



Intel[®] 440FX PCIset 82441FX (PMC) and 82442FX (DBX)

Specification Update

January 2001

Notice: The Intel[®] 440FX PCIset may contain design defects or errors known as errata which may cause the product to deviate from published specifications. Current characterized errata are documented in this Specification Update.

Document Number: 297654-006

Information in this document is provided in connection with Intel® products.

No license, express or implied, by estoppel or otherwise, to any intellectual property rights is granted by this document.

Except as provided in Intel's Terms and Conditions of Sale for such products, Intel assumes no liability whatsoever, and Intel disclaims any express or implied warranty, relating to sale and/or use of Intel products including liability or warranties relating to fitness for a particular purpose, merchantability, or infringement of any patent, copyright or other intellectual property right.

Intel products are not intended for use in medical, life saving, or life sustaining applications.

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

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

The Intel® 440FX PCIset may contain design defects or errors known as errata which may cause the product to deviate from published specifications. Current characterized errata are available on request.

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

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

Intel Corporation

www.intel.com

or call 1-800-548-4725

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

Copyright © 1996–2001, Intel Corporation

Contents

Revision History.....	4
Preface.....	5
82441FX (PMC) Specification Changes.....	9
Errata.....	11
Specification Clarifications	15
Documentation Changes.....	17
82442FX (DBX) Specification Changes	19

Revision History

Rev.	Draft/Changes	Date
-001	This document is the first revision of the Intel 440FX PCIset Specification Update.	May 15, 1996
-002	Added Errata 1-2 and Specification Changes 1-3 for the 82441FX and Specification Change 1 for the 82442FX.	Aug. 14, 1996
-003	Added Errata 3-4 and Specification Clarification 2 and Documentation Change 2 for the 82441FX.	March 18, 1997
-004	Added Documentation Changes #3 and #4 for the 82441FX. In 82441 (PMC) Errata 4 And 82441 (PMC) Specification Clarification 2 Changed all occurrences of "Pentium® Pro" to "Pentium® Pro Processor"	May 6, 1997
-005	Added Specification Clarification #3 for the 82441FX (PMC)	April, 2000
-006	Editorial and formatting changes.	January, 2001

Preface

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

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

Affected Documents/Related Documents

Document Title	Document Number
Intel® 440FX 82441FX (PMC) and 82442FX (DBX) PCIset datasheet	290549-001

Nomenclature

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

Errata are design defects or errors. Errata may cause the Intel® 440FX PCIset 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.

Component Identification via Programming Interface

Component Identification via Programming Interface (PMC Only)

PCI Register	PCI Offset	Value
Vendor ID	00-01h	8086h
Device ID	02-03h	1237h
Revision Number	08h	02h (A-1 stepping)

General Information

Basic Intel 440FX PCIset Identification Information

Product No.	Vendor ID	Device ID	Rev. ID	Product Stepping	Freq.	S-Spec	V _{CC}	TCASE	Notes
SB82441FX	8086h	1237h	2	A1	66.7 MHz	S U053	3.3 V ± 5%	0°C - 85°C	
SB82442FX	n/a ¹	n/a ¹	n/a ¹	A1	66.7 MHz	S U054	3.3 V ± 5%	0°C - 85°C	1

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.

Summary Table of Changes

The following table indicates the Specification Changes, Errata, Specification Clarifications, or Documentation Changes which apply to the Intel® 440FX PCIset. Intel intends to fix some of the errata in a future stepping 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 that applies to a stepping or to this product line.
Doc:	Document change or update that will be implemented.
Fix:	This erratum is intended to be fixed in a future stepping of the component.
Fixed:	This erratum has been previously fixed.
NoFix	There are no plans to fix this erratum.
(No mark) or (Blank Box):	This erratum is fixed in listed stepping or specification change does not apply to listed stepping.
Shaded:	This item is either new or modified from the previous version of the document.

82441FX (PMC)

NO.	A1	PIIX3	EISA	SPECIFICATION CHANGES
1	X	X		Long ISA Master Latency for Read Cycles
2	X		X	Long EISA Master Latency
3	X	X	X	Repeated Retry State on PCI Bus with PCI 2.0 Compliant Devices

NO.	A1	PIIX3	EISA	PLANS	ERRATA
1	X		X	NoFix	Temporary Retry State on the PCI Bus for PCI 2.0 Devices
2	X	X	X	NoFix	RAS Precharge Timing When One Row of 70 ns DRAM Is Installed
3	X	X	X	NoFix	PCI Master Memory Operation with Non-Linear Addressing Mode Writes to Incorrect Address
4	X			NoFix	Code Prefetch following Configuration Write Causes Lockup

NO.	A1	PIIX3	EISA	SPECIFICATION CLARIFICATIONS
1	X			Reading and Writing Configuration Registers
2	X	X		WC Posting in PMC Must Be Enabled with Passive Release Enabled in PIIX3
3	X			Multi-Bit Error Memory Clarification

NO.	A1	PIIX3	EISA	DOCUMENTATION CHANGES
1	X			Special Cycle Enable (SCE) Bit in the PCI Command Register Implemented
2	X			RAS# to CAS# Delay Should Always Be Set to '1' for Proper Operation
3	X			PMC Configuration Register: DBC-DBX Buffer Control Register, Section 3.2.13
4	X			PMC Configuration Register: DRAM Timing Register, Section 3.2.17

82442FX (DBX)

NO.	A1	PIIX3	EISA	SPECIFICATION CHANGES
1	X	X	X	DBX Internal Pulldown Resistors Do Not Pulldown on PD[5:2] Signals

NO.	A1	PIIX3	EISA	PLANS	ERRATA
					There are no errata for the 82442FX (DBX) in this Specification Update revision

NO.	A1	PIIX3	EISA	SPECIFICATION CLARIFICATIONS
				There are no specification clarifications for the 82442FX (DBX) in this Specification Update revision

NO.	A1	PIIX3	EISA	DOCUMENTATION CHANGES
				There are no documentation changes for the 82442FX (DBX) in this Specification Update revision

82441FX (PMC) Specification Changes

1. Long ISA Master Latency for Read Cycles

ISA LAN cards with legacy data transfer mechanisms can experience long DMA read latencies. The Intel 440FX PCIset has been designed to provide higher overall system concurrency by optimizing for processor and PCI transfers. ISA LAN cards that use a small amount of buffering, or PIO mode rather than bus mastering, can experience long latencies in demanding environments where there are multiple PCI devices.

Network interface cards have a minimum latency requirement once the transmission has commenced. If there is not enough buffering on board, or if the programmed I/O mode of transfer is used, these minimum latencies cannot be met. For server designs using ISA LAN cards, using these types of data transfer mechanisms may cause high numbers of CRC underrun errors, leading to eventual client dropoffs.

To allow higher system concurrency and performance, it is necessary to use high throughput peripherals. PCI LAN cards provide higher throughput and reduced latencies by using bus mastering. These devices transfer data to and from memory with minimal processor intervention, allowing the processor to work on other tasks. These devices increase the throughput by supporting 32-bit transfers, as well as optimized PCI commands (such as Memory Read Line, Memory Read Multiple, and Memory Write and Invalidate). Additionally, these devices tend to have more buffering to support the advanced PCI commands.

2. Long EISA Master Latency

The Intel 440FX PCIset has been designed to operate in non-GAT mode. Additionally, the Intel 440FX PCIset is optimized for processor and PCI performance. In systems that have high PCI traffic, EISA master devices can experience long latencies. These latencies result in the PCI-EISA bridge generating an NMI after the NMI timer times out.

Intel recommends turning off the NMI timer (configuration offset 0461h set to 00h). If reducing the long EISA latency is desirable, setting the IOQ depth to 1 can decrease processor traffic. Setting the IOQ depth to 1 requires putting a pullup resistor on PC[0]. The EISA line buffers in the PCI-EISA bridge should also be enabled to reduce EISA latency.

3. Repeated Retry State on the PCI Bus with PCI 2.0 Compliant Devices

This Specification Change affects both ISA and EISA systems. For PCI 2.0 compliant devices, the PCI bus can enter a repeated retry state in which PMC and the device may retry each other incessantly. This was specifically observed on a PCI-to-PCI bridge. If a write cycle from the processor to PCI occurs coincidentally with a PCI to memory transaction, then the PMC will behave as follows:

- If the PCI to memory transaction is a read, then it will be taken as a delayed transaction. Any subsequent PCI to memory read cycles will be retried until the processor-to-PCI write cycle has been completed.
- If the PCI to memory transaction is a write, then it will be completed.

The system may remain in a repeated retry state until the processor-to-PCI write has been completed. For devices that are compliant with the PCI 2.1 Local Bus Specification, the processor-to-PCI write will be completed and operation will continue without the repeated retry state. For PCI devices that are not PCI 2.1 compliant and do not accept processor-to-PCI writes during a pending outbound cycle, this situation could result in a repeated retry state.

This issue has been specifically seen on PCI to PCI bridges that are PCI 2.0 Local Bus Specification compliant. To resolve this issue, program the PCI-to-PCI bridge such that primary PCI bus traffic has precedence over secondary PCI bus traffic. This will assure that the processor-to-PCI write is completed before the PCI to memory read transaction is retried.

Errata

1. Temporary Retry State on the PCI Bus for PCI 2.0 Devices

Problem: This erratum can occur in EISA systems with at least two PCI masters. On EISA systems it is recommended that delayed transactions be turned off, if the PCI-EISA bridge is PCI 2.0 Local Bus Specification compliant. With delayed transactions disabled (as recommended), if one PCI master's read occurs coincidentally with a write cycle from the processor to PCI, then any subsequent writes from the second PCI master will be retried until the processor to PCI write cycle has been completed. This problem can also occur in ISA systems under the following conditions:

- Delayed transactions are disabled (this mode of operation is not supported for ISA systems).
- The PMC Discard Timer expires - Assume a PCI master's read occurs coincidentally with a write cycle from the processor-to-PCI and the PCI master does not return with the delayed transaction. The processor-to-PCI write cycle will be taken as a deferred transaction. Upon the expiration of the discard timer for the pending PCI read transaction, any subsequent write from a second PCI master will be retried until the processor-to-PCI write cycle has been completed. During the time that the discard timer is alive, the PMC will operate normally.

The events occur as described above, with the exception that the PCI read cycle is a locked transaction.

Implication: The system may remain in a temporary retry state until the processor-to-PCI write has been completed. For devices that are compliant with the PCI 2.1 Local Bus Specification, the processor-to-PCI write will be completed and operation will continue without the temporary retry state. For PCI devices that are not PCI 2.1 Local Bus Specification compliant and do not accept processor-to-PCI writes during a pending outbound cycle, it could result in a livelock.

Workaround: Intel has observed this issue only on PCI 2.0 Local Bus Specification compliant PCI to PCI bridges. This issue can be fixed by ensuring that the I/O write from the PMC is prioritized over the secondary PCI bus transaction. Most PCI-to-PCI bridges support this. Contact your PCI-to-PCI bridge vendor for exact register programming. Using devices which are PCI 2.1 compliant would also work around this erratum.

Status: There are no plans to fix this erratum.

2. RAS# Precharge Timing When One Row of 70-ns DRAM Is Installed

Problem: 70 ns DRAMs require 50 ns of RAS# precharge. The RPE bit in the PMC (configuration offset 54h, bit 7) controls whether there are 3 or 4 clocks of RAS# precharge (3 clks = 45 ns, 4 clks = 60 ns at 66 MHz). However, if there is only one row of memory populated, the PMC will only generate 3 clocks of RAS# precharge during a refresh cycle.

Implication: Page miss cycles work correctly and have 4 clocks of RAS# precharge when RPE = 1; only refresh cycles are affected. Since 70 ns DRAMs have a specification of 50 ns for RAS# precharge, there is a potential for failure if only one row of 70 ns DRAMs are populated in the system.

Workaround: To avoid the possibility of failure, the BIOS can program additional DRT registers (configuration offset 55h-56h) from empty row to a non-empty (i.e., FPM type or EDO type). Note that the DRB registers (configuration offset 60h-67h) must still be programmed to show there is no memory in the row. The PMC only uses the DRT registers when generating refresh cycles, so by programming one of the DRT registers to non-empty, the PMC will generate refresh cycles for 2 rows and avoid the potential RAS# precharge spec violation.

Status: There are no plans to fix this erratum.

3. PCI Master Memory Operation with Non-Linear Addressing Mode Writes to Incorrect Address

Problem: Data from the memory write operations are written to the wrong memory address when ALL the following conditions are met:

- The PCI master device issues a non-linear addressing mode memory operation
- The non-linear addressing mode memory operation is not retried immediately by the PMC
- The previous PCI bus operation was a “memory write” or “memory write invalidate”
- The internal data buffer of the PMC is between 70% and 100% full because other PCI masters have been moving large amount of data into memory.

This will result in a system lock-up when a subsequent processor memory read or a PCI memory read is initiated,

Implication: Systems having PCI master devices that utilize non-linear addressing mode will lock-up when all of the above conditions occur.

Workaround: If using a PCI master device that supports non-linear and linear addressing modes, direct the PCI master to only use the linear addressing mode during memory operations. Or, only use PCI master devices that do not support the non-linear addressing mode. An external workaround, for systems using a PCI master device that supports non-linear addressing, is to drive zeros onto PCI AD[1:0] lines during the address phase of non-linear addressing mode memory operations which is initiated by the PCI master. This can be done with an inexpensive PLD and an inexpensive TTL switch buffer.

Status: There are no plans to fix this erratum.

4. Code Prefetch following Configuration Write Causes Lockup

Problem: A system lockup can occur if a PMC configuration write cycle occurs coincidentally with the assertion of a high priority refresh request and a pipelined memory page hit request. This is an issue for Intel® Pentium® Pro processor, B1 stepping only.

Implication: With the refresh queue disabled, all refresh cycles are considered high priority. When a configuration write is generated, normal operation disables processing of any further memory transactions including high priority refresh. High priority refresh also precludes any further memory cycle processing. In the worst case, it is possible that both of these requests go active within a few clocks of each other. If these occur coincident with an outstanding memory read page hit, the conflict of multiple priority requests results in a system lockup. The worst case scenario is a configuration write followed by full IOQ back to back memory code fetches.

Workaround: Set the DRAM Control Register (address offset 57h), bit 6 to 1, which insures that all refresh requests are queued, and set the PMC Configuration Register (address offset 50-51h) bit 2 to 0, which prevents any further memory read traffic being pipelined behind the configuration write.

Status: There are no plans to fix this erratum..

This page is intentionally left blank.

Specification Clarifications

1. Reading and Writing Configuration Registers

When changing values in configuration registers that contain one or more reserved bits, care must be taken to avoid changing the value of the reserved bits. It is not safe to assume that because a bit is reserved that it has no function or meaning. The proper procedure for changing a register value is to first read the register, change only the required bits (without changing the value of any reserved bits), then write the new value back out to the register.

2. WC Posting in PMC Must Be Enabled with Passive Release Enabled in PIIX3

Optimization can be obtained during concurrent PCI-DRAM and processor-Frame buffer traffic by setting chipset registers as described below:

- The USWC Write Posting during I/O bridge access enable (UWPIO) bit is set to 1. This bit is located in DBC (DBX Buffer Control) register (bit #5) in PMC at address offset 53h.
- The Passive Release Enable bit is set to 1. This bit is located in DLC (Deterministic Latency Control) register (bit #1) in PIIX3 at address offset 82h.

Care must be taken to ensure that both bits are set to avoid deadlock condition.

Buffer management rules require that prior to the assertion of PHOLDA to the PIIX3, processor-to-PCI posting buffers are flushed and disabled. This is required due to the shared DRAM read and processor write paths. If processor posts a cycle targeted to PCI ahead of the DRAM read then a deadlock would occur. The deadlock is: PMC retries the PIIX3 cycle and PIIX3 continues to retry the same cycle. PIIX3 implements passive release and alleviates the issue by de-asserting PHOLD for one clock in which PMC would de-assert PHLDA and get back the control of the PCI bus. PMC would then be allowed to run its PCI cycle and effectively flush the processor-to-PCI (say WC cycle).

Simply stated, WC enabled graphics must be located on PCI and passive release must be enabled if WC posting under PHLDA is enabled.

To allow for high speed write capability for graphics, Intel® Pentium® Pro processor has introduced WC memory type. The WC memory type provides a write-combining buffering mechanism for write operations within the Intel Pentium Pro processor. A high percentage of graphics transactions are writes to the graphics region, normally known as linear frame buffer. In case of graphics, current 32-bit drivers (without modifications) would use Partial Write protocol to update the frame buffer. The highest performance write transaction on the Intel Pentium Pro processor bus is line write. By combining the several back-to-back partial write transactions (internal to the Intel Pentium Pro processor) into a line write transaction on the Intel Pentium Pro processor bus, the performance of the frame buffer accesses is greatly improved. To extend this capability to the current drivers, it is necessary to set up the linear frame buffer address range to be WC memory type. This can be done by programming the MTRR registers in the Intel Pentium Pro processor. For more details on MTRR programming, refer to Section 11.3 in the *Intel® Pentium® Pro Processor Family Developers Guide*.

3. **Multi-Bit Error Memory Clarification**

When an Intel® 440FX PCIset platform is configured for ECC support, if a multi-bit uncorrectable memory error is detected during a memory read by a system device, an SERR, SCI, or SMI will be generated. This typically results in an NMI. However, bad data may still reach the intended target before the NMI can be generated or before NMI interrupt handler can service the problem. This may result in bad data being returned to the target and may be permanently stored, resulting in system data corruption. This chipset was not architected or designed to ensure that targets are protected from corrupted data in this situation.

Documentation Changes

1. Special Cycle Enable (SCE) Bit in the PCI Command Register Implemented

The PCEB B-0 step implemented the Special Cycle enable bit (Bit 3) in the PCI Command (04-05h). This bit must be set to a 1 for the PCEB to recognize the STOP_GNT special cycle on the PCI bus and to assert STPGNT# to the ESC for one PCI clock. The STPCLK# power management feature in the 82374SB will not function without this bit set in the 82375SB.

2. RAS# to CAS# Delay Should Always Be Set to '1' for Proper Operation

The default value of this bit is '0', which allows for the fastest leadoff timing. With this bit set to '0', the current 50 ns and 60 ns DRAMs will have negative timing margins and will result in problems for sizing and detecting DRAMs. RAS# to CAS# Delay, RCD, should be always set to 1, which allows for better tRAC and tCSH timing margin.

3. PMC Configuration Register: DBC-DBX Buffer Control, Section 3.2.13

The DBC-DBX Buffer Control register (address offset 53h), bit 4, should read: PCI Delayed Transaction Timer Disable (DTD) 1=Disable. 0=Enable (default).

4. PMC Configuration Register: DRAM Timing Register, Section 3.2.17

The DRAM timing register (address offset 58h), bit 1, should read: RASx# to CASx# Delay (RCD). 1=Three clocks between the assertion of RASx# and CASx#. 0=Two clocks (default).

This page is intentionally left blank.

82442FX (DBX) Specification Changes

1. DBX Internal Pulldown Resistors Do Not Pulldown on the PD[5:2] Signals

The DBX has integrated pulldown resistors on the PD[5:2] pins. The DBX samples the PD[5:2] pins at the rising edge of the CRESET# signal, and will use the sampled values to determine which mode of operation to enter according to the table below:

PD5	PD4	PD3	PD2	Mode Selected
0	0	0	0	Normal Mode (default)
0	0	1	0	NAND Tree
0	0	1	1	Drive all 1s
0	1	0	0	Drive all 0s
1	0	0	1	Tristate
				All other combinations reserved

If any of the PD[5:2] pins are sampled high at the rising edge of CRESET#, then the DBX will enter a test mode and will not transfer data properly. The internal pulldown resistors in the DBX are not enabled as they should be, so the PD[5:2] signals are left floating during reset.

The system may not reset properly (the system locks up) when a reset occurs and one of the PD[5:2] signals is sampled high. Add 20 KΩ external pulldown resistors to the PD[5:2] signals.

Errata

There are no errata for the 82442FX (DBX) in this Specification Update revision.

Specification Clarifications

There are no specification clarifications for the 82442FX (DBX) in this Specification Update revision.

Documentation Changes

There are no documentation changes for the 82442FX (DBX) in this Specification Update revision..



This page is intentionally left blank.