upstream. Once the number of links has been established, they will negotiate to train at the highest common width, and will step down in its supported link widths in order to succeed in training. The ultimate result may be that the link trains as a x1 link. Although the bandwidth of this link size is substantially lower than a x8 link or even a x4 link, it will allow communication between the two devices. Software will then be able to interrogate the device at the other end of the link to determine why it failed to train at a higher width (something that would not be possible without support for the x1 link width). It should be noted that width negotiation is only done during training or retraining, but not recovery.

5.4.2 PCI Express* Retry

The PCI Express* interface incorporates a link level retry mechanism. The hardware detects when a transmission packet is corrupted and performs a retry of that particular packet and all following packets. Although this will cause a temporary interruption in the delivery of packets, it does so in order to maintain the link integrity.

5.4.3 PCI Express* Link Recovery

When enough errors occur, the hardware may determine that the quality of the connection is in question, and the end points can enter what amounts to a quick training sequence known as recovery. The width of the connection will not be renegotiated, but the adjustment of skew between lanes of the link may occur. This occurs without any software intervention, but the software may be notified.

5.4.4 PCI Express* Data Protection

The PCI Express* high-speed serial interface makes use of traditional CRC protection. The data packets will utilize a 32-bit CRC protection scheme, specifically the same CRC-32 used by Ethernet – 0x04C11DB7. The smaller link packets will utilize a 16-bit CRC scheme. Since packets utilize 8B/10B encoding, and not all encodings are used; this provides further data protection, as illegal codes can be detected. Also, if errors are detected on the reception of data packets due to various transients, these data packets can be retransmitted. Hardware logic will support this link-level retry without software intervention.

5.4.5 PCI Express* Retrain

If the hardware is unable to perform a successful recovery, then the link will automatically revert to the polling state, and initiate a full retraining sequence. This is a drastic event with an implicit reset to the downstream device and all subordinate devices, and is logged by the MCH as a “Link Down” error. If escalation of this event is enabled, software is notified of the link DL_DOWN condition. Once software is involved, data will likely be lost, and processes will needed to be restarted, but this is still preferred to having to take the system down, or go offline for an extended period of time.

5.5 Hub Interface 1.5

The MCH interfaces with the ICH via a dedicated Hub Interface 1.5 supporting a peak bandwidth of 266 MB/s using a x4 base clock of 66 MHz.

This 8-bit Hub Interface (HI) provides a parity protection scheme for the data signals. The MCH HI can be configured to either perform or not perform error checking on incoming data, heading towards the MCH. When enabled, the MCH will mark incoming data as “poisoned” before passing it on to its destination when parity errors are detected on the HI. Additionally, the MCH HI logic has the capability to prevent the propagation of outbound data that has been marked as “poisoned” at one of the other MCH interfaces when parity errors were detected on incoming data. See Section 5.10, “Exception Handling” on page 5-236 for a complete discussion of the MCH error handling capabilities.

5.6 MCH Clocking

Figure 5-7 details the operational frequency range for all MCH interfaces, and specifies all supported relative interface frequencies. Operation of the MCH outside the frequency domains stipulated here may result in system instability and failure.

Table 5-9. MCH Clocking Interfaces

Interface	Clock Type	Mnemonic	Fmin	Fmax	Comments
System Bus (host)	Differential	HCLKIN	100 MHz	200 MHz	Input. Address at 2x, Data at 4x
DDR (mem)	Differential	CMDCLK	100 MHz	200 MHz	Output, geared to System Bus.
Hub Interface (ICH)	Single-ended	HICLK	50 MHz	66 MHz	Input. Asynchronous to System Bus
PCI Express*	Differential	EXP_CLK	100 MHz	100 MHz	Input. SSC tolerant, transmit at 25x, asynchronous to System Bus
JTAG	Single-ended	TCLK	10 kHz	16.7 MHz	Input. TAP clock, asynchronous to System Bus
SMBus	Single-ended	SMBSCL	10 kHz	100 kHz	Input. Asynchronous to System Bus

External clock synthesizers and/or distribution buffers are responsible for generating the differential host clock, PCI Express* reference clock, HI clock, SMBus clock, and the JTAG clock (when utilized). All these clocks are phase-independent, and no mode is provided within the MCH to recover the added latency incurred from crossing the resultant asynchronous boundary within the design (i.e., if system design steps are taken to reduce the number of oscillators and thereby provide phase-alignment on otherwise independent interfaces, no improvement in latency will result).

5.6.1 DDR Geared Clocking

The MCH will generate, fan-out, and phase-align the DDR clock at either a gear ratio of 5:6 (DDR333) or 1:1 (DDR2-400) from the host clock reference (assuming 200 MHz FSB reference frequency). The ratio in use is under software control, and must be set properly via configuration register accesses during MCH initialization (accessible from host, JTAG, and SMBus alike). Note that an incorrect setting of the DDR gearing ratio will result in unreliable memory operation.

5.6.2 PCI Express* Clocking

The PCI Express* clocking solution is in many ways analogous to that used by Infiniband (IBA) technology. The transceiver and directly associated hardware operate on a very high-speed clock derived by multiplying a reference by 25 (although IBA systems commonly use a multiple of 20), while the internal hardware that transfers information to and from those transceivers operates at one tenth of the transmit frequency. Also like IBA, the agents at either end of the bus may operate on references generated by completely independent oscillators with no specified phase relationship to each other, provided the frequencies are identical within tight tolerances. (This clocking relationship has been designated “plesiochronous” within this document.) In both PCI Express*

and IBA, the receiver extracts the (remote) clock from the data bit stream, and utilizes specialized hardware to then move the resultant data back into the (local) transmit clock domain for transfer into the core of the device.

PCI Express* differs from IBA in that spread-spectrum clocking (SSC) is allowed in cases where the reference frequency is distributed across both ends of the port (rather than independent as just described). The reference frequencies supplied to all endpoints in the PCI Express* subsystem must be generated from the same oscillator when SSC is desired, such that the frequency drift introduced is identical at both ends of every link. The MCH introduces this limitation by sharing a single reference across all its PCI Express* interfaces, thus requiring the platform designer to provide a board level solution that distributes clocks to all endpoints attached to the MCH such that each reference meets the jitter requirements stipulated in the *PCI Express Interface Specification, Rev 1.0a*. Note that the "plesiochronous" nature of the interface provides a break on distribution skew, as the platform architect need not maintain any particular phase relationship between reference clocks delivered to the various PCI Express* endpoints.

The MCH supports a completely independent reference clock input for its PCI Express* interfaces, allowing for either SSC or non-SSC platform solutions. This also allows the FSB and PCI Express* frequencies to be varied independent of one another, which can be very useful in system debug situations. If any attached PCI Express* device in a MCH-based platform operates on its own local reference, then SSC must be disabled throughout the PCI Express* subsystem.

5.6.3 Spread-Spectrum Clocking Limitations

A fairly expensive clock generation and distribution network is required if SSC is desired throughout the platform, including the PCI Express* subsystem. The clock generator must produce 200MHz reference frequencies for the processor and the MCH, as well as very low jitter 100MHz reference frequencies for the MCH and each of the PCI Express* endpoint devices. Given a fully populated performance MCH platform, this implies either a clock generation component with relatively high pin-count (over 100 pins), or a more modest generator with a somewhat exotic independent low jitter differential clock buffer component.

The MCH can tolerate a non-SSC PCI Express* subsystem, where multiple endpoints receive independent and potentially differing reference frequencies. However, the MCH cannot tolerate a "split personality" PCI Express* subsystem. Either all PCI Express* endpoints utilize the same reference as the MCH, and SSC may be enabled; or all PCI Express* endpoints are assumed to be independent, and the entire PCI Express* subsystem must operate without SSC. (It is strictly prohibited to create a configuration in which the two ends of any given PCI Express* link receive independent SSC enabled reference clocks.)

5.7 System Reset

The MCH is the root of the I/O subsystem tree, and is therefore responsible for general propagation of system reset throughout the platform. The MCH must also facilitate any specialized synchronization of reset mechanisms required by the various system components.

5.7.1 MCH Reset Types

The MCH differentiates among five types of reset as defined in table Table 5-10.

Table 5-10. MCH Reset Classes

Type	Mechanism	Effect / Description
Power-Good	**PwrGd** input pin	Propagated throughout the system hierarchy. Resets all logic and state machines, and initializes all registers to their default states (sticky and non-sticky). Tri-states all MCH outputs, or drives them to "safe" levels.
Hard	**RSTIN#** input pin, Configuration Write	Propagated throughout the system hierarchy. Resets all logic and state machines, and initializes all non-sticky registers to their default states. Tri-states all MCH outputs, or drives them to "safe" levels.
Processor-only	Configuration Write	Propagated to all processors via the CPURST# pin on the FSB. The MCH does not undergo an internal reset.
Targeted	Configuration Write	Propagated down the targeted PCI Express* port hierarchy. Treated as a "Hard" reset by all affected components, clearing all machine state and non-sticky configuration registers.
BINIT#	Internal error handling propagated via FSB **BINIT#** pin	Propagated to all FSB attached components (the MCH and up to two processors). Clears the IOQ, and resets all FSB arbiters and state machines to their default states. Not recoverable.

5.7.1.1 Power-Good Mechanism

The initial boot of a MCH platform is facilitated by the Power-Good mechanism. The voltage sources from all platform power supplies are routed to a system component which tracks them as they ramp-up, asserting the platform "PwrGd" signal a fixed interval (nominally 2mS) after the last voltage reference has stabilized.

Both the MCH and the ICH receive the system PwrGd signal via dedicated pins as an asynchronous input, meaning that there is no assumed relationship between the assertion or deassertion of PwrGd and any system reference clock. When PwrGd is deasserted all platform subsystems are held in their reset state. This is accomplished by various mechanisms on each of the different interfaces. The MCH will hold itself in a power-on reset state when PwrGd is deasserted. The ICH is expected to assert its PCIRST# output and maintain its assertion for 1mS after power is good. The PCIRST# output from ICH is expected to drive the RSTIN# input pin on the MCH, which will in turn hold the processor complex in reset via assertion of the CPURST# FSB signal.

The PCI Express* attached devices and any hierarchy of components underneath them are held in reset via implicit messaging across the PCI Express* interface. The MCH is the root of the hierarchy, and will not engage in link training until power is good, the internal "hard" reset has deasserted, and the port has been turned on.

A PwrGd reset will clear all internal state machines and logic, and initialize all registers to their default states, including "sticky" error status bits that are persistent through all other reset classes. To eliminate potential system reliability problems, all devices are also required to either tri-state their outputs or to drive them to "safe" levels during a power-on reset.

The only system information that will "survive" a PwrGd reset is battery-backed or otherwise non-volatile storage (Flash, ROM, PROM, etc.).

5.7.1.2 Hard Reset Mechanism

Once the platform has been booted and configured, a full system reset may still be required to recover from system error conditions related to various device or subsystem failures. The “hard” reset mechanism is provided to accomplish this recovery without clearing the “sticky” error status bits useful to track down the cause of system reboot.

A hard reset is typically initiated by the ICH component via the PCIRST# output pin, which is commonly connected directly to the MCH RSTIN# input pin. The ICH may be caused to assert PCIRST# via both software and hardware mechanisms. Refer to the *Intel® 82801ER I/O Controller Hub 5-R (ICH5R) Datasheet* for details. The MCH will recognize a hard reset any time RSTIN# is asserted while PwrGd remains asserted.

The MCH will propagate a hard reset to the FSB and to all subordinate PCI Express* subsystems. The FSB components are reset via the CPURST# signal, while the PCI Express* subsystems are reset implicitly when the root port links are taken down.

A hard reset will clear all internal state machines and logic, and initialize all “non-sticky” registers to their default states. Note that although the error registers will remain intact to facilitate root-cause of the hard reset, the MCH platform in general will require a full configuration and initialization sequence to be brought back on-line (all other volatile configuration information is lost).

5.7.1.3 Processor-Only Reset Mechanism

For power management and other reasons, the MCH supports a targeted processor only reset semantic. This mechanism was added to the platform architecture to eliminate double-reset to the system at large when reset-signaled processor information (such as clock gearing selection) must be updated during initialization bringing the system back to the S0 state after power had been removed from the processor complex.

5.7.1.4 Targeted Reset Mechanism

A targeted reset may be requested by setting bit 6 (Secondary Bus Reset) of the Bridge Control Register (offset 3Eh) in the target root port device. This reset will be identical to a general hard reset from the perspective of the destination PCI Express* device; it will not be differentiated at the next level down the hierarchy. Sticky error status will survive in the destination device, but software will be required to fully configure the port and all attached devices once reset and error interrogation have completed. After clearing bit 6, software may determine when the downstream targeted reset has effectively completed by monitoring the state of bit 1 (Link Active) of the VS_STS1 register (offset 47h) in the target root port device. This bit will remain deasserted until the link has regained “link up” status, which implies that the downstream device has completed any internal and downstream resets, and successfully completed a full training sequence.

Under normal operating conditions it should not be necessary to initiate targeted resets to downstream devices, but the mechanism is provided to recover from combinations of fatal and uncorrectable errors which compromise continued link operation.

5.7.1.5 BINIT# Mechanism

The BINIT# mechanism is provided to facilitate processor handling of system errors which result in a hang on the FSB. MCA code responding to an error indication (typically IERR# or MCERR#) will attempt to interrogate the MCH for error status, and if that FSB transaction fails to complete the processor will automatically time out and respond by initiating a BINIT# sequence on the FSB.

When BINIT# is asserted on the FSB, all agents (the MCH and all CPUs) are required to reset their internal FSB arbiters and all FSB tracking state machines and logic to their default states. This will effectively "un-hang" the bus to provide a path into chipset configuration space. Note that the MCH and PXH devices implement "sticky" error status bits, providing the platform software architect with free choice between BINIT# and a general hard reset to recover from a hung system.

Although BINIT# will not clear any configuration status from the system, it is not a recoverable event from which the platform may continue normal execution without first running a hard reset cycle. To guarantee that the FSB is cleared of any hang condition, the MCH will clear all pending transaction state within its internal traffic structures. This applies to outstanding FSB cycles as required, but also to in-flight memory transactions and inbound transactions. The resulting state of the platform will be highly variable depending upon what precisely got wiped-out due to the BINIT# event, and it is not possible for hardware to guarantee that the resulting state of the machine will support continued operation. What the MCH can guarantee is that no subordinate device has been reset due to this event (PCI Express* links remain "up"), and that no internal configuration state (sticky or otherwise) has been lost. The MCH will also continue to maintain main memory via the refresh mechanism through a BINIT# event, thus machine-check software will have access not only to machine state, but also to memory state in tracking-down the source of the error.

5.7.2 Power Sequencing Requirement

There are 3.3V decoupling structures in the miscellaneous I/O pads which create a parasitic diode path that will be forward-biased if the core 1.5V supply is brought up while the 3.3V supply remains at ground. For this reason, the 3.3V supply must be brought up before the 1.5V core supply, or at least the 3.3V supply must not be allowed to get more than 0.5V below the level of the 1.5V core supply. This sequencing requirement must also be enforced when the two supplies are ramping down as they shut off.

5.7.3 Reset Sequencing

Figure 5-6 contains a timing diagram illustrating the progression through the power-on reset sequence. This is intended as a quick reference for system designers to clarify the requirements of the MCH.

Note the breaks in the clock waveform at the top of Figure 5-6, which are intended to illustrate further elapsed time in the interest of displaying a lengthy sequence in a single picture. Each of the delays in the reset sequence is of fixed duration, enforced by either the MCH or the ICH. In the case of a power-on sequence, the MCH internal "hard" and "core" resets deassert simultaneously. The two lines marked with names beginning "HLA" illustrate the Hub Interface special cycle handshake between the MCH and the ICH to coordinate across the deasserting edge of the CPURST# output from the MCH.

Table 5-11 summarizes the durations of the various reset stages illustrated above, and attributes the delays to the component that enforces them.

The fixed delays provide time for subordinate PLL circuitry to lock on interfaces where the clock is withheld or resynchronized during the reset sequence.

Figure 5-6. Power-On Reset Sequence

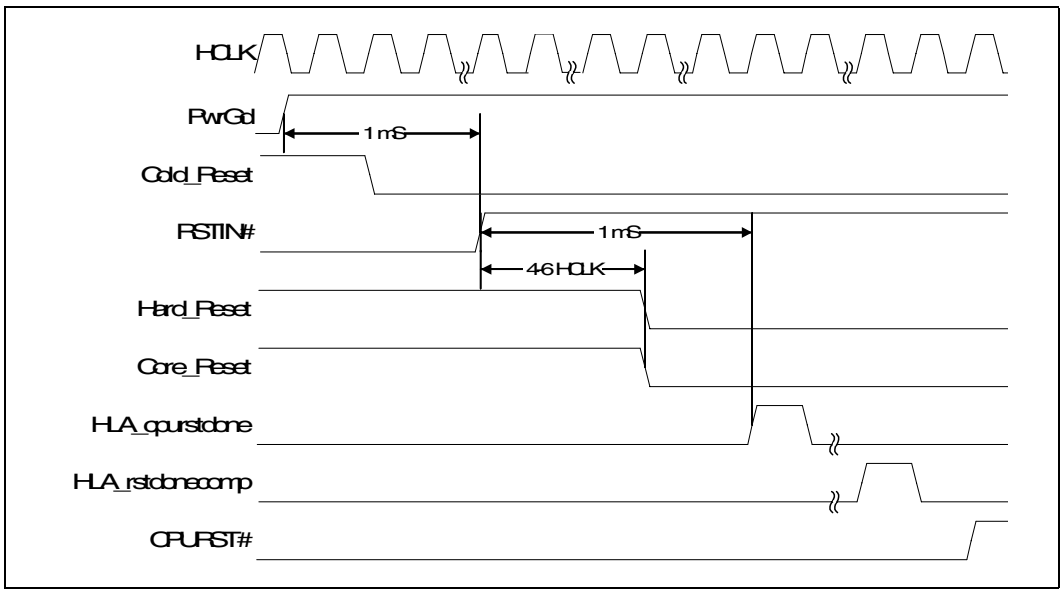

Table 5-11. Reset Sequences and Durations

From	To	Duration	Source	Comment
Power on	PwrGd	>2mS	Platform	Control logic on the platform must ensure that there are at least 2mS of stable power before PwrGd is asserted.
PwrGd	RSTIN# deassertion	1mS	ICH	ICH enforces delay between detecting PwrGd asserted and releasing PCIRST# (note that ICH PCIRST# is directly connected to MCH RSTIN#).
RSTIN# deassertion	Hard/Core deassertion	4-6 HCLK	MCH	MCH waits for a common rising edge on all internal clocks, then releases core reset(s).
RSTIN# deassertion	CPURST# deassertion	1mS	MCH	MCH enforces delay between RSTIN# and CPURST# deassertion. Hublink handshake is incremental to the timer.

5.8 Platform Power Management Support

The MCH is compatible with the *PCI Bus Power Management Interface Specification, Rev 1.1* (referred to here as PCI-PMI). It is also compatible with the *Advanced Configuration and Power Interface Specification (ACPI)*. The MCH is designed to operate seamlessly with operating systems employing these specifications.

5.8.1 Supported System Power States

The MCH and the system power states are analogous, thus no “device” power states are defined for the MCH component. As a result, the MCH power state may be directly inferred from the system power state.

Like all systems, the platform must support the S0 (fully active) state at a minimum. The MCH also supports a system level S1 (idle) state, as well as the S4 (suspend to disk) and S5 (soft off) powered-down idle sleep states. In the latter two states platform power and clocking are disabled, leaving only one or more auxiliary power domains functional. Exit from the S4 and S5 states requires a full system reset and initialization sequence.

The MCH also supports the S3 (suspend to RAM) system power state. A request to enter the S3 power state is communicated to the MCH by the ICH. In response, the MCH will flush all data from the internal coherent write buffer, sequence all active DIMM rows on both channels into their "self refresh" state. Upon completion of this sequence, the MCH will tolerate the removal of all clock references and power sources, save the DDR interface power. DDR interface power must be supplied so that the MCH may hold the DIMMs in self refresh. A full system initialization and configuration sequence is required upon system exit from the S3 state, as all (non-AUX) internal configuration information has been lost throughout the platform, but latency is much lower than it would be from S4 or S5, as the memory image has been maintained.

5.8.1.1 Supported Processor Power States

The MCH supports the C0, C1, and C3 states as defined by the *Advanced Configuration and Power Interface Specification (ACPI)*. This implies that the core logic anchored by the MCH properly understands and handles messaging between the BMC and the system bus, to facilitate transitions into and out of these states.

5.8.1.2 Supported Device Power States

The MCH supports all PCI-PMI and PCI Express* messaging required to place any subordinate device on any of its PCI Express* ports into any of the defined device low power states. Peripherals attached to the PCI segments provided via the PXH/PXH-D component may be placed in any of their supported low power states via messaging directed from the MCH through the intervening PCI Express* hierarchy. Directly attached native PCI Express* devices are not limited in their available low power states, although not all available states support the downstream device "wake-up" semantic.

Further detail on PCI Express* power management support and accompanying PCI Express* and subordinate device power management support are provided in Section 5.8.4.

5.8.1.3 Supported Bus Power States

No low power bus states are supported by the MCH on its Hub Interface to the ICH component, nor does the MCH support placement of the ICH component itself into any low power state below D0. (Other than as a side effect of placing the entire system into one of the S3, S4, or S5 states.)

5.8.2 System Bus Interface Power Management

No low power bus states are supported by the MCH on its processor interface other than those achieved incidentally via placement of the processors themselves into sleep states. The only specialized power management hardware functionality on the system bus interface is the data bus inversion (DBI) logic integral to the processor system bus specification.

5.8.3 DDR Interface Power Management

Hardware in the MCH detects idle conditions on a per-chip-select basis on the DDR subsystem, and places the idle DIMM (or DIMM side) in a clock-disabled low-power state. When in dual-channel operating mode, the MCH simultaneously executes identical control on the corresponding chip selects and clock enables across the two channels. This is true also for the power saving feature described here, because a two-channel configuration necessarily runs in lock-step. When a new access decodes to a sleeping chip select, a single extra clock of latency is incurred to re-enable the clock prior to issuing an activate cycle.

This power saving feature is above and beyond any software power management, and need only be enabled by system BIOS or firmware. No further software direction or interaction is required to realize the power savings from this feature.

5.8.4 PCI Express* Interface Power Management

In PCI Express*, the traditional bus (B*) power states assigned to system buses are replaced by link (L*) power states, which are largely managed by hardware without software intervention. Entry into and out of these states may be initiated by two distinct mechanisms: traditional PCI-PMI type software managed state changes, and non-traditional PCI Express* autonomous hardware state changes. The latter transition type is designated “Active State Power Management,” (ASPM) and is new with the *PCI Express Interface Specification, Rev 1.0a*.

5.8.5 PCI Express* Link Power State Definitions

The *PCI Express Specification* defines the following PCI Express* link power states:

- L0 – Active state with all operations enabled (default state after platform initialization)
- L0s – Low latency, energy saving standby state, disabling exchange of both transaction layer packets and device link layer messages. This state is used exclusively by the ASPM PCI Express* function, with entry and exit managed autonomously by PCI Express* interface hardware. Transitions into and out of the L0s state are accomplished with latencies under 4 ms.
- L1 – Moderate to high latency, very low power, standby state, disabling exchange of both transaction layer packets and device link layer messages. Entered when the downstream device is programmed to a device power state below the D0 active state, or optionally under hardware control during ASPM. The clock remains active in L1, and exit from this state may be initiated by either the upstream or the downstream device.
- L2/L3 Ready – Staging point for removal of main power and clocking. New intermediate state not directly related to PCI PM D-state transitions, nor to ASPM. Hand-shaking will land the link in this state in anticipation of power removal, at which point the link will move to either L2 or L3 depending upon the presence of Vaux.
- L2 – High latency, very low deep sleep state, disabling exchange of transaction layer packets and device link layer messages. L2 is characterized by removal of clocking and main power, but presence of Vaux power. Exit is initiated by restoring clocking and power, and full initialization.
- L3 – High latency, link off state with power, Vaux, and clock reference removed. Exit is initiated by restoring clocking and power, and full initialization.

The MCH supports the *PCI Express Interface Specification, Rev 1.0a*, but does not support the L0s and L1 state via the ASPM mechanism. The reader is referred to the *PCI Express Interface Specification, Rev 1.0a* for further detail on the link states and specific information on entry and exit mechanisms.

5.8.5.1 Software Controlled PCI Express* Link States

Software managed device power state changes do not explicitly control the power L-state of PCI Express* links. Instead the L-state is inferred by hardware from the PCI-PMI power state of the devices attached to that link. When PM software transitions a PCI Express* device to a low power state, that device will automatically negotiate in hardware to bring its upstream link into the appropriate link power state.

Several new semantics are introduced with PCI Express* to support PCI-PMI compatible software managed device and link power state transitions. The majority of the new functionality is to accommodate an essentially edge-triggered in-band message scheme supporting multi chassis cabled system topologies, which must replace the function of traditional level-sensitive board traces for PM event and wake signaling. Further detail on PME signaling appears in Section 5.8.5.3, below.

The MCH supports messaging to facilitate transition of attached PCI Express* devices to power states D0, D1, D2, and D3 (both $D3_{hot}$ and $D3_{cold}$). All attached devices are required by the *PCI Express Interface Specification, Rev 1.0a* to support the D0 and both D3 states, while D1 and D2 support are optional. It is up to software to confirm device support of the optional D1 and D2 states prior to attempting their use on any attached PCI Express* device.

In the D1, D2, and $D3_{hot}$ states the attached device is required to suppress initiation of any link traffic other than PME initiation (if enabled) as a master, and must only accept configuration transactions as a target. Functional context is maintained in the D1 and D2 states, such that full initialization of the attached device is not required upon the wake-up transition back to the D0 state. In both D3 states, functional context is not maintained, and full initialization is required after a transition back to D0.

Placing an attached device into a low power state will result in automatic transition of the associated PCI Express* port to its L1, L2 or L3 link state (depending upon the device power state). To save additional power in L2 state, the platform power manager must remove the reference clock from the link. The MCH does not provide the necessary internal clock generation and distribution control to allow clock removal from one PCI Express* port interface without impacting the operation of its peer ports on the MCH. Nor does the MCH provide support on its PCI Express* link interfaces for the in-band “tone” required to wake from such a state.

5.8.5.2 Hardware Controlled PCI Express* Link States

The MCH does not support ASPM into the L0s and L1 state under hardware control. MCH configuration registers reflect this level of support for ASPM.

All ASPM functionality is disabled by default upon system power-up, and it is the responsibility of software to verify a viable platform clocking configuration prior to enabling ASPM functionality within the MCH or in any attached PCI Express* devices. In topologies where independent clock references are used at any point within the PCI Express* subsystem hierarchy, the “fast training” sequence associated with ASPM is not guaranteed to successfully revive the associated link, and ASPM must remain disabled in the devices at both ends of that link.

5.8.5.3 System Clocking Solution Dependencies

The topology of the platform clocking solution will dictate the viability of ASPM on each of the PCI Express* links, because the nature of the clocks directly impacts the amount of time required to reacquire bit and symbol lock in the receiver after an arbitrarily long non-communicative period. When both ends of a link share a clock source, they will "wander" together over the period they are out of communication with each other, and accordingly will require a relatively brief period of training to reacquire lock. When the two ends of a link utilize completely independent clock references, they may become arbitrarily out of phase with each other while they are in low power states, and will therefore require a significantly longer amount of time to reacquire lock upon waking. For this reason, the *PCI Express Interface Specification, Rev 1.0a* provides for software discovery and communication of the actual clocking topology within the system prior to enabling the ASPM feature on any link within the system.

There are two primary components to the clocking discovery mechanism. Firstly all downstream ports, such as those on the MCH root device, must report whether they use the same clock source as that provided to the slot (or down-device) connected to that port in the platform. This information is recorded in the Slot Clock Configuration bit of the Link Status register for each port, and system BIOS is required to initialize these bits accordingly. Secondly, all add-in devices must report whether they utilize the clock reference provided on the add-in slot via the same bit in the same register of their capability structure.

System software may examine the settings of the Slot Clock Configuration bits of both the upstream and downstream devices for each port in the system, and determine whether a common clock reference is in use. This information is then communicated to both the upstream and the downstream devices via programming of the Common Clock Configuration bit of the Link Status register. The setting of this bit determines the reported exit latency requirements for the L0s and L1 states. System software may then compare the exit latency requirements with the tolerated exit latencies of the attached device, and determine whether or not to enable ASPM for each link the system. (All ASPM functionality defaults to disabled at power-on, and will remain so unless system software determines it may be enabled.)

Note that the "N_FTS" parameters exchanged during initial training will correspond to the "long" exit latencies associated with independent clocks, so if software later sets the Common Clock Configuration bits, it is also necessary to force link retraining in order to update the exchanged N_FTS information.

5.8.5.4 Device and Link PM Initialization

All PCI Express* devices will power-on into the D0 uninitialized state, and will remain in that non-communicative state until they have been configured and at least one of the Memory Space Enable, I/O Space Enable, or Bus Master Enable bits has been set by system software, at which point the device will automatically transition to the D0 active state indicative of normal operation.

5.8.5.5 Device and Slot Power Limits

All add-in devices must power-on to a state in which they limit their total power dissipation to a default maximum according to their form factor (10W for add-in edge-connected cards). When BIOS updates the slot power limit register of the root ports within the MCH, the MCH will automatically transmit a Set_Slot_Power_Limit message with corresponding information to the attached device. It is the responsibility of platform BIOS to properly configure the slot power limit registers in the MCH, and failure to do so may result in attached endpoints remaining completely disabled in order to comply with the default power limitations associated with their form factors.

5.8.6 PME Support

In MCH systems, only the system power manager or a device within the PCI Express* hierarchy may initiate a power state change. Thus the only Power Management Event (PME) signaling support required in the MCH is that associated with PCI Express*. (Note that a device bridging to another technology, such as a PXH bridging to PCI-X, may convert traditional PME signaling into PCI Express* in-band PME messaging and thereby meet this requirement of the MCH.)

PME signaling in PCI Express* is crafted to accomplish two distinct functions. Firstly, it provides a signaling mechanism for devices requiring service to propagate a wake-up request to the power management controller. Secondly, it provides a messaging mechanism for devices requesting a power state change to pass their unique location within the PCI Express* hierarchy to the power management controller. The combination of these two functions provides great flexibility and controllability for the power manager.

5.8.6.1 PME Wake Signaling

Wake signaling is only required to provide for device-initiated transition out of low power states where clock and/or power have been removed from the sleeping device. The PME mechanism does not require a wake-up function for attached devices still powered and receiving an interface reference clock, as devices in this state may simply initiate PME messaging directly. Wake is only required if the device wishing to initiate a PME message cannot do so without first requesting a change to the system clocking and power profile from the power management controller.

The wake signaling aspect of the system power management solution may vary in elegance and granularity. Depending upon the support level provided by the power management controller, a wake-up request from any given device may cause power and clocking to be restored to the entire system, to just the affected branch of the PCI Express* hierarchy, or in some cases only to the requesting device.

While the *PCI Express Interface Specification, Rev 1.0a* provides for two distinct wake signaling mechanisms, the MCH supports only the legacy mechanism described below.

5.8.6.2 Legacy Wake Mechanism

The legacy wake signaling mechanism is directly analogous to that used in historical PCI-based system designs. In this case the platform architect is responsible for crafting paths routing collected wake signals between wake-capable devices and the management controller without participation from the MCH and ICH equivalent devices. The collection of wake logic must run on auxiliary power, and must comprehend the potential for devices both with and without supplied auxiliary power coexisting on the same branch of the PCI Express* hierarchy. Refer to the *PCI Express Interface Specification, Rev 1.0a* for further details on legacy wake signaling.

While familiar and well understood, this mechanism does not provide for device-initiated wake-up in the fully A/C coupled implementation of a multi-chassis PCI Express*-based system solution. The limitation imposed upon the MCH-based system is that remote chassis devices must not be placed in a low-power state with the PCI Express* link clocking and/or power disabled if legacy style wake signaling is desired for any peripheral in that remote chassis. It is still possible for software to place peripheral devices in low power sleep states, and to manage the device state of the PCI Express* device attached to the inter-chassis cable. Devices in the remote chassis may still initiate power state changes via PME messaging provided the inter-chassis link has not been placed into an uncommunicative state. (An alternative for the platform architect would be to place a

compatible switch device between the MCH and the remote chassis that supports the in-band wake mechanism described below, and rely on the switch to forward wake events to the management controller.)

It is up to the platform designer to ensure that the power management controller can adequately isolate the source of a PME wake request as required to take appropriate power management wake-up action.

5.8.6.3 In-Band PCI Express* Wake Mechanism

The PCI Express* In-Band Wake Mechanism is not supported by this MCH.

5.8.6.4 PME Messaging

Once the link requesting a power state change has a communicative upstream link, it sends the PM_PME packet upstream towards the root device (MCH), which in turn is responsible for notifying the management controller. This constitutes an in-band "virtual wire" signaling mechanism to replace the historical solution that involved multiple independent board traces routing PME requests to the power manager. Because the PM_PME propagates "in-band" on the PCI Express* interface without any sideband signaling support, PME functionality is made available to multi-chassis system solutions.

The MCH will collect PME requests from all logical PCI Express* ports, and forward a single MCHPME# output signal directly to the PME# input of the ICH component reserved for power management events. The ICH will then generate a specified interrupt to wake the power manager, and invoke power management software. The interrupt service routine may then interrogate the various PM status registers to determine the source(s) of PME. Note that the ICH PME# pin utilized for PME signaling may not be shared by any other runtime function within the platform.

5.8.6.5 Limitations and Exceptions for Legacy PME

The PME messaging support just outlined for the MCH-based platform is rather inelegant, in that very little information is available to the power management controller. There are further limitations in terms of both lost and spurious PME events as a result of imperfect translation between the PCI Express* in-band "virtual wire" PME mechanism, and the sideband level-sensitive physical wire from MCH to power manager.

This state of affairs is necessitated by the lead intercept nature of a MCH, where the MCH is PCI Express*-aware prior to availability of similarly aware ICH and power management control devices. Mechanisms are provided within the PCI Express* PME semantic to render both spurious and lost PME messages benign from a system architecture perspective.

5.8.7 BIOS Support for PCI Express* PM Messaging

The *PCI Express Interface Specification, Rev 1.0a* stipulates hierarchical messaging semantics enforced by the root device (the MCH) to guarantee proper entry into and exit from unpowered device states. The MCH ACPI BIOS must make special allowances for support of these semantics due to the lack of PCI Express*-aware ICH and power management devices in MCH-based platforms. There are two sets of messages that must be software-assisted in MCH-based platforms to support power-off device states within the PCI Express* hierarchy.

5.8.7.1 PCI Express* PME_TURN_OFF Semantic

Prior to removing power from any attached PCI Express* links anywhere in the hierarchy, the root device must broadcast a PCI Express* "PME_TURN_OFF" message to all downstream devices on the affected PCI Express* port. The receiving devices will propagate this message to all subordinate PCI Express* ports (if any), collect "PME_TO_ACK" acknowledgement packets, and finally return a "PME_TO_ACK" transaction layer packet back to the root device. Once all active ports have acknowledged, the power management device may be notified that it is cleared to modify the collective power state of the PCI Express* hierarchy. These message packets have posted semantics on the interface, thus the turn-off will "push" all prior packets to their endpoints, and the acknowledge will "push" any pending inbound traffic all the way to the root. This prevents "trapping" transactions or PME messages somewhere in the hierarchy at the time power is dropped, ultimately causing them to be lost.

In a pure PCI Express* design, the PME_TURN_OFF packet would originate directly at the power manager, or perhaps at the ICH equivalent device providing connection between the power manager and the remainder of the core logic. Neither the power manager nor the ICH is aware of the PCI Express* messaging mechanism, thus the MCH provides device-specific control and status bits for use by its ACPI BIOS. The sequence of events to place an PCI Express* device in an unpowered state within the platform is as follows:

- PCI-PM or ACPI compliant O/S software is called to place the system into a low power sleep state (S3, S4 or S5), prepares for suspension, and calls ACPI BIOS to carry out the platform power transition.
- The BIOS then communicates to the root complex that all PCI Express* devices should prepare for power-off. This is accomplished through the device-specific configuration space of the internal virtual PCI-to-PCI bridges with subordinate PCI Express* hierarchies. BIOS must configure each active root port to power down. When the configuration write is received to set the "PM Turn Off" bit, the associated root port will transmit a PM_Turn_Off message downstream. At this point, any traffic in-flight continues to be handled normally by the MCH – routed outbound, and completed inbound.
- The target PCI Express* device ceases generation of new transactions inbound, waits for all pending transactions to complete, and prepares to lose power and clocking. If the target device has a subordinate hierarchy of its own, it will propagate the PM_Turn_Off message downstream and wait for acknowledges from all subordinate ports. Once ready to be brought off-line, the target device issues a PM_TO_Ack TLP cycle in acknowledgement back to the root. Note that the link is still communicative at this point, with both power and clock available.
- After issuing the PM_TO_Ack cycle, the downstream device then issues a PM_Enter_L23 DLLP continuously upstream until it receives an acknowledge. In response to the PM_TO_Ack, the root port will set its "Turn Off Back" status bit. In response to the PM_Enter_L23 DLLP, the root will transition its downstream link to the electrical idle state.
- ACPI BIOS, which has been spinning waiting for all of the "Turn Off Ack" status bits to assert, now clears all the command and status bits associated with the PME_TURN_OFF. The routine then informs the power manager to go ahead with the change to the system power state.
- The power manager drops power and clocking to the target device(s), and all associated links automatically transition to either the L2 or L3 uncommunicative power states. The links will enter L2 if Vaux is supplied by the platform, otherwise they will enter L3. The platform will remain in the low power state until a wake event is signaled.

In a fully PCI Express* aware core logic implementation, the ACPI BIOS would not need to act as the interlock between the MCH and the power manager, as all that functionality would be handled in hardware via direct messaging. The combination of the added registers and the added software support compensate for the schizophrenic nature.

5.9 Debug Interface (JTAG)

IEEE 1149.1 standard support is included in the MCH. In addition to the standard JTAG interface, the MCH also supports the XDP (Extended Debug Port) architecture via the **DEBUG[7:0]** pins. JTAG may be utilized as an alternate path into the MCH configuration ring.

5.10 Exception Handling

The MCH recognizes a variety of exception conditions. Some are internally detected; some are detected on input pins; some are passed on behalf of other devices. All recognized exceptions eventually cause the MCH to do one of the following: Send an SERR message, send an SCI message to the ICH, send an SMI message to the ICH, assert MCERR# on the system bus, or do nothing. The MCH, makes no determination of which errors go to which of the three error message schemes; it merely provides the capability for all combinations. It is the BIOS responsibility to determine the ultimate error reporting scheme. The MCH provides a default classification of errors to whether they are at least fatal and non-fatal to more closely match the enterprise error presentation.

5.10.1 Data Error Propagation between Interfaces/Units

Due to the nature of having various data protection schemes; ECC, parity, and CRC, it is necessary to be able to convert between the separate schemes. Beyond this requirement, it is necessary to indicate whether or not incoming data is corrupted. To accomplish this, the MCH implements a functionality referred to as “data poisoning.” Each of the MCH’s external interface units, as well as the internal Posted Memory Write Buffer (PMWB) implements a “Data Poisoning Enable” bit. When this bit is set, and errors are detected on data incoming to the MCH from the external interface (or poisoned data is written to the PMWB), the MCH reports the error via the enabled mechanism for the interface, and also marks the data as poisoned before propagating it on the internal data path towards its destination. This could result in the MCH generating a series of error messages when an error is detected on incoming data via one of the external interfaces, and the associated poisoned data propagates towards its destination through the MCH. Diagnostic software could examine the various error status bits in order to track the errant data through the system.

If the Data Poisoning Enable bit is clear, the error condition is not reported and the data is propagated as if no error was detected.

5.10.2 FERR/NERR Global Register Scheme

Figure 5-7. Global FERR/NERR Register Representation

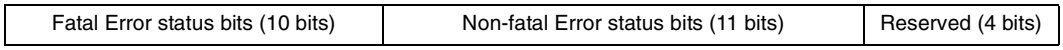

The Global FERR consists of three fields. The first (Fatal) field has 10 bits that will indicate the first signaled fatal error from 10 different sources. The second (Non-fatal) field will indicate the first non-fatal error that occurs from the same 10 sources. A non-fatal error may be either correctable or uncorrectable, but not fatal to the system. These two fields will usually have at most one bit asserted in each field. In the event of simultaneous errors occurring in the same core clock, more than one bit in a field may be set. These registers also contain several bits that are reserved for future enhancements.

The Global NERR will consist of these same three fields with slightly differently functionality. Instead of just the first fatal or non-fatal errors recorded, this register will indicate the second, third, fourth, etc. errors that are reported by the MCH.

5.10.2.1 FERR/NERR Unit Registers

Each major unit has a minimum of a pair of registers, known as the first error (FERR) and next error (NERR). Each unit has different and specific error bit definitions, and provides the specific type of error, information not found in the global registers. It is important to note that the unit FERR/NERR registers are simpler than the global for purposes of reuse and ease of implementation. While the global FERR register has a fatal and a non-fatal field, which lock down separately, the unit FERR register only has one field. However, the units still send out separate fatal and non-fatal indications to the global FERR register if they detect both classifications of errors. Some units will support only one type.

5.10.2.2 Clearing FERR/NERR Registers

When clearing errors, software must clear the FERR/NERR bits in the global registers before clearing the local interface FERR/NERR registers. Software must clear the FERR register and then the NERR register of each interface in that order. After clearing FERR and then clearing NERR, FERR should be read to ensure that remains clear (indicating no more errors have occurred during the clearing process).

5.10.2.3 SERR/SMI/SCI Enabling Registers

Each error reported has a full matrix of direction as to what error message it generates. For each unit FERR/NERR pair there are three more registers that enable each error for one of the three specific error messages. The logic does not preclude the generation of all three messages for a single error, but this is not a supported configuration. SERR stands for System Error, and should be used for reporting address and data parity errors, or any other catastrophic system error. SCI stands for System Control Interrupt, and is a shareable interrupt used to notify the Operating System (OS) of ACPI events. SMI stands for System Management Interrupt, and is an OS-transparent interrupt generated by events on legacy systems.

5.10.2.4 MCERR Enabling Registers

An additional entry to the matrix of error signaling paths is the MCERR (Machine Check Error) enabling register. In addition to the SERR, SMI, and SCI enabling registers, the MCERR enabling register allows the occurrence of an error to result in the MCERR# signal to be asserted on the system bus. MCERR# is asserted to indicate an unrecoverable error, which is not a bus protocol violation.

5.10.2.5 Error Escalation Register

Since all error bits in the error registers are fully configurable, meaning that a given error can be configured to go to any of the four messaging methods, no global error escalation mechanism is required. Although, the errors occurrence is accumulated in the global FERR/NERR registers, all error messaging is initiated from the units themselves, and not from a central location.

5.10.2.6 Error Masking

A new feature being added for the MCH is the concept of an error masking register. Each unit has a mask register, which blocks the recognition/logging/reporting of each specific error type. Since the error will not be recognized when the corresponding mask bit is set, no error messages can be generated. This feature allows intelligent software to ignore specific error types during critical areas of code, where it does not want to be informed of errors that it will create, without ignoring other error types that it doesn't expect to happen. These mask bits will default to unmasked, and must be set by software or BIOS to take affect.

5.10.2.7 Locking DRAM Address and Syndrome on Errors

The two error logging registers for correctable errors, DRAM_SEC_ADD and DRAM_SEC_SYNDROME are locked when bits 0 or 8 of either the DRAM_FERR or DRAM_NERR is set. If both of these bits are '0', then the two logging registers may be updated. If either is set to '1', then the two registers will retain their value even if new correctable errors are found. This allows the first error to be captured and held instead of retaining the last. Corrected data errors as a result of scrubber-initiated traffic will be reflected in these error registers.

The logging register for uncorrectable errors, DRAM_DED_ADD is locked when bits 1 or 9 in either the DRAM_FERR or DRAM_NERR is set. This register holds the address of uncorrectable errors on data reads not initiated by the scrubber.

The logging register for Scrub detected errors, DRAM_SCRUB_ADD should be locked when bits 2 or 10 of either the DRAM_FERR or DRAM_NERR is set. This register holds the address for scrubber-initiated transactions for either periodic memory scrubbing or sparing.

The logging register for Retry detected errors, DRAM_RETRY_ADD should be locked when bits 5 or 13 of either the DRAM_FERR or DRAM_NERR is set. This register is locked when the determination is made that a retry to this address must be performed.

When the FERR/NERR registers are cleared, the logging registers are free to update their contents until such time that either of these FERR/NERR registers again lock.

5.10.2.8 Memory Error Counters

The MCH provides error counters for each DIMM (single-channel mode) or DIMM pair (dual channel mode). There is a correctable error counter and an uncorrectable error counter for each DIMM or DIMM pair. The intent of these counters is primarily for support of the DIMM sparing feature. Compared against configurable threshold values, the values of these counters determine when sparing is invoked. The values in these counters decay over time, since the total number of errors over time is not interesting, but the count of errors within a window of time.

5.10.2.9 PCI Express* Errors and Errors on Behalf of PCI Express*

MCH-specific error detection, masking, and escalation mechanisms operate on a parallel path to their standardized counterparts included in the *PCI Express Interface Specification, Rev 1.0a*. PCI Express* errors are classified as either correctable or uncorrectable. Uncorrectable errors are further broken down as fatal or non-fatal.

PCI Express* specified correctable errors are logged in the Correctable Error Status Register (Device 2-4, Function 0, Offset 110 - 113h), unless they are masked by a corresponding bit in the Correctable Error Detect Mask Register (Device 2-4, Function 0, Offset 150-153h).

PCI Express* specified uncorrectable errors are logged in the Uncorrectable Error Status Register (Device 2-4, Function 0, Offset 104-107h), unless they are masked by a corresponding bit in the Uncorrectable Error Detect Mask Register (Device 2-4, Function 0, Offset 14C-14Fh). The Uncorrectable Error Severity Register (Device 2-4, Function 0, Offset 10C - 10Fh) determines if bits in the Uncorrectable Status register are treated as uncorrectable fatal or uncorrectable non-fatal errors. The Device Status register (6Eh) bits are set when the corresponding category of bit is set in the uncorrectable and correctable status registers.

Reporting of non-masked error bits to the root complex hierarchy of PCI Express* error registers is controlled on three different levels. Individual errors are masked for reporting by the Uncorrectable Error Mask (Device 2-4, Function 0, Offset 108-10Bh) and the Correctable Error Mask (Device 2-4, Function 0, Offset 114-117h) registers. Individual error category (fatal, non-fatal, correctable, or unsupported) reporting is enabled in the Device Control Register (Device 2-4, Function 0, Offset 6Ch) bits 3:0. Finally, uncorrectable error reporting (fatal or non-fatal) reporting may also be enabled by setting the SERR Enable bit in the PCI Command Register (Device 2-4, Function 0, Offset 04-05h).

There is an error pointer, in the Advanced Error Capability and Control Register (Device 2-4, Function 0, Offset 118-11Bh) which will log the first uncorrectable error that is enabled for reporting. Also some uncorrectable errors, when they are the first uncorrectable error, will log their corresponding header log in the Header Log Registers (Device 2-4, Function 0, Offset 11C-12Bh). An error pointer for unmasked correctable errors has been added in the Error Do Command Register (Device 2-4, Function 0, Offset 148-14Bh).

These internally detected errors when they are reported are referred to as virtual error messages. These are different from errors which are detected by the downstream device which then sends an error message to the root complex, which are referred to as externally detected or “received’ error messages. The received system error bit in the PCI Secondary Status Register (Device 2-4, Function 0, Offset 1E-1F) is set when either fatal or non-fatal messages are received at the root complex.

At this point in the PCI Express* error hierarchy, these virtual error messages are logically OR’ed with the received error messages, and will just be referred to as fatal, non-fatal, or correctable error messages, no reference to either virtual or received.

When enabled by the enable system error bit in the PCI Command Register, any fatal or non-fatal messages will set the signaled system error bit in the PCI Status register (Device 2-4, Function 0, Offset 06-07h). The Root Port Error Message Status Register (Device 2-4, Function 0, Offset 130-133h) will indicate first and multiple errors of each error message category, and the corresponding error source IDs of the first correctable and uncorrectable error messages will be the logged in the Root Port Error Source ID register (134h).

These errors that have been reported to the root complex can now be reported to the system, via the category enables in the Root Port Error Command Register (Device 2-4, Function 0, Offset 12C-12Fh) for interrupts. These interrupts can be in the form of legacy type interrupts if so enabled in the PCI command register and MSI is not enabled, or message signaled interrupts if so enabled in the MSI Capabilities register (Device 2-4, Function 0, Offset 5A-5Bh).

The Root Port Control register (Device 2-4, Function 0, Offset 80-83h) enables errors to be reported to the system via other MCH specific methods, again on a category basis. The Error Do Command register, selects between the four methods of system signaling, SERR, SCI, SMI, and MCERR.

The error model outside of PCI Express* includes a local FERR/NERR pair of registers in each unit and a global FERR/NERR pair of registers that indicates which unit had problems. The Local FERR/NERR register pair (Device 2-4, Function 0, Offset 160-163h & 164-167h) includes PCI Express* defined errors and additional detected errors within the PCI Express* unit. This register pair has three sets of error bits for the three categories of errors: the first set for received messages, the second set for internally detected errors (virtual messages need not have been generated), and unit specific errors outside of the *PCI Express Specification*, and the third set for device errors. The error scheme sets FERR/NERR error bits regardless whether or not they were reported via interrupt or other signaling method.

The signaling due to unit specific errors has its logic dependent on the PCI Express* Unit Error Register (Device 2-4, Function 0, Offset 140-143h). The errors flagged in this register must be cleared before exiting the error service routine.

The signaling due to received messages has its logic dependent on the Root Error Status register (Device 2-4, Function 0, Offset 130-133h). The Root Error Status register must be cleared before exiting the error service routine.

The signaling due to internally detected PCI Express* errors has its logic dependent on the Device Status register (Device 2-4, Function 0, Offset 6E-6Fh). The Device Status register must be cleared before exiting the error service routine.

Software must clear the global FERR first, and then the global NERR. Software then clears the local FERR register and the local NERR register of each unit in that order. After clearing FERR and then clearing NERR, the local FERR should be read to make sure that remains ‘0’ indicating no more errors have occurred during the clearing of these registers. After all units’ FERR & NERR registers have been cleared, the global FERR is again read to ensure that no additional errors occurred during the clearing sequence.

Since the PCI Express* units have more hierarchy than other units, more registers must be cleared other than just the local FERR and NERR registers. After clearing the local FERR & NERR, one must also clear the Root Error Status, Unit Error Status, Device Status, Uncorrectable Error Status, and Correctable Error Status registers. One only needs to clear the PCI Status and Secondary Status registers if these are being utilized in a given particular error model. No logic depends on the state of any of these status bits. If not utilized, they can be ignored.

If a PCI Express* error handler is used, with no knowledge of the FERR/NERR registers, then clear the PCI Express*-specific registers: Device Status, Uncorrectable Error Status, Correctable Error Status, and Root Error Status. The MCH specific unit errors would not be enabled for reporting errors.

Figure 5-8. PCI Express* Error Handling

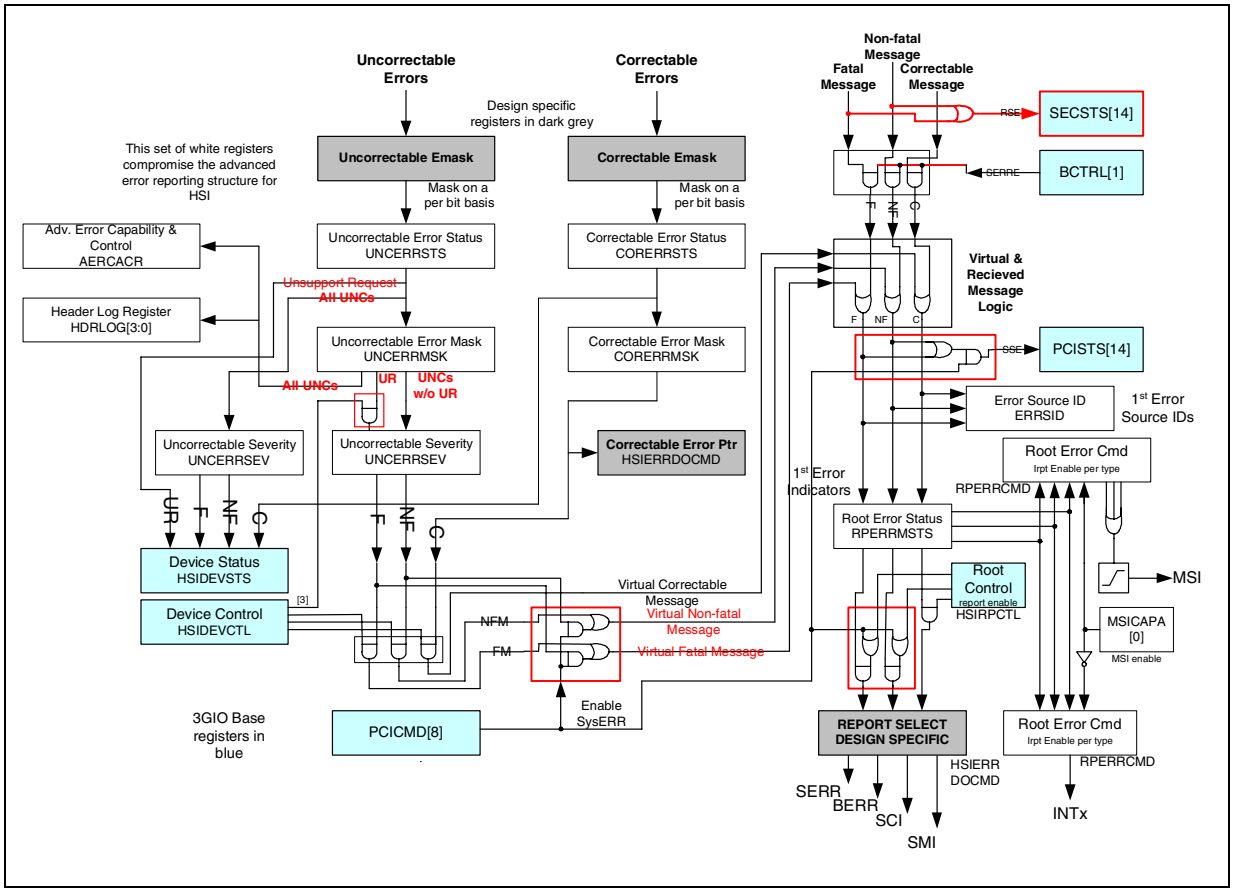

5.10.2.10 Configurable Error Containment at the Legacy Interface

Depending on the I/O devices in use, data errors could have catastrophic effects when allowed to propagate. The legacy Hub Interface to the ICH has two configurable options for error detection and containment. Data Poisoning allows the MCH to mark incoming data (ICH to MCH) as "poisoned" if parity errors are detected on that data. This "poisoned" data will subsequently result in the generation of additional error detection/logging/reporting as it travels through the various functional units within the MCH. The MCH Hub Interface logic also implements a "Stop and Scream" function to prevent "poisoned" outgoing data from transferring at all. When this function is enabled, and "poisoned" data is detected on an outbound transaction destined for the Hub Interface, the transaction is terminated internally (the corrupt data is never sent to the ICH), and the error is escalated to the system. This is extreme behavior, which can be enabled or disabled, in order to prevent data corruption on a critical device.

5.10.2.11 Error Signaling Mechanisms

The notification action via the ICH taken by the MCH upon detection of an error is controlled through three registers. The SERRCMD register enables the generation of the SERR message, the SCICMD register enables the generation of the SCI message, and the SMICMD register enables

the generation of SMI messages. These messages are transmitted to the ICH via special cycles on the Hub Interface. The ICH receives the exception notification from the MCH and may be configured to notify the processor of the condition.

Once the processor has been interrupted, it polls the system to determine the cause of the exception. If the MCH initiated the exception condition by sending a message over Hub Interface, then the processor is so informed by the ICH. At this point, the processor may read the MCH's error status registers to determine the exact cause of the condition. The processor explicitly clears the status bit that points to the exception condition.

The MCH, in addition to signaling errors to the ICH for further handling, has added the capability of signaling the processor directly by use of the system bus error signal MCERR#. The processor, upon observing this signal active, traditionally enters into special error handling code known as machine check code.

For each type of error detected in a given unit, there is a bit that corresponds to that error in the unit_FERR, unit_NERR, unit_SERRCMD, unit_SMICMD, unit_SCICMD, and unit_MCERRCMD registers. (Note that one and only one xCMD bit can be enabled per error type.) The first occurrence of an error type will be indicated by the bit assertion in the unit_FERR. If that error occurs again then the corresponding bit will be set in the unit_NERR register. When a bit is asserted in either the unit_FERR or unit_NERR, and if the corresponding enable bit is set in one of the named CMD registers, then the enabled error signal will be transmitted to the ICH via, the Hub Interface. The transmission of the SERR cycle also requires that the SERR enable in the PCICMD register is set. The assertion of the SERR signal also causes the appropriate SERR signaled status bit to be set in the PCISTS register.

5.11 SMBus Port Description

The MCH provides a fully functional System Management Bus (SMBus) Revision 2.0 compliant target interface, which provides direct access to all internal MCH configuration register space. SMBus access is available to all internal configuration registers, regardless of whether the register in question is normally accessed via the memory-mapped mechanism or the standard configuration mechanism. This provides for highly flexible platform management architectures, particularly given a baseboard management controller (BMC) with an integrated network interface controller (NIC) function.

The SMBus interface consists of two interface pins; one a clock, and the other serial data. Multiple initiator and target devices may be electrically present on the same pair of signals. Each target recognizes a start signaling semantic, and recognizes its own 7-bit address to identify pertinent bus traffic. The MCH address is hard-coded to 0110000b (60h).

The protocol allows for traffic to stop in "mid sentence," requiring all targets to tolerate and properly "clean up" in the event of an access sequence that is abandoned by the initiator prior to normal completion. The MCH is compliant with this requirement.

The protocol comprehends "wait states" on read and write operations, which the MCH takes advantage of to keep the bus busy during internal configuration space accesses.

5.11.1 Internal Access Mechanism

All SMBus accesses to internal register space are initiated via a write to the CMD byte. Any register writes received by the MCH while a command is already in progress will receive a NAK to prevent spurious operation. The master is no longer expected to poll the CMD byte to prevent the obliteration a command in progress prior to issuing further writes. The SMBus access will be delayed by stretching the clock until such time that the data is delivered. Note that per the *System Management Bus (SMBus) Specification, Rev 2.0*, this can not be longer than 25 ms. To set up an internal access, the four ADDR bytes are programmed followed by a command indicator to execute a read or write. Depending on the type of access, these four bytes indicate either the Bus number, Device, Function, Extended Register Offset, and Register Offset, or the memory-mapped region selected and the address within the region. The configuration type access utilizes the traditional bus number, device, function, and register offset; but in addition, also uses an extended register offset which expands the addressable register space from 256 bytes to 4 Kbytes. The memory-mapped type access redefines these bytes to be a memory-mapped region selection byte, a filler byte which today is all zeroes, and then the memory address within the region. Refer to the earlier tables, which display this information. Note that the filler byte is today not utilized but enforces that both types of accesses have the same number of address bytes, and does allow for future expansion.

It is perfectly legal for an SMBus access to be requested while an PSB-initiated access is already in progress. The MCH supports “wait your turn” arbitration to resolve all collisions and overlaps, such that the access that reaches the configuration ring arbiter first will be serviced first while the conflicting access is held off. An absolute tie at the arbiter will be resolved in favor of the FSB. Note that SMBus accesses must be allowed to proceed even if the internal MCH transaction handling hardware and one or more of the other external MCH interfaces are hung or otherwise unresponsive.

5.11.2 SMBus Transaction Field Definitions

The SMBus target port has it’s own set of fields which the MCH sets when receiving an SMBus transaction. They are not directly accessible by any means for any device.

Table 5-12. SMBus Transaction Field Summary

Position	Mnemonic	Field Name
1	CMD	Command
2	BYTCNT	Byte Count
3	ADDR3	Bus Number (Register Mode) or Destination Memory (Memory Mapped Mode)
4	ADDR2	Device / Function Number (Register Mode) or Address Offset [23:16] (Memory Mapped Mode)
5	ADDR1	Extended Register Number (Register Mode) or Address Offset [15:8] (Memory Mapped Mode)
6	ADDR0	Register Number (Register Mode) or Address Offset [7:0] (Memory Mapped Mode)
7	DATA3	Fourth Data Byte [31:24]
8	DATA2	Third Data Byte [23:16]
9	DATA1	Second Data Byte [15:8]
10	DATA0	First Data Byte [7:0]
11	STS	Status, only for reads

Table 5-12 indicates the sequence of data as it is presented on the SMBUS following the byte address of the MCH itself. This is not to necessarily indicate any specific register stack or array implemented in the MCH. Note that the fields can take on different meanings depending on whether it is a configuration or memory-mapped access type. The command indicates how to interpret the bytes.

5.11.2.1 Command Field

The command field indicates the type and size of transfer. All configuration accesses from the SMBus port are initiated by this field. While a command is in progress, all future writes or reads will be NACK'd by the MCH to avoid having registers overwritten while in use. The two command size fields allows for more flexibility on how the data payload is transferred, both internally and externally. The begin and end bits support the breaking of the transaction up into smaller transfers, by defining the start and finish of an overall transfer.

Position	Description
7	**Begin Transaction Indicator.** 0 = Current transaction is NOT the first of a read or write sequence. 1 = Current transaction is the first of a read or write sequence. On a single transaction sequence this bit is set along with the End Transaction Indicator.
6	**End Transaction Indicator.** 0 = Current transaction is NOT the last of a read or write sequence. 1 = Current transaction is the last of a read or write sequence. On a single transaction sequence this bit is set along with the Begin Transaction Indicator.
5	**Address Mode.** Indicates whether memory or configuration space is being accessed in this SMBus sequence. 0 = Memory Mapped Mode 1 = Configuration Register Mode
4	**Packet Error Code (PEC) Enable.** When set, each transaction in the sequence ends with an extra CRC byte. The MCH would check for CRC on writes and generate CRC on reads. PEC is not supported by the MCH. 0 = Disable 1 = Not Supported
3:2	**Internal Command Size.** All accesses are naturally aligned to the access width. This field specifies the internal command to be issued by the SMBus slave logic to the MCH core. 00 = Read Dword 01 = Write Byte 10 = Write Word 11 = Write Dword
1:0	**SMBus Command Size.** This field specifies the SMBus command to be issued on the SMBus. This field is used as an indication of the length of the transfer so that the slave knows when to expect the PEC packet (if enabled). 00 = Byte 01 = Word 10 = DWord 11 = Reserved

5.11.2.2 Byte Count Field

The byte count field indicates the number of bytes following the byte count field when performing a write or when setting up for a read. The byte count is also used when returning data to indicate the following number of bytes (including the status byte) which are returned prior to the data. Note that the byte count is only transmitted for block type accesses on SMBus. SMBus word or byte accesses do not use the byte count.

Position	Description
7:0	**Byte Count.** Number of bytes following the byte count for a transaction.

5.11.2.3 Address Byte 3 Field

This field should be programmed with the Bus Number of the desired configuration register in the lower 5 bits for a configuration access. For a memory-mapped access, this field selects which memory-map region is being accessed. In contrast to how some earlier MCHs operated, there is no status bit to poll to see if a transfer is currently in progress, because by definition if the transfer completed than the task is done. The clock stretch is used to guarantee the transfer is truly complete.

The MCH does not support access to other logical bus numbers via the SMBus port. All registers "attached" to the configuration mechanism SMBus has access to are all on logical bus#0. The MCH makes use of this knowledge to implement a modified usage of the Bus Number register providing access to internal registers outside of the PCI compatible configuration window.

Position	Configuration Register Mode Description	Memory Mapped Mode Description
7:5	Ignored.	Memory map region to access. 01h = DMA 08h = DDR 09h = CHAP Others = Reserved
4:0	**Bus Number.** Must be zero: the SMBus port can only access devices on the MCH and all devices are bus zero.	

5.11.2.4 Address Byte 2 Field

This field indicates the Device Number & Function Number of the desired configuration register if for a configuration type access, otherwise it should be set to zero.

Position	Configuration Register Mode Description	Memory Mapped Mode Description
7:3	**Device Number.** Can only be devices on the MCH.	Zeros used for padding.
2:0	**Function Number.**	

5.11.2.5 Address Byte 1 Field

This field indicates the upper address bits for the register with the 4K region. Whether it is a configuration or memory-map type of access, only the lower bit positions are utilized, the upper four bits are ignored.

Position	Description
7:4	Ignored.
3:0	**Extended Register Number.** Upper address bits for the 4K region of register offset.

5.11.2.6 Address Byte 0 Field

This field indicates the lower eight address bits for the register with the 4K region, regardless whether it is a configuration or memory-map type of access.

Position	Description

7:0	**Register Offset.**

5.11.2.7 Data Field

This field is used to receive read data or to provide write data associated with the desired register.

At the completion of a read command, this field will contain the data retrieved from the selected register. All reads will return an entire aligned DWord (32 bits) of data.

The appropriate number of byte(s) of this 32 bit field should be written with the desired write data prior to issuing a write command. For a byte write only bits 7:0 will be used, for a Word write only bits 15:0 will be used, and for a DWord write all 32 bits will be used.

Position	Description
31:24	**Byte 3 (DATA3).** Data bits [31:24] for DWord.
23:16	**Byte 2 (DATA2).** Data bits [23:16] for DWord.
15:8	**Byte 1 (DATA1).** Data bits [15:8] for DWord and Word.
7:0	**Byte 0 (DATA0).** Data bits [7:0] for DWord, Word and Byte.

5.11.2.8 Status Field

For a read cycle, the data is preceded by a byte of status. The following table shows how these bits are defined.

Position	Description
7	**Internal Timeout.** **0 = SMBus request is completed within 2 ms internally** 1 = SMBus request is not completed in 2 ms internally.
6	Ignored.
5	**Internal Master Abort.** 0 = No Internal Master Abort Detected. 1 = Detected an Internal Master Abort.
4	**Internal Target Abort.** 0 = No Internal Target Abort Detected. 1 = Detected an Internal Target Abort.
3:1	Ignored.
0	**Successful.** 0 = The last SMBus transaction was not completed successfully. 1 = The last SMBus transaction was completed successfully.

5.11.3 Unsupported Access Addresses

It is possible for an SMBus master to program an unsupported bit combination into the ADDR registers. The MCH does not support such usage, and may not gracefully terminate such accesses.

5.11.4 SMB Transaction Pictograms

Since the new SMB target interface is of enterprise origin, it is more complex than the original SMB target interface of desktop origin. The following drawings are included which should be better than words to demonstrate the different types of transactions, especially how they can be broken up into multiple smaller transfers.

Figure 5-9. DWORD Configuration Read Protocol (SMBus Block Write / Block Read, PEC Disabled)

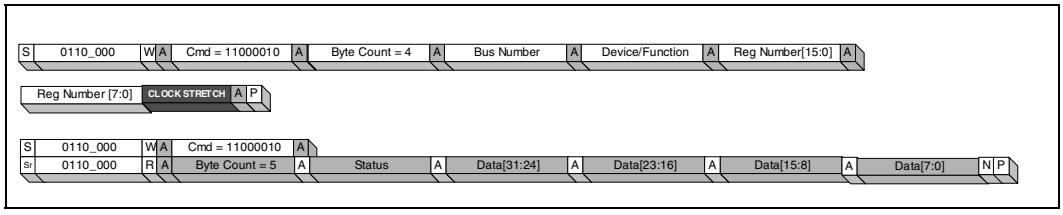

Figure 5-10. DWORD Configuration Write Protocol (SMBus Block Write, PEC Disabled)

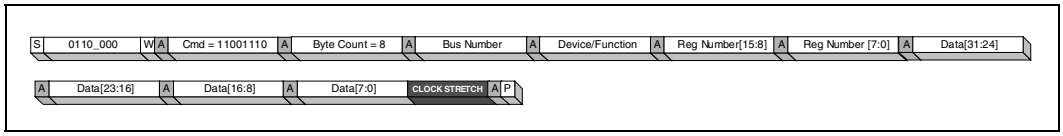

Figure 5-11. DWORD Memory Read Protocol (SMBus Block Write / Bock Read, PEC Disabled)

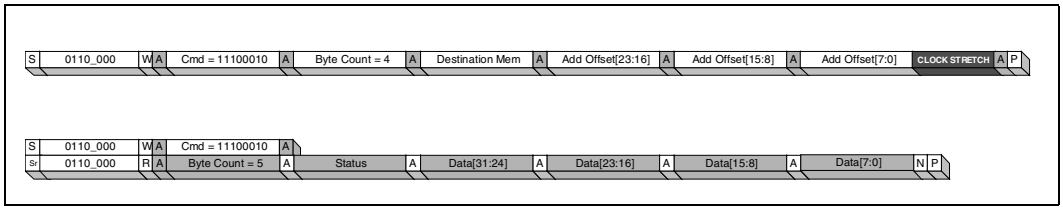

Figure 5-12. DWORD Memory Write Protocol

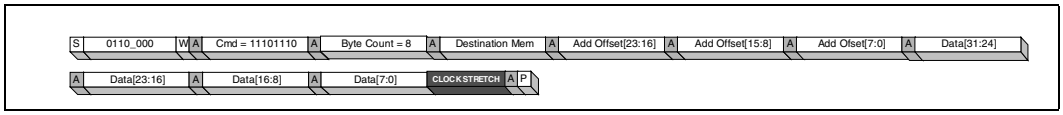

Figure 5-13. DWORD Configuration Read Protocol (SMBus Word Write / Word Read, PEC Disabled)

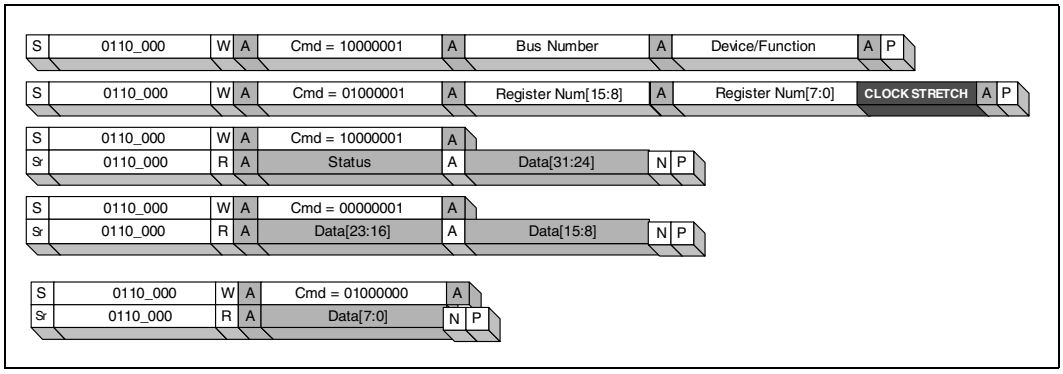

Figure 5-14. DWORD Configuration Write Protocol (SMBus Word Write, PEC Disabled)

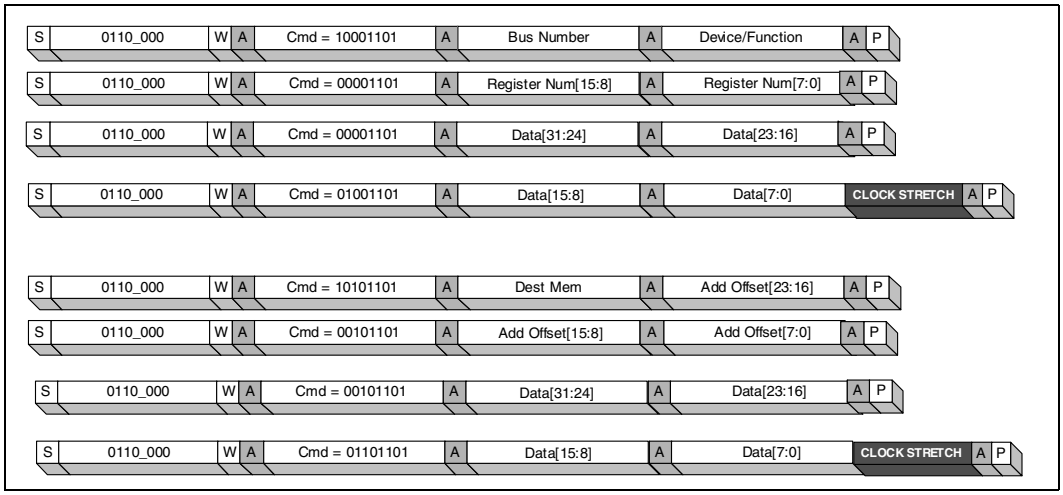

Figure 5-15. DWORD Memory Read Protocol (SMBus Word Write / Word Read, PEC Disabled)

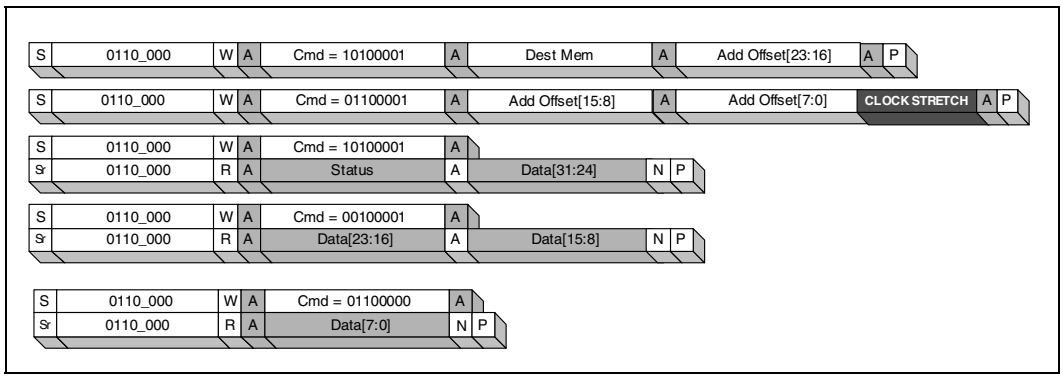

Figure 5-16. WORD Configuration Wrote Protocol (SMBus Byte Write, PEC Disabled)

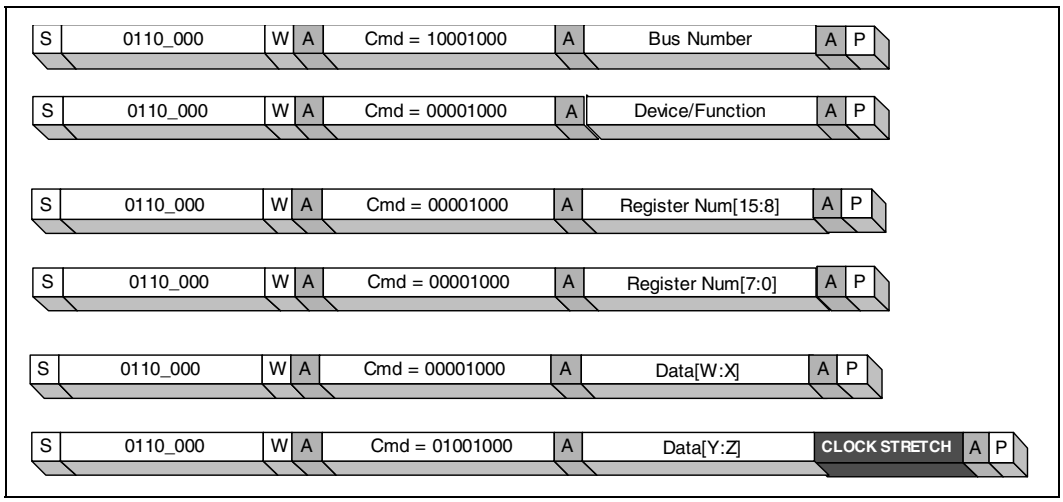

Electrical Characteristics 6

This chapter provides the absolute maximum ratings and DC characteristics for the MCH.

6.1 Absolute Maximum Ratings

Table 6-1 lists the MCH's maximum environmental stress ratings. Functional operation at the absolute maximum and minimum is neither implied nor guaranteed. Functional operating parameters are listed in the AC and DC tables.

Warning: Stressing the device beyond the "Absolute Maximum Ratings" may cause permanent damage. These are stress ratings only. Operating beyond the "operating conditions" is not recommended and extended exposure beyond "operating conditions" may affect reliability.

Table 6-1. Absolute Maximum Ratings

Symbol	Parameter	Min	Max	Unit
Tstorage	Storage Temperature	-40.0	85.0	°C
Vcc_MCH	Core Supply Voltage with respect to Vss	-0.50	1.65	V
VTT_AGTL	Supply Voltage input with respect to Vss	-0.50	1.40	V
Vdd_DDR	DDR Buffer Supply Voltage	-0.50	2.80	V

6.2 Power Characteristics

Table 6-2. Operating Condition Power Supply Rails

Symbol	Parameter	Minimum		Nominal	Maximum		Unit	Notes
		AC	DC		AC	DC		
Vtt	Host AGTL+ Termination Voltage	1.140	1.164	1.2	1.236	1.260	V	1
Itt	Host AGTL+ Termination Current				3.2		A	
Vdd_DDR	DDR Buffer Voltage	2.325	2.375	2.5	2.625	2.675	V	1
Idd_DDR	DDR Buffer Current				8.8		A	3, 4
Vdd_DDR2	DDR2 Buffer Voltage	1.674	1.71	1.8	1.89	1.926	V	1
Idd_DDR	DDR2 Buffer Current				6.7		A	3, 4
Vcc_MCH	1.5 V Supply Voltage	1.425	1.455	1.5	1.545	1.575	V	1
Icc_MCH	1.5V Supply Current on Vcc				3.5		A	5
Icc_EXP	1.5V Supply Current on VccEXP				2.0		A	5, 6
VccA_SB Vcca_HI Vcca_EXP Vcca_DDR	Analog Supply Voltages		1.406	1.5	1.545		V	2
Vcc_BGHS	Analog Bandgap Voltage		2.425	2.5	2.575		V	
V3REF	V3REF Voltage	3.135	3.201	3.3	3.399	3.465	V	

NOTES:
1. The DC min/max window describes the power supply behavior required at frequencies below 5 MHz. The AC min/max window describes the power supply behavior required at frequencies from 5 MHz to 2 GHz.
2. The analog voltage is intended to be a filtered copy of the 1.5V core supply voltage.
3. DDR and DDR2 operation are mutually exclusive; thus the MCH will draw only one of these Idd currents in a given design.
4. Idd_DDR x Vdd_DDR does not equal power dissipated in the MCH from the DDR rail. Most of the current supplied by the Idd rail at max current draw is sourced out the memory signal pins at some voltage above ground. See the memory interface DC chapters for details.
5. Total current drawn off the 1.5V rail is separated into current drawn by the VCC balls and the VCCEXP balls as shown in the component ballout in the EDS.
6. Note that the VCCEXP balls should be inductively isolated from the VCC balls to avoid noise coupling from the core onto the VCCEXP rail, which is particularly sensitive to AC noise.

6.3 I/O Interface Signal Groupings

Table 6-3. Signal Groups FSB Interface (Sheet 1 of 2)

Signal Group	Signal Type	Signals	Notes
(a)	AGTL+ I/O	AP [1:0]#, ADS#, BNR#, DBSY#, DP [3:0]#, DRDY#, HA [35:3]#, HADSTB [1:0] #, HD [63:0] #, HDSTBP [3:0]#, HDSTBN [3:0]#, HIT#, HITM#, HREQ [4:0]#, BREQ[0]#, DINV[3:0]#, MCERR#	
(b)	AGTL+ Output	BPRI#, CPURST#, DEFER#, HTRDY#, RS [2:0]#, RSP#	
(c)	AGTL+ Input	HLOCK#, BINIT#, BREQ[1]#	

Table 6-3. Signal Groups FSB Interface (Sheet 2 of 2)

Signal Group	Signal Type	Signals	Notes
(d)	Host Reference Voltage (Analog)	HDVREF [1:0], HACVREF	
(e)	Host Compensation (Analog)	HODTCRES, HSLWCRES, HCRES0	
(f)	Clock CMOS Inputs	HCLKINN, HCLKINP	

Table 6-4. Signal Groups Memory (DDR and DDR2) Interface

Signal Group	Signal Type	Signals	Notes
(h)	DDR and DDR2 I/O	DDRx_DQ[63:0], DDRx_CB[7:0], DDRx_DQSP[17:0]	1
(i)	DDR and DDR2 Output	DDRx_CMDCLKP[3:0], DDRx_CMDCLKN[3:0], DDRx_MA[13:0], DDRx_BA[2:0], DDRx_RAS#, DDRx_CAS#, DDRx_WE#, DDRx_CS[7:0]#, DDR_CKE[7:0]	1
(j)	Analog I/O	DDR_SLWCRES, DDR_IMPCRES, DDR_RES[2:1], DDR_CRES0	
(k)	Memory Reference Voltage (Analog)	DDRA_DDRVREF, DDRB_DDRVREF	

NOTES:
1. x = A, B DDR Channel.

Table 6-5. Signal Groups PCI Express* Interface

Signal Group	Signal Type	Signals	Notes
(l)	PCI Express* Input	EXP_A0_RXP[7:4] /EXP_A1_RXP[3:0], EXP_A0_RXN[7:4]/EXP_A1_RXN[3:0], EXP_B0_RXP[7:4]/EXP_B1_RXP[3:0], EXP_B0_RXN[7:4]/EXP_B1_RXN[3:0], EXP_C0_RXP[7:4]/EXP_C1_RXP[3:0], EXP_C0_RXN[7:4]/EXP_C1_RXN[3:0]	
(m)	PCI Express* Output	EXP_A0_TXP[7:4]/EXP_A1_TXP[3:0], EXP_A0_TXN[7:4]/EXP_A1_TXN[3:0], EXP_B0_TXP[7:4]/EXP_B1_TXP[3:0], EXP_B0_TXN[7:4]/EXP_B1_TXN[3:0], EXP_C0_TXP[7:4]/EXP_C1_TXP[3:0], EXP_C0_TXN[7:4]/EXP_C1_TXN[3:0]	
(n)	PCI Express* Clock Pair	EXP_CLKP, EXP_CLKN	
(o)	PCI Express* Buffer Compensation	EXP_COMP[1:0]	

Table 6-6. Signal Groups Hub Interface

Signal Group	Signal Type	Signals	Notes
(p)	Hub Interface CMOS I/O	HI[11:0], HI_STBF, HI_STBS	
(q)	Hub Interface CMOS Input Clock	CLK66	
(r)	Hub Interface Reference Voltage Input	HIVREF	
(s)	Hub Interface Voltage Swing	HIVSWING	
(t)	Hub Interface Compensation CMOS I/O	HICOMP	

Table 6-7. Signal Groups Reset and Miscellaneous

Signal Group	Signal Type	Signals	Notes
(u)	SMBus I/O Buffer	SMBSCL, SMBSDA, PWRGD, EXP_HPINT#, RSTIN#	
(v)	JTAG I/O	TRST#, TMS, TDI, TCK, TDO	
(w)	Misc. CMOS Input	TEST#	
(x)	SMBus Output Buffer, OD	MCHPME#, MCHGPE#	
(y)	Misc. Input	PLLSEL[1]	
(z)	Misc. Input	PLLSEL[0]	

6.4 DC Characteristics

Table 6-8. FSB Interface DC Characteristics

Symbol	Signal Group	Parameter	Min	Nom	Max	Unit	Notes
V_{IL_H}	(a), (c)	Host AGTL+ Input Low Voltage			GTLREF – 0.10 x Vtt	V	1
V_{IH_H}	(a), (c)	Host AGTL+ Input High Voltage	GTLREF + 0.10 x Vtt			V	1
V_{OL_H}	(a), (b)	Host AGTL+ Output Low Voltage		0.4		V	
V_{OH_H}	(a), (b)	Host AGTL+ Output High Voltage	Vtt - 0.1	Vtt		V	
R_{TT}		Host Termination Resistance	45	50	55	Ω	
I_{OL_H}	(a), (b)	Host AGTL+ Output Low		17		mA	
I_{L_H}	(a), (c)	Host AGTL+ Input Leakage Current			100	µA	
C_{PAD}	(a), (c)	Host AGTL+ Input Capacitance	1		2.5	pF	
HACVREF	(d)	Host Address and Common Clock Reference Voltage	0.98 x Nominal	0.63 x Vtt	1.02 x Nominal	V	
HDVREF	(d)	Host Data Reference Voltage	0.98 x Nominal	0.63 x Vtt	1.02 x Nominal	V	

NOTE:
1. GTLREF is either HACVREF or HDVREF, depending on whether the input is an address or control signal, or a data signal.

Table 6-9. DDR SDRAM Interface DC Characteristics (1)

Symbol	Signal Group	Parameter	Min	Nom	Max	Unit	Notes
$V_{IL(DC)}$	(h)	DDR Input Low Voltage (DC)			DDRVREF – 0.075	V	3
$V_{IH(DC)}$	(h)	DDR Input High Voltage (DC)	DDRVREF + 0.075			V	3
$V_{IL(AC)}$	(h)	DDR Input Low Voltage (AC)			DDRVREF – 0.175	V	3
$V_{IH(AC)}$	(h)	DDR Input High Voltage (AC)	DDRVREF + 0.175			V	3
V_{OL}	(h), (i)	DDR Output Low Voltage	0		0.414	V	4
V_{OH}	(h), (i)	DDR Output High Voltage	2.0		Vdd_DDR	V	4
C_{PIN}	(h), (i)	DDR Pin Capacitance	2.5		3.75	pF	
I_{OL}	(h), (i)	DDR Output Low Current			20.7	mA	5
I_{OH}	(h), (i)	DDR Output High Current			18	mA	5
I_{Leak}	(h)	Input Leakage Current			1	µA	
DDRVREF	(k)	DDR Reference Voltage		Vdd_DDR / 2		V	

NOTES:
1. Note that these input voltages apply only when the signals are inputs to the MCH. When the signals are inputs to the SDRAM, the SDRAM input voltage specifications apply.
2. Output voltage level extremes are specified with into a 30 ohm termination load to the DDR bus termination voltage, and are guaranteed by design (not 100% tested) for DDR operation.
3. Maximum currents are specified into a 30 ohm termination load to the DDR bus termination voltage, and are guaranteed by design (not 100% tested) for DDR operation.

Table 6-10. DDR2 SDRAM Interface DC Characteristics (Sheet 1 of 2) (1,4)

Symbol	Signal Group	Parameter	Min	Nom	Max	Unit	Notes
$V_{IL(DC)}$	(h)	DDR Input Low Voltage (DC)			DDRVREF – 0.075	V	3
$V_{IH(DC)}$	(h)	DDR Input High Voltage (DC)	DDRVREF + 0.075			V	3
$V_{IL(AC)}$	(h)	DDR Input Low Voltage (AC)			DDRVREF – 0.175	V	3
$V_{IH(AC)}$	(h)	DDR Input High Voltage (AC)	DDRVREF + 0.175			V	3
V_{OL}	(h), (i)	DDR Output Low Voltage	0		0.414	V	
V_{OH}	(h), (i)	DDR Output High Voltage	1.3			V	
C_{PIN}	(h), (i)	DDR Pin Capacitance	2.5		3.75	pF	
I_{OL}	(h), (i)	DDR Output Low Current			18	mA	

Table 6-10. DDR2 SDRAM Interface DC Characteristics (Sheet 2 of 2) (1,4)

Symbol	Signal Group	Parameter	Min	Nom	Max	Unit	Notes
I_{OH}	(h), (i)	DDR Output High Current			20.7	mA	
I_{Leak}	(h)	Input Leakage Current			1	µA	
DDRVREF	(k)	DDR Reference Voltage	0.49 x Vdd_DDR2	0.50 x Vdd_DDR2	0.51 x Vdd_DDR2	V	

NOTES:
1. Refer to the SSTL_18 specification for further details.
2. Note that these input voltages apply only when the signals are inputs to the MCH. When the signals are inputs to the SDRAM, the SDRAM input voltage specifications apply.
3. DDR2 DC parameters are specified into a 43 ohm test load to Vdd/2

Table 6-11. PCI Express* Differential Transmitter (TX) Output DC Specifications

Symbol	Parameter	Min	Nom	Max	Unit	Notes
$V_{TX-CM-DC-ACTIVE-IDLE-DELTA}$	Absolute Delta of DC Common Mode Voltage During L0 and Electrical Idle	0		100	mV	1
$V_{TX-CM-DC-LINE-DELTA}$	Absolute Delta of DC Common Mode Voltage between D+ and D–	0		25	mV	1
$V_{TX-DC-CM}$	TX DC Common Mode Voltage	0		3.6	V	2
$I_{TX-SHORT}$	Short Circuit Current Limit			90	mA	4
$Z_{TX-DIFF-DC}$	DC Differential TX Impedance	80	100	120	Ω	
Z_{TX-DC}	Transmitter DC impedance	40			Ω	
$L_{TX-SKEW}$	Lane-to-Lane Output Skew			500 ps + 2UI		3
C_{TX}	AC Coupling Capacitor	75		200	nF	

NOTES:
1. Specified at the measurement point into a timing and voltage compliance test load and measured over any 250 consecutive Unit Intervals. Also refer to the Transmitter Compliance Eye Diagram.
2. The allowed DC Common Mode voltage under any conditions. Refer to Section 4.3.1.8 in the *PCI Express Specification* for further details.
3. Static skew between any two Transmitter Lanes within a single link
4. The allowed current when any output is shorted to ground.

Table 6-12. PCI Express* Differential Receiver (RX) Input DC Specifications

Symbol	Parameter	Min	Nom	Max	Unit	Notes
$Z_{RX-DIFF-DC}$	DC Differential Input Impedance	80	100	120	Ω	2
Z_{RX-DC}	DC Input Common Mode Impedance	40	50	60	Ω	1, 2
$Z_{RX-HIGH-IMP-DC}$	Powered Down DC Input Common Mode Impedance	200			kΩ	3
$L_{RX-SKEW}$	Total Skew			20	ns	

NOTES:
1. Specified at the measurement point and measured over any 250 consecutive UIs. The test load in Figure 8 17 (not the MCH itself) should be used as the RX device when taking measurements (also refer to the Receiver Compliance Eye Diagram as shown in Figure 8 19). Refer to document 300312-001 at http://www.pciexpressdevnet.org/apps/org/workgroup/dgl/document.php?document_id=89 for more information on

the difference between taking pin and pad measurements. If the clocks to the RX and TX are not derived from the same clock chip the TX UI must be used as a reference for the eye diagram.
2. Impedance during all LTSSM states. When transitioning from a Fundamental Reset to Detect (the initial state of the LTSSM) there is a 5ms transition time before Receiver termination values must be met on all un-configured Lanes of a Port.
3. The RX DC Common Mode Impedance that exists when no power is present or Fundamental Reset is asserted. This helps ensure that the Receiver Detect circuit will not falsely assume a Receiver is powered on when it is not. This term must be measured at 300mV above the RX ground.

Table 6-13. Hub Interface DC Characteristics

Symbol	Signal Group	Parameter	Min	Nom	Max	Unit	Notes
V_{IL_HI}	(p)	Hub Interface Input Low Voltage		0	HIVREF – 0.1	V	
V_{IH_HI}	(p)	Hub Interface Input High Voltage	HIVREF + 0.1	0.7		V	
V_{OL_HI}	(p)	Hub Interface Output Low Voltage	-0.03	0	0.05	V	
V_{OH_HI}	(p)	Hub Interface Output High Voltage	HISWNG – 0.05	HISWNG	HISWNG + 0.05	V	
I_{IL_HI}	(p)	Hub Interface Input Leakage Current			25	µA	
$C_{IN\_H}I$	(p)	Hub Interface Input Pin Capacitance			5.0	pF	
ΔC_{IN}		Strobe to Data Pin Capacitance delta	– 0.5		0.5	pF	
Z_{PD}	(p)	Pull-down Impedance	45	50	55	Ω	1
Z_{PU}	(p)	Pull-up Impedance	40	45	50	Ω	1
HIVREF	(r)	Hub Interface Reference Voltage	0.98 x Nominal	0.236 x Vcc_MCH	1.02 x Nominal		2
HISWNG	(s)	Hub Interface Swing Reference Voltage	0.98 x Nominal	0.536 x Vcc_MCH	1.02 x Nominal		2
HIRCOMP	(t)	Hub Interface Compensation Resistance	0.99 x Nominal	43.2	1.01 x Nominal		

NOTES:
1. These values depend on the HIRCOMP value chosen; the numbers specified here rely on the HIRCOMP value specified.

Table 6-14. Clock DC Characteristics (Sheet 1 of 2) (1)

Symbol	Signal Group	Parameter	Min	Nom	Max	Unit	Notes
FSB Clock (HCLKINN/HCLKINP), 100 MHz Reference Clock (EXP_CLKN/EXP_CLKP)							
V_{IL}	(f), (m)	Input Low Voltage	–0.150	0		V	
V_{IH}	(f), (m)	Input High Voltage	0.660	0.710	0.850	V	
$V_{CROSS(abs)}$	(f), (m)	Absolute Crossing Point	0.250		0.550	V	2, 8
$V_{CROSS(rel)}$	(f), (m)	Relative Crossing Point	0.250 + 0.5x(V_{Havg} – 0.710)		0.550 + 0.5x(V_{Havg} – 0.710)		
ΔV_{CROSS}	(f), (m)	Range of Crossing Points			0.140		
V_{OS}	(f), (m)	Overshoot			V_{IH} + 0.300	V	4
V_{US}	(f), (m)	Undershoot	–0.300			V	5

Table 6-14. Clock DC Characteristics (Sheet 2 of 2) (1)

Symbol	Signal Group	Parameter	Min	Nom	Max	Unit	Notes
V_{RBM}	(f), (m)	Ringback Margin	0.200			V	6
V_{TR}	(f), (m)	Threshold Region	V_{CROSS} – 0.100		V_{CROSS} + 0.100	V	7
Hub Interface Clock							
V_{IL}	(q)	CLK66 Input Low Voltage			0.8	V	
V_{IH}	(q)	CLK66 Input High Voltage	2.4			V	

NOTES:
1. Refer to Figure 8 1. Differential Clock Waveform and Figure 8 2. Differential Clock Crosspoint Specification.
2. Crossing voltage is defined as the instantaneous voltage when the rising edge of BCLK0 is equal to the falling edge of BCLK1.
3. VHavg is the statistical average of the VH measured by the oscilloscope.
4. Overshoot is defined as the absolute value of the maximum voltage.
5. Undershoot is defined as the absolute value of the minimum voltage.
6. Ringback Margin is defined as the absolute voltage difference between the maximum Rising Edge Ringback and the maximum Falling Edge Ringback.
7. Threshold Region is defined as a region entered around the crossing point voltage in which the differential receiver switches. It includes input threshold hysteresis.
8. The crossing point must meet the absolute and relative crossing point specifications simultaneously.
9. VHavg can be measured directly using “Vtop” on Agilent* scopes and “High” on Tektronix* scopes.

Table 6-15. SMBus I/O DC Characteristics

Symbol	Signal Group	Parameter	Min	Max	Unit	Notes
V_{IH}	(u)	Input High Voltage	2.1	V3REF+0.3	V	
V_{IL}	(u)	Input Low Voltage	-0.3	0.8	V	
V_{OL}	(u)	Output Low Voltage		0.4	V	Ipullup = 4 mA
I_{OL}	(u)	Output Low Current		4	mA	Vol = 0.4 V
I_{Leak}	(u)	Leakage Current		10	µA	
C_{pad}	(u)	Pad Capacitance		5	pF	Pad Only

Table 6-16. JTAG I/O DC Characteristics

Symbol	Signal Group	Parameter	Min	Max	Unit	Notes
V_{IH}	(v)	Input High Voltage	1		V	
V_{IL}	(v)	Input Low Voltage		0.5	V	
V_{OL}	(v)	Output Low Voltage		0.4	V	Rpullup = 50 Ω
I_{Leak}	(v)	Leakage Current		2.9	µA	

Table 6-17. Miscellaneous I/O DC Characteristics

Symbol	Signal Group	Parameter	Min	Max	Unit	Notes
V_{IH}	(w)	Input High Voltage	1	Vcc_MCH + 0.3	V	
V_{IL}	(w)	Input Low Voltage	–0.3	0.5	V	
I_{Leak}	(w)	Leakage Current		90	µA	
V_{IH_PLLSEL}	(y)	Input High Voltage		0.35	V	
V_{IL_PLLSEL}	(y)	Input Low Voltage	1		V	
I_{L_PLLSEL}	(y)	Leakage Current		10	µA	

Ballout and Package Specifications 7

This chapter provides the ballout and package dimensions for the MCH. In addition, internal component package trace lengths to enable trace length compensation are listed.

7.1 Ballout

Figure 7-1 shows a top view of the 1077 ball MCH ballout footprint. Figure 7-1 and Figure 7-2 expand the detail of the ballout footprint to a list of signal names for each ball. Table 7-3 lists the MCH ballout with the listing organized alphabetically by signal name.

Figure 7-1. MCH Ballout Diagram (Top View)

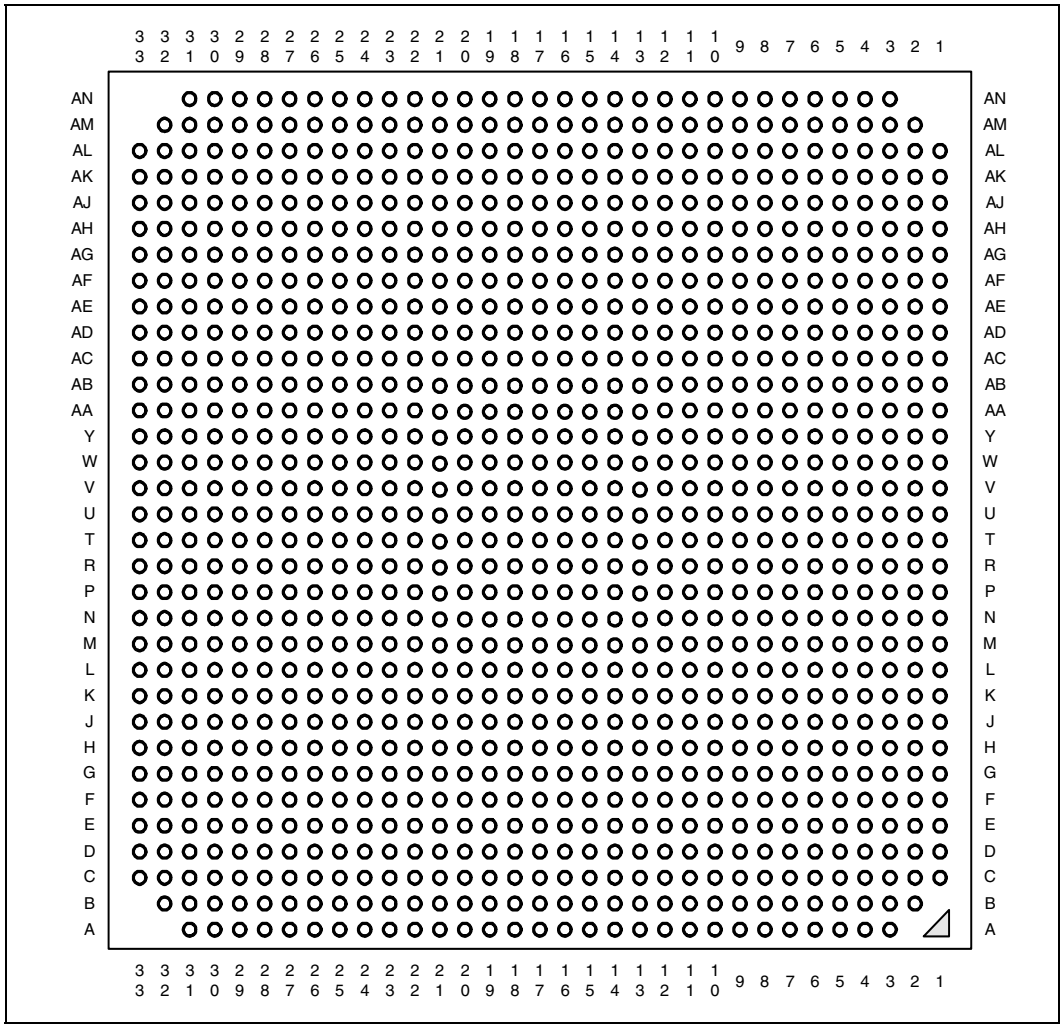

Table 7-1. MCH Ballout (Left Half – Top View)

	33	32	31	30	29	28	27	26	25	24	23	22	21	20	19	18	17
AN			VSS	DDRB_DQ[1]	DDRB_DQSN[0]	VCCDDR	DDRB_DQ[2]	DDRCKE[1]	VSS	DDRB_DQ[9]	DDRB_DQSN[1]	VCCDDR	DDRB_DQ[10]	DDRA_MA[7]	VSS	DDRB_DQSP[11]	DDRB_DQSN[11]
AM		VSS	DDRB_DQ[5]	DDRB_DQ[0]	VSS	DDRB_DQSP[0]	DDRB_DQ[3]	VSS	DDRB_BA[2]	DDRB_DQ[8]	VSS	DDRB_DQSP[1]	DDRB_DQ[11]	VSS	DDRB_DQ[21]	DDRB_DQ[17]	VSS
AL	VSS	DDRA_DQS[9]	DDRA_DQSN[9]	VCCDDR	DDRB_DQSN[9]	DDRB_DQ[6]	VSS	DDRCKE[2]	DDRB_DQ[12]	VCCDDR	DDRB_DQSN[10]	DDRB_DQ[14]	VSS	DDRB_MA[8]	DDRB_DQ[20]	VCCDDR	DDRB_DQSN[2]
AK	DDRA_DQ[5]	DDRA_DQ[0]	VSS	DDRB_DQ[4]	DDRB_DQSP[9]	VSS	DDRB_DQ[7]	DDRCKE[3]	VSS	DDRB_DQ[13]	DDRB_DQSP[10]	VSS	DDRB_DQ[15]	DDRA_MA[8]	VSS	DDRB_DQ[16]	DDRB_DQSP[2]
AJ	DDRA_DQ[4]	VSS	DDRA_DQSN[0]	DDRA_DQSP[0]	VSS	DDRCKE[6]	DDRCKE[5]	VCCDDR	DDRA_DQSN[1]	DDRA_DQSP[1]	VSS	DDRA_MA[9]	DDRB_MA[9]	VCCDDR	RESERVED	DDRA_DQ[22]	VSS
AH	VCCEXP	EXP_B_TXN[0]	DDRA_DQ[1]	VSS	DDRA_DQ[2]	DDRCKE[7]	VSS	DDRCKE[4]	DDRA_DQ[14]	VSS	DDRB_MA[12]	DDRB_MA[11]	VSS	DDRA_DQSN[2]	DDRA_DQSP[2]	VSS	DDRA_MA[5]
AG	EXP_B_RXP[0]	EXP_B_TXP[0]	VSS	DDRA_DQ[6]	DDRA_DQ[7]	VCCDDR	DDRA_DQ[8]	DDRA_DQ[9]	VSS	DDRA_DQ[15]	DDRA_MA[12]	VCCDDR	DDRA_DQ[17]	DDRA_DQ[23]	VSS	DDRA_DQ[19]	DDRB_MA[6]
AF	EXP_B_RXN[0]	VSS	EXP_B_TXP[1]	RESERVED	VSS	DDRA_DQ[3]	DDRA_DQ[13]	VSS	DDRA_DQSP[10]	DDRA_DQSN[10]	VSS	DDRA_MA[11]	DDRA_DQ[16]	VSS	DDRA_DQ[18]	DDRA_MA[6]	VSS
AE	VCCEXP	EXP_B_RXP[1]	EXP_B_TXN[1]	VCCEXP	EXP_B_RXN[4]	DDRA_DQ[12]	VSS	DDRCKE[0]	DDRA_BA[2]	VCCDDR	RESERVED	DDRA_DQ[20]	VSS	DDRA_DQSP[11]	DDRA_DQSN[11]	VCCDDR	DDRA_DQ[28]
AD	EXP_B_TXN[2]	EXP_B_RXN[1]	VSS	EXP_B_RXN[2]	EXP_B_RXP[4]	VSS	EXP_B_TXP[4]	EXP_B_TXN[4]	VSS	DDRA_DQ[10]	DDRA_DQ[11]	VSS	DDRA_DQ[21]	RESERVED	VSS	DDRB_MA[7]	DDRA_DQ[31]
AC	EXP_B_TXP[2]	VSS	EXP_B_RXP[3]	EXP_B_RXP[2]	VCCEXP	EXP_B_TXN[5]	EXP_B_TXP[5]	VSS	EXP_B_RXP[5]	EXP_B_RXN[5]	VCCDDR	VSS	VCCDDR	VSS	VCCDDR	VSS	VCCDDR
AB	VCCEXP	EXP_B_TXP[3]	EXP_B_RXN[3]	VSS	EXP_B_TXP[6]	EXP_B_TXN[6]	VCCEXP	EXP_B_RXP[6]	EXP_B_RXN[6]	VCCEXP	VSS	VCCDDR	VSS	VCCDDR	VSS	VCCDDR	VSS
AA	EXP_B_TXN[10]	EXP_B_TXN[3]	VSS	EXP_B_RXP[10]	EXP_B_RXN[10]	VSS	EXP_B_TXP[7]	EXP_B_TXN[7]	VSS	RESERVED	VCCEXP	VSS	VCC	VSS	VCC	VSS	VCC
Y	EXP_B_TXP[10]	VSS	EXP_B_RXN[9]	EXP_B_RXP[9]	VCCEXP	EXP_B_RXP[8]	EXP_B_RXN[8]	VSS	EXP_B_RXP[7]	EXP_B_RXN[7]	VSS	VCCEXP	VSS	VCC	VSS	VCC	VSS
W	VCCEXP	EXP_B_TXP[11]	EXP_B_TXN[11]	VSS	EXP_B_TXN[9]	EXP_B_TXP[9]	VCCEXP	EXP_B_TXP[8]	EXP_B_TXN[8]	VSS	VCCEXP	VSS	VCC	VSS	VCC	VSS	VCC
V	EXP_B_RXP[11]	EXP_B_RXN[11]	VSS	EXP_B_TXP[13]	EXP_B_TXN[13]	VSS	EXP_B_RXP[14]	EXP_B_RXN[14]	VSS	EXP_B_RXP[15]	VSSA_EXP	VCCEXP	VSS	VCC	VSS	VCC	VSS
U	EXP_COMP[0]	VSS	EXP_B_TXP[12]	EXP_B_TXN[12]	VCCEXP	EXP_VSSBG	EXP_VCCBG	VSS	EXP_COMP[1]	EXP_B_RXN[15]	VCCA_EXP	VSS	VCC	VSS	VCC	VSS	VCC
T	VCCEXP	EXP_B_RXP[12]	EXP_B_RXN[12]	VSS	EXP_B_TXP[14]	EXP_B_TXN[14]	VCCEXP	EXP_B_TXP[15]	EXP_B_TXN[15]	VCCEXP	EXP_CLK_P	VCCEXP	VSS	VCC	VSS	VCC	VSS
R	EXP_A_RXP[0]	RESERVED	VSS	EXP_B_RXP[13]	EXP_B_RXN[13]	VSS	EXP_A_RXN[4]	EXP_A_RXP[4]	VSS	EXP_CLK_N	VCCEXP	VSS	VCC	VSS	VCC	VSS	VCC
P	EXP_A_RXN[0]	VSS	EXP_A_TXN[0]	EXP_A_TXP[0]	VCCEXP	EXP_A_TXN[5]	EXP_A_TXP[5]	VSS	EXP_A_TXN[4]	EXP_A_TXP[4]	VSS	VCC	VSSA_HI	VCCA_HI	VSS	VCC	VSS
N	VCCEXP	EXP_A_TXN[1]	EXP_A_TXP[1]	VSS	EXP_A_RXN[1]	EXP_A_RXP[1]	VSS	EXP_A_RXN[5]	EXP_A_RXP[5]	VSS	VCC	VSS	VCC	VSS	VCC	VSS	VCC
M	EXP_A_TXP[2]	EXP_A_TXN[2]	VSS	EXP_A_TXP[6]	EXP_A_TXN[6]	VCCEXP	EXP_A_RXP[6]	EXP_A_RXN[6]	VSS	MCHPME#	VSS	VCC	VSS	VTT	VSS	VTT	VSS
L	RESERVED	VSS	EXP_A_RXP[2]	EXP_A_RXN[2]	VSS	EXP_A_TXP[7]	EXP_A_TXN[7]	VSS	MCHGPE#	HICLK	VCC	VSS	VTT	VSS	VTT	VSS	VTT
K	VCCEXP	EXP_A_TXP[3]	EXP_A_TXN[3]	VCCEXP	EXP_A_RXP[7]	EXP_A_RXN[7]	VCC	HI[6]	HICOMP	VSS	HREQ[4]#	HA[3]#	VSS	HREQ[0]#	HA[8]#	VSS	HD[21]#
J	EXP_A_RXP[3]	EXP_A_RXN[3]	VSS	HI[0]	HI[10]	VSS	HI[7]	RSP#	VSS	CPURST#	HREQ[2]#	VSS	HREQ[1]#	HA[4]#	VTT	HD[17]#	HD[23]#
H	V3REF	VSS	HI_VSWING	HI[1]	VCC	HI[5]	DBSY#	VTT	AP[1]#	MCERR#	VTT	HREQ[3]#	HA[7]#	VSS	HA[9]#	HD[18]#	VSS
G	VCC	HI[11]	HI[3]	VSS	HI[4]	RS[2]#	VSS	BINIT#	AP[0]#	VSS	HA[5]#	HA[6]#	VSS	HADSTB[0]#	HA[10]#	VSS	HD[20]#
F	RESERVED	HIVREF	VSS	HI[8]	RS[0]#	VSS	DEP[3]#	HSLWCRES	VSS	BREQ[0]#	HACVREF	VSS	HA[15]#	HA[16]#	VSS	HA[13]#	HD[19]#
E	HI[9]	VSS	HI_STBS	HIT#	VSS	DEP[1]#	HODTCRES	VSS	DEP[2]#	VSS	VTT	HA[11]#	HA[12]#	VTT	HA[14]#	HD[16]#	VTT
D	RESERVED	HI_STBF	RS[1]#	VTT	BREQ[1]#	HITM#	VTT	HA[17]#	HA[24]#	VSS	HA[28]#	HA[20]#	VSS	HA[35]#	HA[32]#	VSS	HD[8]#
C	VCC	HI[2]	PLLSEL[0]#	HLOCK#	DEP[0]#	VSS	HCRES0	HA[18]#	VSS	HA[25]#	HADSTB[1]#	VSS	HA[34]#	HA[33]#	VSS	HD[0]#	HD[6]#
B		VSS	BNR#	DRDY#	VSS	DEFER#	ADS#	VSS	HA[23]#	HA[30]#	VSS	HA[21]#	HA[27]#	VSS	HD[1]#	HD[7]#	VSS
A			VTT	HTRDY#	PLLSEL[1]#	BPRI#	VSS	HA[19]#	HA[22]#	VSS	HA[29]#	HA[26]#	VSS	HA[31]#	HD[4]#	VTT	HD[3]#
	33	32	31	30	29	28	27	26	25	24	23	22	21	20	19	18	17

Table 7-2. MCH Ballout (Right Half – Top View)

16	15	14	13	12	11	10	9	8	7	6	5	4	3	2	1	
VCCDDR	DDRB_DQ[18]	DDRB_MA[2]	VSS	DDRB_DQSP[12]	DDRB_DQSN[12]	VCCDDR	DDRB_DQ[26]	DDRB_CB[4]	VSS	DDRB_DQSN[17]	DDRB_CB[6]	DDRB_VREF	VCCDDR			AN
DDRB_DQ[22]	DDRB_DQ[19]	VSS	DDRB_DQ[29]	DDRB_DQ[25]	VSS	DDRB_DQ[30]	DDRB_DQ[27]	VSS	DDRB_CB[0]	DDRB_DQSP[17]	VSS	DDRB_CB[2]	DDRA_VREF	VSS		AM
DDRB_DQ[23]	VSS	DDRA_MA[2]	DDRB_DQ[28]	VCCDDR	DDRB_DQSN[3]	DDRB_DQ[31]	VSS	DDRB_CMDCLKN[3]	DDRB_CB[1]	VCCDDR	DDRB_CB[7]	DDRB_CB[3]	VSS	DDR_IMPCRES	VCCDDR	AL
VSS	DDRA_MA[3]	DDRB_MA[3]	VSS	DDRB_DQ[24]	DDRB_DQSP[3]	VSS	DDRB_CMDCLKP[3]	DDRB_CB[5]	VSS	DDRB_DQSN[8]	DDRB_DQSP[8]	VSS	DDRB_DQ[36]	DDRB_DQ[37]	DDR_SLWCRES	AK
DDRA_MA[4]	DDRA_DQ[29]	VCCDDR	DDRA_DQSN[12]	DDRA_CMDCLKN[1]	VSS	DDRA_CB[4]	DDRA_CB[0]	VCCDDR	DDRA_DQSN[17]	DDRB_CMDCLKN[0]	VSS	DDRB_DQ[33]	DDRB_DQ[32]	VSS	DDRB_DQSP[13]	AJ
DDRB_MA[5]	VSS	DDRA_DQSP[12]	DDRA_CMDCLKP[2]	VSS	DDRA_CMDCLKP[1]	DDRA_CB[5]	VSS	DDRA_DQSP[17]	DDRB_CMDCLKP[0]	DDRB_CMDCLKP[1]	DDRA_MA[0]	DDRA_DQ[32]	VCCDDR	DDRB_DQSN[4]	DDRB_DQSN[13]	AH
VCCDDR	DDRA_DQSN[3]	DDRA_DQSP[3]	VSS	DDRA_CMDCLKN[2]	DDRA_CB[1]	VCCDDR	DDRA_DQSN[8]	DDRB_CMDCLKP[2]	VSS	DDRB_CMDCLKN[1]	DDRA_DQ[33]	VSS	DDRB_DQ[38]	DDRB_DQSP[4]	VSS	AG
DDRA_DQ[24]	DDRA_DQ[25]	VSS	DDRA_CMDCLKP[0]	DDRA_CMDCLKN[0]	VSS	DDRA_CB[6]	DDRA_DQSP[8]	VSS	DDRB_MA[0]	DDRA_BA[1]	VCCDDR	DDRB_DQ[35]	DDRB_DQ[39]	VSS	DDRB_DQ[34]	AF
DDRA_DQ[30]	VSS	DDRB_MA[1]	DDRA_DQ[26]	VCCDDR	DDRA_CB[2]	DDRA_CB[7]	VSS	DDRB_CMDCLKN[2]	DDRB_BA[1]	VSS	DDRA_DQ[38]	DDRA_MA[10]	VSS	DDRRES[1]	DDRRES[2]	AE
VSS	DDRB_MA[4]	DDRA_MA[1]	VSS	DDRA_DQ[27]	DDRA_CB[3]	VSS	DDRA_CMDCLKN[3]	DDRA_DQSP[13]	VCCDDR	DDRA_DQSN[4]	DDRA_DQ[39]	VSS	DDRB_DQ[45]	DDRB_DQ[44]	VCCDDR	AD
VSS	VCCDDR	VSS	VCCDDR	VSS	VCCDDR	DDRA_CMDCLKP[3]	DDR_CRES0	VSS	DDRA_DQSN[13]	DDRA_DQSP[4]	VSS	DDRB_MA[10]	DDRB_DQ[40]	VSS	DDRB_DQ[41]	AC
VCCDDR	VSS	VCCDDR	VSS	VCCDDR	VSS	DDRA_DQ[36]	VCCDDR	DDRA_DQ[34]	DDRA_DQ[35]	VSS	DDRA_BA[0]	DDRB_DQSN[5]	VCCDDR	DDRB_DQSP[14]	DDRB_DQSN[14]	AB
VSS	VCC	VSS	VCC	VSS	VCCDDR	VSS	DDRA_DQ[37]	DDRB_BA[0]	VSS	DDRA_RAS#	DDRA_DQ[41]	VSS	DDRB_DQSP[5]	DDRB_DQ[46]	VSS	AA
VCC	VSS	VCC	VSS	VCCDDR	VSS	DDRA_WE#	DDRB_RAS#	VSS	DDRA_DQSP[14]	DDRA_DQSN[14]	VCCDDR	DDRB_DQ[43]	DDRB_DQ[42]	VSS	DDRB_DQ[47]	Y
VSS	VCC	VSS	VCC	VSS	VCCDDR	DDRA_DQ[44]	VSS	DDRA_CAS#	DDRA_DQSP[5]	VSS	DDRA_DQ[46]	DDRB_WE#	VSS	DDRA_CS[0]#	DDRB_CAS#	W
VCC	VSS	VCC	VSS	VCCDDR	VSS	VSS	DDRB_CS[0]#	DDRA_DQSN[5]	VCCDDR	DDRA_DQ[42]	DDRA_DQ[47]	VSS	DDRA_CS[1]#	DDRB_CS[1]#	VCCDDR	V
VSS	VCC	VSS	VCC	VSS	VCCDDR	DDRA_DQ[45]	DDRA_DQ[40]	VSS	DDRA_DQ[43]	DDRA_MA[13]	VSS	DDRB_MA[13]	DDRB_DQ[52]	VSS	DDRB_DQ[53]	U
VCC	VSS	VCC	VSS	VCCDDR	VSS	DDRA_CS[3]#	VCCDDR	DDRA_CS[2]#	DDRB_CS[2]#	VSS	DDRA_DQ[53]	DDRB_DQ[48]	VCCDDR	DDRB_DQSP[15]	DDRB_DQ[49]	T
VSS	VCC	VSS	VCC	VSS	VCCDDR	RESERVED	RESERVED	RESERVED	VSS	DDRA_DQ[48]	DDRA_DQ[49]	VSS	DDRB_DQSN[15]	DDRB_DQSN[6]	VSS	R
VCC	VSS	VCC	VSS	VCCDDR	VSS	DDRA_DQSP[15]	DDRA_DQ[52]	VSS	DDRA_DQSN[6]	DDRB_CS[3]#	VCCDDR	DDRB_DQ[55]	DDRB_DQ[54]	VSS	DDRB_DQSP[6]	P
VSS	VCC	VSS	VCC	VSS	VCCDDR	DDRA_DQSN[15]	VSS	DDRA_DQ[54]	DDRA_DQSP[6]	VSS	DDRA_CS[4]#	DDRB_CS[4]#	VSS	DDRB_DQ[51]	DDRB_DQ[50]	N
VTT	VSS	VTT	VSS	VCC	VSS	VSS	DDRA_DQ[55]	RESERVED	VCCDDR	DDRB_CS[6]#	DDRA_CS[5]#	VSS	DDRA_CS[6]#	DDRB_CS[5]#	VCCDDR	M
VSS	VTT	VSS	VTT	TEST#	DEBUG[6]	DDRA_DQ[61]	DDRA_DQ[60]	VSS	DDRA_DQ[50]	DDRA_DQ[51]	VSS	DDRA_CS[7]#	DDRB_CS[7]#	VSS	DDRB_DQ[60]	L
HD[28]#	VSS	HD[31]#	HD[46]#	VSS	HCLKINP	DDRA_DQ[59]	VCC	DDRA_DQ[58]	DDRA_DQ[62]	VSS	DDRA_DQ[56]	DDRB_DQ[61]	VCCDDR	DDRB_DQ[56]	DDRB_DQ[57]	K
VTT	HDSTBP[1]#	HD[24]#	VTT	HD[44]#	HCLKINN	VSS	TRST#	DEBUG[0]	VSS	DDRA_DQSP[16]	DDRA_DQ[57]	VSS	DDRB_DQSP[16]	DDRB_DQSN[16]	VSS	J
HD[29]#	HDSTBN[1]#	VSS	HD[34]#	HD[47]#	VSS	HD[42]#	DEBUG[3]	VSS	DDRA_DQ[63]	DDRA_DQSN[16]	VCCDDR	DDRA_DQSPPN[7]	DDRB_DQSP[7]	VSS	DDRB_DQSN[7]	H
HD[27]#	VSS	HD[30]#	HD[36]#	VSS	HD[45]#	HD[40]#	VSS	DEBUG[2]	DEBUG[1]	TDO	TDI	DDRA_DQSPP[7]	VSS	DDRB_DQ[63]	DDRB_DQ[62]	G
VSS	HD[26]#	HD[25]#	VSS	HD[39]#	HD[35]#	VSS	DBI[2]#	HD[43]#	VSS	VCCA_CORE	VSSA_CORE	VSS	TMS	DDRB_DQ[58]	VCCDDR	F
HD[22]#	DBI[1]#	VTT	HDVREF[1]	HD[32]#	VTT	HDSTBP[2]#	HD[38]#	VTT	HD[56]#	RESERVED	VSS	VCCA_DDR	PWRGD	VSS	DDRB_DQ[59]	E
DBI[0]#	VSS	HD[13]#	HDVREF[0]	VSS	HD[37]#	HDSTBN[2]#	VSS	HD[41]#	HD[62]#	VSS	HD[60]#	SMBDATA	DEBUG[5]	TCK	DEBUG[7]	D
VSS	HDSTBP[0]#	HD[2]#	VSS	HD[14]#	HD[33]#	VSS	HD[51]#	HD[53]#	VSS	HD[61]#	HD[63]#	VSS	SMBCLK	RSTIN#	VSS	C
HD[5]#	HDSTBN[0]#	VSS	HD[10]#	HD[15]#	VSS	HD[48]#	HD[52]#	VSS	HD[55]#	HD[54]#	VSS	HD[57]#	HD[59]#	DEBUG[4]		B
HD[9]#	VTT	HD[11]#	HD[12]#	VTT	HD[50]#	HD[49]#	VTT	HDSTBN[3]#	HDSTBP[3]#	VTT	DBI[3]#	HD[58]#	VTT			A
16	15	14	13	12	11	10	9	8	7	6	5	4	3	2	1	

Table 7-3. MCH Signal Listing (Alphabetical by Signal Name) (Sheet 1 of 8)

Ball #	Signal Name
B27	ADS#
G25	AP[0]#
H25	AP[1]#
G26	BINIT#
B31	BNR#
A28	BPRI#
F24	BREQ[0]#
D29	BREQ[1]#
J24	CPURST#
D16	DBI[0]#
E15	DBI[1]#
F9	DBI[2]#
A5	DBI[3]#
H27	DBSY#
AC9	DDR_CRES0
AL2	DDR_IMPCRES
AK1	DDR_SLWCRES
AB5	DDRA_BA[0]
AF6	DDRA_BA[1]
AE25	DDRA_BA[2]
W8	DDRA_CAS#
AJ9	DDRA_CB[0]
AG11	DDRA_CB[1]
AE11	DDRA_CB[2]
AD11	DDRA_CB[3]
AJ10	DDRA_CB[4]
AH10	DDRA_CB[5]
AF10	DDRA_CB[6]
AE10	DDRA_CB[7]
AF12	DDRA_CMDCLKN[0]
AJ12	DDRA_CMDCLKN[1]
AG12	DDRA_CMDCLKN[2]
AD9	DDRA_CMDCLKN[3]
AF13	DDRA_CMDCLKP[0]
AH11	DDRA_CMDCLKP[1]
AH13	DDRA_CMDCLKP[2]
AC10	DDRA_CMDCLKP[3]
W2	DDRA_CS[0]#
V3	DDRA_CS[1]#
T8	DDRA_CS[2]#
T10	DDRA_CS[3]#
N5	DDRA_CS[4]#
M5	DDRA_CS[5]#
M3	DDRA_CS[6]#
L4	DDRA_CS[7]#

Ball #	Signal Name
AK32	DDRA_DQ[0]
AH31	DDRA_DQ[1]
AH29	DDRA_DQ[2]
AF28	DDRA_DQ[3]
AJ33	DDRA_DQ[4]
AK33	DDRA_DQ[5]
AG30	DDRA_DQ[6]
AG29	DDRA_DQ[7]
AG27	DDRA_DQ[8]
AG26	DDRA_DQ[9]
AD24	DDRA_DQ[10]
AD23	DDRA_DQ[11]
AE28	DDRA_DQ[12]
AF27	DDRA_DQ[13]
AH25	DDRA_DQ[14]
AG24	DDRA_DQ[15]
AF21	DDRA_DQ[16]
AG21	DDRA_DQ[17]
AF19	DDRA_DQ[18]
AG18	DDRA_DQ[19]
AE22	DDRA_DQ[20]
AD21	DDRA_DQ[21]
AJ18	DDRA_DQ[22]
AG20	DDRA_DQ[23]
AF16	DDRA_DQ[24]
AF15	DDRA_DQ[25]
AE13	DDRA_DQ[26]
AD12	DDRA_DQ[27]
AE17	DDRA_DQ[28]
AJ15	DDRA_DQ[29]
AE16	DDRA_DQ[30]
AD17	DDRA_DQ[31]
AH4	DDRA_DQ[32]
AG5	DDRA_DQ[33]
AB8	DDRA_DQ[34]
AB7	DDRA_DQ[35]
AB10	DDRA_DQ[36]
AA9	DDRA_DQ[37]
AE5	DDRA_DQ[38]
AD5	DDRA_DQ[39]
U9	DDRA_DQ[40]
AA5	DDRA_DQ[41]
V6	DDRA_DQ[42]
U7	DDRA_DQ[43]
W10	DDRA_DQ[44]

Ball #	Signal Name
U10	DDRA_DQ[45]
W5	DDRA_DQ[46]
V5	DDRA_DQ[47]
R6	DDRA_DQ[48]
R5	DDRA_DQ[49]
L7	DDRA_DQ[50]
L6	DDRA_DQ[51]
P9	DDRA_DQ[52]
T5	DDRA_DQ[53]
N8	DDRA_DQ[54]
M9	DDRA_DQ[55]
K5	DDRA_DQ[56]
J5	DDRA_DQ[57]
K8	DDRA_DQ[58]
K10	DDRA_DQ[59]
L9	DDRA_DQ[60]
L10	DDRA_DQ[61]
K7	DDRA_DQ[62]
H7	DDRA_DQ[63]
AJ31	DDRA_DQSN[0]
AJ25	DDRA_DQSN[1]
AH20	DDRA_DQSN[2]
AG15	DDRA_DQSN[3]
AD6	DDRA_DQSN[4]
V8	DDRA_DQSN[5]
P7	DDRA_DQSN[6]
H4	DDRA_DQSN[7]
AG9	DDRA_DQSN[8]
AL31	DDRA_DQSN[9]
AF24	DDRA_DQSN[10]
AE19	DDRA_DQSN[11]
AJ13	DDRA_DQSN[12]
AC7	DDRA_DQSN[13]
Y6	DDRA_DQSN[14]
N10	DDRA_DQSN[15]
H6	DDRA_DQSN[16]
AJ7	DDRA_DQSN[17]
AJ30	DDRA_DQSP[0]
AJ24	DDRA_DQSP[1]
AH19	DDRA_DQSP[2]
AG14	DDRA_DQSP[3]
AC6	DDRA_DQSP[4]
W7	DDRA_DQSP[5]
N7	DDRA_DQSP[6]
G4	DDRA_DQSP[7]

Table 7-3. MCH Signal Listing (Alphabetical by Signal Name) (Sheet 2 of 8)

Ball #	Signal Name
AF9	DDRA_DQSP[8]
AL32	DDRA_DQSP[9]
AF25	DDRA_DQSP[10]
AE20	DDRA_DQSP[11]
AH14	DDRA_DQSP[12]
AD8	DDRA_DQSP[13]
Y7	DDRA_DQSP[14]
P10	DDRA_DQSP[15]
J6	DDRA_DQSP[16]
AH8	DDRA_DQSP[17]
AH5	DDRA_MA[0]
AD14	DDRA_MA[1]
AL14	DDRA_MA[2]
AK15	DDRA_MA[3]
AJ16	DDRA_MA[4]
AH17	DDRA_MA[5]
AF18	DDRA_MA[6]
AN20	DDRA_MA[7]
AK20	DDRA_MA[8]
AJ22	DDRA_MA[9]
AE4	DDRA_MA[10]
AF22	DDRA_MA[11]
AG23	DDRA_MA[12]
U6	DDRA_MA[13]
AA6	DDRA_RAS#
AM3	DDRA_VREF
Y10	DDRA_WE#
AA8	DDRB_BA[0]
AE7	DDRB_BA[1]
AM25	DDRB_BA[2]
W1	DDRB_CAS#
AM7	DDRB_CB[0]
AL7	DDRB_CB[1]
AM4	DDRB_CB[2]
AL4	DDRB_CB[3]
AN8	DDRB_CB[4]
AK8	DDRB_CB[5]
AN5	DDRB_CB[6]
AL5	DDRB_CB[7]
AJ6	DDRB_CMDCLKN[0]
AG6	DDRB_CMDCLKN[1]
AE8	DDRB_CMDCLKN[2]
AL8	DDRB_CMDCLKN[3]
AH7	DDRB_CMDCLKP[0]
AH6	DDRB_CMDCLKP[1]
AG8	DDRB_CMDCLKP[2]
AK9	DDRB_CMDCLKP[3]
V9	DDRB_CS[0]#
V2	DDRB_CS[1]#
T7	DDRB_CS[2]#
P6	DDRB_CS[3]#
N4	DDRB_CS[4]#
M2	DDRB_CS[5]#
M6	DDRB_CS[6]#
L3	DDRB_CS[7]#
AM30	DDRB_DQ[0]
AN30	DDRB_DQ[1]
AN27	DDRB_DQ[2]
AM27	DDRB_DQ[3]
AK30	DDRB_DQ[4]
AM31	DDRB_DQ[5]
AL28	DDRB_DQ[6]
AK27	DDRB_DQ[7]
AM24	DDRB_DQ[8]
AN24	DDRB_DQ[9]
AN21	DDRB_DQ[10]
AM21	DDRB_DQ[11]
AL25	DDRB_DQ[12]
AK24	DDRB_DQ[13]
AL22	DDRB_DQ[14]
AK21	DDRB_DQ[15]
AK18	DDRB_DQ[16]
AM18	DDRB_DQ[17]
AN15	DDRB_DQ[18]
AM15	DDRB_DQ[19]
AL19	DDRB_DQ[20]
AM19	DDRB_DQ[21]
AM16	DDRB_DQ[22]
AL16	DDRB_DQ[23]
AK12	DDRB_DQ[24]
AM12	DDRB_DQ[25]
AN9	DDRB_DQ[26]
AM9	DDRB_DQ[27]
AL13	DDRB_DQ[28]
AM13	DDRB_DQ[29]
AM10	DDRB_DQ[30]
AL10	DDRB_DQ[31]
AJ3	DDRB_DQ[32]
AJ4	DDRB_DQ[33]
AF1	DDRB_DQ[34]
AF4	DDRB_DQ[35]
AK3	DDRB_DQ[36]
AK2	DDRB_DQ[37]
AG3	DDRB_DQ[38]
AF3	DDRB_DQ[39]
AC3	DDRB_DQ[40]
AC1	DDRB_DQ[41]
Y3	DDRB_DQ[42]
Y4	DDRB_DQ[43]
AD2	DDRB_DQ[44]
AD3	DDRB_DQ[45]
AA2	DDRB_DQ[46]
Y1	DDRB_DQ[47]
T4	DDRB_DQ[48]
T1	DDRB_DQ[49]
N1	DDRB_DQ[50]
N2	DDRB_DQ[51]
U3	DDRB_DQ[52]
U1	DDRB_DQ[53]
P3	DDRB_DQ[54]
P4	DDRB_DQ[55]
K2	DDRB_DQ[56]
K1	DDRB_DQ[57]
F2	DDRB_DQ[58]
E1	DDRB_DQ[59]
L1	DDRB_DQ[60]
K4	DDRB_DQ[61]
G1	DDRB_DQ[62]
G2	DDRB_DQ[63]
AN29	DDRB_DQSN[0]
AN23	DDRB_DQSN[1]
AL17	DDRB_DQSN[2]
AL11	DDRB_DQSN[3]
AH2	DDRB_DQSN[4]
AB4	DDRB_DQSN[5]
R2	DDRB_DQSN[6]
H1	DDRB_DQSN[7]
AK6	DDRB_DQSN[8]
AL29	DDRB_DQSN[9]
AL23	DDRB_DQSN[10]
AN17	DDRB_DQSN[11]
AN11	DDRB_DQSN[12]
AH1	DDRB_DQSN[13]
AB1	DDRB_DQSN[14]
R3	DDRB_DQSN[15]
J2	DDRB_DQSN[16]
AN6	DDRB_DQSN[17]
AM28	DDRB_DQSP[0]

Table 7-3. MCH Signal Listing (Alphabetical by Signal Name) (Sheet 3 of 8)

Ball #	Signal Name
AM22	DDRB_DQSP[1]
AK17	DDRB_DQSP[2]
AK11	DDRB_DQSP[3]
AG2	DDRB_DQSP[4]
AA3	DDRB_DQSP[5]
P1	DDRB_DQSP[6]
H3	DDRB_DQSP[7]
AK5	DDRB_DQSP[8]
AK29	DDRB_DQSP[9]
AK23	DDRB_DQSP[10]
AN18	DDRB_DQSP[11]
AN12	DDRB_DQSP[12]
AJ1	DDRB_DQSP[13]
AB2	DDRB_DQSP[14]
T2	DDRB_DQSP[15]
J3	DDRB_DQSP[16]
AM6	DDRB_DQSP[17]
AF7	DDRB_MA[0]
AE14	DDRB_MA[1]
AN14	DDRB_MA[2]
AK14	DDRB_MA[3]
AD15	DDRB_MA[4]
AH16	DDRB_MA[5]
AG17	DDRB_MA[6]
AD18	DDRB_MA[7]
AL20	DDRB_MA[8]
AJ21	DDRB_MA[9]
AC4	DDRB_MA[10]
AH22	DDRB_MA[11]
AH23	DDRB_MA[12]
U4	DDRB_MA[13]
Y9	DDRB_RAS#
AN4	DDRB_VREF
W4	DDRB_WE#
AE26	DDRCKE[0]
AN26	DDRCKE[1]
AL26	DDRCKE[2]
AK26	DDRCKE[3]
AH26	DDRCKE[4]
AJ27	DDRCKE[5]
AJ28	DDRCKE[6]
AH28	DDRCKE[7]
AE2	DDRRES[1]
AE1	DDRRES[2]
J8	DEBUG[0]
G7	DEBUG[1]
G8	DEBUG[2]
H9	DEBUG[3]
B2	DEBUG[4]
D3	DEBUG[5]
L11	DEBUG[6]
D1	DEBUG[7]
B28	DEFER#
C29	DEP[0]#
E28	DEP[1]#
E25	DEP[2]#
F27	DEP[3]#
B30	DRDY#
P33	EXP_A_RXN[0]
N29	EXP_A_RXN[1]
L30	EXP_A_RXN[2]
J32	EXP_A_RXN[3]
R27	EXP_A_RXN[4]
N26	EXP_A_RXN[5]
M26	EXP_A_RXN[6]
K28	EXP_A_RXN[7]
R33	EXP_A_RXP[0]
N28	EXP_A_RXP[1]
L31	EXP_A_RXP[2]
J33	EXP_A_RXP[3]
R26	EXP_A_RXP[4]
N25	EXP_A_RXP[5]
M27	EXP_A_RXP[6]
K29	EXP_A_RXP[7]
P31	EXP_A_TXN[0]
N32	EXP_A_TXN[1]
M32	EXP_A_TXN[2]
K31	EXP_A_TXN[3]
P25	EXP_A_TXN[4]
P28	EXP_A_TXN[5]
M29	EXP_A_TXN[6]
L27	EXP_A_TXN[7]
P30	EXP_A_TXP[0]
N31	EXP_A_TXP[1]
M33	EXP_A_TXP[2]
K32	EXP_A_TXP[3]
P24	EXP_A_TXP[4]
P27	EXP_A_TXP[5]
M30	EXP_A_TXP[6]
L28	EXP_A_TXP[7]
AF33	EXP_B_RXN[0]
AD32	EXP_B_RXN[1]
AD30	EXP_B_RXN[2]
AB31	EXP_B_RXN[3]
AE29	EXP_B_RXN[4]
AC24	EXP_B_RXN[5]
AB25	EXP_B_RXN[6]
Y24	EXP_B_RXN[7]
Y27	EXP_B_RXN[8]
Y31	EXP_B_RXN[9]
AA29	EXP_B_RXN[10]
V32	EXP_B_RXN[11]
T31	EXP_B_RXN[12]
R29	EXP_B_RXN[13]
V26	EXP_B_RXN[14]
U24	EXP_B_RXN[15]
AG33	EXP_B_RXP[0]
AE32	EXP_B_RXP[1]
AC30	EXP_B_RXP[2]
AC31	EXP_B_RXP[3]
AD29	EXP_B_RXP[4]
AC25	EXP_B_RXP[5]
AB26	EXP_B_RXP[6]
Y25	EXP_B_RXP[7]
Y28	EXP_B_RXP[8]
Y30	EXP_B_RXP[9]
AA30	EXP_B_RXP[10]
V33	EXP_B_RXP[11]
T32	EXP_B_RXP[12]
R30	EXP_B_RXP[13]
V27	EXP_B_RXP[14]
V24	EXP_B_RXP[15]
AH32	EXP_B_TXN[0]
AE31	EXP_B_TXN[1]
AD33	EXP_B_TXN[2]
AA32	EXP_B_TXN[3]
AD26	EXP_B_TXN[4]
AC28	EXP_B_TXN[5]
AB28	EXP_B_TXN[6]
AA26	EXP_B_TXN[7]
W25	EXP_B_TXN[8]
W29	EXP_B_TXN[9]
AA33	EXP_B_TXN[10]
W31	EXP_B_TXN[11]
U30	EXP_B_TXN[12]
V29	EXP_B_TXN[13]
T28	EXP_B_TXN[14]
T25	EXP_B_TXN[15]

Table 7-3. MCH Signal Listing (Alphabetical by Signal Name) (Sheet 4 of 8)

Ball #	Signal Name
AG32	EXP_B_TXP[0]
AF31	EXP_B_TXP[1]
AC33	EXP_B_TXP[2]
AB32	EXP_B_TXP[3]
AD27	EXP_B_TXP[4]
AC27	EXP_B_TXP[5]
AB29	EXP_B_TXP[6]
AA27	EXP_B_TXP[7]
W26	EXP_B_TXP[8]
W28	EXP_B_TXP[9]
Y33	EXP_B_TXP[10]
W32	EXP_B_TXP[11]
U31	EXP_B_TXP[12]
V30	EXP_B_TXP[13]
T29	EXP_B_TXP[14]
T26	EXP_B_TXP[15]
R24	EXP_CLKN
T23	EXP_CLKP
U33	EXP_COMP[0]
U25	EXP_COMP[1]
U27	EXP_VCCBG
U28	EXP_VSSBG
K22	HA[3]#
J20	HA[4]#
G23	HA[5]#
G22	HA[6]#
H21	HA[7]#
K19	HA[8]#
H19	HA[9]#
G19	HA[10]#
E22	HA[11]#
E21	HA[12]#
F18	HA[13]#
E19	HA[14]#
F21	HA[15]#
F20	HA[16]#
D26	HA[17]#
C26	HA[18]#
A26	HA[19]#
D22	HA[20]#
B22	HA[21]#
A25	HA[22]#
B25	HA[23]#
D25	HA[24]#
C24	HA[25]#
A22	HA[26]#

Ball #	Signal Name
B21	HA[27]#
D23	HA[28]#
A23	HA[29]#
B24	HA[30]#
A20	HA[31]#
D19	HA[32]#
C20	HA[33]#
C21	HA[34]#
D20	HA[35]#
F23	HACVREF
G20	HADSTB[0]#
C23	HADSTB[1]#
J11	HCLKINN
K11	HCLKINP
C27	HCRES0
C18	HD[0]#
B19	HD[1]#
C14	HD[2]#
A17	HD[3]#
A19	HD[4]#
B16	HD[5]#
C17	HD[6]#
B18	HD[7]#
D17	HD[8]#
A16	HD[9]#
B13	HD[10]#
A14	HD[11]#
A13	HD[12]#
D14	HD[13]#
C12	HD[14]#
B12	HD[15]#
E18	HD[16]#
J18	HD[17]#
H18	HD[18]#
F17	HD[19]#
G17	HD[20]#
K17	HD[21]#
E16	HD[22]#
J17	HD[23]#
J14	HD[24]#
F14	HD[25]#
F15	HD[26]#
G16	HD[27]#
K16	HD[28]#
H16	HD[29]#
G14	HD[30]#

Ball #	Signal Name
K14	HD[31]#
E12	HD[32]#
C11	HD[33]#
H13	HD[34]#
F11	HD[35]#
G13	HD[36]#
D11	HD[37]#
E9	HD[38]#
F12	HD[39]#
G10	HD[40]#
D8	HD[41]#
H10	HD[42]#
F8	HD[43]#
J12	HD[44]#
G11	HD[45]#
K13	HD[46]#
H12	HD[47]#
B10	HD[48]#
A10	HD[49]#
A11	HD[50]#
C9	HD[51]#
B9	HD[52]#
C8	HD[53]#
B6	HD[54]#
B7	HD[55]#
E7	HD[56]#
B4	HD[57]#
A4	HD[58]#
B3	HD[59]#
D5	HD[60]#
C6	HD[61]#
D7	HD[62]#
C5	HD[63]#
B15	HDSTBN[0]#
H15	HDSTBN[1]#
D10	HDSTBN[2]#
A8	HDSTBN[3]#
C15	HDSTBP[0]#
J15	HDSTBP[1]#
E10	HDSTBP[2]#
A7	HDSTBP[3]#
D13	HDVREF[0]
E13	HDVREF[1]
J30	HI[0]
H30	HI[1]
C32	HI[2]

Table 7-3. MCH Signal Listing (Alphabetical by Signal Name) (Sheet 5 of 8)

Ball #	Signal Name
G31	HI[3]
G29	HI[4]
H28	HI[5]
K26	HI[6]
J27	HI[7]
F30	HI[8]
E33	HI[9]
J29	HI[10]
G32	HI[11]
D32	HI_STBF
E31	HI_STBS
H31	HI_VSWING
L24	HICLK
K25	HICOMP
E30	HIT#
D28	HITM#
F32	HIVREF
C30	HLOCK#
E27	HODTCRES
K20	HREQ[0]#
J21	HREQ[1]#
J23	HREQ[2]#
H22	HREQ[3]#
K23	HREQ[4]#
F26	HSLWCRES
A30	HTRDY#
H24	MCERR#
L25	MCHGPE#
M24	MCHPME#
C31	PLLSEL[0]#
A29	PLLSEL[1]#
E3	PWRGD
D33	RESERVED
F33	RESERVED
E6	RESERVED
L33	RESERVED
M8	RESERVED
R8	RESERVED
R9	RESERVED
R10	RESERVED
R32	RESERVED
AA24	RESERVED
AD20	RESERVED
AE23	RESERVED
AF30	RESERVED
AJ19	RESERVED
F29	RS[0]#
D31	RS[1]#
G28	RS[2]#
J26	RSP#
C2	RSTIN#
C3	SMBSCL
D4	SMBSDA
D2	TCK
G5	TDI
G6	TDO
L12	TEST#
F3	TMS
J9	TRST#
H33	V3REF
C33	VCC
G33	VCC
H29	VCC
K9	VCC
K27	VCC
L23	VCC
M12	VCC
M22	VCC
N13	VCC
N15	VCC
N17	VCC
N19	VCC
N21	VCC
N23	VCC
P14	VCC
P16	VCC
P18	VCC
P22	VCC
R13	VCC
R15	VCC
R17	VCC
R19	VCC
R21	VCC
T14	VCC
T16	VCC
T18	VCC
T20	VCC
U13	VCC
U15	VCC
U17	VCC
U19	VCC
U21	VCC
V14	VCC
V16	VCC
V18	VCC
V20	VCC
W13	VCC
W15	VCC
W17	VCC
W19	VCC
W21	VCC
Y14	VCC
Y16	VCC
Y18	VCC
Y20	VCC
AA13	VCC
AA15	VCC
AA17	VCC
AA19	VCC
AA21	VCC
F6	VCCA_CORE
E4	VCCA_DDR
U23	VCCA_EXP
P20	VCCA_HI
F1	VCCDDR
H5	VCCDDR
K3	VCCDDR
M1	VCCDDR
M7	VCCDDR
N11	VCCDDR
P5	VCCDDR
P12	VCCDDR
R11	VCCDDR
T3	VCCDDR
T9	VCCDDR
T12	VCCDDR
U11	VCCDDR
V1	VCCDDR
V7	VCCDDR
V12	VCCDDR
W11	VCCDDR
Y5	VCCDDR
Y12	VCCDDR
AA11	VCCDDR
AB3	VCCDDR
AB9	VCCDDR
AB12	VCCDDR
AB14	VCCDDR

Table 7-3. MCH Signal Listing (Alphabetical by Signal Name) (Sheet 6 of 8)

Ball #	Signal Name
AB16	VCCDDR
AB18	VCCDDR
AB20	VCCDDR
AB22	VCCDDR
AC11	VCCDDR
AC13	VCCDDR
AC15	VCCDDR
AC17	VCCDDR
AC19	VCCDDR
AC21	VCCDDR
AC23	VCCDDR
AD1	VCCDDR
AD7	VCCDDR
AE12	VCCDDR
AE18	VCCDDR
AE24	VCCDDR
AF5	VCCDDR
AG10	VCCDDR
AG16	VCCDDR
AG22	VCCDDR
AG28	VCCDDR
AH3	VCCDDR
AJ8	VCCDDR
AJ14	VCCDDR
AJ20	VCCDDR
AJ26	VCCDDR
AL1	VCCDDR
AL6	VCCDDR
AL12	VCCDDR
AL18	VCCDDR
AL24	VCCDDR
AL30	VCCDDR
AN3	VCCDDR
AN10	VCCDDR
AN16	VCCDDR
AN22	VCCDDR
AN28	VCCDDR
K30	VCCEXP
K33	VCCEXP
M28	VCCEXP
N33	VCCEXP
P29	VCCEXP
R23	VCCEXP
T22	VCCEXP
T24	VCCEXP
T27	VCCEXP
T33	VCCEXP
U29	VCCEXP
V22	VCCEXP
W23	VCCEXP
W27	VCCEXP
W33	VCCEXP
Y22	VCCEXP
Y29	VCCEXP
AA23	VCCEXP
AB24	VCCEXP
AB27	VCCEXP
AB33	VCCEXP
AC29	VCCEXP
AE30	VCCEXP
AE33	VCCEXP
AH33	VCCEXP
A21	VSS
A24	VSS
A27	VSS
B5	VSS
B8	VSS
B11	VSS
B14	VSS
B17	VSS
B20	VSS
B23	VSS
B26	VSS
B29	VSS
B32	VSS
C1	VSS
C4	VSS
C7	VSS
C10	VSS
C13	VSS
C16	VSS
C19	VSS
C22	VSS
C25	VSS
C28	VSS
D6	VSS
D9	VSS
D12	VSS
D15	VSS
D18	VSS
D21	VSS
D24	VSS
E2	VSS
E5	VSS
E24	VSS
E26	VSS
E29	VSS
E32	VSS
F4	VSS
F7	VSS
F10	VSS
F13	VSS
F16	VSS
F19	VSS
F22	VSS
F25	VSS
F28	VSS
F31	VSS
G3	VSS
G9	VSS
G12	VSS
G15	VSS
G18	VSS
G21	VSS
G24	VSS
G27	VSS
G30	VSS
H2	VSS
H8	VSS
H11	VSS
H14	VSS
H17	VSS
H20	VSS
H32	VSS
J1	VSS
J4	VSS
J7	VSS
J10	VSS
J22	VSS
J25	VSS
J28	VSS
J31	VSS
K6	VSS
K12	VSS
K15	VSS
K18	VSS
K21	VSS
K24	VSS

Table 7-3. MCH Signal Listing (Alphabetical by Signal Name) (Sheet 7 of 8)

Ball #	Signal Name
L2	VSS
L5	VSS
L8	VSS
L14	VSS
L16	VSS
L18	VSS
L20	VSS
L22	VSS
L26	VSS
L29	VSS
L32	VSS
M4	VSS
M10	VSS
M11	VSS
M13	VSS
M15	VSS
M17	VSS
M19	VSS
M21	VSS
M23	VSS
M25	VSS
M31	VSS
N3	VSS
N6	VSS
N9	VSS
N12	VSS
N14	VSS
N16	VSS
N18	VSS
N20	VSS
N22	VSS
N24	VSS
N27	VSS
N30	VSS
P2	VSS
P8	VSS
P11	VSS
P13	VSS
P15	VSS
P17	VSS
P19	VSS
P23	VSS
P26	VSS
P32	VSS
R1	VSS
R4	VSS
R7	VSS
R12	VSS
R14	VSS
R16	VSS
R18	VSS
R20	VSS
R22	VSS
R25	VSS
R28	VSS
R31	VSS
T6	VSS
T11	VSS
T13	VSS
T15	VSS
T17	VSS
T19	VSS
T21	VSS
T30	VSS
U2	VSS
U5	VSS
U8	VSS
U12	VSS
U14	VSS
U16	VSS
U18	VSS
U20	VSS
U22	VSS
U26	VSS
U32	VSS
V4	VSS
V10	VSS
V11	VSS
V13	VSS
V15	VSS
V17	VSS
V19	VSS
V21	VSS
V25	VSS
V28	VSS
V31	VSS
W3	VSS
W6	VSS
W9	VSS
W12	VSS
W14	VSS
W16	VSS
W18	VSS
W20	VSS
W22	VSS
W24	VSS
W30	VSS
Y2	VSS
Y8	VSS
Y11	VSS
Y13	VSS
Y15	VSS
Y17	VSS
Y19	VSS
Y21	VSS
Y23	VSS
Y26	VSS
Y32	VSS
AA1	VSS
AA4	VSS
AA7	VSS
AA10	VSS
AA12	VSS
AA14	VSS
AA16	VSS
AA18	VSS
AA20	VSS
AA22	VSS
AA25	VSS
AA28	VSS
AA31	VSS
AB6	VSS
AB11	VSS
AB13	VSS
AB15	VSS
AB17	VSS
AB19	VSS
AB21	VSS
AB23	VSS
AB30	VSS
AC2	VSS
AC5	VSS
AC8	VSS
AC12	VSS
AC14	VSS
AC16	VSS
AC18	VSS
AC20	VSS

Table 7-3. MCH Signal Listing (Alphabetical by Signal Name) (Sheet 8 of 8)

Ball #	Signal Name
AC22	VSS
AC26	VSS
AC32	VSS
AD4	VSS
AD10	VSS
AD13	VSS
AD16	VSS
AD19	VSS
AD22	VSS
AD25	VSS
AD28	VSS
AD31	VSS
AE3	VSS
AE6	VSS
AE9	VSS
AE15	VSS
AE21	VSS
AE27	VSS
AF2	VSS
AF8	VSS
AF11	VSS
AF14	VSS
AF17	VSS
AF20	VSS
AF23	VSS
AF26	VSS
AF29	VSS
AF32	VSS
AG1	VSS
AG4	VSS
AG7	VSS
AG13	VSS
AG19	VSS
AG25	VSS
AG31	VSS
AH9	VSS
AH12	VSS
AH15	VSS
AH18	VSS
AH21	VSS
AH24	VSS
AH27	VSS
AH30	VSS
AJ2	VSS
AJ5	VSS
AJ11	VSS
AJ17	VSS
AJ23	VSS
AJ29	VSS
AJ32	VSS
AK4	VSS
AK7	VSS
AK10	VSS
AK13	VSS
AK16	VSS
AK19	VSS
AK22	VSS
AK25	VSS
AK28	VSS
AK31	VSS
AL3	VSS
AL9	VSS
AL15	VSS
AL21	VSS
AL27	VSS
AL33	VSS
AM2	VSS
AM5	VSS
AM8	VSS
AM11	VSS
AM14	VSS
AM17	VSS
AM20	VSS
AM23	VSS
AM26	VSS
AM29	VSS
AM32	VSS
AN7	VSS
AN13	VSS
AN19	VSS
AN25	VSS
AN31	VSS
F5	VSSA_CORE
V23	VSSA_EXP
P21	VSSA_HI
A3	VTT
A6	VTT
A9	VTT
A12	VTT
A15	VTT
A18	VTT
A31	VTT
D27	VTT
D30	VTT
E8	VTT
E11	VTT
E14	VTT
E17	VTT
E20	VTT
E23	VTT
H23	VTT
H26	VTT
J13	VTT
J16	VTT
J19	VTT
L13	VTT
L15	VTT
L17	VTT
L19	VTT
L21	VTT
M14	VTT
M16	VTT
M18	VTT
M20	VTT

7.2 Package Specifications

Figure 7-2. MCH Package Dimensions (Bottom View)

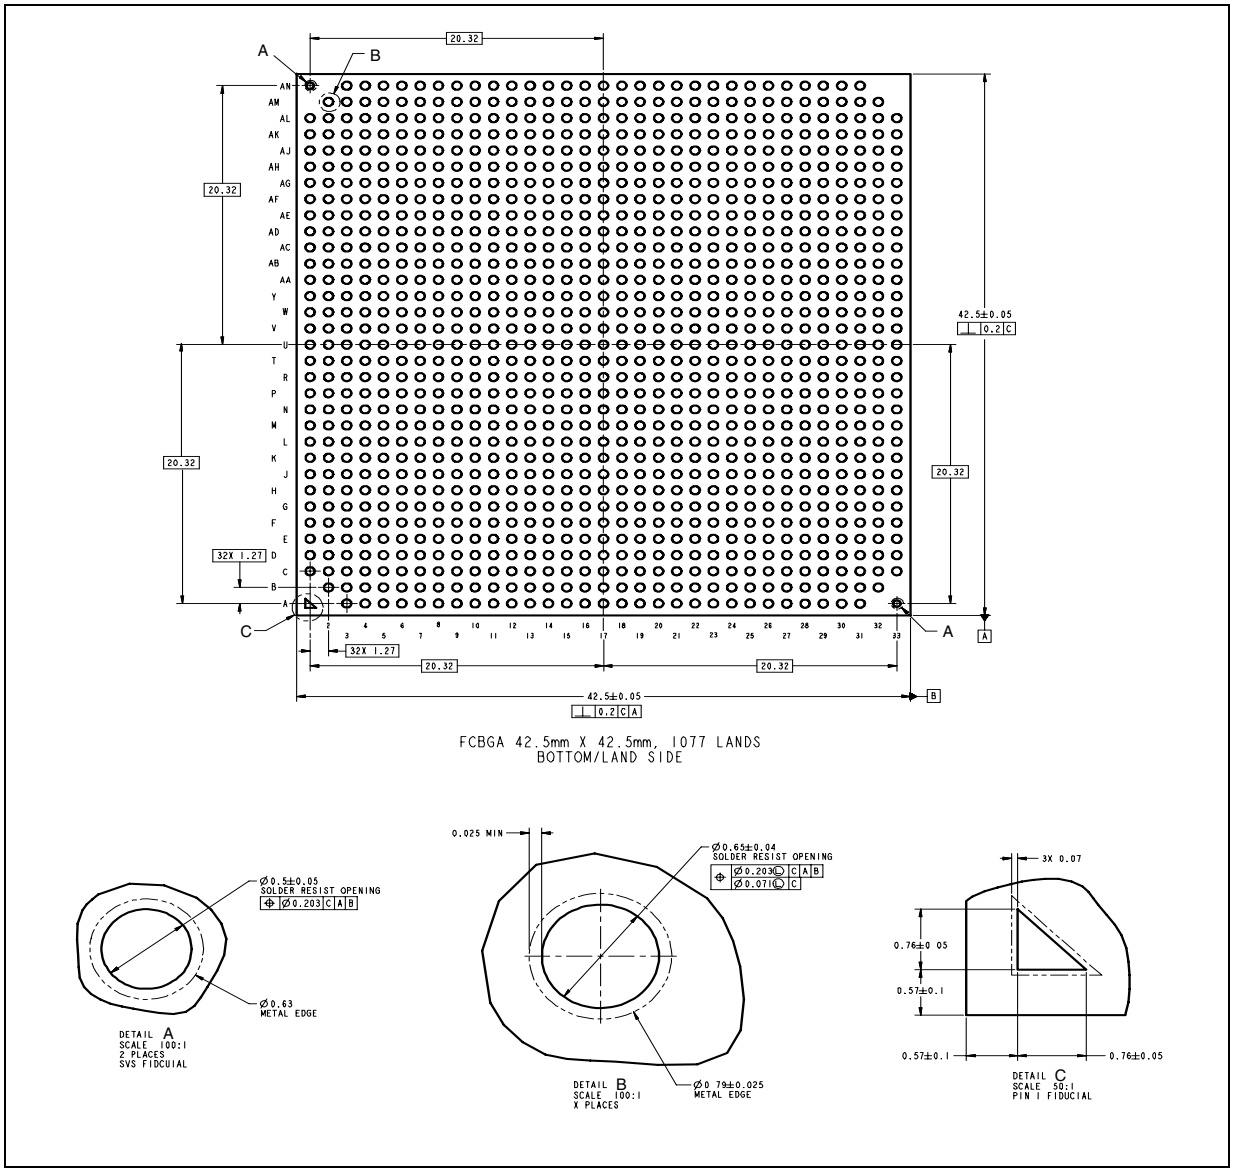

Figure 7-3. MCH Package Dimensions (Side View)

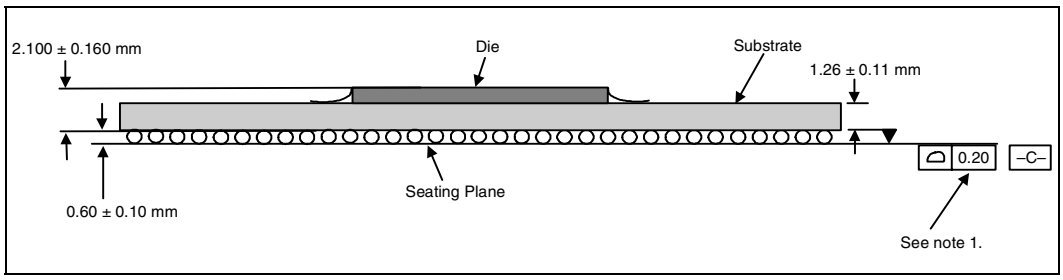

NOTES:
1. Primary datum -C- and seating plane are defined by the spherical crowns of the solder balls.
2. All dimensions and tolerances conform to ANSI Y14.5M-1982.

7.3 Chipset Interface Trace Length Compensation

In this section, detailed information is given about the internal component package trace lengths to enable trace length compensation. Trace length compensation is required for platform design. These lengths must be considered when matching trace lengths. Note that these lengths represent the actual lengths from pad to ball.

7.3.1 System Bus Signal Package Trace Length Data

Table 7-4 provides the MCH package trace length information for the system bus.

Table 7-4. MCH L_{PKG} Data for the System Bus (Sheet 1 of 2)

Signal	Ball No.	L_{PKG} (mils)
HADSTB0#	G20	434.83
HA3#	K22	437.54
HA4#	J20	340.49
HA5#	G23	523.93
HA6#	G22	508.86
HA7#	H21	448.47
HA8#	K19	282.73
HA9#	H19	399.55
HA10#	G19	441.88
HA11#	E22	657.91
HA12#	E21	631.09
HA13#	F18	490.21
HA14#	E19	585.54
HA15#	F21	585.62
HA16#	F20	562.86
HREQ0#	K20	316.39
HREQ1#	J21	470.22
HREQ2#	J23	659.76
HREQ3#	H22	442.37

Signal	Ball No.	L_{PKG} (mils)
HDSTBN0#	B15	708.84
HDSTBP0#	C15	685.07
HD0#	C18	622.20
HD1#	B19	723.28
HD2#	C14	640.91
HD3#	A17	779.33
HD4#	A19	756.98
HD5#	B16	681.35
HD6#	C17	617.5
HD7#	B18	713.25
HD8#	D17	582.94
HD9#	A16	746.55
HD10#	B13	733.81
HD11#	A14	739.39
HD12#	A13	813.60
HD13#	D14	579.54
HD14#	C12	635.65
HD15#	B12	710.56
DBI0#	D16	593.83

Table 7-4. MCH L_{PKG} Data for the System Bus (Sheet 2 of 2)

Signal	Ball No.	L_{PKG} (mils)
HREQ4#	K23	557.13
HADSTB1#	C23	749.76
HA17#	D26	740.04
HA18#	C26	779.70
HA19#	A26	895.88
HA20#	D22	678.27
HA21#	B22	802.17
HA22#	A25	852.39
HA23#	B25	799.72
HA24#	D25	685.58
HA25#	C24	726.14
HA26#	A22	825.57
HA27#	B21	771.24
HA28#	D23	651.69
HA29#	A23	849.39
HA30#	B24	797.98
HA31#	A20	764.35
HA32#	D19	617.06
HA33#	C20	697.36
HA34#	C21	707.58
HA35#	D20	623.62
HCLKINN	J11	401.19
HCLKINP	K11	401.24
HDSTBN2#	D10	750.98
HDSTBP2#	E10	726.39
HD32#	E12	549.94
HD33#	C11	677.29
HD34#	H13	401.08
HD35#	F11	517.91
HD36#	G13	415.18
HD37#	D11	625.95
HD38#	E9	690.71
HD39#	F12	519.70
HD40#	G10	448.02
HD41#	D8	715.80
HD42#	H10	422.33
HD43#	F8	608.17
HD44#	J12	387.05
HD45#	G11	461.09
HD46#	K13	394.56
HD47#	H12	391.33
DBI2#	F9	616.50

Signal	Ball No.	L_{PKG} (mils)
HDSTBN1#	H15	381.21
HDSTBP1#	J15	360.19
HD16#	E18	534.21
HD17#	J18	380.80
HD18#	H18	482.26
HD19#	F17	476.51
HD20#	G17	429.13
HD21#	K17	433.63
HD22#	E16	532.54
HD23#	J17	322.78
HD24#	J14	306.69
HD25#	F14	481.59
HD26#	F15	486.34
HD27#	G16	435.47
HD28#	K16	364.76
HD29#	H16	357.74
HD30#	G14	436.72
HD31#	K14	273.96
DBI1#	E15	626.10
HDSTBN3#	A8	988.22
HDSTBP3#	A7	968.76
HD48#	B10	922.17
HD49#	A10	839.04
HD50#	A11	785.72
HD51#	C9	771.01
HD52#	B9	793.75
HD53#	C8	749.34
HD54#	B6	837.17
HD55#	B7	813.55
HD56#	E7	673.39
HD57#	B4	891.84
HD58#	A4	955.63
HD59#	B3	885.23
HD60#	D5	761.97
HD61#	C6	811.14
HD62#	D7	796.21
HD63#	C5	826.43
DBI3#	A5	1017.37

7.3.2 Other MCH Interface Signal Package Trace Length Data

The MCH does not require package trace length compensation on any other interfaces (DDR, Hub Interface, and PCI Express*). As such, the signal package trace lengths are not disclosed here. Because package trace lengths are required for simulation, they are all documented in a spreadsheet packaged with the signal integrity models.