Intel 450KX/GX PCIset

82454KX/GX PCI Bridge (PB)

82453KX/GX DRAM Controller (DC)

82452KX/GX Data Path (DP)

82451KX/GX Memory Interface Component (MIC)

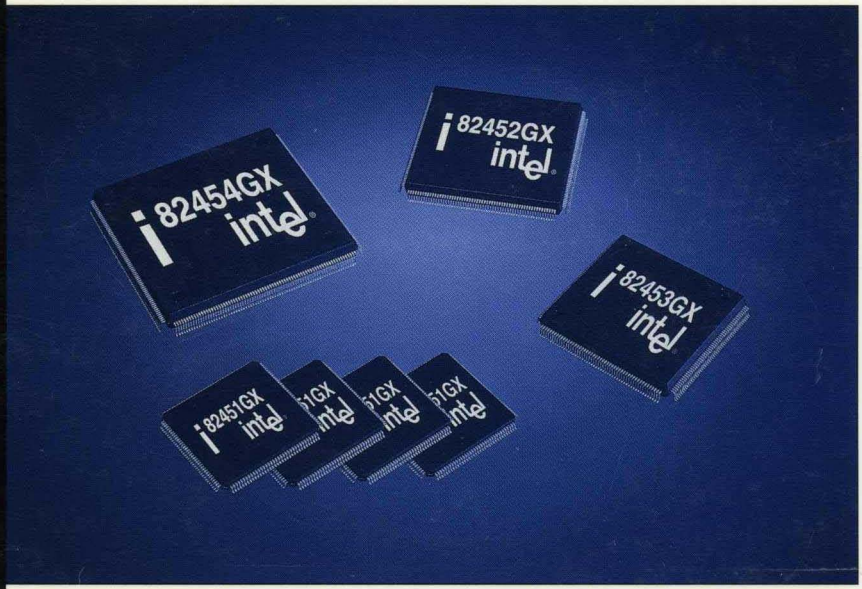

PCI LOCAL BUS

intel®

intel

Intel 450KX/GX PCIset

The information in this document is subject to change.

PRELIMINARY

Copies of documents which have an ordering number and are referenced in this document, or other Intel literature, may be obtained from:

Intel Corporation
P.O. Box 7641
Mt. Prospect, IL 60056-7641

or call 1-800-879-4683

intel

Contents

Chapter 1 Intel 450KX/GX PCIset Overview

1.0 Intel 450KX PCIset 5
2.0 Intel 450GX PCIset 6
3.0 Host Bus Efficiency 6
4.0 System Memory Map 7
4.1 Compatibility Area 8
4.2 Extended Memory (ISA) 9
4.3 Extended Memory (EISA) 10
4.4 Extended Memory (above 4 Gbytes) 12
4.5 System Management Mode (SMM) 12
5.0 I/O Space (PB Only) 12
6.0 Memory Mapped I/O 13

Chapter 2 82454KX/GX PCI Bridge (PB)

1.0 PB Signal Descriptions 19
1.1 PB Signals 19
1.2 Signal State During Reset 25
2.0 PB Register Description 26
2.1 Initialization and Configuration 26
2.2 I/O Space Registers 27
2.2.1 CONFADD—Configuration Address Register 28
2.2.2 TRC—Turbo and Reset ControL 29
2.2.3 CONFDATA—Configuration Data Register 30
2.3 PCI Configuration Space 30
2.4 PB PCI Configuration Registers 32
2.4.1 VID—Vendor Identification Register 34
2.4.2 DID—Device Identification Register 34
2.4.3 PCICMD—PCI Command Register 34
2.4.4 PCISTS—PCI Status Register 35
2.4.5 RID—Revision Identification Register 36
2.4.6 CLASSC—Class Code Register 36
2.4.7 CLSIZE—Cache Line Size Register 36
2.4.8 PLTMR—PCI Latency Timer 37
2.4.9 HEADT—Header Type Register 37
2.4.10 BIST—Bist register 37
2.4.11 TSM—Top of System Memory Register 38
2.4.12 PDM—PCI Decode Mode 38

2.4.13 BDNUM—Bridge Device Number Register 39
2.4.14 PBNUM—PCI Bus Number Register 39
2.4.15 PSBNUM—Subordinate Bus Number Register 40
2.4.16 PBC—PB Configuration Register 40
2.4.17 DCC—Deturbo Counter Register 41
2.4.18 CRWC—CPU Read/Write Control Register 41
2.4.19 PRWC—PCI Read/Write Control 42
2.4.20 SMME—SMRAM Enable Register 43
2.4.21 VBAE—Video Buffer Area Enable Register 43
2.4.22 PAM[0:6]—Programmable Attribute MAp Register 44
2.4.23 ERRCMD—Error Reporting Command Register 45
2.4.24 ERRSTS—Error Reporting Status Register 45
2.4.25 MGR—Memory Gap Range Register 46
2.4.26 MGUA—Memory Gap Upper Address Register 46
2.4.27 PFB—PCI Frame Buffer Register 47
2.4.28 HMGSA—High Memory Gap Range Start Address Register 48
2.4.29 HMGEA—High Memory Gap End Address Register 48
2.4.30 IOSR1—I/O Space Range 1 Register (82454GX Only) 49
2.4.31 PCIRSR—PCI Reset Register 49
2.4.32 IOSR2—I/O Space Range 2 Register (82454GX Only) 50
2.4.33 APICR—I/O APIC Range Register 50
2.4.34 CONFVR—Configuration Values Driven on Reset Register 51
2.4.35 CSCONFV—Captured System Configuration Values Register 52
2.4.36 SMMR—SMRAM Range Register 53
2.4.37 HBIOSR—High BIOS Range Register 53
2.4.38 EXERRCMD—PB Extended Error Reporting Command Register 53
2.4.39 EXERRSTS—PB Extended Error Reporting Status 55
2.4.40 PBRTMR—PB Retry Timers 56
3.0 PB Functional Description 57
3.1 Memory and I/O Map 57
3.1.1 Memory Address Map 57
3.1.2 I/O Address Map 59
3.2 Host Bus Interface 60
3.3 PCI Bus Interface 61
3.4 Data Integrity and Error Handling 62
3.4.1 Host Bus Errors 62
3.4.2 PCI Bus Errors 62
3.4.2.1 PB Master Operation on PCI 63
3.4.2.2 PB Target Operation on PCI 63
3.5 Dual PB Architectures (82454GX Only) 65

3.6 Peripheral Operation and Performance 68
3.6.1 Matching Peripherals to the 450KX/GX 68
3.6.2 Distributing Peripherals Within the I/O Subsystem 69
3.6.3 PCI-to-PCI Bridges 69
3.6.4 BIOS Performance Tuning 69
3.7 Clock, Reset, and Configuration 70
3.7.1 System clocking 70
3.7.1.1 Host Bus Clock 70
3.7.1.2 PCI Clock 71
3.7.2 System Reset 71
3.7.3 System Initialization 72
3.7.4 Dual PB Configuration (82454GX only) 72
3.7.5 Using the 82379AB SIO.A PCI-to-ISA Bridge with the 450KX/GX 73
3.8 Host to PCI Bus Command Translation 76
3.9 PCI to Host Bus Command Translation 77
4.0 PB Pinout and Package Information 79
4.1 Pin Assignment 79
4.2 Package Information 87

Chapter 3 Memory Controller (MC)

1.0 MC Signal Description 93
1.1 DC Signals 93
1.2 DP Signals 96
1.3 MIC Signals 98
1.4 Signal State During Reset 100
2.0 MC Register Description 101
2.1 Initialization and Configuration 101
2.2 I/O Space Registers 102
2.2.1 CONFADD—Configuration Address Register 103
2.2.2 CONFDATA—Configuration Data Register 103
2.3 MC Configuration Registers 104
2.3.1 VID—Vendor Identification Register 105
2.3.2 DID—Device Identification Register 105
2.3.3 PCICMD—PCI Command Register 106
2.3.4 PCISTS—PCI Status Register 106
2.3.5 RID—Revision Identification Register 106
2.3.6 CLASSC—Class Code Register 107
2.3.7 BASEADD—MC Base Address Register (450GX only) 107
2.3.8 CDNUM—Controller Device Number Register 108
2.3.9 CMD—Command Register 108

2.3.10 SMME—SMRAM Enable Register 110
2.3.11 VBRE—Video Buffer Region Enable Register 110
2.3.12 PAM[0:6]—Programmable Attribute Map Registers 110
2.3.13 DRL—DRAM Row Limit (0 to 7) 111
2.3.14 SBCERRADD—Single Bit Correctable Error Address Register 112
2.3.15 MG—Memory Gap Register 113
2.3.16 MGUA—Memory Gap Upper Address Register 114
2.3.17 LMG—Low Memory Gap Register 114
2.3.18 HMGSA—High Memory Gap Start Address Register 115
2.3.19 HMGEA—High Memory Gap End Address Register 115
2.3.20 APICR—I/O APIC Range Register 115
2.3.21 UERRADD—Uncorrectable Error Address Register 116
2.3.22 MEMTIM—Memory Timing Register 116
2.3.23 SMMR—SMRAM Range Register 119
2.3.24 HBIOSR—High BIOS Gap Range Register 119
2.3.25 MERRCMD—Memory Error Reporting Command 120
2.3.26 MERRSTS—Memory Error Status Register 120
2.3.27 SERRCMD—System Error Reporting Command Register 121
2.3.28 SERRSTS—System Error Status Register 122
2.4 Memory Configuration Determination Algorithm 122
3.0 MC Functional Description 123
3.1 Memory and I/O Map 123
3.2 Host Bus Interface 124
3.3 DRAM Interface 125
3.3.1 DRAM Configurations 127
3.3.1.1 Memory Interface Component (MIC) 127
3.3.1.2 4-Way DRAM Configuration (450GX Only) 128
3.3.1.3 2-Way DRAM configuration 129
3.3.1.4 Non-Interleaved DRAM configuration 130
3.4 Clocks and Reset 131
3.4.1 Clocks 131
3.4.2 Reset 131
4.0 MC Pinout and Package Information 132
4.1 82453KX/82453GX (DC) Pin Assignment 132
4.2 82452GX/82452KX (DP) Pin Assignment 135
4.3 82451GX/82451KX (MIC) Pin Assignment 142
4.4 82453GX/82453KX (DC) Package Dimensions 145
4.5 82452GX/82452KX (DP) Package Dimensions 146
4.6 82451GX/82451KX (MIC) Package Dimensions 148

intel

Chapter 4 PCIset Electrical Specifications

1.0 Electrical Characteristics 151
1.1 Test Pins and Unused Pins 151
1.2 Signal Groups 151
1.2.1 The Power Good Signal— PWRGD 152
1.3 Maximum Ratings 153
1.4 DC Specifications 154
1.5 AC Specifications 157
1.5.1 Waveforms 160
2.0 Signal Quality 162
2.1 I/O Signal Simulations—Ensuring I/O Timings 162
2.2 Signal Quality Specifications 163
2.3 PCIset Ringback Specification 163
2.4 450KX/GX Undershoot Specification 165
3.0 Thermal Specifications 166

Index

Alphabetized Index 169

Chapter 1

Intel 450KX/GX PCIset Overview

PCIset Product Overview

- PCIset Host Bus Support
  - Supports Pentium® Pro Processor at 60 MHz, and 66 MHz Bus Speeds
  - 64-Bit Data and 36-Bit Address Bus
  - Parity Protection on Control Signals
  - ECC Protection on Host Data Bus (450GX)
  - Dual-Processor Support (450KX)
  - Quad-Processor Support (450GX)
  - Up to Eight Deep In-Order Queue
  - Four Deep Outbound Request Queue
  - Four Cache Line Read and Write Buffers
  - GTL+ Bus Driver Technology
- Host-to-PCI Bridge (PB)
  - Combines Both the Control and Data Path in a Single Chip
  - Internal Bridge Arbiter For Two PBs in a system (450GX)
  - Synchronous PCI Interface
  - 32-bit Address/Data PCI Bus (64-bit Dual Cycle Address Support)
  - Parity Protection on All PCI Bus Signals
  - Four Deep Inbound Request Queue
  - Data Collection/Write Assembly of Line Bursts.
  - Support for 3.3V & 5V PCI Devices
  - Available in 304 Pin QFP or 352 pin BGA
- Memory Controller (MC)
  - 1 GB Maximum Memory (450KX)
  - 4 GBs Maximum Main Memory (per 82453GX)
  - 2-Way interleaved and Non-Interleaved Memory Organizations
  - 4-Way and 2-Way interleaved, and Non-Interleaved Memory Organizations (450GX)
  - Up to Two MCs in a System (450GX)
  - Supports 3.3V and 5V SIMMs
  - Supports Standard 32- or 36-bit SIMMs or 72-bit DIMMs
  - Supports 4 Mbit, 16 Mbit, and 64 Mbit DRAM Technology
  - Single Bit Error Correction, Double Bit and Nibble Error Detection
  - Memory Array Power Management
  - Recovers DRAM Memory Behind Programmable Memory Gaps
  - Read Page Hit 8-1-1-1 (at 66 MHz, 60 ns DRAM)
  - Read Page Miss 11-1-1-1 (66 MHz, 60 ns DRAM)
  - Read Page Miss + Precharge 14-1-1-1 (66 MHz, 60 ns DRAM)
  - Available in 208-Pin QFP for the DC; 240-Pin QFP or 256-Pin BGA for the DP; 144-Pin QFP for the MIC
- On-Chip Digital PLL (Both PB and MC)
- Test Support (JTAG) (Both PB and MC)

The Intel 450KX/GX PCIsets provide a high-performance system solution for Pentium® Pro processor-based PCI systems by combining high integration, high performance technology with a scalable architecture that is capable of high throughput for up to four Pentium Pro processors. Scalability provides a wide range of system solutions from cost-effective uniprocessor systems to high-end multiprocessor systems without sacrificing performance. For systems requiring extensive I/O (e.g., file servers), a second PB can be easily added providing two high-performance PCI bus structures. The flexibility of the memory controller permits easy expansion from a simple non-interleaved organization to a 2-way or 4-way interleaved organization to increase performance. Extended error checking and logging, ECC, and the ability to build in redundancy (e.g, multiple processors and dual PCI bridges) provides a comprehensive solution for systems requiring high reliability.

The PCIset may contain design defects or errors known as errata. Current characterized errata are available upon request.

This document describes both the Intel 450KX and 450GX PCIsets. Unshaded areas apply to both the PCIsets. Shaded areas, like this one, describe the 450GX operations that differ from the 450KX.

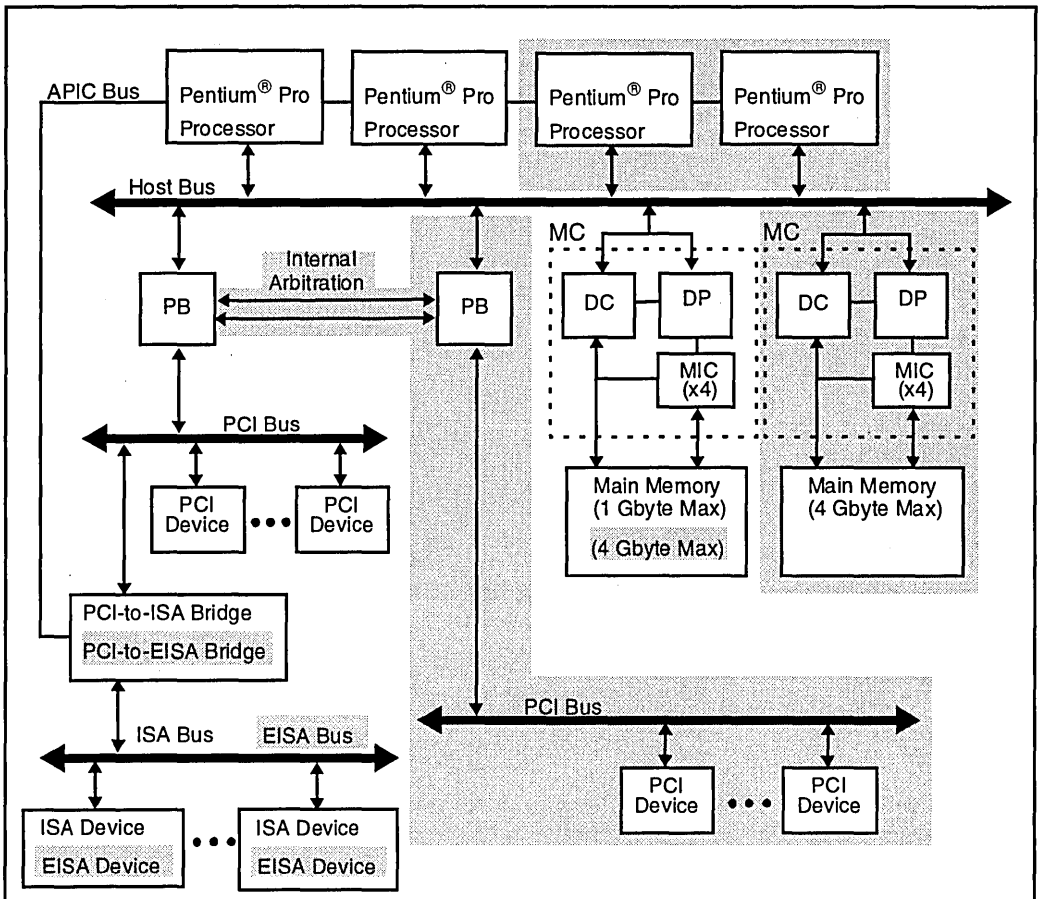

Note: 450KX: Supports one PB, one MC, and up to two processors.
450GX: Supports up to two PBs, two MCs, and four processors on the host bus.

Figure 1. 450KX/GX Simplified System Block Diagram

1.0 INTEL 450KX PCISET

The 450KX desktop PCIset consists of the 82454KX PCI Bridge (PB) and the Memory Controller (MC). The MC consists of the 82453KX DRAM Controller (DC), the 82452KX Data Path (DP), and four 82451KX Memory Interface Components (MIC). The system configuration using the Intel 450KX PCIset supports one PB, one MC and up to two Pentium Pro processors (Figure 1). An ISA subsystem is also located below the PB. For Pentium Pro processor bus error detection, the 450KX generates and checks parity over the address and request/response signal lines. This feature can be enabled/disabled during system configuration.

KX PCI Bridge (PB)

The PB is a single-chip host-to-PCI Bridge. A rich set of CPU-to-PCI and PCI-to-CPU bus transaction translations optimize bus bandwidth and improve system performance. All ISA and EISA regions are supported. Three programmable memory gaps can be created—a PCI Frame Buffer Region with specialized frame buffer attributes and two general-purpose memory gaps (called the Memory Gap Region and the High Memory Gap Region).

The PB takes advantage of the Pentium Pro processor ratio clocking scheme to assure modularity now and upgradability in the future. The PB has a synchronous interface to the Pentium Pro processor bus and supports a derived clock for the synchronous PCI interface. The PB derives either a 30 or 33 MHz PCI clock output from the Pentium Pro processor bus clock. The PB PCI signals are 5 volt tolerant and can be used with either 5 volt or 3.3 volt PCI devices.

KX Memory Controller (MC)

The combined MC (DC, DP, and four MICs) act as one physical load on the Pentium Pro processor bus. The DC provides control for the DRAM memory subsystem, the DP provides the data path, and the four MICs are used to interface the MC datapath with the DRAM memory subsystem.

The memory configuration can be either 2-way interleaved or non-interleaved. Both single-sided and double-sided SIMMs are supported. DRAM technologies up to 64 Mbits at speeds of 50ns, 60ns, and 70ns can be used. Asymmetric DRAM is supported up to two bits of asymmetry (e.g., 12 row address lines and 10 column address lines). The maximum memory size is 1 Gbyte for the 2-way interleaved configuration and 512 Mbytes for the non-interleaved configuration using 16 Mbit technology. In addition to these memory configurations, the MC provides data integrity features including ECC in the memory array. These features, as well as a set of error reporting mechanisms, can be selected via configuration of the MC. Each interleave provides a 64-bit data path to main memory (72-bits including ECC).

The MC is PC compatible. All ISA and EISA regions are decoded and shadowed based on programmable configurations. Regions above 1 Mbyte with size 1 Mbyte or larger that are not mapped to memory may be reclaimed by setting the appropriate configuration in the MC. Three programmable memory gaps can be created and are called the *Low Memory Gap Region*, the *Memory Gap Region* and the *High Memory Gap Region*.

2.0 INTEL 450GX PCISET

The Intel 450GX PCIset includes the features discussed for the Intel 450KX PCIset and provides the additional capabilities described in this section. This PCIset consists of the 82454GX PCI Bridge (PB) and the Memory Controller (MC). The MC for the 450GX consists of the 82453GX DRAM Controller (DC), the 82452GX Data Path (DP), and four 82451GX Memory Interface Controllers (MIC). The 450GX permits two PBs and two MCs in a system. In addition to parity support on the host bus described for the 450KX, the 450GX generates and checks ECC over the host data lines. This feature can be enabled/disabled during configuration.

One aspect of the 450GX is that it can be used as a drop-in replacement for an 450KX design. Additional pins are added in such a way that proper wiring of 450KX test pins (GTLHI, TESTLO, and TESTHI) will allow an 450GX to operate in the same system while functioning exactly as an 450KX.

GX PCI Bridge (PB)

Two 82454GX PBs can be used in a system. Dual PBs provide a modular approach to I/O performance improvements. Compatibility versus speed are addressed with an optional compatibility operating mode to guarantee standard bus compatible operation when needed, and allow bus concurrency when possible.

In a dual PB system, one PB is configured by strapping options at power-up to be the *Compatibility PB*. This PB provides the PC compatible path to Boot ROM and the ISA/EISA bus. The second PB is configured by the strapping options to be the *Auxiliary PB*. The Compatibility PB is the highest priority bridge to ensure a proper response time for ISA bus masters. When two PBs are on the host bus, the Compatibility PB handles arbitration with an internal arbiter.

GX Memory Controller (MC)

The memory configuration can be either 4-way interleaved, 2-way interleaved, or non-interleaved. Both single-sided and double-sided SIMMs are supported. DRAM technologies up to 64Mbit at speeds of 50ns, 60ns, and 70ns can be used. Asymmetric DRAM is supported up to two bits of asymmetry (e.g., 12 row address lines and 10 column address lines). The maximum memory size is 4 Gbytes for the 4-way interleaved configuration, 2 Gbytes for the 2-way interleaved configuration, and 1 Gbyte for the non-interleaved configuration using 64 Mbit technology. The MC provides a 64-bit data path to main memory (72-bits including ECC) for each interleave (288 bits for a 4-way interleave design).

3.0 HOST BUS EFFICIENCY

The Pentium Pro processor bus achieves high bus efficiency by providing support for multiple, pipelined transactions. A single Pentium Pro processor may have up to four transactions outstanding at the same time, and can be configured to support up to eight transactions active on the Pentium Pro processor bus at any one time. The PB and MC support a choice of one or eight active transactions on the Pentium Pro processor system bus at one time (In-Order Queue depth).

The number of transactions that can target a particular bus client is configured separately from the total number of transactions allowed on the bus. Each PB can accept up to four transactions into its Outbound Request Queue that target its associated PCI bus. The PB also contains a four deep Inbound Queue that holds PCI initiated requests directed to the Pentium Pro processor bus. Each MC can accept up to four transactions that target its associated memory space.

Both the PB and MC provide four 32-byte buffers for outbound data and four 32-byte buffers for inbound data. For the PB, the outbound data refers to CPU-to-PCI writes or PCI reads from the CPU bus and inbound data refers to PCI-to-CPU writes or CPU reads from PCI. For the MC outbound data refers to CPU writes to main memory and inbound data refers to CPU reads of main memory.

The maximum data transfer that is supported by the Pentium Pro processor bus is four 64-bit wide transfers. This transfer satisfies the 32-byte cache line size of the Pentium Pro processor interface. The Pentium Pro processor supports operations that are not completed in the order in which they were requested. This 'deferred response' capability allows the Pentium Pro processor bus to be freed to execute other requests while waiting for the response from a request to a device with relatively long latency. Note that the 450 PCIset does not defer requests to itself, nor does it (the PB) allow its transactions to be deferred.

4.0 SYSTEM MEMORY MAP

A Pentium Pro processor system can have up to 64 Gbytes of addressable memory. The lower 1 Mbyte of this memory address space is divided into regions that can be individually controlled with programmable attributes such as disable, read/write, write only, or read only.

At the highest level, the address space is divided into four conceptual regions as shown in Figure 2. These are the 0–1 Mbyte Compatibility Area, the 1 Mbyte to 16 Mbyte Extended Memory region used by ISA, the 16 Mbyte to 4 Gbyte Extended Memory region used by EISA, and the 4 Gbyte to 64 Gbyte Extended Memory introduced by 36 bit addressing. Each of the regions are divided into subregions, as described in the following sections.

For the 450GX, up to two MCs can be placed in the address space spanned by these regions. In a PC architecture, the only restrictions on memory placement are that there be memory starting at address 0 and that there be enough memory to operate a system. The MCs in a system need not have contiguous address spaces. Each MC also supports two memory ranges for the memory connected to the MC, by providing a high memory gap range register that defines the space between the two ranges of memory. This range effectively defines the top address for the lower memory range and the base address for the upper memory range.

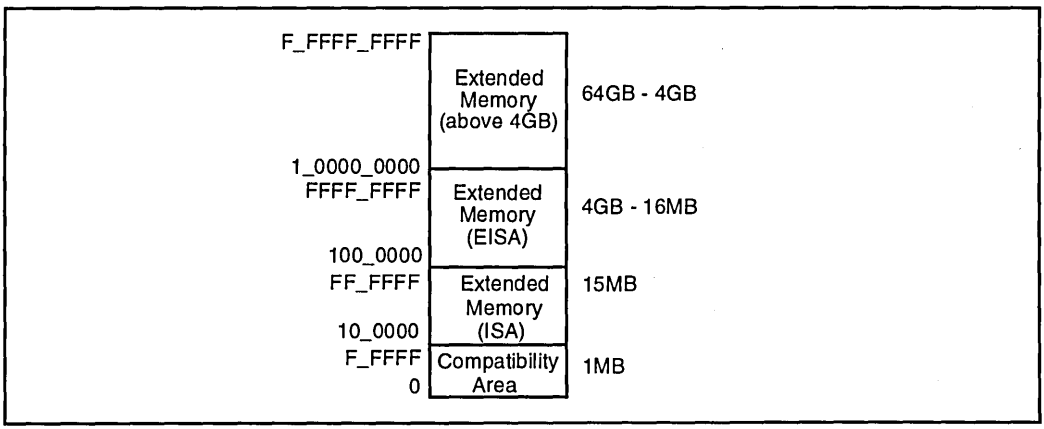

Figure 2. Pentium Pro Processor Memory Address Space.

4.1 Compatibility Area

The first region of memory is called the Compatibility Area because it was defined for early PCs. This region is divided into 5 subregions, as shown in Figure 3.

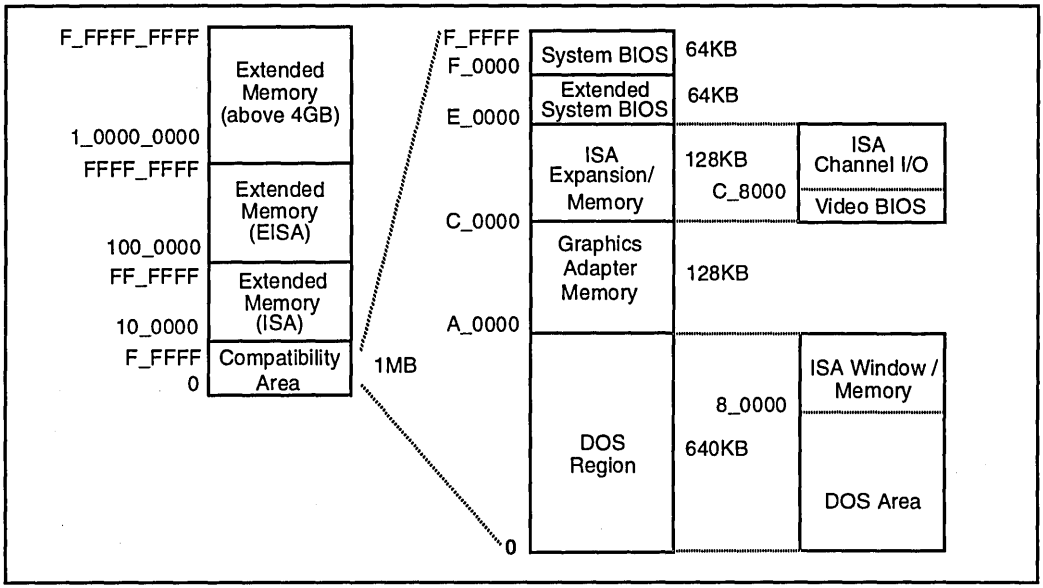

Figure 3. Expanded View of Compatibility Area.

DOS Region

The DOS Region is 640 Kbytes in the address range 00000h–9FFFFh. DOS applications execute here. This region is further divided into two parts. The 512 Kbyte area at 00000h–7FFFFh is always mapped to memory on the Pentium Pro processor bus (enabled in the MC), while the 128 Kbyte area from 80000h–9FFFFh can be mapped to memory on the Pentium Pro processor bus or PCI memory (enabled in the PB). This region can be programmed as disabled, read/write, write only, or read only.

Graphics Adapter Memory

The 128 Kbyte Graphics Adapter Memory region at A0000h–BFFFFh is normally mapped to a video device on the PCI bus. Typically, this is a VGA controller. If there are no graphics compatible devices, this region can be used as system memory. The range A0000h–AFFFFh (64 Kbytes) is also the default region for SMM space. The SMM region can be re-mapped by programming the SMM Range Register in the PB and MC.

ISA Expansion

The 128 Kbyte ISA expansion region is divided into eight 16 Kbyte blocks that can be independently programmed as disabled, read/write, write only, or read only providing the capability to "shadow" these regions in main memory. Typically, these blocks are mapped through the PB to ISA space.

Historically, the 32 Kbyte region from C0000h–C7FFFh has contained the video BIOS located on a video card in the ISA Expansion Area. However, in the high integration portable and desktop market video BIOS is more likely to be located in the Extended System BIOS or System BIOS regions that start at E0000h.

The 96 Kbyte area from C8000h–DFFFFh has usually been made available to expand memory windows in 16 Kbyte blocks, depending on the requirements of other channel devices in the corresponding ISA space. More recently, PCMCIA devices for the portable market have been assigned within this region.

This region could also be used as System Management Mode (SMM) memory.

Extended System BIOS

This 64 Kbyte region from E0000h–EFFFFh is divided into four 16 Kbyte blocks and may be mapped either to the memory controller or the PCI bridge. This region can be programmed as disabled, read/write, write only, or read only, providing the capability to shadow these regions in main memory. Typically, this area is used for RAM or ROM.

System BIOS

The 64 Kbyte region from F0000h–FFFFFh is treated as a single block. After power-on reset, the PB (Compatibility PB in an 450GX dual PB system) has this area R/W enabled to respond to fetches during system initialization. The MC(s) and Auxiliary PBs (450GX PCIset) have this area R/W disabled. This region can be programmed as disabled, read/write, write only, or read only, providing the capability to shadow these regions in main memory.

4.2 Extended Memory (ISA)

The ISA Extended Memory region in Figure 4 covers 15 Mbytes ranging from 100000h–FFFFFFh. There are three programmable ranges that may be mapped to the ISA Extended Memory region of the MC—the Low Memory Gap range, the Memory Gap Range, and the High memory Gap Range. Memory in these ranges, that would normally be "lost", is recovered by the MC by extending the effective top of system memory, if reclaiming is enabled. The Memory Gap Range and High Memory Gap range are also programmable ranges in the PB. The PB also has a programmable PCI Frame Buffer Range.

Low Memory Gap Range (MC Only)

The Low Memory Gap range can start on any 1 Mbyte boundary in the ISA or EISA Extended Memory region, and can be 1, 2, 4, 8, 16, or 32 Mbytes. This region defines a "hole" in system DRAM space where accesses can be directed to the PCI bus. The Low Memory Gap Range is used by ISA devices such as LAN or linear frame buffers which are mapped into the ISA Extended region, or by any EISA or PCI device. The Low Memory Gap Range must reside at the lowest address of the three memory gaps, if it is enabled.

PCI Frame Buffer Range (PB Only)

The PCI Frame Buffer range can start on any 1 Mbyte boundary in either the ISA Extended Memory region or the EISA Extended Memory Region, and can be 1, 2, 4, 8, 16,or 32 Mbytes.

Memory Gap Range (MC and PB)

The Memory Gap Range can start on any 1 Mbyte boundary, above 1 Mbyte, and can 1, 2, 4, 8, 16, or 32 Mbytes. This region defines a "hole" in system DRAM space where accesses can be directed to the PCI bus. The Memory Gap Range is used by ISA devices such as LAN or linear frame buffers which are mapped into the ISA Extended region, or by any EISA or PCI device. The Memory Gap Range must reside above the Low Memory Gap Range and below the High Memory Gap Range, if it is enabled.

High Memory Gap Range (MC and PB)

The High Memory Space Gap can start on any 1 Mbyte boundary in either the ISA Extended Memory region, EISA Extended Memory Region, or the Extended Memory Region above 4Gbyte, and can extend up to 64 Gigabytes. It is defined by specifying a start and end address, both on 1 Mbyte boundaries. The High Memory Gap Range is provided as additional support for memory mapped I/O. The High Memory Gap Range must reside at the highest address of the three memory gap range registers, if it is enabled.

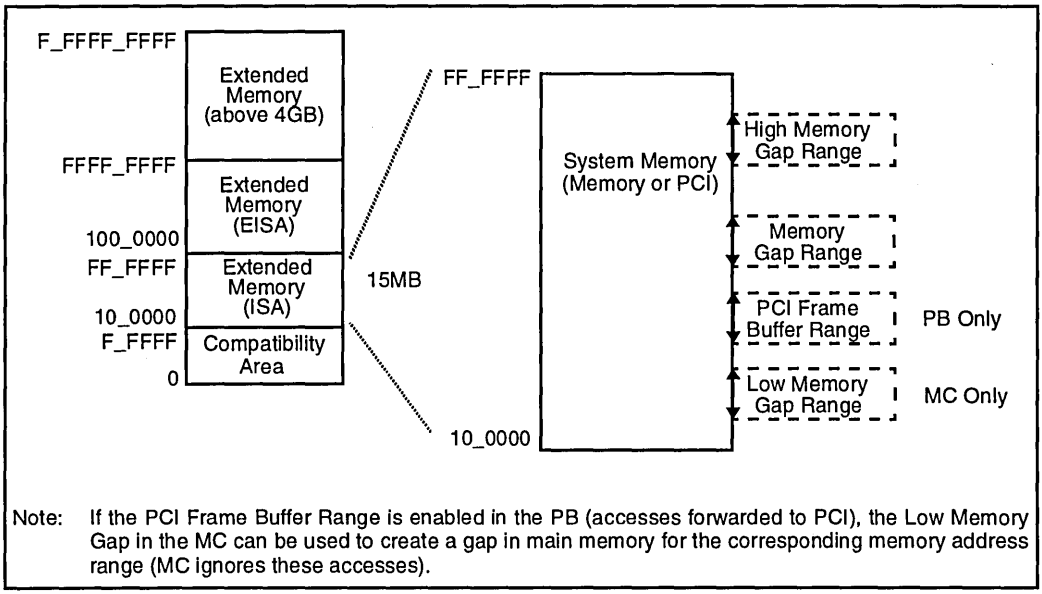

Note: If the PCI Frame Buffer Range is enabled in the PB (accesses forwarded to PCI), the Low Memory Gap in the MC can be used to create a gap in main memory for the corresponding memory address range (MC ignores these accesses).

Figure 4. Expanded View of Extended Memory (ISA)

4.3 Extended Memory (EISA)

The EISA Extended Memory region covers the 16 Mbyte to 4 Gbyte range (1000000h–FFFFFFFFh). This region is divided into three sections—System BIOS, APIC configuration space, and system memory. The APIC configuration space is contained within the system memory region (Figure 5). The Low Memory Gap, Memory Gap, and High Memory Gap ranges can also be enabled in this region.

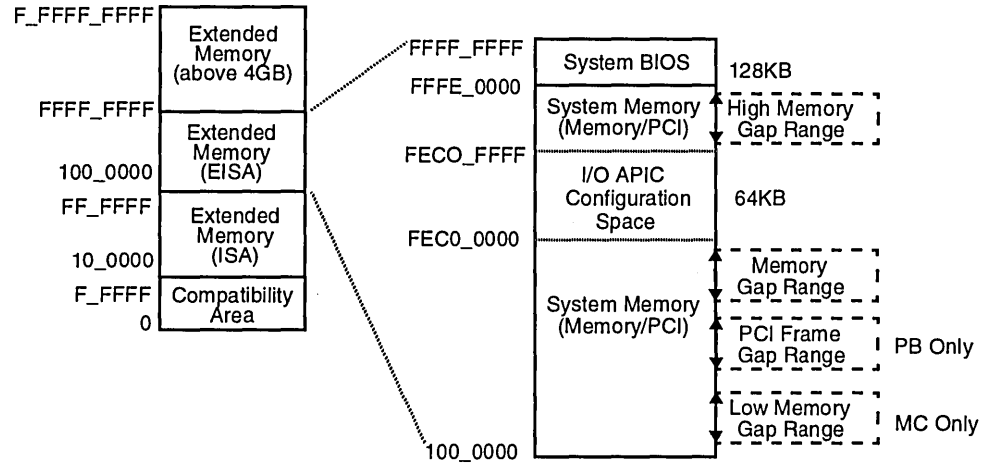

Note: If the PCI Frame Buffer Range is enabled in the PB (accesses forwarded to PCI), the Low Memory Gap in the MC can be used to create a gap in main memory for the corresponding memory address range (MC ignores these accesses).

Figure 5. Expanded View of Extended Memory (EISA).

System BIOS

The top 2 Mbytes of the EISA Extended Memory region is used for System BIOS (High BIOS). This is where the Pentium Pro processor begins execution after reset. If the PCI bus is bridged to an ISA bus, this region is aliased to the top 128 Kbyte of the ISA Extended Memory range.

The actual address space required for system BIOS is less than 2 Mbytes. However, the minimum Pentium Pro processor MTTR range for this region is 2 Mbytes. This establishes the minimum size for this gap. The MC supports enabling or disabling this region for access to the MC memory via the HBIOSR Register.

I/O APIC Configuration Space

The FEC00000h (4GB minus 20 MB) to FEC0FFFFh range is reserved for APIC configuration space which includes the default I/O APIC configuration space. Note that there is no I/O APIC unit in either the MC or PB. The default Local APIC configuration space is FEE00000h–FEE00FFFh.

Pentium Pro processor accesses to the Local APIC configuration space do not result in external bus activity since the Local APIC configuration space is internal to the Pentium Pro processor. However, an MTRR must be programmed to make the Local APIC range uncacheable (UC). The Local APIC base address in each Pentium Pro processor should be relocated to the FEC00000h (4GB minus 20 MB) to FEC0FFFFh range so that one MTRR can be programmed to 64 Kbyte for the Local and I/O APICs.

I/O APIC units (there should be at least one for each I/O subsystem) are located beginning at the default base address FEC00000h. The first I/O APIC (unit #0) is at FEC00000h. Each I/O APIC unit is located at FEC0x000h where x is I/O APIC unit 0 through F.

The address range between the APIC Configuration space and the High BIOS (FED00000h–FFDFFFFFh) is always mapped to local memory unless the range is above top of physical memory or The High BIOS and APIC Range are disabled in the PB and the range falls within a memory gap range. The MC supports enabling or disabling this region for access to the MC memory via the I/O APIC Range Register.

4.4 Extended Memory (above 4 Gbytes)

The Extended Memory region is from 4 Gbyte to 64 Gbyte (100000000h–FFFFFFFFFh). The PB and MC can be mapped into this range. The Memory Gap Range and High Memory Gap Range are both available for use within the Extended memory region (above 4 Gbyte).

4.5 System Management Mode (SMM)

A Pentium Pro processor asserts SMMEM# in its Request Phase if it is operating in System Management Mode. SM code resides in SM memory space. SM memory can overlap with memory residing on the Pentium Pro processor bus or memory normally residing on the PCI bus. The MC and PB determine where SM memory space is located through the value programmed in their respective SMM Range Registers.

5.0 I/O SPACE (PB ONLY)

The PB optionally supports ISA expansion aliasing (Figure 6). When ISA expansion aliasing is enabled, the ranges designated as I/O Expansion are internally aliased to the 100–3FFh range before the I/O Space Range registers are checked. Note that all devices on the Pentium Pro processor bus that are mapped into I/O space must have I/O aliasing consistently enabled/disabled.

For the Intel 450GX PCIset, the PB allows I/O addresses to be mapped to the Pentium Pro processor bus or through designated bridges in a multi bridge system. Two I/O Space Range registers allow the PB to decode two I/O address ranges. If the address range is enabled, transactions targeting that range are forwarded to the PCI bus. If the address range is disabled, the transaction is ignored.

6.0 MEMORY MAPPED I/O

The PB allows memory addresses to be mapped to the host bus, or for the 450GX PCIset, through the other bridge in a dual PB system. Memory mapped I/O devices can be located anywhere in the 64 Gbyte address space. The Frame Buffer Range allows the PB to decode memory mapped I/O space extending up to 4 Gbyte. The Memory Space Gap and High Memory Gap registers allow the PB to decode two address ranges extending up to 64 Gbytes.

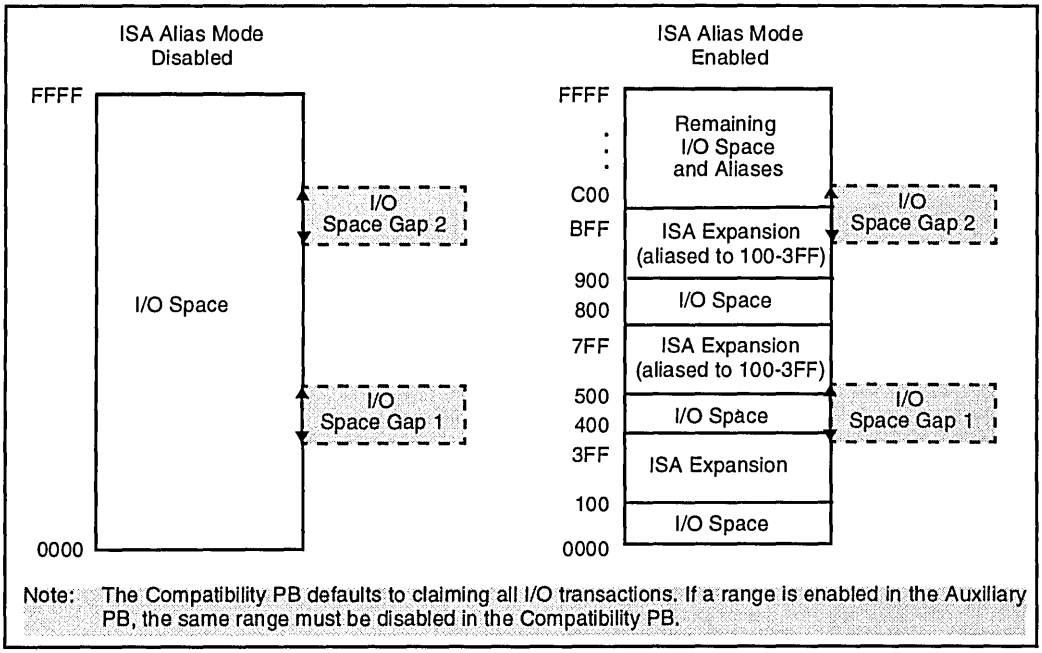

Note: The Compatibility PB defaults to claiming all I/O transactions. If a range is enabled in the Auxiliary PB, the same range must be disabled in the Compatibility PB.

Figure 6. View of I/O Space

Chapter 2

82454KX/GX PCI Bridge (PB)

82454KX/GX PCI Bridge (PB)

- Supports the Pentium® Pro Processor at 60 MHz and 66 MHz Bus Speeds
- PCI Specification 2.0 Compliant
- 64-Bit Data Bus and 36-Bit Address Bus
- Parity Protection on Control Signals
  - ECC Protection on Data Bus (450GX)
- Up to Eight Deep In-Order Queue
- Four Deep Outbound Request Queue
- Dual-Processor Support (450KX)
  - Quad-processor Support (450GX)
- Four Cache Line Size Read and Write Buffers
- GTL+ Host Bus Interface
- Synchronous PCI Interface
- 32-bit Address/Data PCI Bus (64-bit Dual Cycle Address Support)
- Parity Protection on All PCI Bus Signals
- Four Deep Inbound Request Queue
- Data Collection/Write Assembly of Line Bursts.
- Single Chip: Combined Controller and Data Path in a 304-Pin QFP or 352 BGA
  - Internal Bridge Arbiter For Two PBs in a system (450GX)
- Support for 3.3V and 5V PCI Devices
- On-Chip Digital PLL (DPLL)
- Component and In-System Connectivity Test Support (JTAG)

The 82454KX/GX PB are single-chip PC-compatible host-to-PCI bridges. A rich set of Host-to-PCI and PCI-to-Host bus transaction translations optimize bus bandwidth and improve system performance. All ISA and EISA regions are supported. Three programmable memory gaps can be created—a PCI Frame Buffer Region and two general-purpose memory gaps (the Memory Gap Region and the High Memory Gap Region). The PB has a synchronous interface to the Pentium Pro processor bus and supports a derived clock for the synchronous PCI interface. The PB generates and checks ECC over the host data bus (82545GX only), and generates and checks parity over the address and request/response signal lines (both 82454KX and 82454GX). The PB also checks address and data parity on the PCI bus. For the 82454GX, two PBs can be used in a system.

The Intel 450KX/GX PCIsets may contain design defects or errors known as errata. Current characterized errata are available upon request.

This document describes both the 82454KX and 82454GX PBs. Unshaded areas describe features common to the 82454KX and 82454GX. Shaded areas, like this one, describe the 82454GX operations that differ from the 82454KX.

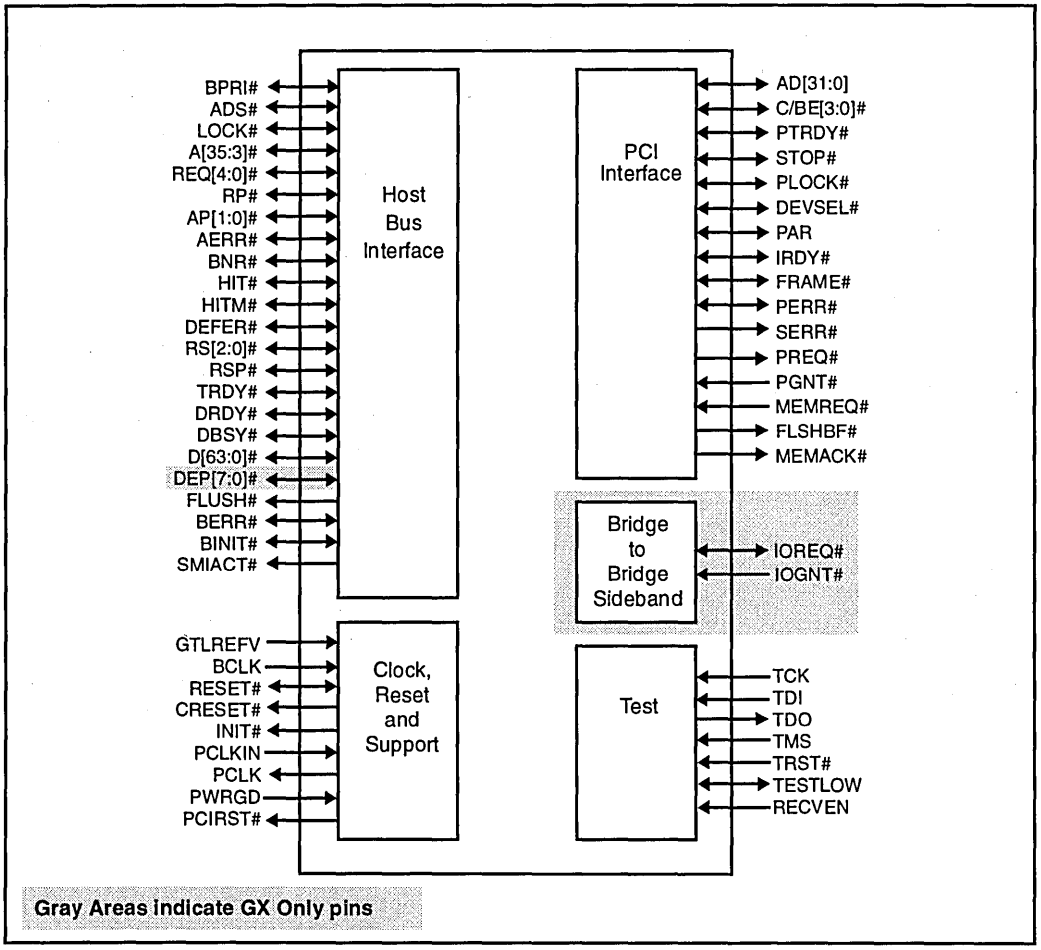

Figure 1. 82454KX/GX Simplified Block Diagram

1.0 PB SIGNAL DESCRIPTIONS

This section contains a detailed description of each signal. The signals are arranged in functional groups according to their interface.

Note that the '#' symbol at the end of a signal name indicates that the active, or asserted state occurs when the signal is at a low voltage level. When '#' is not present at the end of a signal name, the signal is asserted when at the high voltage level.

The terms assertion and negation are used extensively. This is done to avoid confusion when working with a mixture of 'active-low' and 'active-high' signals. The term **assert**, or **assertion** indicates that a signal is active, independent of whether that level is represented by a high or low voltage. The term **negate**, or **negation** indicates that a signal is inactive.

The following notations are used to describe the signal type.

I	*Input* is a standard input-only signal.
O	*Totem* Pole Output is a standard active driver.
I/O	*Input/Output* is bi-directional, tri-state signal.
GTL+	*GTL+* Processor bus signal defined for 1.5V operation.
CMOS	Rail-to-Rail *CMOS* Tolerant to 5V levels.
PCI	CMOS signal specifically meeting PCI Specification 2.0.
Analog	Reference Voltage.

1.1 PB Signals

Table 1. Host Bus Interface Signals

Signal	Type	Description
A[35:3]#	I/O, GTL+	**ADDRESS BUS.** A[35:3]# contains the transaction address on the clock cycle with ADS# asserted. Byte enables, deferred ID, and additional transaction information are encoded on these lines during the cycle following ADS#. Note that the PB never asserts Defer Enable when it is a bus master.
ADS#	I/O, GTL+	**ADDRESS STROBE.** ADS# is asserted during the first cycle of the Request Phase to indicate valid address and command signals.
AERR#	I/O, GTL+	**ADDRESS ERROR.** AERR# is asserted by any agent that detects an address parity error, If enabled in the EXERRCMD Register.
AP[1:0]#	I/O, GTL+	**ADDRESS PARITY.** AP1# covers A[35:24]# and AP0# covers A[23:3]#. AP[1:0]# is valid on both cycles of the request.
BERR#	I/O, GTL+	**BUS ERROR.** BERR# is asserted by any agent that observes an unrecoverable bus protocol violation, if enabled in the EXERRCMD Register.
BINIT#	I/O, GTL+	**BUS INITIALIZATION.** BINIT# is asserted to re-initialize the bus. The PB terminates any ongoing PCI transaction at this time and resets its inbound and outbound queues. No configuration registers or error logging registers are affected.
BNR#	I/O, GTL+	**BLOCK NEXT REQUEST.** BNR# is asserted by an agent to prevent the request bus owner from issuing further requests.
BPRI#	I/O, GTL+	**PRIORITY AGENT BUS REQUEST.** BPRI# is issued by the high priority bus agent to acquire the request bus. The high priority agent is always the next bus owner.

Table 1. Host Bus Interface Signals (Continued)

Signal	Type	Description
D[63:0]#	I/O, GTL+	**DATA BUS.** The data bus consists of eight bytes. All bytes are valid for line transfers. The valid bytes are determined by the byte enables that are asserted in the second cycle of the request phase.
DBSY#	I/O, GTL+	**DATA BUSY.** DBSY# is asserted by the data bus owner to hold the data bus for the next cycle. DBSY# is not asserted for single cycle transfers.
DEFER#	I/O, GTL+	**DEFER.** DEFER# is driven by the addressed agent to indicate that the transaction cannot be guaranteed bus completion.
DEP[7:0]#	I/O, GTL+	**DATA ECC.** On the host bus, DEP[7:0]# are used for ECC on the D[63:0]# signals.
DRDY#	I/O, GTL+	**DATA READY.** DRDY# is driven by the data bus owner for each cycle that contains valid data. DRDY# is negated to indicate idle cycles during the data phase.
FLUSH#	O, CMOS	**FLUSH.** The PB asserts FLUSH# to cause the processor to stop caching new lines, writeback all cache lines in the Modified state, and disable further caching until FLUSH# is negated. In an 82454GX dual PB system this signal is only available on the Compatibility PB and is not available on the Auxiliary PB.
HIT#	I/O, GTL+	**HIT.** The PB asserts HIT# and HITM# together to extend the snoop window of a transaction targeting its PCI bus. Since the PB is not a caching agent, it never asserts HIT# alone.
HITM#	I/O, GTL+	**HIT MODIFIED.** The PB asserts HIT# and HITM# together to extend the snoop window of a transaction targeting its PCI bus. Since the PB is not a caching agent, it never asserts HITM# alone.
LOCK#	I/O, GTL+	**LOCK.** The LOCK# signal is asserted for an indivisible sequence of transactions.
REQ[4:0]#	I/O, GTL+	**REQUEST TYPE.** REQ[4:0]# contain the command on the clock with ADS# asserted and data size/length information on the next clock.
RP#	I/O, GTL+	**REQUEST PARITY.** RP# is even parity that covers REQ[4:0]# and ADS#. RP# is valid on both cycles of the request.
RS[2:0]#	I/O, GTL+	**RESPONSE.** RS[2:0]# encode the response to a request.
RSP#	I/O, GTL+	**RESPONSE PARITY.** RSP# provides response parity for RS[2:0]#.
SMIACT#	O, CMOS	**SMI ACKNOWLEDGE.** SMIACT# is asserted when the PB detects a host SMI Acknowledge special transaction (regardless of its initiator) with SMMEM# asserted. Once asserted, SMIACT# remains asserted until the PB detects a host SMI Acknowledge special transaction with SMMEM# negated. In an 82454GX dual PB system this signal is only available on the Compatibility PB and is not available on the Auxiliary PB.
TRDY#	I/O, GTL+	**TARGET READY.** TRDY# is driven by the target of the data to indicate it is ready to receive data.

Table 2. PCI Interface Signals

Signal	Type	Description
AD[31:0]	I/O, PCI	**PCI ADDRESS/DATA.** Addresses and data are multiplexed on this bus. The physical byte address is output during the address phase and the data follows in the subsequent data phase(s).
C/BE[3:0]#	I/O, PCI	**BUS COMMAND AND BYTE ENABLES.** C/BE[3:0]# contains commands during the address phase and byte enables during the data phase.
DEVSEL#	I/O, PCI	**DEVICE SELECT.** DEVSEL# is driven by the device that has decoded its address as the target of the current access.
FLSHBF#	I, CMOS	**FLUSH BUFFERS.** This sideband signal is typically generated by a standard PCI bus bridge (e.g., ISA or EISA bridge) to command the PB to flush all write post buffers pointed toward the PCI bus and disable further posting. Once all buffers are flushed, the PB asserts MEMACK# until FLSHBF# is negated. **FLSHBF# MEMREQ# Function** 0 0 No Action. 0 1 Reserved. 1 0 APIC Flush. Flush buffers pointing toward PCI. 1 1 Guaranteed Access Time (GAT) mode. Guarantee PCI bus immediate access to the CPU bus. Flush all buffers, request queues, empty in-order queue, and retain host bus ownership.
FRAME#	I/O, PCI	**PCI FRAME.** FRAME# is driven by a master to indicate the beginning and end of a transaction.
IRDY#	I/O, PCI	**PCI INITIATOR READY.** IRDY# is asserted by the master to indicate that it is able to complete the current data transfer.
MEMACK#	O, CMOS	**MEMORY ACKNOWLEDGE.** MEMACK# is generated in response to FLSHBF# or MEMREQ# generated by a standard bus bridge.
MEMREQ#	I, CMOS	**MEMORY REQUEST.** This sideband signal is typically generated by a standard bridge (e.g., ISA or EISA bridge) to guarantee access latency from standard bus masters to main memory (see FLSHBF# description). Once all buffers have been flushed, the PB asserts MEMACK# continuously until MEMREQ# is negated.
PAR	I/O, PCI	**PCI PARITY.** PAR is driven to even parity across AD[31:0] and C/BE[3:0]# by the master during address and write data phases. The target drives PAR during read data phases.
PERR#	I/O, PCI	**PCI PARITY ERROR.** PERR# is pulsed by an agent receiving data with bad parity one clock after PAR is asserted.
PGNT#	I, CMOS	**PCI GRANT.** PGNT# indicates to the PB that it has been granted the PCI bus.
PLOCK#	I/O, PCI	**PCI LOCK.** PLock# is asserted by an agent requiring exclusive access to a target.

Table 2. PCI Interface Signals (Continued)

Signal	Type	Description
PREQ#	O, CMOS	**PCI REQUEST.** The PB asserts PREQ# to the PCI arbiter requesting the PCI bus.
PTRDY#	I/O, PCI	**PCI TARGET READY.** PTRDY# is asserted by the target to indicate that it is able to complete the current data transfer.
SERR#	O, PCI	**PCI SYSTEM ERROR.** SERR# is asserted by the PCI bridge to alert the system of serious errors. Several events (e.g., address parity errors, data parity errors, etc.) can optionally result in an SERR#. In a typical system, SERR# causes an NMI (e.g., by a PCI-to-ISA or PCI-to-EISA bridge).
STOP#	I/O, PCI	**STOP.** Stop# is a request from the target to stop the current transaction.

Bridge to Bridge Sideband Signals

The IOREQ# and IOGNT# signals are not driven or sampled in a single bridge system. During a power-on reset, IOREQ# and IOGNT# provide part of the PB's PCI Bridge Device Number. See Section 3.7 for details.

Table 3. Bridge to Bridge Sideband Signals

Signal	Type	Description
IOGNT#	I, CMOS	**I/O GRANT (82454GX ONLY).** The Compatibility PB is the bridge arbiter and IOGNT# is an input from Auxiliary PB requesting ownership of the host bus.
IOREQ#	I/O, CMOS	**I/O REQUEST (82454GX ONLY).**The Compatibility PB is the bridge arbiter and this signal is a host bus grant from the Compatibility PB to the Auxiliary PB.

Table 4. Clock, Reset, and Support Signals

Signal	Type	Description
BCLK	I, CMOS	**BUS CLOCK.** BCLK is the host bus clock input to the PB. All host bus timings are referenced to the rising edge of this clock. Note that the BCLK input to the PB must be running for 10 clocks before the assertion of PWRGD.
CRESET#	O, CMOS	**CMOS RESET.** CRESET# is a CMOS version of RESET#. RESET# and CRESET# are asserted simultaneously. The negation of CRESET# is delayed two clocks from the negation of RESET#. CRESET# can be used to control an external mux to select the Pentium Pro processor clock ratio during RESET#. In an 82454GX dual PB system, this signal is only available on the Compatibility PB and is not available on the Auxiliary PB.
GTLREFV	I, Analog	**GTL REFERENCE VOLTAGE.** This voltage is the 1.0 Volt reference for the GTL+ receivers. This should be created by a voltage divider from V_{TT} (1.5V)

Table 4. Clock, Reset, and Support Signals (Continued)

Signal	Type	Description
INIT#	O, CMOS	**INITIALIZATION.** INIT is asserted by the PB (Compatibility PB in an 82454GX dual PB system) to generate a soft reset to the processor. If INIT is asserted on the falling edge of reset, BIST executes in the CPU before the processor boots from ROM. In an 82454GX dual PB system, this signal is only available on the Compatibility PB and is not available on the Auxiliary PB.
PCIRST#	O, CMOS	**PCI RESET.** PCIRST# is asserted by the PB to reset PCI bus devices for power-on reset, programmed hard reset (TRC Register), and programmed PCI reset (PCIRST Register).
PCLK	O, CMOS	**PCI CLOCK.** This signal is an output that is derived from the processor clock (derived frequency is 1/2 the host bus frequency). The derived PCI clock should be externally buffered with a low skew clock driver. An external pull-down resistor is required on this signal.
PCLKIN	I, CMOS	**PCI CLOCK INPUT.** PB reference clock for all PCI bus transactions in both PB PCI clock modes. All PCI timing is referenced to the rising edge of this clock. PCLKIN is provided by an external low skew clock driver and should be coincident with PCLK at the PCI slots. This can be achieved by adjusting trace lengths.
PWRGD	I, CMOS	**POWER GOOD.** PWRGD provides a power-on reset to the PB (see Section 3.7). The PB asserts PCIRST# when PWRGD is not asserted to tri-state the busses to prevent contention of active output buffers on the PCI bus. In an 82454GX dual PB system, all PBs assert PCIRST# when PWRGD is not asserted. Only the Compatibility PB uses PWRGD as a power-on reset.
RESET#	I/O, GTL+	**RESET.** The PB resets the host bus devices (asserts RESET#) on power-up or when programmed through the TRC Register. The PB initializes its internal registers to the default values, except for the Bridge Device Number Register and the Configuration Driven on Reset Register. In an 82454GX dual PB system, only the Compatibility PB drives this signal. For Auxiliary PBs, this signal is an input.

Table 5. Test Signals

Signal	Type	Description
GTLHI	I/O GTL+	**GTL TEST HI.** These signals must be tied to V_{TT} using a 10KΩ resistor for proper operation in both test and normal operating modes.
TCK	I CMOS	**JTAG Test Clock.** When TMS is tied low, this signal has no affect on normal operation.
TDI	I CMOS	**JTAG Test Data In.** When TMS is tied low, this signal has no affect on normal operation.
TDO	O CMOS	**JTAG Test Data Out.** When TMS is tied low, this signal has no affect on normal operation.
TESTHI	I/O	**TEST HIGH.** These signals must be tied high using a 10KΩ resistor for proper operation in both test and normal operating modes.
TESTLO	I/O	**TEST LOW.** These signals must be tied low using a 1KΩ resistor for proper operation in both test and normal operating modes.
TMS	I CMOS	**JTAG Test Mode Select.** This signal must be tied low for normal operation.
TRST#	I CMOS	**JTAG Test Reset.** When TMS is tied low, this signal has no affect on normal operation.
RECVEN	I	**RECEIVER ENABLE.** Useful for component test. This signal is negated with PWRGOOD to disable GTL+ receivers and tri-state outputs for board test.

1.2 Signal State During Reset

Table 6 shows the state of all PB output and bi-directional signals during a hard reset (RESET# asserted).

Table 6. Signal State During Reset

Signal	State
A[35:3]#	Not Driven[1]
ADS#	Not Driven
AP[1:0]	Not Driven
BINIT#	Not Driven
BPRI#	Not Driven
CRESET#	Low[5]
DBSY#	Not Driven
DEP[7:0]#	Not Driven
DRDY#	Not Driven
FRAME#	Not Driven
HITM#	Not Driven
INIT#	High[5]
LOCK#	Not Driven
PAR	Not Driven
PCLK	Driven[6]
PLOCK	Not Driven
PTRDY#	Not Driven
REQ[4:0]#	Not Driven
RS[2:0]#	Not Driven
SERR#	Not Driven
STOP#	Not Driven
TRDY#	Not Driven

Signal	State
AD[31:0]	Not Driven
AERR#	Not Driven
BERR#	Not Driven
BNR#	Not Driven
C/BE[3:0]#	Not Driven
D[63:0]#	Not Driven
DEFER#	Not Driven
DEVSEL#	Not Driven
FLUSH#	High[5]
HIT#	Not Driven
IOREQ#	Input[2]
IRDY#	Not Driven
MEMACK#	Not Driven
PCIRST#	Low
PERR#	Not Driven
PREQ#	Not Driven
RESET#	Low[3]
RP#	Not Driven
RSP#	Not Driven
SMIACT#	High[5]
TDO	Tri-state during TRST#

NOTES:

1. During a power-on reset, A[12:5]# are inputs providing configuration information. For the 82454KX/GX, during a programmed hard reset (via the Compatibility PB's TRC register), the Compatibility PB drives these signals and the all other host bus devices sample these signals.
2. For the 82454GX during a power-on reset, IOGNT# and IOREQ# are inputs used to set the PB configuration mode.
3. For the 82454GX after a power-on reset, RESET# is an output from all PBs until the PBs have read in their PBID from the IOGNT# and IOREQ# signals. After the PBs receive their PBID, RESET# is an output from the Compatibility PB and an input to the Auxiliary PBs.
4. During a power-on reset, INIT# is driven inactive. The PB can be programmed (via the TRC Register) to drive this signal low during a programmed hard reset to invoke CPU Built-In Self Test (BIST).
5. These signals not used in the Auxiliary bridge in 82450GX systems.
6. Tri-state during PWRGD inactive.

2.0 PB REGISTER DESCRIPTION

The PB contains two sets of registers (I/O space registers and PCI configuration registers) that are accessed via the host CPU I/O address space. The I/O space registers provide access to the PCI configuration registers through an indirect address scheme.

The PB internal registers (both I/O space registers and PCI configuration registers) are only accessible by the host bus and cannot be directly accessed by PCI masters. The registers can be accessed as Byte, Word (16-bit), or Dword (32-bit) quantities, with the exception of CONFADD which can only be accessed as a Dword. The following nomenclature is used for access attributes.

RO *Read Only.* If a register is read only, writes to this register have no effect.
R/W *Read/Write.* A register with this attribute can be read and written.
R/WC *Read/Write Clear.* A register bit with this attribute can be read and written. However, a write of a 1 clears (sets to 0) the corresponding bit and a write of a 0 has no effect.

Some of the PB registers contain reserved bits. 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, unless otherwise specified in the individual register descriptions, 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 new values for other bit positions and written back. In some cases, software must program reserved bit positions to a particular value. This value is defined in the individual bit descriptions.

In addition to reserved bits within a register, the PB contains address locations in the PCI configuration space that are marked "Reserved". The PB responds to accesses to these address locations by completing the host transaction. Software should not write to reserved PB configuration locations in the device-specific region (above address offset 3Fh).

If RESET# is asserted (via either a power-on reset or by programming the TRC Register), the PB initializes its registers to the **default** value (except for the BDNUM and CONFVR Registers). The default state 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 operating parameters and optional system features that are applicable, and to program the PB registers accordingly. The PB (Compatibility PB in an 450GX system) can generate a programmed hard reset via the TRC Register.

In dual PB systems, during a hard reset (via a power-on hard reset or by programming the Compatibility PB's TRC Register), both PBs set their internal configuration registers to predetermined **default** conditions.

2.1 Initialization and Configuration

The PB (and MC) contain a configuration space that uses the same access mechanism as described in the PCI bus specification. With the exception of address decoding for BIOS accesses, the PB does not respond to host-initiated memory accesses until the associated registers are initialized. The device number for the PB is hardwired to 11001 for the 82454KX.

For dual PB systems, there are two PB configurations (Compatibility and Auxiliary PB). These configurations are defined by values on the IOGNT# and IOREQ# signal lines (high or low voltage levels) on the rising edge of PWRGD. The values on IOGNT# and IOREQ# define the PB Identification (PBID) and are reported in the DBNUM Register (offset 49h). Physical connections for the IOGNT# and IOREQ# signals are shown in the Section 3.5. The PBID value defines the lower two bits of the five-bit device number.

The bridge that is in the path to the Boot ROM is always the PB with device number=11001, and is referred to as the Compatibility PB. The Compatibility PB always decodes BIOS addresses after power-on reset.

NOTE:

When the address decode ranges of the 450KX/GX devices are being updated, no other host bus traffic is allowed. This means that the code that updates initial configuration must be non-cached (to prevent speculative reads). Further, in a multiprocessor system, precautions should be taken to assure that only one CPU is accessing configuration space at a time.

2.2 I/O Space Registers

The PB has three registers located in I/O Space—the Configuration Address (CONFADD) Register, the Turbo and Reset Control (TRC) Register, and the Configuration Data (CONFDATA) Register.

Note that in a dual PB system (82454GX only), the TRC Register is only in the Compatibility PB and the Auxiliary PB ignores this address.

The CONFADD and CONFDATA Registers provide a window into the PB's configuration space registers (see Section 2.3 for additional details). A specific PCI bus, device, and register are selected by writing to the CONFADD Register. Data is read from or written to the selected register by accessing the CONFDATA Register. Note that the CONFADD Register is only selected by DWord accesses to CF8h. This allows the CONFADD Register to overlap other byte registers (e.g., the TRC Register at CF9h). The CONFDATA Register is not selected unless configuration accesses are enabled in the CONFADD Register. This allows the CONFDATA Register to overlap other registers as well.

Table 7. I/O Space Registers

I/O Address	Mnemonic	Register Name	Access
CF8h	CONFADD	Configuration Address	R/W
CF9h	TRC	Turbo and Reset Control (Compatibility PB only)	R/W
CFCh	CONFDATA	Configuration Data	R/W

2.2.1 CONFADD—CONFIGURATION ADDRESS REGISTER

Address Offset: 0CF8h (Dword access only)
Default: 00000000h
Attribute: Read/Write

The CONFADD Register contains the Bus Number, Device Number, Function Number, and Register Number for which a subsequent configuration access is intended. For example, a write of 8000C8B8h to the CONFADD register can be used to access the 82453KX SMM Range Register.

Bits	Description
31	**Configuration Enable (CSE).** 1=Enable. 0=Disable.
30:24	**Reserved.**
23:16	**Bus Number (BUSNUM).** This field selects 1 of 255 possible buses in a system. When BUSNUM=00h, the target of the configuration cycle is a host bus device or the PCI bus directly connected to the PB (Compatibility PB in an 450GX system), depending on the Device Number field. The bus number for the PCI bus directly connected to a PB is programmed into the PBNUM Register (offset 4Ah). If the PB (or MC) is not the target (DEVNUM≤15) and the bus number matches the number in the PB's PBNUM Register, a type 0 configuration cycle is generated on the PCI bus. If the bus number is to a hierarchical PCI bus below the PB's PCI bus (BUSNUM is between the values programmed into the PCI Bus Number Register and the Subordinate PCI Bus Number Register), a type 1 configuration cycle is generated on the PCI bus with the Bus Number mapped to AD[23:16] during the address phase.
15:11	**Device Number (DEVNUM).** This field selects either an agent on the host bus (BUSNUM=00h and DEVNUM≥16) or an agent on the PCI bus selected by the bus number field. For the 82454KX, the PB device number is hardwired to 11001. This number is reported in the BDNUM Register. A device on the host bus has a device number greater than 16. Note that logically a PB can support up to 16 physical devices on the PCI bus connected directly to it. Thus, when BUSNUM=00h and DEVNUM≤15, a type 0 or type 1 configuration cycle is generated by the Compatibility PB. During a type 0 configuration cycle, this field is decoded and one of AD[31:16] is driven to a 1. For device number n (0≤n≤15), AD[16+n] is driven to a 1. During a type 1 configuration cycle, this field is mapped to AD[15:11]. In an 82454GX dual PB system, the device number for each PB is determined at power-up and is reported in the BDNUM Register (offset 49h). The PB device number is always equal to or greater than 16.

Bits	Description
10:8	**Function Number (FUNCNUM).** This field is mapped to AD[10:8] during PCI configuration cycles. This allows the configuration registers of a particular function in a multi-function device to be accessed. The PB responds to configuration cycles with a function number of 000b; all other function number values attempting access to the PB (BUSNUM=00h and DEVNUM matching PB device number) generate a type 0 configuration cycle on the PCI bus with no IDSEL asserted, which results in a master abort.
7:2	**Register Number (REGNUM).** This field selects one 32-bit register within a particular bus, device, and function as specified by the other fields in the CONFADD Register. This field is mapped to AD[7:2] during PCI configuration cycles.
1:0	**Reserved.**

2.2.2 TRC—TURBO AND RESET CONTROL

Address Offset: CF9h
Default: 00h
Attribute: Read/Write

This register enables/disables BIST, provides software generation of hard and soft resets, and enables/disables deturbo mode.

For the 82454GX in a dual PB system, this register is only available in the Compatibility PB and is not part of the Auxiliary PB.

Bits	Description
7:4	**Reserved.**
3	**CPU BIST Enable.** 1=Enable. 0=Disable. When enabled, the PB invokes CPU BIST when the CPU is reset (the value of this bit overrides the value of the CPU Hard Reset bit). Subsequent initiation of hard reset (through bit 2 of this register) causes the PB to perform a hard CPU reset, leaving INIT# asserted when RESET# is released initiating CPU BIST.
2	**Reset CPU.** 1=hard reset, soft reset, or hard reset with BIST (type is controlled by bits[3,1] of this register). The transition from 0 to 1 of this bit triggers the PB to initiate the CPU reset. Therefore, bits[3,1] should be programmed before this bit is set. In addition, bit 0 must be 0 before programming this register.
1	**Hard Reset Enable.** 1=Hard reset. 0=Soft reset. Reset occurs when the Reset CPU bit transitions from 0 to 1.
0	**Deturbo Enable.** 1=Enable. 0=Disable. Note that this bit must be set to 0 before setting bit 2 to 1.

2.2.3 CONFDATA—CONFIGURATION DATA REGISTER

Address Offset: 0CFCh
Default: 00000000h
Attribute: Read/Write

CONFDATA is a 32-bit read/write window into configuration space. The 32-bit portion of configuration space that is referenced by CONFDATA is determined by the contents of CONFADD. The byte enables during the 0CFCh access select which bytes of the 32-bit window are updated.

Bits	Description
31:0	**Configuration Data Window.** If bit 31 of CONFADD=1, an access to CONFDATA I/O space is mapped to configuration space using the contents of CONFADD.

2.3 PCI Configuration Space

The PB fully supports mechanism #1 for host accesses to PCI Configuration Space Registers (refer to the PCI Specification for details on mechanism 1). The PB can perform three types of configuration cycles.

1. An internal access is performed if the Bus Number is 0, and the Device Number selects this PB. No PCI cycles are generated.
2. A Type 0 translation is performed if the PCI device being configured is on this PB's PCI bus (the Bus Number matches the number in the PB's PCI Bus Number register), and the Device Number is less than or equal to 15.
3. A Type 1 translation is performed if the device being configured is on another hierarchical PCI bus below the PB's PCI bus (the Bus Number is between the PB's PCI Bus Number and Subordinate PCI Bus Number).

The PCI Configuration Space protocol requires that all PCI buses in a system be assigned a Bus Number. Furthermore, bus numbers must be assigned in ascending order within hierarchical buses. Each bridge must have a register that contains its PCI Bus Number and a register that contains its Subordinate PCI Bus Number. The PCI Bus Number and Subordinate PCI Bus Number must be loaded by POST code. The Subordinate PCI Bus Number is the bus number of the last hierarchical PCI bus under the current bridge. (The PCI Bus Number and Subordinate PCI Bus Number are the same in the last hierarchical bridge.) At the top of the hierarchy, peer bridges continue the ascending bus numbering scheme. Refer to the PCI specification for additional examples.

For the 450KX/GX, the implementation of the PCI configuration protocol logically maps the configuration registers of the PB (and MC) to bus number 0. These devices, which are on the host bus, use Device Numbers 16 through 30 (Figure 2). Device numbers below 15 can be used on the PCI bus that uses bus number 0. This allows a system to be designed with hierarchical PCI buses starting with bus number 0. All bridges have programmable PCI bus numbers and programmable subordinate PCI bus numbers as described in the PCI CSE protocol for dual PCI bridge systems. A PB's bus number register should be programmed to the number of the PCI bus immediately beneath it. However, the PB's configuration registers remain at Bus number 0.

The PB is the response agent for CPU accesses to the CONFADD location. The MC snoops writes to this location. The device selected by the CONFADD Register responds to CONFDATA accesses.

For the 82454GX, the Compatibility PB is the response agent for CPU accesses to the CONFADD location and the Auxiliary PB as well as the MCs snoop writes to this location.

The PB (both PBs in an 82454GX dual PB system) defaults to a PCI Bus Number equal to 0 after power-on reset. The MC is hard coded to bus number 0.

For the 450GX, each PB and MC must have a unique ID assigned at power-on reset (via strapping options). The relation between the PB and MC number, the Device Number, and the Host Bus Agent Number is shown in Table 8. The PBID is also used as the PB Agent ID when it is a host bus master.

Note that the 82454 does not support *programmable special cycles* of the PCI specification 2.0.

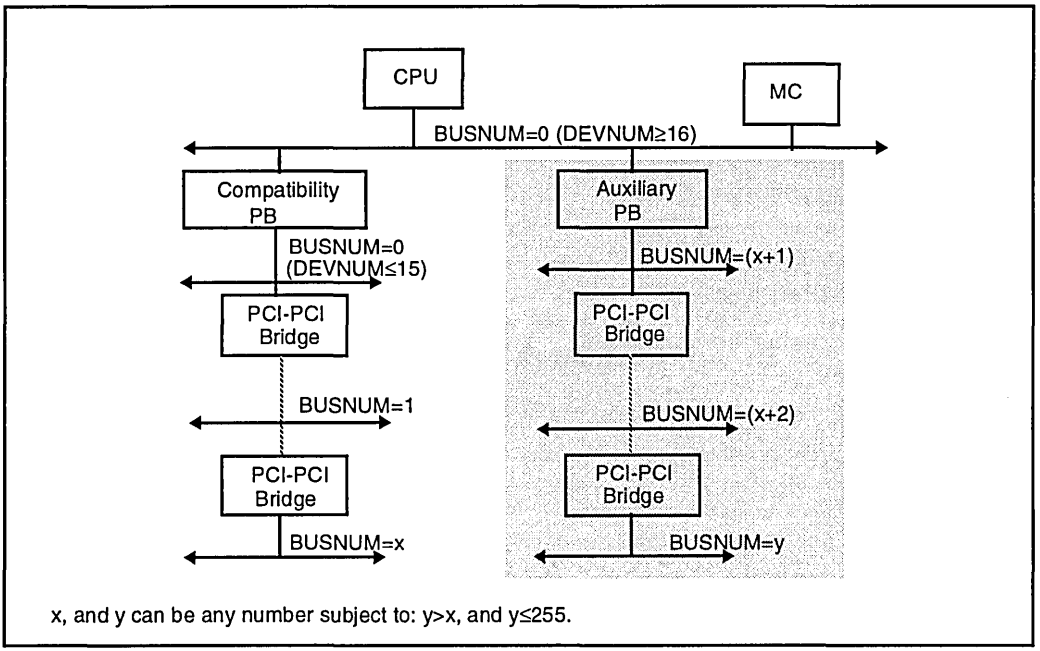

Figure 2. Bus Number Assignments

Table 8. Device Numbers for Bus Number 0

Device	Unique ID Loaded at Reset	PCI Device Number	Host Bus Agent ID
CPU	—	N/A	0000–0111
MC	00	10100	NA
MC	01	10101	NA
PB	01	11001	1001
PB	10	11010	1010
Reserved	All Others	All Others	All Others

2.4 PB PCI Configuration Registers

Table 9. PCI Configuration Space Registers

Address Offset	Mnemonic	Register Name	Access
00–01h	VID	Vendor Identification	RO
02–03h	DID	Device Identification	RO
04–05h	PCICMD	PCI Command	R/W
06–07h	PCISTS	PCI Status	RO, R/WC
08h	RID	Revision Identification	RO
09–0Bh	CLASSC	Class Code	RO
0Ch	CLSIZE	PCI Cache Line Size	RO
0Dh	PLTMR	PCI Latency Timer	R/W
0E	HEADT	Header Type	RO
0F	BIST	BIST Register	R/W
10–3Fh	—	Reserved	—
40–43h	TSM	Top of System Memory	R/W
44–47h	—	Reserved	—
48h	PDM	PCI Decode Mode	R/W
49h	BDNUM	Bridge Device Number	RO
4Ah	PBNUM	PCI Bus Number	R/W
4Bh	PSBNUM	PCI Subordinate Bus Number	R/W
4C	PBC	PB Configuration	R/W
4D–50h	—	Reserved	—
51h	DCC	Deturbo Counter Control	R/W
52h	—	Reserved	—
53h	CRWC	CPU Read/Write Control	R/W
54–55h	PRWC	PCI Read/Write Control	R/W
56h	—	Reserved	—
57h	SMME	SMM Enable	R/W
58h	VBAE	Video Buffer Area Enable	R/W
59–5Fh	PAM[0:6]	Programmable Attribute Map (7 Registers)	R/W
60–6Fh	—	Reserved	—
70h	ERRCMD	Error Reporting Command	R/W
71h	ERRSTS	Error Reporting Status	R/WC

Table 9. PCI Configuration Space Registers (Continued)

Address Offset	Mnemonic	Register Name	Access
72–77h	—	Reserved	—
78–79h	MGR	Memory Gap Range	R/W
7A–7B	MGUA	Memory Gap Upper Address	R/W
7C–7Fh	PFB	PCI Frame Buffer	R/W
80–87h	—	Reserved	—
88–8Bh	HMGSA	High Memory Gap Start Address	R/W
8C–8Fh	HMGEA	High Memory Gap End Address	R/W
90–97h	—	Reserved	—
98–9Bh	—	Reserved (450KX)	—
	IOSR1	I/O Space Range #1 (450GX)	R/W
9C	PCIRSR	PCI Reset	R/W
9D–9Fh	—	Reserved	—
A0–A3h	—	Reserved (450KX)	—
	IOSR2	I/O Space Range #2 (450GX)	R/W
A4h–A7h	APICR	I/O APIC Range	R/W
A8–AFh	—	Reserved	—
B0–B1h	CONFVR	Configuration Values Driven on Reset	R/W
B2–B3h	—	Reserved	—
B4–B5h	CSCONFV	Captured System Configuration Values	RO
B6–B7h	—	Reserved	—
B8–BBh	SMMR	SMM Range	R/W
BC	HBIOSR	High BIOS Register	R/W
BD–BFh	—	Reserved	—
C0–C3h	EXERRCMD	PB Extended Error Reporting Command	R/W
C4–C7h	EXERRSTS	PB Extended Error Reporting Status	R/WC
C8–CBh	PBRTMR	PB Retry Timers	R/W
CC–FFh	—	Reserved	—

2.4.1 VID—VENDOR IDENTIFICATION REGISTER

Address Offset: 00–01h
Default: 8086h
Attribute: Read Only

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

Bits	Description
15:00	**Vendor Identification.** This is a16-bit value (8086) assigned to Intel.

2.4.2 DID—DEVICE IDENTIFICATION REGISTER

Address Offset: 02–03h
Default: 84C4h
Attribute: Read Only

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

Bits	Description
15:00	**Device Identification.** This is a16-bit value (84C4) assigned to the PB.

2.4.3 PCICMD—PCI COMMAND REGISTER

Address Offset: 04–05h
Default: 0007h
Attribute: Read/Write

This register controls the PB's ability to respond to PCI cycles. See PCISTS Register for corresponding error reporting. See ERRCMD Register (70h) for additional controls.

Bits	Description
15:9	**Reserved.**
8	**SERR# Enable.** 1=Enable. 0=Disable. When enabled, the PB asserts SERR#, if the corresponding bits in the ERRCMD Register are enabled.
7	**Wait Cycle Control.** (Not Implemented). This bit is hardwired to 0.
6	**Parity Error Response Enable.** 1=Enable PCI parity error checking (See ERRCMD Register for generation of PERR# signal.). 0=Disable. Note that PCI parity errors will not be reported using SERR# unless both this bit and bit 8 are set to 1.
5	**Reserved.**
4	**Memory Write and Invalidate Enable.** 1=Enable. 0=Disable. When disabled, Memory Write commands are used.
3	**Reserved.**

Bits	Description
2	**Bus Master Enable—RO.** The PB does not support disabling its bus master capability. This bit is hardwired to 1.
1	**Memory Space Enable.** 1=Enable PCI memory accesses to the host bus. 0=Disable.
0	**I/O Space Enable.** 1=Enable PCI I/O accesses to the host bus. 0=Disable.

2.4.4 PCISTS—PCI STATUS REGISTER

Address Offset: 06–07h
Default: 0240h
Attribute: Read Only and Read/Write Clear

The PCISTS Register reports the occurrence of a PCI master abort/PCI target abort, system error, and parity errors. This register also indicates the DEVSEL# timing that has been set by the PB hardware. Software sets the bits labeled R/WC to 0 by writing a 1 to them.

Bits	Description
15	**Parity Error Detected—R/WC.** 1=PB detected a PCI address or data parity error. The PB checks all address cycles, regardless of the intended target, for address parity errors. When the PB is involved in a PCI transaction (as either master or target), it checks all data cycles for data parity errors. The Parity Error Detected bit is set independent of whether parity error reporting (bit 6 in the PCICMD Register) is enabled.
14	**Signaled System Error—R/WC.** 1=PB asserted the SERR# signal.
13	**Received Master Abort—R/WC.** 1=PB is PCI bus master and terminates its transaction (other than Special Cycle commands) with a master-abort.
12	**Received Target Abort—R/WC.** 1=PB as a PCI bus master received a target abort.
11	**Signaled Target Abort—R/WC.** 1=PB issued a target abort. This only happens for invalid byte enables during an I/O access or a Hard Failure from a host bus agent.
10:9	**DEVSEL# Assertion—RO.** Bits[10:9]=01 (indicates medium timing when the PB responds as a target).
8	**Data Parity Error Reported—R/WC.** This bit is set to 1 when all of the following conditions are met: 1.) The PB asserted PERR# or sampled PERR# asserted. 2.) The PB was the bus master for the transaction in which the error occurred. 3.) The Parity Error Response bit is set to 1 in the PCICMD Register.
7	**Fast Back-to-Back Capable—RO.** This bit is hardwired to 0 to indicate that the PB is not capable of accepting fast back-to-back transactions that are not to the same agent.
6:0	**Reserved.**

2.4.5 RID—REVISION IDENTIFICATION REGISTER

Address Offset: 08h
Default: See stepping information in the *450 GX/KX Specification Update*
Attribute: Read Only

This register contains the revision number of the PB.

Bits	Description
7:0	**Revision Identification Value.** This is an 8-bit value that indicates the revision identification number for the PB.

2.4.6 CLASSC—CLASS CODE REGISTER

Address Offset: 09–0Bh
Default: 060000h
Attribute: Read Only

This register contains the device programming interface information related to the Sub-Class Code and Base Class Code definition for the PB. This register also identifies the Base Class Code and the function sub-class in relation to the Base Class Code.

Bits	Description
23:16	**Base Class Code (BCC).** 06h=Bridge device.
15:8	**Sub-Class Code (SCC).** 00h=Host Bridge.
7:0	**Programming Interface (PI).** 00h=Hardwired as a Host-to-PCI Bridge.

2.4.7 CLSIZE—CACHE LINE SIZE REGISTER

Address Offset: 0Ch
Default: 08h
Attribute: Read Only

This register indicates the system cache line size. The value equals the number of 32-bit dwords in the cache line.

Bits	Description
7:0	**System Cache Line Size.** 08h=32 Byte cache line size.

2.4.8 PLTMR—PCI LATENCY TIMER

Address Offset: 0Dh
Default: 20h
Attribute: Read/Write

This register controls the duration of a burst cycle.

Bits	Description
7:0	**PCI Master Latency Timer value.** If PGNT# is negated during a PB initiated PCI burst cycle, the PB limits the duration of the burst cycle to the number of PCI clocks specified in this field. Masters capable of bursting multiple lines perform better with a higher value than the default. A value of 40h, for example, will allow a bursting master to always transfer at least four cache lines before a disconnect occurs.

2.4.9 HEADT—HEADER TYPE REGISTER

Address Offset: 0Eh
Default: 00h
Attribute: Read Only

This register indicates the header type for the PB.

Bits	Description
7:0	**Header Type (HTYPE).** 00h=Basic configuration space format.

2.4.10 BIST—BIST REGISTER

Address Offset: 0Fh
Default: 00h
Attribute: Read/Write

The Built-In Self Test (BIST) function is not supported by the PB. Writes to this register have no effect.

Bits	Description
7	**BIST Supported.** This read only bit is set to 0 indicating that the 82454 does not support BIST.
6	**Start BIST.** This function is not supported.
5:4	**Reserved.**
3:0	**Completion Code.** This read only field returns 0 when read.

2.4.11 TSM—TOP OF SYSTEM MEMORY REGISTER

Address Offset: 40–43h
Default: 0000h
Attribute: Read/Write

This register permits the PB to respond to memory transactions above the main memory range of the MC(s) on the host bus.

Bits	Description
31	**Host Bus Top of Main Memory Default Enable.** 1=Enable. 0=Disable. When enabled, the PB forwards all host bus memory space transactions between the Top of Memory (determined by bits [15:0] of this register) and 64 Gbytes to the PCI bus, except regions defined by the memory gap registers (MGR/MGUA and HMGSA/HMGEA Registers). When disabled, the PB ignores these transactions. Note that when memory accesses are enabled to be forwarded from the host bus to PCI, the PB blocks (ignores the transaction) the corresponding memory accesses initiated on the PCI bus from being forwarded to the host bus.
30:16	**Reserved.** Must be programmed to 0s when writing this register.
15:0	**Top of Host Bus System Memory Address.** Bits[15:0] of this register are compared to A[35:20]. The top of system memory is programmed in units of 1 Mbyte (i.e., 00001h=1 Mbyte, 00002h= 2 Mbytes, 00003=3 Mbytes, etc.).

2.4.12 PDM—PCI DECODE MODE

Address Offset: 48h
Default: 06h
Attribute: Read/Write

This register masks AD[31:16] for host I/O transactions. In addition, this register enables/disables ISA aliasing for I/O addresses in the range 100–3FFh. Note that the PB never forwards PCI I/O addresses above 64 Kbytes to the host bus.

Bits	Description
7:3	**Reserved.**
2	**I/O Address Mask Enable.** 1=Enable (default). 0=Disable. When enabled, the PB forces PCI AD[31:16] to zero for host bus to PCI I/O transactions. (The processor may assert A16 during I/O in real mode.) In all cases, the PB only decodes the lower 64 Kbytes of the host bus I/O address.
1	**I/O Aliasing Enable.** 1=Enable ISA expansion aliasing (default). 0=Disable. **Aliasing Algorithm (bit 1=1)** If A[9:8]=00, the address does not fall into an I/O alias range and A[15:4] are compared to the I/O space ranges defined by the IOSR1 and IOSR2 Registers (offsets 98–9Bh and A0–A3h, respectively). If A[9:8]≠00h, the address is in an alias range so A[15:10] are masked (the address is aliased for decoding purposes) before comparing the address to the I/O space range registers. Note that, when I/O aliasing is enabled (bit 1=1) and the I/O address mask enable feature is disabled (bit 2=0), the PB decoder aliases any bus I/O address above 64 Kbytes. In an 82454GX dual PB system, both PBs must have this bit set the same. Otherwise, both PBs may respond to host bus transactions targeting an aliased ISA expansion I/O address.
0	**Reserved.**

2.4.13 BDNUM—BRIDGE DEVICE NUMBER REGISTER

Address Offset: 49h
Default: 0001 1001h (82454KX)
0001 1001h (Compatibility 82454GX)
0001 1010h (Auxiliary 82454GX)
Attribute: Read Only

This register contains the bridge device number. For the 82454KX this value is hardwired to 11001.

For the 82454GX, this number is loaded from the IOGNT# and IOREQ# pins during power-up. Only the Compatibility PB (PBID = 01) responds to initial BIOS code fetches (defaults to this area enabled).

Bits	Description
7:5	**Reserved.**
4:2	**Fixed Value.** The upper three bits of the PB Bridge Device Number are always 110.
1:0	**82454KX: Fixed Value.** The lower two bits of the PB Bridge Device number are always 01 **82454GX: PB Identification (PBID).** The lower two bits of the Bridge Device Number encoding is defined as follows: **Bits[1:0] Function** 0 0 Reserved. 0 1 Used by the Compatibility PB. 1 0 Used by the Auxiliary PB in a two PB system. 1 1 Reserved.

2.4.14 PBNUM—PCI BUS NUMBER REGISTER

Address Offset: 4Ah
Default: 00h
Attribute: Read/Write

This register contains the bus number of the PCI bus that is immediately behind the PB. Note, this does not affect the Bus Number for the PB configuration registers. The PB configuration registers are always addressed as Bus Number 0.

For the 82454GX, both PBs default to Bus Number 0. The auxiliary PB must be changed before configuration accesses are made to the PCI bus.

Bits	Description
7:0	**Bus Number.** The PCI Bus Number of the PCI bus immediately behind the PB.

2.4.15 PSBNUM—SUBORDINATE BUS NUMBER REGISTER

Address Offset: 4Bh
Default: 00h
Attribute: Read/Write

This register contains the bus number of the last hierarchical bridge under the current bridge.

For the 82454GX, both PBs default to Subordinate Bus Number 0. The auxiliary PB must be changed before configuration accesses are made to the PCI bus.

Bits	Description
7:0	**Last Bus Number.** The bus number of the last hierarchical bridge under the current bridge.

2.4.16 PBC—PB CONFIGURATION REGISTER

Address Offset: 4Ch
Default: 39h (82454KX)
39h (Compatibility PB)
3Ah (Auxiliary PB)
Attribute: Read/Write

This register configures the PB for various operations.

Bits	Description
7	**Long Watchdog Timer Enable (LWTE).** 1=30 ms. 0=1.5 ms
6	**Lock Atomic Reads.** 1=Reads that cross a PCI Dword boundary are issued as locked reads. This bit must be enabled for systems containing a PCI to PCI bridge.
5	**Reserved.**
4	**Branch Trace Message Response Enable.** 1=PB responds to the Branch Trace Message host bus command. (default). Used by external development tools that need to capture the processor Branch Trace Message. 0=PB ignores this command. In a 450GX dual PB system, the auxiliary bridge ignores this bit.
3	**INIT on Shutdown Enable.** 1=PB asserts the INIT# signal when receiving a Shutdown command. (default). For the 82454GX in a dual PB system, this bit is only used in the Compatibility PB and is not used in the Auxiliary PB.
2	**Reserved.**

Bits	Description
1:0	**82454KX: Reserved.** Must be set to 01. **82454GX: Bridge Arbitration Mode.** These bits determine the arbitration mode the PB uses before taking ownership of the processor BPRI# signal. **Bits[1:0] Function** 00 No arbitration (single bridge system) 01 Arbitration mode (This PB provides the arbitration unit for an Auxiliary PB) 10 External arbiter mode. (This setting is for the Auxiliary bridge) 11 Reserved. Note that, in a single PB system where the internal arbiter is not needed, this field should be changed from its default value to 00.

2.4.17 DCC—DETURBO COUNTER REGISTER

Address Offset: 51h
Default: 80h
Attribute: Read/Write

In deturbo mode this register determines how many clocks BPRI# is asserted in order to throttle the host bus. The host bus pipeline is stalled at a rate determined by this register. The Deturbo Counter value is compared to an 8 bit counter running at the host bus clock (BCLK) divided by 8.

For the 82454GX in a dual PB system, this register is only available in the Compatibility PB and is not available in the Auxiliary PB.

Bits	Description
7:0	**Deturbo Count Value.** When the counter value is greater than this register value, BPRI# is asserted by the PB. BPRI# is negated when the count value is less than or equal to this register value. Smaller values in this register result in slower deturbo emulation speeds.

2.4.18 CRWC—CPU READ/WRITE CONTROL REGISTER

Address Offset: 53h
Default: 00h
Attribute: Read/Write

This register enables/disables processor writes to PCI to be posted in the PB.

Bits	Description
7:2	**Reserved.**
1	**Host-to-PCI Write (Outbound) Posting Enable.** 1=Enable. 0=Disable.
0	**Reserved.**

2.4.19 PRWC—PCI READ/WRITE CONTROL

Address Offset: 54–55h
Default: 00h
Attribute: Read/Write

The PRWC Register enables/disables read pre-fetching on the host bus. This register also enables/disables the assembly of back-to-back sequential host-to-PCI memory space cache line writes into PCI burst cycles and enables/disables PCI-to-host (inbound) write posting.

Bits	Description
15:10	**Reserved.**
9	**CPU Line Read Pre-Fetch for PCI Memory Read Commands Enable.** 1=Enable. 0=Disable. When enabled, PCI Memory Read commands cause a fetch of a CPU cache line plus a pre-fetch of three or more CPU cache lines (Pre-fetching does not cross 4 Kbyte address boundaries). This bit has no affect unless CPU Line Read Enable (bit 8) is also set.
8	**CPU Line Read for PCI Memory Read Commands Enable.** 1=Enable. 0=Disable. This bit is set to enable PCI Memory Read commands to fetch full CPU cache lines. When disabled, a PCI Memory Read command results in read partials on the host bus.
7	**Reserved.**
6	**CPU Line Read Multiple Pre-Fetch for PCI Memory Read Multiple Commands Enable.** 1=Enable. 0=Disable. When enabled, PCI Memory Read Multiple commands cause a fetch of a CPU cache line plus a pre-fetch of three or more CPU cache lines (Pre-fetching does not cross 4 Kbyte address boundaries). This bit has no affect unless CPU Line Read Multiple Enable (bit 5) is also set.
5	**CPU Line Read Multiple for PCI Memory Read Multiple Commands Enable.** 1=Enable. 0=Disable. When enabled, PCI Memory Read Multiple commands fetch full CPU cache lines. When disabled, a PCI Memory Read Multiple command results in read partials on the host bus.
4	**CPU Line Read Pre-Fetch for PCI Memory Read Line Commands Enable.** 1=Enable. 0=Disable. When enabled, PCI Memory Read Line commands cause a fetch of a CPU cache line plus a pre-fetch of three or more full CPU cache lines. Pre-fetching does not cross 4 Kbyte address boundaries. This bit has no affect unless CPU Line Read Enable (bit 3) is also set.
3	**CPU Line Read for PCI Memory Read Line Commands Enable.** 1=Enable. 0=Disable. When enabled, PCI Memory Read Line commands fetch full CPU cache lines. When disabled, a PCI Memory Read Line command results in read partials on the host bus.
2	**Reserved.**
1	**Burst Write Assembly Enable.** 1=Enable. 0=Disable. When enabled, back-to-back sequential CPU-to-PCI memory space cache line writes (USWC memory type) are converted to continuous PCI write bursts (write combining). This feature should only be enabled if the cache line writes are guaranteed not to cross component address boundaries.
0	**PCI-to-Host Bus Write (Inbound) Posting Enable.** 1=Enable. 0=Disable. (Caution: Do not enable if CPU to PCI locks split across component boundaries on the host bus.)

2.4.20 SMME—SMRAM ENABLE REGISTER

Address Offset: 57h
Default: 00h
Attribute: Read/Write

This register enables the address range programmed into the SMMR Register (B8–BBh) and permits System Management Mode (SMM) RAM to overlay memory space normally mapped to the PCI bus.

Bits	Description
7:4	**Reserved.**
3	**SMM RAM Normal Decode Range Override Enable.** 1=Enable. 0=Disable. When enabled,. the PB ignores the SMM RAM Gap Range for all transactions and uses normal decode ranges (i.e., the SMMR Register has no affect on PB address decoding). When disabled, accesses to the SMM Range made when SMMEM# is asserted are ignored by the PB, even if the address is within the positive decode range of a different range register in the PB. Accesses outside of the SMM Range are not affected by the state of the SMMEM# signal. SMM RAM may overlay memory space normally mapped to the PCI bus.
2:0	**Reserved.**

2.4.21 VBAE—VIDEO BUFFER AREA ENABLE REGISTER

Address Offset: 58h
Default: 02h (82454KX)
02h (Compatibility PB), 00h (Auxiliary PB)
Attribute: Read/Write

The VBAE Register selects where VGA Buffer requests are directed. This register defaults to Read/Write accesses directed to the PB (Compatibility PB only for the 82454GX). The MC and Auxiliary PB (82454GX only) default to this range disabled.

Bits	Description
7:2	**Reserved.**
1	**Video Buffer Area Enable (A0000–BFFFFh).** 1 = Host bus requests to the VGA Buffer range are forwarded to the PCI bus. 0=Host bus requests to the VGA Buffer range are ignored.
0	**Reserved.**

2.4.22 PAM[0:6]—PROGRAMMABLE ATTRIBUTE MAP REGISTER

Address Offset: PAM0 (59h)—PAM6 (5Fh)
Default: PAM0—30h; PAM[1:6]—33h (82454KX)
PAM0—30h (Compatibility PB), 00h (Auxiliary PB)
PAM[1:6]—33h (Compatibility PB), 00h (Auxiliary PB)
Read/Write

These seven registers select read only (RE=1, WE=0), write only (RE=0, WE=1), or read/write (RE=1, WE=1) access attributes for 14 memory regions between the 512 Kbyte and 1 Mbyte address range. The individual memory regions can also be disabled (RE=0, WE=0). Each register controls two regions; bits [7:4] control one region and bits [3:0] control the other region. Note that the default for the system BIOS region is read/write enabled. The default for all other regions is read/write disabled.

When a region is enabled, the corresponding host bus access is forwarded to PCI; the corresponding PCI access is ignored (not forwarded to the host bus). When a region is disabled, the corresponding host bus access is ignored (not forwarded to the PCI bus); the corresponding PCI access is forwarded to the host bus.

For the 450GX, the default for the system BIOS region in the Compatibility PB is read/write enabled. The default for the auxiliary PB is read/write disabled.

Note that the MC has corresponding PAM registers. Only one device (MC/PB) should have the same space enabled at one time to avoid access conflicts.

PAM	Attribute Bits				Memory Segment	Comments	Offset
Register	7,3	6,2	5,1	4,0			
PAM0[7:4]	Reserved		WE	RE	0F0000–0FFFFFh	BIOS	59h
PAM0[3:0]	Reserved		WE	RE	080000–09FFFFh	512–640 KB	59h
PAM1[7:4]	Reserved		WE	RE	0C4000–0C7FFFh	ISA Expansion	5Ah
PAM1[3:0]	Reserved		WE	RE	0C0000–0C3FFFh	ISA Expansion	5Ah
PAM2[7:4]	Reserved		WE	RE	0CC000–0CFFFFh	ISA Expansion	5Bh
PAM2[3:0]	Reserved		WE	RE	0C8000–0CBFFFh	ISA Expansion	5Bh
PAM3[7:4]	Reserved		WE	RE	0D4000–0D7FFFh	ISA Expansion	5Ch
PAM3[3:0]	Reserved		WE	RE	0D0000–0D3FFFh	ISA Expansion	5Ch
PAM4[7:4]	Reserved		WE	RE	0DC000–0DFFFFh	ISA Expansion	5Dh
PAM4[3:0]	Reserved		WE	RE	0D8000–0DBFFFh	ISA Expansion	5Dh
PAM5[7:4]	Reserved		WE	RE	0E4000–0E7FFFh	BIOS Extension	5Eh
PAM5[3:0]	Reserved		WE	RE	0E0000–0E3FFFh	BIOS Extension	5Eh
PAM6[7:4]	Reserved		WE	RE	0EC000–EFFFFh	BIOS Extension	5Fh
PAM6[3:0]	Reserved		WE	RE	0E8000–0EBFFFH	BIOS Extension	5Fh

2.4.23 ERRCMD—ERROR REPORTING COMMAND REGISTER

Address Offset: 70h
Default: 00h
Attribute: Read/Write

This register provides control for generating PCI SERR# and PERR# error signals. Note that for bits[7:4], SERR# must be enabled in the PCICMD Register. For bit 3, PCI parity error checking must be enabled in the PCICMD Register.

Bits	Description
7	**SERR# on Receiving Target Abort (PB is PCI bus master).** 1=Enable. 0=Disable.
6	**SERR# on Transmitted Data Parity Error (Detected via PERR#).** 1=Enable. 0=Disable.
5	**SERR# on Received Data Parity Error (Detected via PAR).** 1=Enable. 0=Disable. The PB is the master.
4	**SERR# on Address Parity Error Enable.** 1=Enable. 0=Disable.
3	**PERR# on Data Parity Error Enable.** 1=Enable. 0=Disable. The received data can be the result of a PB read or another PCI master write to the PB.
2:0	**Reserved.**

2.4.24 ERRSTS—ERROR REPORTING STATUS REGISTER

Address Offset: 71h
Default: 00h
Attribute: Read/Write Clear

This register reports certain PCI data and address parity errors and for detection of a CPU shutdown cycle. Software sets these bits to 0 by writing a 1 to them.

Bits	Description
7	**Reserved.**
6	**PCI Data Parity Error When Writing PCI Data (PERR# was asserted).** 1=Detected parity error.
5	**Data Parity Error When Reading PCI Data (PAR was incorrect).** 1=Detected parity error.
4	**Detected PCI Address Parity Error.** 1=PAR was incorrect when receiving address and C/BE[3:0]# from another PCI master.
3:1	**Reserved.**
0	**Shutdown Cycle from Host Bus Detected.** 1=Detected. The PB optionally asserts INIT# as per the setting in the PBC register (4Ch). For the 82454GX in a dual PB system, this bit is only used in the Compatibility PB and is not used in the Auxiliary PB.

2.4.25 MGR—MEMORY GAP RANGE REGISTER

Address Offset: 78–79h
Default: 00h
Attribute: Read/Write

This register is used with the MGUA Register (7A–7Bh) to define the memory gap range. When enabled, the corresponding host bus access is forwarded to PCI; the corresponding PCI access is ignored (not forwarded to the host bus). When disabled, the corresponding host bus access is ignored (not forwarded to the PCI bus); the corresponding PCI access is forwarded to the host bus.

Bits	Description
15	**Memory Space Gap Enable.** 1=Enable. 0=Disable.
14:10	**Memory Gap Size.** This field defines the memory gap size as follows: **Bits[14:10] Size Bits[14:10] Size** 00000 1 MB 11100 8 MB 00100 2 MB 11110 16 MB 01100 4 MB 11111 32 MB Note that all other combinations are reserved.
9:8	**Reserved.**
7:4	**Gap Starting Address.** These bits correspond to address bits A[23:20]. The remaining address bits are specified in the MGUA Register. When bits[7:4]=0000, this memory gap is disabled.
3:0	**Reserved.**

2.4.26 MGUA—MEMORY GAP UPPER ADDRESS REGISTER

Address Offset: 7A–7Bh
Default: 00h
Attribute: Read/Write

This register is used with the MG Register (78–79h) to define the memory gap range.

Bits	Description
15:12	**Reserved.** Should be set to 0.
11:0	**Memory Gap Range Upper Address.** These bits correspond to the memory space gap starting address bits A[35:24].

2.4.27 PFB—PCI FRAME BUFFER REGISTER

Address Offset: 7C–7Fh
Default: 00h
Attribute: Read/Write

This register defines a PCI Frame Buffer range. When enabled, the corresponding host bus access is forwarded to PCI; the corresponding PCI access is ignored (not forwarded to the host bus). When disabled, the corresponding host bus access is ignored (not forwarded to the PCI bus); the corresponding PCI access is forwarded to the host bus.

Bits	Description
31:20	**Frame Buffer Starting Address.** These bits correspond to address bits A[31:20] and select the starting address in 1 Mbyte increments. All 0s=Disable range (overrides bit 11).
19:13	**Reserved.**
12	**VGA Performance Mode Enable.** 1=Enable. 0=Disable. When enabled, the VGA memory range of A0000–BFFFFh uses the same features that are enabled for the frame buffer address range (i.e., disable locks).
11	**Frame Buffer Range Enable.** 1=Enable. 0=Disable. This bit does not affect the VGA memory range defined by the VBAE Register (offset 58h).
10	**Reserved.**
9	**Frame Buffer Lock Disable.** 1=Disable. 0=Enable. When locks are disabled, CPU locked requests (i.e., exchange instructions) do not result in a PCI exclusive access. Note that, locking is usually not necessary when accessing frame buffer data.
8	**Reserved.**
7	**Flush Inbound Data Buffer on Non-deferred Frame Buffer Reads.** 1=Do not Flush. 0=Flush.
6:5	**Reserved.**
4:0	**Frame Buffer Size.** This field defines the frame buffer size as follows: **Bits[14:10] Size** — **Bits[14:10] Size** 00000 1 MB — 00111 8 MB 00001 2 MB — 01111 16 MB 00011 4 MB — 11111 32 MB Note that all other combinations are reserved.

2.4.28 HMGSA—HIGH MEMORY GAP RANGE START ADDRESS REGISTER

Address Offset: 88–8Bh
Default: 00h
Attribute: Read/Write

This register is used with the HMGEA Register (8C–8F) to define the high memory gap. When enabled, the corresponding host bus access is forwarded to PCI; the corresponding PCI access is ignored (not forwarded to the host bus). When disabled, the corresponding host bus access is ignored (not forwarded to the PCI bus); the corresponding PCI access is forwarded to the host bus.

Bits	Description
31	**High Memory Gap Enable.** 1=Enable. 0=Disable.
30:16	**Reserved.** These bits should be set to 0.
15:0	**High Memory Gap Start Address.** These bits correspond to address bits A[35:20] and select the starting address in 1 Mbyte increments. All 0s=Disable range.

2.4.29 HMGEA—HIGH MEMORY GAP END ADDRESS REGISTER

Address Offset: 8C–8Fh
Default: 00h
Attribute: Read/Write

This register is used with the HMGSA Register (88–8Bh) to define the high memory gap.

Bits	Description
31:16	**Reserved.** These bits should be set to 0.
15:0	**High Memory Gap End Address.** These bits correspond to address bits A[35:20] and select the ending address in 1 MB increments.

2.4.30 IOSR1—I/O SPACE RANGE 1 REGISTER (82454GX ONLY)

Address Offset: 98–9Bh
Default: FFF0 0001h (Compatibility PB)
FFF0 0000h (Auxiliary PB)
Attribute: Read/Write

This register defines an I/O space range. A second I/O space range is defined by the IOSR2 Register. Except for the ranges defined by these two registers, the Compatibility PB forwards all host bus accesses to PCI (and ignores PCI bus accesses) and the Auxiliary PB ignores all host bus accesses (and forwards PCI bus accesses to the host bus).

Bits	Description
31:20	**I/O Space Range 1 End Address.** Bits [31:20] correspond to A[15:4]. Must be set to the same value in both bridges.
19:16	**Reserved.**
15:4	**I/O Space Range 1 Start Address.** Bits [15:4] correspond to A[15:4]. Must be set to the same value in both bridges.
3:1	**Reserved.**
0	**I/O Space Range 1 Enable.** 1=Forward host bus accesses in the range to PCI and ignore PCI bus accesses in the range. 0=Ignore host bus accesses in the range and forward PCI bus accesses in the range to the host bus. Compatibility bridge: 1=default. To open a gap in the compatibility bridge I/O space, this bit must set to 0. Auxiliary bridge: 0=default. To claim an I/O range in the auxiliary bridge, this bit must be set to 1.

2.4.31 PCIRSR—PCI RESET REGISTER

Address Offset: 9Ch
Default: 00h
Attribute: Read/Write

This register permits software to reset the PCI bus without also resetting the CPU bus. Note, the PCI bus is always reset when the host bus is reset through a hard or power-on reset.

Bits	Description
7:1	**Reserved.**
0	**Reset PCI Bus.** Setting this bit from 0 to 1 causes the PB to assert PCIRST# for at least one millisecond. Resetting the PCI bus could cause unwanted system signals to drive into the processor. Be sure to understand the state of any signals going from the PCI bus back to the processor during reset. See Section 3.7.5.

2.4.32 IOSR2—I/O SPACE RANGE 2 REGISTER (82454GX ONLY)

Address Offset: A0–A3h
Default: FFF0 0001h (compatibility PB)
FFF0 0000h (Auxiliary PB)
Attribute: Read/Write

This register defines an I/O space range. A second I/O space range is defined by the IOSR1 Register. Except for the ranges defined by these two registers, the Compatibility PB forwards all host bus accesses to PCI (and ignores PCI bus accesses) and the Auxiliary PB ignores all host bus accesses (and forwards PCI bus accesses to the host bus)

Bits	Description
31:20	**I/O Space Range 2 End Address.** Bits[31:20] correspond to A[15:4]. Must be set to the same value in both bridges.
19:16	**Reserved.**
15:4	**I/O Space Range 2 Start Address.** Bits[15:4] correspond to A[15:4]. Must be set to the same value in both bridges.
3:1	**Reserved.**
0	**I/O Range 2 Enable.** 1=Forward host bus accesses in the range to PCI and ignore PCI bus accesses in the range. 0=Ignore host bus accesses in the range and forward PCI bus accesses in the range to the host bus. Compatibility bridge: 1=default. To open a gap in the compatibility bridge I/O space, this bit must set to 0. Auxiliary bridge: 0=default. To claim an I/O range in the auxiliary bridge, this bit must be set to 1.

2.4.33 APICR—I/O APIC RANGE REGISTER

Address Offset: A4–A7h
Default: 00FE C001h (82454KX)
00FE C001h (Compatibility PB)
00FE C000h (Auxiliary PB)
Attribute: Read/Write

This range defines an I/O APIC range. There can be up to 16 APICs, with contiguous ascending unit IDs below a PB. One of the 16 APIC 4 Kbyte blocks must be reserved for all CPU Local APIC units. (Multiple CPU(s) may use the same Local APIC address since Local APIC transactions are not visible on the host bus.) The PB responds to I/O APIC address range (base + x000h) through (base + yFFFh) where x is the I/O APIC Starting Unit ID and y is the highest unit ID number.

Note that a 64 Kbyte range is allocated to APIC space. Local APIC transactions are not visible on the host bus, but still require UC MTRR attributes. The Local APIC base address register in each processor should be programmed to point to one of the 4 Kbyte blocks in the 64 Kbyte APIC range so that one MTRR may be used for Local and I/O APIC configuration ranges. The MC does not reclaim any 64 Kbyte memory gaps created for the APIC range.

If there is an I/O APIC behind more than one PB, each PB must use the same APIC base address and all 64KB of APIC range must be accounted for among the PBs.

Bits	Description
31:28	**Reserved.** Must be set to zero.
27:12	**I/O APIC Base Address.** Bits[27:12] correspond to A[35:20] and select the I/O APIC base address in 1 MB increments.
11:8	**I/O APIC Starting Unit ID.** This field contains the lowest unit ID (0–Fh) of any APICs located below this bridge.
7:4	**Highest Unit ID Number.** This field contains the highest unit ID (0–Fh) of any APICs located below this bridge.
3:1	**Reserved.**
0	**I/O APIC Range Enable.** 1=Enable. 0=Disable.

2.4.34 CONFVR—CONFIGURATION VALUES DRIVEN ON RESET REGISTER

Address Offset: B0–B1h
Default: 00h
Attribute: Read/Write

During a programmed hard reset (via the PB's TRC Register), this register provides the processors and host bus agents with certain configuration details that have been programmed into the PB (Compatibility PB only for the 82454GX).

During a power-on reset, this register is set to its default values and these values are driven on the appropriate host bus signals. After initialization, software programs this register. During a programmed hard reset this register retains its programmed values and these values are driven on the host bus when the PB asserts RESET#.

For the 82454GX in a dual PB system, this register is only available in the Compatibility PB and is not available in the Auxillary PB.

Bits	Description
15:13	**Reserved.**
12:11	**APIC Cluster ID.** Software programs this field with the APIC cluster ID. The value in these bits are driven to the processors on A[12:11]#.
10	**82454KX: Reserved.** **82454GX: BINIT# Input Enable.** 1=Enable. 0=Disable. The value in this bit is driven on A10#. All host bus agents enable BINIT# if this bit is 1. See EXERRCMD Register for additional signal details.
9	**BERR# Input Enable.** 1=Enable. 0=Disable. This bit value is driven on A9#. All host bus agents enable BERR# reporting if this bit is 1. See the EXERRCMD Register for additional signal details.
8	**82454KX: Reserved.** **82454GX: AERR# Input Enable.** 1=Enable. 0=Disable. Used to enable the reporting of Address parity errors. The value in this bit is driven on A8#. All host bus agents cancel erroneous requests if this bit is 1. Expected use is to enable this bit and then map AERR# to NMI in the EXERRCMD Register.

Bits	Description
7	**In-Order Queue Depth 1 Select.** 1=Depth of 8. 0=Depth of 1. The value in this bit is driven on A7#. Pentium Pro processors use an in-order queue depth of 8 if this bit is 1.
6	**1M Power-on Reset Vector Select.** 1=1 Mbyte. 0=4 Gbyte. The value in this bit is driven on A6#. Pentium Pro processors use the 1 Mbyte reset vector if this bit is 1.
5	**FRC Mode Enable.** 1=Enable. 0=Disable. The value in this bit is driven on A5#. Pentium Pro processors enter FRC enabled mode if this bit is 1.
4:0	**Reserved.**

2.4.35 CSCONFV—CAPTURED SYSTEM CONFIGURATION VALUES REGISTER

Address Offset: B4–B5h
Default: XXh (X=Captured During hard reset)
Attribute: Read Only

The register reports how the system is set up for certain functions. The values in this register are captured on the rising edge of RESET#.

Bits	Description
15:13	**Reserved.**
12:11	**APIC Cluster ID.** The PB captures this value from A[12:11]#.
10	**BINIT# Input Enable.** 1=Enable. 0=Disable. The PB captures this value from A10#. For the 82454GX, see the EXERRCMD Register (offset C0–C3h) for additional details. **Caution:** Programming the corresponding output enable bit in the EXERRCMD Register must be consistent with the value in this bit captured from the host bus. Otherwise, incorrect system operations will result.
9	**BERR# Input Enable.** 1=Enable. 0=Disable. The PB captures this value from A9#. For the 82454GX, see the EXERRCMD Register (offset C0–C3h) for additional details.
8	**AERR# Input Enable.** 1=Enable. 0=Disable. The PB captures this value from A8#. For the 82454GX, see the EXERRCMD Register (offset C0–C3h) for additional details. **Caution:** Programming the corresponding output enable bit in the EXERRCMD Register must be consistent with the value in this bit captured from the host bus. Otherwise, incorrect system operations will result.
7	**In-Order Queue Depth 1 Select.** 1=Depth of 8. 0=Depth of 1. The PB captures this value from A7#. See EXERRCMD Register for additional signal details.
6	**1M Power-on Reset Vector.** 1=1 Mbyte. 0=4 Gbyte. The PB captures this value from A6#.
5	**FRC Mode Enable.** 1=Enable. 0=Disable. The PB captures this value from A5#.
4:0	**Reserved.**

2.4.36 SMMR—SMRAM RANGE REGISTER

Address Offset: B8–BBh
Default: 0000 0005h
Attribute: Read/Write

The range programmed into this register is required when SMRAM addresses overlap addresses normally mapped to the PCI bus.

Bits	Description
31:28	**SMM Range Size.** The size (in 64 Kbyte increments) is selected as follows: **Bits[31:28]** **Size** **Bits[31:28]** **Size** 0000 64 Kbyte 0011 256 Kbytes 0001 128 Kbytes ... ... 0010 192 Kbytes 1111 1 Mbytes
27:16	**Reserved.** Must be set to zero.
15:0	**SMM Range Start Address.** Bits[15:0] correspond to A[31:17] and select the starting address in 64 Kbyte increments beginning at 64 Kbytes. The default starting address is A0000h and ranges to BFFFFh.

2.4.37 HBIOSR—HIGH BIOS RANGE REGISTER

Address Offset: BCh
Default: 01h (82454KX)
01h (Compatibility PB)
00h (Auxiliary PB)
Attribute: Read/Write

This register enables/disables the 0–512 Kbyte and the high BIOS ranges. When enabled, the PB forwards host bus memory accesses in this range to PCI and ignores PCI memory accesses to the same range. When disabled, the PB ignores host bus memory accesses to this range and forwards PCI memory accesses to the same range to the host bus.

Bits	Description
7:5	**Reserved.**
4	**0–512 Kbyte Range Enable.** 1=Enable. 0=Disable.
3:1	**Reserved.**
0	**2 Mbyte High BIOS Range Enable (00 FFE0 0000–00 FFFF FFFFh).** 1=Enable. 0=Disable.

2.4.38 EXERRCMD—PB EXTENDED ERROR REPORTING COMMAND REGISTER

Address Offset: C0–C3h
Default: 0000 0010h
Attribute: Read/Write

This register enables/disables the reporting of certain error conditions.

Bits	Description
31:14	**Reserved.**
13	**Report Errors using Hard Fail/SERR# Enable.** 1=Enable (Hard Fail mechanism). 0=Disable (PB uses SERR#). These errors are reported when the PB is servicing a host bus request and detects an error. Note that the PB does not report PERR# through Hard Fail.
12	**Report PCI Master Abort Errors.** 1=Enable. 0=Disable. When enabled, the PB normally returns all 1's to CPU read transactions that receive a master abort time-out. Note that bit 13 in the PCISTS Register is always set for master abort time-outs. The error reporting mechanism (Hard Fail or SERR#) is determined by bit 13 of this register.
11	**82454KX: Reserved.** **82454GX: Report Uncorrectable Host Data Bus ECC Errors.** 1=Report by BERR# signal (bit 2 must be set to 1). 0=Disable
10	**82454KX: Reserved.** **82454GX: Single-bit ECC Error Correcting of Host Data Bus Enable.** 1=Enable. 0=Disable.
9	**Report Host Bus Time-out Errors Enable.** 1=Enable. 0=Disable. When enabled, the error reporting mechanism (Hard Fail or SERR#) is determined by bit 13 of this register. Note that the PB normally returns all 1's to CPU read transactions that receive a time-out. Bit 9 in the EXERRSTS Register is set, regardless of whether the error is reported.
8	**Host Bus Time-out Enable.** 1=The PB responds to unclaimed host bus transactions when the Bus watchdog timer expires. The time-out value can be programmed to either 1.5 ms or 30 ms. 0=Disable watchdog timer. For the 82454GX in a dual PB system, this bit only has affect in the Compatibility PB and has no affect in the Auxiliary PB
7	**AERR# to NMI Enable.** 1=Enable. 0=Disable. When enabled (and bit 8=1 in the Captured System Configuration Values Register and SERR# is enabled in the PCICMD Register), the PB (Compatibility PB in an 82454GX dual PB system) asserts the SERR# signal when detecting AERR# signal asserted. Note that, depending on the system architecture, the SERR# signal can result in the generation of an NMI. The NMI signal is not part of the PB and is typically provided by a PCI-to-ISA or PCI-to-EISA bridge.
6	**BERR# to NMI Enable.** 1=Enable. 0=Disable. When enabled (and bit 9=1 in the CSCONFV Register, offset B4–B5h), the PB (Compatibility PB in an 82454GX dual PB system) asserts the SERR# signal (which can result in an NMI) when BERR# is asserted. Note that the NMI signal is not part of the PB. NMI is typically provided by a PCI-to-ISA or PCI-to-EISA bridge.
5	**Reserved.**
4	**BERR# to BINIT# Enable.** 1=Enable. 0=Disable. When enabled, the PB asserts BINIT# when BERR# is asserted.
3	**Assert BINIT# on Detection of Host Bus Protocol Violations Enable.** 1=Enable. 0=Disable. **Caution:** Programming this bit must be consistent with the value in the corresponding bit of the CSCONFV Register captured from the host bus. Otherwise, incorrect system operations will result.
2	**Assert BERR# on Bus Errors Enable.** 1=Enable. 0=Disable.
1	**Reserved.** Planned use is AERR# to BERR# Enable. 1=Enable. 0=Disable.
0	**Assert AERR# on Request Phase Signal Parity Errors Enable.** 1=Enable. 0=Disable. **Caution:** Programming this bit must be consistent with the value in the corresponding bit of the CSCONFV Register captured from the host bus. Otherwise, incorrect system operations will result.

2.4.39 EXERRSTS—PB EXTENDED ERROR REPORTING STATUS

Address Offset: C4–C7h
Default: 0000 0000h
Attribute: Read/Write Clear

This register reports certain host bus error conditions. Software sets each error condition to 0 by writing 1 to it.

Bits	Description
31:21	**Reserved.**
20	**Received Hard Failure Response.** 1=Host bus hard failure response or a second AERR# assertion was received for a PB initiated transaction. (Exception—The PB does not log a Hard Fail Response by another host bus agent to a PB posted write. The host bus agent error registers must be checked to determine the cause of the error.)
19	**Host Bus Address Parity (AP#) Error Detected.** This bit is set when an AP# parity error is reported by the 82454.
18	**Host Bus Request Parity (RP#) Error Detected.** This bit is set when an RP# parity error is reported by the 82454.
17	**82454KX: Reserved.** **82454GX: Host Bus Correctable Error Detected.** 1=Logs a single-bit ECC error detected on the data bus. No error is reported when a host bus correctable error is detected.
16	**Host Bus Protocol Violation Detected.** This bit is set when a protocol violation, including RS parity errors, is reported by the 82454. Both OPBs check for protocol violations in dual-bridge 82454GX systems.
15:12	**Reserved.**
11	**82454KX: Reserved.** **82454GX: Host Bus Uncorrectable Error Detected.** 1=Logs a multiple-bit ECC error detected on the data bus. Note that this bit is set independent of whether error reporting is enabled via bit 11 of the EXERRCMD Register. If BERR# is enabled in the CSCONFV and EXERRCMD Registers, this error is reported by generating a BERR#.
10	**Reserved.**
9	**Time-out on Host Bus Detected.** 1=The PB detected a time-out (no response phase within the time-out value programmed into the PBC resister (4Ch) and In-Order Queue not empty) on the host bus. This bit is set, regardless of whether the event is reported. This bit is not used in the auxiliary bridge of a dual PB system.
8:4	**Reserved.**
3	**BINIT# on Host Bus Detected.** 1=BINIT# was detected on the host bus.
2	**BERR# on Host Bus Detected.** 1=BERR# was detected on the host bus. The PB (Compatibility PB in an 82454GX dual PB system) generates a SERR# if enabled via bit 7 of the EXERRCMD Register, or the PB generates BINIT if BERR# to BINIT# is enabled via bit 4 of the EXERRCMD Register.
1	**Reserved.**
0	**AERR# on Host Bus Detected.** 1=AERR# was detected on the host bus. The PB (Compatibility PB in an 82454GX dual PB system) generates SERR#, if enabled via bit 7 of the EXERRCMD Register. No recovery is possible as there is a chance of data corruption.

2.4.40 PBRTMR—PB RETRY TIMERS

Address Offset: C8h
Default: 0000 0003h
Attribute: Read/Write

This register configures the host bus retry counter operation and the PCI retry counter operation.

Bits	Description
31:16	**Host Retry Counter Value.** This field is programmed with the retry count value. The count is in host bus clocks. 0000h disables the counter. The PB re-enables inbound posting after an outbound read transaction is retried if the retried agent does not return before this count expires.
15:5	**Reserved.**
4:2	**PCI Retry Count Value.** Posting is re-enabled when this count expires. Note that this count should match the count in the external PCI arbiter, if the arbiter has a retry masking counter. **Bits[4:2] PCI Clocks Bits[4:2] PCI Clocks** 000 0 011 64 001 16 100 128 010 32 101–111 Reserved
1	**Re-Enable Posting After an Inbound Transaction is Retried Enable.** 1=Enable posting when count expires. 0=Disable (Ignore PCI retry counter). This bit may be set along with bit 0.
0	**Re-Enable Posting After an Inbound Transaction is Retried.** 1=Re-enable posting as soon as PCI bus parking is observed. 0=Ignore state of PCI bus parking.

3.0 PB FUNCTIONAL DESCRIPTION

This section describes the PB functions and hardware interfaces including the I/O and Memory Map, Host bus, PCI bus, and Dual-bridge Architectures. Data Integrity and Error Handling are covered. Clock, Reset, and PB configuration are also covered.

3.1 Memory and I/O Map

The 82454KX/GX PB provides the interface between the host bus and the PCI bus. Memory transactions can be sent from the PCI bus to the host bus and from the host bus to the PCI bus. Gaps and positive decode ranges can be programmed via the configuration registers. For the 82454KX, I/O transactions can be sent from the host bus to the PCI bus. However, I/O transactions can not be sent from the PCI bus to the host bus.

For the 82454GX, both memory and I/O transactions can be sent from the PCI bus to the host bus and from the host bus to the PCI bus. Memory and I/O gaps and positive decode ranges can be programmed via the configuration registers.

If an access is enabled to be forwarded from the host bus to the PCI bus, the corresponding access on the PCI bus is ignored (not forwarded to the host bus). Conversely, if an access is enabled to be forwarded from the PCI bus to the host bus, the corresponding access on the host bus is ignored (not forwarded to the PCI bus).

The PB and MC perform a positive address decode of each host transaction and one default device handles the unclaimed transactions. In a standard PC system, unclaimed transactions are sent to the ISA bus. Thus, the PB (Compatibility PB in an 82454GX dual PB system) is the default responder on the host bus.

3.1.1 MEMORY ADDRESS MAP

The Pentium Pro processor memory address space is 64 Gigabytes (36-bit addressing). The PB does not support transactions of address size larger than 36-bits whether directed to the PB or not. The PB registers that control the memory space access are:

- **Programmable Attribute Map (PAM[6:0]) Registers.** These registers provide Read Only, Write Only, and Read/Write Disable for fixed memory regions in the PC compatibility area.
- **Video Buffer Area Enable (VBA) Register.** This register enables the A0000–BFFFF fixed region.
- **Top System Memory (TSM) Register.** This register permits the PB (Compatibility PB in an 82454GX dual PB system) to claim memory transactions above the top of main memory (top of memory to 64 Gbytes) and forward these transactions to the PCI Bus.
- **Memory Gap Range Registers (MGR and MGUA Registers).** The Memory Gap Range can start on any 1 Mbyte boundary from 1 Mbyte to 64 Gbytes and can be 1, 2, 4, 8,16, or 32 Mbytes.
- **High Memory Gap Range Registers (HMGSA and HMGEA Registers).** The High Memory Gap can start on any 1 Mbyte boundary from 1 Mbyte to 64 Gbytes.
- **High BIOS (HBIOS) Register.** The 64 KByte region from F0000–FFFFFh is treated as a single block and is normally read/write disabled in the MC(s) and Read/Write enabled in the PB.
  After power-on reset, this region is read/write enabled in the Compatibility PB and read/write disabled in the Auxiliary PB. Thus, the Compatibility PB responds to fetches during system initialization.

- **I/O APIC Range (APICR) Register.** This register provides an I/O APIC configuration space. There is no I/O APIC in the PB or the MC. Note that, the address range between the APIC configuration space and the High BIOS range (FED00000h–FFDFFFFFh) is always mapped to local memory unless: 1) The range is above top of physical memory or 2) The High BIOS and APIC ranges are disabled in the PB and the range falls within a memory gap range.
- **PCI Frame Buffer (PFB) Register.** The PCI Frame Buffer range can start on any 1 MByte boundary from 1–4 Gbytes and can be 1, 2, 4, 8, 16,or 32 Mbytes.
- **SMM Range (SMMR) Register along with the SMM Enable (SMME) Register (only when SMMEM# is asserted).** A Pentium Pro processor asserts SMMEM# in its Request Phase if it is operating in System Management Mode. The default SMRAM area is an address range that is normally mapped through the PB to the PC compatible video graphics adapter. The PB ignores accesses to this overlaid address range when the SMMR Register is enabled and SMMEM# is asserted during host bus transactions.

SMM Support. The PB supports System Management Mode by allowing the SMRAM region in the MC to overlay addresses that are normally mapped to the PCI bus. For cases where 64 Kbytes is insufficient for a given application, SMRAM can be relocated by the SMMR Register to a different start address set in 64 Kbyte increments and a maximum range of 1 Mbyte. The SMMR Register should also be used if the Top Of Memory Register is enabled and SMRAM is placed above normal memory. The SMMR Register is programmed in this instance to ignore the SMRAM range during SMMEM# accesses, but claim this range for normal accesses (SMRAM Range overrides Top of Memory).

When the processor receives an SMI#, it invokes an SMI Acknowledge Transaction before entering the SMI handler routine. The Compatibility PB generates the response phase for an SMI Acknowledge transaction and also asserts the SMIACT# signal, if SMMEM# is asserted. Once asserted SMIACT# remains asserted until an SMI Acknowledge transaction occurs with SMMEM# negated. See the Host Bus Interface section for additional information on SMM mode.

Memory Mapped I/O. The PB allows memory addresses to be mapped to the host bus or to a PCI bus below the PB. Memory mapped I/O devices can be located anywhere in the PB's 64 Gbyte address space. The Frame Buffer Range allows the PB to decode memory mapped I/O space extending up to 4 Gbytes. The Memory Space Gap and High Memory Gap Registers allow the PB to decode two address ranges extending up to 64 Gbytes.

Host Transactions to Memory Space. If a memory space address is in one of the above ranges, and that range is enabled, the PB claims the transaction and forwards it to the PCI bus. Accesses that are not in one of the enabled ranges and below the top of main memory are assumed to be accesses to main memory and are not claimed by the PB. The PB (Compatibility PB in an 82454GX dual PB system) is responsible for any unclaimed transactions on the host bus. Therefore, any memory space access that is above the top of main memory is claimed by this PB and forwarded to its PCI bus, if enabled in the TSM Register. Otherwise, transactions that are not mapped to any host bus device will time-out. Transactions that time-out on the host bus are handled by the PB (Compatibility PB in an 82454GX dual PB system) to remove them from the In-Order Queue. These transactions are not forwarded to PCI.

PCI Transactions to Memory Space. All PCI memory space accesses below the top of main memory (as programmed in the TSM Register) are forwarded to the host bus, unless they are specifically directed to PCI by one of the memory space access registers listed at the beginning of this section.

In a dual PB system, the Compatibility and Auxiliary PB default to forwarding all PCI memory space accesses above the top of memory to the host bus.

3.1.2 I/O ADDRESS MAP

The Pentium Pro processor I/O address space is 64 Kbytes. For the 82454KX, the PB maps all host bus I/O accesses to the PCI bus, except for the CONFADD, CONFDATA, and TRC Register locations.

For the 82454GX, the Compatibility PB maps all host bus I/O accesses to the PCI bus, except for I/O address ranges programmed into the IOSR[2:1] Registers (and the CONFADD, CONFDATA, and TRC Register locations). In a dual PB system, the Auxiliary PB ignores all host bus I/O accesses (except for the CONFADD and CONFDATA Register locations), unless forwarding is programmed into the IOSR[2:1] Registers.

The PB registers that control the I/O space accesses are:

- **CONFADD, CONFDATA, and TRC Registers.** These three PB registers are located in the processors I/O address space. See the Register Description section for details.
- **I/O Space Range Registers (IOSR[2:1]).** Two I/O Space Range Registers (IOSR1 and IOSR2) permit the PB to forward transactions targeting that range to the PCI bus.
- **PCI Decode Mode (PDM) Register.** The PB optionally supports ISA expansion aliasing. When ISA expansion aliasing is enabled (via the PDM Register), the ranges designated as I/O Expansion are internally aliased to the 100–3FFh range before the I/O Space Range registers are checked.

CPU Transactions to I/O Space. For the 82454KX, the PB claims all host bus I/O accesses and forwards the accesses to the PCI bus, except for the CONFADD, CONFDATA, and TRC Register locations. Accesses to CONFADD (OCF8) must be Dword aligned. I/O Transactions targeting 0CF8h are treated as normal I/O transactions when they are not Dword aligned. Accesses to CONFDATA (0CFCh) are treated as normal I/O transactions when the Configuration Space Enable bit of the CONFADD Register is not set.

If an I/O space address is in either of the I/O ranges, and that range is enabled, the PB claims the transaction and forward it to the PCI bus. For the 82454GX in a dual PB system, the Compatibility PB is the default I/O response agent responsible for claiming all I/O transactions on the host system bus. Therefore, any I/O address range that is mapped to an Auxiliary PB must be disabled by an I/O range register in the Compatibility PB.

When using the I/O Space Range Registers, the CONFADD, CONFDATA, and TRC Registers (0CF8h, 0CF9h, and 0CFCh) are treated differently than other I/O space addresses. I/O Transactions targeting 0CF8h are treated as normal I/O transactions by the IOSR[2:1] Registers when they are not Dword aligned transactions. Accesses to 0CFCh are treated as normal I/O transactions by the IOSR[2:1] Registers when the Configuration Space Enable bit of the CONFADD Register is not set. Byte address OCF9h is recognized only by the Compatibility PB, and is never affected by the IOSR[2:1] Registers.

ISA Expansion Board Aliasing. In PCs the I/O address range 100–3FFh is reserved for ISA Expansion boards. Many ISA Expansion boards only decode address bits [9:0] which results in aliases of the decode range of these boards. The PB provides a method to route the alias of an address in the 100–3FFh range through the appropriate PB when I/O space has been split between dual PBs. See PCI Decode Mode (PDM) Register (offset 48h).

PCI Transactions to I/O Space. For the 82454KX, I/O space accesses are never forwarded to the host bus. For the Auxiliary PB in a dual PB system, all PCI I/O space accesses are forwarded to the host bus unless they are specifically directed to PCI by one of the I/O Range registers. For the Compatibility PB, I/O space accesses are never forwarded to the host bus, unless specifically directed to the host bus by one of the I/O Range registers. The PB never forwards PCI I/O accesses greater than 64 Kbytes to the host bus.

3.2 Host Bus Interface

The Pentium Pro processor bus provides an efficient, reliable interconnect between multiple Pentium Pro processors and the PB and MC. The bus provides 36 bits of address, 64 bits of data, protection signals needed to support data integrity, and the control signals to maintain a coherent shared memory in the presence of multiple caches.

The Pentium Pro processor bus achieves high bus efficiency by providing support for multiple, pipelined transactions and deferred replies. A single Pentium Pro processor may have up to four transactions outstanding at the same time, and can be configured to support a total of either one or eight transactions active on the Pentium Pro processor bus at any one time. The PB supports up to eight active transactions on the host bus (In-Order Queue depth of 8). During the host bus reset and configuration, all host bus devices are configured to support either one or eight transactions in their In-Order Queue.

The number of transactions that can target a particular bus client is configured separately from the total number of transactions allowed on the bus. The PB accepts up to four transactions into the Outbound Request Queue that target its associated PCI bus.

The PB provides four 32-byte buffers for outbound data (host-to-PCI writes or PCI reads from the host bus), and four 32-byte buffers for inbound data (PCI-to-host writes or CPU reads from PCI).

As a host bus master, the PB does not support deferred responses. The EXF1# extended function signal (Defer Enable) will never be asserted for a host transaction initiated by the PB.

The host bus supports ECC over the data bus, and parity protection over the address, request, and response lines. The PB generates and checks ECC over the data lines (82454GX only), and generates and checks parity over the address and request/response signal lines (both 82454KX/GX). Note, ECC generation and checking on the data lines and parity generation and checking on the request/response lines can be enabled or disabled during system configuration.

NOTE:

1. The PB is a non-caching agent and does not participate in the Snoop phase. The Write Back (WB) memory types can not be mapped through the PB (snoop write-back data is ignored by the PB for implicit writebacks initiated by other agents). No WB memory types should be mapped to PCI. For PCI Frame Buffers, the Write Combining (WC) memory type is recommended.
2. The PB is a non-caching agent; however all Pentium Pro processor commands are defined for the PB. Therefore Read Invalidate transactions are treated as reads by the PB. Write Invalidate cycles are treated as writes of length 0 by the PB. Write-backs initiated by other agents are ignored by the PB.
3. When the processor receives an SMI#, it invokes an SMI Acknowledge Transaction before entering the SMI handler routine. The Compatibility PB generates the response phase for an SMI Acknowledge transaction and also asserts the SMIACT# signal, if SMMEM# is asserted. Once asserted SMIACT# remains asserted until an SMI Acknowledge transaction occurs with SMMEM# negated. The other System Management Mode transaction that is supported on the processor interface is Stop Clock Acknowledge. The Stop Clock Acknowledge is an indication from the processor to the system that the processor is powering down the internal caches to save power. For Stop Clock Acknowledge Transactions, the Compatibility PB is the responding agent and generates a Stop Clock Grant special cycle on its PCI bus.
4. If the SMRAM space is set up as writeback memory, A WBINVD instruction must be executed in the SMM handler immediately before execution of the RSM instruction that exits SMM mode.

3.3 PCI Bus Interface

The PB has a standard master/slave PCI bus interface. All legal PCI (PCI specification 2.0) bus transactions are supported. PCI cycle termination and error logging/reporting are discussed in the Data Integrity and Error Handling section. The PCI arbitration unit is not implemented in the PB.

PCI Locks. Systems which support PCI initiate locks (either inbound locks or peer-to-peer) must configure the arbiter for full bus locks rather than resource locks. The PB will not recognize resource locks made by peer-to-peer accesses. When a PCI master asserts LOCK# while targeting the PB, the locked PCI transactions are converted to locked host bus transactions. The host bus lock continues as long as the PCI master asserts LOCK# for exclusive access to the PB. The host bus lock is assisted by the bridge continuing to assert BPRI# as long as the PCI bus is asserting resource lock to the bridge. Additional locked CPU transactions are issued if the PCI master continues to burst.

In systems in which target abort reporting is disabled, the write portion of a lock will be committed even when the read portion is aborted.

NOTE:

Locks that cross cache line boundaries initiated on the PCI bus will not generate a SPLCK# signal on the host bus. This should be understood by all host bus agents. Neither the PB nor the MC require SPLCK# assertion.

Host Bus Locks. Any transactions that target the bridge during a host bus lock are converted into a similar PCI lock transaction. The lock on the PCI bus is held until the host bus lock is released. Locks over the Frame Buffer region can be disabled through a mode bit in the PCI Frame Buffer Range Register.

NOTE:

Locks that split across PCI host bus device boundaries (originate to one device and complete to another) are only supported for shadowed memory, and then only behind the compatibility PB. Shadowed memory is memory mapped for read only or write only in the MC and the opposite way in the PB. An update may be required for older non-PCI 2.0 compliant device drivers to comply with this. Since the revision 2.0 of the PCI specification does not allow locks to cross device boundaries, this will not be an issue with new device drivers.

Indivisible Operations. CPU initiated read operations that cross a Dword boundary (e.g., Read 8 Bytes, Read 16 Bytes, etc.) are indivisible operations on the host bus. However, since the PCI protocol allows a target device to disconnect at any point in a transfer sequence, these operations must be locked indivisible on the PCI bus. The PB optionally locks all CPU initiated reads that cross a Dword boundary. This mode is enabled by setting the Lock Atomic Reads in the PB Configuration Register. CPU initiated Write operations (e.g., Write 8 Bytes, Write 16 Bytes, etc.) are indivisible operations on the host bus. However, these accesses can not be made indivisible on the PCI bus because the PCI Specification states that the first transaction of a locked operation must be a read. Therefore, software must not rely upon the atomicity of CPU initiated write transactions greater then 32 bits once they are translated to the PCI bus.

Software Generated Special Cycles. This optional feature is not supported by the 450KX/GX PCIset.

3.4 Data Integrity and Error Handling

Several data integrity features are included in the PB. These include ECC on the host data bus (450GX only), parity on the host address, parity on the CPU Request/Response signals, and parity on the PCI bus. Error logging (setting a status bit) and reporting (generating an error signal) are controlled by the PCICMD Register (04–05h), PCISTS Register (06–07h), ERRCMD Register (70h), ERRSTS Register (71h), EXERRCMD Register (C0–C3h), and EXERRSTS Register (C4–C7h).

3.4.1 HOST BUS ERRORS

The PB detects errors on the host bus by checking the ECC provided with data (450GX only) and the parity provided with control signals. In turn, the PB will generate ECC with data (450GX only) and parity with control signals so that bus errors can be detected by receiving clients.

Request Parity (RP#) is the parity signal for ADS# and REQ[4:0]# and is computed as even parity. AP[1:0]# are the parity signals for A[35:3]# and are computed as even parity (AP1# is for A[35:24]# and AP0# is for A[23:3]#. RSP# is the parity signal for RS[2:0]# and is computed as even parity. In addition, certain host bus protocol violations are detected by the PB.

On the 450GX, ECC error checking is used for the data bus. The ECC check bits are provided by the DEP[7:0]# signals.

The PB (Compatibility PB in a 82454GX dual PB system) is responsible for responding to any unclaimed transactions on the host bus. The PB uses a watchdog timer to monitor host response phases. The timer is started at the end of a response phase if the In-Order Queue is not empty. If the timer expires before the next host response phase, a host bus time-out has occurred. The time-out window for such an event is programmable to 1.5 or 30 milliseconds via the PBC Register (4Ch). This allows for several host to PCI transactions, which may be blocking the progress of the In-Order Queue, to undergo multiple retries. When a host bus time-out occurs, the PB (Compatibility PB in an 450GX dual PB system) claims the transaction by returning all 1's to a read transaction or "pretending" to accept data for a write transaction. This event is logged in the EXERRSTS Register and can generate a hard fail or SERR#, if enabled in the EXERRCMD Register.

AERR#. If AERR# observation is enabled, then AERR# to NMI should be enabled in the EXERRCMD register (C0-C3h). This allows software to accept an NMI to log or recover from the event.

BINIT#. A BINIT# on the Host bus creates a PCIRST# and resets the 450KX/GX PCIset host bus state machines. This allows for logging or recovery from catastrophic bus errors.

3.4.2 PCI BUS ERRORS

The PB always detects address parity errors when it is not the PCI master, even if it is not the selected target. The PB detects data parity errors if it is either the master or the target of a transaction, and optionally reports them to the system. Address parity errors are reported using the SERR# signal. Data parity errors are reported using the PERR# signal.

3.4.2.1 PB Master Operation on PCI

Master Abort. When the PB performs a master abort, if the command was not a Special Cycle, the event is logged by setting the Received Master Abort bit (bit 13) in the PCISTS Register. An interrupt can be generated on this event. Special Cycle commands, which are broadcast to all PCI targets, are always terminated with master abort. Therefore, master aborts during Special Cycle commands are not considered errors, and are never logged or reported.

Target Disconnect and Target Retry. Target disconnects and target retries are not errors, and are not logged or reported.

Target Abort. The PB logs a target abort by setting the Received Target Abort bit (bit 12) in the PCISTS Register. If the SERR# enable bit (bit 8) of the PCICMD Register is set, and the SERR# on Receiving Target Abort bit (bit 7) of the ERRCMD Register is set, this event is reported by asserting SERR#. When the PB asserts SERR#, the Signalled System Error bit (bit 14) in the PCISTS Register is set. Optionally, the PB reports a hard failure response to the host bus transaction (PB EXERRST Register). Note that this is not possible for posted writes because the response phase has already occurred.

Data Parity Errors. As a PCI bus master, the PB checks the data parity provided during read data cycles and monitors PERR# during write data cycles. The errors are logged by setting the appropriate status bits. If a parity error is detected, the Detected Parity Error bit (bit 15) in the PCISTS Register is set. To distinguish between read data parity errors and write data parity errors, the appropriate bit (bit 6 for writes, bit 5 for reads) is set in the ERRSTS Register. Errors are reported via the SERR# and PERR# signals. The conditions causing the assertion of SERR# due to data parity errors are summarized in Figure 5. The conditions causing the assertion of PERR# and the Detected Parity Error Status bit are summarized in Figure 3 and Figure 4. Note that for read data parity errors, the PB returns the corrupted data (with good parity/ECC) as the CPU read response data. For write data parity errors, the corrupted data has already been delivered to the target; it is not retried by the PB.

3.4.2.2 PB Target Operation on PCI

Target Disconnect. PB generated target disconnect is not considered an error and is not logged or reported.

Target Retry. Target retry is not an error and is not logged or reported. The PCI master is responsible for determining the maximum number of retries.

Target Abort. When the PB issues a target abort it sets the Signaled Target Abort bit (bit 11) in the PCISTS Register. No further reporting or logging is done by the PB. The PCI initiator logs the target abort and may report the error.

Data Parity Errors. As a target on the PCI bus, the PB checks the data parity provided during write data cycles. If a parity error is detected during write data cycle, PERR# is asserted and bit 15 of the PCISTS Register is set. No further reporting or logging is done by the PB.

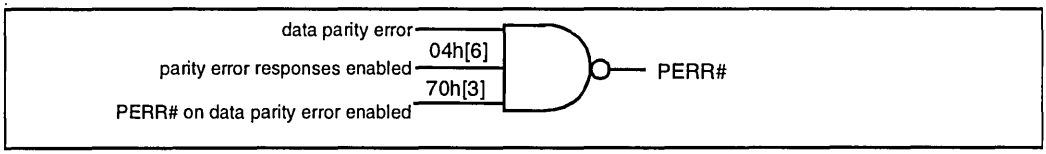

Figure 3. Logic Diagram of the Assertion of PERR#

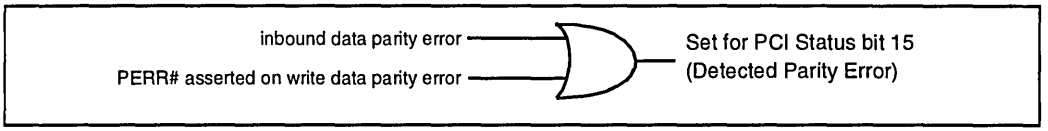

Figure 4. Logic Diagram of the Setting of the Detected Parity Error Bit0

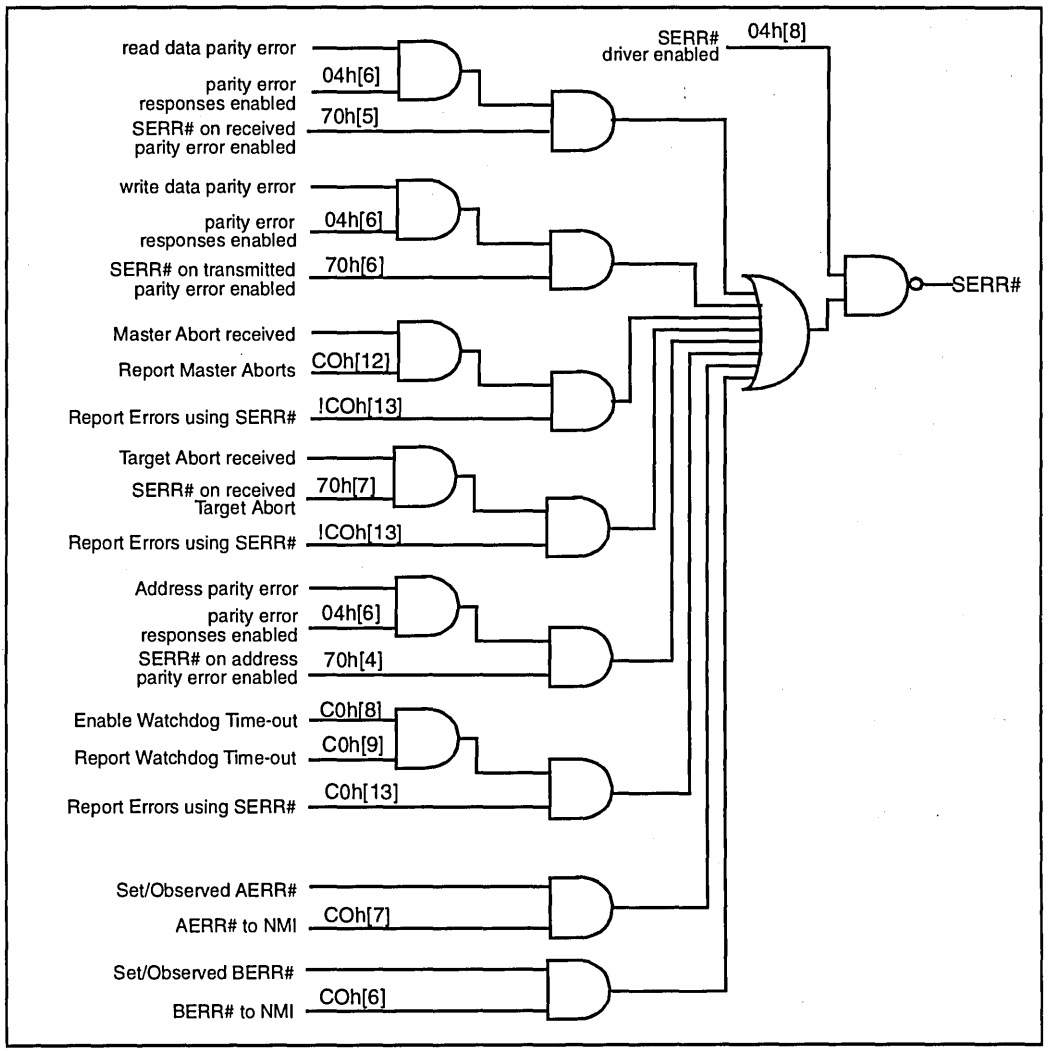

Figure 5. Logic Diagram of the Assertion of SERR#

3.5 Dual PB Architectures (82454GX Only)

In a dual bridge system, one PB is configured as the default bridge (Compatibility PB) after power-on RESET. The Compatibility PB provides a path to the ISA bus devices needed in a PC-compatible system such as the boot ROM. The Compatibility PB is the highest priority bridge in a dual bridge system to ensure a fast enough response time for ISA bus masters. See the Clocks, Reset, and Configuration section for details on configuring a PB as the Compatibility PB.

Multiple I/O APICs

In a dual PB system, the auxiliary PCI bus interrupt requests are routed to the auxiliary bus I/O APIC. When booting the system with one processor, the IRQ control logic is enabled, feeding the interrupt request to the standard interrupt controller in the ESC. When the system is in multiprocessor mode, the routing logic is disabled after ensuring PB buffer coherency, and interrupt requests are forwarded to the processors via the APIC bus. The Intel 82379AB (SIO.A) may be utilized as a stand-alone I/O APIC device. However, the additional logic for interrupt/memory consistency and the interrupt steering logic is not provided in the SIO.A and must be implemented externally.

Dual Bridge Arbitration for the Host Address Bus

The PB requests the host address bus with BPRI#. However, only one bridge is allowed to drive BPRI# at a time. With two PBs, an internal arbiter is used to establish bus ownership. This arbitration is transparent to the CPU and other symmetric bus agents.

In a two PB system, the compatibility PB acts as the arbitration unit between it and the other PB, as shown in Figure 6. When a PB is programmed to be the arbitration unit, its IOGNT# is the input for the IOREQ# from the other bridge and IOREQ# is the output to IOGNT# of the other bridge.

Figure 7 shows the minimum arbitration timing in a two bridge system. IOGNT# may assert later than shown and IOREQ# may negate later than the two clocks after IOGNT# negates.

The arbiter bridge can assert BPRI# as long as it has not asserted its IOREQ# (Grant to the other bridge) and BPRI# is not currently driven. In turn, the other bridge, after receiving it's IOGNT#, samples BPRI# released before assuming ownership of BPRI#. This allows the BPRI# arbitration to be performed in parallel with another bridge transfer. This timing is shown in Figure 8.

Bridge-to-bridge misaligned (split) locks are not recommended and could cause deadlock in systems.

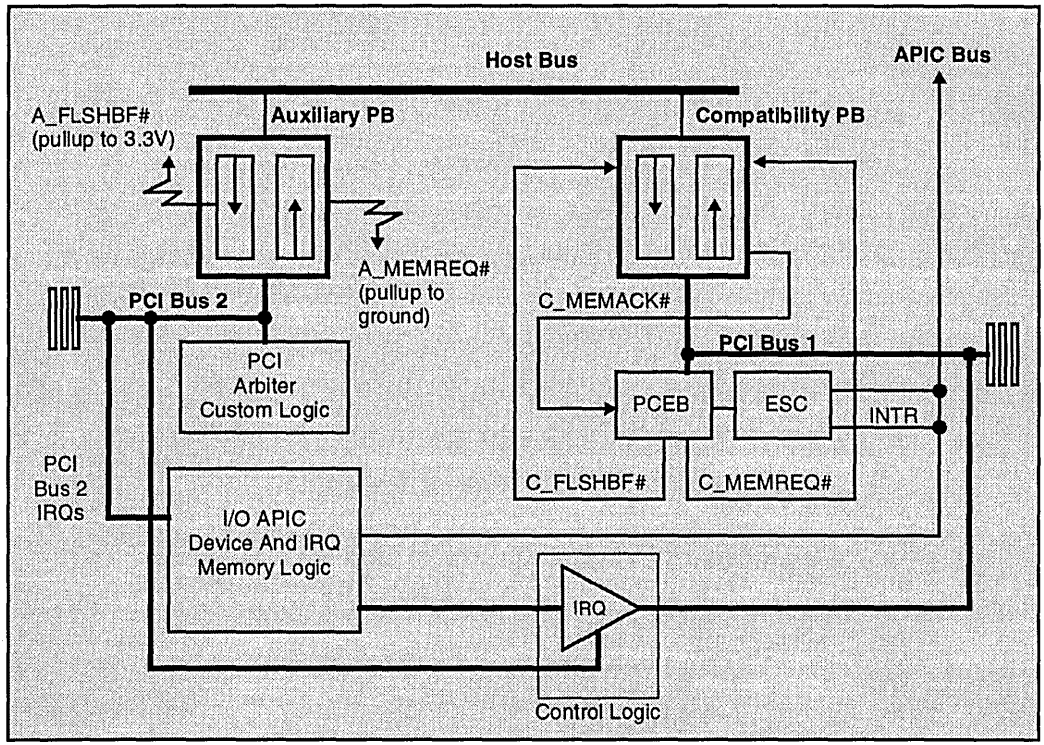

Figure 6. Dual Bridge System Configuration

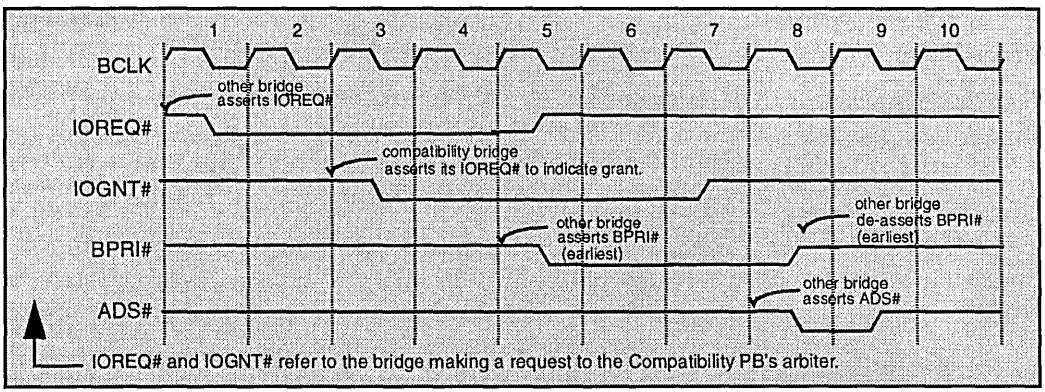

Figure 7. BPRI# Arbitration Timing

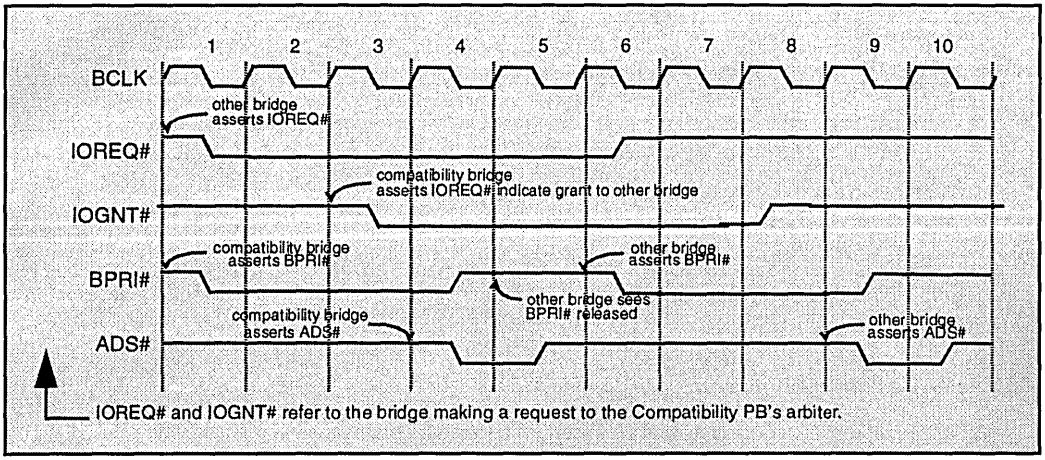

Figure 8. BPRI# Arbitration Overlapped with Bridge Transfer

Bridge-to-Bridge Communication

PB-to-PB communication is supported by the PB, but is not recommended for optimal performance.

PB-to-PB transactions involving a standard bus bridge (SIO, PCEB/ESC) require special precautions to avoid deadlock and latency problems. The PB does NOT support PB-to-PB transactions from agents that cannot be backed off such as those originating on an ISA or EISA bus and targeting a device on a different PB's PCI bus. Any device that asserts FLSHBUF# must be targeting a device on the local PCI bus or the host bus.

Dual PB Configuration (82454GX only)

During a power-on reset (PWRGD asserted), IOREQ# and IOGNT# provide a unique identification number for each PB (PBID). The PBID is part of the PB's PCI Bridge Device Number and is available to programmers via the BDNUM Register (offset 49h). The Dual PB system must have a pull-up and a pull-down as shown in Figure 9. The encoding for these signals is shown in Table 10

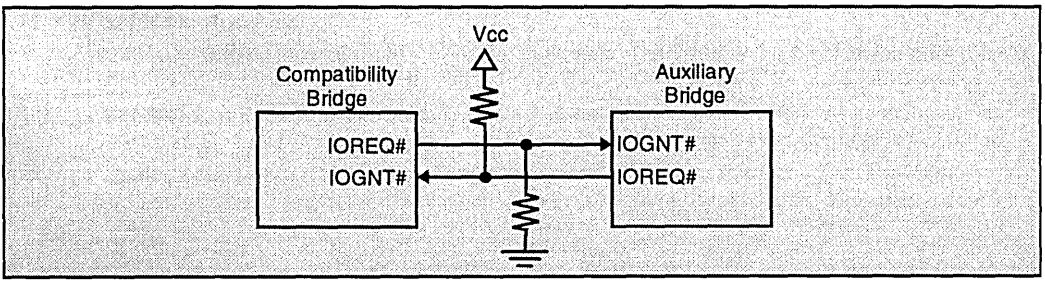

Figure 9. IOGNT# & IOREQ# Pull-Ups/Pull-Downs (Dual Bridge System)

Table 10. Bridge Device Number Encoding

IOGNT#	IOREQ#	PBID (BDNUM Register)	Description
High	Low	01	Compatibility PB
Low	High	10	Auxiliary PB in a two PB system

3.6 Peripheral Operation and Performance

The 82454 PB is designed for optimum processor performance to get the most out of a Pentium Pro processor's capabilities. In systems with multiple PCI devices, one must take into account the architecture of the 82454 PB in order to maximize overall system performance.

3.6.1 MATCHING PERIPHERALS TO THE 450KX/GX

The 82454 PB is optimized for use with high performance PCI peripherals. Support for multiple CPUs and multiple 82454 PBs comes at the cost of greater latency to system memory, which must be offset by more efficient use of the PCI bus to achieve high bandwidth I/O throughput.

For best system performance, only devices meeting the following criteria should be used in conjunction with the 82454 PB:

1. **High throughput peripherals should be PCI bus masters that control their own DMA.** Peripherals which act as bus masters transfer data to and from memory with minimal intervention from the CPU. The available bandwidth for such devices is considerably greater than that available to *programmed-I/O* devices, which require the CPU to transfer data on their behalf. Bus mastering devices also allow the CPU to pursue other work in parallel with I/O transfers from PCI, resulting in higher overall system efficiency. Finally, outbound traffic from the CPU interferes with inbound bus mastering transactions, as they both compete for ownership of the 82454 PB. The latter effect implies that one non-mastering device can adversely impact the performance of several other mastering devices.
2. **Peripherals should support the advanced PCI command subset.** The advanced PCI commands are Memory Read Line (command encoding **E**), Memory Read Multiple (command encoding **C**), and Memory Write and Invalidate (command encoding **F**). Devices utilizing these commands differentiate between long data transfers and short overhead transfers, and use appropriate PCI commands for each. Further, such devices tend to implement sufficient on-board data FIFO space to support full-speed PCI burst transfers greater than a cache line in length.
3. **Latency-sensitive peripherals should provide adequate data buffering.** Peripherals such as network interface cards have a latency requirement once transmission has begun. If they cannot buffer sufficient data on board prior to initiating a transfer, they are subject to transmission under-runs when competing I/O subsystem activity causes the bandwidth across the wire to exceed the bandwidth into system memory. A similar scenario occurs in the opposite direction if the bandwidth into system memory falls below the bandwidth across the wire. For example, a PCI card supporting *fast ethernet* at 100 Mbit/sec, should provide at least 128 bytes of data buffering for transfers in each direction.

3.6.2 DISTRIBUTING PERIPHERALS WITHIN THE I/O SUBSYSTEM

While this is not necessary for system operation, systems implementing dual 82454 PBs have additional latitude to isolate high speed I/O devices from competing system traffic initiated by the CPU.

All graphics and the vast majority of I/O space communication (such as keyboard controller, system timer, and interrupt support) will be directed to the *primary* PCI bus behind the Compatibility 82454 PB. (This is the bus with a subsequent connection via another bridge to an ISA or EISA bus.) This processor traffic will compete with bus mastering peripheral devices attempting to move data to and from system memory. It is desirable then to place latency sensitive devices behind the *Auxiliary* 82454 PB, to isolate them from competing CPU traffic.

In a full system configuration, in which all PCI slots are occupied, it is preferable to segregate peripherals intelligently. Limit the primary PCI bus to graphics accelerators and SCSI RAID controllers, leaving Auxiliary 82454 PB PCI slots free for latency-sensitive devices such as network adapters. In systems connecting a large number of network adapters, divide them evenly between the two busses to minimize the amount of latency-sensitive competition at any one point in the system.

3.6.3 PCI-TO-PCI BRIDGES

Since PCI-to-PCI bridge (P2P) components are a popular mechanism for increasing the connectivity of a PCI subsystem, the issues associated with using them should be understood. Note that these components are not only used on motherboards, but are sometimes used on PCI adapters as well.

The hierarchical bus added into the system in this manner must compete with all other devices on the primary bus for bandwidth. Further, the devices sharing the additional PCI bus connected via the P2P must compete with each other for serial service across the P2P bridge. This means that peripherals placed behind a P2P device will perceive higher latency to memory and will be limited to shorter burst transfers; a condition which may cause errors in latency-sensitive peripherals.

Finally, if a P2P device in use is not fully compliant with the PCI 2.1 specification, the system is exposed to unresolvable conflicts between multiple bus masters issuing transactions attempting to cross between the hierarchical PCI busses. To eliminate the possibility of a resulting livelock failure, the system must operate with CPU-PCI write posting disabled. This will degrade the performance of outbound traffic such as graphics, but will not adversely affect the performance of bus mastering I/O devices.

3.6.4 BIOS PERFORMANCE TUNING

Specific system configurations each have an optimum set of performance feature settings, but the following recommendations establishes a good baseline to begin system tuning.

The system designer should tune the read prefetch enable bits in the 82454 to avoid wasted host bus bandwidth due to short reads that do not make use of prefetched data. Most PCI peripherals which implement the advanced PCI command set also use these commands as recommended in the PCI specification. Specifically, PCI masters should use the PCI memory read command for transfers less than a cache line in length, the PCI memory read line command for transfers of one or two cache lines, and the PCI memory read multiple command for transfers of two or more cache lines. Given no specific data on the peripherals to be used in the system, the BIOS should default to a configuration which assumes that PCI peripherals will behave as described above. That is, enable the line read alias bits for all PCI read command types, but only enable the read prefetch bit for the PCI memory read multiple command. This configuration may be modified, perhaps in a setup utility, if that provides better performance for a given set of devices.

The relatively high latency to memory in 450KX/GX-based systems will require larger PCI master latency timer values than the typical 32 clock default. In order to allow each master the opportunity to burst multiple cache lines per transfer, the master latency timer (MLT) of each PCI master in the system should be set to a value between 48h and 60h. (Note that an MLT setting that is arbitrarily larger than 60h will allow a master capable of extremely long PCI bursts to adversely impact the performance of other masters with more limited burst capability.)

3.7 Clock, Reset, and Configuration

3.7.1 SYSTEM CLOCKING

The PB operates in two clock domains. The PB interface to the host bus operates at the host bus clock frequency. The host bus clock is generated externally and distributed to host bus components by a low skew clock driver. The clock driver provides multiple copies of the bus clock. The PB receives its copy of the host bus clock through the BCLK input pin.

The PB interface to the PCI bus operates at the PCI bus clock frequency. The PCI bus clock is generated internally by the PB and is ½ the frequency of the host bus clock frequency. This output is designed to drive a single load and must be distributed by an external low skew clock driver. The external clock driver provides multiple copies of the bus clock. The PB receives a matching copy of the skewed PCI bus clock through its PCLKIN pin.

3.7.1.1 Host Bus Clock

Host Bus clock distribution is shown in Figure 10. The loading on the host bus clock lines must be balanced in order to minimize clock skew among the components on the host bus. This may require adjustment of clock line lengths. Note that the BCLK input to the PB must be running for 10 clocks before the assertion of PWRGD.

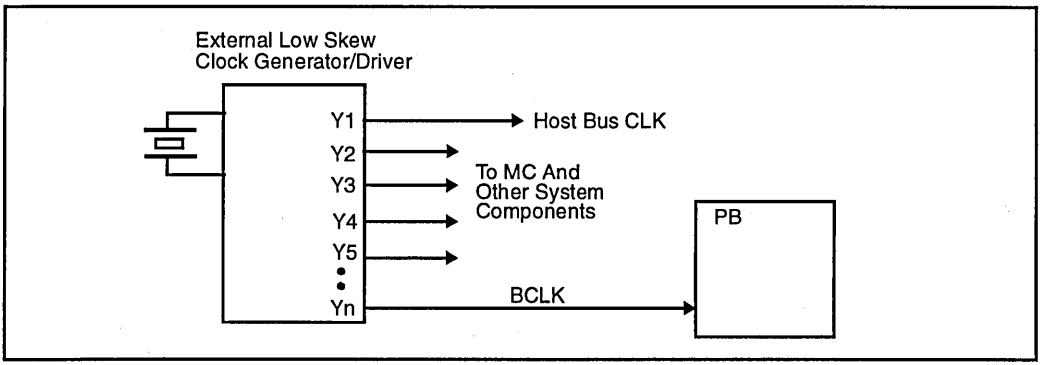

Figure 10. Host Bus Clock Distribution

3.7.1.2 PCI Clock

PCI clock distribution is illustrated in Figure 11. An external 10KΩ pull-up resistor is required to place the PB in derived clock mode (only mode supported). The PB provides a PCI bus clock that is generated by dividing the processor clock frequency by two. The phase of the PCLK signal is matched to the host clock. Externally, this PCI clock drives a low skew clock driver which in turn supplies multiple copies of the PCI clock to the PCI bus. One of the outputs of the external clock driver is fed back to the PB. This copy is expected to meet the skew requirements of the PCI specification. A PLL in the PB forces the external PCI clock to phase lock to the internal PCI clock tree.

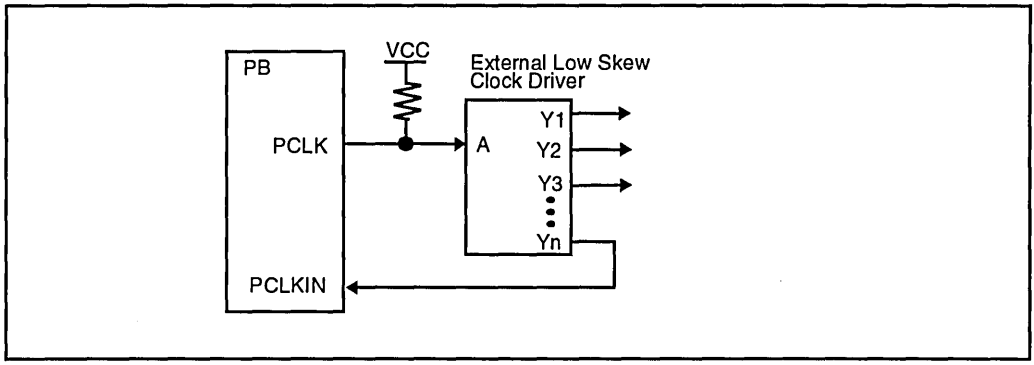

Figure 11. PCI Clock Distribution

3.7.2 SYSTEM RESET

Power-On Reset

When the system is initially powered, the power supply must wait until all voltages are stable for at least one millisecond, and then assert the PWRGD signal. Note that BCLK must be running to the PB for 10 clocks before the assertion of PWRGD. The PB captures their Bridge Device Number identification when PWRGD is asserted (see PB Configuration section).

While RESET# is asserted, the PB resets and initializes its internal registers to the default state. The PB also initializes the PCI busses by asserting PCIRST# for a minimum of one millisecond. While RESET# is asserted, the PB (Compatibility PB in an 82454GX dual PB system) drives the appropriate host data bus signals with the values specified in the Configuration Values Driven on Reset Register. In addition to asserting RESET#, the PB (Compatibility PB in an 82454GX dual PB system) also asserts CRESET# and continues to assert CRESET# two clocks longer than RESET#. CRESET# may be used to select a Mux that drives the host bus clock to core clock ratio onto pins LINT[1:0], IGNNE#, and A20M# of the CPU during RESET#.

Initially all PBs assert RESET# until the ID is captured and the Compatibility PB is established. The Compatibility PB continues to assert RESET# for a minimum of one millisecond and RESET# becomes an input for the Auxiliary PB; however, it does not affect their captured ID.

Programmed Hard Reset

The PB (Compatibility PB in an 82454GX dual PB system) can be programmed to deliver a hard reset (assert RESET#) to the host bus through the TRC Register. Note that the internal register values are reset.

Programmed Soft Reset (INIT#)

The PB (Compatibility PB in an 82454GX dual PB system) can be programmed to deliver a soft reset (INIT#) to the processors through the TRC Register. Note that the internal register values are preserved.

Programmed PCI Bus Reset

The PB (both PBs in an 82454GX dual PB system) can be programmed to reset their PCI buses (assert PCIRST#) without resetting the host bus (via the PCI Reset Register). Note that internal register values are preserved.

Programmed CPU BIST

The PB (Compatibility PB in an 82454GX dual PB system) can be programmed to put the processor into BIST mode via the TRC Register. CPU BIST is triggered by performing a hard reset and having the INIT# signal asserted on the edge that RESET# is released. Note all 450KX/GX PCIset devices are reset during the hard reset portion of this operation.

3.7.3 SYSTEM INITIALIZATION

All host bus devices must sample the following configuration options at reset:

- Address/request/response parity checking: Enabled or Disabled
- AERR detection enable
- BERR detection enable
- BINIT detection enable
- FRC mode: Enabled or Disabled
- Power-on reset vector: 1M or 4G
- In-Order Queue depth: 1 or 8
- APIC cluster ID: 0, 1, 2, or 3
- Symmetric agent arbitration ID: 0, 1, 2, 3

The MC provides the Symmetric Arbitration ID parameter. The PB provides some of the other parameters. See Configuration Values Driven on Reset Register.

3.7.4 DUAL PB CONFIGURATION (82454GX ONLY)

During a power-on reset (PWRGD asserted), IOREQ# and IOGNT# provide a unique identification number for each PB (PBID). See Dual PB Architectures Section for details.

3.7.5 USING THE 82379AB SIO.A PCI-TO-ISA BRIDGE WITH THE 450KX/GX

There is an anomaly with systems that use the 82379AB (SIO.A) during targeted PCI Resets. In addition, 450GX/KX systems can boot improperly at power-up and react improperly to the assertion of the Pentium Pro bus signal BINIT# signal (due to the assertion of PCIRST# via BINIT#).

The SIO.A drives SMI#, ALT_A20, INT, NMI, IGNNE#, ALT_RST#, and STPCLK# low while PCIRST# is asserted low, and does not drive them high until after PCI reset is released. An anomaly can exist with these seven signals remaining low during and immediately after PCIRST# is negated. The three instances in which this can cause an anomaly are: during a targeted PCI Reset, and in a 450GX/KX-Pentium Pro processor system, both during power-up and when BINIT# is asserted on the Pentium Pro processor bus.

Power-Up

During power up of an 450GX/KX-Pentium Pro system, the OPB negates PCIRST# and RESET# (to the Pentium Pro) simultaneously. The delay in negating these seven signals after PCIRST# is driven inactive can cause these signals to be sampled active low by the Pentium Pro when RESET# (to the Pentium Pro) is released.

- SMI#: When SMI# is sampled low at power-up, the Pentium Pro attempts to jump to the SMI handler instead of to the boot vector. An external solution is necessary to avoid this erroneous power-up condition. Essentially, SMI# must be blocked from being sampled low by the CPU when RESET# is driven inactive. If SMI# is not being used, the SMI# input to the CPU can be pulled high. For systems using SMI#, the solution shown below on the SMI# signal ensures that SMI# is high when RESET# transitions inactive.

  This solution is not needed in systems using Pentium Pro B0 stepping (or later). These steppings of the Pentium Pro will not sample these seven inputs for at least 300ns after RESET# is negated.
- ALT_RST#: The CPU will reset again (generates INIT# to the Pentium Pro). This signal must be blocked. This solution is not needed in systems using Pentium Pro B0 stepping (or later). These steppings of the Pentium Pro will not sample these seven inputs for at least 300ns after RESET# is negated.
- INT: Not an issue since this signal remains low following PCIRST#.
- NMI: Not an issue since this signal remains low following PCIRST#.
- IGNNE#: Has no affect on the processor when sampled low during power-up.
- ALT_A20: Has no affect on the processor when sampled low during power-up.
- STPCLK#: Has no affect on the processor when sampled low during power-up.

BINIT# Assertion

When BINIT# is asserted by an agent on the Pentium Pro processor bus, the 82454KX/GX asserts PCIRST# to reset the PCI Bus. The SIO.A drives SMI#, ALT_A20, INT, NMI, IGNNE#, ALT_RST#, and STPCLK# low while PCIRST# is asserted low, and does not drive these signals high until after PCI reset is released. Several of these signals must be blocked with external logic if the system architecture cannot handle them going low as a result of BINIT#. Architectural Considerations consist primarily of how the system's Pentium Pro processors have been configured to handle the assertion of BINIT# and how any external error handling logic might influence the need for blocking logic. All steppings of the Pentium Pro processor need these solutions (where appropriate) if BINIT# is to be handled.

- ALT_RST#: The low assertion of this signal causes the Pentium Pro processor to be reset each time BINIT# is asserted on the Pentium Pro processor bus (this signal is combined with INIT# from the PB to generate the INIT# signal to the Pentium Pro processor). This signal MUST be blocked with external logic if the system architecture cannot handle this.

- INT: Not a problem since this signal remains low following PCIRST#. Note that any pending interrupts on this pin will be lost when the SIO receives the PCI reset.
- NMI: Not a problem since this signal remains low following PCIRST#. Note that any pending interrupts on this pin will be lost when the SIO receives the PCI reset.
- IGNNE#: Unless software or the system architecture can ensure that no floating point errors are generated during the PCI reset, this signal must be blocked.
- ALT_A20: Designers must be aware that this signal will be driven (and remain) low during and following PCIRST#. Additionally, RSTDRV is asserted during PCIRST#. This sill cause the RSTAR# output of the keyboard controller to also be set low. These factors must be considered in any design implementing BINIT# functions.
- STPCLK#: If the STPCLK# function is enabled in the design, then this signal must incorporate blocking logic.
- SMI#: If SMI# is enabled in the design, then this signal must incorporate blocking logic.

Targeted PCI Resets

Systems that support targeted PCI Resets (Resetting the PCI bus via Software control without resetting the microprocessor) may have a problem with some of the seven signals being asserted low during the targeted PCI reset. Since the microprocessor can not know when a PCIRST is occurring, this fix must be incorporated in order to reset the PCI bus via the register. This affects all designs using the SIO.A.

- ALT_RST#: The low assertion of this signal will cause the microprocessor to be reset each time the PCI bus is reset (this signal is normally combined with a CPU soft reset pin from another component to generate the INIT# signal to the microprocessor). This signal MUST be blocked with external logic.
- INT: Not an issue since this signal remains low following PCIRST#. Note that any pending interrupts on this pin will be lost when the SIO receives the PCI reset.
- NMI: Not an issue since this signal remains low following PCIRST#. Note that any pending interrupts on this pin will be lost when the SIO receives the PCI reset.
- IGNNE#: Unless software or the system architecture can ensure that no floating point errors are generated during the PCI bus reset, this signal must be blocked.
- ALT_A20: Designers must be aware that this signal will be driven (and remain) low during and following PCIRST#. Additionally, RSTDRV is asserted during PCIRST#. This sill cause the RSTAR# output of the keyboard controller to also be set low. These factors must be considered in any design implementing targeted PCI resets.
- STPCLK#: If the STPCLK# function is enabled in the design and it is not desirable to have the CPU "shutdown" throughout the PCI reset, then this signal must incorporate blocking logic.
- SMI#: If SMI# is enabled in the design, then this signal must incorporate blocking logic.

The RESET MASK blocking circuit shown in Figure 12 will block signal A from being seen by the CPU during the PCI Reset. The second flip flop is necessary to avoid a glitch on the Z output to the CPU which can happen if signal A is asserted simultaneously with PCIRST#.

The blocking circuitry for all of the signals should be incorporated into a PLD. This will ease loading on PCICLK and PCIRST#.

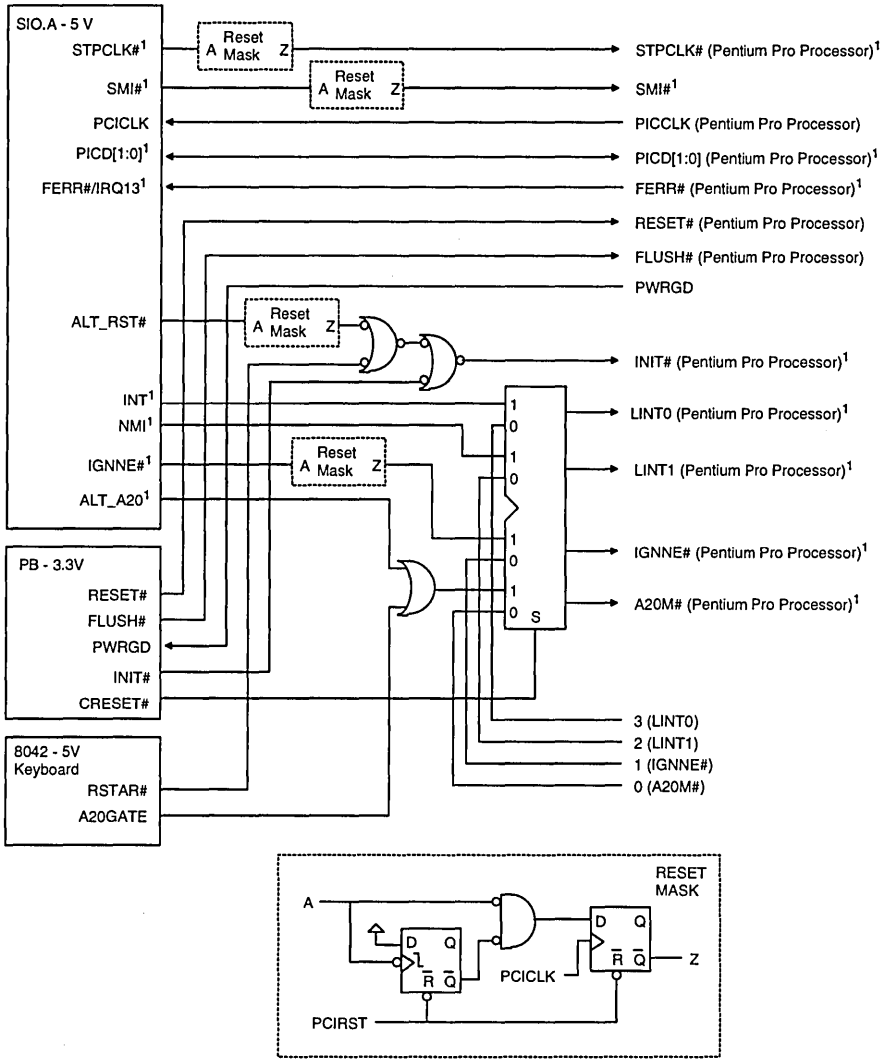

052053_4.ds4

NOTES:
1. 5V to 3.3V translation is not shown. The system designer must ensure correct logic levels are provided. All PB outputs are 3.3 volt.

Figure 12. Blocking Logic For 82420KX/GX Designs Using the 82379AB SIO.A

3.8 Host to PCI Bus Command Translation

If a valid Pentium Pro processor bus command is directed at the bridge, the PB generates a PCI bus request. When the PB is granted the PCI bus, it issues a PCI command after the commit point of the Pentium Pro processor bus transaction. Pentium Pro processor bus commands that are directed at the PB and consequently to the PCI bus must be converted into appropriate PCI bus commands. The PB is a non-caching agent on the host bus; however, it must respond appropriately to Pentium Pro processor commands that are cache oriented.

Table 11. Host to PCI Bus Command Translation

Host Bus Command (ASZ = 36, DSZ=64)	Other Encoded Information	PCI Bus Command
Deferred Reply	don't care	none
INTA	LEN: <= 8 bytes	Interrupt Acknowledge with BE[0] asserted.
Special Cycles	BE: Shutdown	none
	BE: Stop Clock Acknowledge	PCI Special Cycle - Stop Clock Grant
	BE: all others	none
I/O Read	LEN: <= 8 bytes up to 4 BEs asserted	I/O Read (one or two transactions)
Branch Trace Message		none
I/O Write	LEN: <= 8 bytes up to 4 BEs asserted	I/O Write (one or two transactions)
Read Invalidate	don't care	Memory Read Line (8 Dword burst starting with the low address)
Code Read	LEN: <= 8 bytes without all byte enables asserted	Memory Read (one or two transactions)
	LEN: <= 8 bytes with all byte enables asserted	Memory Read (2 Dword burst starting with the low address)
	LEN: 16 bytes	Memory Read (4 Dword burst starting with the low address)
	LEN: 32 bytes	Memory Read Line (8 Dword burst starting with the low address)
Memory Read	LEN: <= 8 bytes without all byte enables asserted	Memory Read (one or two transactions)
	LEN: <= 8 bytes with all byte enables asserted	Memory Read (2 Dword burst starting with the low address)
	LEN: 16 bytes	Memory Read (4 Dword burst starting with the low address)
	LEN: 32 bytes	Memory Read Line (8 Dword burst starting with the low address)

Table 11. Host to PCI Bus Command Translation (Continued)

Host Bus Command (ASZ = 36, DSZ=64)	Other Encoded Information	PCI Bus Command
Central Agent Reserved Transactions with no data.		None
Reserved Encodings	These encodings are ignored. A bus time-out will complete the cycle	
Memory Write	LEN: <= 8 bytes without all byte enables asserted	Memory Write (one or two transactions)
	LEN: <= 8 bytes with all byte enables asserted	Memory Write (2 Dword burst starting with the low address)
	LEN: 16 bytes	Memory Write (4 Dword burst starting with the low address)
	LEN: 32 bytes	Memory Write and Invalidate or Memory Write (8 Dword burst starting with the low address)

3.9 PCI to Host Bus Command Translation

When a PCI bus command is directed at the bridge, the PB generates a BPRI# request or arbitrates for BPRI# ownership if there are two bridges. The actual point in time when the BPRI# is issued depends on a number of factors including whether the bridge accepted the PCI request or forced a retry to the PCI master, and when a complete cache line is filled during a write command.

PCI bus commands that are directed at the PB and consequently to the host bus are converted into the following appropriate host bus commands.

Table 12. PCI to Host Bus Command Translation

PCI Bus Command	Host Bus command
Memory Read	Memory Read LEN: ≤ 8 or LEN: 32 (When *CPU Line Read for PCI memory Read Commands* [Bit 8] is enabled in the PCI Read/Write Control Register: 54-55h)
Memory Read Line	Memory Read LEN: ≤ 8 or LEN: 32 (When *CPU Line Reads for PCI Memory Read Line Commands* [Bit 3] is enabled in the PCI Read/Write Control Register: 54-55h)
Memory Read Multiple	Memory Read LEN: ≤ 8 or LEN: 32 (When *CPU Line Read Multiple for PCI Memory Read Multiple Commands* [Bit 5] is enabled in the PCI Read/Write Control Register: 54-55h)

Table 12. PCI to Host Bus Command Translation (Continued)

PCI Bus Command	Host Bus command
Memory Write	Memory Write LEN: ≤ 8
Memory Write and Invalidate	Line Write LEN: 32 (If *PCI-to-CPU Write Posting* is enabled in bit 0 of the PCI Read/Write Control Register.
I/O Read	I/O Read LEN: ≤ 8
I/O Write	I/O Write LEN: ≤ 8
Dual Address	Translated up to a 36-bit host bus address.
Other Command Encodings	All other command encodings[1] are ignored by the PB.

NOTES:

This includes Special Cycles. There is no mechanism to pass Special Cycles originating on the PCI bus to the host bus. Accesses to the CONFADD and CONFDATA Registers can be passed to the host bus if the PB is not programmed to accept CF8 and CFC transactions from the host bus. For the 82454GX in a dual PB system, this allows the auxiliary bridge to pass configuration commands to the host bus from its PCI bus.

In general, the length of PCI master bursts is indeterminate unless the master can be programmed for specific burst lengths. The PCI bridge takes advantage of the PCI bursting capability by always trying to assemble the most efficient host bus size transfers. The bridge selectively asserts STOP# (retry) to the PCI master as posting buffers become unavailable.

During PCI master reads the bridge will attempt to make the most efficient use of the host bus by generating line reads or partial reads based upon the type of PCI read occurring and the options programmed into the bridge configuration registers.

4.0 PB PINOUT AND PACKAGE INFORMATION

4.1 Pin Assignment

There are two packages for the 82454KX/GX—304 pin QFP and 352 pin BGA. Pins unique to the 450GX are shown in a list at the center of the figure. In the tables, the first name is the 450GX name. **Note that TESTLO pins must be pulled-low with a 150Ω resistor and TESTHI pins must be pulled to 3.3V with a 10KΩ resistor. GTLHI pins should be pulled up with 10KΩ to V_{TT}.**

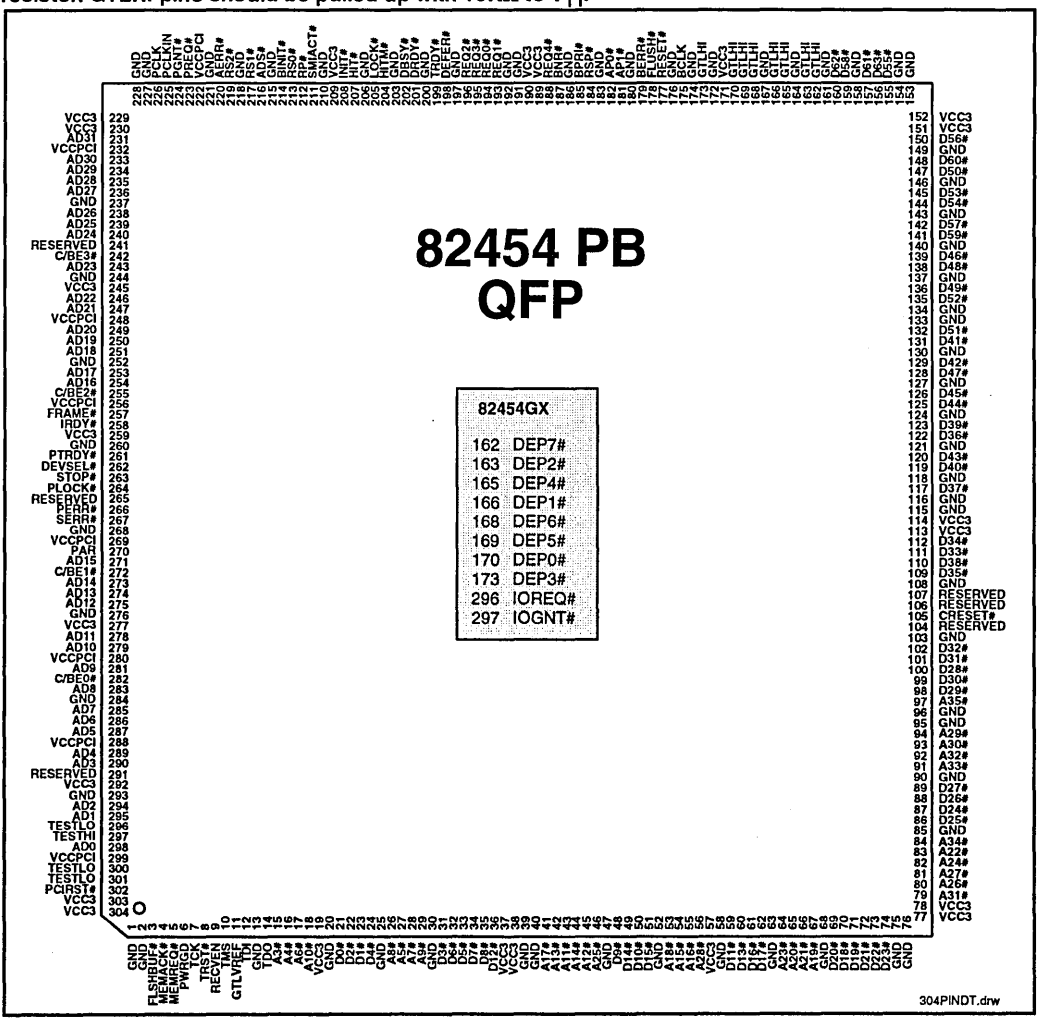

Figure 13. 82454 Pin Assignment (304-Pin QFP)

Table 13. 82454KX/GX PB Alphabetical Pin List (304-Pin QFP)

Signal	Pin #	Type
A3#	15	I/O
A4#	16	I/O
A5#	27	I/O
A6#	17	I/O
A7#	28	I/O
A8#	26	I/O
A9#	29	I/O
A10#	18	I/O
A11#	43	I/O
A12#	45	I/O
A13#	42	I/O
A14#	44	I/O
A15#	54	I/O
A16#	55	I/O
A17#	41	I/O
A18#	53	I/O
A19#	67	I/O
A20#	65	I/O
A21#	66	I/O
A22#	83	I/O
A23#	64	I/O
A24#	82	I/O
A25#	46	I/O
A26#	80	I/O
A27#	81	I/O
A28#	56	I/O
A29#	94	I/O
A30#	93	I/O
A31#	79	I/O
A32#	92	I/O
A33#	91	I/O
A34#	84	I/O
A35#	97	I/O
AD0	298	I/O
AD1	295	I/O
AD2	294	I/O

Table 13. 82454KX/GX PB Alphabetical Pin List (304-Pin QFP) (Continued)

Signal	Pin #	Type
AD3	290	I/O
AD4	289	I/O
AD5	287	I/O
AD6	286	I/O
AD7	285	I/O
AD8	283	I/O
AD9	281	I/O
AD10	279	I/O
AD11	278	I/O
AD12	275	I/O
AD13	274	I/O
AD14	273	I/O
AD15	271	I/O
AD16	254	I/O
AD17	253	I/O
AD18	251	I/O
AD19	250	I/O
AD20	249	I/O
AD21	247	I/O
AD22	246	I/O
AD23	243	I/O
AD24	240	I/O
AD25	239	I/O
AD26	238	I/O
AD27	236	I/O
AD28	235	I/O
AD29	234	I/O
AD30	233	I/O
AD31	231	I/O
ADS#	216	I/O
AERR#	220	I/O
AP0#	182	I/O
AP1#	181	I/O
BCLK	175	I
BERR#	179	I/O
BINIT#	214	I/O

Table 13. 82454KX/GX PB Alphabetical Pin List (304-Pin QFP) (Continued)

Signal	Pin #	Type
BNR#	187	I/O
BPRI#	185	I/O
C/BE0#	282	I/O
C/BE1#	272	I/O
C/BE2#	255	I/O
C/BE3#	242	I/O
CRESET#	105	O
D0#	21	I/O
D1#	23	I/O
D2#	22	I/O
D3#	31	I/O
D4#	24	I/O
D5#	33	I/O
D6#	32	I/O
D7#	34	I/O
D8#	35	I/O
D9#	48	I/O
D10#	50	I/O
D11#	59	I/O
D12#	36	I/O
D13#	60	I/O
D14#	49	I/O
D15#	51	I/O
D16#	61	I/O
D17#	62	I/O
D18#	70	I/O
D19#	71	I/O
D20#	69	I/O
D21#	72	I/O
D22#	73	I/O
D23#	74	I/O
D24#	87	I/O
D25#	86	I/O
D26#	88	I/O
D27#	89	I/O
D28#	100	I/O

Table 13. 82454KX/GX PB Alphabetical Pin List (304-Pin QFP) (Continued)

Signal	Pin #	Type
D29#	98	I/O
D30#	99	I/O
D31#	101	I/O
D32#	102	I/O
D33#	111	I/O
D34#	112	I/O
D35#	109	I/O
D36#	122	I/O
D37#	117	I/O
D38#	110	I/O
D39#	123	I/O
D40#	119	I/O
D41#	131	I/O
D42#	129	I/O
D43#	120	I/O
D44#	125	I/O
D45#	126	I/O
D46#	139	I/O
D47#	128	I/O
D48#	138	I/O
D49#	136	I/O
D50#	147	I/O
D51#	132	I/O
D52#	135	I/O
D53#	145	I/O
D54#	144	I/O
D55#	155	I/O
D56#	150	I/O
D57#	142	I/O
D58#	159	I/O
D59#	141	I/O
D60#	148	I/O
D61#	157	I/O
D62#	160	I/O
D63#	156	I/O
DBSY#	202	I/O

Table 13. 82454KX/GX PB Alphabetical Pin List (304-Pin QFP) (Continued)

Signal	Pin #	Type
DEFER#	198	I/O
DEP0#/ GTLHI	170	I/O
DEP1#/ GTLHI	166	I/O
DEP2#/ GTLHI	163	I/O
DEP3#/ GTLHI	173	I/O
DEP4#/ GTLHI	165	I/O
DEP5#/ GTLHI	169	I/O
DEP6#/ GTLHI	168	I/O
DEP7#/ GTLHI	162	I/O
DEVSEL#	262	I/O
DRDY#	201	I/O
FLSHBUF#	3	I
FLUSH#	178	O
FRAME#	257	I/O
GND	1	V
GND	2	V
GND	13	V
GND	20	V
GND	25	V
GND	30	V
GND	39	V
GND	40	V
GND	47	V
GND	52	V
GND	58	V
GND	63	V
GND	68	V
GND	75	V

Table 13. 82454KX/GX PB Alphabetical Pin List (304-Pin QFP) (Continued)

Signal	Pin #	Type
GND	76	V
GND	85	V
GND	90	V
GND	95	V
GND	96	V
GND	103	V
GND	108	V
GND	115	V
GND	116	V
GND	118	V
GND	121	V
GND	124	V
GND	127	V
GND	130	V
GND	133	V
GND	134	V
GND	137	V
GND	140	V
GND	143	V
GND	146	V
GND	149	V
GND	153	V
GND	154	V
GND	158	V
GND	161	V
GND	164	V
GND	167	V
GND	172	V
GND	174	V
GND	176	V
GND	180	V
GND	183	V
GND	186	V
GND	191	V
GND	192	V
GND	197	V

Table 13. 82454KX/GX PB Alphabetical Pin List (304-Pin QFP) (Continued)

Signal	Pin #	Type
GND	200	V
GND	203	V
GND	206	V
GND	210	V
GND	215	V
GND	218	V
GND	221	V
GND	227	V
GND	228	V
GND	237	V
GND	244	V
GND	252	V
GND	260	V
GND	268	V
GND	276	V
GND	284	V
GND	293	V
GTLVREF	11	I
HIT#	207	I/O
HITM#	204	I/O
INIT#	208	O
IOGNT#/ TESTHI	297	I
IOREQ#/ TESTLO	296	I/O
IRDY#	258	I/O
LOCK#	205	I/O
MEMACK#	4	O
MEMREQ#	5	I
PAR	270	I/O
PCIRST#	302	O
PCLK	226	I/O
PCLKIN	225	I

Table 13. 82454KX/GX PB Alphabetical Pin List (304-Pin QFP) (Continued)

Signal	Pin #	Type
PERR#	266	I/O
PGNT#	224	I
PLOCK#	264	I/O
PREQ#	223	O
PTRDY#	261	I/O
PWRGD	6	I
RECVEN	9	I
REQ0#	194	I/O
REQ1#	193	I/O
REQ2#	196	I/O
REQ3#	195	I/O
REQ4#	188	I/O
RESERVED	104	NC
RESERVED	106	NC
RESERVED	107	NC
RESERVED	241	NC
RESERVED	265	NC
RESERVED	291	NC
RESET#	177	I/O
RP#	212	I/O
RS0#	213	I/O
RS1#	217	I/O
RS2#	219	I/O
RSP#	184	I/O
SERR#	267	O
SMIACT#	211	O
STOP#	263	I/O
TCK	7	I
TDI	12	I
TDO	14	O
TESTLO	300	I/O
TESTLO	301	I/O
TMS	10	I

Table 13. 82454KX/GX PB Alphabetical Pin List (304-Pin QFP) (Continued)

Signal	Pin #	Type
TRDY#	199	I/O
TRST#	8	I
VCC3	19	V
VCC3	37	V
VCC3	38	V
VCC3	57	V
VCC3	77	V
VCC3	78	V
VCC3	113	V
VCC3	114	V
VCC3	151	V
VCC3	152	V
VCC3	171	V
VCC3	189	V
VCC3	190	V
VCC3	209	V
VCC3	229	V
VCC3	230	V
VCC3	245	V
VCC3	259	V
VCC3	277	V
VCC3	292	V
VCC3	303	V
VCC3	304	V
VCCPCI	222	V
VCCPCI	232	V
VCCPCI	248	V
VCCPCI	256	V
VCCPCI	269	V
VCCPCI	280	V
VCCPCI	288	V
VCCPCI	299	V

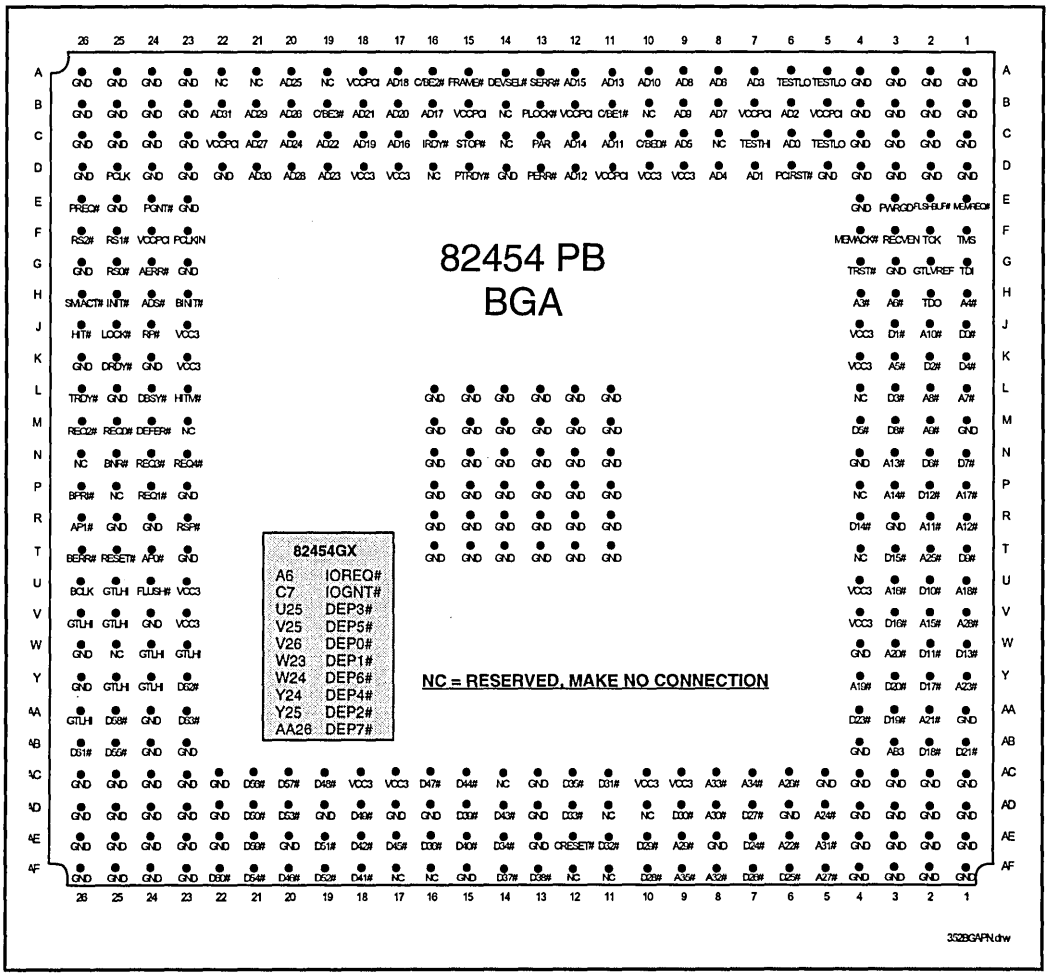

Figure 14. PB Pin Assignment (352 BGA)

Table 14. 82454 KX/GX PB Alphabetical Ball List (352 BGA)

Signal	Ball#	Type
A3#	H4	I/O
A4#	H1	I/O
A5#	K3	I/O
A6#	H3	I/O
A7#	L1	I/O
A8#	L2	I/O
A9#	M2	I/O
A10#	J2	I/O
A11#	R2	I/O
A12#	R1	I/O
A13#	N3	I/O
A14#	P3	I/O
A15#	V2	I/O
A16#	U3	I/O
A17#	P1	I/O
A18#	U1	I/O
A19#	Y4	I/O
A20#	W3	I/O
A21#	AA2	I/O
A22#	AE6	I/O
A23#	Y1	I/O
A24#	AD5	I/O
A25#	T2	I/O
A26#	AC6	I/O
A27#	AF5	I/O
A28#	V1	I/O
A29#	AE9	I/O
A30#	AD8	I/O
A31#	AE5	I/O
A32#	AF8	I/O
A33#	AC8	I/O
A34#	AC7	I/O
A35#	AF9	I/O
AD0	C6	I/O
AD1	D7	I/O
AD2	B6	I/O

Table 14. 82454 KX/GX PB Alphabetical Ball List (352 BGA) (Continued)

Signal	Ball#	Type
AD3	A7	I/O
AD4	D8	I/O
AD5	C9	I/O
AD6	A8	I/O
AD7	B8	I/O
AD8	A9	I/O
AD9	B9	I/O
AD10	A10	I/O
AD11	C11	I/O
AD12	D12	I/O
AD13	A11	I/O
AD14	C12	I/O
AD15	A12	I/O
AD16	C17	I/O
AD17	B16	I/O
AD18	A17	I/O
AD19	C18	I/O
AD20	B17	I/O
AD21	B18	I/O
AD22	C19	I/O
AD23	D19	I/O
AD24	C20	I/O
AD25	A20	I/O
AD26	B20	I/O
AD27	C21	I/O
AD28	D20	I/O
AD29	B21	I/O
AD30	D21	I/O
AD31	B22	I/O
ADS#	H24	I/O
AERR#	G24	I/O
AP0#	T24	I/O
AP1#	R26	I/O
BCLK	U26	I
BERR#	T26	I/O
BINIT#	H23	I/O

Table 14. 82454 KX/GX PB Alphabetical Ball List (352 BGA) (Continued)

Signal	Ball#	Type
BNR#	N25	I/O
BPRI#	P26	I/O
C/BE0#	C10	I/O
C/BE1#	B11	I/O
C/BE2#	A16	I/O
C/BE3#	B19	I/O
CRESET#	AE12	O
D0#	J1	I/O
D1#	J3	I/O
D2#	K2	I/O
D3#	L3	I/O
D4#	K1	I/O
D5#	M4	I/O
D6#	N2	I/O
D7#	N1	I/O
D8#	M3	I/O
D9#	T1	I/O
D10#	U2	I/O
D11#	W2	I/O
D12#	P2	I/O
D13#	W1	I/O
D14#	R4	I/O
D15#	T3	I/O
D16#	V3	I/O
D17#	Y2	I/O
D18#	AB2	I/O
D19#	AA3	I/O
D20#	Y3	I/O
D21#	AB1	I/O
D22#	AB3	I/O
D23#	AA4	I/O
D24#	AE7	I/O
D25#	AF6	I/O
D26#	AF7	I/O
D27#	AD7	I/O
D28#	AF10	I/O

Table 14. 82454 KX/GX PB Alphabetical Ball List (352 BGA) (Continued)

Signal	Ball#	Type
D29#	AE10	I/O
D30#	AD9	I/O
D31#	AC11	I/O
D32#	AE11	I/O
D33#	AD12	I/O
D34#	AE14	I/O
D35#	AC12	I/O
D36#	AE16	I/O
D37#	AF14	I/O
D38#	AF13	I/O
D39#	AD15	I/O
D40#	AE15	I/O
D41#	AF18	I/O
D42#	AE18	I/O
D43#	AD14	I/O
D44#	AC15	I/O
D45#	AE17	I/O
D46#	AF20	I/O
D47#	AC16	I/O
D48#	AC19	I/O
D49#	AD18	I/O
D50#	AD21	I/O
D51#	AE19	I/O
D52#	AF19	I/O
D53#	AD20	I/O
D54#	AF21	I/O
D55#	AB25	I/O
D56#	AC21	I/O
D57#	AC20	I/O
D58#	AA25	I/O
D59#	AE21	I/O
D60#	AF22	I/O
D61#	AB26	I/O
D62#	Y23	I/O
D63#	AA23	I/O
DBSY#	L24	I/O

Table 14. 82454 KX/GX PB Alphabetical Ball List (352 BGA) (Continued)

Signal	Ball#	Type
DEFER#	M24	I/O
DEP0#/ GTLHI	V26	I/O
DEP1#/ GTLHI	W23	I/O
DEP2#/ GTLHI	Y25	I/O
DEP3#/ GTLHI	U25	I/O
DEP4#/ GTLHI	Y24	I/O
DEP5#/ GTLHI	V25	I/O
DEP6#/ GTLHI	W24	I/O
DEP7#/ GTLHI	AA26	I/O
DEVSEL#	A14	I/O
DRDY#	K25	I/O
FLSHBUF#	E2	I
FLUSH#	U24	O
FRAME#	A15	I/O
GTLVREF	G2	I
HIT#	J26	I/O
HITM#	L23	I/O
INIT#	H25	O
IOGNT#/ TESTHI	C7	I
IOREQ#/ TESTLO	A6	I/O
IRDY#	C16	I/O
LOCK#	J25	I/O
MEMACK#	F4	O
MEMREQ#	E1	I
PAR	C13	I/O
PCIRST#	D6	O

Table 14. 82454 KX/GX PB Alphabetical Ball List (352 BGA) (Continued)

Signal	Ball#	Type
PCLK	D25	I/O
PCLKIN	F23	I
PERR#	D13	I/O
PGNT#	E24	I
PLOCK#	B13	I/O
PREQ#	E26	O
PTRDY#	D15	I/O
PWRGD	E3	I
RECVEN	F3	I
REQ0#	M25	I/O
REQ1#	P24	I/O
REQ2#	M26	I/O
REQ3#	N24	I/O
REQ4#	N23	I/O
RESET#	T25	I/O
RP#	J24	I/O
RS0#	G25	I/O
RS1#	F25	I/O
RS2#	F26	I/O
RSP#	R23	I/O
SERR#	A13	O
SMIACT#	H26	O
STOP#	C15	I/O
TCK	F2	I
TDI	G1	I
TDO	H2	O
TESTLO	C5	I
TESTLO	A5	I
TMS	F1	I
TRDY#	L26	I/O
TRST#	G4	I

Table 15. Power, Ground, and No Connect (NC) Signals (352, BGA)

Signal	BGA Ball #
VCC_3[1,3]	D9, D10, D17, D18, J4, K4, U4, V4, AC9, AC10, AC17, AC18, U23, V23, J23, K23
VCC_{PCI}[4]	A18, B5, B7, B12, B15, C22, D11, F24
GND[2]	A1, A2, A3, A4, D14, B1, B2, B3, B4, E4, C1, C2, C3, C4, AC13, D1, D2, D3, D4, D5, N4, P23, A23, A24, A25, A26, B23, B24, B25, B26, C23, C24, C25, C26, D22, D23, D24, D26, E23, AC1, AC2, AC3, AC4, AC5, AB4, AD1, AD2, AD3, AD4, AE1, AE2, AE3, AE4, AF1, AF2, AF3, AF4, AC22, AC23, AC24, AC25, AC26, AB23, AD23, AD24, AD25, AD26, AE23, AE24, AE25, AE26, AF23, AF24, AF25, AF26, G3, M1, R3, W4, AA1, AD6, AE8, AE13, AD13, AF15, AD16, AD17, AE20, AD19, AE22, AD22, AB24, AA24, Y26, W26, V24, T23, R25, R24, L25, K26, K24, G26, G23, E25, [L11, L12, L13, L14, L15, L16, M11, M12, M13, M14, M15, M16, N11, N12, N13, N14, N15, N16, P11, P12, P13, P14, P15, P16, R11, R12, R13, R14, R15, R16, T11, T12, T13, T14, T15, T16][2]
NC, RESERVED	L4, P4, T4, C8, B10, A19, A21, A22, B14, D16, C14, W25, P25, N26, M23, AD10, AC14, AF16, AF17, AF11, AF12, AD11

NOTES:

1. For the BGA package, external VCC_3 balls connect to a VCC_3 "ring" internal to the BGA Package, VCC_3 die pads are bonded to this ring.
2. For the BGA package, the ground pins between brackets are thermal GND pads under the die. External GND balls connect to a GND ring internal to the BGA Package, GND die pads are bonded to this ring.
3. VCC_3 is the 3.3V supply.
4. VCC_{PCI} is the PCI bus voltage.

4.2 Package Information

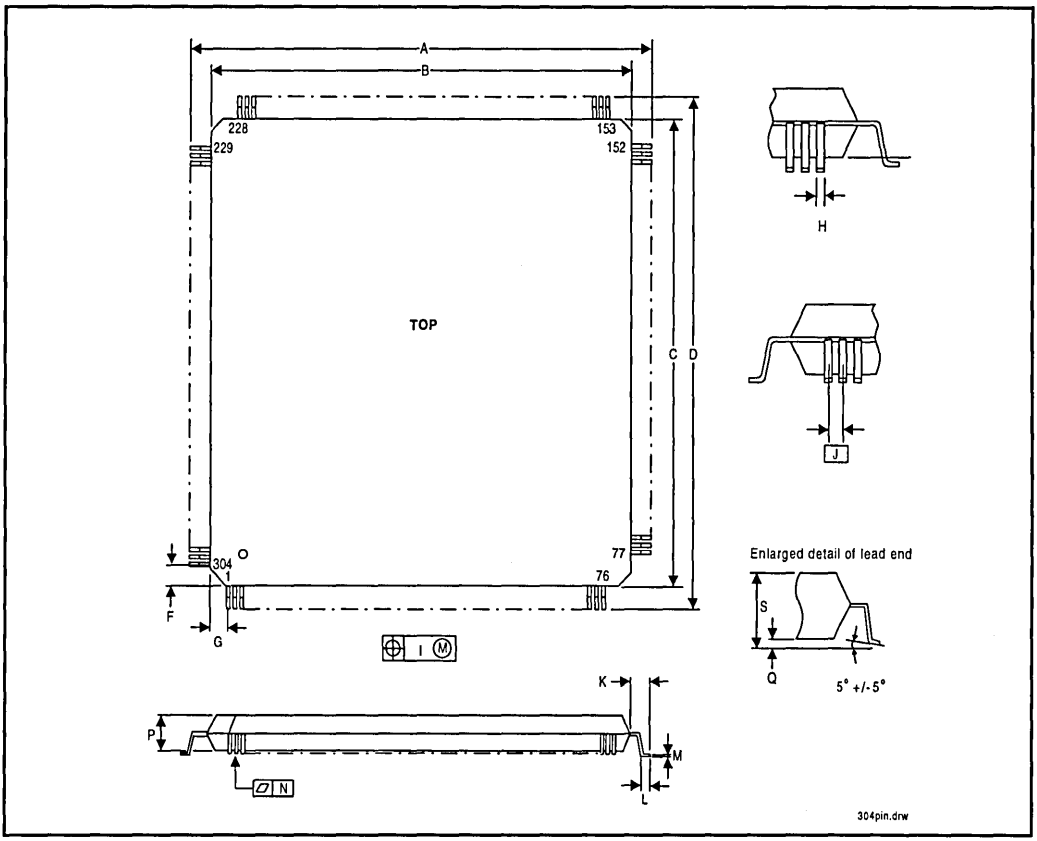

Figure 15. 82454KX/GX PB Package Physical Dimensions (304-Pin QFP)

Table 16. 82454KX/GX PB Package Physical Dimensions (304-Pin QFP)

Item	Millimeters	Inches
A	42.6 ± 0.3	1.677 ±0.012
B	40.0 ± 0.2	1.575 ±0.008
C	40.0 ± 0.2	1.575 ±0.008
D	42.6 ± 0.3	1.677 ±0.012
F	1.25	0.049
G	1.25	0.049
H	0.20 ± 0.10	0.008 ±0.004
I	0.08	0.003

Item	Millimeter	Inches
J	0.5 (T.P.)	0.020 (T.P.)
K	1.3 ± 0.2	0.051 ±0.008
L	0.5 ± 0.2	0.020 ±0.008
M	0.125 ± 0.05	0.005 ±0.002
N	0.10	0.004
P	3.7	0.146
Q	0.4 ± 0.1	0.016 ±.004
S	4.3 max	0.170 max

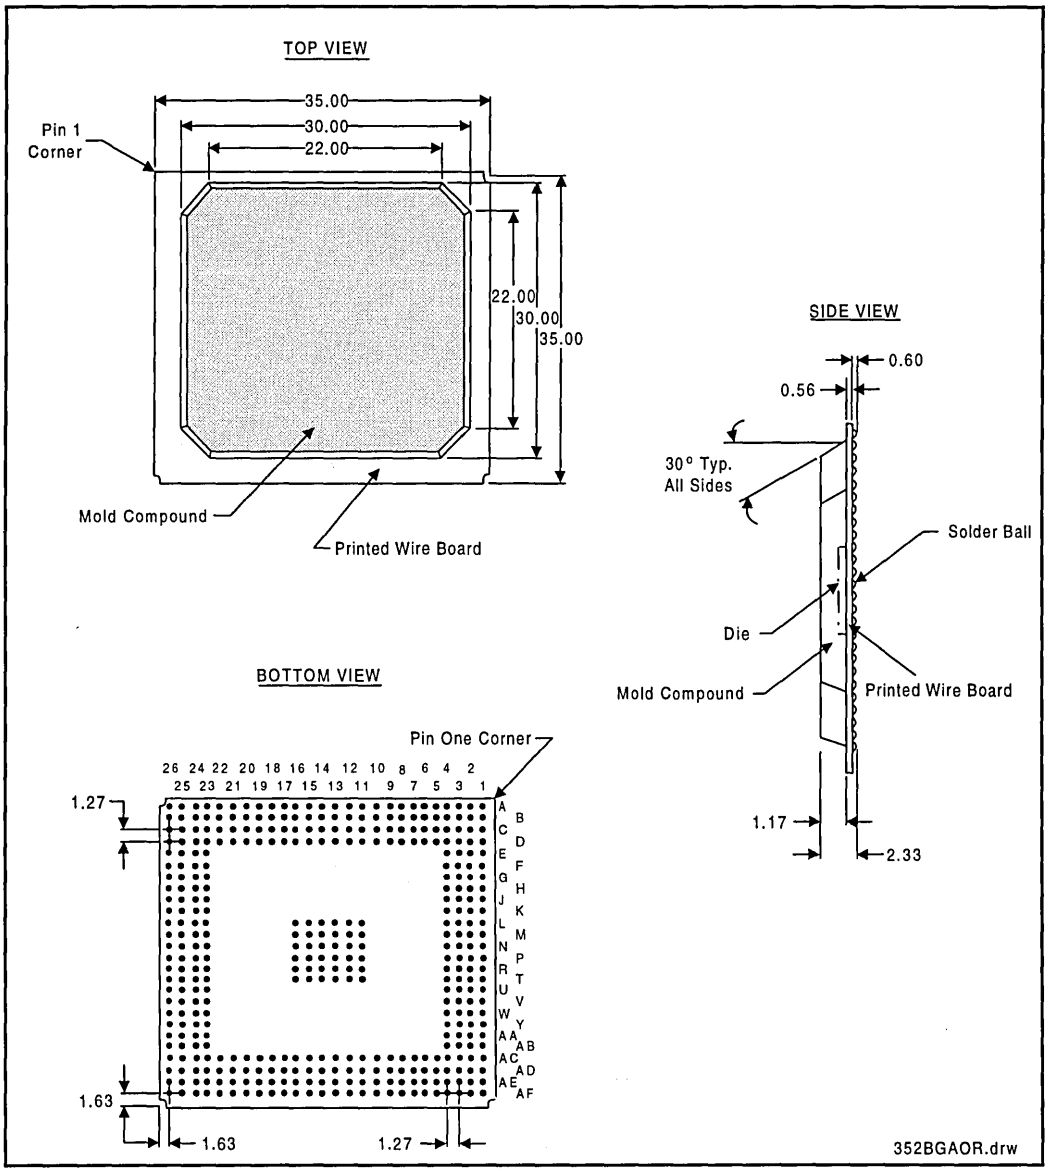

Figure 16. 82454KX/GX PB Package Physical Dimensions (352 BGA)

Chapter 3

Memory Controller (MC)

82453KX/GX DRAM Controller (DC)
82452KX/GX Data Path (DP)
82451KX/GX Memory Interface Component (MIC)

Memory Controller (MC)
82453KX/GX DRAM Controller (DC)
82452KX/GX Memory Data Path (DP)
82451KX/GX Memory Interface Component (MIC)

- Supports Pentium® Pro Processor 60 MHz and 66 MHz Bus Speeds
- Supports 64-Bit Data Bus and 36-Bit Address Bus
- Parity Protection on Control Signals
  - ECC on Data Bus (450GX)
- Dual-Processor Support (450KX)
  - Quad-Processor Support (450GX)
- Eight Deep In-Order Queue
- Four Deep Outbound Request Queue
- Four Cache Line Read Buffer
- Four Cache Line Write Buffer
- GTL+ Bus Driver Technology
- Supports 3.3V and 5V SIMMs
- Read Access, Page Hit 8-1-1-1 (at 66 MHz, 60 ns DRAM)
- Read Access, Page Miss 11-1-1-1 (at 66 MHz, 60 ns DRAM)
- Read Access, Page Miss + Precharge 14-1-1-1 (at 66 MHz, 60 ns DRAM)
- 1 GB Maximum Main Memory (450KX)
  - 4 Gbytes Maximum Main Memory (per 82453GX)
- 2-Way interleaved and Non-Interleaved Memory Organizations (450KX)
  - 4-Way, 2-Way interleaved, and Non-Interleaved Memory Organizations (450GX)
  - Supports Two MCs (450GX) System
- Supports Standard 32 or 36 bit SIMMs
- Supports 72 bit DIMMs
- 4 Mbit, 16 Mbit and 64 Mbit DRAM
- Power Management of Memory Array
- Recovers DRAM Memory Behind Programmable Memory Gaps
- Available in 208-Pin QFP for the DC; 240-Pin QFP or 256-BGA for the DP; 144-Pin QFP for the MIC
- On-Chip Digital PLL
- JTAG Boundary Scan Support

The MC consists of the 82453KX/GX DRAM Controller (DC), the 82452KX/GX Data Path (DP), and four 82451KX/GX Memory Interface Components (MIC). The combined MC uses one physical load on the Pentium Pro processor bus. The memory configuration can be either non-interleaved (450KX/GX), 2-way interleaved (450KX/GX), or 4-way interleaved (450GX only). Both single-sided and double-sided SIMMs are supported at 3.3 and 5 volts. DRAM technologies of 512kx8, 1Mx4, 2Mx8, 4Mx4, 8Mx8, and 16Mx4 at speeds of 50ns, 60ns, 70ns, and 80ns can be used. The maximum memory size is 4 Gbytes for the 4-way interleaved configuration (450GX only), 1 Gbyte (2 Gbytes for the 450GX) for the 2-way interleaved configuration, and 512 Mbytes (1 Gbyte for the 450GX) for the non-interleaved configuration. The MC provides data integrity including ECC in the memory array, and parity on the host bus control signals. The 450GX also provides ECC on the host data bus. The MC is PC compatible. All ISA and EISA regions are decoded and shadowed based on programmable configurations. Regions above 1 Mbyte with size 1 Mbyte or larger that are not mapped to memory may be reclaimed. Three programmable memory gaps can be created. For the 450GX, two MCs can be used in a system.

The Intel 450KX/GX PCIsets may contain design defects or errors known as errata. Current characterized errata are available upon request.

This document describes both the 82454KX and 82454GX PCIsets. Unshaded areas describe features common to the 450KX and 450GX. Shaded areas, like this one, describe the 450GX operations that differ from the 450KX.

PRELIMINARY

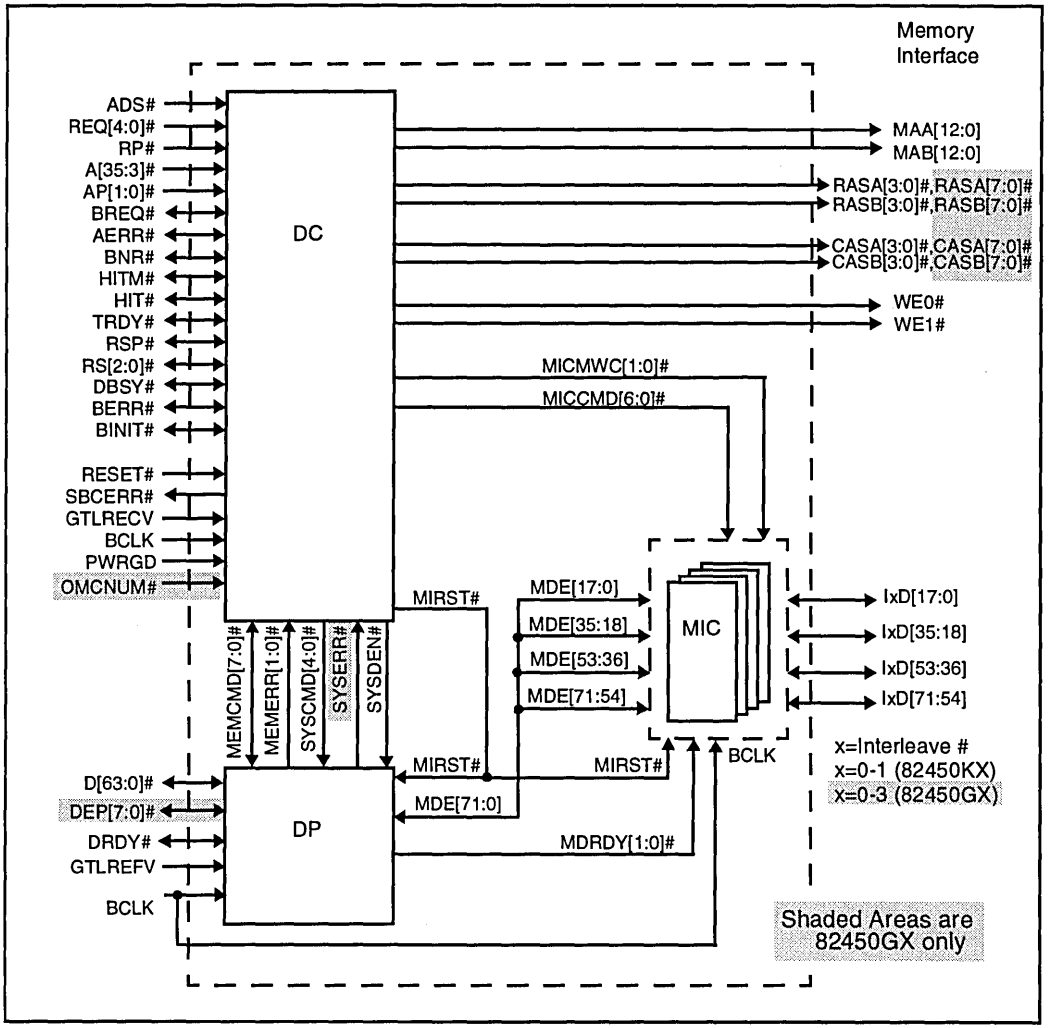

Figure 1. Memory Controller (MC) Simplified Block Diagram

1.0 MC SIGNAL DESCRIPTION

This section contains a detailed description of each signal. The signals are arranged in functional groups according to their interface.

Note that the '#' symbol at the end of a signal name indicates that the active, or asserted state occurs when the signal is at a low voltage level. When '#' is not present after the signal name, the signal is asserted when at the high voltage level.

The terms assertion and negation are used extensively. This is done to avoid confusion when working with a mixture of 'active-low' and 'active-high' signals. The term **assert**, or **assertion** indicates that a signal is active, independent of whether that level is represented by a high or low voltage. The term **negate**, or **negation** indicates that a signal is inactive.

The following notations are used to describe the signal type.

I	*Input* is a standard input-only signal.
O	*Totem* Pole Output is a standard active driver.
I/O	*Input/Output* is bi-directional, tri-state signal.
GTL+	*GTL+* Processor bus signal defined for 1.5V operation.
CMOS	Rail-to-Rail *CMOS* Tolerant to 5V levels.
PCI	Rail-to-Rail CMOS signal specifically for PCI bus connection.
Analog	Reference Voltage.

1.1 DC Signals

Table 1. Host Bus Address/Control Interface Signals (DC)

Signal	Type	Description
A[35:3]#	I GTL+	**ADDRESS BUS.** Upper address bits issued with the current request.
ADS#	I GTL+	**ADDRESS STROBE.** ADS# indicates that the current cycle is the first of two cycles of a request.
AERR#	I/O GTL+	**ADDRESS PARITY CHECK.** Asserted when either an address or request parity error occurs.
AP[1:0]#	I GTL+	**ADDRESS PARITY.** Parity computed over the address; AP1# covers A[35:24]#, and AP0# covers A[23:3]#.
BERR#	I/O GTL+	**BUS ERROR.** BERR# indicates an unrecoverable bus error.
BNR#	I/O GTL+	**BLOCK NEXT REQUEST.** BNR# is asserted by an agent to prevent the request bus owner from issuing further requests.
BREQ0#	O GTL+	**BUS REQUEST 0.** Asserted at reset to set agent IDs in all processors.
DBSY#	I/O GTL+	**DATA BUS BUSY.** DBSY# is used by the data bus owner to hold the data bus for transfers requiring more than one cycle.
HIT#	I/O GTL+	**HIT.** Indicates that a caching agent holds an unmodified version of the requested line. HIT# is also driven in conjunction with HITM# to extend the snoop window.

Table 1. Host Bus Address/Control Interface Signals (DC) (Continued)

Signal	Type	Description
HITM#	I/O GTL+	**HIT MODIFIED.** HITM# indicates that a caching agent holds a modified version of the requested line and that this agent assumes responsibility for providing the line. Also, driven in conjunction with HIT# to extend the snoop window.
REQ[4:0]#	I GTL+	**REQUEST.** In the first cycle of a request these signals carry the request type. In the second cycle they carry the data size and transfer length.
RP#	I GTL+	**REQUEST PARITY.** RP# is even parity that covers REQ[4:0]# and ADS#. RP# is valid on both cycles of the request.
RS[2:0]#	I/O GTL+	**RESPONSE.** RS[2:0]# encode the response to a request.
RSP#	I/O GTL+	**RESPONSE PARITY.** RSP# provides response parity for RS[2:0]#.
TRDY#	I/O GTL+	**TARGET READY.** TRDY# is driven by the target of the data to indicate it is ready to receive data.

Table 2. Memory Address/Control Interface Signals (DC)

Signal	Type	Description
CASA[3:0]# CASB[3:0]# CASA[7:0]# CASB[7:0]#	O CMOS	**COLUMN ADDRESS STROBE (TWO COPIES).** Indicates that the address on MA[12:0] is the column address. There is one CAS# per logical row of memory. Two copies are provided to support external loading.
MAA[12:0] MAB[12:0]	O CMOS	**MEMORY ADDRESS (TWO COPIES).** Multiplexed row and column memory address. Two copies are provided to support external loading.
RASA[3:0]# RASB[3:0]# RASA[7:0]# RASB[7:0]#	O CMOS	**ROW ADDRESS STROBE (TWO COPIES).** Indicates that the address on MA[12:0] is the row address. There is one RAS# per logical row of memory. Two copies are provided to support external loading.
WE0# WE1#	O CMOS	**WRITE ENABLE (TWO COPIES).** Indicates that the current memory request is a write. Two copies are provided to support external loading.

Table 3. DC/DP Interchip Signals (DC)

Signal	Type	Description
MEMCMD[7:0]#	I/O CMOS	**MEMORY SIDE COMMAND.** These signals transfer command and configuration information between the DC and DP.
MEMERR[1:0]#	I CMOS	**MEMORY ERROR.** These signals transfer memory error information from the DP to the DC.

Table 3. DC/DP Interchip Signals (DC) (Continued)

Signal	Type	Description
SYSCMD[4:0]#	O CMOS	**SYSTEM SIDE COMMAND.** These signals send commands and other information from the DC to the DP.
SYSDEN#	O CMOS	**SYSTEM SIDE DATA ENABLE.** This signal permits the DC to control the enabling of DP data information onto the host bus.
SYSERR#	I CMOS	**SYSTEM ERROR.** This signal sends system error data conditions from the DP to the DC.

Table 4. DC/MIC Interchip Signals (DC)

Signal	Type	Description
MICCMD[6:0]#	O CMOS	**MIC COMMAND.** Sends read/write/configuration commands to the MIC.
MICMWC[1:0]#	O CMOS	**MIC MEMORY WRITE COMMAND (TWO COPIES).** Instructs the MIC to drive write data held in its internal buffers on the memory data bus.

Table 5. Reset and Error Signals (DC)

Signal	Type	Description
BINIT#	I/O GTL+	**BUS INITIALIZATION.** BINIT# is asserted to initialize the host bus. Configuration registers are not affected.
MIRST#	O CMOS	**MEMORY INTERFACE RESET.** The DC uses this signal to reset the DP and MIC.
RESET#	I GTL+	**RESET.** This is a hard reset to the DC. The DC sets its internal registers to their default conditions and asserts the MIRST# to the DP and MICs.
SBCERR#	O CMOS	**SINGLE BIT CORRECTED ERROR.** When SBC error reporting is enabled in the MERRCMD or SERRCMD Register, this signal is asserted to indicate that a single bit error was detected and corrected in the memory array.

Table 6. Clock, Power, and Support Signals (DC)

Signal	Type	Description
BCLK	I CMOS	**BUS CLOCK.** This is the input clock for the DC.
GTLREFV	I Analog	**GTL REFERENCE VOLTAGE.** GTLREFV sets the voltage level used by the GTL input receivers for comparison against incoming GTL level signals.
OMCNUM	I CMOS	**MEMORY CONTROLLER NUMBER.** During a power-on reset, this signal provides the MC device number (see MCNUM Register).
PWRGD	I CMOS	**POWER GOOD.** PWRGD is provided by the power supply when all voltages have stabilized for at least 1 ms.

Table 7. Test Signals (DC)

Signal	Type	Description
GTLHI	I/O GTL+	These signals must be tied to V_{TT} using a 10KΩ resistor for proper operation in both test and normal operating modes.
TCK	I CMOS	**JTAG Test Clock.** When TMS is tied low, this signal has no effect on normal operation.
TDI	I CMOS	**JTAG Test Data In.** When TMS is tied low, this signal has no effect on normal operation.
TDO	O CMOS	**JTAG Test Data Out.** When TMS is tied low, this signal has no effect on normal operation.
TESTHI	I/O	**TEST HIGH.** These signals must be tied high using a 10KΩ resistor for proper operation in both test and normal operating modes.
TESTLO	I/O	**TEST LOW.** These signals must be tied low using a 1KΩ resistor for proper operation in both test and normal operating modes.
TMS	I CMOS	**JTAG Test Mode Select.** This signal must be tied low for normal operation.
TRST#	I CMOS	**JTAG Test Reset.** When TMS is tied low, this signal has no effect on normal operation.
RECVEN	I	**RECEIVER ENABLE.** This function is useful for component test. This signal is negated with PWRGOOD to disable GTL+ receivers and tri-state outputs for board test.

1.2 DP Signals

Table 8. Host Bus Interface Signals (DP)

Signal	Type	Description
D[63:0]#	I/O GTL+	**DATA BUS.** The data bus consists of eight bytes.
DEP[7:0]#	I/O GTL+	**DATA ECC/PARITY.** DEP[7:0]# provides ECC for the D[63:0]# signals. ECC is computed over the 64 data bits. Parity is not generated or checked by the MC.
DRDY#	I/O GTL+	**DATA READY.** Asserted for each cycle that data is transferred.

Table 9. Data Path Interface Signals (DP)

Signal	Type	Description
MDE[71:0]	I/O CMOS	**MEMORY DATA AND ECC.** Common to all types and sizes of memory supported, these signals include the 64 bits of data and 8 ECC check bits. ECC is computed over 64-bit data words. Parity is computed as byte-parity over a 64-bit word.
MDRDY0# MDRDY1#	O CMOS	**MEMORY DATA READY (TWO COPIES).** Asserted when write data on the MDE bus is valid. Two copies are provided to support external loading.

Table 10. DC/DP Interchip Signals (DP)

Signal	Type	Description
MEMCMD[7:0]#	I/O CMOS	**MEMORY SIDE COMMAND.** These signals transfer command and configuration information between the DC and DP.
MEMERR[1:0]#	O CMOS	**MEMORY ERROR.** These signals transfer memory error information from the DP to the DC.
SYSCMD[4:0]#	I CMOS	**SYSTEM SIDE COMMAND.** These signals send commands and other information from the DC to the DP.
SYSDEN#	I CMOS	**SYSTEM SIDE DATA ENABLE.** This signal permits the DC to control the enabling of DP data information onto the host bus.
SYSERR#	O CMOS	**SYSTEM ERROR.** This signal sends system error data conditions from the DP to the DC.

Table 11. Clock, Power, Reset Signal (DP)

Signal	Type	Description
BCLK	I CMOS	**BUS CLOCK.** This is the clock input for the DP.
GTLREFV	I Analog	**GTL REFERENCE VOLTAGE.** GTLREFV sets the voltage level used by the GTL input receivers for comparison against incoming GTL level signals.
MIRST#	I CMOS	**MEMORY INTERFACE RESET.** This signal is driven by MIRST# signal from the DC.

Table 12. Test Signals (DP)

Signal	Type	Description
TCK	I CMOS	**JTAG Test Clock.** When TMS is tied low, this signal has no effect on normal operation.
TDI	I CMOS	**JTAG Test Data In.** When TMS is tied low, this signal has no effect on normal operation.
TDO	O CMOS	**JTAG Test Data Out.** When TMS is tied low, this signal has no effect on normal operation.
TESTHI	I/O	**TEST HIGH.** These signals must be tied high using a 10KΩ resistor for proper operation in both test and normal operating modes.
TESTLO	I/O	**TEST LOW.** These signals must be tied low using a 1KΩ resistor for proper operation in both test and normal operating modes.
TMS	I CMOS	**JTAG Test Mode Select.** This signal must be tied low for normal operation.
TRST#	I CMOS	**JTAG Test Reset.** When TMS is tied low, this signal has no effect on normal operation.

1.3 MIC Signals

Table 13. Control Interface Signals (MIC)

Signal	Type	Description
MICCMD[6:0]#	I CMOS	**MIC COMMAND.** These signals send read/write/configuration commands to the MIC.
MICMWC#	I CMOS	**MIC MEMORY WRITE COMMAND.** This signal instructs the MIC to drive write data held in its internal buffers on the memory data bus.

Table 14. Data Path Interface Signals (MIC)

Signal	Type	Description
MDE[17:0]	I/O CMOS	**MEMORY DATA AND ECC.** MDE[17:0] is one fourth of a Qword and is connected to one of the word portions of the DP memory data bus. ECC is computed over 64-bit data words. MDE[17:0] is one fourth of a Qword.
MDRDY#	I CMOS	**MEMORY DATA READY.** The DP asserts this signal to the MIC when data on the MDE bus is valid.

Table 15. Memory Interface Signals (MIC)

Signal	Type	Description
I0D[17:0]	I/O CMOS	**MEMORY DATA.** I0D[17:0] is one fourth of a QWord that is connected to interleave zero of the memory.
I1D[17:0]	I/O CMOS	**MEMORY DATA.** I1D[17:0] is one fourth of a QWord that is connected to interleave one of the memory.
I2D[17:0]	I/O CMOS	**MEMORY DATA.** I2D[17:0] is one fourth of a QWord that is connected to interleave two of the memory.
I3D[17:0]	I/O CMOS	**MEMORY DATA.** I3D[17:0] is one fourth of a QWord that is connected to interleave three of the memory.

Table 16. Clock and Reset Signals (MIC)

Signal	Type	Description
BCLK	I CMOS	**BUS CLOCK.** This is the clock input to the device.
MIRST#	I CMOS	**MEMORY INTERFACE RESET.** This signal is connected to the MIRST# signal on the DC.

Table 17. Test Signals (MIC)

Signal	Type	Description
TCK	I CMOS	**JTAG Test Clock.** When TMS is tied low, this signal has no effect on normal operation.
TDI	I CMOS	**JTAG Test Data In.** When TMS is tied low, this signal has no effect on normal operation.
TDO	O CMOS	**JTAG Test Data Out.** When TMS is tied low, this signal has no effect on normal operation.
TESTLO	I/O	**TEST LOW.** These signals must be tied low using a 1KΩ resistor for proper operation in both test and normal operating modes.
TMS	I CMOS	**JTAG Test Mode Select.** This signal must be tied low for normal operation.
TRST#	I CMOS	**JTAG Test Reset.** When TMS is tied low, this signal has no effect on normal operation.

1.4 Signal State During Reset

Table 18 shows the state of all MC output and bi-directional signals during a hard reset (RESET# and MIRST# asserted).

Table 18. Signal State During Reset

Signal	State
DC	
AERR#	Not Driven
BNR#	Not Driven
BERR#	Not Driven
BINIT#	Not Driven
BREQ0#	Not Driven
CASAx#	High
CASBx#	High
DBSY#	Not Driven
HIT#	Not Driven
HITM#	Not Driven
MAA[12:0]	Not Driven
MAB[12:0]	Not Driven
MEMCMD[7:0]#	High
MICCCMD[6:0]#	High
MICMWC[1:0]#	High
MIRST#	Low
RASAx#	High
RASBx#	High
RS[2:0]#	Not Driven
RSP#	Not Driven
SBCERR#	High

Table 18. Signal State During Reset (Continued)

Signal	State
SYSCMD[4:0]#	High
SYSDEN#	High
TDO	Hi-Z during TRST#
TRDY#	Not Driven
WEx#	High
DP	
D[63:0]#	Not Driven
DEP[7:0]#	Not Driven
DRDY#	Not Driven
MDE[71:0]	Not Driven
MDMERR[1:0]#	High
MDRDY[1:0#]	High
MEMCMD[7:0]#	High
SYSERR#	High
TDO	Hi-Z during TRST#
MIC	
I0D[17:0]	Low
I1D[17:0]	Low
I2D[17:0]	Low
I3D[17:0]	Low
MDE[71:0]	Low
TDO	Hi-Z during TRST#

2.0 MC REGISTER DESCRIPTION

The MC contains two sets of registers (I/O space registers and configuration registers) that are accessed via the host CPU I/O address space. Accessing MC configuration registers uses the same procedure as is defined for accessing PCI device configuration registers (e.g., for the PB). The required PCI Header register set is provided permitting the MC to respond to initialization software in the same manner as an actual PCI device (i.e., the MC register set is PCI compliant). In some cases, (e.g., PCICMD register), the register may only be appropriate for a device attached to the PCI bus. In these cases the register is shown as a read only register with the bits hardwired appropriately and writes have no effect.

The I/O space registers (CONFADD and CONFDATA Registers) provide indirect access to the MC configuration registers. The MC's internal registers can be accessed as Byte, Word (16-bit), or Dword (32-bit) quantities, with the exception of CONFADD which can only be accessed as a Dword. 0-length I/O reads are not supported by the MC. The following nomenclature is used for access attributes.

RO *Read Only.* If a register is read only, writes to this register have no effect.
R/W *Read/Write.* A register with this attribute can be read and written.
R/WC *Read/Write Clear.* A register bit with this attribute can be read and written. However, a write of a 1 clears (sets to 0) the corresponding bit and a write of a 0 has no effect.

Some of the MC registers contain reserved bits. 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, unless otherwise specified in the individual register descriptions, 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 new values for other bit positions and written back. In some cases, software must program reserved bit positions to a particular value. This value is defined in the individual bit descriptions.

In addition to reserved bits within a register, the MC contains address locations in the MC's configuration space that are marked "Reserved". The MC responds to accesses to these address locations by completing the host transaction. Software should not write to reserved MC configuration locations in the device-specific region (above address offset 3Fh).

2.1 Initialization and Configuration

After a power-on reset, the type of memory assumed is non-interleaved. The default DRAM timing values for the non-interleaved memory are set to reasonable values for 66 MHz and slow memory. It is expected that the BIOS will change these default memory timing values as necessary before memory is accessed. For the 450KX, the memory base address is hardwired to zero. For the 450GX, each MC memory base address is determined at reset according to its controller number (strapping option on the OMCNUM signal). After the power-on initialization, system BIOS determines memory size and configuration and programs the configuration registers accordingly. After a power-on reset, the MC is set for a 4 Mbyte main memory size. The 512 Kbyte DOS RAM region and accesses from 1 Mbyte to 4 Mbytes are enabled (MC accepts accesses). Until initialized, the MC does not respond to any other memory locations.

The MC and PB residing on the host bus contain a configuration space that is compliant with the configuration space in the PCI bus specification. While the MC is not a true PCI device, it uses the same configuration register access mechanism. The VID Register (00–01h) and the DID Register (02–03h) both return legitimate values.

The MC has two registers located in I/O Space—The Configuration Address (CONFADD) Register and the Configuration Data (CONFDATA) Register. The compatibility PB is the only response agent for host accesses to CONFADD and the MC snoops writes to this register. CONFADD is first written to select the MC. A

subsequent read or write of CONFDATA causes the MC to generate the requested configuration cycle on the host bus. During accesses to CONFDATA, only the host device selected by CONFADD responds. If no valid device is selected in CONFADD, a subsequent CONFDATA transaction is treated like a normal I/O transaction to I/O address CFCh. The MC only snoops 32-bit writes to the CONFADD Register. All other access sizes to the CF8h location are ignored by the MC. .

I/O Address	Mnemonic	Register Name	Access
CF8h	CONFADD	Configuration Address	WO
CFCh	CONFDATA	Configuration Data	R/W

There can be up to two MC's in the system. After power-on reset, the MC designated as controller number 0 (Device Number 10100 in the CDNUM Register) is the default response agent for accesses to the DOS RAM region (0–512 Kbytes). The MC controller number is determined at reset by latching an ID into the MC from the OMCNUM pin.

Both PBs default to Bus Number 0 after power-on reset. The MC is hard coded to Bus Number 0. Each PB and MC must have a unique Device Number assigned at power-on reset. The relation between the PB and MC number and the Device Number that is assigned is shown in Table 19.

Table 19. PCI Device Numbers for Bus Number 0

Device	Unique ID loaded at Reset	PCI Device Number
PB	01	11001
PB	10	11010
MC	0	10100
MC	1	10101
Reserved	—	All other

2.2 I/O Space Registers

The MC has two registers located in I/O Space—the Configuration Address (CONFADD) Register, and the Configuration Data (CONFDATA) Register.

The CONFADD and CONFDATA Registers provide a window into the MC's configuration space registers. The specific device, and register are selected by writing to the CONFADD Register. Data is read from or written to the selected register by accessing the CONFDATA Register. Note that the CONFADD Register is only selected by DWord accesses to CF8h. The CONFDATA Register is not selected unless configuration accesses are enabled in the CONFADD Register. This allows the CONFDATA Register to overlap other registers.

Table 20. I/O Space Registers

I/O Address	Mnemonic	Register Name	Access
CF8h	CONFADD	Configuration Address	R/W
CFCh	CONFDATA	Configuration Data	R/W

2.2.1 CONFADD—CONFIGURATION ADDRESS REGISTER

Address Offset: 0CF8h (Dword access only)
Default: 00000000h
Attribute: Write Only

The CONFADD Register contains the Bus Number, Device Number, Function Number, and Register Number for which a subsequent configuration access is intended. For example, a write of 8000A0ACh to the CONFADD register can be used to access the 82453KX Memory Timing Register.

Bits	Description
31	**Configuration Enable (CSE).** 1=Enable. 0=Disable.
30:24	**Reserved.** Write 0's.
23:16	**Bus Number (BUSNUM).** The MC bus number is 00h.
15:11	**Device Number (DEVNUM).** The device number is reported in the CDNUM Register (offset 49h). **450KX:** DEVNUM=10100 (hardwired) **450GX:** The device number for an MC is determined at power-up and is reported in the CDNUM Register (offset 49h). DEVNUM=10100, or 10101 (depends on OMCNUM signal strapping).
10:8	**Function Number (FUNCNUM).** The MC responds to configuration cycles with a function number of 000b. All other function number values attempting access to the MC (BUSNUM=00h and DEVNUM matching MC device number) are ignored by the MC.
7:2	**Register Number (REGNUM).** This field selects one register within MC, if the MC configuration space is selected by the other fields in the CONFADD Register.
1:0	**Reserved.** Write 0's.

2.2.2 CONFDATA—CONFIGURATION DATA REGISTER

Address Offset: 0CFCh
Default: 00000000h
Attribute: Read/Write

CONFDATA is a 32-bit read/write window into configuration space. The portion of configuration space that is reference by CONFDATA is determined by the contents of CONFADD. The BE signals determine which bytes get written..

Bits	Description
31:0	**Configuration Data Window.** If bit 31 of CONFADD=1, an access to CONFDATA I/O space is mapped to configuration space using the contents of CONFADD.

2.3 MC Configuration Registers

Table 21. MC Configuration Registers

Address Offset	Mnemonic	Register Name	Access
00–01h	VID	PCI Vendor Identification	RO
02–03h	DID	PCI Device Identification	RO
04–05h	PCICMD	PCI Command	RO
06–07h	PCISTS	PCI Status	RO
08h	RID	Revision Identification	RO
09–0Bh	CLASSC	Class Code	RO
0C–3Fh	—	Reserved	—
40–43h	—	Reserved (450KX)	—
	BASEADD	MC Base Address (450GX)	R/W
44–48h	—	Reserved	—
49h	CDNUM	Controller Device Number	RO
4A–4Bh	—	Reserved	—
4C–4Fh	CMD	MC Command	R/W
50–56h	—	Reserved	—
57h	SMME	SMRAM Enable	R/W
58h	VBRE	Video Buffer Region Enable	R/W
59–5Fh	PAM[0:6]	Programmable Attribute Map (7 registers)	R/W
60-63h	DRL[0:3]	DRAM Row Limit for rows 0-3	R/W
64–67h	—	Reserved (450KX)	—
	DRL[4:7]	DRAM Row Limit for rows 4-7 (450GX)	R/W
70–73h	—	Reserved	—
74–77h	SBCERRADD	First Single Bit Correctable Error Address	R/W
78–79h	MG	Memory Gap Register	R/W
7A–7Bh	MGUA	Memory Gap Upper Address	R/W
7C–7Fh	LMG	Low Memory Gap Register	R/W
80–87h	—	Reserved	R/W
88–8Bh	HMGSA	High Memory Gap Start Address	R/W
8C–8Fh	HMGEA	High Memory Gap End Address	R/W
90–A3h	—	Reserved	—
A4–A7h	APICR	I/O APIC Range Register	R/W

Table 21. MC Configuration Registers (Continued)

Address Offset	Mnemonic	Register Name	Access
A8–ABh	UERRADD	First Uncorrectable Error Address	RO
AC–AFh	MEMTIM	Memory Timing	R/W
B0–B7h	—	Reserved	—
B8–BBh	SMMR	SMRAM Range Register	R/W
BCh	HBIOSR	High BIOS Range Register	R/W
BD–BFh	—	Reserved	—
C0–C1h	MERRCMD	Memory Error Reporting Command Register	R/W
C2–C3h	MERRSTS	Memory Error Status Register	R/W
C4–C5h	SERRCMD	System Error Reporting Command Register	R/W
C6–C7h	SERRSTS	System Error Status Register	R/W
C8–FFh	—	Reserved	—

2.3.1 VID—VENDOR IDENTIFICATION REGISTER

Address Offset: 00–01h
Default: 8086h
Attribute: Read Only

The VID Register contains the vendor identification number.

Bits	Description
15:0	**PCI Vendor ID.** This 16-bit value (8086) is assigned to Intel.

2.3.2 DID—DEVICE IDENTIFICATION REGISTER

Address Offset: 02–03h
Default: 84C5h
Attribute: Read Only

The DID Register contains the PCI device identification.

Bits	Description
15:0	**PCI Device ID.** This 16-bit value (84C5) is assigned to the MC.

2.3.3 PCICMD—PCI COMMAND REGISTER

Address Offset: 04–05h
Default: 0000h
Attribute: Read Only

The PCICMD register controls the device's ability to respond to PCI cycles. This register location is provided to maintain PCI Header Register set compliance.

Bits	Function
15:0	**Reserved.** When read, this register returns 0000h. Writes have no effect.

2.3.4 PCISTS—PCI STATUS REGISTER

Address Offset: 06–07h
Default: 0080h
Attribute: Read Only

The PCISTS register reports various PCI error conditions and indicates the device's response to various transactions on the PCI Bus. This register location is provided to maintain PCI Header Register set compliance.

Bits	Function
15:0	**Reserved.** When read, this register returns 0080h indicating that the MC is capable of fast back-to-back cycles. Writes have no effect.

2.3.5 RID—REVISION IDENTIFICATION REGISTER

Address Offset: 08h
Default: See stepping information in the 450GX/KX Specification Update.
Attribute: Read Only

This register contains the revision number of the MC.

Bits	Function
7:0	**Revision Identification Value.** This is an 8-bit value that indicates the revision identification number for the MC.

2.3.6 CLASSC—CLASS CODE REGISTER

Address Offset: 09–0Bh
Default: 050000h
Attribute: Read Only

This register contains the device programming interface information related to the Sub-Class Code and Base Class Code definition for the MC. This register also identifies the function sub-class in relation to the Base Class Code and the Based Class Code.

Bits	Description
23:16	**Base Class Code (BCC).** 05h=Memory controller device.
15:8	**Sub-Class Code (SCC).** 00h=RAM device.
7:0	**Programming Interface (PI).** 00h=Standard register format.

2.3.7 BASEADD—MC BASE ADDRESS REGISTER (450GX ONLY)

Address Offset: 40–43h
Default: 0000 0000 0000 0000 000x 0000 0000 0000b (x=MC Device Number)
Attribute: Read/Write

This register sets the base memory address for the MC. Since there can be up to two MCs in a system, the default base address is determined at power-up (hard reset), assuming 4 Gbytes of main memory behind each MC. The strapping option on the OMCNUM pin determines the MC's device number. This number affects A[33:32]# (bits[13:12]) of the default base address.

Bits	Description
31:16	**Reserved.**
15:0	**MC Main Memory Base Address (In units of 1 Mbyte).** Bits [15:0] correspond to address bits A[35:20]#.

2.3.8 CDNUM—CONTROLLER DEVICE NUMBER REGISTER

Address Offset: 49h
Default: 00010100 (450KX)
0001010x (x=loaded at reset)
Attribute: Read Only

This register contains the MC device number. This value is hard coded in the 450KX.

This number is loaded from the OMCNUM pin. This pin must be strapped (VCC3 or GND) appropriately. MC Number 0 defaults to responding to DOS accesses (0–512 Kbytes). MC Number 1 defaults to this range disabled (ignore these accesses).

Bits	Description
7:5	**Reserved.**
4:1	**Fixed value.** The upper four bits of the MC Device Number are always 1010 for the MC.
0	**450KX: Fixed value.** The lower bit of the MC Device Number is always 0. **450GX: MC Device Number.** The lower bit of the MC device number is encoded as follows: **Bits[1:0] Function** 0 MC #0 1 MC #1

2.3.9 CMD—COMMAND REGISTER

Address Offset: 4C–4Fh
Default: 0000 0000 0000 0000 x000 1000 0000 1011b
(x=captured from address bus on reset).
Attribute: Read/Write

This register controls DRAM configuration, various memory controller operations, and reports the in-order queue depth.

Bits	Description
31:16	**Reserved.**
15	**In-order Queue Depth 1 Select.** 1=Depth of 1. 0=Depth of 8. Value captured from A7#.
14:11	**Active Interleaves.** 1=Active. 0=Inactive. Each bit enables/disables a single interleave. For example, bits[14:11]=0101 defines a 2-way interleaved system with interleaves 2 and 0 active. Default is Bits[14:11]=0001. **Bit** **Active Interleave (450KX)** **Active Interleave (450GX)** 14 Reserved 3 13 Reserved 2 12 1 1 11 0 0 For a 4-way interleaved system that is functionally reduced to a non-interleaved or 2-way interleaved system, these bits define the currently active interleaves out of the four possible interleaves. For a 4-way interleave, all bits must be set to 1.

<table>
<tr><th>Bits</th><th>Description</th></tr>
<tr><td>10</td><td>Page Open Policy. 1=Hold page open. 0=Close page (default).</td></tr>
<tr><td>9</td><td>Common CAS# Enable. 1=Enable (common CAS# signals). 0=Disable (independent CAS# signals) (default). When enabled, this bit indicates when adjacent rows of memory have a common CAS# connection. (This is typical of double-sided byte parity SIMMs.)</td></tr>
<tr><td>8</td><td>Extended Read-Around-Write Enable. 1=Enable. 0=Disable (default). When enabled, reads by-pass writes within the MICs, provided their addresses do not match. Note, this does not change the order in which data transfers occur on the host bus. This bit can not be enabled unless Read-around-write is also enabled in bit 7.</td></tr>
<tr><td>7</td><td>Read-Around-Write Enable. 1=Enable. 0=Disable (default). When enabled, reads by-pass writes within the MC, provided their addresses do not match. Note, this does not change the order in which data transfers occur on the host bus.</td></tr>
<tr><td>6</td><td>Memory Address Bit Permuting Enable. 1=Enable. 0=Disable (default). When enabled, the MC permutes memory addresses to obtain alternate memory row selection bits. Note that when enabled, there must be a power of 2 number of rows, all rows must be the same size, and all populated rows must be adjacent and start at row 0.</td></tr>
<tr><td>5</td><td>Reserved.</td></tr>
<tr><td>4:3</td><td>Memory Configuration. The memory configuration is as follows:
<table>
<tr><th>Bit[4:3]</th><th>Function (450KX)</th><th>Function (450GX)</th></tr>
<tr><td>00</td><td>Reserved</td><td>Reserved</td></tr>
<tr><td>01</td><td>Non-interleaved (default)</td><td>Non-interleaved (default)</td></tr>
<tr><td>10</td><td>2-way interleaved</td><td>2-way interleaved</td></tr>
<tr><td>11</td><td>Reserved</td><td>4-way interleaved</td></tr>
</table></td></tr>
<tr><td>2:0</td><td>Read Burst Delay. This field selects the number of delay cycles between data QWords in a read burst.
• For a non-interleaved memory configuration, the number of delay cycles can be set to 3, 4, 5, or 6 clocks (Read Burst Delay = (T_{RCAS} + T_{CP}) -1).
• For a 2-way interleaved configuration, the number of delay cycles can be set to 1, 2, or 3 clocks (Read Burst Delay = ((T_{RCAS} + T_{CP})/2) - 1 and (T_{RCAS} + T_{CP}) must be even).
• For a 4-way interleaved configuration (450GX), the number of delay cycles must be set to 0 clocks.
<table>
<tr><th>Bits[2:0]</th><th>Delay Cycles</th><th>Bits[2:0]</th><th>Delay Cycles</th></tr>
<tr><td>000</td><td>0 (burst rate=1,1,1)</td><td>100</td><td>4 (burst rate=5,5,5)</td></tr>
<tr><td>001</td><td>1 (burst rate=2,2,2)</td><td>101</td><td>5 (burst rate=6,6,6)</td></tr>
<tr><td>010</td><td>2 (burst rate=3,3,3)</td><td>110</td><td>6 (burst rate=7,7,7)</td></tr>
<tr><td>011</td><td>3 (burst rate=4,4,4) (default)</td><td>111</td><td>7 (burst rate=8,8,8)</td></tr>
</table></td></tr>
</table>

2.3.10 SMME—SMRAM ENABLE REGISTER

Address Offset: 57h
Default: 00h
Attribute: Read/Write

This register enables/disables the SMM range specified in the SMMR Register.

Bits	Description
7:4	**Reserved.**
3	**SMM RAM Enable.** 1=Enable. 0=Disable.
2:0	**Reserved.**

2.3.11 VBRE—VIDEO BUFFER REGION ENABLE REGISTER

Address Offset: 58h
Default: 00h
Attribute: Read/Write

This register enables/disables the video buffer area.

Bits	Description
7:2	**Reserved.**
1	**Video Buffer Area Enable (A0000–BFFFFh).** 1=Enable. 0=Disable.
0	**Reserved.**

2.3.12 PAM[0:6]—PROGRAMMABLE ATTRIBUTE MAP REGISTERS

Address Offset: PAM0 (59h)—PAM6 (5Fh)
Default: PAM0=03h, PAM[1:6]=00h
Attribute: Read/Write

These seven registers select read only (RE=1, WE=0), write only (RE=0, WE=1), or read/write (RE=1, WE=1) access attributes for 14 memory regions between the 512 Kbyte and 1 Mbyte address range. The individual memory regions can also be disabled (RE=0, WE=0). Each register controls two regions; bits [7:4] control one region and bits [3:0] control the other region. Note that the default for the 512–640 Kbyte region is read/write enabled. The default for all other regions is read/write disabled.

Note that the PB has corresponding PAM registers. Only one device (MC/PB) should have the same space enabled at one time to avoid access conflicts

Read Enable (RE) When RE=1 (enabled), CPU read accesses to the corresponding memory region are directed to main memory. When RE=0 (disabled), CPU read accesses are ignored.

Write Enable (WE) When WE=1 (enabled), CPU write accesses to the corresponding memory region are directed to main memory. When WE=0 (disabled), CPU write accesses are ignored.

PAM	Attribute Bits				Memory Segment	Comments	Offset
Register	7,3	6,2	5,1	4,0			
PAM0[7:4]	Reserved		WE	RE	0F0000–0FFFFFh	BIOS	59h
PAM0[3:0]	Reserved		WE	RE	080000–09FFFFh	512–640 KB	59h
PAM1[7:4]	Reserved		WE	RE	0C4000–0C7FFFh	ISA Expansion	5Ah
PAM1[3:0]	Reserved		WE	RE	0C0000–0C3FFFh	ISA Expansion	5Ah
PAM2[7:4]	Reserved		WE	RE	0CC000–0CFFFFh	ISA Expansion	5Bh
PAM2[3:0]	Reserved		WE	RE	0C8000–0CBFFFh	ISA Expansion	5Bh
PAM3[7:4]	Reserved		WE	RE	0D4000–0D7FFFh	ISA Expansion	5Ch
PAM3[3:0]	Reserved		WE	RE	0D0000–0D3FFFh	ISA Expansion	5Ch
PAM4[7:4]	Reserved		WE	RE	0DC000–0DFFFFh	ISA Expansion	5Dh
PAM4[3:0]	Reserved		WE	RE	0D8000–0DBFFFh	ISA Expansion	5Dh
PAM5[7:4]	Reserved		WE	RE	0E4000–0E7FFFh	BIOS Extension	5Eh
PAM5[3:0]	Reserved		WE	RE	0E0000–0E3FFFh	BIOS Extension	5Eh
PAM6[7:4]	Reserved		WE	RE	0EC000–EFFFFh	BIOS Extension	5Fh
PAM6[3:0]	Reserved		WE	RE	0E8000–0EBFFFH	BIOS Extension	5Fh

2.3.13 DRL—DRAM ROW LIMIT (0 TO 7)

Address Offset: 60–6Fh
Default: 0001h
Attribute: Read/Write

The 450KX MC supports 4 rows of DRAM. DRL[0:3] define the upper and lower addresses for each DRAM row. The addresses are relative to the memory space of the MC and do not include any programmed memory gaps. The contents of these 16-bit registers represent the boundary addresses in 4 Mbyte granularity. Rows with no memory are programmed with the upper limit of the previous row. The default after reset reflects the requirement that the first row be populated.

Note that for the 450KX, DRL[4:7] must be programmed with the value programmed in DRL3.

The 450GX supports 8 rows of DRAM. DRL[0:7] define the upper and lower addresses for each DRAM row.

DRL0 = Total memory in row 0 (in 4 MB)
DRL1 = Total memory in row 0 + row 1 (in 4 MB)
DRL2 = Total memory in row 0 + row 1 + row 2 (in 4 MB)
DRL3 = Total memory in row 0 + row 1 + row 2 + row 3 (in 4 MB)
DRL4 = Total memory in row 0 + row 1 + row 2 + row 3 +row 4 (in 4 MB)
DRL5 = Total memory in row 0 + row 1 + row 2 + row 3 +row 4 +row 5 (in 4 MB)
DRL6 = Total memory in row 0 + row 1 + row 2 + row 3 +row 4 +row 5 + row 6 (in 4 MB)
DRL7 = Total memory in row 0 + row 1 + row 2 + row 3 +row 4 +row 5 + row 6 +row 7 (in 4 MB)

Bits	Description
15:7	**Reserved.**
6:0	**DRAM Row Upper Address Limit (in units of 4 Mbytes).** This 7-bit field determines the upper address limit of a particular row (i.e., DRL minus previous DRL = row size). Note that the allowed minimum and maximum values for row size depend on the memory configuration (non-interleaved row size = 4 Mbytes minimum and 128 Mbytes maximum; 2-way interleaved row size = 8 Mbytes minimum and 256 Mbytes maximum). **Example:** Row 0 has 16 Mbytes, row 1 is unpopulated, and row 2 has 32 Mbytes. The DRL Registers would be programmed as follows: DRL0 = 04h DRL1 = 04h DRL[2:7]= 0Ch (each register)

Bits	Description
15:11	**Reserved.**
10:0	**DRAM Row Upper Address Limit (in units of 4 Mbytes).** This 11-bit field determines the upper address limit of a particular row (i.e., DRL minus previous DRL = row size). Note that the allowed minimum and maximum values for row size depend on the memory configuration (non-interleaved row size = 4 Mbytes minimum and 128 Mbytes maximum; 2-way interleaved row size = 8 Mbytes minimum and 256 Mbytes maximum; 4-way interleaved = 16 Mbytes minimum and 512 Mbytes maximum). **Example:** In a 4-way interleaved configuration with all rows containing 512 Mbytes, the DRL Registers would be programmed as follows: DRL0 = 80h DRL4 = 280h DRL1 = 100h DRL5 = 300h DRL2 = 180h DRL6 = 380h DRL3 = 200h DRL7 = 400h

2.3.14 SBCERRADD—SINGLE BIT CORRECTABLE ERROR ADDRESS REGISTER

Address Offset: 74–77h
Default: 0000h
Attribute: Read Only

This register provides the effective address of the memory access that caused a single bit correctable error on the memory side. The value in this register is only valid if the SBC correctable error bit is set in the MERRSTS Register (C2–C3h).

Bits	Description
31:3	**Address of First Single-bit Correctable ECC Error.** This is the effective DRAM address used in the MC and must be converted back to the original physical address by software. The programmed memory gaps and MC base address must be taken into account for proper calculation of the address in which the DRAM side ECC error occurred.
2:1	**QWord Number Error Detect.** When a single-bit error occurred in a transfer from the DRAM array, this field indicates which QWord in the transfer contained the error. Note that this field reports the QWord number relative to the order of the transfer (0 to 3), even if the transfer does not begin with the first QWord of a cache line. In addition, in a single QWord transfer, if an error is detected, this field will be set to 00. For further granularity in determining the address of the error, bit 0 of this register reports which half of the QWord contained the error. **Bits [2:1] — QWord Number of the Transfer** 00 — First QWord Transferred (QWord 0) 01 — Second QWord Transferred (QWord 1) 10 — Third QWord Transferred (QWord 2) 11 — Fourth QWord Transferred (QWord 3)
0	**DWord Number Error Detect (In a QWord).** 1=Error in upper Dword. 0=Error in lower Dword. When a single-bit error occurred in a transfer from the DRAM array, this bit identifies which half of the QWord (upper 2 Words or lower 2 Words) contains the error. Single-bit Error Correcting of Memory Data must enabled (bit 2 of MERRCMD Register, C0–C1h) to obtain this functionality.

2.3.15 MG—MEMORY GAP REGISTER

Address Offset: 78–79h
Default: 0010h
Attribute: Read/Write

This register, along with the MGUA Register, defines the Memory Gap. Note that the Memory Gap must be located above the Low Memory Gap and below the High Memory Gap.

Bits	Description
15	**Memory Gap Enable.** 1=Enable. 0=Disable (default).
14:10	**Memory Gap Size.** This field defines the memory gap size as follows: **Bits[14:10] — Size — Bits[14:10] — Size** 00000 — 1 MB — 11100 — 8 MB 00100 — 2 MB — 11110 — 16 MB 01100 — 4 MB — 11111 — 32 MB Note that all other combinations are reserved.
9	**Reclaim Enable.** 1=Enable. 0=Disable (default). When enabled, the physical memory in this gap is reclaimed.
8	**Reserved.**
7:4	**Memory Gap Starting Address (Lower Nibble).** Bits [7:4] correspond to A[23:20]# and are used with bits [11:0] of the MGRUA Register to form the complete starting address.
3:0	**Reserved.**

2.3.16 MGUA—MEMORY GAP UPPER ADDRESS REGISTER

Address Offset: 7A–7Bh
Default: 0000h
Attribute: Read/Write

This register is used, along with the MG Register, to define the Memory Gap.

Bits	Description
15:12	**Reserved.**
11:0	**Memory Gap Starting Address (Upper Part).** These bits correspond to address bits A[35:24]# and are used with bits [7:4] of the MG Register to form the complete starting address.

2.3.17 LMG—LOW MEMORY GAP REGISTER

Address Offset: 7C–7Fh
Default: 00100000h
Attribute: Read/Write

This register defines the Low Memory Gap range. Note that the Low Memory Gap must be located below the Memory Gap and the High Memory Gap.

Bits	Description
31:20	**Low Memory Gap Starting Address (in 1 Mbyte increments).** Bits [31:20] correspond to address bits A[31:20]#. A[35:32]# are zero for this range (i.e., this range is limited to the lower 4 Gbytes).
19:12	**Reserved.**
11	**Low Memory Gap Enable.** 1=Enable. 0=Disable (default).
10	**Reclaim Enable.** 1=Enable. 0=Disable (default). When enabled, the physical memory in this gap is reclaimed.
4:0	**Low Memory Gap Size.** This field defines the memory gap size as follows: **Bits[4:0] Size Bits[4:0] Size** 00000 1 MB 11100 8 MB 00100 2 MB 11110 16 MB 01100 4 MB 11111 32 MB Note that all other combinations reserved.

2.3.18 HMGSA—HIGH MEMORY GAP START ADDRESS REGISTER

Address Offset: 88–8Bh
Default: 0000h
Attribute: Read/Write

This register, along with the HMGEA Register, define the High Memory Gap.

Bits	Description
31	**High Memory Gap Enable.** 1=Enable. 0=Disable (default).
29:16	**Reserved.**
15:0	**High Memory Gap Start Address (In 1 Mbyte increments).** Bits[15:0] correspond to A[35:20]#.

2.3.19 HMGEA—HIGH MEMORY GAP END ADDRESS REGISTER

Address Offset: 8C–8Fh
Default: 0000h
Attribute: Read/Write

This register, along with the HMGSA Register, define the High Memory Gap.

Bits	Description
31:16	**Reserved.**
15:0	**High Memory Range End Address (in 1 Mbyte increments).** Bits[15:0] correspond to A[35:20]# of the last 1 Mbyte region *within* the memory gap.

2.3.20 APICR—I/O APIC RANGE REGISTER

Address Offset: A4–A7h
Default: 00FEC001h
Attribute: Read/Write

This register defines a 64 Kbyte I/O APIC range.

Bits	Description
31:28	**Reserved.** Must be set to zero.
27:12	**I/O APIC Base Address (located on 1 MB increments).** Bits[27:12] correspond to A[35:20]#.
11:4	**Reserved.** Used in PB for unit ID numbers.
3:1	**Reserved.**
0	**I/O APIC Range Enable.** 1=Enable (default). 0=Disable. The I/O APIC range is 64 Kbytes. When this region is enabled, accesses to the region are ignored by the MC.

2.3.21 UERRADD—UNCORRECTABLE ERROR ADDRESS REGISTER

Address Offset: A8–ABh
Default: 0000h
Attribute: Read Only

This register provides the effective address of the *memory access* that caused an uncorrectable ECC error. The value in this register is only valid if the SBC error bit is set in the Error Reporting Register.

Bits	Description
31:3	**Address of First Uncorrectable ECC Error.** This is the effective address used in the MC and must be converted to the original physical address by software. MC base address and any programmed memory gaps must be taken into account for proper calculation of the address.
2:1	**QWord Number Error Detect.** When an uncorrectable error occurred in a transfer, this field indicates which QWord in the transfer contained the error. Note that this field reports the QWord number relative to the order of the transfer (0 to 3), even if the transfer does not begin with the first QWord of a cache line. In addition, in a single QWord transfer, if an error is detected, this field will be set to 00. **Bits [2:1] — QWord Number of the Transfer** 00 — First QWord Transferred (QWord 0) 01 — Second QWord Transferred (QWord 1) 10 — Third QWord Transferred (QWord 2) 11 — Fourth QWord Transferred (QWord 3)
0	**Reserved.**

2.3.22 MEMTIM—MEMORY TIMING REGISTER

Address Offset: AC–AFh
Default: 30DF3516h
Attribute: Read/Write

The memory timing register has two main functional sections—refresh timing and memory timing. The refresh timing portion of the memory timing register includes selections for time between refreshes (refresh counter) and time between refreshing rows in the memory array (refresh stagger). An enable bit for refreshing is also provided.

Most of the Asynchronous DRAM timing parameters are programmable in the MC to achieve maximum performance across a wide range of system operating frequencies. Each field in the memory timing register that pertains to DRAM timing is referenced by the most common DRAM timing parameter as published in the major DRAM vendors data books. Each field provides enough values to cover a wide range of operating frequencies. Care must be taken in programming the memory timing parameters so that the proper system timing is achieved and no conflicts are induced.

The memory timing register allows the memory controller to be adjusted for maximum performance when accessing Asynchronous DRAMs. The MC generates all control signals synchronously to the system clock. This limits the granularity of the generated control signals to a single clock period. The memory timing register allows the selection of the number of clocks to the most optimal value for a wide range of system clock frequencies.

On reset, the Memory Timing Register fields are set to values that allow operation in the range 50 MHz to 66 MHz with 70 ns DRAMs. The refresh counter is set such that refreshes occur assuming that the operating frequency is 50 MHz, which is faster than required at 60 or 66 MHz. All other parameters are set assuming that the operating frequency is 66 MHz which adds more clocks than required if the real frequency is 60 MHz. For optimal performance, the values in this register may have to be reprogrammed after reset.

<table>
<tr><th>Bits</th><th>Description</th></tr>
<tr><td>31</td><td>Reserved.</td></tr>
<tr><td>30:20</td><td>Refresh Count (in cycles) (REFRC). The refresh counter must be set so that refreshes occur often enough that the entire DRAM array is refreshed before DRAM data loss occurs. The eleven bit counter can be programmed from 1 to 2047. The counter time base is equal to one system clock period (15 ns for a 66 MHz clock, etc.). The value is chosen to give a refresh every 15.625 usec (or less). For example, 30Dh=15.620 usec at 50 MHz (default) and 411h = 15.615 usec at 66 MHz</td></tr>
<tr><td>19:17</td><td>Refresh Stagger (REFRS). The refresh stagger sets the time, in clock cycles, from the start of one row's refresh to the start of the next row's refresh. Refresh in the DRAMs causes the DRAMs to become active which draws considerable power. Refreshing all rows at once may not be possible for the system power supply. The refresh stagger field of the memory timing register allows the power surge to be spread evenly across the refresh cycle. The amount of time necessary to stagger the refreshes is system design and DRAM memory type dependent. 000=All rows refreshed at once.
<table>
<tr><th>Bits[19:17]</th><th>Stagger</th><th>Bits[19:17]</th><th>Stagger</th></tr>
<tr><td>000</td><td>0 cycles</td><td>100</td><td>4 cycles</td></tr>
<tr><td>001</td><td>1 cycle</td><td>101</td><td>5 cycles</td></tr>
<tr><td>010</td><td>2 cycles</td><td>110</td><td>6 cycles</td></tr>
<tr><td>011</td><td>3 cycles</td><td>111</td><td>7 cycles (default)</td></tr>
</table></td></tr>
<tr><td>16</td><td>Refresh Enable (REFRE). 1=Enable (default). 0=Disable.</td></tr>
<tr><td>15</td><td>CAS# Setup Time To RAS# for CAS-Before-RAS Refresh Cycles. 1=1 Cycles. 0=2 Cycle (default). Typically, 1 cycle is sufficient. However, in some cases the combination of DRAM timings, clock speed, and system level skew between CASx# and RASx# may require 2 cycles.</td></tr>
<tr><td>14:13</td><td>Last Write to CAS# (LWC). Number of cycles from when the last data is asserted to the MIC to when CAS# is asserted. This determines data setup time before CAS# (i.e., data is driven for LWC cycles, but delayed by one cycle).
<table>
<tr><th>Bits[14:13]</th><th>Cycles</th></tr>
<tr><td>00</td><td>Reserved</td></tr>
<tr><td>01</td><td>2 (default)</td></tr>
<tr><td>10</td><td>3</td></tr>
<tr><td>11</td><td>4</td></tr>
</table>
NOTES:
1. Write data setup time to CAS# asserted is LWC minus 1.
2. Write data hold time from CAS# asserted is 1 cycle if WCAS (bits[7:6]) equals 2 and is 2 cycles if WCAS is greater than 2.
3. The following are the legal combinations of the WCAS, LWC, and CP fields for non-interleaved and 2-way interleaved memory configurations. There are no restrictions for 4-way interleaved.
<table>
<tr><th>WCAS</th><th>LWC</th><th>CP</th><th>WCAS</th><th>LWC</th><th>CP</th></tr>
<tr><td>2</td><td>2</td><td>1,2</td><td>4</td><td>2</td><td>1,2</td></tr>
<tr><td>2</td><td>3</td><td>2</td><td>4</td><td>3</td><td>1,2</td></tr>
<tr><td>3</td><td>2</td><td>1,2</td><td>4</td><td>4</td><td>2</td></tr>
<tr><td>3</td><td>3</td><td>2</td><td></td><td></td><td></td></tr>
</table></td></tr>
</table>

Bits	Description
12:11	**RAS# Pulse Width (RASPW).** This field selects the minimum cycles that RASx# is asserted. RASPW should be set to the larger of the following values—(RCD + CAH + 1) or (RCD + WCAS - 1) or (RCD +RCAS - 1) **Bits[12:11] Cycles** 00 4 01 5 10 6 (default) 11 7
10	**Column Address Hold Time (CAH).** 0=1 cycle. 1=2 cycles (default). CAH is the number of cycles from the time CAS# is asserted to when the column address can be changed. Note that CAH must be set to 1 cycle if RCAS=2 or RASPW minus RCD = 2.
9:8	**Read CAS# Pulse Width (RCAS).** Number of cycles CAS# is asserted for read cycles. RCAS must be set to ensure data setup to the DP from CAS# asserted. **Bits[9:8] Cycles** 00 2 01 3 (default) 10 4 11 5
7:6	**Write CAS# Pulse Width (WCAS).** WCAS selects the number of cycles CAS# is asserted for write cycles. See notes for LWC field description. WCAS should be set to RCAS or RCAS minus 1, depending on data hold time requirements (see LWC field). **Bits[7:6] Cycles** 00 2 (default) 01 3 10 4 11 5
5	**CAS# Precharge Time (CP).** 0=1 cycle (default). 1=2 cycles. CP selects the number of cycles for CAS# precharge. See notes for LWC field description.
4	**RAS# to Column Address Delay (RCAD).** RCAD selects the number of cycles from the time RAS# is asserted to when Column address is asserted. 0 = 1 cycle. 1 = 2 cycles (default). RCAD must equal 1 cycle if RCD equals 2 cycles.
3:2	**RAS# to CAS# Delay (RCD).** RCD selects the number of cycles from the time RAS# is asserted to when CAS# is asserted. **Bits[3:2] Cycles** 00 Reserved 01 3 (default) 10 4 11 Reserved
1:0	**RAS# Precharge Time (RP).** RP selects the number of cycles for RAS# pre-charge. **Bits[1:0] Cycles** 00 3 01 4 10 5 (default) 11 6

2.3.23 SMMR—SMRAM RANGE REGISTER

Address Offset: B8–BBh
Default: 000Ah
Attribute: Read/Write

This register controls the size and location of SMRAM.

Bits	Description
31:28	**SMM Range Size.** The size (in 64 Kbyte increments) is selected as follows: **Bits[31:28] Size Bits[31:28] Size** 0000 64 KB (default) 0100 512 KB 0001 128 KB ... ... 0010 192 KB 1111 1 MB 0011 256 KB
27:16	**Reserved.**
15:0	**SMM Range Start Address (in 64 Kbyte increments).** Bits [15:0] correspond to A[31:16]#. Default=A0000.

2.3.24 HBIOSR—HIGH BIOS GAP RANGE REGISTER

Address Offset: BCh
Default: 01h
Attribute: Read/Write

This register enables/disables the high BIOS range. This range is 2 Mbytes extending from 00_FFE0_0000h to 00_FFFF_FFFFh).

Bits	Description
7:1	**Reserved.**
0	**High BIOS Range Enable.** 1=Enable (default). 0=Disable. When enabled, the MC ignores accesses to this range.

2.3.25 MERRCMD—MEMORY ERROR REPORTING COMMAND

Address Offset: C0–C1h
Default: 0000h
Attribute: Read /Write

This register enables/disables the correcting of single bit errors in memory data. This register also controls the reporting of correctable and uncorrectable memory errors.

Bits	Description
15:3	**Reserved.**
2	**Single-bit Error Correcting of Memory Data Enable.** 1=Enable. 0=Disable (default). This bit must be enabled for the "identify SIMM" functionality described for bit 0 in the SBCERRADD Register (detecting which half of a QWord generated the error).
1	**Report Correctable Memory Errors Enable.** 1=Enable. 0=Disable (default). When enabled, the MC reports correctable errors in data read from memory in the MERRSTS Register. The MC also asserts the sideband single SBCERR#.
0	**Report Uncorrectable Memory Errors Enable.** 1=Enable. 0=Disable (default). When enabled, the MC logs uncorrectable errors in the MC Memory Error Status Register. If BERR# output is enabled (SERRCMD Register), BERR# is also be asserted.

2.3.26 MERRSTS—MEMORY ERROR STATUS REGISTER

Address Offset: C2–C3h
Default: 0000h
Attribute: Read /Write Clear

This register logs correctable and uncorrectable memory errors. Software sets these bits to 0 by writing a 1 to them.

Bits	Description
15:2	**Reserved.**
1	**Correctable Memory Error.** 1=Detected correctable memory error.
0	**Uncorrectable Memory Error.** 1=Detected uncorrectable memory error.

2.3.27 SERRCMD—SYSTEM ERROR REPORTING COMMAND REGISTER

Address Offset: C4–C5h
Default: 0000_0000_000x_0x0xb (x=captured at reset)
Attribute: Read/Write

This register controls the reporting of system errors. Note that when bits[9:7] of this register are disabled, the MC forces all ECC bits written to memory to 0. This mechanism is used to force ECC errors in the memory array for debugging the memory error correcting/detecting circuits.

Bits	Description
15:10	**Reserved.**
9	**450KX: Reserved.** **450GX: Single-bit Error Correcting of Host Data Enable.** 1=Enable. 0=Disable (default).
8	**450KX: Reserved.** **450GX: Logging Correctable Errors on the Host Data Bus Enable.**1=Enable. 0=Disable (default). When enabled, the MC logs correctable errors in data read from the host bus in the SERRSTS Register. The MC also asserts the sideband single SBCERR#.
7	**450KX: Enable Memory ECC.** 1=Enable. 0=Disable. This bit must be set to enable ECC on the memory array. **The memory array must be initialized before enabling memory ECC.** **450GX: Logging Uncorrectable Errors on the Host Data Bus Enable.** 1=Enable. 0=Disable (default). When enabled, the MC logs uncorrectable errors in the SERRSTS Register. If BERR# is enabled, BERR# is also be asserted. **The memory array must be initialized before enabling memory ECC.**
6	**Reserved.**
5	**AERR# Driver Enable.** 1=Enable. 0=Disable (default). This bit enables/disables reporting of parity errors on request signals.
4	**AERR# Input Enable.** 1=Enable. 0=Disable. The MC captures this value from A8#.
3	**BERR# Driver Enable.** 1=Enable. 0=Disable (default). This enables/disables reporting of uncorrectable errors on the data bus or memory interface.
2	**BERR# Input Enable.** 1=Enable. 0=Disable. The MC captures this value from A9#.
1	**BINIT# Driver Enable.** 1=Enable. 0=Disable. When enabled, protocol violations are reported on BINIT#.
0	**BINIT# Input Enable.** 1=Enable. 0=Disable. The MC captures this value from A10#.

2.3.28 SERRSTS—SYSTEM ERROR STATUS REGISTER

Address Offset: C6–C7h
Default: 0000h
Attribute: Read /Write Clear

This register logs system errors. Software sets these bits to 0 by writing a 1 to them.

Bits	Function
15:5	**Reserved.**
4	**Host Address Parity Error Detected. (via AP[1:0]#).** 1=Logs parity errors on A[35:3], regardless of whether the event is reported. If AERR# Input Enable (bit 4) of the SERRCMD Register is set, the event is reported during the error phase.
3	**Host Bus Request Parity Error Detected (via RP#).** 1=Logs parity errors on the ADS# and REQ[4:0]# signals, regardless of whether the event is reported. If enabled in the SERRCMD Register (bit 4), this error is reported by generating an AERR#.
2	**450KX: Reserved.** **450GX: Host Bus Correctable Error Detected.** 1=Logs a single-bit ECC error detected on the host data bus. No error is reported when a host bus correctable error is detected.
1	**450KX: Reserved.** **450GX: Host Bus Uncorrectable Error Detected.** 1=Logs a multiple-bit ECC error detected on the host data bus. Note that this bit is set independent of whether error reporting is enabled via bit 7 of the SERRCMD Register. If BERR# is enabled in the SERRCMD Register, this error is reported by generating a BERR#.
0	**Host Bus Protocol Violation Detected (via RSP#).** 1=Logs protocol violations, regardless of whether event is reported. If BINIT# is enabled in the SERRCMD Register, these errors are reported by generating a BINIT#.

2.4 Memory Configuration Determination Algorithm

The number of rows of memory and the size of the memory in each row must be determined by power-on self test (POST) code prior to programming the configuration registers for the true system configuration.

After reset, each MC is configured for a non-interleaved memory configuration operating with the default values given in the Memory Timing Register. Base addresses are set assuming maximum memory. However, row limits are set at 4 Mbytes.

To complete the configuration of the MCs in a system the BIOS must perform a complete setup as described in the *Pentium Pro Processor BIOS Writer's Guide* (Order #649733).

3.0 MC FUNCTIONAL DESCRIPTION

This section describes the MC functions and hardware interfaces including the Memory and I/O Mapping, Host Bus Interface, DRAM Interface, and Clocks and Reset.

3.1 Memory and I/O Map

The MC provides the interface between the host bus and main memory. The processor memory space is 64 Gbytes (36-bit addressing). An MC can control up to 1 Gbyte of memory for the 450KX and 4 Gbytes of memory for the 450GX. The MC registers that control memory space access are:

- **Programmable Attribute Map (PAM[6:0]) Registers.** These registers provide Read Only, Write Only, and Read/Write Disable for fixed memory regions in the PC compatibility area.
- **Video Buffer Area Enable (VBA) Register.** This register enables the A0000–BFFFFh fixed region.
- **Low Memory Gap (LMG) Register.** This register defines a hole in memory located from 1 to 4 Gbytes on any 1 Mbyte boundary where accesses can be directed to the PCI bus (via the PB). The size can be 1, 2, 4, 8, 16, or 32 Mbytes. This gap must be located below the Memory Gap and High Memory Gap. The Low Memory Gap is used by ISA devices such as LAN or linear frame buffers that are mapped into the ISA Extended region, or by any EISA or PCI device.
- **Memory Gap Registers (MG and MGUA) Registers.** These two registers define a hole in memory located from 1 to 64 Gbytes on any 1 Mbyte boundary where accesses can be directed to the PCI bus (via the PB). This gap (1, 2, 4, 8,16, or 32 Mbytes in size) must be located above the Low Memory Gap and below the High Memory Gap areas. The Memory Gap is used by ISA devices (e.g., LAN or linear frame buffers) that are mapped into the ISA Extended region, or by any EISA or PCI device.
- **High Memory Gap Registers (HMGSA and HMGEA) Registers.** These two registers define a gap in memory that can be located from 1 to 64 Gbytes on any 1 Mbyte boundary where accesses can be directed to the PCI bus (via the PB). The size ranges from 1 Mbyte to 64 Gbytes. This gap must be located above the Memory Gap and the Low Memory Gap areas. The High Memory Gap provides additional support for memory mapped I/O.
- **Base Address (BASEADD) Register.** An 82453GX responds to memory accesses between the address programmed into this register and the calculated top of its memory range (calculated top of MC memory address = base + memory size + Low Memory Gap size + Memory Gap size + High Memory Gap size). Note that the DRAM memory behind the memory gaps can be reclaimed.
- **SMMRAM Range (SMMR) Register and the SMMRAM Enable (SMME) Register (Only when SMMEM# is asserted by the processor.).** SM memory can overlap with memory residing on the host bus or memory normally residing on the PCI bus. When the SM range is enabled, SM accesses are handled by the MC. If the SMMEM# signal is not asserted, accesses to the MC's enabled SM Range are ignored (this allows the SM memory to overlap with memory normally residing on the host bus, since the SMM Range may also be mapped through another MC range register). The RSMI# signal may be asserted in the Response Phase by a device in SMM power-down mode. The MC does not assert this signal.

NOTE:

Since leaving system management mode effectively remaps the system memory space, one must take care with SMM memory that is cached. If SMMRAM is cast as writeback memory, a WBINVD instruction must be executed immediately prior to the execution of the RSM instruction which exits SMM mode. This will force all modified data to be written back while memory is still mapped for SMM.

- **High BIOS (HBIOS) Register.** The 64 Kbyte region from F0000–FFFFFh is treated as a single block and is normally Read/Write disabled in the MC(s) and Read/Write enabled in the PB. After power-on reset, this region is R/W enabled in the PB (Compatibility PB only in the 450GX and R/W disabled in the Auxiliary PB). Thus, the PB can respond to fetches during system initialization. The Read/Write attributes for this region may be used in conjunction with the Read/Write attributes in the PB to "shadow" BIOS into RAM.
- **I/O APIC Range (APICR) Register.** This register provides an I/O APIC configuration space. There is no I/O APIC in the PB or the MC.
- **DRAM Row Limit (DRL) Registers.** These registers define the upper and lower addresses for each DRAM row and represent the boundary addresses in 4 Mbyte granularity.

If a memory space access is in one of the above ranges, and that range is enabled for memory access, the MC claims the transaction and becomes the response agent.

The MC performs memory recovery on gap ranges greater than or equal to 1 Mbyte that are created by the Low Memory Gap, Memory Gap, and the High Memory Gap areas. This memory is relocated to the top of the MC's memory. The MC performs a subtraction of the size of the hole in the memory map to generate an effective memory address.

For the 450GX, the base address for the MC that is not MC #0 must include the size of any memory gaps programmed in the previous (or lower base address) MC.

There can be up to two MCs in a system permitting up to 8 Gbytes of system main memory. The portion of the processor's memory space controlled by an MC is determined by the Base Address Register and memory size. In a PC architecture, the only restrictions on MC placement are that there be memory starting at address 0 and that there be enough memory to operate a system. The MCs in a system need not have contiguous address spaces. The High Memory Gap in one MC could be used to span the gap between the top of its memory map and the base address of the other MC.

Note that the PB (Compatibility PB in an 450GX dual PB system) is responsible for claiming any unclaimed transactions on the host system bus. Therefore, any memory space access that is above the top of system main memory is claimed by the PB.

The MC has two registers located in the processor's I/O space (0CF8h and 0CFCh) that are used to configure the MC. See the Register Descripiton section for details.

3.2 Host Bus Interface

The Pentium Pro processor bus provides an efficient, reliable interconnect between multiple Pentium Pro processors and the PB and MC. The bus provides 36 bits of address, 64 bits of data, protection signals needed to support data integrity, and the control signals to maintain a coherent shared memory in the presence of multiple caches.

The Pentium Pro processor bus achieves high bus efficiency by providing support for multiple, pipelined transactions and deferred replies. A single Pentium Pro processor may have up to four transactions outstanding at the same time, and can be configured to support up to eight transactions active on the Pentium Pro processor bus at any one time. The MC supports up to four transactions that target its associated memory space. The MC contains read and write buffers for memory accesses.

NOTES:

1. The MC does not generate deferred responses.
2. The MC does not provide the ability to abort a transaction during the response phase.
3. On the host bus, a Hard Failure Response is generated for failures in accessing a resource. Such a failure could be a time-out after requesting a device that is not available. Note that data failures do not fall into a hard failure class. The MC does not generate Hard Failure responses.
4. All transactions in the MC are processed in "address" order with respect to when they are received on the host bus. There is reordering of read-around-writes, but only when the address of the read is different from the address of the write. If there is an address conflict, the transactions are processed in the order they are received. (Note, responses to transactions still occur in the order in which they were received, only the processing of the requests is reordered.)
5. The MC does not respond to an SMI Acknowledge Transaction or Stop Clock Acknowledge Transaction, even though they are encoded as memory type operations on the host bus.

AERR#. An AERR# on the host bus stops traffic in the memory controller. Reporting is done by the 82454 (PB).

BINIT#. A BINIT# on the Host bus resets the 450KX/GX host bus state machines. This allows for logging or recovery from catastrophic bus errors. Note that during the last clock of a BINIT# pulse, ADS# may not be asserted as this will start the host bus state machine prematurely.

3.3 DRAM Interface

In the following discussion the term *row* refers to the set of memory devices that are simultaneously selected by a RAS# signal. A row may be composed of two or more single-sided SIMMs, or one side (the same side) from two or more double-sided SIMMs. An *interleave* is 72-bits wide (64 data bits plus 8 bits of ECC) and requires two 36 bit SIMMs. The term *page* refers to the data within a row that is selected by a row address and is held active in the device waiting for a column address to be asserted.

The MC interfaces the main memory DRAM to the host bus. For the 450KX, two basic DRAM configurations are supported—2-way interleaved (or 2:1 interleaved), and non-interleaved (or 1:1 interleaved). In the 2-way and non-interleaved configurations, a row is made up of 4 SIMM sides and 2 SIMM sides respectively. There can be up to 1 Gbyte of DRAM for a 2-way interleaved configuration and 512 Mbytes of DRAM for a non-interleaved configuration as shown in Table 22. The MC is fully configurable through the MC's configuration registers.

For the 450GX, three basic DRAM configurations are supported—4-way interleaved (4:1 interleaved), 2-way interleaved, and non-interleaved. In the 4-way interleaved configuration, a row is made up of 8 36-bit SIMM sides. In the 2-way interleaved and non-interleaved configurations, a row is made up of 4 SIMM sides and 2 SIMM sides respectively. There can be up to 4 Gbytes of DRAM for a 4-way interleaved configuration, 2Gbytes for a 2-way interleaved configuration, and 1Gbyte for a non-interleaved configuration.

Configurations cannot be mixed. The MC does not support portions of the memory being 2-way interleaved and other portions being non-interleaved. The system does, however, support a 2-way interleaved design in which one interleave is populated (operates as a non-interleaved configuration). There is no restriction on which interleave is populated (0 or 1) to form a non-interleaved configuration, as long as all rows are populated in the same way.

The 450GX MC does not support portions of the memory being 4-way interleaved and other portions being non-interleaved or 2-way interleaved. The system does, however, support a 4-way or 2-way interleaved design in which one interleave is populated (operates as a non-interleaved configuration) or a 4-way interleaved design in which two interleaves are populated (operates as a 2-way configuration). There is no restriction on which interleaves are populated to form a non-interleaved or 2-way interleaved configuration, as long as all rows are populated in the same way.

Table 22 provides a summary of the characteristics of memory configurations supported by the 450KX/GX MC. Minimum values listed are obtained with single-sided SIMMs, and maximum values are obtained with double-sided SIMMs.

Table 22. Minimum and Maximum Memory Sizes for Each Configuration

Device	Non-Interleaved		2-Way Interleaved		4-Way Interleaved	
	Min (Inc)	Max	Min (Inc)	Max	Min (Inc)	Max
512k x 8	4 MB	16 MB 32 MB	8 MB	32 MB 64 MB	16 MB	128 MB
1M x 4	8 MB	32 MB 64 MB	16 MB	64 MB 128 MB	32 MB	256 MB
2M x 8	16 MB	64 MB 128 MB	32 MB	128 MB 256 MB	64 MB	512 MB
4M x 4	32 MB	128 MB 256 MB	64 MB	256 MB 512 MB	128 MB	1 GB
8M x 8	64 MB	256 MB 512 MB	128 MB	512 MB 1 GB	256 MB	2 GB
16M x 4	128 MB	512 MB 1 GB	256 MB	1 GB 2 GB	512 MB	4 GB

For the 450GX, two MCs can be used in a system permitting a maximum of 8 Gbytes of main memory. When two controllers are present in a system, the memory configuration and operation of an MC is independent of the other MC. The OMCNUM signal determines the configuration address and the default base address of a controller. Beyond this, the general behavior of each memory subsystem is identical.

Refresh Operation

Refresh for the memory array is handled automatically by the MC. The rate of refresh cycles is programmable in the MEMTIM register (AC-AFh).

An alternative to a single refresh cycle is to stagger refreshes across the DRAM rows. Refresh stagger allows the refresh power surge to be tailored to the system. This allows the system to select staggering of row refreshes by one clock increments for zero to seven cycles. While refreshing fewer rows at once increases the chance of a refresh request collision with a host request, it enables the system to handle the power surge caused by refresh. Staggering refreshes within a group increases the time for the group refresh, but spreads the power demands over time, and thus allows larger groups to be refreshed. Refresh Staggering provides substantial power surge reduction over refreshing all rows simultaneously.

3.3.1 DRAM CONFIGURATIONS

The memory supported by the MC is arranged as 4 rows with 1 or 2 interleaves (8 rows with 1, 2, or 4 interleaves for the 450GX). This can be implemented with discrete memory devices, single-sided SIMMs, or double-sided SIMMs. Systems in which adjacent rows of memory have a common CAS# connection are a special case, and are selected via the CMD Register (offset 4C–4Fh). The primary example of this is a system constructed with double-sided SIMMs having a common CAS# connection between the two sides.

For all the memory configuration types, the MC provides 4 logical RAS# signals (8 for the 450GX); one per row. Two copies of the RAS# signals (RASAx# and RASBx#) are provided for fanout. The MC provides 4 logical CAS# signals (8 for the 450GX). Two copies (CASAx# and CASBx#) are provided for fanout.

In the case of a common CAS# connection between adjacent rows, there are only 2 logical CAS# signals (one per pair of rows) for the 450KX and 4 logical CAS# signals for the 450GX. In addition, the loading per CAS# signal is doubled. To accommodate this, the MC combines the CAS# signals for two rows (e.g. CASA[1:0#] and CASB[1:0]#) are driven with the same value and are used to drive the first two memory rows.

The descriptions of the supported configurations that follow assume that the DRAM in the system is implemented with double-sided SIMMs that do not have a common CAS# connection and that do not have buffers on the SIMMs. Figure 2 shows the connections required for each double-sided SIMM (DSSIMM). Note that these are SIMM connections and do not map one-to-one to MC signals. Also shown is the symbol used to represent the 72-bit wide memory formed from two DSSIMMs.

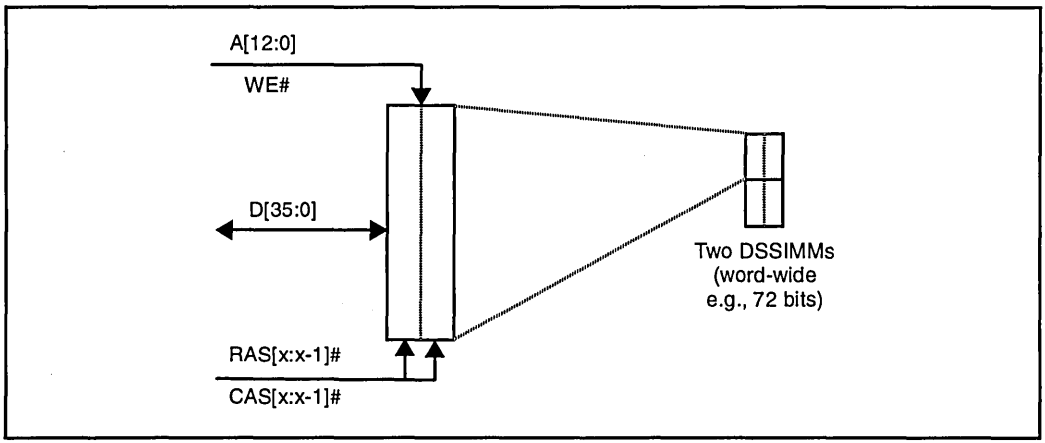

Figure 2. Signal Connections to a Double-Sided SIMM

3.3.1.1 Memory Interface Component (MIC)

To interface with the data signals from the devices in the memory array, the MC utilizes a set of four Memory Interface Components (MICs), each 18 bits wide. These components multiplex data read from the interleaved memory, register data being written to memory, and provide the buffering required to drive the memory devices. All configurations utilize four of these devices. The interconnection of the MICs and the memory devices is shown in Figure 3.

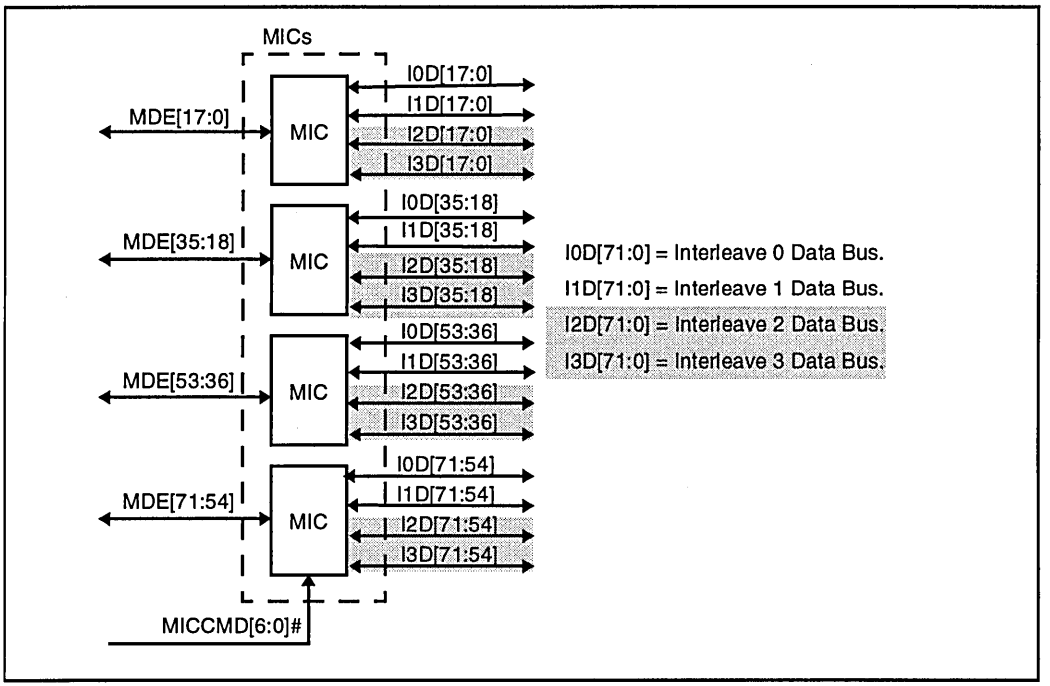

Figure 3. MIC to Memory Interconnections

3.3.1.2 4-Way DRAM Configuration (450GX Only)

In the 4-way interleaved DRAM configuration, the memory controller supports up to 8 rows of conventional DRAM. Each of these rows can be up to 512 Mbytes, using 64-Mbit technology. This configuration is illustrated in Figure 4, as implemented with DSSIMMs.

The basic structure of the 4-way interleaved memory (Figure 4) is four 72-bit word wide connections from the DRAM time multiplexed to the MC. This multiplexing allows the MC to read or write memory at the rate of one 72-bit word each clock cycle and to hide much of the access latency of the DRAM devices.

Logically, the 4-way configuration requires eight RAS# lines and eight CAS# lines (one for each row). The RAS# signals latch the row address in the four interleaves, and the CAS# signals latch the column address in each interleave. MA[12:0] and WE# are broadcast to all devices, and must be buffered to each DSSIMM. The exact buffer type used is system design dependent.

The minimum memory size for this configuration is 4 Mbytes using 4-Mbit technology organized as 512kx8 devices (1 row, 4 Mbytes, operating as non-interleaved). The maximum size is 4 Gbytes using 64-Mbit technology. Mixing of row sizes is supported; however, within a row, all SIMMs must be the same size.

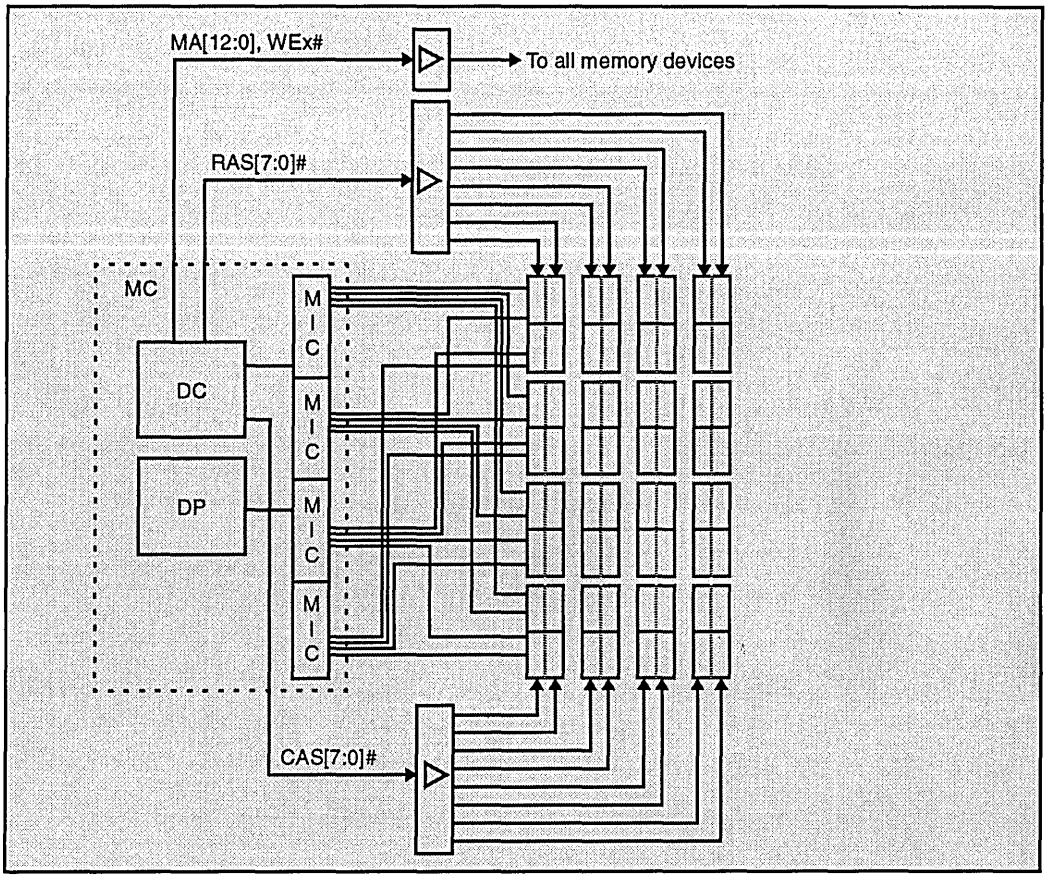

Figure 4. 4-Way Interleaved Configuration

3.3.1.3 2-Way DRAM configuration

In a 2-way interleaved DRAM configuration (Figure 5), the memory controller supports up to 4 rows (8 rows for an 450GX) of 2-way interleaved DRAM. The minimum and maximum main memory sizes are listed in Table 22.The MC supports mixing different row sizes; however, within a row, all SIMMs must be the same size. MA[12:0] and WE# are broadcast to all devices, and must be buffered to each DSSIMM. The exact buffer type used is system design dependent.

A 2-way interleaved configuration may result from populating two interleaves of the 4:1 memory system shown in the previous section, or as a 2-way interleaved only system. The signal connections for 2-way interleaved memory are the same as those described for the 4-way interleaved configuration. There is no restriction on the connection of the two interleaves to the MIC ports.

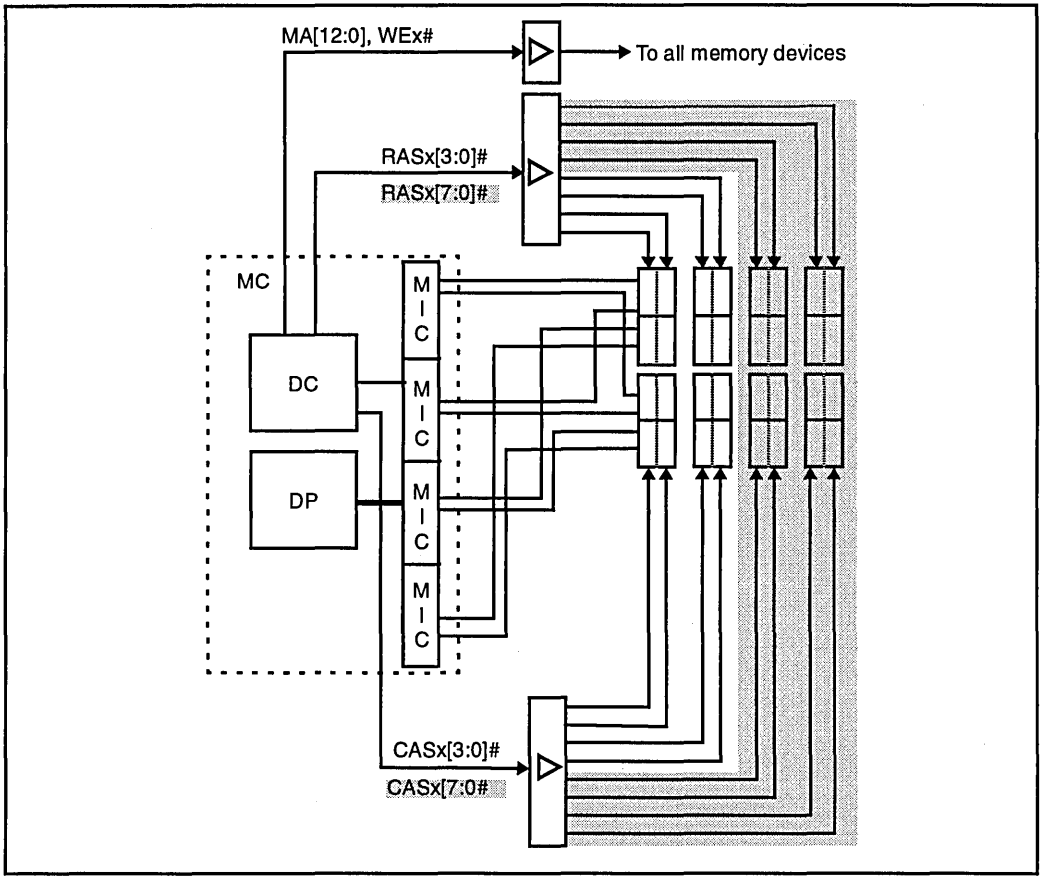

Figure 5. 2-Way Interleaved Configuration

3.3.1.4 Non-Interleaved DRAM configuration

In this configuration, the MC supports up to 4 rows (8 rows for an 450GX) of non-interleaved DRAM. This configuration may result from populating any one interleave of a 2-way interleaved system (or 4-way interleaved system for the 450GX) described previously. Note, however, that in practice, the non-interleaved configuration is used only for minimum memory sizes. Memory expansion usually occurs across interleaves before extending the number of rows. The MC supports mixing different row sizes.

3.4 Clocks and Reset

3.4.1 CLOCKS

The host bus clock is used for the MC and is input on the BCLK pin. This clock is generated externally and distributed to host bus components by a low skew clock driver (Figure 6). The clock driver provides multiple copies of the bus clock. The loading on the clock lines must be balanced to minimize clock skew among the components on the bus.

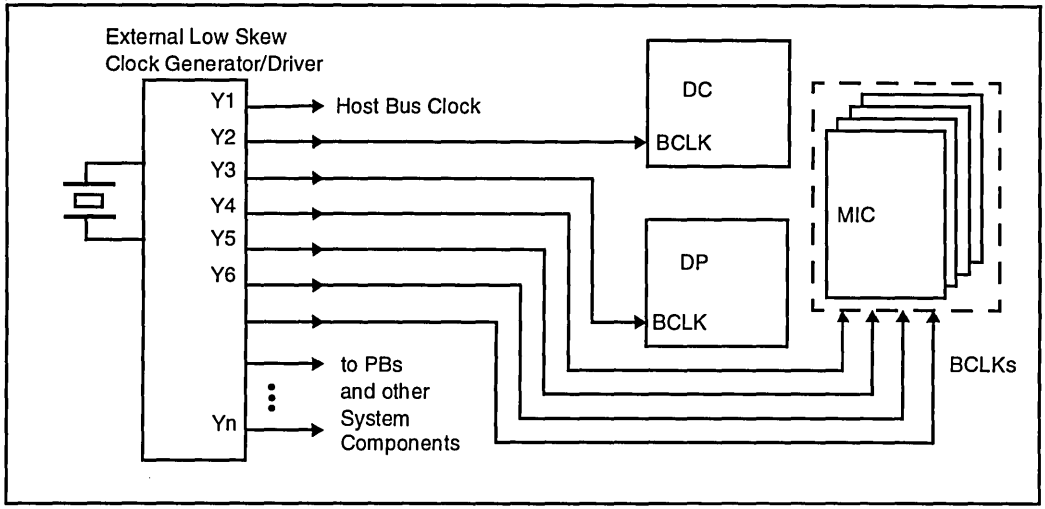

Figure 6. Clock Distribution

3.4.2 RESET

When the system is initially powered up, the power supply must wait until all voltages are stable for at least one millisecond, and then assert the PWRGD signal. A transition from 0 to 1 on PWRGD resets the PCI Bridge (PB) and portions of the DC. The PB is responsible for resetting and configuring the DC and other host devices.

During a hard reset, the MC initalizes its internal registers. When the DC receives a hard reset on RESET#, it resets the DP and MIC by asserting MIRST#. MIRST# on the DC should be connected to the MIRST# on the DP and MIC.

When the PB (Compatibility PB in an 450GX dual PB system) generates a hard reset, it also drives the appropriate host data bus signals with the values specified in its Configuration Values Driven on Reset Register. The MC captures the values it needs (see Capture System Configuration Register description).

4.0 MC PINOUT AND PACKAGE INFORMATION

The pinout and package information for the 82453GX/KX, 82452GX/KX, and 82451GX/KX is shown in this section. Pins unique to the 82450GX are shown in a list at the center of the figure. In the tables, the first name is the 450GX name. **Note that TESTLO pins must be pulled low with a 150Ω resistor and TESTHI pins must be pulled to 3.3V with a 10KΩ resistor. GTLHI pins should be pulled up with 10KΩ to V_{TT}.**

4.1 82453KX/82453GX (DC) Pin Assignment

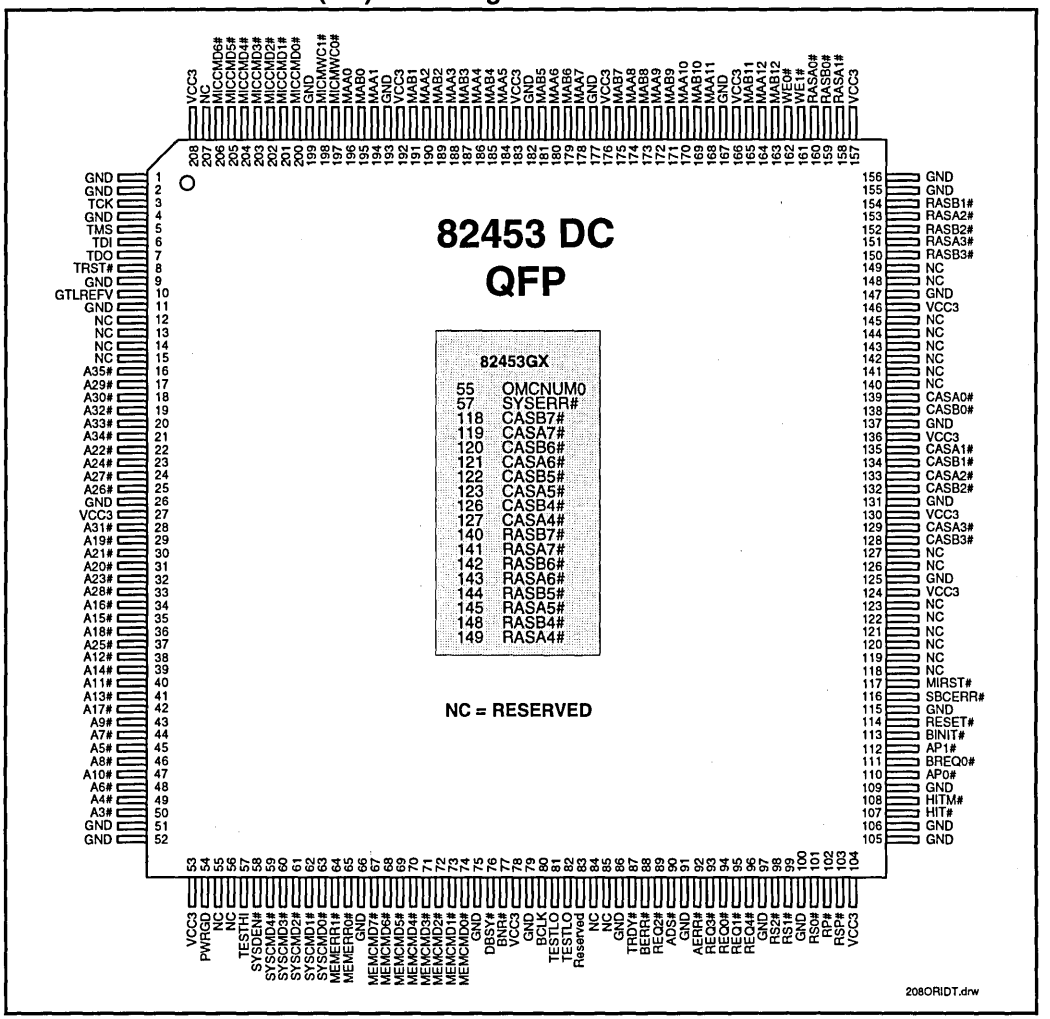

Figure 7. 82453KX/GX DC Pin Assignment (208-Pin QFP)

Table 23. 82453 KX/GX DC Alphabetical Pin List (208-Pin QFP)

Name	Pin#	Type
A3#	50	I
A4#	49	I
A5#	45	I
A6#	48	I
A7#	44	I
A8#	46	I
A9#	43	I
A10#	47	I
A11#	40	I
A12#	38	I
A13#	41	I
A14#	39	I
A15#	35	I
A16#	34	I
A17#	42	I
A18#	36	I
A19#	29	I
A20#	31	I
A21#	30	I
A22#	22	I
A23#	32	I
A24#	23	I
A25#	37	I
A26#	25	I
A27#	24	I
A28#	33	I
A29#	17	I
A30#	18	I
A31#	28	I
A32#	19	I
A33#	20	I
A34#	21	I
A35#	16	I
ADS#	90	I
AERR#	92	I/O
AP0#	110	I

Table 23. 82453 KX/GX DC Alphabetical Pin List (208-Pin QFP) (Continued)

Name	Pin#	Type
AP1#	112	I
BCLK	80	I
BERR#	88	I/O
BINIT#	113	I/O
BNR#	77	I/O
BREQ0#	111	O
CASA0#	139	O
CASA1#	135	O
CASA2#	133	O
CASA3#	129	O
CASA4#/NC	127	O
CASA5#/NC	123	O
CASA6#/NC	121	O
CASA7#/NC	119	O
CASB0#	138	O
CASB1#	134	O
CASB2#	132	O
CASB3#	128	O
CASB4#/NC	126	O
CASB5#/NC	122	O
CASB6#/NC	120	O
CASB7#/NC	118	O
DBSY#	76	I/O
GND	1	V
GND	2	V
GND	4	V
GND	9	V
GND	11	V
GND	26	V
GND	51	V
GND	52	V
GND	66	V
GND	75	V
GND	79	V
GND	86	V
GND	91	V

Table 23. 82453 KX/GX DC Alphabetical Pin List (208-Pin QFP) (Continued)

Name	Pin#	Type
GND	97	V
GND	100	V
GND	105	V
GND	106	V
GND	109	V
GND	115	V
GND	125	V
GND	131	V
GND	137	V
GND	147	V
GND	155	V
GND	156	V
GND	167	V
GND	177	V
GND	182	V
GND	193	V
GND	199	V
GTLREFV	10	I
HIT#	107	I/O
HITM#	108	I/O
MAA0	196	O
MAA1	194	O
MAA2	190	O
MAA3	188	O
MAA4	186	O
MAA5	184	O
MAA6	180	O
MAA7	178	O
MAA8	174	O
MAA9	172	O
MAA10	170	O
MAA11	168	O
MAA12	164	O
MAB0	195	O
MAB1	191	O
MAB2	189	O

Table 23. 82453 KX/GX DC Alphabetical Pin List (208-Pin QFP) (Continued)

Name	Pin#	Type
MAB3	187	O
MAB4	185	O
MAB5	181	O
MAB6	179	O
MAB7	175	O
MAB8	173	O
MAB9	171	O
MAB10	169	O
MAB11	165	O
MAB12	163	O
MEMCMD0#	74	I/O
MEMCMD1#	73	I/O
MEMCMD2#	72	I/O
MEMCMD3#	71	I/O
MEMCMD4#	70	I/O
MEMCMD5#	69	I/O
MEMCMD6#	68	I/O
MEMCMD7#	67	I/O
MEMERR0#	65	I
MEMERR1#	64	I
MIRST#	117	O
MICCMD0#	200	O
MICCMD1#	201	O
MICCMD2#	202	O
MICCMD3#	203	O
MICCMD4#	204	O
MICCMD5#	205	O
MICCMD6#	206	O
MICMWC0#	197	O
MICMWC1#	198	O
OMCNUM/ TESTLO	55	I
PWRGD	54	I
RASA0#	160	O
RASA1#	158	O
RASA2#	153	O

Table 23. 82453 KX/GX DC Alphabetical Pin List (208-Pin QFP) (Continued)

Name	Pin#	Type
RASA3#	151	O
RASA4#/NC	149	O
RASA5#/NC	145	O
RASA6#/NC	143	O
RASA7#/NC	141	O
RASB0#	159	O
RASB1#	154	O
RASB2#	152	O
RASB3#	150	O
RASB4#/NC	148	O
RASB5#/NC	144	O
RASB6#/NC	142	O
RASB7#/NC	140	O
REQ0#	94	I
REQ1#	95	I
REQ2#	89	I
REQ3#	93	I
REQ4#	96	I
RESERVED	12	NC
RESERVED	13	NC
RESERVED	14	NC
RESERVED	15	NC
RESERVED	83	NC
RESERVED	207	NC
RESET#	114	I
RP#	102	I/O
RS0#	101	I/O
RS1#	99	I/O
RS2#	98	I/O
RSP#	103	I/O
SBCERR#	116	O
SYSCMD0#	63	O
SYSCMD1#	62	O
SYSCMD2#	61	O
SYSCMD3#	60	O
SYSCMD4#	59	O

Table 23. 82453 KX/GX DC Alphabetical Pin List (208-Pin QFP) (Continued)

Name	Pin#	Type
SYSDEN#	58	O
SYSERR#/ TESTHI	57	I
TCK	3	I
TDI	6	I
TDO	7	O
TESTLO	56	I
TESTLO	81	I
TESTLO	82	I
TESTLO	84	I
TESTLO	85	I
TMS	5	I
TRDY#	87	I/O
TRST#	8	I
VCC3	27	V
VCC3	53	V
VCC3	78	V
VCC3	104	V
VCC3	124	V
VCC3	130	V
VCC3	136	V
VCC3	146	V
VCC3	157	V
VCC3	166	V
VCC3	176	V
VCC3	183	V
VCC3	192	V
VCC3	208	V
WE#0	162	O
WE#1	161	O

4.2 82452GX/82452KX (DP) Pin Assignment

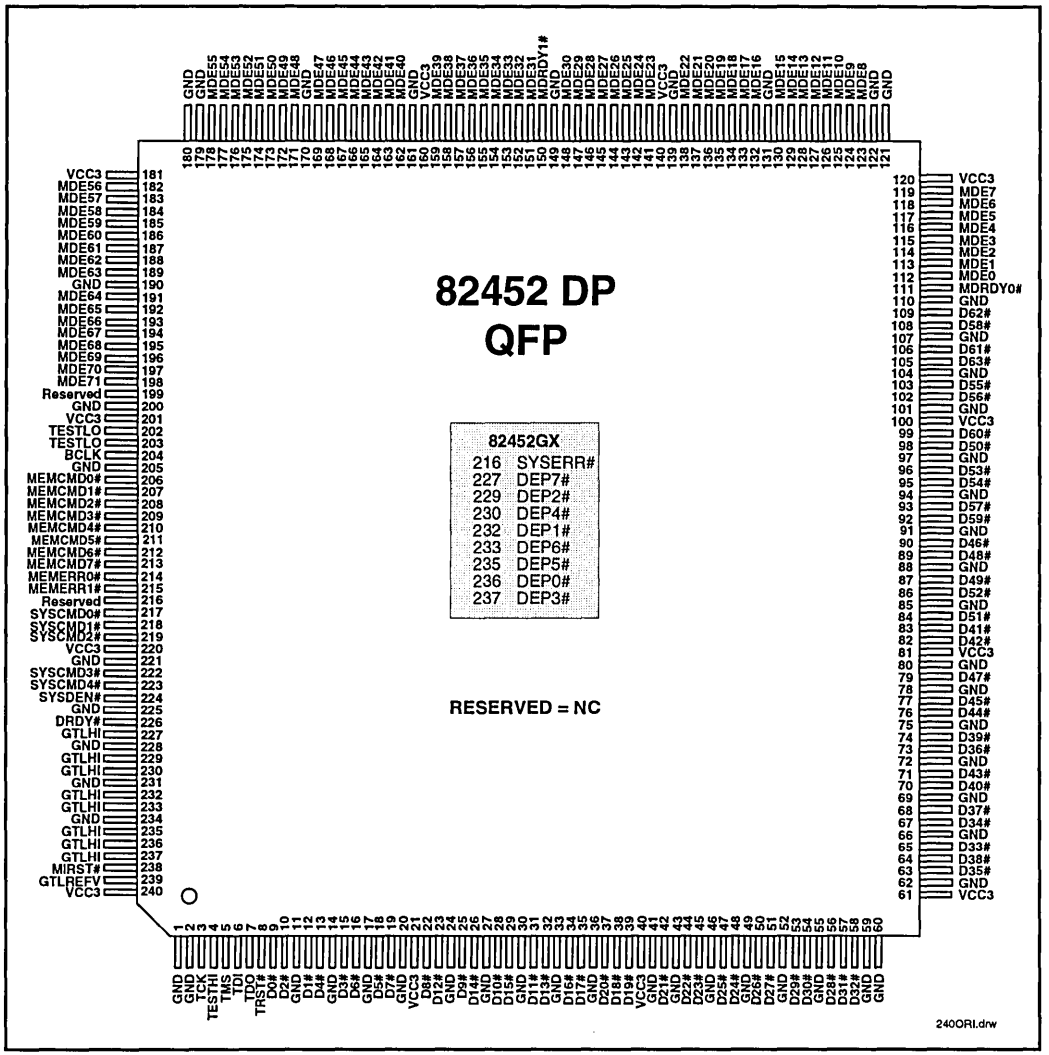

Figure 8. 82452 DP Pin Assignment (240-Pin QFP)

Table 24. 82452 KX/GX DP Alphabetical Pin List (240-Pin QFP)

Signal	Pin#	Type
BCLK	204	I
D0#	9	I/O
D1#	12	I/O
D2#	10	I/O
D3#	15	I/O
D4#	13	I/O
D5#	18	I/O
D6#	16	I/O
D7#	19	I/O
D8#	22	I/O
D9#	25	I/O
D10#	28	I/O
D11#	31	I/O
D12#	23	I/O
D13#	32	I/O
D14#	26	I/O
D15#	29	I/O
D16#	34	I/O
D17#	35	I/O
D18#	38	I/O
D19#	39	I/O
D20#	37	I/O
D21#	42	I/O
D22#	44	I/O
D23#	45	I/O
D24#	48	I/O
D25#	47	I/O
D26#	50	I/O
D27#	51	I/O
D28#	56	I/O
D29#	53	I/O
D30#	54	I/O
D31#	57	I/O

Table 24. 82452 KX/GX DP Alphabetical Pin List (240-Pin QFP) (Continued)

Signal	Pin#	Type
D32#	58	I/O
D33#	65	I/O
D34#	67	I/O
D35#	63	I/O
D36#	73	I/O
D37#	68	I/O
D38#	64	I/O
D39#	74	I/O
D40#	70	I/O
D41#	83	I/O
D42#	82	I/O
D43#	71	I/O
D44#	76	I/O
D45#	77	I/O
D46#	90	I/O
D47#	79	I/O
D48#	89	I/O
D49#	87	I/O
D50#	98	I/O
D51#	84	I/O
D52#	86	I/O
D53#	96	I/O
D54#	95	I/O
D55#	103	I/O
D56#	102	I/O
D57#	93	I/O
D58#	108	I/O
D59#	92	I/O
D60#	99	I/O
D61#	106	I/O
D62#	109	I/O
D63#	105	I/O
DEP0#/ GTLHI	236	I/O

Table 24. 82452 KX/GX DP Alphabetical Pin List (240-Pin QFP) (Continued)

Signal	Pin#	Type
DEP1#/ GTLHI	232	I/O
DEP2#/ GTLHI	229	I/O
DEP3#/ GTLHI	237	I/O
DEP4#/ GTLHI	230	I/O
DEP5#/ GTLHI	235	I/O
DEP6#/ GTLHI	233	I/O
DEP7#/ GTLHI	227	I/O
DRDY#	226	I/O
GND	1	V
GND	2	V
GND	11	V
GND	14	V
GND	17	V
GND	20	V
GND	24	V
GND	27	V
GND	30	V
GND	33	V
GND	36	V
GND	41	V
GND	43	V
GND	46	V
GND	49	V
GND	52	V
GND	55	V

Table 24. 82452 KX/GX DP Alphabetical Pin List (240-Pin QFP) (Continued)

Signal	Pin#	Type
GND	59	V
GND	60	V
GND	62	V
GND	66	V
GND	69	V
GND	72	V
GND	75	V
GND	78	V
GND	80	V
GND	85	V
GND	88	V
GND	91	V
GND	94	V
GND	97	V
GND	101	V
GND	104	V
GND	107	V
GND	110	V
GND	121	V
GND	122	V
GND	131	V
GND	139	V
GND	149	V
GND	161	V
GND	170	V
GND	179	V
GND	180	V
GND	190	V
GND	200	V
GND	205	V
GND	221	V
GND	225	V
GND	228	V

Table 24. 82452 KX/GX DP Alphabetical Pin List (240-Pin QFP) (Continued)

Signal	Pin#	Type
GND	231	V
GND	234	V
GTLREFV	239	I
MDE0	112	I/O
MDE1	113	I/O
MDE2	114	I/O
MDE3	115	I/O
MDE4	116	I/O
MDE5	117	I/O
MDE6	118	I/O
MDE7	119	I/O
MDE8	123	I/O
MDE9	124	I/O
MDE10	125	I/O
MDE11	126	I/O
MDE12	127	I/O
MDE13	128	I/O
MDE14	129	I/O
MDE15	130	I/O
MDE16	132	I/O
MDE17	133	I/O
MDE18	134	I/O
MDE19	135	I/O
MDE20	136	I/O
MDE21	137	I/O
MDE22	138	I/O
MDE23	141	I/O
MDE24	142	I/O
MDE25	143	I/O
MDE26	144	I/O
MDE27	145	I/O
MDE28	146	I/O
MDE29	147	I/O

Table 24. 82452 KX/GX DP Alphabetical Pin List (240-Pin QFP) (Continued)

Signal	Pin#	Type
MDE30	148	I/O
MDE31	151	I/O
MDE32	152	I/O
MDE33	153	I/O
MDE34	154	I/O
MDE35	155	I/O
MDE36	156	I/O
MDE37	157	I/O
MDE38	158	I/O
MDE39	159	I/O
MDE40	162	I/O
MDE41	163	I/O
MDE42	164	I/O
MDE43	165	I/O
MDE44	166	I/O
MDE45	167	I/O
MDE46	168	I/O
MDE47	169	I/O
MDE48	171	I/O
MDE49	172	I/O
MDE50	173	I/O
MDE51	174	I/O
MDE52	175	I/O
MDE53	176	I/O
MDE54	177	I/O
MDE55	178	I/O
MDE56	182	I/O
MDE57	183	I/O
MDE58	184	I/O
MDE59	185	I/O
MDE60	186	I/O
MDE61	187	I/O
MDE62	188	I/O

Table 24. 82452 KX/GX DP Alphabetical Pin List (240-Pin QFP) (Continued)

Signal	Pin#	Type
MDE63	189	I/O
MDE64	191	I/O
MDE65	192	I/O
MDE66	193	I/O
MDE67	194	I/O
MDE68	195	I/O
MDE69	196	I/O
MDE70	197	I/O
MDE71	198	I/O
MDRDY0#	111	O
MDRDY1#	150	O
MEMCMD0#	206	I/O
MEMCMD1#	207	I/O
MEMCMD2#	208	I/O
MEMCMD3#	209	I/O
MEMCMD4#	210	I/O
MEMCMD5#	211	I/O

Table 24. 82452 KX/GX DP Alphabetical Pin List (240-Pin QFP) (Continued)

Signal	Pin#	Type
MEMCMD6#	212	I/O
MEMCMD7#	213	I/O
MEMERR0#	214	O
MEMERR1#	215	O
MIRST#	238	I
RESERVED	199	
SYSCMD#0	217	I
SYSCMD#1	218	I
SYSCMD#2	219	I
SYSCMD#3	222	I
SYSCMD#4	223	I
SYSDEN#	224	I
SYSERR#/ RESERVED	216	O
TCK	3	I
TD0	7	O
TDI	6	I

Table 24. 82452 KX/GX DP Alphabetical Pin List (240-Pin QFP) (Continued)

Signal	Pin#	Type
TESTHI	4	
TESTLO	202	
TESTLO	203	
TMS	5	I
TRST#	8	I
VCC3	21	V
VCC3	40	V
VCC3	61	V
VCC3	81	V
VCC3	100	V
VCC3	120	V
VCC3	140	V
VCC3	160	V
VCC3	181	V
VCC3	201	V
VCC3	220	V
VCC3	240	V

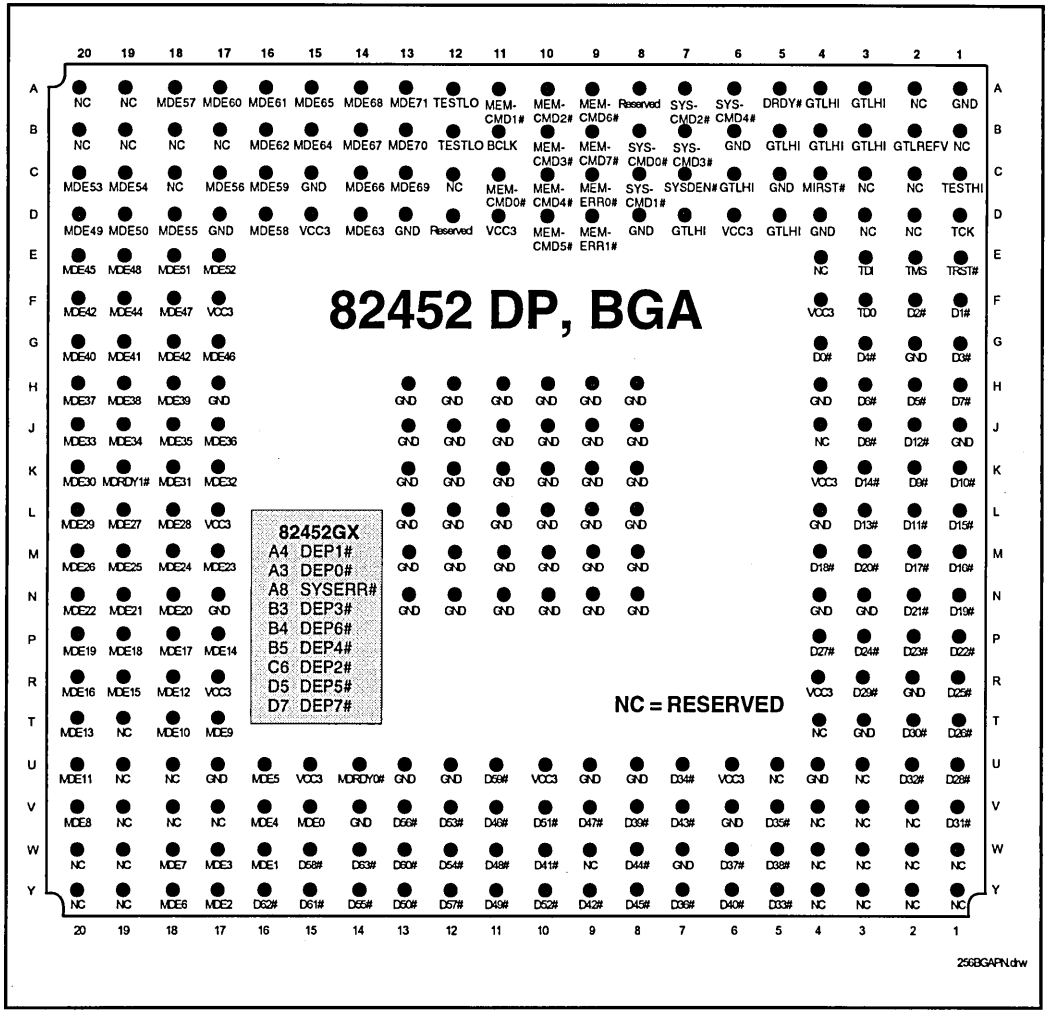

Figure 9. 82452 DP Pin Assignment (256 BGA)

Table 25. 82452KX/GX DP Alphabetical Pin List (256 BGA)

Signal	Ball#	Type
BCLK	B11	I
D0#	G4	I/O
D1#	F1	I/O
D2#	F2	I/O
D3#	G1	I/O
D4#	G3	I/O
D5#	H2	I/O
D6#	H3	I/O
D7#	H1	I/O
D8#	J3	I/O
D9#	K2	I/O
D10#	K1	I/O
D11#	L2	I/O
D12#	J2	I/O
D13#	L3	I/O
D14#	K3	I/O
D15#	L1	I/O
D16#	M1	I/O
D17#	M2	I/O
D18#	M4	I/O
D19#	N1	I/O
D20#	M3	I/O
D21#	N2	I/O
D22#	P1	I/O
D23#	P2	I/O
D24#	P3	I/O
D25#	R1	I/O
D26#	T1	I/O
D27#	P4	I/O
D28#	U1	I/O
D29#	R3	I/O
D30#	T2	I/O
D31#	V1	I/O

Table 25. 82452KX/GX DP Alphabetical Pin List (256 BGA) (Continued)

Signal	Ball#	Type
D32#	U2	I/O
D33#	Y5	I/O
D34#	U7	I/O
D35#	V5	I/O
D36#	Y7	I/O
D37#	W6	I/O
D38#	W5	I/O
D39#	V8	I/O
D40#	Y6	I/O
D41#	W10	I/O
D42#	Y9	I/O
D43#	V7	I/O
D44#	W8	I/O
D45#	Y8	I/O
D46#	V11	I/O
D47#	V9	I/O
D48#	W11	I/O
D49#	Y11	I/O
D50#	Y13	I/O
D51#	V10	I/O
D52#	Y10	I/O
D53#	V12	I/O
D54#	W12	I/O
D55#	Y14	I/O
D56#	V13	I/O
D57#	Y12	I/O
D58#	W15	I/O
D59#	U11	I/O
D60#	W13	I/O
D61#	Y15	I/O
D62#	Y16	I/O
D63#	W14	I/O
DRDY#	A5	I/O

Table 25. 82452KX/GX DP Alphabetical Pin List (256 BGA) (Continued)

Signal	Ball#	Type
DEP0#/ GTLHI	A3	I/O
DEP1#/ GTLHI	A4	I/O
DEP2#/ GTLHI	C6	I/O
DEP3#/ GTLHI	B3	I/O
DEP4#/ GTLHI	B5	I/O
DEP5#/ GTLHI	D5	I/O
DEP6#/ GTLHI	B4	I/O
DEP7#/ GTLHI	D7	I/O
GTLREFV	B2	I
MDRDY0#	U14	O
MDRDY1#	K19	O
MDE0	V15	I/O
MDE1	W16	I/O
MDE2	Y17	I/O
MDE3	W17	I/O
MDE4	V16	I/O
MDE5	U16	I/O
MDE6	Y18	I/O
MDE7	W18	I/O
MDE8	V20	I/O
MDE9	T17	I/O
MDE10	T18	I/O
MDE11	U20	I/O
MDE12	R18	I/O
MDE13	T20	I/O

Table 25. 82452KX/GX DP Alphabetical Pin List (256 BGA) (Continued)

Signal	Ball#	Type
MDE14	P17	I/O
MDE15	R19	I/O
MDE16	R20	I/O
MDE17	P18	I/O
MDE18	P19	I/O
MDE19	P20	I/O
MDE20	N18	I/O
MDE21	N19	I/O
MDE22	N20	I/O
MDE23	M17	I/O
MDE24	M18	I/O
MDE25	M19	I/O
MDE26	M20	I/O
MDE27	L19	I/O
MDE28	L18	I/O
MDE29	L20	I/O
MDE30	K20	I/O
MDE31	K18	I/O
MDE32	K17	I/O
MDE33	J20	I/O
MDE34	J19	I/O
MDE35	J18	I/O
MDE36	J17	I/O
MDE37	H20	I/O
MDE38	H19	I/O
MDE39	H18	I/O
MDE40	G20	I/O
MDE41	G19	I/O
MDE42	F20	I/O

Table 25. 82452KX/GX DP Alphabetical Pin List (256 BGA) (Continued)

Signal	Ball#	Type
MDE43	G18	I/O
MDE44	F19	I/O
MDE45	E20	I/O
MDE46	G17	I/O
MDE47	F18	I/O
MDE48	E19	I/O
MDE49	D20	I/O
MDE50	D19	I/O
MDE51	E18	I/O
MDE52	E17	I/O
MDE53	C20	I/O
MDE54	C19	I/O
MDE55	D18	I/O
MDE56	C17	I/O
MDE57	A18	I/O
MDE58	D16	I/O
MDE59	C16	I/O
MDE60	A17	I/O
MDE61	A16	I/O
MDE62	B16	I/O
MDE63	D14	I/O
MDE64	B15	I/O
MDE65	A15	I/O
MDE66	C14	I/O
MDE67	B14	I/O
MDE68	A14	I/O
MDE69	C13	I/O
MDE70	B13	I/O
MDE71	A13	I/O

Table 25. 82452KX/GX DP Alphabetical Pin List (256 BGA) (Continued)

Signal	Ball#	Type
MEMCMD0#	C11	I/O
MEMCMD1#	A11	I/O
MEMCMD2#	A10	I/O
MEMCMD3#	B10	I/O
MEMCMD4#	C10	I/O
MEMCMD5#	D10	I/O
MEMCMD6#	A9	I/O
MEMCMD7#	B9	I/O
MEMERR0#	C9	O
MEMERR1#	D9	O
MIRST#	C4	I
RESERVED	D12	
SYSDEN#	C7	I
SYSCMD0#	B8	I
SYSCMD1#	C8	I
SYSCMD2#	A7	I
SYSCMD3#	B7	I
SYSCMD4#	A6	I
SYSERR#/ RESERVED	A8	O
TCK	D1	I
TD0	F3	O
TDI	E3	I
TESTHI	C1	I
TESTLO	A12	I
TESTLO	B12	I
TMS	E2	I
TRST#	E1	I

Table 26. 82452KX/GX DP VCC, Ground, and No Connect Pins (256 BGA)

Signal	BGA Ball#
VCC3	F4, K4, R4, U6, U10, U15, R17, L17, F17, D6, D11, D15
GND	A1, D4, D8, D13, D17, H4, N4, U4, U8, U13, U17, H17, N17, G2, J1, L4, N3, R2, T3, V6, W7, U9, U12, V14, C15, B6, C5, [H8, H9, H10, H11, H12, H13, J8, J9, J10, J11, J12, J13, K8, K9, K10, K11, K12, K13, L8, L9, L10, L11, L12, L13, M8, M9, M10, M11, M12, M13, N8, N9, N10, N11, N12, N13](1)
NC	B1, C2, D2, D3, E4, J4, T4, U3, V2, W1, V3, W2, Y1, W3, Y2, W4, V4, U5, Y3, Y4, W9, V17, Y19, V18, W19, Y20, T19, W20, V19, U19, U18, B20, C18, B19, A20, C3, A2, C12, A19, B17, B18

1. Ground pins within the brackets are thermal connections.

4.3 82451KX/82451GX (MIC) Pin Assignment

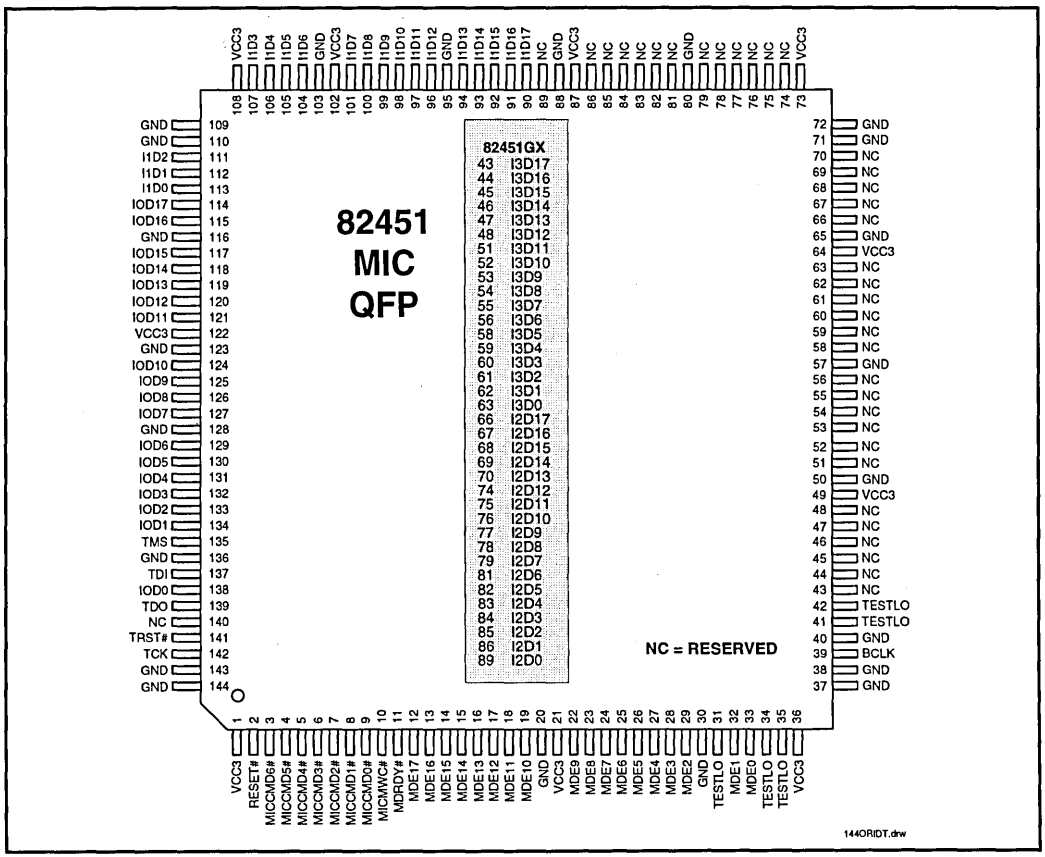

Figure 10. 82451KX/GX MIC Pin Assignment (144-Pin QFP)

Table 27. 82451 KX/GX MIC Alphabetical Pin List (144-Pin QFP)

Signal	Pin#	Type
BCLK	39	I
GND	20	V
GND	30	V
GND	37	V
GND	38	V
GND	40	V
GND	50	V
GND	57	V
GND	65	V
GND	71	V
GND	72	V
GND	80	V
GND	88	V
GND	95	V
GND	103	V
GND	109	V
GND	110	V
GND	116	V
GND	123	V
GND	128	V
GND	136	V
GND	143	V
GND	144	V
I0D0	138	I/O
I0D1	134	I/O
I0D2	133	I/O
I0D3	132	I/O
I0D4	131	I/O
I0D5	130	I/O
I0D6	129	I/O
I0D7	127	I/O
I0D8	126	I/O
I0D9	125	I/O

Table 27. 82451 KX/GX MIC Alphabetical Pin List (144-Pin QFP) (Continued)

Signal	Pin#	Type
I0D10	124	I/O
I0D11	121	I/O
I0D12	120	I/O
I0D13	119	I/O
I0D14	118	I/O
I0D15	117	I/O
I0D16	115	I/O
I0D17	114	I/O
I1D0	113	I/O
I1D1	112	I/O
I1D2	111	I/O
I1D3	107	I/O
I1D4	106	I/O
I1D5	105	I/O
I1D6	104	I/O
I1D7	101	I/O
I1D8	100	I/O
I1D9	99	I/O
I1D10	98	I/O
I1D11	97	I/O
I1D12	96	I/O
I1D13	94	I/O
I1D14	93	I/O
I1D15	92	I/O
I1D16	91	I/O
I1D17	90	I/O
I2D0/NC	89	I/O
I2D1/NC	86	I/O
I2D2/NC	85	I/O
I2D3/NC	84	I/O
I2D4/NC	83	I/O
I2D5/NC	82	I/O
I2D6/NC	81	I/O

Table 27. 82451 KX/GX MIC Alphabetical Pin List (144-Pin QFP) (Continued)

Signal	Pin#	Type
I2D7/NC	79	I/O
I2D8/NC	78	I/O
I2D9/NC	77	I/O
I2D10/NC	76	I/O
I2D11/NC	75	I/O
I2D12/NC	74	I/O
I2D13/NC	70	I/O
I2D14/NC	69	I/O
I2D15/NC	68	I/O
I2D16/NC	67	I/O
I2D17/NC	66	I/O
I3D0/NC	63	I/O
I3D1/NC	62	I/O
I3D2/NC	61	I/O
I3D3/NC	60	I/O
I3D4/NC	59	I/O
I3D5/NC	58	I/O
I3D6/NC	56	I/O
I3D7/NC	55	I/O
I3D8/NC	54	I/O
I3D9/NC	53	I/O
I3D10/NC	52	I/O
I3D11/NC	51	I/O
I3D12/NC	48	I/O
I3D13/NC	47	I/O
I3D14/NC	46	I/O
I3D15/NC	45	I/O
I3D16/NC	44	I/O
I3D17/NC	43	I/O
MDE0	33	I/O
MDE1	32	I/O
MDE2	29	I/O
MDE3	28	I/O

Table 27. 82451 KX/GX MIC Alphabetical Pin List (144-Pin QFP) (Continued)

Signal	Pin#	Type
MDE4	27	I/O
MDE5	26	I/O
MDE6	25	I/O
MDE7	24	I/O
MDE8	23	I/O
MDE9	22	I/O
MDE10	19	I/O
MDE11	18	I/O
MDE12	17	I/O
MDE13	16	I/O
MDE14	15	I/O
MDE15	14	I/O
MDE16	13	I/O
MDE17	12	I/O
MDRDY#	11	I

Table 27. 82451 KX/GX MIC Alphabetical Pin List (144-Pin QFP) (Continued)

Signal	Pin#	Type
MICCMD0#	9	I
MICCMD1#	8	I
MICCMD2#	7	I
MICCMD3#	6	I
MICCMD4#	5	I
MICCMD5#	4	I
MICCMD6#	3	I
MICMWC#	10	I
MIRST#	2	I
RESERVED	140	NC
TCK	142	I
TDI	137	I
TDO	139	O
TESTLO	31	I
TESTLO	34	I

Table 27. 82451 KX/GX MIC Alphabetical Pin List (144-Pin QFP) (Continued)

Signal	Pin#	Type
TESTLO	35	I
TESTLO	41	I
TESTLO	42	I
TMS	135	I
TRST#	141	I
VCC3	1	V
VCC3	21	V
VCC3	36	V
VCC3	49	V
VCC3	64	V
VCC3	73	V
VCC3	87	V
VCC3	102	V
VCC3	108	V
VCC3	122	V

4.4 82453KX/82453GX (DC) Package Dimensions

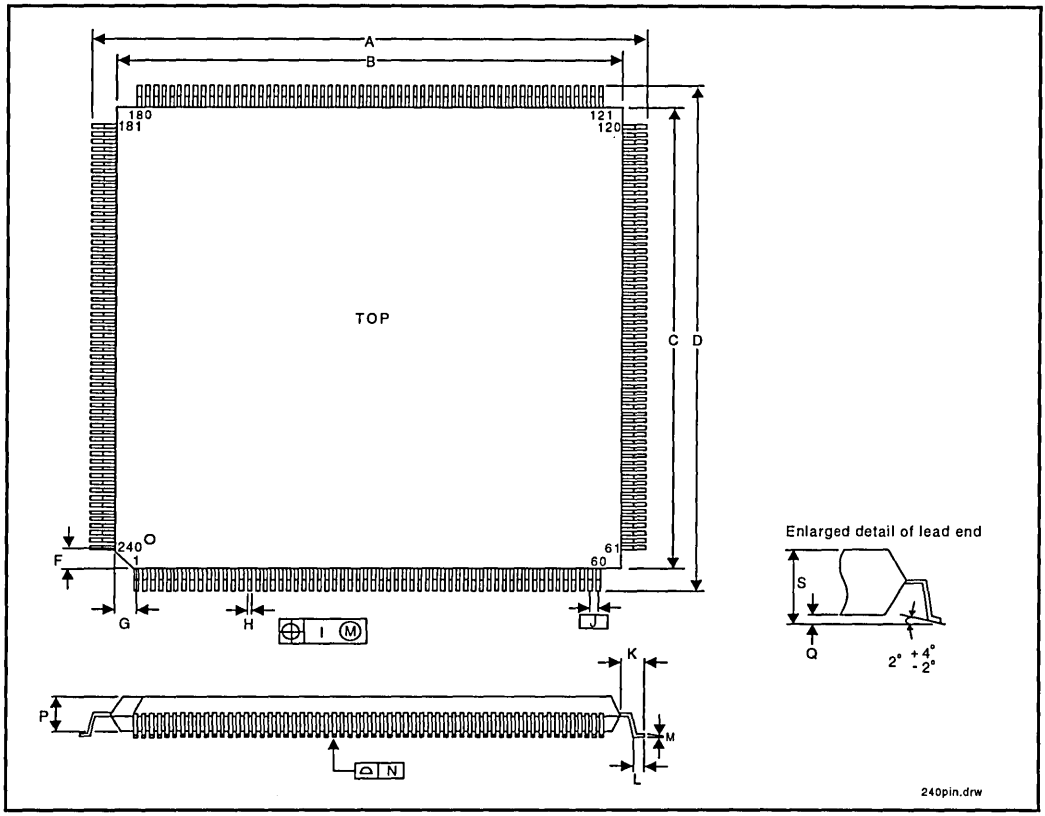

Figure 11. DC Package Physical Dimensions (208-Pin QFP)

Table 28. DC Package Physical Dimensions (208-Pin QFP)

Dim	Millimeters	Inches
A	30.6 ± 0.4	1.205 ± 0.016
B	28.0 ± 0.2	1.102 ± 0.008
C	28.0 ± 0.2	1.102 ± 0.008
D	30.6 ± 0.4	1.205 ± 0.016
F	1.25	0.49
G	1.25	0.049
H	0.20 ± 0.05	0.008 ± 0.002
I	0.08	0.003

Dim	Millimeters	Inches
J	0.5 (T.P.)	0.020 (T.P.)
K	1.3 ± 0.2	0.051 ± 0.008
L	0.5 ± 0.2	0.020 ±0.008
M	0.15 ± 0.05	0.006 ± 0.002
N	0.075	0.003
P	3.7	0.0126
Q	0.4 ± 0.1	0.016 ± 0.004
S	3.8 max	0.150 max

4.5 82452KX/82452GX (DP) Package Dimensions

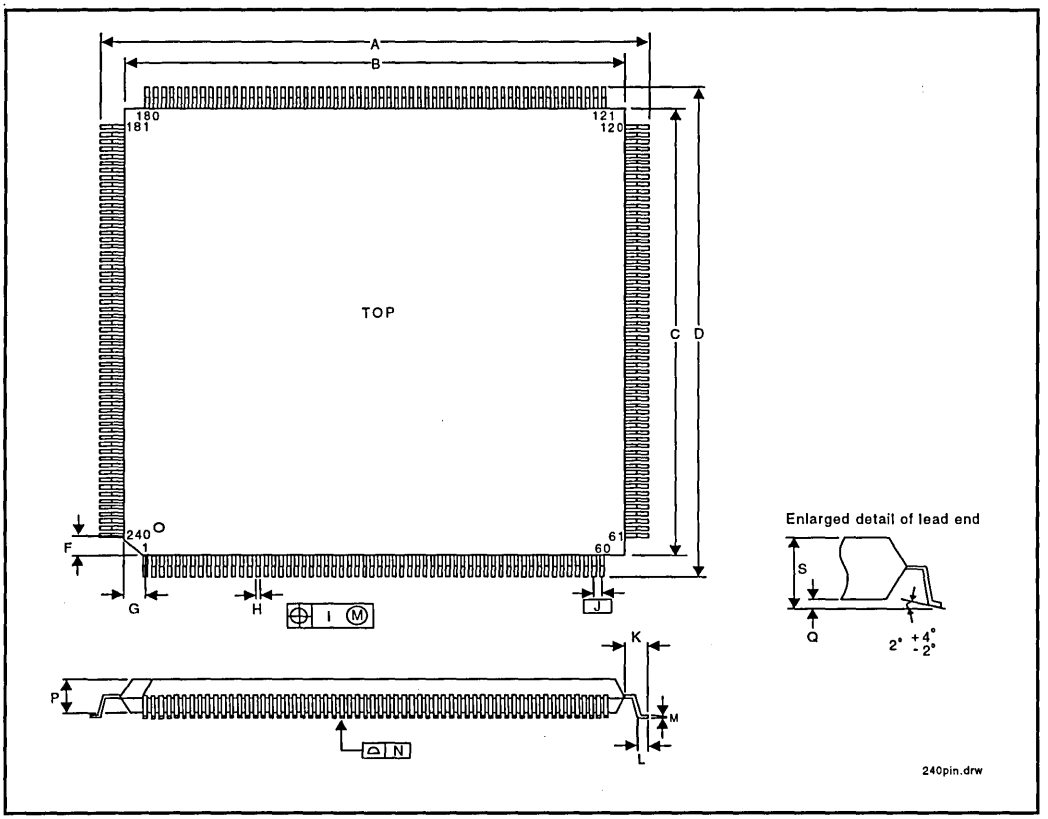

Figure 12. DP Package Physical Dimensions (240-Pin QFP)

Table 29. DP Physical Package Dimensions (240-Pin QFP)

Dim	Millimeters	Inches
A	34.6 ± 0.3	1.362 ±0.012
B	32.0 ± 0.2	1.260 ±0.008
C	32.0 ± 0.2	1.260 ±0.008
D	34.6 ± 0.3	1.362 ±0.012
F	1.25	0.049
G	1.25	0.049
H	0.22 ± 0.08	0.009 ± 0.003
I	0.08	0.003

Dim	Millimeters	Inches
J	0.5 (T.P.)	0.020 (T.P.)
K	1.3 ± 0.2	0.051 ±0.008
L	0.5 ± 0.2	0.020 ±0.008
M	0.15 ± 0.05	0.006 ±0.002
N	0.1	0.004
P	3.2	0.126
Q	0.4 ± 0.1	0.016 ± 0.004
S	3.8 max	0.150 max

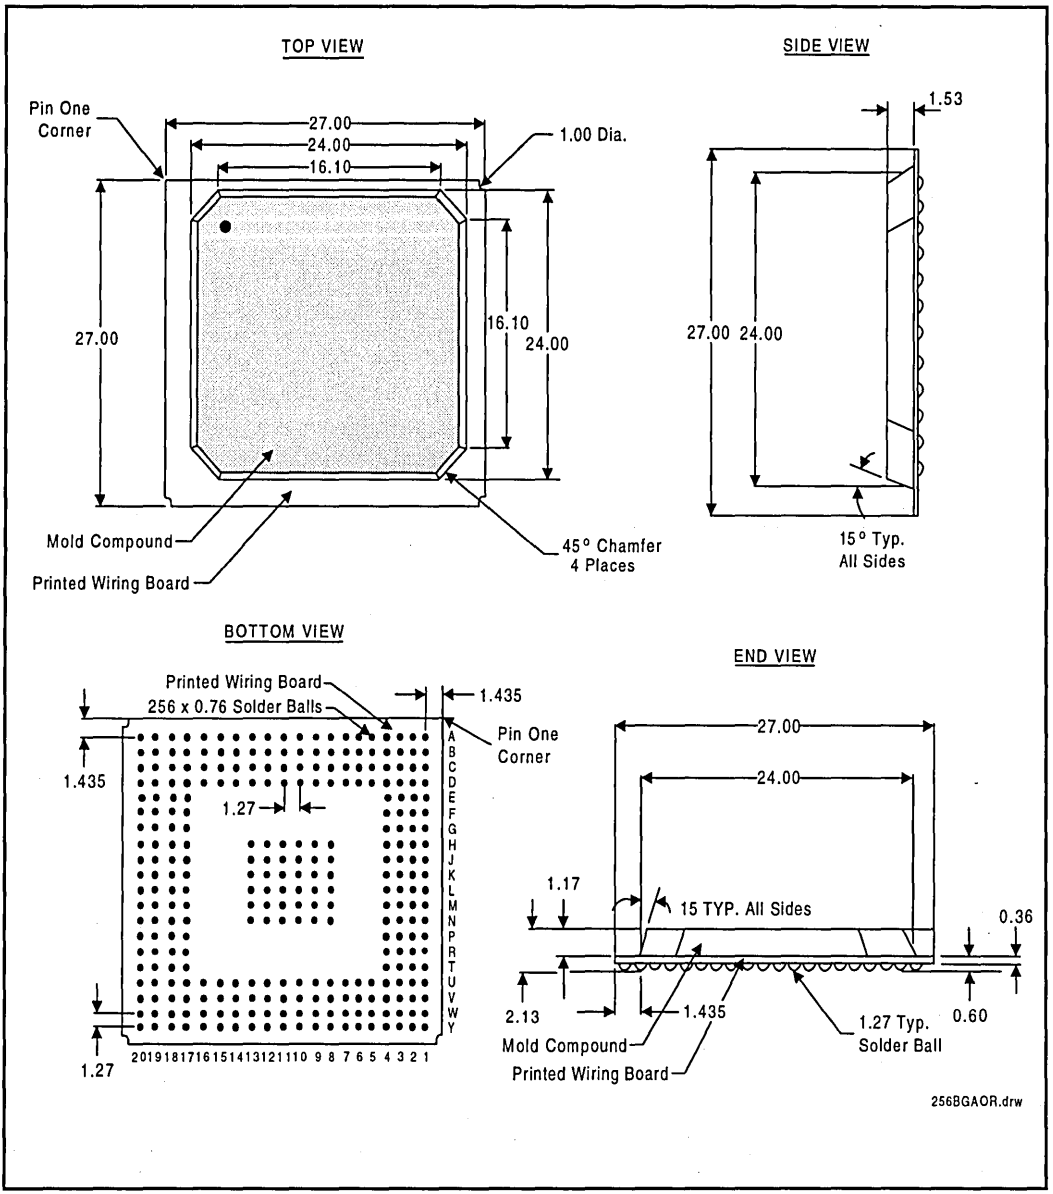

Figure 13. DP Package Dimensions (256 BGA)

4.6 82451KX/82451GX (MIC) Package Dimensions

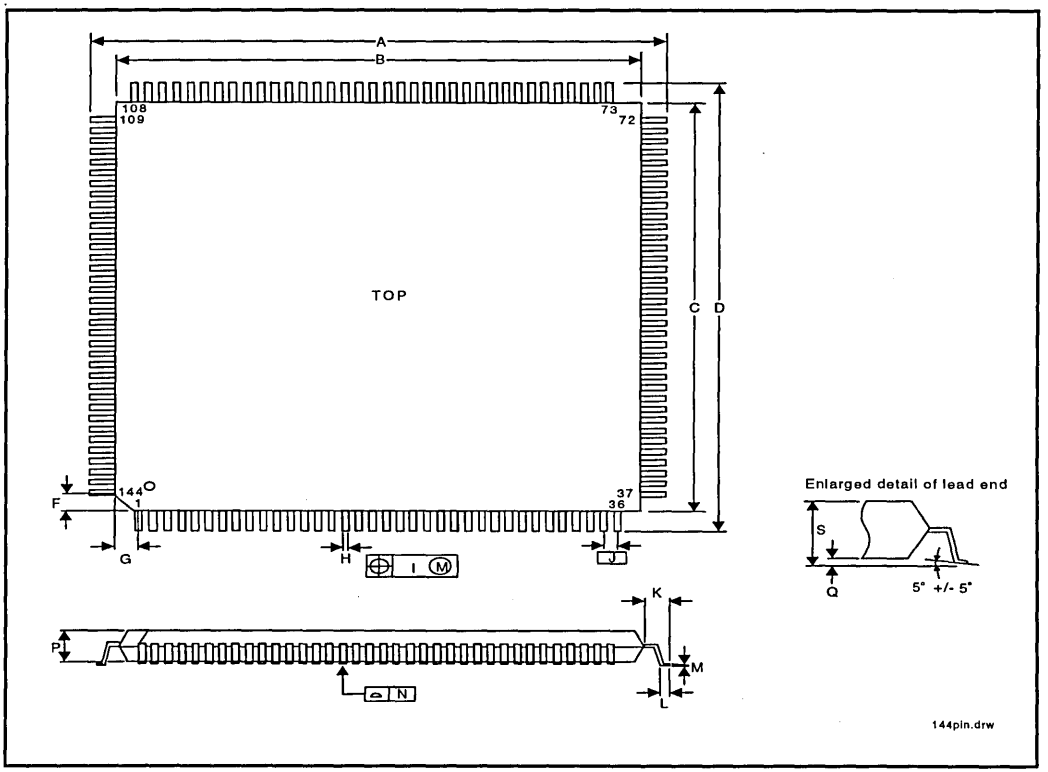

Figure 14. MIC Package Physical Dimensions (144-Pin QFP)

Table 30. MIC Package Physical Dimensions (144-Pin QFP)

Dim	Millimeters	Inches
A	22.0 ± 0.4	0.866 ± 0.016
B	20.0 ± 0.2	0.787 ± 0.008
C	20.0 ± 0.2	0.787 ± 0.008
D	22.0 ± 0.4	0.866 ± 0.016
F	1.25	0.049
G	1.25	0.049
H	0.20 ± 0.05	.008
I	0.08	0.003

Dim	Millimeters	Inches
J	0.5 (T.P.)	0.020 (T.P.)
K	1.0 ± 0.2	0.039 ± 0.008
L	0.5 ± 0.2	0.020 ± 0.008
M	0.15 ± 0.05	0.006 ± 0.002
N	0.10	0.004
P	2.7	0.106
Q	0.1 ± 0.1	0.004 ± 0.004
S	3.0 max	0.119 max

Chapter 4

PCIset Electrical Specifications

1.0 ELECTRICAL CHARACTERISTICS

This section contains the electrical characteristics associated with the Intel 450KX/GX PCIsets. This includes information on power consumption, AC and DC signal timing information for GTL+ and non-GTL+ signals. To ensure functionality and reliability, the Intel 450KX/GX PCIsets are specified for proper operation when T_C (case temperature) is within the specified range of 0°C to 85°C. For more information on measuring T_C in your system, please refer to the *Pentium Pro Family Developer's Manual: Specifications*. Most of the Pentium Pro Processor signals use a variation of the low voltage GTL (Gunning Transceiver Logic). For more information on this technology refer to the *Pentium Pro Family Developer's Manual: Specifications*.

1.1 Test Pins and Unused Pins

For reliable operation, always connect unused inputs to an appropriate signal level. Unused GTL+ inputs should be pulled-up to V_{TT} with a 25Ω to 10KΩ resistor. Unused active low 3.3V inputs should be pulled-up to 3.3V with a 10KΩ resistor. Unused active high inputs should be pulled to ground (VSS) with a 1KΩ resistor. When connecting bidirectional signals to power or ground, a resistor must be used. When tying any signal to power or ground, a resistor will also allow for fully testing the processor and PCIset after board assembly. It is suggested that ~10KΩ resistors be used for pull-ups and ~1KΩ resistors be used as pull-downs.

In the 450KX/GX QFP and BGA Pinlists, in Section 3.0 of Chapter 2 and Chapter 3, all signals labeled "TESTLO" should be pulled to GND with a 1KΩ resistor. Signals labeled "TESTHI" should be pulled up to $V_{CC}3$ with a 10KΩ resistor. Signals labeled "GTLHI" should be pulled up to V_{TT} with a 10KΩ resistor. All RESERVED and NC pins must remain unconnected.

1.2 Signal Groups

In order to simplify the following discussion, signals have been combined into groups of like characteristics in Table 1. See Section 3.0 in Chapter 2 and Chapter 3 for a description of the signals and their functions. All pins of the Intel 450GX PCIset are listed here. Not all of these signals are available on the Intel 450KX PCIset. The following pins do not exist on the Intel 450KX PCIset: CASA[7:4]#, CASB[7:4]#, DEP#[7:0], I2D[17:0], I3D[17:0], IOGNT#, IOREQ#, OMCNUM, RASA[7:4]#, RASB[7:4]#, SYSERR#.

Table 1. Signal Groups

Pin Group	Signals	Notes
GTL+ Input	*A[35:3]#, ADS#, AP[1:0]#, REQ[4:0]#, RP#, RESET#*	(1)
GTL+ Output	BREQ0#	
GTL+ I/O	*A[35:3]#, ADS#,* AERR#, *AP[1:0]#,* BERR#, BINIT#, BNR#, BPRI#, D[63:0]#, DBSY#, DEFER#, DEP[7:0]#, DRDY#, HIT#, HITM#, LOCK#, *REQ[4:0]#,* RESET#, *RP#,* RS[2:0]#, RSP#, TRDY#	(1)
CMOS Input, 5V Tolerant	IOGNT#, FLSHBF#, MEMREQ#, PCLKIN, PGNT#, RECVEN	
CMOS Input, 3.3V	BCLK, MDRDY[1:0]#, MEMCMD[7:0]#, MEMERR[1:0]#, *MIRST#,* OMCNUM, PWRGD, SYSCMD[4:0]#, SYSDEN#, TCK, TDI, TMS, TRST#	
CMOS Output, 6mA, 5V Tolerant	FLUSH#, INIT#, PCIRST#, SMIACT#	
CMOS Output, 6mA, 3.3V	MEMERR[1:0]#, SYSCMD[4:0]#, SYSERR#, TDO	

Table 1. Signal Groups (Continued)

Pin Group	Signals	Notes
CMOS Output, 12mA, 5V Tolerant	CRESET#, MEMACK#, PREQ#	
CMOS Output, 12mA, 3.3V	CASA[7:0]#, CASB[7:0]#, MA0[12:0], MA1[12:0], MICMWC[1:0]#, RASA[7:0]#, RASB[7:0]#, RESET#, SBCERR#, SYSDEN#, WE[1:0]#	(2)
CMOS Output, 18mA, 3.3V	MDRDY[1:0]#, MICCMD[6:0]#, *MIRST#*	
CMOS I/O, 6mA, 5V Tolerant	PCLK	
CMOS I/O, 6mA, 3.3V	MEMCMD[7:0]#	
CMOS I/O, 12mA, 5V Tolerant	I0D[17:0], I1D[17:0], I2D[17:0], I3D[17:0]	
CMOS I/O, 12mA, 3.3V	IOREQ#, MDE[71:0]	
PCI Signals, 24mA	AD[31:0], C/BE[3:0]#, DEVSEL#, FRAME#, IRDY#, PAR, PERR#, PLOCK#, PTRDY#, SERR#, STOP#	(3)
Power	GND, GTLVREF, VCC_3, VCC_{PCI}	

1. *Italicized* signals are inputs on one device and I/O on another device: A[35:3]#, ADS#, AP[1:0]#, REQ[4:0]#, RP#, and RESET# are inputs to the DC, and I/O on the PB. MIRST# is an output from the DC and an input to the DP and MICs.
2. SBCERR# is an open-drain signal.
3. PCI signals are both 3.3V and 5V tolerant. The drive and receive strength for the PCI signals is set by the VCC_{PCI} input (PCIBus voltage). For additional details, see the PCI Local Bus Specification, Rev 2.0.

1.2.1 THE POWER GOOD SIGNAL— PWRGD

PWRGD is a 3.3V input to the PCI Bridge and memory controller components. It is expected that this signal is a clean indication that the clocks and the 3.3V, VCC_{PCI} supplies are within their specifications. 'Clean' implies that PWRGD will remain low, (capable of sinking leakage current) without glitches, from the time that the power supplies are turned on until they become valid. The signal will then transition monotonically to a high (3.3V) state with the transition not taking longer than 100ns. PWRGD needs to be negated for at least 10 BCLKs before this transition from low to high can take place. Figure 1 illustrates the relationship of PWRGD to BCLK and the system reset signals.

The PWRGD inputs to the Intel 450KX/GX PCIsets and to the Pentium Pro Processor(s) should be driven with an "AND" of 'Power-Good' signals from the 5V, 3.3V and VCCP supplies. The output of this logic should be a 3.3V level and should have a pull-down resistor at the output to cover the period when this logic is not receiving power.

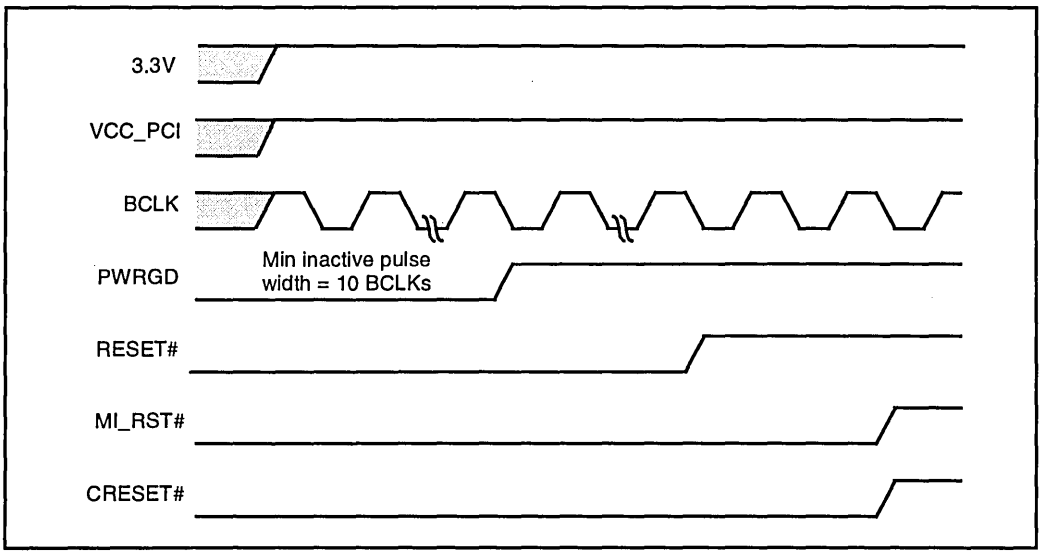

Figure 1. PWRGD Relationship

1.3 Maximum Ratings

Table 2 contains stress ratings for the Intel 450KX/GX PCIsets. Functional operation at the absolute maximum and minimum is neither implied nor guaranteed. The PCIsets should not receive a clock while subjected to these conditions. Functional operating conditions are given in the AC and DC tables. Extended exposure to the maximum ratings may affect device reliability. Furthermore, although the Intel 450KX/GX PCIsets contain protective circuitry to resist damage from static discharge, one should always take precautions to avoid high static voltages or electric fields

Table 2. Absolute Maximum Ratings

Sym	Parameter	Min	Max	Unit	Notes
$V_{CC}3$	3.3V Supply Voltage with respect to VSS	-0.5	4.3	V	
V_{IN}	GTL+ Buffer DC Input Voltage with respect to VSS	-0.5	$V_{CC}3$ + 0.5 (not to exceed 4.3)	V	(1)
$V_{IN}3$	3.3V DC Input Voltage with respect to VSS	-0.5	$V_{CC}3$ + 0.9 (not to exceed 4.7)	V	(2)
$V_{IN}5$	5V Tolerant DC Input Voltage with respect to VSS	-0.5	$V_{CC\text{-}PCI}$ + 0.5	V	(3)
T_{STOR}	Storage Temperature	-65	150	°C	

1. Parameter applies to GTL+ Signal Group Only
2. Parameter applies to 3.3V and JTAG signal groups only
3. Parameter applies to 5V tolerant signal groups and PCI signals only. $V_{CC\text{-}PCI}$ is the voltage level of the PCI Bus.

1.4 DC Specifications

The following tables list the DC specifications associated with the Intel 450KX/GX PCIsets. Care should be taken to read any notes associated with each parameter listed.

Table 3. Voltage and Temperature Specifications

Sym	Parameter	Min	Typ	Max	Unit	Notes
$V_{CC}3$	Device VCC	3.13	3.3	3.46	V	± 5%
$V_{CC\text{-}PCI}$ (3.3)	PCI VCC for 3.3 V PCI Operation	3.0	3.3	3.6	V	± 10%
$V_{CC\text{-}PCI}$ (5)	PCI VCC for 5.0 V PCI Operation	4.5	5.0	5.5	V	± 10%
$I_{CC\text{-}PCI}$	Clamping Diode Leakage Current			2	mA	At 33MHz
T_C	Operating Case Temperature	0		85	°C	

Table 4. Power Specifications

Sym	Parameter	Max	Unit	Notes
P_{MAX}	PB Max Power Dissipation	2.4	W	(1)
P_{MAX}	DC Max Power Dissipation	1.13	W	(1)
P_{MAX}	DP Max Power Dissipation	2.3	W	(1)
P_{MAX}	MIC Max Power Dissipation	0.69	W	
$I_{CC}3$	PB Max Power Supply Current	640	mA	
$I_{CC}3$	DC Max Power Supply Current	370	mA	
Symbol	**Parameter**	**Max**	**Unit**	**Notes**
$I_{CC}3$	DP Max Power Supply Current	460	mA	
$I_{CC}3$	MIC Max Power Supply Current	200	mA	
I_{SS}	PB Max V_{SS} Current	4800	mA	(2)
I_{SS}	DC Max V_{SS} Current	600	mA	(2)
I_{SS}	DP Max V_{SS} Current	4400	mA	(2)
I_{SS}	MIC Max V_{SS} Current	200	mA	

1. Includes power dissipated in the GTL+ buffers.
2. I_{SS} is the maximum supply current consumption when all GTL+ signals are low. It is the sum of ICC3 and GTL+ current.

Table 5 lists the specification for the GTL+ termination voltage (V_{TT}) and the GTL+ reference voltage (V_{REF}).

Table 5. GTL+ Bus DC Specifications

Sym	Parameter	Min	Typ	Max	Unit	Notes
V_{TT}	Bus Termination Voltage	1.35	1.5	1.65	V	± 10%
V_{REF}	Input Reference Voltage	2/3 V_{TT} - 2%	2/3 V_{TT}	2/3 V_{TT} +2%	V	± 2%

Some of the signals on the PB, DC, and DP are in the GTL+ signal group. These signals are specified to be terminated to 1.5V. The DC specifications for these signals are listed in Table 6.

Table 6. DC Specifications (GTL+ signal groups)

Sym	Parameter	Min	Max	Unit	Notes
V_{IL}	Input Low Voltage	-0.3	V_{REF} - 0.1	V	(1)
V_{IH}	Input High Voltage	V_{REF} + 0.1	1.8	V	(1)
V_{OL}	Output Low Voltage	0.30	0.55	V	(2)
V_{OH}	Output High Voltage	--	--	V	(3)
I_{OL}	Output Low Current	38	48	mA	(2)
I_{LI}	Input Leakage Current		+ 15	µA	(4)
I_{REF}	Reference Voltage Current		+ 15	µA	
I_{LO}	Output Leakage Current		+ 15	µA	(5)
C	Pin Capacitance		10	pF	(6)

1. Refer to the 450KX and 450GX Ringback Specification for additional information regarding noise limits.
2. Parameter measured into a 25Ω resistor to V_{TT} (1.5V).
3. The output high voltage level is determined by V_{TT}.
4. ($0 \le V_{IN} \le V_{TT}$)
5. ($0 \le V_{OUT} \le V_{TT}$)
6. Total of buffer and package parasitics.

The DC specifications for the non-GTL+ signal groups are listed in Table 7. A Specification covers 3.3V and 5V tolerant buffers, unless otherwise stated.

Table 7. DC Specifications (non-GTL+ groups)

Sym	Pin Group	Parameter	Min	Max	Unit	Notes
V_{IL}	CMOS Input	Input Low Voltage		0.8	V	
V_{IL-PCI}	PCI	Input Low Voltage	- 0.5	0.8	V	
V_{IH}	3.3V CMOS Input	Input High Voltage	2.0	3.6	V	(1)
V_{IH-5V}	5V Tolerant Input	Input High Voltage	2.0	5.0	V	
V_{IH-PCI}	PCI	Input High Voltage	2.0	V_{CC-PCI} +0.5	V	
V_{OL}	CMOS 6, 12, 18mA	Output Low Voltage		0.40	V	At 4mA
V_{OL-PCI}	PCI	Output Low Voltage		0.55	V	
V_{OH}	CMOS 6, 12, 18mA	Output High Voltage	V_{CC}3-0.1		V	
V_{OH-PCI}	PCI	Output High Voltage	2.4		V	
I_{OL-6}	CMOS 6mA	Output Low Current	6.0		mA	At V_{OL} = 0.4V
I_{OL-12}	CMOS 12mA	Output Low Current	12.0		mA	At V_{OL} = 0.4V
I_{OL-18}	CMOS 18mA	Output Low Current	18.0		mA	At V_{OL} = 0.4V
I_{OL-PCI}	PCI	Output Low Current	6.0		mA	
I_{IL}	CMOS Input	Input Leakage Current		±100	µA	(2)
I_{LO}	CMOS 6, 12, 18mA	Output Leakage Current		±10	µA	(3)
I_{IL-PCI}	PCI	Input Leakage Current		±70	µA	
I_{OL-PCI}	PCI	Output Leakage Current		±10	µA	
C_{IN}	All	Input Capacitance		10	pF	(4)
C_o	All	Output Capacitance		10	pF	
$C_{I/O}$	All	I/O Capacitance		10	pF	
C_{BCLK}	BCLK	BCLK Capacitance		8	pF	
C_{TCK}	TCK	TCK Capacitance		8	pF	

1. The Interleave databus signals (IxDxx) can interface to 3V or 5V DRAM. These signals are 5V tolerant.

2. ($0 \le V_{IN} \le V_{CC}3$)
   -100µA for pins with 50KΩ pullups, +100µA for pins with 50KΩ pulldowns.
   Pins with 50KΩ internal pullups: MICCMD#, MICMWC#, MDRDY#, RESET, MEMCMD#, SYSCMD#, MEMERR#, SYSERR#
   Pins with 50KΩ internal pulldowns: MDE[71:0]#, IxD[17:0]
3. ($0 \le V_{OUT} \le V_{CC}3$)
4. Except BCLK, TCK

1.5 AC Specifications

The following tables list the AC specifications associated with the Intel 450GX PCIsets. Care should be taken to read any notes associated with each parameter listed.

Table 8. AC Specifications (Clock signal groups)

Symbol	Parameter	Min	Max	Unit	Figure	Notes
f	BCLK Frequency	50.00	66.67	MHz		
T1	BCLK Period	15	20	ns	2.2	
T2	BCLK Period Stability		300	ps		(1)
T3	BCLK High Time	6		ns	2.2	At > 2.0V
T4	BCLK Low Time	6		ns	2.2	At < 0.8V
T5	BCLK Rise Time	0.3	1.5	ns	2.2	0.8-2.0V
T6	BCLK Fall Time	0.3	1.5	ns	2.2	2.0-0.8V
T7	PCLKIN Period	30		ns		(2)
T44	TCK Rise/Fall Time	0	50	ns		
T44a	TCK Frequency		16	MHz		

1. Measured at rising edge of BCLK crossing 1.5V.
2. Please refer to the PCI Specification 2.0 for more details.

Table 9. AC Specifications (GTL+ Signals)

Symbol	Parameter	Setup Min	Hold Min	Delay Min	Delay Max	Unit	Notes
T8	GTL+ Inputs	4.5	0.3			ns	(1)
T9	GTL+ Outputs			1.0	6.0	ns	(2)
T10	GTL+ I/O	4.5	0.3	1.0	6.0	ns	(1) (2)

1. Values correspond to input waveforms as specified in the GTL+ Interface Specification in *Pentium Pro Processor Developer's Manual: Specifications*. A waveform with less than 200mV of overdrive is acceptable if it meets the criteria in Figure 5. However, in this case, the minimum setup specification will be impacted by 1ns, making Setup Min 5.5ns. An example of a signal that would benefit from this derating is shown in Figure 6.
2. Outputs measured into 25Ω tied to 1.5V. Measurement made at 1.0V crossing.

Table 10, contains AC specifications for the CMOS signals. Please note that the PCI signals are specified with reference to the PCLKIN; the testability signals are specified with reference to the TCK, and all other signals reference the system bus clock, BCLK

Table 10. AC Specifications (CMOS signals)

Sym	Parameter	Setup Min	Hold Min	Delay Min	Delay Max	Unit	Notes
	PCI Side Signals						
T11	AD[31:0], C/BE[3:0]#, PTRDY#, STOP#, PLOCK#, DEVSEL#, PAR, IRDY#, FRAME#, PERR#, SERR#	7.0	0.0	2.0	11.0	ns	(1)
T11a	PREQ#			2.0	12.0	ns	(1)
T12	PGNT#	10.0	0.0			ns	(1)
	Compatibility Bridge Signals						
T13	FLSHBF#, MEMREQ#	12.0	0.0			ns	
T14	MEMACK#			2.0	12.0	ns	(2)
	Bridge To Bridge Signals						
T15	I/OREQ#	3.0	0.6	2.0	7.0	ns	(3)
T16	I/OGNT#	3.0	0.6			ns	(3)
	Memory Interface Signals—DP						
T17a	MDE[71:0] — Write Cycle	2.0	0.9	1.2	7.1	ns	(3)
T18	MDRDY[1:0]#			1.2	5.9	ns	(3)
T19	MEMCMD[7:0]#	1.9	0.8	1.4	5.9	ns	(3)
	Memory Interface Signals—DC						
T20	MICMWC[0:1]#			1.1	5.3	ns	(3)
T21	MICCMD[6:0]#			1.0	5.0	ns	(3)
T22	CASA[7:0]#, CASB[7:0]#			1.1	5.3	ns	(3) (4)
T23	RASA[7:0]#, RASB[7:0]#			1.1	5.3	ns	(3) (4)
T24	MA0[12:0], MA1[12:0]			1.1	5.3	ns	(3)
T25	WE0#, WE1#			1.1	5.3	ns	(3)
T26	MEMCMD[7:0]#	2.0	0.8	1.4	7.1	ns	(3)
T27	SBCERR#			1.2	5.5	ns	(3)
	Memory Interface Signals—MIC						
T28	MICMWC# — Write Cycle	4.3	0.5			ns	
T29	MICCMD[6:0]#	5.2	0.6			ns	
T30	MDRDY#	5.6	0.4			ns	
T31a	MDE[17:0] — Write Cycle	4.7	0.5			ns	
T31b	MDE[17:0] Delay — Read Cycle			1.8	8.0	ns	(3) (5)

Table 10. AC Specifications (CMOS signals) (Continued)

Sym	Parameter	Setup Min	Hold Min	Delay Min	Delay Max	Unit	Notes
T31c	MDE[17:0] flow through during window — Read Cycle			1.7	7.5	ns	(3) (5)
T31d	MDE[17:0] flow through pre-window — Read Cycle			1.9	9.0	ns	(3) (5)
T31e	Window start time from BCLK rising — Read Cycle				2.0	ns	(3) (5)
T31f	Window end time after BCLK rising — Read Cycle			0.5		ns	(3) (5)
T32a	I[3:0]Dx	1.2	1.3			ns	(5)
T32b	I[3:0]Dx delay — Write Cycle			1.8	10.4	ns	(3)
T32c	I[3:0]Dx float delay			2.1	9.7	ns	(3)
	Memory Control Signals						
T33	SYSDEN#	6.3	0.6	1.4	6.8	ns	(3)
T34	DC-to-DP: SYSCMD[4:0]#	2.3	0.8	1.1	7.4	ns	(3)
T35	DP-to-DC: MEMERR[1:0]#, SYSERR#	3.2	0.8	1.2	5.3	ns	(3)
T36	OMCNUM	3.3	0.3			ns	
	Reset Signals						
T39	MIRST#	2.9	0.6	1.3	5.5	ns	(3)
T40	CRESET#			2.0	10.0	ns	(3)
T41	SMIACT#			2.0	12.0	ns	(3)
T42	PWRGD Inactive pulse	10				Clks	Bus Clocks (6)
	Testability Signals						
T45	TRST#					ns	Async.
T46	TMS	4.6	0.8			ns	(7)
T47	TDI	6.0	2.0			ns	(7)
T48a	TDO			1.2	16.0	ns	(3) (7)
T48b	TDO on/off delay			1.4	16.0	ns	(3) (7)
	Non-GTL+ Host Bus Signals						
T49	INIT#, FLUSH#			2.0	12.0	ns	(3)

1. Min timings are measured with 0pF load, Max timings are measured with 50pF load.
2. Min and Max timings are measured with a 50pF load.
3. Min and Max timings are measured with 0pF load.
4. T22 and T23 apply to the Intel 450KX PCIset, signals CASB[3:0]#, CASA[3:0]#, RASB[3:0]#, RASA[3:0]#.
5. See Figure 7.
6. The power supply must wait until all voltages are stable for at least 1ms, and then assert the PWRGD signal.
7. Inputs are referenced to TCK rising, outputs are referenced to TCK falling.

1.5.1 WAVEFORMS

GTL+ Signal Waveforms: Please see the *Pentium Pro Family Developer's Manual: Specifications* for more information on GTL+ Clock, Setup, Hold, and Valid Delay waveforms.

PCI Bus Signal Waveforms: All PCI Bus signals are referenced to the PCLKIN Rising edge. For more information on the PCI Bus signals and waveforms, please refer to the PCI Specification 2.0

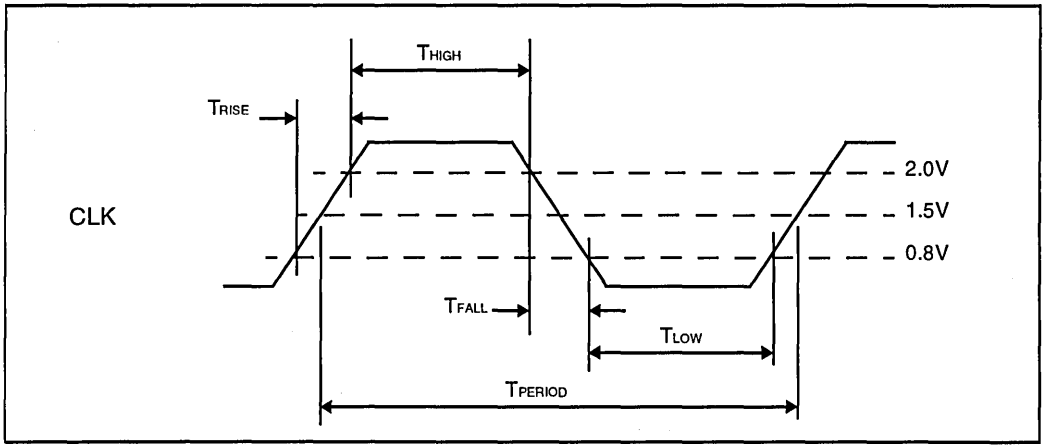

Figure 2. CLock Waveform

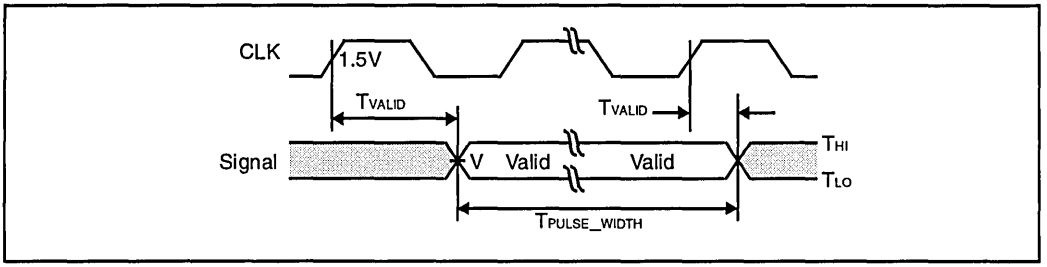

Figure 3. Valid Delay Timings

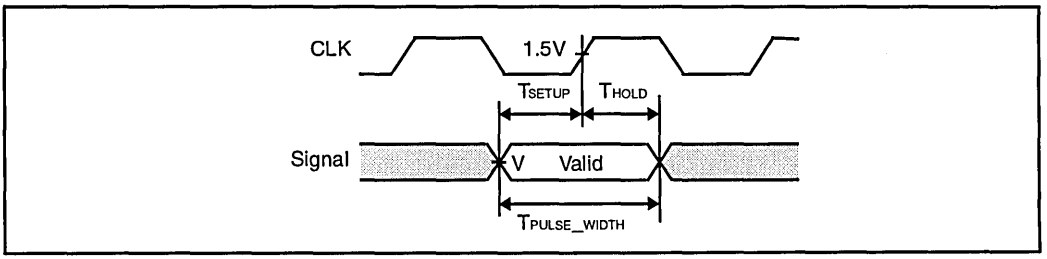

Figure 4. Setup & Hold Timings

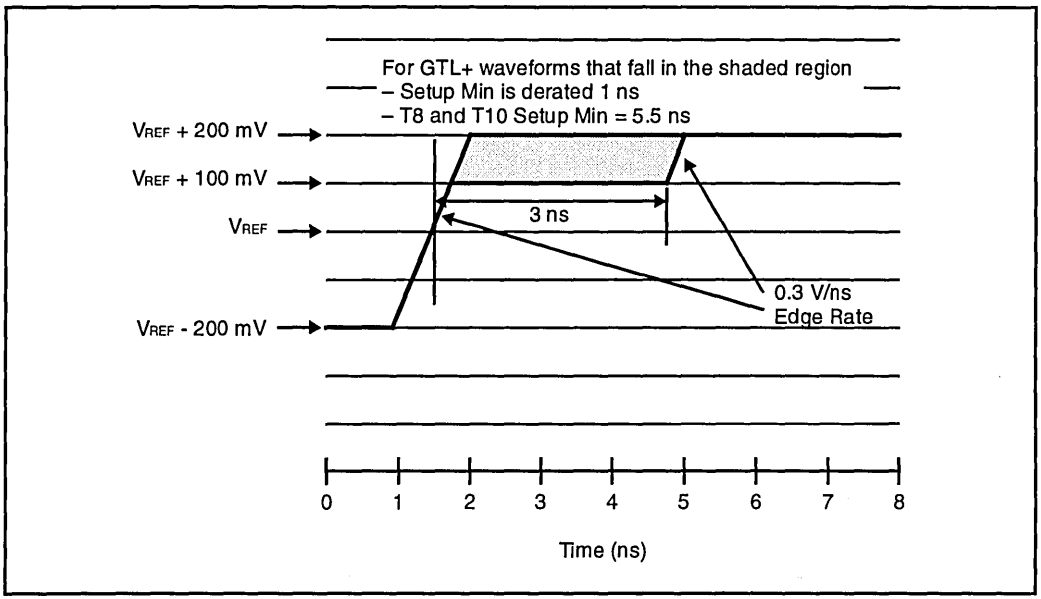

Figure 5. Exception to GTL+ Minimal Input Waveform With Setup Impact

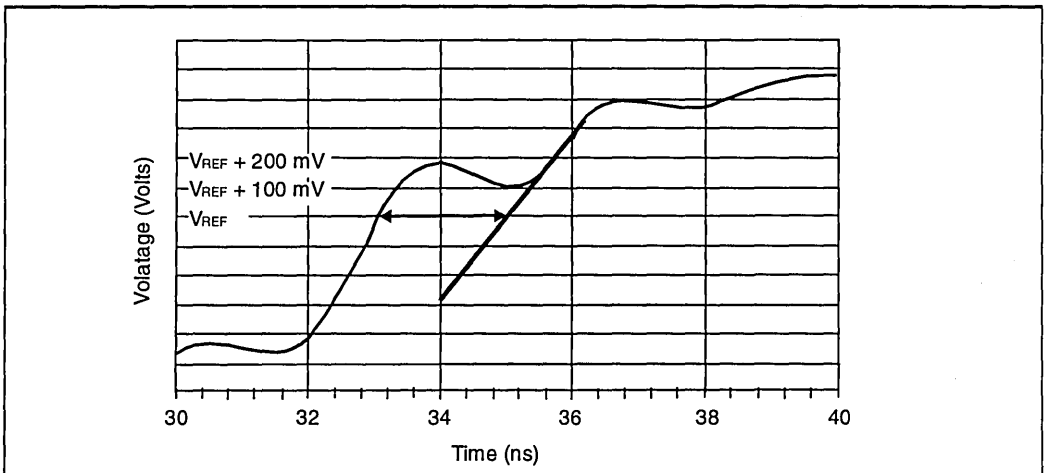

Note: 2ns Additional Flight Time as per the GTL+ I/O specification. However, another option is to derate setup by 1ns as this meets the requirements for the exception.

Figure 6. Example of 450KX/GX Receiver Waveform that Benefits from T_{SU} Derating

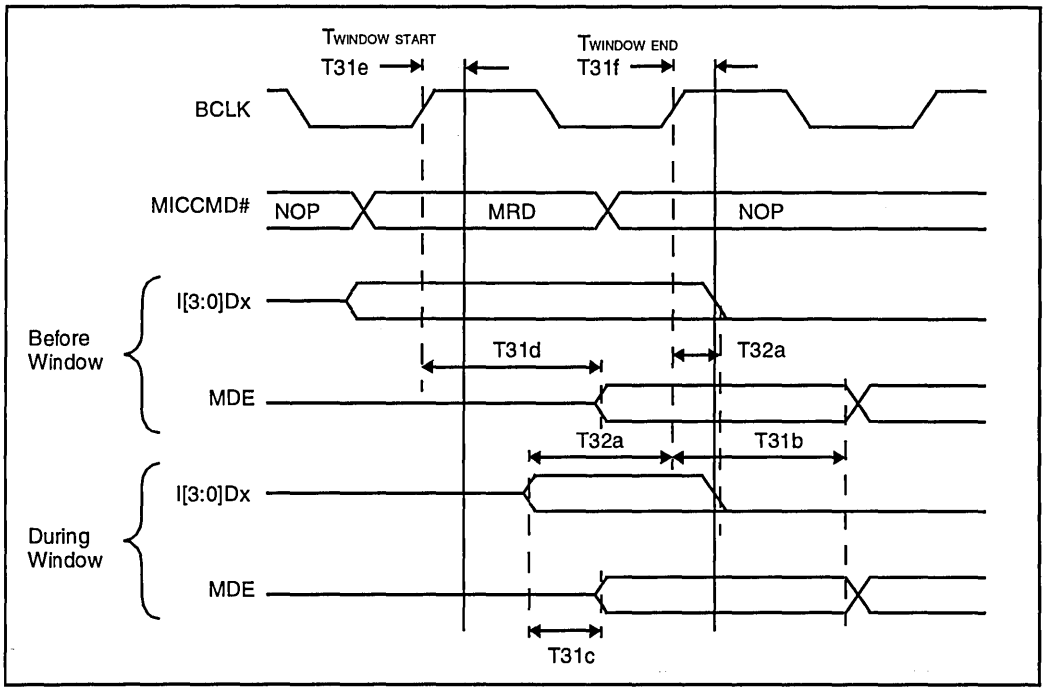

Figure 7. System Memory Signals [Memory Read (MIC)]

The Read Data Window is shown in Figure 7. At the first clock of a memory read cycle, data is valid to the MICs either before or during the MICs Transparent Latch Window. If data is valid before the window, T31d timings apply. If data is valid during the window, T31c timings apply. T31b is T_{CO} of the 2nd, 3rd, and 4th quad words.

2.0 SIGNAL QUALITY

The GTL+ I/O buffer specification defines new terms for use by buffer and system designers. These 450KX/GX PCIset buffers are described in these terms within this section. See the *Pentium Pro Family Developer's Manual: Specifications* for the complete explanation of these terms.

2.1 I/O Signal Simulations—Ensuring I/O Timings

It is highly recommended that system designers run extensive simulations on their Pentium Pro Processor/450GX- or 450KX-based designs. In addition, these simulations should include the memory subsystem designl. Please refer to the *GTL+ Guidelines Application Note*, AP-524, and the *Pentium Pro Family Developer's Manual: Specifications* for more information. For simulations, your field representative can provide the *Intel 450KX/GX PCIset I/O Buffer Model Specification.*

2.2 Signal Quality Specifications

Signals driven by any component on the Pentium Pro Processor bus must meet signal quality specifications to guarantee that the components read data properly, and to ensure that incoming signals do not affect the long term reliability of the components. There are three signal quality parameters defined: Overshoot/Undershoot, Ringback, and Settling Limit. These parameters are discussed in the next sections.

2.3 PCIset Ringback Specification

This section discusses the ringback specification for the parameters in the GTL+ signal groups on the Intel 450KX/GX PCIsets.

Case A requires less time than Case B from the V_{REF} crossing until the ringback into the "overdrive" region. The longer time from V_{REF} crossing until the ringback into the "overdrive" region required in Case B allows the ringback to be closer to V_{REF} for a defined period.

NOTE:

Specified for an edge rate of 0.3–0.8 V/ns. See the Pentium Pro Family Developer's Manual for the definition of these terms. See the figures below for the generic waveforms. All values determined by design/characterization.

Table 11. 450KX/GX GTL+ Signal Groups Ringback Tolerance—Case A

Parameter	Min	Unit	Figure
α: Overshoot	100	mV	8, 9
τ: Minimum Time at High or Low	2.25	ns	8, 9
ρ: Amplitude of Ringback	-100	mV	8, 9
δ: Duration of Square-wave Ringback	N/A	ns	8, 9
ϕ: Final Settling Voltage	100	mV	8, 9

NOTES:

1. Specified for an edge rate of 0.3–0.8 V/ns.
   See the Pentium Pro Family Developer's Manual for the definition of these terms.
   See the figures below for the generic waveforms.
   All values determined by design/characterization.

Table 12. 450KX/GX GTL+ Signal Groups Ringback Tolerance—Case B

Parameter	Min	Unit	Figure
α: Overshoot	100	mV	8, 9
τ_1:Minimum Time at High	2.7	ns	8, 9
τ_2:Minimum Time at Low	3.7	ns	8, 9
ρ: Amplitude of Ringback	-0	mV	8, 9
δ: Duration of Square-wave Ringback	2	ns	8, 9
ϕ: Final Settling Voltage	100	mV	8, 9

NOTES:

1. Specified for an edge rate of 0.3–0.8 V/ns.
   See the Pentium Pro Family Developer's Manual for the definition of these terms.
   See the figures below for the generic waveforms.
   All values determined by design/characterization.

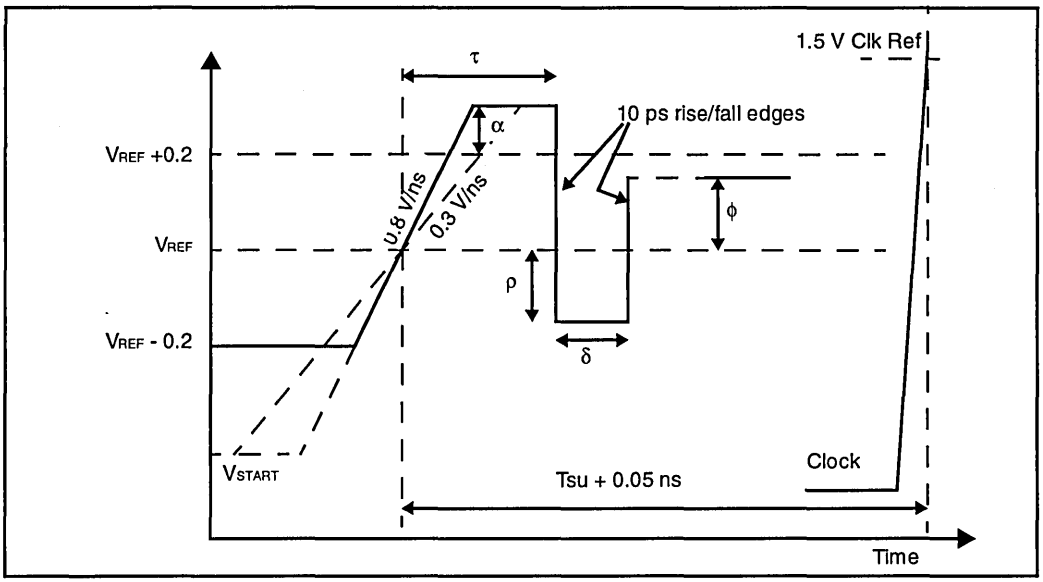

Figure 8. Standard Input Lo-to-Hi Waveform for Characterizing Receiver Ringback Tolerance

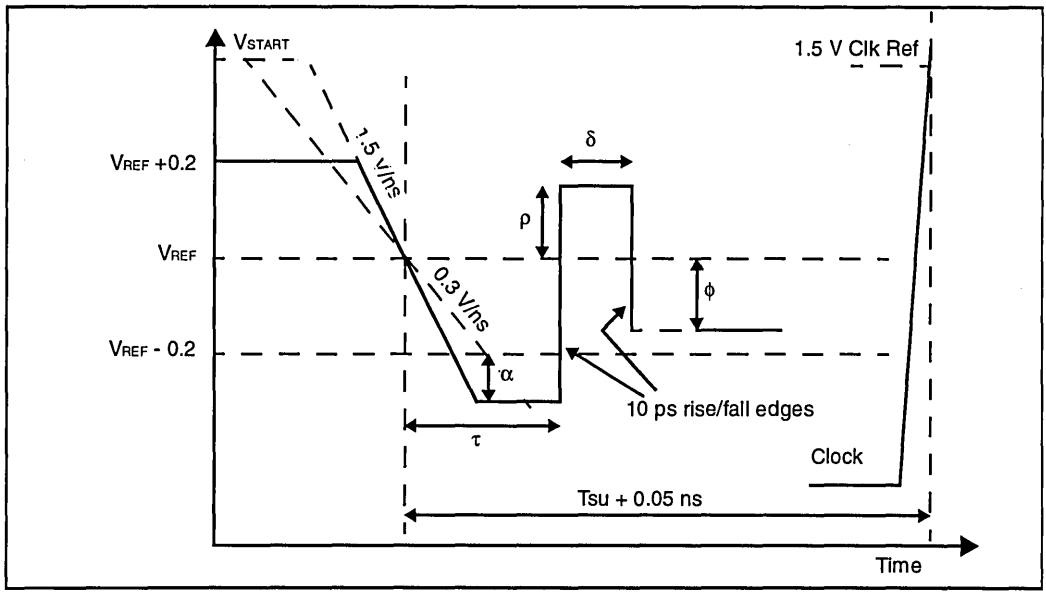

Figure 9. Standard Input Hi-to-Lo Waveform for Characterizing Receiver Ringback Tolerance

2.4 450KX/GX Undershoot Specification

The Pentium Pro processor bus signals AERR#, BERR#, BINIT#, BNR#, HIT#, and HITM# (only) are capable of sinking an 85mA current pulse at a 2.4% average time duty cycle. This is equivalent to -1.7V applied to a 20Ω source in series with the device pin for 8 ns at 66 MHz with a utilization of 5%.

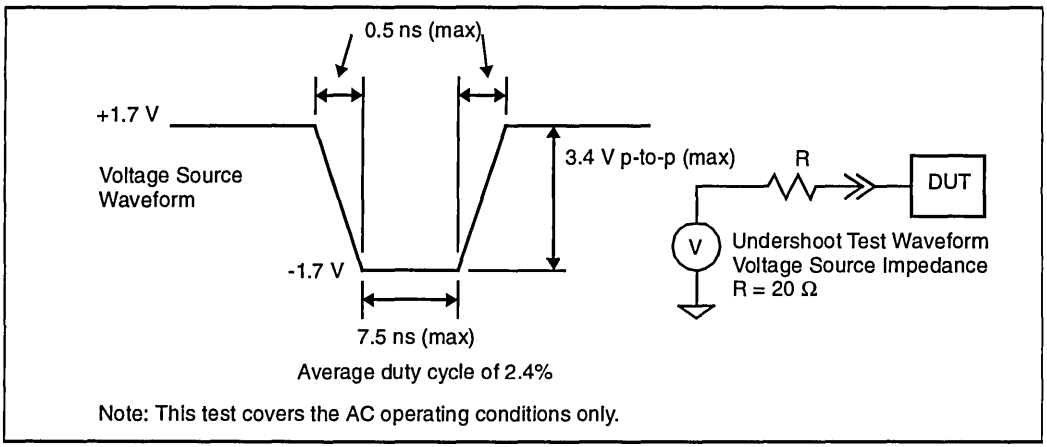

Figure 10. Undershoot

3.0 THERMAL SPECIFICATIONS

These tables show the power dissipation for the Intel 450KX/GX PCIsets in QFP and BGA packages.

Table 13. QFP Package Thermal Specifications

Device	Max Power in Watts	θja °C/W	θjc °C/W
PB — 304 QFP	2.40	18	5
DP — 240 QFP	2.30	25	6
DC — 208 QFP	1.13	35	7
MIC — 144 QFP	0.69	48	6

1. θja and θjc at 0 LFM air flow.
2. θja would be reduced by 12% at 100 LFM air flow.

Table 14. BGA Package Thermal Specifications

Device	Max Power in Watts	θja °C/W	θjc °C/W
PB — 352 BGA	2.40	18	6
DP — 256 BGA	2.30	20	5

3. θja and θjc at 0 LFM air flow.
4. θja would be reduced by 10% at 100 LFM air flow.
5. Both components are 4 layers.
6. Thermal data extracted from a 4"x4" four layer board with two planes.

intel

Index

ALPHABETIZED INDEX

Numerics
1M Power-on Reset Vector 52
2-Way DRAM configuration 129
4-Way DRAM Configuration 128

A
Absolute Maximum Ratings 153
AC Specifications 157
  CMOS 158
  GTL+ 157
Address Bit Permuting 109
AERR# 19, 52, 54, 62, 121
Aliasing 12, 38, 59
Anomaly 73
APIC Cluster ID 51, 52
APICR—I/O APIC Range Register 50, 115

B
BASEADD—MC Base Address Register 107
BCLK 22, 70, 157
BDNUM—Bridge Device Number Register 39
BERR# 19, 52, 54, 121
BINIT# 19, 54, 62, 73, 121
BIOS 9, 69
BIOS Range 53
BIST 29, 37, 72
  BIST register 37
Block Diagram 4
  PCI bridge 18
Blocking Logic 75
BPRI# 19, 66
Bridge Arbitration 41
Bridge to Bridge Sideband Signals 22
Bridge-to-Bridge Communication 67
Burst Write Assembly 42
Bus Number 28, 30, 39, 103
Bus Protocol Violation 55

C
Cache Line Size 36
Capacitance 155, 156
CDNUM—Controller Device Number Register 108
CLASSC—Class Code Register 36, 107
Clock 70, 131, 157
Clock Distribution 70, 131
CLSIZE—Cache Line Size Register 36
CMD—Command Register 108
Command Translation 76
Common CAS# 109, 127
Compatibility 8
CONFADD Register 27, 28, 101, 102
CONFDATA Register 27, 30, 101, 102
Configuration 26, 70, 72, 101
  Memory 122
  PCI Bridge Registers 32
  System 68
CONFVR—Configuration Values Driven on Reset Register 51
Controller Number 101
CRWC—CPU Read/Write Control Register 41
CSCONFV—Captured System Configuration Values Register 52
Current 154

D
Data Integrity 62
DC Specifications 154
DCC—Deturbo Counter Register 41
Deturbo 29, 41
Device Number 28, 31, 39, 103, 108
Device Numbers 102
DID—Device Identification Register 34, 105
Dimensions 87, 145
DRAM Interface 125
DRAM Row Limit 111, 124
DRL—DRAM Row Limit 111
Dual Address 78
Dual PB 65

E
ECC 62, 116
Efficiency 6
EISA 5, 10
Electrical Specifications 151
ERRCMD—Error Reporting Command Register 45
Error 54, 113, 120
Error Handling 62
ERRSTS—Error Reporting Status Register 45
EXERRCMD—PB Extended Error Reporting Command Register 53
EXERRSTS—PB Extended Error Reporting Status 55
Extended Memory 12

F
Fast Back-to-Back 35
FLSHBF# 21
FLUSH# 20
Frame Buffer 9, 47, 58
FRC Mode 52
Function Number 29, 103

G
Graphics Adapter Memory Region 8
GTL+ 155
GTLREFV 22

H
Hard Failure 55
Hard Reset 26, 29, 72
HBIOSR—High BIOS Gap Range Register 119
HBIOSR—High BIOS Range Register 53
HEADT—Header Type Register 37
HMGEA—High Memory Gap End Address Register 48, 115
HMGSA—High Memory Gap Range Start Address Register 48
HMGSA—High Memory Gap Start Address Register 115
Host Bus Interface Signals 19
Host Bus Time-out 54
Host-to-PCI Write Posting 41

I
I/O APIC 51, 58, 115, 124
  Configuration Space 11
  Multiple 65
I/O Map 57, 123
I/O Space 12, 27, 49, 50, 102
Inbound Posting 42
Inbound Queue 6
INIT on Shutdown 40
Initialization 26, 72, 101
In-Order Queue 6, 62
In-Order Queue Depth 52, 108
Interleave 5, 125
IOGNT# 22, 65
IOREQ# 22, 65
IOSR1—I/O Space Range 1 Register 49
IOSR2—I/O Space Range 2 Register 50
ISA 5, 9
ISA Alias Mode 12, 38, 59

L
Latency-Sensitive Peripherals 68
Leakage 155, 156
Line Read 42
Line Read Pre-Fetch 42
Livelock 69
LMG—Low Memory Gap Register 114
Lock Atomic Reads 40
Locks 20, 61

M
Master Abort 29, 35, 54, 63
Master Latency Timer 37, 70
Maximum Ratings 153
MEMACK# 21
Memory Gap 5, 57, 113, 123
Memory Interface Component 127
Memory Map 7, 57, 123
Memory Mapped I/O 13, 58
Memory Read Line 77
Memory Size Algorithm 122
Memory Sizes 126
Memory Space Gap 46
Memory Timing Register 116
Memory Write and Invalidate 34, 78
MEMTIM—Memory Timing Register 116
MERRCMD—Memory Error Reporting Command 120

MERRSTS—Memory Error Status Register 120
MG—Memory Gap Register 113
MGR—Memory Gap Range Register 46
MGUA—Memory Gap Upper Address Register 46, 114

N
NMI 54
Non-Interleaved DRAM configuration 130

O
OMCNUM 101, 108
Outbound Posting 41
Outbound Request Queue 6
Overshoot 163, 164

P
P2P Bridges 69
Package 79, 132
Package Dimensions 145
Page Open Policy 109
PAM—Programmable Attribute Map Register 44, 110
Parity Error 34, 35, 45, 55, 62, 122
PBC—PB Configuration Register 40
PBNUM—PCI Bus Number Register 39
PBRTMR—PB Retry Timers 56
PCI Bridge 17
PCI Bus Reset 72
PCI Configuration Space 30
PCICMD—PCI Command Register 34, 106
PCIRSR—PCI Reset Register 49
PCISTS—PCI Status Register 35, 106
PCI-to-Host Bus Write Posting 42
PCI-to-PCI Bridge 69
PCI-to-PCI Bridges 69
PCLK 157
PCLKIN 70
PDM—PCI Decode Mode 38
Performance 68, 69
Peripherals 68
PFB—PCI Frame Buffer Register 47
Pinout 79, 132
PLTMR—PCI Latency Timer 37
Posting 56
Power Specifications 154
Prefetch Enable 69
Protocol Violation 122
PRWC—PCI Read/Write Control 42
PSBNUM—Subordinate Bus Number Register 40
PWRGD 70, 152

R
Read Burst Delay 109
Read Multiple 77
Read/Write Clear 26
Read-Around-Write 109
Reclaim Enable 113
Reference Voltage 155
Refresh 117, 126
Register Number 29, 103
Registers 26, 101
Request Parity 62
Reserved Bits 26, 101
Reset 23, 25, 26, 29, 49, 51, 70, 71, 117, 131
Resource Locks 61
Retry Counter 56
Retry Timers 56
Revision Identification 36
RID—Revision Identification Register 36, 106
Ringback 163
Row 125

S
SERR# 34, 45, 54, 62, 64
SERRCMD—System Error Reporting Command Register 121
SERRSTS—System Error Status Register 122
Settling Voltage 163, 164
Shutdown 45
Signal Groups 151
Signal Quality 163
Signals 19
Simulations 162
SMI# 60
SMIACT# 20
SMM 8, 12, 43, 53, 58, 110, 119, 123
SMME—SMRAM Enable Register 110
SMMR—SMRAM Range Register 53, 119

SMRAM—SMM Enable Register 43
Snoop phase 60
Soft Reset 29, 72
Special Cycles 61

T
Target Abort 35, 45, 61, 63
Target Disconnect 63
Target Retry 63
Targeted PCI Resets 74
TCK 157
Temperature Specification 154
Termination Voltage 155
Test Signals 151
Thermal Specifications 166
Time-out 54, 62
TRC—Turbo and Reset Control 29
TSM—Top of System Memory Register 38

U
UERRADD—Uncorrectable Error Address Register 116
Undershoot 165
Unused Pins 151

V
VBAE—Video Buffer Area Enable Register 43
VBRE—Video Buffer Region Enable Register 110
Video Buffer 57, 123
Video Buffer Area 43, 110
VID—Vendor Identification Register 34, 105
Voltage 154

W
Watchdog Timer 40, 62
Write Back 60
Write Combining 60

UNITED STATES
Intel Corporation
2200 Mission College Boulevard
P.O. Box 58119
Santa Clara, CA 95052-8119

JAPAN
Intel Japan K.K.
5-6 Tokodai, Tsukuba-shi
Ibaraki-ken, 300-26

FRANCE
Intel Corporation S.A.R.L.
1, Quai De Grenelle, BP543
75725 Paris Cedex

UNITED KINGDOM
Intel Corporation (U.K.) Ltd.
Pipers Way
Swindon SN3 1RJ
Wiltshire, England

GERMANY
Intel GmbH
Dornacher Strasse 1
D-85622 Felkirchen
Muenchen

HONG KONG
Intel Semiconductor Ltd.
32/F Two Pacific Place
88 Queensway, Central

CANADA
Intel Semiconductor of Canada, Ltd.
190 Attwell Drive, Suite 500
Rexdale, Ontario M9W 6H8

To learn more about Intel Corporation,
visit our site on the World Wide Web at
http://www.intel.com/

Order Number: 290523-001
Printed in USA/0596/5K/MRR MS
Peripheral Components

Intel 450KX/GX PCIset Specification Update 82454KX/GX (PB) 82453KX/GX (MC) 82452KX/GX (DP) 82451KX/GX (MIC)

Release Date: January, 1997

Order Number: 243109-009

The Intel 450KX/GX PCIset may contain design defects or errors known as errata. Characterized errata that may cause the Intel 450KX/GX PCIset's behavior to deviate from published specifications are documented in this Specification Update.

Copies of documents which have an ordering number and are referenced in this document, or other Intel literature, may be obtained from:

Intel Corporation
P.O. Box 7641
Mt. Prospect, IL 60056-7641

or call 1-800-879-4683

CONTENTS

REVISION HISTORY v
PREFACE vi
Part I: Specification Update for the Intel 450KX PCIset
GENERAL INFORMATION 3
SPECIFICATION CHANGES 7
ERRATA 10
SPECIFICATION CLARIFICATIONS 25
DOCUMENTATION CHANGES 26
Part II: Specification Update for the Intel 450GX PCIset
GENERAL INFORMATION 31
SPECIFICATION CHANGES 33
ERRATA 37
SPECIFICATION CLARIFICATIONS 50
DOCUMENTATION CHANGES 52

REVISION HISTORY

Date of Revision	Version	Description
May, 1996	-001	This document is the first Specification Update for the Intel 450KX/GX PCIset.
June, 1996	-002	Updated 450KX PCIset Errata 6, 20, and 21, and 450GX PCIset Erratum 4. Added 450KX PCIset Errata 30-33 and 450GX PCIset Errata 12-15. Added 450KX and 450GX PCIset Specification Change 1.
July, 1996	-003	Added 450KX and 450GX Specification Change 2. Added 450KX Errata 34-36 and 45GX Errata 16-19.
August, 1996	-004	Updated 450KX and 450GX Specification Change 2, 450KX Erratum 35, 450GX Errata 13, 17, and 19, and 450KX and 450GX Documentation Change 1. Added 450KX and 450GX Specification Change 3, 450KX Errata 37 and 38, and 450GX Errata 20-22. Updated plans for 450KX Errata 4, 20, 21, 23-25, 27, and 30-34.
September, 1996	-005	Updated 450KX and 450GX Documentation Change 1.
October, 1996	-006	Added Documentation Change 2.
November, 1996	-007	Updated 450KX and 450GX Specification Change 2, and 450GX Erratum 20.
December, 1996	-008	Added 450GX Specification Changes 4-7 and 450GX Errata 23 and 24.
January, 1997	-009	Added 450KX Errata 39 and Specification Clarification 2. Added 450GX Errata 25 and Specification Clarification 2.

PREFACE

This document is an update to the specifications contained in the *Intel 450KX/GX PCIset* datasheet, (Order Number 290523). It is intended for hardware system manufacturers and software developers of applications, operating systems, or tools. It contains Specification Changes, S-Specs, Errata, Specification Clarifications, and Documentation Changes.

Nomenclature

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

S-Specs are exceptions to the published specifications, and apply only to the units assembled under that s-spec.

Errata are design defects or errors. Errata may cause the Intel 450KX/GX PCIset's behavior to deviate from published specifications. Hardware and software designed to be used with any given stepping must assume that all errata documented for that stepping are present on all devices.

Specification Clarifications describe a specification in greater detail or further highlight a specification's impact to a complex design situation. These clarifications will be incorporated in the next release of the specifications.

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

Identification Information

The 82453 Memory Controller can be identified by the following values:

Vendor ID[1]	Device ID[1]
8086h	84C5h

The 82454 PCI Bridge can be identified by the following values:

Vendor ID[1]	Device ID[2]
8086h	84C4h

NOTES:

1. The Vendor ID corresponds to the value in the Vendor ID register of the device's PCI configuration space, at address offset 00-01h (the A2 stepping of the 82453KX MC located the Vendor ID register at address offset 04-05h).
2. The Device ID corresponds to the value in the Device ID register of the device's PCI configuration space, at address offset 02-03h (the A2 stepping of the 82453KX MC located the Device ID register at address offset 06-07h).

Part I: Specification Update for the Intel 450KX PCIset

GENERAL INFORMATION

Top Markings

A- and B-step Production Units, BGA:

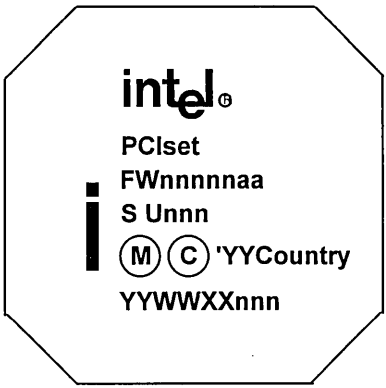

A- and B-step Production Units, QFP:

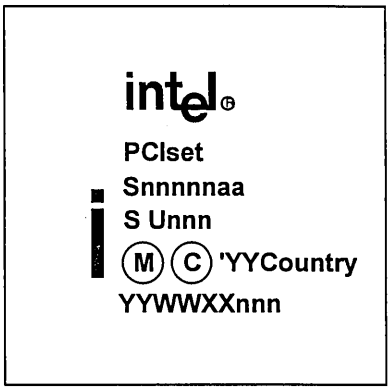

NOTES:

- nnnnnaa = Product Number.
- Q nnn = Sample Specification Number
- S Unnn = S-spec Number.
- 'YY Country = Copyright Dates and Country of Origin.
- YYWWXXnnn = Alternative Identification Number.

Basic Intel 450KX PCIset Identification Information

Product Number	Vendor ID	Device ID	Revision ID	Product Stepping	Kit Steppings	S-Spec	V_{CC}	T_{CASE}	Notes
S82451KX	n/a[1]	n/a[1]	n/a[1]	A1	A2, B0	S U025	3.3V ± 5%	0°C - 85°C	
S82451KX	n/a[1]	n/a[1]	n/a[1]	A1	A2, B0	S U039	3.3V ± 5%	0°C - 85°C	
S82452KX	n/a[1]	n/a[1]	n/a[1]	A1	A2	S U026	3.3V ± 5%	0°C - 85°C	
FW82452KX	n/a[1]	n/a[1]	n/a[1]	A1	A2	S U029	3.3V ± 5%	0°C - 85°C	2
S82452KX	n/a[1]	n/a[1]	n/a[1]	A1	A2	S U040	3.3V ± 5%	0°C - 85°C	
FW82452KX	n/a[1]	n/a[1]	n/a[1]	A1	A2	S U043	3.3V ± 5%	0°C - 85°C	2
S82452KX	n/a[1]	n/a[1]	n/a[1]	A3	B0	S U061	3.3V ± 5%	0°C - 85°C	
S82453KX	8086h[3]	84C5h[3]	2	A2	A2	S U027	3.3V ± 5%	0°C - 85°C	
S82453KX	8086h[3]	84C5h[3]	2	A2	A2	S U041	3.3V ± 5%	0°C - 85°C	
S82453KX	8086h	84C5h	4	A4	B0	S U062	3.3V ± 5%	0°C - 85°C	
S82454KX	8086h	84C4h	2	A2	A2	S U022	3.3V ± 5%	0°C - 85°C	
FW82454KX	8086h	84C4h	2	A2	A2	S U024	3.3V ± 5%	0°C - 85°C	2
S82454KX	8086h	84C4h	2	A2	A2	S U028	3.3V ± 5%	0°C - 85°C	
FW82454KX	8086h	84C4h	2	A2	A2	S U030	3.3V ± 5%	0°C - 85°C	2
S82454KX	8086h	84C4h	2	A2	A2	S U042	3.3V ± 5%	0°C - 85°C	
FW82454KX	8086h	84C4h	2	A2	A2	S U044	3.3V ± 5%	0°C - 85°C	2
S82454KX	8086h	84C4h	4	A4	B0	S U064	3.3V ± 5%	0°C - 85°C	

NOTES:

1. These components are not visible from the PCI bus, and so do not have Vendor, Device, or Revision ID registers at the PCI specification-defined locations.
2. These components have BGA (Ball Grid Array) packaging.
3. The A2 82453KX MC's Vendor and Device ID registers are not at the PCI specification-defined locations; the Vendor ID register is located at register offset 04-05h, and the Device ID is located at 06-07h.

Summary Table of Changes

The following table indicates the Specification Changes, Errata, Specification Clarifications, or Documentation Changes which apply to the Intel 450KX PCIset. Intel intends to fix some of the errata in future steppings of the component(s), and to account for the other outstanding issues through documentation or specification changes as noted. This table uses the following notations:

CODES USED IN SUMMARY TABLE

X:	Specification Change, Erratum, Specification Clarification, or Documentation Change applies to the given stepping.
Doc:	Intel intends to update the appropriate documentation in a future revision.
Fix:	Intel is investigating the possibility of fixing this erratum in a future stepping of the component(s).
Fixed:	This erratum has been previously fixed.
NoFix:	Intel is currently not investigating a fix for this erratum.
(No mark) or (blank box):	This item is fixed in or does not apply to the given kit stepping.
Shaded:	This erratum is either new or modified from the previous version of the document.

NO.	A2	B0	Plans	SPECIFICATION CHANGES
1	X	X	Doc	PLLRST pin added
2	X	X	Doc	Valid memory timing parameters
3	X	X	Doc	CMOS overshoot/undershoot specification

NO.	A2	B0	Plans	ERRATA
1	X		Fixed	Stop Clock Acknowledge cycles may confuse buffers
2	X		Fixed	0-byte length write may cause subsequent write failure
3	X		Fixed	Inbound read may be accepted despite posted outbound write
4	X	X	NoFix	SMRAM addresses may not be decoded correctly
5	X		Fixed	Mixed read lines and partials may corrupt data
6	X		Fixed	Inbound write posting may cause write failure
7	X		Fixed	RMW with line write may cause data corruption
8	X		Fixed	Inbound write may collide with some special cycles
9	X		Fixed	Extended read-around-write may corrupt write data
10	X		Fixed	Outbound posted write after inbound read prefetch may hang
11	X		Fixed	Inbound posted write with 1:1 interleaving may corrupt data
12	X		Fixed	RAW may cause data corruption
13	X		Fixed	Data transfer order 3 may cause data to be issued out of order
14	X		Fixed	PCI address parity error may cause dropped transaction
15	X		Fixed	System hang with inbound write posting enabled
16	X		Fixed	RAW may cause data corruption during refresh

NO.	A2	B0	Plans	ERRATA
17	X		Fixed	GAT devices may time out during inbound read prefetch
18	X		Fixed	Data may be corrupted if RCD = 4 and LWC = 3
19	X		Fixed	T_{CO_MAX} specification not met for GTL+ signals
20	X	X	NoFix	RAW may hang 1:1 or 2:1 interleaved MP systems
21	X	X	NoFix	Mixed interleave increments may cause data corruption
22	X	X	NoFix	Parity error may occur for ADS# during BINIT#
23	X	X	NoFix	Hang with PCI-to-PCI bridges in MP systems
24	X	X	NoFix	PCI_RST# not asserted asynchronously
25	X	X	NoFix	BERR# to BINIT# conversion may prevent recovery from BINIT#
26	X		Fixed	Error reporting registers may not record error information correctly
27	X	X	NoFix	Combination of ECC errors may cause one error to be undetected
28	X		Fixed	16-Byte write may hang system
29	X	X	NoFix	16-Byte read with two ECC errors may not be reported correctly
30	X	X	NoFix	Inbound locked PCI transactions may hang system
31	X	X	NoFix	Retry on inbound read may corrupt outbound data
32	X	X	NoFix	ADS# in last clock of BINIT# prevents recovery
33	X	X	NoFix	Some signals indeterminate after RESET# deassertion
34	X	X	NoFix	Delayed read from PCI-to-PCI bridge may corrupt data
35	X	X	NoFix	Page Open Policy of "hold page open" may corrupt write data
36	X	X	NoFix	T_{CASE} drop plus voltage swing may cause DPLL failure
37	X	X	NoFix	Memory gap reclaiming may corrupt data
38	X	X	NoFix	RAW may corrupt write data in 1:1 interleaving
39	X	X	NoFix	BINIT# Assertion May Cause Active RAS# Negation

NO.	A2	B0	Plans	SPECIFICATION CLARIFICATIONS
1	X	X	Doc	Explicit writebacks claimed by 82454KX PB
2	X	X	Doc	Supported Configurations For MC Row Limit Register Programming

NO.	A2	B0	Plans	DOCUMENTATION CHANGES
1	X	X	Doc	Register offset and default value correction
2	X	X	Doc	CMOS definition should be 3.3V or 5V

SPECIFICATION CHANGES

1. PLLRST Pin Added

A PLLRST pin will be added to the definition of each device in the Intel 450KX PCIset. The pin numbers will be assigned as follows:

Device	Pin Number
82454KX PB QFP	301
82454KX PB BGA	A5
82453KX DC QFP	81
82452KX DP QFP	202
82452KX DP BGA	B12
82451KX MIC QFP	42

This signal will be added to the block diagram in Figure 1 of both Chapter 2 and Chapter 3.

Each of the specified PLLRST pins are 5V tolerant signals.

The signal will be added to Chapter 2, Section 1.4 and to Chapter 3, sections 1.1, 1.2 and 1.3 as signal "PLLRST", type "I, CMOS", and described as "This pin must be driven high for at least 2 clocks to reset the internal DPLL (Digital Phase Lock Loop). The DPLL should be reset after (or until) the clock input pins are stable at their final operating frequency. This pin does not have an edge rate requirement."

The following sentences will be added to Chapter 2, section 3.7.2 and Chapter 3, section 3.4: "The PLLRST pin must be driven high for at least 2 clocks to reset the internal DPLL. The DPLL should be reset after (or until) the clock input pins are stable at their final operating frequency."

Chapter 4 will be updated to include this pin information throughout.

2. Valid Memory Timing Parameters

The following is a list of timing values which have been validated by Intel. The list is the result of applying the rules set forth in the *Intel 450KX/GX PCIset* datasheet plus a set of filters to eliminate settings that Intel believes would not or could not be used in practice. Note that OEMs must still ensure that the timing parameters used meet the timing constraints for their system design, applicable clock rates, and supported DRAM speeds. See below for a list of acronyms used in the table.

RCD = 3, RCAD = 2, and CSR = 1 for all setting options listed below.

						1:1			2:1			4:1		
LWC	RASPW	CAH	RCAS	WCAS	RP	CP	RBD	ACh	CP	RBD	ACh	CP	RBD	ACh
2	5	1	2	2	3	1	2	2814	2	1	2834	1	0	2814
2	5	1	3	2	3	1	3	2914	1	1	2914	1	0	2914
2	5	1	3	3	3	1	3	2954	1	1	2954	1	0	2954
2	6	1	2	2	3	1	2	3014	2	1	3034	1	0	3014

						1:1			2:1			4:1		
LWC	RASPW	CAH	RCAS	WCAS	RP	CP	RBD	ACh	CP	RBD	ACh	CP	RBD	ACh
2	6	1	3	2	3	1	3	3114	1	1	3114	1	0	3114
2	6	1	3	2	4	1	3	3115	1	1	3115	1	0	3115
2	6	1	3	3	3	1	3	3154	1	1	3154	1	0	3154
2	6	1	3	3	4	1	3	3155	1	1	3155	1	0	3155
2	6	2	3	2	3	1	3	3514	1	1	3514	1	0	3514
2	6	2	3	2	4	1	3	3515	1	1	3515	1	0	3515
2	6	2	3	3	3	1	3	3554	1	1	3554	1	0	3554
2	6	2	3	3	4	1	3	3555	1	1	3555	1	0	3555
2	6	2	4	3	3	1	4	3654	2	2	3674	1	0	3654
2	6	2	4	3	4	1	4	3655	2	2	3675	1	0	3655
3	5	1	2	2	3	2	3	4834	2	1	4834	1	0	4814
3	5	1	3	2	3	2	4	4934				1	0	4914
3	5	1	3	3	3	2	4	4974				1	0	4954
3	6	2	3	2	3	2	4	5534				1	0	5514
3	6	2	3	2	4	2	4	5535				1	0	5515
3	6	2	3	3	3	2	4	5574				1	0	5554
3	6	2	3	3	4	2	4	5575				1	0	5555
3	6	2	4	3	3	2	5	5674	2	2	5674	1	0	5654
3	6	2	4	3	4	2	5	5675	2	2	5675	1	0	5655
3	6	2	4	4	4	1	4	5695	2	2	56B5	1	0	5695

RCD = 3, RCAD = 2, and CSR = 2 for all setting options listed below.

						1:1			2:1			4:1		
LWC	RASPW	CAH	RCAS	WCAS	RP	CP	RBD	ACh	CP	RBD	ACh	CP	RBD	ACh
3	6	2	3	4	3	1	3	D594	1	1	D594	1	0	D594
3	6	2	4	3	4	2	5	D675	2	2	D675	1	0	D655
2	6	2	3	3	3	1	3	B554	1	1	B554	1	0	B554

NOTE:

All above values are in number of cycles, not actual bit settings, except ACh (which is the actual bit setting, in hexadecimal).

RCD: RAS# to CAS# Delay: bits 3:2 of the Memory Timing Register (AC-AFh).

RCAD: RAS# to Column Address Delay: bit 4 of the Memory Timing Register (AC-AFh).

CSR: CAS# Setup to RAS# for CAS# before RAS# refresh: bit 15 of the Memory Timing Register (AC-AFh).

LWC: Last Write to CAS#: bits 14:13 of the Memory Timing Register (AC-AFh).

RASPW: RAS# Pulse Width: bits 12:11 of the Memory Timing Register (AC-AFh).

CAH: Column Address Hold Time: bit 10 of the Memory Timing Register (AC-AFh).

RCAS: Read CAS# Pulse Width: bits 9:8 of the Memory Timing Register (AC-AFh).

WCAS: Write CAS# Pulse Width: bits 7:6 of the Memory Timing Register (AC-AFh).

RP: RAS# Precharge Time: bits 1:0 of the Memory Timing Register (AC-AFh).
CP: CAS# Precharge Time: bit 5 of the Memory Timing Register (AC-AFh).
RBD: Read Burst Delay: bits 2:0 of the Command Register (4C-4Fh).
ACh: Actual hexadecimal value programmed into bits 15:0 of the Memory Timing Register (AC-AFh).

3. CMOS Overshoot/Undershoot Specification

The following will be added to Chapter 4, Section 2.4, *Intel 450KX/GX Undershoot Specification*, in the *Intel 450KX/GX PCIset Data Sheet,* and the section will be retitled *Intel 450KX/GX Overshoot/Undershoot Specifications.*

The 3.3V tolerant CMOS signals of the processor bus allow for the following maximum AC waveforms:

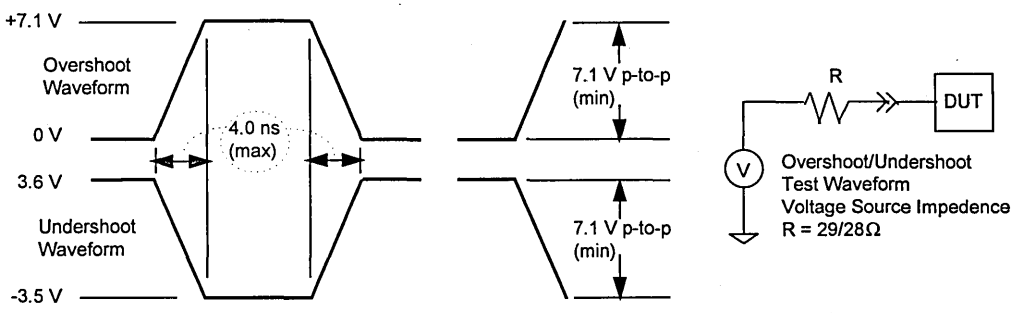

Average duty cycle of 17.6%

Note: This test covers the AC operating conditions only.

ERRATA

1. Stop Clock Acknowledge Cycles May Confuse Buffers

PROBLEM: If a Stop Clock Acknowledge special cycle is issued from a processor bus agent, the 82454 PB may pipeline an outbound (host-to-PCI bus) read behind it if the IOQ (In Order Queue) depth is set to 8. The read and write buffer pointers which track these transactions may become confused, and the data buffer which contains the data for the outbound read may be deallocated incorrectly. A similar situation arises if the special cycle is issued concurrent with an inbound read. Even if the IOQ depth is 1, this may occur, if inbound (PCI-to-host bus) write-posting is enabled in the 82454 PB.

IMPLICATION: A Stop Clock Acknowledge special cycle may cause premature deallocation of a buffer holding outbound read (or inbound write) data, resulting in a system hang with DBSY# asserted.

WORKAROUND: BIOS code can contain a workaround for this erratum.

2. 0-Byte Length Write May Cause Subsequent Write Failure

PROBLEM: If a 0-byte write occurs, a subsequent write to memory may fail, corrupting memory, or a resulting implicit writeback or line read may not complete, causing a system hang. This can happen if:

- Back-to-back 8-byte writes occur after the 0-byte write.
- The 0-byte write receives an implicit writeback response with RAW enabled in the 82453KX MC.
- The 0-byte write is followed by a line read and preceded by a 0-byte read.

IMPLICATION: The Pentium® Pro processor will not perform 0-byte writes. The only configurations identified by Intel where the failing sequence is issued are with 824731FB (PIIX) or 824731SB (PIIX3) controllers used as IDE bus masters behind the 82454 PB, with inbound write-posting enabled and an IOQ depth of 8. In this case, the 82454KX PB may issue 0-byte writes as well as pipelined back-to-back 8-byte writes (which may be issued by the 82454KX PB normally).

WORKAROUND: Ensure that no system bus agent issues 0-byte writes. If the 824731FB (PIIX) or 824731SB (PIIX3) controllers are used, they must not be used as IDE bus masters.

3. Inbound Read May Be Accepted Despite Posted Outbound Write

PROBLEM: The 82454 PB may accept an inbound read even if a posted write is queued. This is an ordering violation.

IMPLICATION: An outbound access which follows the outbound posted write may be snoop stalled indefinitely. If a new, non-posted, inbound request is issued, BPRI# will be asserted even though the outbound posted write has not been issued to the PCI bus, causing a deadlock, and the system will hang.

WORKAROUND: Do not enable outbound write posting in the 82454 PB.

4. SMRAM Addresses May Not Be Decoded Correctly

PROBLEM: While executing in SMM (System Management Mode), certain sequences of transactions may allow the CAS# signal to be asserted without a corresponding RAS# signal during a memory access to SMRAM. In the A2 stepping of the Intel 450KX PCIset, there are many such sequences. For the B0 stepping

of the Intel 450KX PCIset, the sequences must be pipelined SMRAM requests (with the IOQ depth set to 8) as follows:

- An SMRAM request which opens a page (i.e. the transaction address accesses a block of DRAM with a new row address).
- An SMRAM request which is a page hit.
- Another SMRAM request which is a page hit.

The address phase of the third request must occur during the first clock of CAS# assertion for the second request to encounter this erratum. A second sequence is:

- An SMRAM request which is a page hit.
- Another SMRAM request whose ADS# assertion comes 5 or 6 clocks later than the ADS# assertion for the previous transaction, and which is also a page hit.

IMPLICATION: If these sequences occur while in SMM, data corruption may result.

WORKAROUND: For the A2 stepping of the Intel 450KX PCIset, SMM cannot be used. For the B0 stepping of the Intel 450KX PCIset, this erratum can be avoided by setting the IOQ depth to 1, thus preventing transactions from being pipelined together.

5. *Mixed Read Lines and Partials May Corrupt Data*

PROBLEM: The 82453KX MC may not deallocate a data buffer correctly after the following sequence of transactions occurs:

1. A read line.
2. A read partial.
3. A read line which receives an implicit writeback response.

IMPLICATION: Stale data may be retrieved for a subsequent read line or read partial, resulting in corrupted data, or a protocol violation may be observed during a read line, causing a system shutdown.

WORKAROUND: If the following conditions are met, this erratum will not occur:

- In an MP system, all processors must have identical memory models, with all MTRRs for a particular region of memory mapped the same way.
- Mixed mode paging must not be used.
- All memory behind the 82453KX MC must be mapped to the same memory type.
- The IOQ depth must be set to 1.

6. *Inbound Write Posting May Cause Write Failure*

PROBLEM: During an inbound posted write burst, the 82454 PB may use the same data buffer for two adjacent transactions. This causes the first write to use and deallocate the data from the second write, corrupting data and incorrectly deallocating a data buffer. When the data phase for the second write arrives, DBSY# may be asserted for some time, with no data phase or DRDY# assertion. When the PCI master completes the write burst, the deallocated buffer is never rewritten, resulting in a system hang with DBSY# asserted.

IMPLICATION: Enabling inbound write posting may cause data corruption followed by a system hang. Disabling inbound write posting may result in significant I/O performance loss.

WORKAROUND: An external arbiter should be implemented which allows inbound posting to be enabled, and guarantees at least two idle cycles between inbound PCI requests.

This arbiter must not perform overlapped arbitration. One way to ensure this is to implement the circuit described below, external to the arbiter. The actual workaround involves waiting for a bus idle before allowing a new grant to be issued. The IOQ depth must still be set to 1, and 2:1 interleaving must be used, to work around other errata which may surface when inbound write posting is enabled (even with this workaround). Note that this latter restriction requires at least 4 SIMM DRAM modules to be loaded at all times.

This following circuit is an implementation of the described workaround:

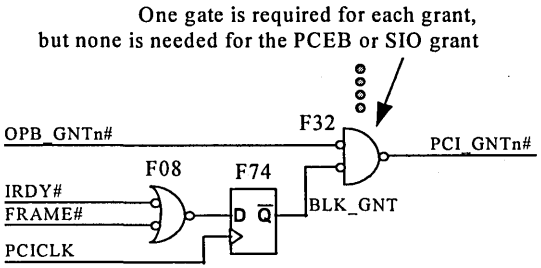

This workaround will allow inbound write posting to be enabled on A2 450KX PCIset silicon, and will not affect the functionality of B0 450KX/GX PCIset silicon, though a slight performance loss may result if this workaround is used with the B0 stepping. Enabling inbound posting by using this workaround is a significant performance enhancement for systems with A2 450KX PCIset silicon; refer to Intel's World Wide Web site for more details on the performance loss incurred by disabling inbound write posting (at URL http://www.intel.com/procs/support/ppro/82450.htm).

7. *RMW with Line Write May Cause Data Corruption*

PROBLEM: If the 82453KX MC is performing a Read-Modify-Write (RMW) transaction due to a partial write to memory, and this is followed by a line write before the RMW completes, the byte enable signals for the RMW will be held too long, and only part of the line write will actually be written to memory.

IMPLICATION: A write may fail, corrupting data, if partial writes and line writes are pipelined together.

WORKAROUND: If the following conditions are met, this erratum will not occur:

- All processors in an MP system must have identical memory models, with all MTRRs for a particular region of memory mapped the same way.
- Mixed mode paging must not be used.
- All memory behind the 82453KX MC must be mapped to the same memory type.
- The IOQ depth must be set to 1.

8. *Inbound Write May Collide with Some Special Cycles*

PROBLEM: A special cycle from a processor bus agent which occurs during a write from the PCI bus to the processor bus may corrupt the byte enable signals for the write, causing incorrect data to be written into memory. Special cycles include Shutdown, Flush, Halt, Sync, Flush Acknowledge, Stop Clock Acknowledge, and SMI Acknowledge. This can only occur with outbound write posting enabled.

IMPLICATION: Special cycles coincident with inbound traffic from the PCI bus may corrupt data. However, only Flush, Sync, and Flush Acknowledge special cycles will normally be issued from Pentium Pro processors using the A2 stepping of the Intel 450KX PCIset.

- Flush special cycles are issued upon execution of an INVD instruction. Note that this instruction intentionally removes modified lines from the cache without writing them back to memory; this results in the contents of memory being unusable. Therefore, the possibility of data corruption due to a Flush special cycle can be ignored.
- Flush Acknowledge special cycles are issued upon completion of flushing the caches after observation of a FLUSH# signal assertion. This signal is only asserted by the 82454KX PB when the system is in Deturbo mode. Therefore, these special cycles can be avoided by not enabling Deturbo mode in the affected systems.
- Sync special cycles are issued upon execution of a WBINVD instruction. This instruction is a privileged instruction which should only be executed by operating system level code. Intel has not identified any such code which uses this instruction during inbound traffic.

Since each of these special cycles are only issued when the contents of the cache are being invalidated, inbound traffic which hits a cached memory location is not valid. If all inbound traffic during these special cycles targets cached memory, no effect will be seen due to this erratum. However, inbound traffic targeting uncached main memory would be corrupted if Deturbo mode, INVD instructions, or WBINVD instructions were used simultaneously. Note that INVD and WBINVD instructions are very rare system events; execution of an INVD instruction removes modified lines from the cache without writing them back to memory, so incorrect data in memory will not cause further system problems.

WORKAROUND: Do not enable outbound write posting in the 82454KX PB.

9. *Extended Read-Around-Write May Corrupt Write Data*

PROBLEM: In a sequence of transactions to two different pages in memory (page X and page Y), the row address strobe (RAS#) and memory address (MA#) signals may be asserted on the same clock, resulting in a protocol violation which corrupts the data as it is written into memory. The sequence of transactions is as follows:

1. A line read to page X.
2. A line read to page X.
3. A line read to page X.
4. A line write to page Y.
5. A line read to page X.
6. A line write to page Y.
7. A line read to page X.
8. A line read to page X.

This sequence will appear at the DRAM in the following order, with Read-Around-Write (RAW) and Extended Read-Around-Write (ERAW) enabled:

1 - 2 - 3 - 5 - 4 - 7 - 8 - 6

In this sequence, transaction 5 is issued to memory before transaction 4 due to RAW/ERAW. Transaction 6 is then detected, and transaction 4 is issued since no reordering may occur around multiple writes. At this point, transaction 6 is at the front of the queue, and is a row-hit/page-hit relative to transaction 4, so the RAS# signal is left asserted in preparation for transaction 6. ERAW then causes transactions 7 and 8 to be reordered in front of transaction 6; these transactions are not row-hit/page-hit transactions, and the MA# lines are updated

on the same clock as RAS# is reasserted, causing the protocol violation. The data returned for transactions 7 and 8 will then be corrupted.

IMPLICATION: ERAW will corrupt data returned for this sequence of memory transactions.

WORKAROUND: Do not enable ERAW in the 82453KX MC.

10. Outbound Posted Write After Inbound Read Prefetch May Hang

PROBLEM: If an outbound write closely follows an inbound read prefetch, the data buffer containing the outbound write's data may be rewritten by the data for the prefetch. This occurs when the outbound write is issued after the PCI master has disconnected, but before the last prefetch has been issued on the processor bus; in this situation, the outbound write will incorrectly be allowed to post.

IMPLICATION: This condition will result in improper deallocation of the data buffer before the outbound write completes, hanging the system.

WORKAROUND: Outbound write posting must be disabled in the 82454KX PB.

11. Inbound Posted Write with 1:1 Interleaving May Corrupt Data

PROBLEM: In systems configured with 1:1 interleaved memory (with only 2 SIMMs, for example) and inbound write posting enabled, if a string of line write transactions receive implicit writeback responses, data may be corrupted for subsequent writes if a refresh cycle occurs during the line writes. This occurs when the refresh backs up the transactions to memory such that the first partial write following the line write sequence is issued to the 82451 MIC with line-size encoding. The SYSCMD# and MEMCMD# signals will be correct, however.

IMPLICATION: Corrupted data will be written to memory in this case, with the data which was supposed to be written shifted by 8 bytes (data chunks 0-1-2-3 will be written as X-0-1-2, where X is most likely all 0's).

WORKAROUND: Do not enable inbound write posting in the 82454KX PB if memory is 1:1 interleaved, and use an IOQ depth of 1.

12. RAW May Cause Data Corruption

PROBLEM: If Read-Around-Write (RAW) is enabled in the 82453KX MC, some sequences of memory transactions may cause the 82453KX MC to issue the memory address (MA#) and row address strobe (RAS#) signals on the same clock edge for a DRAM read cycle, causing data corruption. The sequence which encounters this behavior is as follows:

1. A read to the 82454KX PB is snoop stalled.
2. During the snoop stall, a read-for-ownership (i.e. read and invalidate line) occurs.
3. Also during the snoop stall, a line write which is a row-hit/page-hit occurs.
4. A code line read which is a row-miss/page-miss occurs at the end of the snoop stall, but before snoop status is available for transactions 2 and 3.

IMPLICATION: When this occurs, the 82453KX MC will not start the write cycle until the snoop status is available for the preceding read. By then, the second read is in the internal queues behind the write, MA# will be driven with RAS# still asserted, and incorrect data will be returned for transaction 3.

WORKAROUND: Do not enable RAW in the 82453KX MC.

13. Data Transfer Order 3 May Cause Data to Be Issued Out of Order

PROBLEM: If a line write is issued which uses a data transfer order of 3 (with A[4:3] = 11, so that the chunks are issued in order 3-2-1-0) and receives an implicit writeback response, a line write immediately following this transaction may be issued with a reversed data transfer order, corrupting memory.

IMPLICATION: A line write whose address has A[4:3] = 11 (binary) will have a "critical chunk" of 3. According to the protocol established by the Intel486™ processor, this chunk must be written to memory first. A line write transaction which follows this one may have a different critical chunk (chunk 2, for example, issued in the order 2-3-0-1), but may be written to DRAM with a reversed order (1-0-3-2, in this case, or a data transfer order of 1). This causes the memory location to contain incorrect data. Note that neither the Pentium Pro processor nor the 450KX PCIset will issue line write transactions with this data transfer order (they will always have a data transfer order of 0). Third party agents which issue such transactions with a data transfer order of 3 will encounter this erratum.

WORKAROUND: Do not use third party agents which issue line writes with a data transfer order of 3 to the 82453KX MC.

14. PCI Address Parity Error May Cause Dropped Transaction

PROBLEM: If PCI address parity checking is enabled in the 82454KX PB, an inbound (PCI to processor bus) transaction with bad address parity will cause the 82454KX PB to assert SERR#. The DEVSEL# signal may then remain deasserted incorrectly. If a second inbound transaction occurs after the error, it will not be forwarded to the processor bus.

IMPLICATION: Though a PCI address parity error may be detected and flagged, a second transaction after the one with the error may be dropped, possibly corrupting data.

WORKAROUND: The system can be configured to promote the SERR# resulting from the parity error to an NMI, allowing software to handle the failing I/O system normally.

15. System Hang with Inbound Write Posting Enabled

PROBLEM: After a long write burst from the PCI to the processor bus, an outbound request will dynamically disable inbound write posting. The transaction immediately following the inbound burst may get a very fast retry. If this occurs, the 82454KX PB state machine may encounter a race condition which causes a bus hang on the next inbound request. The conditions for this to occur include:

- An inbound posted line write completes.
- The queue and data buffer for inbound transactions are full, or inbound write posting is dynamically disabled.
- The transaction after the inbound posted write is a non-dual address transaction with A[4:2] = 5 or 6, with only one idle cycle between it and the inbound posted line write.

This has only been observed in sequences containing a WRINV command on the PCI bus.

IMPLICATION: If inbound write posting is enabled in the 82454KX PB and the WRINV command is being used by a PCI device, the PCI bus may hang. The next outbound transaction will then hang the processor bus with BNR# toggling or with an infinite snoop stall.

WORKAROUND: Ensure non-overlapped PCI arbitration via the workaround documented under Erratum 6, *Inbound Write Posting May Cause Write Failure*. If this workaround is not used, inbound write posting must be disabled in the 82454KX PB to avoid this erratum.

16. RAW May Cause Data Corruption During Refresh

PROBLEM: If Read-Around-Write (RAW) is enabled in the 82453KX MC and there are no transactions in its queues, a refresh which should occur next may cause improper transaction ordering. This boundary condition is as follows:

1. A line read transaction is issued.
2. From 2 to 6 clocks after the ADS# assertion for the read, a refresh is requested.
3. A line write transaction is issued.
4. A line read transaction is issued which is a row-miss.

In this case, the refresh will be recognized before the write appears at the head of the queue, and will block the write from being handled. This allows the second line read to be reordered by RAW to occur before the write.

IMPLICATION: This boundary condition may cause data to be corrupted, if the data for the line read is dependent on the line write. Intel has observed this erratum in 1:1 interleaved configurations, and to a lesser extent in 2:1 interleaved configurations.

WORKAROUND: Do not enable RAW in the 82453KX MC.

17. GAT Devices May Time Out During Inbound Read Prefetch

PROBLEM: While in GAT (Guaranteed Access Time) mode, the 82454KX PB does not dynamically disable inbound read prefetching. Thus, a transaction which is issued immediately after MEMACK# is asserted may cause a prefetch to occur, tying up the bus and possibly allowing ISA or EISA devices configured for GAT to time out.

IMPLICATION: The 2.4 μs guaranteed access time may not be met for devices behind the 82454KX PB and an ISA or EISA bridge, resulting in spurious timeouts for such devices if inbound read prefetching is enabled, though most such devices will accept a longer access time.

WORKAROUND: If timeout problems occur while in GAT mode, disable inbound read prefetching in the 82454KX PB.

18. Data May Be Corrupted If RCD = 4 and LWC = 3

PROBLEM: The memory timings RCD (RAS# to CAS# Delay) and LWC (Last Write to CAS#) interact in such a way that if RCD is set to 4 and LWC is set to 3, the following transaction sequence will cause a data read to be corrupted:

1. A line read which triggers an implicit writeback and results in a page miss.
2. A line read which triggers an implicit writeback.
3. A line read to the same line as transaction 2.

IMPLICATION: These timing values may cause read data to be corrupted if the writeback memory type is used.

WORKAROUND: Do not use an RCD value of 4.

19. T_{CO_MAX} Specification Not Met for GTL+ Signals

PROBLEM: The T_{CO_MAX} specification for the GTL+ signals of the Intel 450KX PCIset is 6.0 ns. The actual value is 7.5 ns.

IMPLICATION: This value should be used when designing systems using the affected chipsets.

WORKAROUND: None identified.

20. RAW May Hang 1:1 or 2:1 Interleaved MP Systems

PROBLEM: If Read-Around-Write (RAW) is enabled in the 82453KX MC, there is a potential for two reads being assigned to the same internal data buffer, causing the system to hang. The following transactions must be issued during a snoop stall for a previous read for this erratum to occur:

1. A line read transaction.
2. A 0-byte length read transaction.
3. An explicit writeback transaction.
4. A line read transaction.

Transactions 1 and 4 may target the same buffer (causing the system to hang) in this situation. In A2 450KX PCIset silicon, it is also possible for a read transaction to incorrectly pass a write transaction with the same line address, but different chunk address. Both of these conditions require multiple symmetric bus agents (such as processors).

IMPLICATION: Uniprocessor systems will not be affected by this erratum. Other Intel 450KX PCIset-based systems will hang if this sequence occurs during a snoop stall.

WORKAROUND: Do not enable RAW in the 82453KX MC.

21. Mixed Interleave Increments May Cause Data Corruption

PROBLEM: If some DRAM module sizes are mixed together, some Intel 450KX PCIset-based systems may corrupt user or system data, resulting in incorrect calculations and/or system failure. This can only occur if:

- Interleave increments containing 8-Mbytes or 32-Mbytes of memory exist in conjunction with interleave increments which contain other amounts of memory, or
- One or more interleave increments have 8- or 32-Mbytes of memory, and there are no DRAM modules in the first row (row 0) of the memory subsystem.

Note that in systems with only one row of memory in use (i.e. 2 SIMMs in 1:1 interleave or 4 SIMMs in 2:1 interleave), or with all SIMMs or DIMMs the same size and configuration, only one interleave increment size is possible (since all memory in the same row must be the same size), so this erratum will not affect the system.

Also note that the BIOS of some (usually server) platforms may automatically downsize or eliminate a row of memory if a bad DRAM module is detected in that row. It is possible for this alteration of memory sizes to result in the system entering a vulnerable configuration (if memory is downsized to interleave increments of 8 or 32 Mbytes, or if row 0 is removed from the configuration). This can occur even with logically double-sided DRAM modules, which otherwise always have the same size interleave increments across their two rows.

IMPLICATION: Some systems only support one row of memory (i.e. 1:1 interleaving with 2 SIMMs or 2:1 interleaving with 4 SIMMs). In these systems, only one interleave increment size is possible (since all memory in the same row must be the same size), so they are not affected by this erratum. In systems which support two or more rows of memory (other than systems which support two rows of double-sided SIMM DRAM only),

one must be careful not to mix DRAM of certain sizes. If there are 8 or 32 Mbytes of DRAM in any interleave increment, all other interleave increments must have the same amount of DRAM, and row 0 must be populated with DRAM modules to avoid data corruption due to this erratum.

The chart below gives the interleave increments for some common memory technologies:

SIMMs	DRAM	Technology	Interleave Increment
Two 4-Mbyte, single-sided	1-Mbit x 4	4-Mbit	8 Mbytes*
Two 8-Mbyte, double-sided	1-Mbit x 4	4-Mbit	8 Mbytes*
Two 8-Mbyte, single-sided	2-Mbit x 8	16-Mbit	16 Mbytes
Two 16-Mbyte, double-sided	2-Mbit x 8	16-Mbit	16 Mbytes
Two 16-Mbyte, single-sided	4-Mbit x 4	16-Mbit	32 Mbytes*
Two 32-Mbyte, double-sided	4-Mbit x 4	16-Mbit	32 Mbytes*
Two 32-Mbyte, single-sided	8-Mbit x 8	64-Mbit	64 Mbytes
Two 64-Mbyte, double-sided	8-Mbit x 8	64-Mbit	64 Mbytes
Two 64-Mbyte, single-sided	16-Mbit x 4	64-Mbit	128 Mbytes
Two 128-Mbyte, double-sided	16-Mbit x 4	64-Mbit	128 Mbytes

DIMMs	DRAM	Technology	Interleave Increment
One 8-Mbyte, single-sided	1-Mbit x 4	4-Mbit	8 Mbytes*
One 16-Mbyte, double-sided	1-Mbit x 16	16-Mbit	8 Mbytes*
One 16-Mbyte, single-sided	2-Mbit x 8	16-Mbit	16 Mbytes
One 32-Mbyte, double-sided	2-Mbit x 8	16-Mbit	16 Mbytes
One 32-Mbyte, single-sided	4-Mbit x 4	16-Mbit	32 Mbytes*
One 64-Mbyte, double-sided	4-Mbit x16	64-Mbit	32 Mbytes*
One 64-Mbyte, single-sided	8-Mbit x 8	64-Mbit	64 Mbytes
One 128-Mbyte, double-sided	8-Mbit x 8	64-Mbit	64 Mbytes
One 128-Mbyte, single-sided	16-Mbit x 4	64-Mbit	128 Mbytes

NOTE:

*To avoid corruption of data, do not mix 8-Mbyte or 32-Mbyte interleave increments with other increment sizes (including each other). If 8-Mbyte or 32-Mbyte interleave increments are used (exclusively), ensure that row 0 is populated with DRAM modules.

WORKAROUND: When upgrading memory:

- In systems which support multiple rows of memory, always populate row 0 with DRAM modules.
- Exercise caution when using 4-Mbyte or 16-Mbyte SIMMs (single- or double- sided), or when using 8-Mbyte or 32-Mbyte DIMMs. Note that exclusively using DRAM modules which are all the same size will always avoid this erratum.
- When writing a BIOS for an Intel 450KX-based system, the memory configurations susceptible to this erratum should be detected and flagged to alert the user of the problem. Ensure that such detection mechanisms are placed after any bad DRAM detection mechanisms so that vulnerable configurations are detected after memory downsizing has occurred.

- Refer to your system documentation, or contact your system vendor, for details on your system's support for the various interleaving modes and DRAM styles.

See the white paper *Mixing DRAM Sizes with the 82450KX/GX PCIset* on Intel's World Wide Web site at URL http://www.intel.com/procs/support/ppro/450kxgx.htm for more information.

22. Parity Error May Occur for ADS# During BINIT#

PROBLEM: If BINIT# is asserted due to some catastrophic system event, the system will reset the bus and attempt to recover. If, however, ADS# is asserted for a request during the BINIT#, some of the signals associated with the request may not be driven when they should be. This results in an incorrect request parity and an assertion of AERR# 2 clocks after the BINIT# is complete, which is not a valid error phase.

IMPLICATION: A second (spurious) AERR# will be observed, resulting in a spurious SERR#. No data loss or hang is associated with this erratum, just an extra assertion of AERR# after detection of a catastrophic bus condition.

WORKAROUND: None identified.

23. Hang with PCI-to-PCI Bridges in MP Systems

PROBLEM: If a PCI-to-PCI bridge is present in a system, either on a card or native on the motherboard, a "livelock" hang condition may be possible if pipelined transactions are allowed (i.e. the IOQ depth is set to 8). This may happen if the following events occur:

1. An ordering event (such as I/O reads, and I/O writes for some components) from one processor is followed by a second such event from another processor. These outbound events are directed through the 82454KX PB, and dynamically disable inbound write posting.
2. The PCI-to-PCI bridge is attempting to perform an inbound write through the 82454KX PB.

If this situation arises, the PCI-to-PCI bridge will retry the outbound transactions, but must follow ordering constraints on the I/O operations. The outbound transactions are then reissued on the processor bus faster than the 82454KX PB can prepare to service the inbound write request from the PCI-to-PCI bridge. Consequently, the inbound write cannot complete, and the I/O transactions are continuously retried and re-issued on the processor bus, resulting in a "livelock" hang with neither side making forward progress.

IMPLICATION: PCI-to-PCI bridges present in an MP system may cause the system to hang.

WORKAROUND: Set the IOQ depth to 1. If this is not acceptable (multiprocessor systems may see significant performance degradation with this setting), a hardware workaround may be implemented using a programmable logic component. This workaround allows the outbound transactions to complete after they are retried by the PCI-to-PCI bridge, by putting the 82454KX PB in non-GAT (as opposed to GAT, or Guaranteed Access Time, mode) for 64 processor clocks after an outbound request is retried once, or after an inbound request is retried 8 times.

24. PCI_RST# Not Asserted Asynchronously

PROBLEM: The 82454KX PB requires a valid BCLK to assert the PCI_RST# signal, due to clocked logic in the assertion paths of these signals.

IMPLICATION: This signal is not asynchronous, as was intended.

WORKAROUND: Ensure a valid BCLK is provided before attempting to reset the PCI bus.

25. BERR# to BINIT# Conversion May Prevent Recovery From BINIT#

PROBLEM: If an error on a bus split operation occurs which causes an assertion of both BERR# and BINIT#, and BERR# to BINIT# conversion is enabled in the 82454KX PB, BINIT# will be driven a second time, 2 bus clocks after the first completes. At this time, another 450KX agent may still be driving BNR# for the first BINIT#, and other devices may then try to drive BNR# for the second BINIT#. This results in the system being unable to recover successfully from the BINIT#, hanging the system.

IMPLICATION: After a double assertion of BINIT# due to detection of a catastrophic error, the bus may not be reset properly, resulting in a system hang.

WORKAROUND: Do not enable BERR#-to-BINIT# conversion in the 82454KX PB. BERR#-to-SERR# conversion can be used instead.

26. Error Reporting Registers May Not Record Error Information Correctly

PROBLEM: The error reporting functionality of the 82453KX MC is limited in the following ways:

1. If a transaction occurs which contains an SBC (single-bit correctable) error, the error is detected, corrected, and logged (if ECC error detection, correction, and logging are enabled in the 82453KX MC). If a subsequent transaction contains an SBC error on one chunk (an 8-byte block of data) and an UNC (uncorrectable, double-bit) error on another, both will be detected, and the SBCERR# and BERR# signals will be asserted normally, but only the SBC error will actually be logged in the error reporting registers of the 82453KX MC.
2. If a transaction occurs with an UNC error, the error is tracked via an internal error buffer, and BERR# is asserted (if this feature is enabled). If the system does not shut down due to this error, and an SBC error occurs such that the UNC error is being popped from the error buffer at the same time that the SBC error is pushed into the buffer, the UNC error will not be logged into the error reporting registers of the 82453KX MC.

IMPLICATION: Under some circumstances, some errors will not be reported properly when detected in combination with other errors. This will result in an error not being logged into the error reporting registers of the 82453KX MC. However, this erratum does not affect the detection and correction of SBC errors or the detection of UNC errors.

WORKAROUND: None identified.

27. Combination of ECC Errors May Cause One Error to Be Undetected

PROBLEM: The 82452KX DP may not detect a memory ECC error during pipelined line reads, when there are multiple SBC and/or UNC errors across cache line boundaries. The error combinations which are affected are as follows:

1. An SBC error will detected and corrected, but not reported to the 82453KX MC, if it occurs on chunk 3 of the first line read when there is an UNC error on chunk 2 of the first line read and an SBC on chunk 0 of the second line read. In this case, no error record will exist for the first SBC, but data will not be corrupted.
2. An UNC error will be detected by the 82452KX DP but not reported to the 82453KX MC, if it occurs on chunk 3 of the first line read when there is an SBC error on chunk 2 of the first line read and either type of error on chunk 0 of the second line read. In this case, no error record will exist for the first UNC, and BERR# will not be driven if the third error is an SBC.

Note that these are the only combinations of errors affected by this erratum; other combinations of errors (including cache lines with *more* errors than specified) will all be detected, and any subsequent UNC errors which occur will cause BERR# to be driven by the 82453KX MC (assuming this feature is enabled).

IMPLICATION: An error will not be detected if it occurs with these very specific combinations, possibly resulting in corrupted data. The other errors in these combinations, as well as all other errors which occur on other transactions which do not fall within this pattern, will be detected correctly.

WORKAROUND: None identified.

28. 16-Byte Write May Hang System

PROBLEM: If a 16-byte partial write is issued in the middle of a sequence of line writes, the line write immediately after the 16-byte write may fail. A subsequent read transaction will hang the processor bus. Note that neither the Pentium Pro processor nor any Intel 450KX PCIset agent will generate 16-byte writes. Only systems containing third-party bus agents may be affected by this erratum.

IMPLICATION: If third-party agents which issue 16-byte writes are present in a system, the system may hang after a sequence of write transactions.

WORKAROUND: Do not use third-party agents which issue 16-byte writes.

29. 16-Byte Read with Two ECC Errors May Not Be Reported Correctly

PROBLEM: If a line read contains either type of error in chunk 3, and the line read is followed by a 16-byte read which has either an SBC error in chunk 0 and an UNC error in chunk 1 or an UNC error in chunk 1 and an SBC in chunk 0, a stale error record will be left in the error reporting registers which does not correspond to the last errors detected. All errors in this combination will be detected, however, and BERR# will be driven for UNC errors (assuming this feature is enabled). Note that neither the Pentium Pro processor nor any Intel 450KX PCIset agent will generate 16-byte reads; only systems containing third-party bus agents may be affected by this case.

IMPLICATION: Under some circumstances, some errors will not be reported properly when detected in combination with other errors. This may result in an error not being logged into the error reporting registers of the 82453KX MC at all, or it may result in incorrect error information being logged. This erratum does not affect the detection and correction of SBC errors or the detection of UNC errors.

WORKAROUND: Do not use third-party agents which issue 16-byte reads.

30. Inbound Locked PCI Transactions May Hang System

PROBLEM: This erratum manifests itself in the following sequence:

1. A locked read request is issued on the PCI bus, and targets the processor bus. The 82454KX PB asserts DEVSEL#.
2. The transaction is entered into the 82454KX PB's request queue, which then asserts BPRI#.
3. Before the ADS# for this transaction occurs, a pipelined retry response is received for the previous request.
4. ADS# is asserted on the processor bus for the locked read. In the sequence which causes the system to hang, the 82454KX PB confuses the retry for the previous transaction with a retry for this locked read. LOCK# is then asserted for only one clock, but the PCI LOCK# signal remains asserted.

5. The locked read's response phase completes on the processor bus and TRDY# is asserted on the PCI bus, completing the first part of the lock. Note that BPRI# is still asserted.
6. The locked write which is the second half of the inbound lock is issued on the PCI bus, and the 82454KX PB asserts DEVSEL#.
7. The locked write transaction is entered into the queue. However, the 82454KX PB cannot issue the transaction on the processor bus until PCI LOCK# is deasserted. The PCI bus cannot deassert PCI LOCK# until the write completes. This deadlock results in the system hanging.

IMPLICATION: If a PCI device driver allows the device to issue a locked transaction targeting the processor bus, and the IOQ depth is set to 8, a retry response for a transaction may cause the system to hang.

WORKAROUND: Drivers that are capable of issuing inbound locked PCI transactions are very rare. When using these drivers in a system, use an IOQ depth of 1.

31. Retry on Inbound Read May Corrupt Outbound Data

PROBLEM: If an inbound read is retried by a third-party agent, and DBSY# is still asserted from a normal data response for a pipelined read directed through the 82453KX MC, an outbound write can be issued while the outbound data buffer and status pointer are out of synch. This results in incorrect data being used for the outbound write.

IMPLICATION: Using third-party devices which issue pipelined retry responses may result in data corruption.

WORKAROUND: If third-party devices which issue retry responses are used in the system, use an IOQ depth of 1. Alternatively, a third-party device can delay its retry response until DBSY# is deasserted from the previous response.

32. ADS# in Last Clock of BINIT# Prevents Recovery

PROBLEM: If an ADS# assertion for a transaction targeting the 82453KX MC occurs during the last clock of an assertion of BINIT#, the ADS# will not be canceled correctly. The system will hang instead of recovering from the catastrophic condition which resulted in the BINIT# assertion.

IMPLICATION: BINIT# is only asserted upon detection of a catastrophic bus condition. If this occurs, the system may not generally be able to recover. This erratum decreases the possibility of being able to recover from this condition, but does not cause any additional incorrect behavior.

WORKAROUND: None identified.

33. Some Signals Indeterminate After RESET# Deassertion

PROBLEM: There exists a window of a single bus clock after RESET# has been deasserted where simulation has shown a possibility of the 82453KX MC sending a spurious command via the MEM_CMD# and/or MIC_CMD# signals to the 82452KX DP and/or 82451KX MICs.

IMPLICATION: If this were to occur, it might result in the 82452KX DP and/or 82451KX MICs attempting to execute a false command, most likely resulting in a system hang on startup. However, Intel has not identified any silicon 450KX component that has ever exhibited this condition, either under test or in actual systems; the MEM_CMD# and MIC_CMD# signals have always been observed to come up in a deterministic fashion.

WORKAROUND: None identified.

34. Delayed Read from PCI-to-PCI Bridge May Corrupt Data

PROBLEM: There exists a boundary condition in the 82454KX PB which may cause corruption of read data in an MP system, if the following sequence of events occurs:

1. A processor performs a destructive read directed towards a device below a PCI 2.1 Local Bus Specification compliant PCI-to-PCI bridge. The bridge delays the read and the 82454KX PB gives a retry response to the transaction. When this occurs, the PCI-to-PCI bridge starts a timer. The PCI-to-PCI bridge discards the read data from the target device if this timer expires.
2. A different processor initiates a configuration access (read or write) to a nonexistent device with a device ID of greater than 15. These accesses are allowed to time out on the processor bus (and are claimed by the 82454KX PB's watchdog timer), and are not forwarded to the PCI bus by the 82454KX PB.
3. Another processor issues a non-posted write to the 82454KX PB, or an inbound request targets the 82454KX PB before the first processor can retry the destructive read. If this occurs, the read cannot be forwarded to the PCI bus until the 82454KX PB's watchdog timer expires.

The PCI-to-PCI bridge's timer will expire after 2^{15} PCI bus clocks (which is less than the 82454KX PB's minimum watchdog timer value of 1.5 ms). After this, it will discard the data, as required by the PCI 2.1 Local Bus Specification. Since the read was destructive, the data is lost.

IMPLICATION: This erratum would typically occur when a driver is scanning for populated PCI device numbers in the system. If delayed reads are enabled in the PCI-to-PCI bridge, data corruption may result.

WORKAROUND: Do not enable delayed reads in PCI 2.1 Local Bus Specification compliant PCI-to-PCI bridge devices.

35. Page Open Policy of "Hold Page Open" May Corrupt Write Data

PROBLEM: If Read-Around-Write (RAW) is enabled, the Page Open Policy is set to "hold page open," and the IOQ depth is set to 8, there is a potential for data corruption after a configuration cycle is issued during a sequence of memory transactions, as follows:

1. A read from memory.
2. A write to the Configuration Address register.
3. A write to memory which is a page miss.
4. A read from memory which is a page hit (relative to transaction #1).

Or,

1. A read from memory with an ECC error logged.
2. A write to memory which is a page miss.
3. A read from memory which is a page hit (relative to transaction #1).

These sequences, with the configuration detailed above, will cause the data for the write to memory to be corrupted.

IMPLICATION: With all of these features enabled, data corruption may occur, resulting in unpredictable system failure.

WORKAROUND: Use the default Page Open Policy (close page). Disabling RAW and ERAW or using an IOQ depth of 1 will also prevent this erratum, but may have a larger impact on performance.

36. T_{CASE} Drop Plus Voltage Swing May Cause DPLL Failure

PROBLEM: Analysis of the DPLL (Digital Phase Lock Loop) units of the Intel 450KX PCIset components has shown that a DPLL failure can occur during certain changes in temperature and/or voltage within the Intel 450KX PCIset component specification, resulting in a loss of DPLL functionality for between 2 and 500 bus clocks.

The current specification states that the Intel 450KX/ PCIset is operational over a T_{CASE} range of 0 - 85°C and a V_{CC} range of 3.3V ± 5%.

After a device has undergone a hard reset of the DPLL (via the PLLRST pin), a drop in T_{CASE} accompanied by an increase in V_{CC} may cause this erratum to occur. The magnitude of the temperature drop and supply voltage increase required to cause this failure is graphed below:

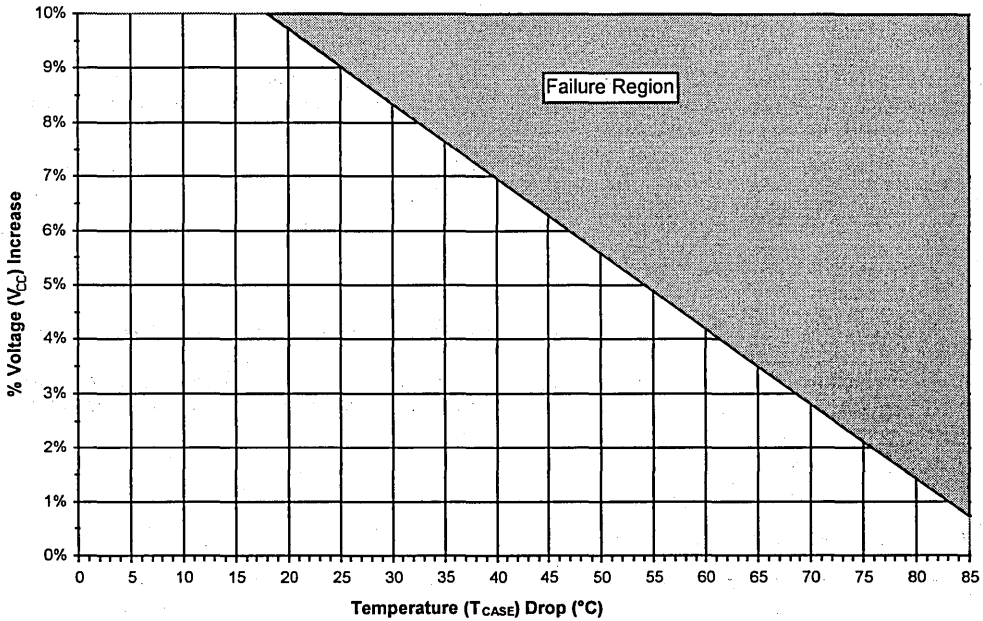

Figure 1. V_{CC} Increase vs. T_{CASE} Drop for DPLL Failure (Worst Case)

IMPLICATION: If the T_{CASE} of any Intel 450KX PCIset component is allowed to decrease by more than 18°C after the point at which DPLL lock is obtained after a DPLL hard reset, a positive voltage swing may cause the component to lose its internal clock for 2 - 500 processor bus clocks. This may result in a system hang, lost data, or other unpredictable system failure (note that if the component is not processing commands or data for the time during which the internal clock is lost, no failure will be observed). The larger the temperature drop and/or voltage swing, the more likely the failure.

If V_{CC} swings from 3.3V -5% to 3.3V +5%, and T_{CASE} does drop more than 18°C past the temperature at which the DPLL was reset, this erratum may occur. Similarly, if T_{CASE} transitions from 85°C (at DPLL reset

time) to 0°C, a swing in V_{CC} of more than 0.5% may cause this erratum. Thus, maintaining a moderately constant T_{CASE} and V_{CC} will prevent this erratum, as per the graph above.

WORKAROUND: Ensure that the DPLL reset occurs during a period where the T_{CASE} is low and/or the V_{CC} is high relative to normal operational conditions (e.g. if the system is running at a high temperature and is powered down, allow it to cool before powering up again). Alternately, ensure that the system meets a restricted specification for changes in T_{CASE} and V_{CC}.

37. Memory Gap Reclaiming May Corrupt Data

PROBLEM: If the memory subsystem is configured using the Memory Gap Register (MG), Low Memory Gap Register (LMG), or High Memory Gap Start Address and End Address registers (HMGSA, HMGEA) to allow a gap in the address map of the 82453KX MC, and if the memory in these gaps is reclaimed (by setting the appropriate bit in the MG, LMG, or HMGSA registers), the row address strobe (RAS#) signals may be corrupted.

IMPLICATION: An incorrect address may be generated for a memory access if any of these three memory gaps are enabled with the memory reclaimed. This would result in data corruption and unpredictable system failure.

WORKAROUND: If these memory gaps are used, do not reclaim the memory in the gaps.

38. RAW May Corrupt Write Data in 1:1 Interleaving

PROBLEM: If a memory subsystem is configured for 1:1 interleaving (i.e. only using one memory interleave, also known as non-interleaved), and read-around-write (RAW) is enabled, a partial write transaction to a memory location marked as Modified in the L2 cache followed by a line read transaction which is re-ordered around the partial write may result in the write data being corrupted.

IMPLICATION: With RAW enabled and the memory subsystem 1:1 interleaved, data corruption may occur, resulting in unpredictable system failure.

WORKAROUND: Do not enable RAW if the system is configured in 1:1 interleaving.

39. BINIT# Assertion May Cause Active RAS# Negation

PROBLEM: BINIT# is asserted on the Pentium Pro processor system bus to indicate the occurrence of a catastrophic system error condition or a situation that prevents reliable future operation. When the 82453GX MC component of the 450GX PCIset recognizes the assertion of BINIT# on the system bus, it prepares to clear certain internal state. In the process, a currently active RAS# signal associated with a read, write or refresh operation in progress may also be negated. If this negation causes the minimum RAS# pulse width timing to be violated then spurious bits may appear in one page of the memory array.

IMPLICATION: The catastrophic system error conditions that led to BINIT# assertion may in turn lead to a page of memory being incompletely written or improperly refreshed before potentially recovering from BINIT#.

WORKAROUND: Recovery from BINIT# assertion may be aided by a Machine Check Exception (MCE) or System Management Interrupt (SMI) handler. MCE or SMI handlers resident in non-volatile memory can protect recovery routines from encountering this situation, allowing them to be reliably executed after BINIT# assertion.

SPECIFICATION CLARIFICATIONS

1. Explicit Writebacks Claimed by 82454KX PB

Note 2 in Section 3.2 of the *Intel 450KX/GX PCIset Data Sheet* states that writebacks initiated by other agents are ignored by the PB. It should be noted that while this is true with respect to PCI bus transactions (i.e. no PCI bus cycles will be generated due to any writeback transaction), if a writeback occurs to memory behind the 82454KX PB, the data will be lost or a violation of processor bus protocol will occur. The 82454KX PB is not a caching agent, and no writeback transactions should be targeted to devices on the PCI bus. Memory behind the 82454KX PB may not be mapped as cacheable (WB type) memory.

2. Supported Configurations For MC Row Limit Register Programming

The following text will be added to the paragraph describing the DRL DRAM Row Limit Register in Section 2.3.13 of *Chapter 3 Memory Controller* of the Intel 450GX PCIset databook.

"The DRAM Row Limit registers may only be programmed in such a fashion that they do not violate the supported memory configurations outlined in Table 22, *Minimum and Maximum Memory Sizes for Each Configuration.* Other configurations of the row limit registers using increments other than those described are not supported."

DOCUMENTATION CHANGES

1. *Register Offset and Default Value Correction*

Some configuration registers are documented unclearly or inconsistently in the *Intel 450KX/GX PCIset* datasheet. A table of the correct offsets and values is given below with the changes in bold (note that this list only contains registers with changes):

Configuration Register	Address Offset	Default Value	Notes
82454KX			
Top of System Memory	40-43h	**0000** 0000h	
Bridge Device Number	49h	0001 1001**b**	
PB Configuration	4Ch	**1**9h	
PCI Read/Write Control	54-55h	**0**000h	
Memory Gap Range	78-79h	**0**000h	
Memory Gap Upper Address	7A-7Bh	**0**000h	
PCI Frame Buffer	7C-7Fh	**0000 0**000h	
High Memory Gap Range Start Address	88-8Bh	**0000 0**000h	
High Memory Gap End Address	8C-8Fh	**0000 0**000h	
Configuration Values Driven on Reset	B0-B1h	**0**000h	Bit 7: 1 = Depth of **1**. 0 = Depth of **8**. Pentium Pro processors use an in-order depth of **1** if this bit is 1.
Captured System Configuration Values	B4-B5h	**000X XXXX XXX0 0000b**	X = captured during hard reset. Bit 7: 1 = Depth of **1**. 0 = Depth of **8**. Pentium Pro processors use an in-order depth of **1** if this bit is 1.
SMRAM Range	B8-BBh	0000 0005h	Bits [15:0] correspond to A[31:**16**]#. The default starting address is **5**0000h and ranges to **5**FFFFh.
PB Retry Timers	C8-**CB**h	0000 0003h	
82453KX			
Controller Device Number	49h	0001010X**b**	X = loaded at reset
Single Bit Correctable Error Address	74-77h	**0000** 0000h	
Low Memory Gap Register	7C-7Fh	0010 0000h	**Bits [9:5]: Reserved.** Bits [4:0]: **Low Memory Gap Size.** This field defines the memory gap size as follows: **Bits [4:0] Size Bits [4:0] Size**

Configuration Register	Address Offset	Default Value	Notes
			00000 1 MB **00111** 8 MB **00001** 2 MB **01111** 16 MB **00011** 4 MB **11111** 32 MB
High Memory Gap Start Address	88-8Bh	**0000** 0000h	**Bit 30: Reclaim Enable.** 1 = Enable. 0 = Disable (default). When enabled, the physical memory in this gap is reclaimed.
High Memory Gap End Address	8C-8Fh	**0000** 0000h	
Memory Timing Register	AC-AFh	30DF3516h	Bit 15: 1 = **2** Cycles. 0 = **1** Cycle (default).
SMRAM Range	B8-BBh	**0000** 000Ah	

2. *CMOS Definition Should be 3.3V or 5V*

In Section 1.0 of both Chapter 2 and Chapter 3 of the *Intel 450KX/GX PCIset* datasheet, CMOS signals are defined as follows:

CMOS Rail-to-rail CMOS tolerant to 5V levels.

This should read:

CMOS Rail-to-rail CMOS tolerant to 3.3V or 5V levels. See Chapter 4, Section 1.2., *Signal Groups*.

Part II: Specification Update for the Intel 450GX PCIset

GENERAL INFORMATION

Top Markings

B- and C-step Production Units, BGA:

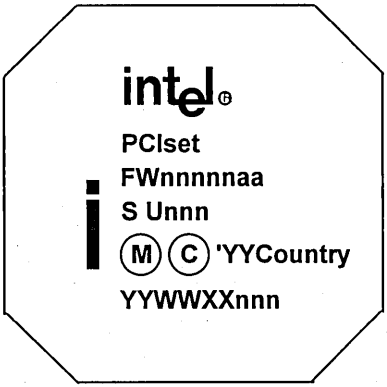

B- and C-step Production Units, QFP:

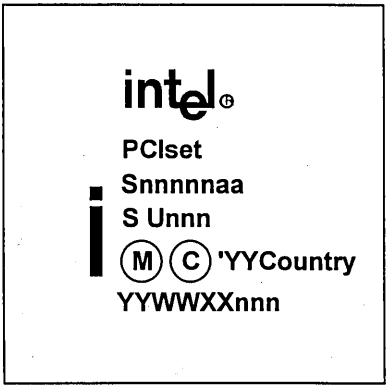

NOTES:

- nnnnnaa = Product Number.
- Q nnn = Sample Specification Number
- S Unnn = S-spec Number.
- 'YY Country = Copyright Dates and Country of Origin.
- YYWWXXnnn = Alternative Identification Number.

Basic Intel 450GX PCIset Identification Information

Product Number	Vendor ID	Device ID	Revision ID	Product Stepping	Kit Steppings	S-Spec	V_{CC}	T_{CASE}	Notes
S82451GX	n/a[1]	n/a[1]	n/a[1]	A1	B0	SU019	3.3V ± 5%	0°C - 85°C	
S82451GX	n/a[1]	n/a[1]	n/a[1]	A1	C0	SU019	3.3V ± 5%	0°C - 85°C	
S82452GX	n/a[1]	n/a[1]	n/a[1]	A3	B0	SU056	3.3V ± 5%	0°C - 85°C	
FW82452GX	n/a[1]	n/a[1]	n/a[1]	A3	B0	SU057	3.3V ± 5%	0°C - 85°C	2
S82452GX	n/a[1]	n/a[1]	n/a[1]	A4	C0	SY050	3.3V ± 5%	0°C - 85°C	
FW82452GX	n/a[1]	n/a[1]	n/a[1]	A4	C0	SY053	3.3V ± 5%	0°C - 85°C	2
S82453GX	8086h	84C5h	4	A4	B0	SU058	3.3V ± 5%	0°C - 85°C	
S82453GX	8086h	84C5h	5	A5	C0	SY051	3.3V ± 5%	0°C - 85°C	
S82454GX	8086h	84C4h	4	A4	B0	SU059	3.3V ± 5%	0°C - 85°C	
FW82454GX	8086h	84C4h	4	A4	B0	SU063	3.3V ± 5%	0°C - 85°C	2
S82454GX	8086h	84C4h	6	A6	C0	SY052	3.3V ± 5%	0°C - 85°C	
FW82454GX	8086h	84C4h	6	A6	C0	SY054	3.3V ± 5%	0°C - 85°C	2

NOTES:

1. These components are not visible from the PCI bus, and so do not have Vendor, Device, or Revision ID registers.
2. These components have BGA (Ball Grid Array) packaging.

Summary Table of Changes

The following table indicates the Specification Changes, Errata, Specification Clarifications, or Documentation Changes which apply to the Intel 450GX PCIset. Intel intends to fix some of the errata in future steppings of the component(s), and to account for the other outstanding issues through documentation or specification changes as noted. This table uses the following notations:

CODES USED IN SUMMARY TABLE

X:	Specification Change, Erratum, Specification Clarification, or Documentation Change applies to the given stepping.
Doc:	Intel intends to update the appropriate documentation in a future revision.
Fix:	Intel is investigating the possibility of fixing this erratum in a future stepping of the component(s).
Fixed:	This erratum has been previously fixed.
NoFix:	Intel is currently not investigating a fix for this erratum.
(No mark) or (blank box):	This item is fixed in or does not apply to the given kit stepping.
Shaded:	This erratum is either new or modified from the previous version of the document.

NO.	B0	C0	Plans	SPECIFICATION CHANGES
1	X	X	Doc	PLLRST pin added
2	X	X	Doc	Valid memory timing parameters
3	X	X	Doc	CMOS overshoot/undershoot specification
4		X	Doc	New features added to PDM register
5		X	Doc	Unused Pentium® Pro Processor Device Log register added
6		X	Doc	Locks to in-line shadowed BIOS not supported in Aliased GAT mode
7		X	Doc	AERR# to BERR# Conversion Enable bit added

NO.	B0	C0	Plans	ERRATA
1	X		Fixed	SMRAM addresses may not be decoded correctly
2	X		Fixed	Processor bus ECC error reporting may not record error information for one error
3	X		Fixed	RAW may hang 1:1 or 2:1 interleaved MP systems
4	X		Fixed	Mixed interleave increments may cause data corruption
5	X	X	NoFix	Parity error may occur for ADS# during BINIT#
6	X		Fixed	Configuration cycle to non-compatibility bridge may collide with inbound posted write
7	X		Fixed	Hang with PCI-to-PCI bridges in MP systems
8	X		Fixed	PCI_RST# not asserted asynchronously
9	X		Fixed	Combination of ECC errors may cause one error to be undetected
10	X		Fixed	BERR# to BINIT# conversion may prevent recovery from BINIT#
11	X	X	NoFix	16-Byte read with two ECC errors may not be reported correctly

NO.	B0	C0	Plans	ERRATA
12	X		Fixed	Inbound locked PCI transactions may hang system
13	X		Fixed	Retry on inbound read may corrupt outbound data
14	X		Fixed	ADS# in last clock of BINIT# prevents recovery
15	X		Fixed	Some signals indeterminate after RESET# deassertion
16	X		Fixed	Delayed read from PCI-to-PCI bridge may corrupt data
17	X		Fixed	Page Open Policy of "hold page open" may corrupt write data
18	X	X	NoFix	T_{CASE} drop plus voltage swing may cause DPLL failure
19	X		Fixed	Processor bus ECC signals may be corrupted
20	X		Fixed	IO_REQ# may not be deasserted for BPRI# during GAT, Non-GAT, or APIC Flush mode request
21	X		Fixed	Memory gap reclaiming may corrupt data
22	X		Fixed	RAW may corrupt write data in 1:1 interleaving
23		X	NoFix	Inbound write may be posted too soon after misaligned or multi-Dword cycle
24	X	X	NoFix	IO_REQ# may not be deasserted for BPRI# during GAT mode request
25	X	X	NoFix	BINIT# Assertion May Cause Active RAS# Negation

NO.	B0	C0	Plans	SPECIFICATION CLARIFICATIONS
1	X	X	Doc	Explicit writebacks claimed by 82454GX PB
2	X	X	Doc	Supported Configurations For MC Row Limit Register Programming

NO.	B0	C0	Plans	DOCUMENTATION CHANGES
1	X	X	Doc	Register offset and default value correction
2	X	X	Doc	CMOS definition should be 3.3V or 5V

SPECIFICATION CHANGES

1. PLLRST Pin Added

A PLLRST pin will be added to the definition of each device in the Intel 450GX PCIset. The pin numbers will be assigned as follows:

Device	Pin Number
82454GX PB QFP	301
82454GX PB BGA	A5
82453GX DC QFP	81
82452GX DP QFP	202
82452GX DP BGA	B12
82451GX MIC QFP	42

This signal will be added to the block diagram in Figure 1 of both Chapter 2 and Chapter 3.

Each of the specified PLLRST pins are 5V tolerant signals.

The signal will be added to Chapter 2, Section 1.4 and to Chapter 3, sections 1.1, 1.2 and 1.3 as signal "PLLRST", type "I, CMOS", and described as "This pin must be driven high for at least 2 clocks to reset the internal DPLL (Digital Phase Lock Loop). The DPLL should be reset after (or until) the clock input pins are stable at their final operating frequency. This pin does not have an edge rate requirement."

The following sentences will be added to Chapter 2, section 3.7.2 and Chapter 3, section 3.4: "The PLLRST pin must be driven high for at least 2 clocks to reset the internal DPLL. The DPLL should be reset after (or until) the clock input pins are stable at their final operating frequency."

Chapter 4 will be updated to include this pin information throughout.

2. Valid Memory Timing Parameters

The following is a list of timing values which have been validated by Intel. The list is the result of applying the rules set forth in the *Intel 450KX/GX PCIset* datasheet plus a set of filters to eliminate settings that Intel believes would not or could not be used in practice. Note that OEMs must still ensure that the timing parameters used meet the timing constraints for their system design, applicable clock rates, and supported DRAM speeds. See below for a list of acronyms used in the table.

RCD = 3, RCAD = 2, and CSR = 1 for all setting options listed below.

						1:1			2:1			4:1		
LWC	RASPW	CAH	RCAS	WCAS	RP	CP	RBD	ACh	CP	RBD	ACh	CP	RBD	ACh
2	5	1	2	2	3	1	2	2814	2	1	2834	1	0	2814
2	5	1	3	2	3	1	3	2914	1	1	2914	1	0	2914
2	5	1	3	3	3	1	3	2954	1	1	2954	1	0	2954
2	6	1	2	2	3	1	2	3014	2	1	3034	1	0	3014

						1:1			2:1			4:1		
LWC	RASPW	CAH	RCAS	WCAS	RP	CP	RBD	ACh	CP	RBD	ACh	CP	RBD	ACh
2	6	1	3	2	3	1	3	3114	1	1	3114	1	0	3114
2	6	1	3	2	4	1	3	3115	1	1	3115	1	0	3115
2	6	1	3	3	3	1	3	3154	1	1	3154	1	0	3154
2	6	1	3	3	4	1	3	3155	1	1	3155	1	0	3155
2	6	2	3	2	3	1	3	3514	1	1	3514	1	0	3514
2	6	2	3	2	4	1	3	3515	1	1	3515	1	0	3515
2	6	2	3	3	3	1	3	3554	1	1	3554	1	0	3554
2	6	2	3	3	4	1	3	3555	1	1	3555	1	0	3555
2	6	2	4	3	3	1	4	3654	2	2	3674	1	0	3654
2	6	2	4	3	4	1	4	3655	2	2	3675	1	0	3655
3	5	1	2	2	3	2	3	4834	2	1	4834	1	0	4814
3	5	1	3	2	3	2	4	4934				1	0	4914
3	5	1	3	3	3	2	4	4974				1	0	4954
3	6	2	3	2	3	2	4	5534				1	0	5514
3	6	2	3	2	4	2	4	5535				1	0	5515
3	6	2	3	3	3	2	4	5574				1	0	5554
3	6	2	3	3	4	2	4	5575				1	0	5555
3	6	2	4	3	3	2	5	5674	2	2	5674	1	0	5654
3	6	2	4	3	4	2	5	5675	2	2	5675	1	0	5655
3	6	2	4	4	4	1	4	5695	2	2	56B5	1	0	5695

RCD = 3, RCAD = 2, and CSR = 2 for all setting options listed below.

						1:1			2:1			4:1		
LWC	RASPW	CAH	RCAS	WCAS	RP	CP	RBD	ACh	CP	RBD	ACh	CP	RBD	ACh
3	6	2	3	4	3	1	3	D594	1	1	D594	1	0	D594
3	6	2	4	3	4	2	5	D675	2	2	D675	1	0	D655
2	6	2	3	3	3	1	3	B554	1	1	B554	1	0	B554

NOTE:

All above values are in number of cycles, not actual bit settings, except ACh (which is the actual bit setting, in hexadecimal).

RCD: RAS# to CAS# Delay: bits 3:2 of the Memory Timing Register (AC-AFh).

RCAD: RAS# to Column Address Delay: bit 4 of the Memory Timing Register (AC-AFh).

CSR: CAS# Setup to RAS# for CAS# before RAS# refresh: bit 15 of the Memory Timing Register (AC-AFh).

LWC: Last Write to CAS#: bits 14:13 of the Memory Timing Register (AC-AFh).

RASPW: RAS# Pulse Width: bits 12:11 of the Memory Timing Register (AC-AFh).

CAH: Column Address Hold Time: bit 10 of the Memory Timing Register (AC-AFh).

RCAS: Read CAS# Pulse Width: bits 9:8 of the Memory Timing Register (AC-AFh).

WCAS: Write CAS# Pulse Width: bits 7:6 of the Memory Timing Register (AC-AFh).

RP: RAS# Precharge Time: bits 1:0 of the Memory Timing Register (AC-AFh).
CP: CAS# Precharge Time: bit 5 of the Memory Timing Register (AC-AFh).
RBD: Read Burst Delay: bits 2:0 of the Command Register (4C-4Fh).
ACh: Actual hexadecimal value programmed into bits 15:0 of the Memory Timing Register (AC-AFh).

3. CMOS Overshoot/Undershoot Specification

The following will be added to Chapter 4, Section 2.4, *Intel 450KX/GX Undershoot Specification*, in the *Intel 450KX/GX PCIset Data Sheet,* and the section will be retitled *Intel 450KX/GX Overshoot/Undershoot Specifications.*

The 3.3V tolerant CMOS signals of the processor bus allow for the following maximum AC waveforms:

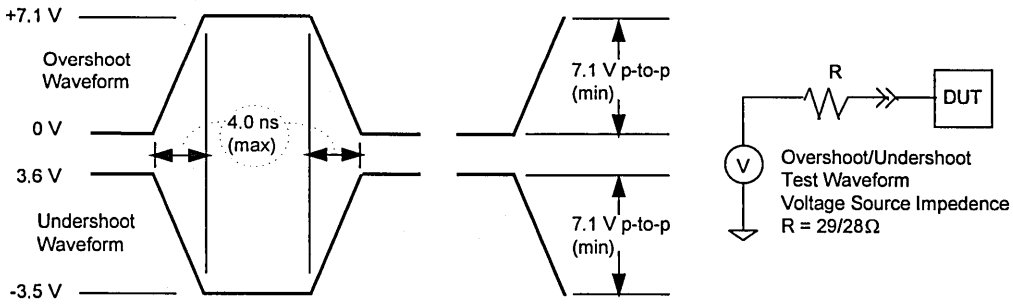

Average duty cycle of 17.6%

Note: This test covers the AC operating conditions only.

4. New Features Added to PDM Register

- GAT/non-GAT Optimization Changes: the optimizations previously associated with non-GAT mode (e.g. less frequent dynamic disabling of inbound write posting) have been added to GAT mode. Thus, EISA performance should see non-GAT-like performance while running in full GAT mode. In addition, support was added for Aliased GAT mode (i.e. putting the 82454 PB in non-GAT mode with the south bridge in GAT mode). This mode allows the inbound write posting optimization to take effect, in addition to allowing the deassertion of BPRI# between requests. Bit 3 of the PCI Decode Mode register at offset 48h is used for this purpose; setting this bit to '1' enables Aliased GAT mode. Table 1 details the optimizations.

Table 1. GAT/non-GAT Optimization Changes in C0 82454GX PB Silicon

Operation	A2 and B0 Steppings	C0 Stepping
82454GX PB receiving GAT mode requests, South bridge in GAT mode	Inbound write posting dynamically disabled, BPRI# held asserted	Inbound write posting kept enabled, BPRI# held asserted

Operation	A2 and B0 Steppings	C0 Stepping
82454GX PB receiving non-GAT mode requests (Aliased GAT mode), South bridge in GAT mode, Both sideband signals aliased to FLUSHBUF	Inbound write posting kept enabled, BPRI# held asserted	Inbound write posting kept enabled, BPRI# deasserted after MEMACK# Locks to shadowed BIOS not supported
82454GX PB receiving any sideband requests, South bridge in non-GAT mode	Not supported	Not supported

- Livelock Prevention Disable: The workaround for Intel 450GX Erratum 7 has been incorporated internally into the C0 stepping of the 82454GX PB. A mechanism is provided to disable this workaround, by setting bit 5 of the PCI Decode Mode register at offset 48h to a '1'.
- Traffic Priority Mode: This bit can be asserted to override the priority normally given to outbound requests. If this bit is zero (the default state), then non-GAT mode is entered for 64 clocks after an inbound request is retried eight times. Setting bit 6 of the PCI Decode Mode register at offset 48h to '1' disables this new feature. Intel recommends that BIOS override the default value and set this bit to '1' in systems with PCI Local Bus Specification 2.1 compliant PCI-to-PCI bridges which may experience extremely heavy inbound traffic; otherwise, temporary processor read starvation could result, reducing performance. Note that in other systems, especially those which experience heavy outbound traffic, this bit should remain at the default setting of '0'; otherwise, inbound read starvation may result, also impacting performance.

Configuration Register	Address Offset	Default Value	Notes
82454GX			
PCI Decode Mode	48h	06h	Bit 6: Traffic priority mode Bit 5: Livelock prevention disable Bit 3: Aliased GAT mode enable

5. *Unused Pentium® Pro Processor Device Log Register Added*

A new 16-bit register, called the "Unused Pentium Pro Processor Device Log", has been added to the 82454GX PB at address offset CCh. This register should be programmed by BIOS to reflect the populated and unpopulated device IDs between 16 and 31 on the host bus (bus 0). The BIOS algorithm should perform a scan of the range from 16 to 31 on bus 0, detect each location that returns all 1's (indicating an unpopulated device ID) from offset 00h in their configuration space, and set the corresponding bit in 82454GX PB configuration register CCh. (Bits corresponding to populated device ID's must remain "0" (deasserted). Bit *n* of this register corresponds to device ID *n* + 16.) Previously, all reads to locations which return all 1's from offset 00h generated by further PCI scans would have timed-out (been claimed by the watchdog timer in the 82454GX PB); the new register will force the 82454GX PB to claim these transactions via positive decode and forward them to the PCI bus, where they will result in a PCI master abort. This will provide much faster missing device access handling than the previous watchdog timer mechanism in the 82454GX PB (which will still be used if this register is left with its default value of 00h). The new treatment of these cycles is necessary to avoid a potential for loss of read data in the presence of PCI 2.1 compliant PCI-to-PCI bridges during PCI device scans (see Intel 450GX Erratum #16).

Configuration Register	Address Offset	Default Value	Notes
82454GX			
Unused Pentium Pro Processor Device Log	CCh	00h	Bit *n* corresponds to device ID *n* + 16. 1 = device ID is unpopulated (as determined by a scan of bus 0). 0 = device present at ID *n* + 16.

6. *Locks to In-Line Shadowed BIOS Not Supported in Aliased GAT Mode*

As described in Intel 450GX Specification Change 4, *New Features Added to PCM Register*, bit 3 of the PCI Decode Mode register in the 82454GX PB is used to put the PCI Bridge into Aliased GAT mode. Section 3.3, *PCI Bus Interface*, of the *Intel 450KX/GX PCIset* datasheet describes the 82454GX PB's support for host bus locks. The following information will be added to this section:

Locked transactions targeting in-line shadowed BIOS regions using conventional write protection (i.e. read-only to memory and write only to the PB) are not supported if Aliased GAT mode is enabled. If locked transactions must be issued to BIOS with Aliased GAT mode enabled, then the relevant BIOS region must be either:

- Mapped read/write in memory (forgoing write protection),
- Read only in memory, write only to a third party device, or
- Read only in memory, allowing the PB watchdog timer to time out write transactions.

7. *AERR# to BERR# Conversion Enable Bit Added*

Bit 1 of the PB Extended Error Reporting Command register (EXERRCMD) at offset C0h is Reserved for Intel 450GX PCIsets of B0 stepping; for the C0 stepping of the Intel 450GX PCIset, this bit can be set to '1' to enable the assertion of the BERR# signal upon detection of an assertion of AERR#.

Configuration Register	Address Offset	Default Value	Notes
82454GX			
PB Extended Error Reporting Command	C0h	0000 0010h	Bit 1: AERR# to BERR# enable.

ERRATA

1. SMRAM Addresses May Not Be Decoded Correctly

PROBLEM: While executing in SMM (System Management Mode), certain sequences of transactions may allow the CAS# signal to be asserted without a corresponding RAS# signal during a memory access to SMRAM. The sequences must be pipelined SMRAM requests (with the IOQ depth set to 8) as follows:

- An SMRAM request which opens a page (i.e. the transaction address accesses a block of DRAM with a new row address).
- An SMRAM request which is a page hit.
- Another SMRAM request which is a page hit.

The address phase of the third request must occur during the first clock of CAS# assertion for the second request to encounter this erratum. A second sequence is:

- An SMRAM request which is a page hit.
- Another SMRAM request whose ADS# assertion comes 5 or 6 clocks later than the ADS# assertion for the previous transaction, and which is also a page hit.

IMPLICATION: If these sequences occur while in SMM, data corruption may result.

WORKAROUND: This erratum can be avoided in the Intel 450GX PCIset by setting the IOQ depth to 1, thus preventing transactions from being pipelined together.

2. Processor Bus ECC Error Reporting May Not Record Error Information for One Error

PROBLEM: If a line transaction returns data to the processor bus with an UNC or SBC ECC error detected on chunk 0, it will be reported correctly. If, however, an error of the opposite type (SBC vs. UNC) is detected on chunk 1 or chunk 2, this error will not be reported in the error reporting registers. However, an SBC error will be detected and corrected appropriately by the 82452GX DP, and an UNC error will cause BERR# to be driven (assuming this feature is enabled).

IMPLICATION: After this specific sequence of errors, one error will not be reported properly when detected in combination with other errors. This will result in an error not being logged into the error reporting registers of the 82453GX MC. However, this erratum does not affect the detection and correction of SBC errors or the detection of UNC errors.

WORKAROUND: None identified.

3. RAW May Hang 1:1 or 2:1 Interleaved MP Systems

PROBLEM: If Read-Around-Write (RAW) is enabled in the 82453GX MC, there is a potential for two reads being assigned to the same internal data buffer, causing the system to hang. The following transactions must be issued during a snoop stall for a previous read for this erratum to occur:

1. A line read transaction
2. A 0-byte length read transaction
3. An explicit writeback transaction
4. A line read transaction

Transactions 1 and 4 may target the same buffer in this situation.

IMPLICATION: Systems which use 4:1 interleaving and uniprocessor systems will not be affected by this erratum. Other systems will hang if this sequence occurs during a snoop stall.

WORKAROUND: Do not enable RAW in the 82453GX MC in a 1:1 or 2:1 interleaved MP configuration.

4. *Mixed Interleave Increments May Cause Data Corruption*

PROBLEM: If some DRAM module sizes are mixed together, some Intel 450GX PCIset-based systems may corrupt user or system data, resulting in incorrect calculations and/or system failure. This can only occur if:

- Interleave increments containing 8-Mbytes or 32-Mbytes of memory exist in conjunction with interleave increments which contain other amounts of memory, or
- One or more interleave increments have 8- or 32-Mbytes of memory, and there are no DRAM modules in the first row (row 0) of the memory subsystem.

Note that in systems with only one row of memory in use (i.e. 2 SIMMs in 1:1 interleave, 4 SIMMs in 2:1 interleave, or 8 SIMMs in 4:1 interleave), or with all SIMMs or DIMMs the same size and configuration, only one interleave increment size is possible (since all memory in the same row must be the same size), so this erratum will not affect the system.

Also note that the BIOS of some (usually server) platforms may automatically downsize or eliminate a row of memory if a bad DRAM module is detected in that row. It is possible for this alteration of memory sizes to result in the system entering a vulnerable configuration (if memory is downsized to interleave increments of 8 or 32 Mbytes, or if row 0 is removed from the configuration). This can occur even with logically double-sided DRAM modules, which otherwise always have the same size interleave increments across their two rows.

IMPLICATION: Some systems only support one row of memory (i.e. 1:1 interleaving with 2 SIMMs, 2:1 interleaving with 4 SIMMs, or 4:1 interleaving with 8 SIMMs). In these systems, only one interleave increment size is possible (since all memory in the same row must be the same size), so they are not affected by this erratum. In systems which support two or more rows of memory (other than systems which support two rows of double-sided SIMM DRAM only), one must be careful not to mix DRAM of certain sizes. If there are 8 or 32 Mbytes of DRAM in any interleave increment, all other interleave increments must have the same amount of DRAM, and row 0 must be populated with DRAM modules to avoid data corruption due to this erratum.

The chart below gives the interleave increments for some common memory technologies:

SIMMs	DRAM	Technology	Interleave Increment
Two 4-Mbyte, single-sided	1-Mbit x 4	4-Mbit	8 Mbytes*
Two 8-Mbyte, double-sided	1-Mbit x 4	4-Mbit	8 Mbytes*
Two 8-Mbyte, single-sided	2-Mbit x 8	16-Mbit	16 Mbytes
Two 16-Mbyte, double-sided	2-Mbit x 8	16-Mbit	16 Mbytes
Two 16-Mbyte, single-sided	4-Mbit x 4	16-Mbit	32 Mbytes*
Two 32-Mbyte, double-sided	4-Mbit x 4	16-Mbit	32 Mbytes*
Two 32-Mbyte, single-sided	8-Mbit x 8	64-Mbit	64 Mbytes
Two 64-Mbyte, double-sided	8-Mbit x 8	64-Mbit	64 Mbytes
Two 64-Mbyte, single-sided	16-Mbit x 4	64-Mbit	128 Mbytes
Two 128-Mbyte, double-sided	16-Mbit x 4	64-Mbit	128 Mbytes

DIMMs	DRAM	Technology	Interleave Increment
One 8-Mbyte, single-sided	1-Mbit x 4	4-Mbit	8 Mbytes*
One 16-Mbyte, double-sided	1-Mbit x 16	16-Mbit	8 Mbytes*
One 16-Mbyte, single-sided	2-Mbit x 8	16-Mbit	16 Mbytes
One 32-Mbyte, double-sided	2-Mbit x 8	16-Mbit	16 Mbytes
One 32-Mbyte, single-sided	4-Mbit x 4	16-Mbit	32 Mbytes*
One 64-Mbyte, double-sided	4-Mbit x16	64-Mbit	32 Mbytes*
One 64-Mbyte, single-sided	8-Mbit x 8	64-Mbit	64 Mbytes
One 128-Mbyte, double-sided	8-Mbit x 8	64-Mbit	64 Mbytes
One 128-Mbyte, single-sided	16-Mbit x 4	64-Mbit	128 Mbytes

NOTE:

*To avoid corruption of data, do not mix 8-Mbyte or 32-Mbyte interleave increments with other increment sizes (including each other). If 8-Mbyte or 32-Mbyte interleave increments are used (exclusively), ensure that row 0 is populated with DRAM modules.

WORKAROUND: When upgrading memory:

- In systems which support multiple rows of memory, always populate row 0 with DRAM modules.
- Exercise caution when using 4-Mbyte or 16-Mbyte SIMMs (single- or double- sided), or when using 8-Mbyte or 32-Mbyte DIMMs. Note that exclusively using DRAM modules which are all the same size will always avoid this erratum.
- When writing a BIOS for an Intel 450GX-based system, the memory configurations susceptible to this erratum should be detected and flagged to alert the user of the problem. Ensure that such detection mechanisms are placed after any bad DRAM detection mechanisms so that vulnerable configurations are detected after memory downsizing has occurred.
- Refer to your system documentation, or contact your system vendor, for details on your system's support for the various interleaving modes and DRAM styles.

See the white paper *Mixing DRAM Sizes with the 82450KX/GX PCIset* on Intel's World Wide Web site at URL http://www.intel.com/procs/support/ppro/450kxgx.htm for more information.

5. *Parity Error May Occur for ADS# During BINIT#*

PROBLEM: If BINIT# is asserted due to some catastrophic system event, the system will reset the bus and attempt to recover. If, however, ADS# is asserted for a request during the BINIT#, some of the signals associated with the request may not be driven when they should be. This results in an incorrect request parity and an assertion of AERR# 2 clocks after the BINIT# is complete, which is not a valid error phase.

IMPLICATION: A second (spurious) AERR# will be observed, resulting in a spurious SERR#. No data loss or hang is associated with this erratum, just an extra assertion of AERR# after detection of a catastrophic bus condition.

WORKAROUND: None identified.

6. Configuration Cycle to Non-Compatibility Bridge May Collide with Inbound Posted Write

PROBLEM: It is possible for a non-compatibility 82454GX PB (nc82454) in a dual-bridge system to respond incorrectly to a CONFDATA cycle. This is due to a boundary condition in the 82454GX PB which may result in the nc82454's CONFADDR register not being updated for a previous CONFADDR cycle. This, in turn, may result in the nc82454 either claiming a subsequent CONFDATA cycle intended for another PCI device, or not responding to a subsequent CONFDATA cycle intended for it. This boundary condition arises only when configuration cycles are directed at the nc82454 during inbound posted write traffic.

IMPLICATION: This erratum may result in corrupted configuration space of bridges or PCI devices in dual-82454GX PB systems, possibly resulting in assertion of BINIT#.

WORKAROUND: Use an IOQ depth of 1 on dual-82454GX PB systems. This will prevent the pipelining necessary for the inbound posted write and outbound configuration cycle to collide. If an IOQ depth of 8 is desired, do not enable inbound write posting in the non-compatibility 82454GX PB. In this case, BIOS level and operating-system level workarounds (if present) will re-enable inbound posting in the non-compatibility 82454GX PB and work around the erratum via BIOS code and OS software. The OS workaround will be incorporated into future revisions of some operating systems; contact your OS vendor for release details. These workarounds require that:

1. There are two 82454GX PB's in the system of stepping B0 or earlier.
2. Inbound write posting is enabled in the compatibility 82454GX PB.

If these conditions are met, the operating system and BIOS must do the following whenever configuration cycles are issued:

1. In a loop, read the Vendor ID from the nc82454's configuration space, until it matches the known value. When this occurs, the OS is pointing at the nc82454's configuration space.
2. In a loop, read the PCI Read/Write Control register. When this value no longer matches the Vendor ID, the OS is pointing at the correct register in configuration space.
3. Write a value into the PCI Read/Write Control register which disables inbound write posting.
4. Issue the configuration cycles as originally intended.
5. Write a value into the nc82454's PCI Read/Write Control register to re-enable inbound write posting.

7. Hang with PCI-to-PCI Bridges in MP Systems

PROBLEM: If a PCI-to-PCI bridge is present in a system, either on a card or native on the motherboard, a "livelock" hang condition may be possible if pipelined transactions are allowed (i.e. the IOQ depth is set to 8). This may happen if the following events occur:

1. An ordering event (such as I/O reads, and I/O writes for some components) from one processor is followed by a second such event from another processor. These outbound events are directed through the 82454GX PB, and dynamically disable inbound write posting.
2. The PCI-to-PCI bridge is attempting to perform an inbound write through the 82454GX PB.

If this situation arises, the PCI-to-PCI bridge will retry the outbound transactions, but must follow ordering constraints on the I/O operations. The outbound transactions are then reissued on the processor bus faster than the 82454GX PB can prepare to service the inbound write request from the PCI-to-PCI bridge. Consequently, the inbound write cannot complete, and the I/O transactions are continuously retried and re-issued on the processor bus, resulting in a "livelock" hang with neither side making forward progress.

IMPLICATION: PCI-to-PCI bridges present in an MP system may cause the system to hang.

WORKAROUND: Set the IOQ depth to 1. If this is not acceptable (multiprocessor systems may see significant performance degradation with this setting), a hardware workaround may be implemented using a programmable logic component. This workaround allows the outbound transactions to complete after they are retried by the PCI-to-PCI bridge, by putting the 82454GX PB in non-GAT (as opposed to GAT, or Guaranteed Access Time, mode) for 64 processor clocks after an outbound request is retried once, or after an inbound request is retried 8 times.

8. PCI_RST# Not Asserted Asynchronously

PROBLEM: The 82454GX PB requires a valid BCLK to assert the PCI_RST# signal, due to clocked logic in the assertion paths of these signals.

IMPLICATION: This signal is not asynchronous, as was intended.

WORKAROUND: Ensure a valid BCLK is provided before attempting to reset the PCI bus.

9. Combination of ECC Errors May Cause One Error to Be Undetected

PROBLEM: The 82452GX DP may not detect a memory ECC error during pipelined line reads, when there are multiple SBC and/or UNC errors across cache line boundaries. The error combinations which are affected are as follows:

1. An SBC error will detected and corrected, but not reported to the 82453GX MC, if it occurs on chunk 3 of the first line read when there is an UNC error on chunk 2 of the first line read and an SBC on chunk 0 of the second line read. In this case, no error record will exist for the first SBC, but data will not be corrupted.
2. An UNC error will be detected by the 82452GX DP but not reported to the 82453GX MC, if it occurs on chunk 3 of the first line read when there is an SBC error on chunk 2 of the first line read and either type of error on chunk 0 of the second line read. In this case, no error record will exist for the first UNC, and BERR# will not be driven if the third error is an SBC.

Note that these are the only combinations of errors affected by this erratum; other combinations of errors will all be detected (including cache lines with *more* errors than specified), and any subsequent UNC errors which occur will cause BERR# to be driven by the 82453GX MC (assuming this feature is enabled).

IMPLICATION: An error will not be detected if it occurs with these very specific combinations, possibly resulting in corrupted data. The other errors in these combinations, as well as all other errors which occur on other transactions which do not fall within this pattern, will be detected correctly.

WORKAROUND: None identified.

10. BERR# to BINIT# Conversion May Prevent Recovery From BINIT#

PROBLEM: If an error on a bus split operation occurs which causes an assertion of both BERR# and BINIT#, and BERR# to BINIT# conversion is enabled in the 82454GX PB, BINIT# will be driven a second time, 2 bus clocks after the first completes. At this time, another 450GX agent may still be driving BNR# for the first BINIT#, and other devices may then try to drive BNR# for the second BINIT#. This results in the system being unable to recover successfully from the BINIT#, hanging the system.

IMPLICATION: After a double assertion of BINIT# due to detection of a catastrophic error, the bus may not be reset properly, resulting in a system hang.

WORKAROUND: Do not enable BERR#-to-BINIT# conversion in the 82454GX PB. BERR#-to-SERR# conversion can be used instead.

11. 16-Byte Read with Two ECC Errors May Not Be Reported Correctly

PROBLEM: If a line read contains either type of error in chunk 3, and the line read is followed by a 16-byte read which has either an SBC error in chunk 0 and an UNC error in chunk 1 or an UNC error in chunk 1 and an SBC in chunk 0, a stale error record will be left in the error reporting registers which does not correspond to the last errors detected. All errors in this combination will be detected, however, and BERR# will be driven for UNC errors (assuming this feature is enabled). Note that neither the Pentium Pro processor nor any 450GX PCIset agent will generate 16-byte reads; only systems containing third-party bus agents may be affected by this case.

IMPLICATION: Under some circumstances, some errors will not be reported properly when detected in combination with other errors. This may result in an error not being logged into the error reporting registers of the 82453GX MC at all, or it may result in incorrect error information being logged. This erratum does not affect the detection and correction of SBC errors or the detection of UNC errors.

WORKAROUND: Do not use third-party agents which issue 16-byte reads.

12. Inbound Locked PCI Transactions May Hang System

PROBLEM: This erratum manifests itself in the following sequence:

1. A locked read request is issued on the PCI bus, and targets the processor bus. The 82454KX PB asserts DEVSEL#.
2. The transaction is entered into the 82454GX PB's request queue, which then asserts BPRI#.
3. Before the ADS# for this transaction occurs, a pipelined retry response is received for the previous request.
4. ADS# is asserted on the processor bus for the locked read. In the sequence which causes the system to hang, the 82454GX PB confuses the retry for the previous transaction with a retry for this locked read. LOCK# is then asserted for only one clock, but the PCI LOCK# signal remains asserted.
5. The locked read's response phase completes on the processor bus and TRDY# is asserted on the PCI bus, completing the first part of the lock. Note that BPRI# is still asserted.
6. The locked write which is the second half of the inbound lock is issued on the PCI bus, and the 82454KX PB asserts DEVSEL#.
7. The locked write transaction is entered into the queue. However, the 82454GX PB cannot issue the transaction on the processor bus until PCI LOCK# is deasserted. The PCI bus cannot deassert PCI LOCK# until the write completes. This deadlock results in the system hanging.

IMPLICATION: If a PCI device driver allows the device to issue a locked transaction targeting the processor bus, and the IOQ depth is set to 8, a retry response for a transaction may cause the system to hang.

WORKAROUND: Drivers that are capable of issuing inbound locked PCI transactions are very rare. When using these drivers in a system, use an IOQ depth of 1.

13. Retry on Inbound Read May Corrupt Outbound Data

PROBLEM: If an inbound read is retried by a third-party agent or second 82454GX PB (in a peer-to-peer PCI transaction), and DBSY# is still asserted from a normal data response for a pipelined read directed through the 82453GX MC, an outbound write can be issued while the outbound data buffer and status pointer are out of synch. This results in incorrect data being used for the outbound write.

IMPLICATION: Using third-party devices which issue pipelined retry responses or drivers which cause PB-to-PB traffic may result in data corruption. Intel has not currently identified any software which is capable of causing PB-to-PB traffic in a dual-82454GX PB system.

WORKAROUND: If third-party devices which issue retry responses are used in the system, or if drivers are used in a dual-82454GX PB system which allow PB-to-PB traffic, use an IOQ depth of 1. Alternatively, a third-party device can delay its retry response until DBSY# is deasserted from the previous response. Also, ensure that drivers which support peer-to-peer PCI transactions only do so for devices under the same 82454GX PB.

14. ADS# in Last Clock of BINIT# Prevents Recovery

PROBLEM: If an ADS# assertion for a transaction targeting the 82453GX MC occurs during the last clock of an assertion of BINIT#, the ADS# will not be canceled correctly. The system will hang instead of recovering from the catastrophic condition which resulted in the BINIT# assertion.

IMPLICATION: BINIT# is only asserted upon detection of a catastrophic bus condition. If this occurs, the system may not generally be able to recover. This erratum decreases the possibility of being able to recover from this condition, but does not cause any additional incorrect behavior.

WORKAROUND: None identified.

15. Some Signals Indeterminate After RESET# Deassertion

PROBLEM: There exists a window of a single bus clock after RESET# has been deasserted where simulation has shown a possibility of the 82453GX MC sending a spurious command via the MEM_CMD# and/or MIC_CMD# signals to the 82452GX DP and/or 82451GX MICs.

IMPLICATION: If this were to occur, it might result in the 82452GX DP and/or 82451GX MICs attempting to execute a false command, most likely resulting in a system hang on startup. However, Intel has not identified any silicon 450GX component that has ever exhibited this condition, either under test or in actual systems; the MEM_CMD# and MIC_CMD# signals have always been observed to come up in a deterministic fashion.

WORKAROUND: None identified.

16. Delayed Read from PCI-to-PCI Bridge May Corrupt Data

PROBLEM: There exists a boundary condition in the 82454GX PB which may cause corruption of read data in an MP system, if the following sequence of events occurs:

1. A processor performs a destructive read directed towards a device below a PCI 2.1 Local Bus Specification compliant PCI-to-PCI bridge. The bridge delays the read and the 82454GX PB gives a retry response to the transaction. When this occurs, the PCI-to-PCI bridge starts a timer. The PCI-to-PCI bridge discards the read data from the target device if this timer expires.
2. A different processor initiates a configuration access (read or write) to a nonexistent device with a device ID of greater than 15. These accesses are allowed to time out on the processor bus (and are claimed by the 82454GX PB's watchdog timer), and are not forwarded to the PCI bus by the 82454GX PB.
3. Another processor issues a non-posted write to the 82454GX PB, or an inbound request targets the 82454GX PB before the first processor can retry the destructive read. If this occurs, the read cannot be forwarded to the PCI bus until the 82454GX PB's watchdog timer expires.

The PCI-to-PCI bridge's timer will expire after 2^{15} PCI bus clocks (which is less than the 82454GX PB's minimum watchdog timer value of 1.5 ms). After this, it will discard the data, as required by the PCI 2.1 Local Bus Specification. Since the read was destructive, the data is lost.

IMPLICATION: This erratum would typically occur when a driver is scanning for populated PCI device numbers in the system. If delayed reads are enabled in the PCI-to-PCI bridge, data corruption may result.

WORKAROUND: Do not enable delayed reads in PCI 2.1 Local Bus Specification compliant PCI-to-PCI bridge devices.

17. *Page Open Policy of "Hold Page Open" May Corrupt Write Data*

PROBLEM: If Read-Around-Write (RAW) is enabled, the Page Open Policy is set to "hold page open," and the IOQ depth is set to 8, there is a potential for data corruption after a configuration cycle is issued during a sequence of memory transactions, as follows:

1. A read from memory.
2. A write to the Configuration Address register.
3. A write to memory which is a page miss.
4. A read from memory which is a page hit (relative to transaction #1).

Or,

1. A read from memory with an ECC error logged.
2. A write to memory which is a page miss.
3. A read from memory which is a page hit (relative to transaction #1).

These sequences, with the configuration detailed above, will cause the data for the write to memory to be corrupted.

IMPLICATION: With all of these features enabled, data corruption may occur, resulting in unpredictable system failure.

WORKAROUND: Use the default Page Open Policy (close page). Disabling RAW and ERAW or using an IOQ depth of 1 will also prevent this erratum, but may have a larger impact on performance.

18. *T_{CASE} Drop Plus Voltage Swing May Cause DPLL Failure*

PROBLEM: Analysis of the DPLL (Digital Phase Lock Loop) units of the Intel 450GX PCIset components has shown that a DPLL failure can occur during certain changes in temperature and/or voltage within the Intel 450GX PCIset component specification, resulting in a loss of DPLL functionality for between 2 and 500 bus clocks.

The current specification states that the Intel 450GX PCIset is operational over a T_{CASE} range of 0 - 85°C and a V_{CC} range of 3.3V ± 5%.

After a device has undergone a hard reset of the DPLL (via the PLLRST pin), a drop in T_{CASE} accompanied by an increase in V_{CC} may cause this erratum to occur. The magnitude of the temperature drop and supply voltage increase required to cause this failure is graphed below:

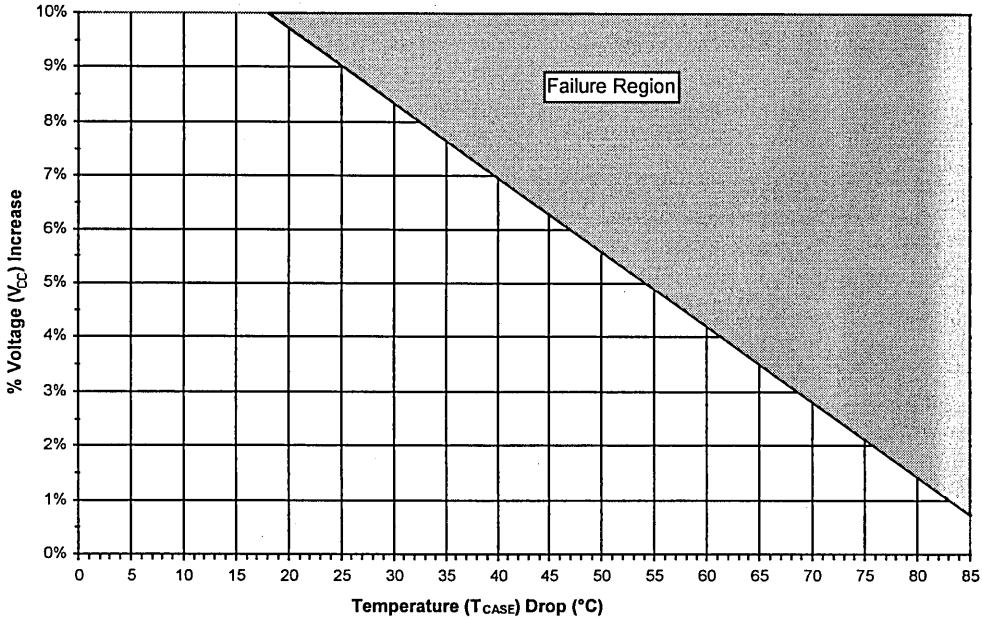

Figure 1. V_{CC} Increase vs. T_{CASE} Drop for DPLL Failure (Worst Case)

IMPLICATION: If the T_{CASE} of any Intel 450GX PCIset component is allowed to decrease by more than 18°C after the point at which DPLL lock is obtained after a DPLL hard reset, a positive voltage swing may cause the component to lose its internal clock for 2 - 500 processor bus clocks. This may result in a system hang, lost data, or other unpredictable system failure (note that if the component is not processing commands or data for the time during which the internal clock is lost, no failure will be observed). The larger the temperature drop and/or voltage swing, the more likely the failure.

If V_{CC} swings from 3.3V -5% to 3.3V +5%, and T_{CASE} does drop more than 18°C past the temperature at which the DPLL was reset, this erratum may occur. Similarly, if T_{CASE} transitions from 85°C (at DPLL reset time) to 0°C, a swing in V_{CC} of more than 0.5% may cause this erratum. Thus, maintaining a moderately constant T_{CASE} and V_{CC} will prevent this erratum, as per the graph above.

WORKAROUND: Ensure that the DPLL reset occurs during a period where the T_{CASE} is low and/or the V_{CC} is high relative to normal operational conditions (e.g. if the system is running at a high temperature and is powered down, allow it to cool before powering up again). Alternately, ensure that the system meets a restricted specification for changes in T_{CASE} and V_{CC}.

19. Processor Bus ECC Signals May Be Corrupted

PROBLEM: The processor bus ECC signals issued by the 82454GX PB may not contain correct ECC information, when processor bus ECC is enabled, the IOQ depth is 8, and inbound posting (PCI bus to

processor bus) is enabled. This erratum occurs on all Intel 450GX components; Intel has observed it more frequently at low voltage and high temperature.

IMPLICATION: With an IOQ depth of 8 and inbound posting enabled, processor bus ECC will not be reliable.

WORKAROUND: Any one of the following three workarounds may be implemented to prevent this erratum:

1. Disable processor bus ECC in all processor bus agents. Memory bus ECC is not affected and can remain enabled. To disable processor bus ECC and keep memory bus ECC enabled, use the following processor bus ECC settings:
    - Data Error Checking Enable:
      Pentium Pro processor, EBL_CR_POWERON MSR, bit 1 = 0 (disabled).
    - Logging Uncorrectable Errors on the Host Data Bus Enable:
      82453GX MC, register offset C4-C5h, bit 7 = 0 (disabled).
    - Logging Correctable Errors on the Host Data Bus Enable:
      82453GX MC, register offset C4-C5h, bit 8 = 1 (enabled).
    - Single-bit Error Correcting of Host Data Bus Enable:
      82453GX MC, register offset C4-C5h, bit 9 = 0 (disabled).
    - Single-bit ECC Error Correcting of Host Data Bus Enable:
      82454GX PB, register offset C0h, bit 10 = 0 (disabled).
    - Report Uncorrectable Host Data Bus ECC Errors:
      82454GX PB, register offset C0h, bit 11 = 0 (disabled).
2. Use an IOQ depth of 1.
3. Disable inbound write posting in the 82454GX PB.

20. *IO_REQ# May Not Be Deasserted for BPRI# During GAT, Non-GAT, or APIC Flush Mode Request*

PROBLEM: During a GAT, non-GAT, or APIC Flush mode sideband request after the reception of inbound posted writes, the 82454GX PB goes into a state where all inbound posted writes are then drained. IO_REQ# is kept asserted during this process. However, if a retry occurs, IO_REQ# should be deasserted soon after BPRI# is asserted, instead of remaining asserted until all inbound posted writes are drained. If inbound writes target a cluster bridge and cannot make progress (due to retries from the cluster bridge), a deadlock occurs. After MEMACK# assertion for a sideband request, if a posted write is retried, the 82454GX PB does not rearbitrate for BPRI# before reissuing the write, and contention on the BPRI# signal may result. It is also possible for traffic originating behind the compatibility 82454GX PB and targeting the non-compatibility 82454GX PB to encounter this boundary condition. However, Intel has not currently identified any software which generates such traffic.

IMPLICATION: Handling of all GAT mode (FLUSHBUF# and MEMREQ# asserted), non-GAT mode (only FLUSHBUF# asserted), and APIC Flush (only MEMREQ# asserted) sideband requests are affected, whether from a EISA master/south bridge or other logic on the board (including the logic workaround for 450GX Erratum #7 or for GAT mode transaction aliasing).

WORKAROUND: Ensure that drivers which support peer-to-peer PCI transactions only do so for PCI devices under the same 82454GX PB. No workaround has been identified for third-party clustering agents.

21. Memory Gap Reclaiming May Corrupt Data

PROBLEM: If the memory subsystem is configured using the Memory Gap Register (MG), Low Memory Gap Register (LMG), or High Memory Gap Start Address and End Address registers (HMGSA, HMGEA) to allow a gap in the address map of the 82453GX MC, and if the memory in these gaps is reclaimed (by setting the appropriate bit in the MG, LMG, or HMGSA registers), the row address signals may be corrupted during the row address strobe (RAS#) signal assertion.

IMPLICATION: An incorrect address may be generated for a memory access if any of these three memory gaps are enabled with the memory reclaimed. This would result in data corruption and unpredictable system failure.

WORKAROUND: If these memory gaps are used, do not reclaim the memory in the gaps.

22. RAW May Corrupt Write Data in 1:1 Interleaving

PROBLEM: If a memory subsystem is configured for 1:1 interleaving (i.e. only using one memory interleave, also known as non-interleaved), and read-around-write (RAW) is enabled, a partial write transaction to a memory location marked as Modified in the L2 cache followed by a line read transaction which is re-ordered around the partial write may result in the write data being corrupted.

IMPLICATION: With RAW enabled and the memory subsystem 1:1 interleaved, data corruption may occur, resulting in unpredictable system failure.

WORKAROUND: Do not enable RAW if the system is configured in 1:1 interleaving.

23. Inbound Write May be Posted Too Soon After Misaligned or Multi-Dword Cycle

PROBLEM: Inbound writes should not be posted by the 82454GX PB between the time an outbound memory read, I/O write, or I/O read cycle is taken from the request queue until it completes on the PCI bus. If such a cycle required multiple transfers on the PCI bus (i.e. is misaligned across a Dword boundary or is a multi-Dword request) and the last transfer of the request receives a retry response from the PCI bus, then a write may be posted before the cycle completes.

IMPLICATION: Misaligned or multi-Dword memory reads, I/O reads, and I/O writes may cause a subsequent inbound write transaction to be posted improperly, resulting in an ordering violation and system hang.

WORKAROUND: Ensure that:

1. The Lock Atomic Reads feature in the 82454GX PB's PB Configuration Register (bit 6 of register offset 4Ch) is set, and that the arbiter uses full PCI bus locks. This will prevent another PCI master from acquiring the PCI bus after a retry on the last transfer of a misaligned or multi-Dword memory or I/O read.
2. Ensure that misaligned I/O writes do not occur. Intel has not currently identified any commercial operating system or application software which contains misaligned I/O writes.

24. IO_REQ# May Not Be Deasserted for BPRI# During GAT Mode Request

PROBLEM: During a GAT mode sideband request after the reception of inbound posted writes, the 82454GX PB goes into a state where all inbound posted writes are then drained. IO_REQ# is kept asserted during this process. However, if a retry occurs, IO_REQ# should be deasserted soon after BPRI# is asserted, instead of

remaining asserted until all inbound posted writes are drained. If inbound writes target a cluster bridge and cannot make progress (due to retries from the cluster bridge), a deadlock occurs. After MEMACK# assertion for a GAT mode sideband request, if a posted write is retried, the 82454GX PB does not rearbitrate for BPRI# before reissuing the write, and contention on the BPRI# signal may result. It is also possible for traffic originating behind the compatibility 82454GX PB and targeting the non-compatibility 82454GX PB to encounter this boundary condition. However, Intel has not currently identified any software which generates such traffic.

IMPLICATION: Handling of all GAT mode (FLUSHBUF# and MEMREQ# asserted) sideband requests are affected. Such requests may result in a deadlock if peer-to-peer transactions or third-party clustering agents are used.

WORKAROUND: Ensure that drivers which support peer-to-peer PCI transactions only do so for PCI devices under the same 82454GX PB. Systems with third-party clustering agents must use Aliased GAT mode to avoid this erratum.

25. *BINIT# Assertion May Cause Active RAS# Negation*

PROBLEM: BINIT# is asserted on the Pentium Pro processor system bus to indicate the occurrence of a catastrophic system error condition or a situation that prevents reliable future operation. When the 82453GX MC component of the 450GX PCIset recognizes the assertion of BINIT# on the system bus, it prepares to clear certain internal state. In the process, a currently active RAS# signal associated with a read, write or refresh operation in progress may also be negated. If this negation causes the minimum RAS# pulse width timing to be violated then spurious bits may appear in one page of the memory array.

IMPLICATION: The catastrophic system error conditions that led to BINIT# assertion may in turn lead to a page of memory being incompletely written or improperly refreshed before potentially recovering from BINIT#.

WORKAROUND: Recovery from BINIT# assertion may be aided by a Machine Check Exception (MCE) or System Management Interrupt (SMI) handler. MCE or SMI handlers resident in non-volatile memory can protect recovery routines from encountering this situation, allowing them to be reliably executed after BINIT# assertion.

SPECIFICATION CLARIFICATIONS

1. Explicit Writebacks Claimed by 82454GX PB

Note 2 in Section 3.2 of the *Intel 450KX/GX PCIset Data Sheet* states that writebacks initiated by other agents are ignored by the PB. It should be noted that while this is true with respect to PCI bus transactions (i.e. no PCI bus cycles will be generated due to any writeback transaction), if a writeback occurs to memory behind the 82454GX PB, the data will be lost or a violation of processor bus protocol will occur. The 82454GX PB is not a caching agent, and no writeback transactions should be targeted to devices on the PCI bus. Memory behind the 82454GX PB may not be mapped as cacheable (WB type) memory.

2. Supported Configurations For MC Row Limit Register Programming

The following text will be added to the paragraph describing the DRL DRAM Row Limit Register in Section 2.3.13 of *Chapter 3 Memory Controller* of the Intel 450GX PCIset databook.

"The DRAM Row Limit registers may only be programmed in such a fashion that they do not violate the supported memory configurations outlined in Table 22, *Minimum and Maximum Memory Sizes for Each Configuration.* Other configurations of the row limit registers using increments other than those described are not supported."

DOCUMENTATION CHANGES

1. *Register Offset and Default Value Correction*

Some configuration registers are documented unclearly or inconsistently in the *Intel 450KX/GX PCIset* datasheet. A table of the correct offsets and values is given below with the changes in bold (note that this list only contains registers with changes):

Configuration Register	Address Offset	Default Value	Notes
82454GX			
Top of System Memory	40-43h	**0000** 0000h	
Bridge Device Number	49h	0001 1001**b** 0001 1010**b**	Compatibility 82454GX Auxiliary 82454GX
PB Configuration	4Ch	**1**9h **1**Ah	Compatibility 82454GX Auxiliary 82454GX
PCI Read/Write Control	54-55h	**0**000h	
Memory Gap Range	78-79h	**0**000h	
Memory Gap Upper Address	7A-7Bh	**0**000h	
PCI Frame Buffer	7C-7Fh	**0000 0**000h	
High Memory Gap Range Start Address	88-8Bh	**0000 0**000h	
High Memory Gap End Address	8C-8Fh	**0000 0**000h	
Configuration Values Driven on Reset	B0-B1h	**0**000h	Bit 7: 1 = Depth of **1**. 0 = Depth of **8**. Pentium Pro processors use an in-order depth of **1** if this bit is 1.
Captured System Configuration Values	B4-B5h	**000X XXXX XXX0 0000b**	X = captured during hard reset. Bit 7: 1 = Depth of **1**. 0 = Depth of **8**. Pentium Pro processors use an in-order depth of **1** if this bit is 1.
SMRAM Range	B8-BBh	0000 0005h	Bits [15:0] correspond to A[31:**16**]#. The default starting address is 50000h and ranges to **5**FFFFh.
PB Retry Timers	C8-**CB**h	0000 0003h	
82453GX			
Controller Device Number	49h	0001010X**b**	X = loaded at reset
Single Bit Correctable Error Address	74-77h	**0000** 0000h	
Low Memory Gap Register	7C-7Fh	0010 0000h	**Bits [9:5]: Reserved.** Bits [4:0]: **Low Memory Gap Size.** This field defines the memory gap

Configuration Register	Address Offset	Default Value	Notes
			size as follows: **Bits [4:0] Size** **Bits [4:0] Size** 00000 1 MB **00111** 8 MB **00001** 2 MB **01111** 16 MB **00011** 4 MB **11111** 32 MB
High Memory Gap Start Address	88-8Bh	**0000** 0000h	**Bit 30: Reclaim Enable.** 1 = Enable. 0 = Disable (default). When enabled, the physical memory in this gap is reclaimed.
High Memory Gap End Address	8C-8Fh	**0000** 0000h	
Memory Timing Register	AC-AFh	30DF3516h	Bit 15: 1 = **2** Cycles. 0 = **1** Cycle (default).
SMRAM Range	B8-BBh	**0000** 000Ah	

2. *CMOS Definition Should be 3.3V or 5V*

In Section 1.0 of both Chapter 2 and Chapter 3 of the *Intel 450KX/GX PCIset* datasheet, CMOS signals are defined as follows:

CMOS Rail-to-rail CMOS tolerant to 5V levels.

This should read:

CMOS Rail-to-rail CMOS tolerant to 3.3V or 5V levels. See Chapter 4, Section 1.2., *Signal Groups*.

UNITED STATES, Intel Corporation
2200 Mission College Blvd., P.O. Box 58119, Santa Clara, CA 95052-8119
Tel: +1 408 765-8080

JAPAN, Intel Japan K.K.
5-6 Tokodai, Tsukuba-shi, Ibaraki-ken 300-26
Tel: + 81-2[illegible]47-8522

FRANCE, Intel Corporation S.A.R.L.
1, Quai de Grenelle, 75015 Paris
Tel: +33 1-45717171

UNITED KINGDOM, Intel Corporation (U.K.) Ltd.
Pipers Way, Swindon, Wiltshire, England SN3 1RJ
Tel: +44 1-793-641440

GERMANY, Intel GmbH
Dornacher Strasse 1
85622 Feldkirchen/ Muenchen
Tel: +49 89/99143-0

HONG KONG, Intel Semiconductor Ltd.
32/F Two Pacific Place, 88 Queensway, Central
Tel: +852 2844-4555

CANADA, Intel Semiconductor of Canada, Ltd.
190 Attwell Drive, Suite 500
Rexdale, Ontario M9W 6H8
Tel: +416 675-2438

Printed in USA/0197/NCG/PG