

Digital Semiconductor 21172 Core Logic Chipset

Technical Reference Manual

Order Number: EC-QUQJA-TE

Revision/Update Information: This is a draft version.

April 1996

While Digital believes the information in this publication is correct as of the date of publication, it is subject to change without notice.

Digital Equipment Corporation makes no representations that the use of its products in the manner described in this publication will not infringe on existing or future patent rights, nor do the descriptions contained in this publication imply granting of licenses to make, use, or sell equipment or software in accordance with the description.

© Digital Equipment Corporation 1996. All rights reserved.
Printed in U.S.A.

DEC, DECchip, Digital, Digital Semiconductor, OpenVMS, VAX DOCUMENT, and the DIGITAL logo are trademarks of Digital Equipment Corporation.

Digital Semiconductor is a Digital Equipment Corporation business.

Intel is a trademark of Intel Corporation.
Windows NT is a trademark of Microsoft Corporation.

All other trademarks and registered trademarks are the property of their respective owners.

This document was prepared using VAX DOCUMENT Version 2.1.

Contents

Preface	xiii
1 21172 Core Logic Chipset Overview	
1.1 Features	1-1
1.2 System Overview	1-2
1.3 21172 Core Logic Chipset Microarchitecture	1-4
1.3.1 CIA Microarchitecture	1-4
1.3.2 DSW Microarchitecture	1-6
1.4 System Clocks	1-7
1.5 21172 Core Logic Chipset Characteristics	1-8
Part I 21172-CA (CIA) Information	
2 21172-CA Pin Descriptions	
2.1 21172-CA Pin List	2-1
2.2 21172-CA Signal Descriptions	2-4
2.2.1 System Bus Signals	2-4
2.2.2 PCI Local Bus Signals	2-4
2.2.3 IOD Bus Signals	2-8
2.2.4 Memory Control Signals	2-11
2.2.4.1 Set Select/Alternate Memory Modes	2-11
2.2.4.2 Memory Control Signal Definitions	2-12
2.2.5 Phase-Locked Loop Signals	2-14
2.2.6 Miscellaneous Power Signals	2-15
2.2.7 Miscellaneous Signals	2-15
2.2.8 Reserved Signals	2-16
2.2.9 21172-CA Pin Name List (Alphabetic)	2-17
2.2.10 21172-CA Pin Assignment List (Alphanumeric)	2-23
2.3 21172-CA Mechanical Specifications	2-30

3 21172–CA Architecture Overview

3.1	CIA Data Paths with Functional Units Description	3–1
3.1.1	System Functional Units	3–1
3.1.2	CIA Functional Units	3–3
3.1.2.1	21164 Instruction and Address Logic	3–4
3.1.2.2	PCI Data Path Logic	3–4
3.1.2.3	Memory Logic	3–6
3.1.2.4	I/O Address Logic	3–6
3.1.2.5	DMA Address Logic	3–7
3.2	CIA Transactions	3–7
3.2.1	21164 Read Miss Transaction	3–8
3.2.2	21164 Read Miss with Victim Transaction	3–8
3.2.3	21164 Read Transaction to Noncacheable Space	3–9
3.2.4	21164 Write Transaction to Noncacheable Space	3–10
3.2.5	DMA Transactions	3–11
3.2.6	DMA Read Transaction	3–13
3.2.7	DMA Read Transaction (Prefetch)	3–14
3.2.8	DMA Write Transaction	3–14
3.3	System Data Coherency Issues	3–15
3.3.1	Basic Properties of the System	3–16
3.3.2	Post and Run I/O Write Data Coherency	3–17
3.4	CIA Memory ras , cas , and Address Operation	3–17
3.4.1	Traditional Mapping of 21164 Address to Memory Address	3–17
3.4.2	CIA Mapping of 21164 Address to Memory Address	3–20
3.5	MB and LOCK Instructions	3–26
3.5.1	MB Instruction	3–26
3.5.2	LOCK Instruction	3–26
3.5.3	Locks to Uncached Space	3–28

4 21172–CA Control and Status Registers

4.1	21172–CA Registers	4–1
4.1.1	21172–CA General Registers	4–2
4.1.2	21172–CA Diagnostic Registers	4–2
4.1.3	21172–CA Performance Monitor Registers	4–3
4.1.4	21172–CA Error Registers	4–3
4.1.5	21172–CA System Configuration Registers	4–4
4.1.6	21172–CA PCI Address and Scatter-Gather Registers	4–5
4.1.7	21172–CA Address Translation Registers	4–6
4.2	21172–CA General Registers	4–8
4.2.1	CIA Revision Register	4–8

4.2.2	PCI Latency Register	4-9
4.2.3	CIA Control Register	4-11
4.2.4	CIA Configuration Register	4-15
4.2.5	Hardware Address Extension Register	4-16
4.2.6	Hardware Address Extension I/O Register	4-18
4.2.7	Configuration Register	4-19
4.2.8	CIA Acknowledgment Enable Register	4-20
4.3	21172-CA Diagnostic Registers	4-21
4.3.1	CIA Diagnostic Control Register	4-21
4.3.2	Diagnostic Check Register	4-23
4.4	21172-CA Performance Monitor Registers	4-24
4.4.1	Performance Monitor Register	4-24
4.4.2	Performance Control Register	4-25
4.5	21172-CA Error Registers	4-28
4.5.1	CPU Error Information Register 0	4-28
4.5.2	CPU Error Information Register 1	4-29
4.5.3	CIA Error Register	4-30
4.5.4	CIA Status Register	4-34
4.5.5	CIA Error Mask Register	4-36
4.5.6	CIA Syndrome Register	4-38
4.5.7	CIA Memory Port Status Register 0	4-40
4.5.8	CIA Memory Port Status Register 1	4-41
4.5.9	PCI Error Register 0	4-44
4.5.10	PCI Error Register 1	4-47
4.5.11	PCI Error Register 2	4-48
4.6	21172-CA System Configuration Registers	4-49
4.6.1	Memory Configuration Register	4-49
4.6.2	Memory Base Address Register <i>n</i>	4-52
4.6.3	Memory Timing Information Register <i>n</i>	4-55
4.7	21172-CA PCI Address and Scatter-Gather Registers	4-63
4.7.1	Scatter-Gather Translation Buffer Invalidate Register	4-63
4.7.2	Window Base Register	4-64
4.7.3	Window Mask Register <i>n</i>	4-66
4.7.4	Translated Base Register <i>n</i>	4-68
4.7.5	Window Base DAC Register	4-71
4.8	21172-CA Address Translation Registers	4-72
4.8.1	Lockable Translation Buffer Tag <i>n</i> Registers	4-72
4.8.2	Translation Buffer Tag <i>n</i> Register	4-74
4.8.3	Translation Buffer <i>m</i> Page <i>n</i> Registers	4-76

5 21172–CA Power-Up and Initialization

5.1	Power-Up	5–1
5.2	Internal Reset	5–1
5.3	State of Pins on Reset Assertion	5–2
5.4	Configuration After Reset Deassertion	5–3

6 21172–CA Address Mapping

6.1	Address Mapping Overview	6–1
6.2	21164 Address Space Configuration Supported by the CIA	6–2
6.2.1	21164 Access to Address Space	6–4
6.2.2	PCI Access to Address Space	6–4
6.3	21164 Address Space	6–8
6.3.1	PCI Dense Memory Space	6–9
6.3.2	PCI Sparse Memory Space	6–12
6.3.3	PCI Sparse I/O Space	6–17
6.3.4	PCI Configuration Space	6–21
6.3.4.1	Device Select (IDSEL)	6–24
6.3.4.2	PCI Special/Interrupt Acknowledge Cycles	6–26
6.3.4.3	Hardware-Specific and Miscellaneous Register Space	6–27
6.3.5	Byte/Word PCI Space	6–27
6.4	PCI to Physical Memory Addressing	6–30
6.4.1	Address Mapping Windows	6–30
6.4.1.1	PCI Device Address Space	6–32
6.4.1.2	Address Mapping Example	6–32
6.4.2	Direct Mapped Addressing	6–35
6.4.3	Scatter-Gather Addressing	6–36
6.4.3.1	Scatter-Gather Translation Lookaside Buffer (TLB)	6–38
6.4.3.1.1	Scatter-Gather TLB Hit Process	6–40
6.4.3.1.2	Scatter-Gather TLB Miss Process	6–40
6.4.4	Suggested Use of a PCI Window	6–42
6.4.4.1	Peripheral Component Architecture Compatibility Addressing and Holes	6–43
6.4.4.2	Memory Chip Select Signal mem_cs_1	6–43

Part II 21172–BA (DSW) Information

7 21172–BA Pin Descriptions

7.1	21172–BA Pin List	7–1
7.2	21172–BA Signal Descriptions	7–3
7.2.1	System Bus Signals	7–3
7.2.2	IOD Bus Signals	7–3
7.2.3	MEMDATA Bus Signals	7–6
7.2.4	Miscellaneous Signals	7–7
7.2.5	Phase-Locked Loop Signals	7–7
7.2.6	Miscellaneous Power Signals	7–8
7.2.7	Reserved Signals	7–8
7.2.8	21172–BA Pin Name List (Alphabetic)	7–10
7.2.9	21172–BA Pin Assignment List (Numeric)	7–13
7.3	21172–BA Mechanical Specifications	7–17

8 21172–BA Architecture Overview

8.1	DSW Functional Units	8–1
8.1.1	Victim Buffer—64 Bytes	8–3
8.1.2	I/O Read Buffer (IOR)—32 Bytes	8–3
8.1.3	I/O Write (IOW) Buffers—4 * 32 Bytes	8–4
8.1.4	DMA Buffer Sets	8–4
8.1.4.1	Flush Buffer	8–5
8.1.4.2	PCI Buffer	8–5

A Technical Support and Ordering Information

A.1	Obtaining Technical Support	A–1
A.2	Ordering Digital Semiconductor Products	A–1
A.3	Ordering Digital Semiconductor Literature	A–2
A.4	Ordering Associated Third-Party Literature	A–3

Index

Examples

3-1	21164 and PCI Device Lock Contention	3-27
-----	--	------

Figures

1-1	21172 Core Logic Chipset System Block Diagram	1-3
1-2	CIA Block Diagram	1-5
1-3	DSW Block Diagram	1-7
2-1	External PLL Filter Components	2-14
2-2	21172-CA 388-Pin EPBGA Package Dimensions	2-31
3-1	System Functional Block Diagram	3-2
3-2	CIA Functional Block Diagram	3-3
3-3	Address Space Overview	3-13
3-4	Traditional Mapping to a Memory Address	3-18
3-5	Traditional Victim Cache to Memory Correspondence	3-19
3-6	Modified Mapping to a Memory Address	3-20
3-7	CIA Victim Cache to Memory Correspondence	3-21
3-8	CIA Mapping to a Memory Address	3-22
3-9	Memory Address Maps for 128-Bit Data Path	3-24
3-10	Memory Address Maps for 256-Bit Data Path	3-25
4-1	Memory Read Timing Sample	4-58
4-2	Memory Write Timing Sample	4-59
4-3	CIA Signal Timing Limits	4-60
4-4	DSW Signal Timing Limits	4-62
6-1	21164 Address Space	6-2
6-2	21164 Address Space Configuration	6-3
6-3	Address Space Overview	6-4
6-4	Address Mapping Overview	6-5
6-5	21164 and DMA Read and Write Transactions	6-7
6-6	Dense Space Address Generation	6-11
6-7	PCI Memory Sparse Space Address Generation (Region 1)	6-15
6-8	PCI Memory Sparse Space Address Generation (Region 2)	6-16
6-9	PCI Memory Sparse Space Address Generation (Region 3)	6-17
6-10	PCI Sparse I/O Space Address Translation (Region A)	6-18

6-11	PCI Sparse I/O Space Address Translation (Region B)	6-19
6-12	PCI Configuration Space Definition	6-22
6-13	21164 System PCI Bus Hierarchy Example	6-26
6-14	Byte/Word PCI Space	6-28
6-15	PCI DMA Addressing Example	6-33
6-16	PCI Target Window Compare	6-34
6-17	Direct Mapped Translation	6-35
6-18	Scatter-Gather PTE Format	6-37
6-19	Scatter-Gather Associative TLB	6-39
6-20	Scatter-Gather Map Translation	6-41
6-21	Default PCI Window Allocation	6-42
6-22	Memory Chip Select Signal (mem_cs_1) Decode Area	6-45
6-23	Memory Chip Select Signal (mem_cs_1) Logic	6-46
7-1	21172-BA 208-Pin PQFP Package Dimensions	7-17
8-1	DSW Functional Block Diagram	8-2

Tables

1	Register Field Notation	xv
2	Unnamed Register Field Notation	xvi
3	Data Units	xvii
1-1	21172-CA (CIA) and 21172-BA (DSW) Chip Characteristics	1-8
2-1	21172-CA Pin List	2-1
2-2	ioc<6:0> Control Functions	2-8
2-3	cmc<8:0> Control Functions	2-9
2-4	Var Encodings for CMC Command 0	2-10
2-5	Var Encodings for CMC Commands 0, 8, A, and B	2-10
2-6	Var Encodings for CMC Commands 2, 3, 4, 5, 6, and 7	2-10
2-7	Alternate Memory Mode Memory Control Signal Redefinition	2-11
2-8	DSW Operating Modes	2-15
2-9	21172-CA Pin Names (Alphabetic)	2-17
2-10	21172-CA Pin Assignment (Alphanumeric)	2-23
2-11	21172-CA 388-Pin EPBGA Package Dimensions	2-32
3-1	PCI Commands Recognized by the CIA	3-5
3-2	CIA Prefetch Strategy	3-14

3-3	Row/Column Map	3-23
4-1	21172-CA General Registers	4-2
4-2	21172-CA Diagnostic Registers	4-2
4-3	21172-CA Performance Monitor Registers	4-3
4-4	21172-CA Error Registers	4-3
4-5	21172-CA System Configuration Registers	4-4
4-6	21172-CA PCI Address and Scatter-Gather Registers	4-5
4-7	21172-CA Address Translation Registers	4-6
4-8	CPU and I/O Flush Request	4-14
4-9	PCI READ Prefetch Algorithm	4-14
4-10	HAE_MEM High-Order Sparse-Space Bits	4-17
4-11	Hardware Address Extension Memory Register	4-17
4-12	PERF_CONTROL Low/High Select Encodings	4-26
4-13	CPU Error Information Register 1	4-30
4-14	ECC Syndromes for Single-Bit Errors	4-39
4-15	Memory Port Command Field (MEM_PORT_CMD)	4-42
4-16	Memory Sequencer State Field (SEQ_ST)	4-43
4-17	Encoded Set Select Field (SET_SEL_ENC)	4-43
4-18	Memory Timing Parameters, Encoded Values	4-57
4-19	PCI Address Translation—Scatter-Gather Mapping Enabled	4-69
4-20	PCI Address Translation—Scatter-Gather Mapping Disabled	4-70
5-1	CIA Pin State During Reset	5-2
5-2	CIA Pin State After Reset	5-3
5-3	Registers Initialized by SROM Code	5-4
6-1	21164 Physical Address Space	6-8
6-2	int4_valid_h<3:0> and addr<4:3> for Sparse Space Write Transactions	6-13
6-3	PCI Memory Sparse Space Read/Write Encodings	6-14
6-4	HAE_MEM High Order Sparse Space Bits	6-15
6-5	PCI I/O Sparse Space Read/Write Encodings	6-20
6-6	CIA PCI and (E)ISA I/O Map	6-21
6-7	PCI Configuration Space Read/Write Encodings	6-23
6-8	Generating IDSEL Pin Signals	6-24
6-9	Hardware-Specific Register Address Space	6-27
6-10	21164 Byte/Word Addressing	6-29

6-11	PCI Target Window MASK Register (<i>W_n_MASK</i>)	6-31
6-12	Direct-Mapped PCI Target Address Translation	6-36
6-13	Scatter-Gather Mapped PCI Target Address Translation	6-38
6-14	PCI Window Power-Up Configuration	6-43
7-1	21172-BA Pin List	7-1
7-2	ioc<6:0> Control Functions	7-4
7-3	cmc<8:0> Control Functions	7-5
7-4	Var Encodings for CMC Command 0	7-5
7-5	Var Encodings for CMC Commands 0, 8, A, and B	7-5
7-6	Var Encodings for CMC Commands 2, 3, 4, 5, 6, and 7	7-6
7-7	IOD Bus and MEMDATA Bus Data Width Selection	7-7
7-8	DSW Operating Modes	7-7
7-9	21172-BA Pin Assignment List (Alphabetic)	7-10
7-10	21172-BA Pin Assignment (Numeric)	7-13
7-11	208-Pin PQFP Package Dimensions	7-18

Preface

Purpose and Audience

This document is a support and reference document for engineers using the 21164 Alpha microprocessor to design uniprocessor systems.

Organization

This document contains the following parts, chapters, and appendix:

- Chapter 1, 21172 Core Logic Chipset Overview, provides an overview of the 21172 core logic chipset features.

Part I—21172–CA (CIA) Information

- Chapter 2, 21172–CA Pin Descriptions, provides control, I/O interface, and address chip (CIA) pin descriptions.
- Chapter 3, 21172–CA Architecture Overview, provides the CIA architecture overview.
- Chapter 4, 21172–CA Control and Status Registers, provides a programmer's reference for control and status registers (CSRs).
- Chapter 5, 21172–CA Power-Up and Initialization, provides information on the behavior of CIA during power-up and initialization.
- Chapter 6, 21172–CA Address Mapping, provides information on system address mapping.

Part II—21172–BA (DSW) Information

- Chapter 7, 21172–BA Pin Descriptions, provides data switch chip (DSW) pin descriptions.
- Chapter 8, 21172–BA Architecture Overview, provides an overview of the DSW architecture.

- Appendix A, Technical Support and Ordering Information, provides information on how to obtain technical support and how to order products and documentation.

Document Conventions

This section describes the abbreviation and notation conventions used throughout this guide.

Processor Chip Type

All references to the 21164 microprocessor specifically refer to the version of the Digital Semiconductor 21164 in 0.35- μm CMOS. For the original version of the 21164 (0.5- μm CMOS), please refer to the companion Digital Semiconductor 21171 core logic chipset documentation.

Numbering

All numbers are decimal or hexadecimal unless otherwise specified. In cases of ambiguity, a subscript indicates the radix of nondecimal numbers. For example, 19 is decimal, but 19_{16} and 19A are hexadecimal.

UNPREDICTABLE and UNDEFINED Definitions

Results specified as UNPREDICTABLE may vary from moment to moment, implementation to implementation, and instruction to instruction within implementations. Software can never depend on results specified as UNPREDICTABLE.

Operations specified as UNDEFINED may vary from moment to moment, implementation to implementation, and instruction to instruction within implementations. The operation may vary from nothing to stopping system operation. UNDEFINED operations must not cause the processor to hang, that is, reach a state from which there is no transition to a normal state where the machine can execute instructions.

Note the distinction between results and operations. Nonprivileged software cannot invoke UNDEFINED operations.

Ranges

Ranges are specified by a pair of numbers separated by two periods (..) and are inclusive. For example, a range of integers 0..4 includes the integers 0, 1, 2, 3, and 4.

Extents

Extents are specified by a pair of numbers in angle brackets (<>) separated by a colon (:), and are inclusive. For example, bits <7:3> specifies an extent including bits 7, 6, 5, 4, and 3.

Must Be Zero

Fields specified as must be zero (MBZ) must never be filled by software with a nonzero value. If the processor encounters a nonzero value in a field specified as MBZ, a reserved operand exception occurs.

Should Be Zero

Fields specified as should be zero (SBZ) should be filled by software with a zero value. These fields may be used at some future time. Nonzero values in SBZ fields produce UNPREDICTABLE results.

Register and Memory Figures

Register figures have bit and field position numbering starting at the right (low-order) and increasing to the left (high-order). Memory figures have addresses starting at the top and increasing toward the bottom.

Register Field Notation

Register figures and tabulated descriptions have a mnemonic that indicates the bit or field characteristic as described in Table 1.

Table 1 Register Field Notation

Notation	Description
RW	A read-write bit or field. The value can be read and written by software, microcode, or hardware.
RO	A read-only bit or field. The value can be read by software, microcode, or hardware. The bit is written by hardware. Software or microcode write operations to this bit are ignored.
WO	A write-only bit. The value can be written by software and microcode. The bit is read by hardware. Read operations to this bit by software or microcode return an UNPREDICTABLE result.
WZ	A write-only bit or field. The value can be written by software or microcode. The bit is read by hardware. Read operations to this bit by software or microcode return a zero.
WC	A write-to-clear bit. The value can be read by software or microcode. Software or microcode write operations with a 1 to this bit cause the bit to be cleared by hardware. Software or microcode write operations with a 0 to this bit do not modify the state of the bit.
RC	A read-to-clear field. The bit is written by hardware and remains unchanged until the bit is read. The bit can be read by software or microcode, at which point hardware can write a new value into the field.

Other register fields that are unnamed may be labeled as specified in Table 2.

Table 2 Unnamed Register Field Notation

Notation	Description
0	A 0 in a bit position indicates a register bit that is read as a 0 and is ignored on a write operation.
1	A 1 in a bit position indicates a register bit that is read as a 1 and is ignored on a write operation.
x	An x in a bit position indicates a register bit that does not exist in hardware. The value is UNPREDICTABLE when read and is ignored on a write operation.

Bit Notation

Multiple bit fields are shown as extents (see Extents).

Warning

Warning provides information to prevent personal injury.

Caution

Caution indicates potential damage to equipment or data.

Note

Note provides general information that could be helpful.

Data Units

Table 3 defines the data unit terminology used throughout this manual.

Table 3 Data Units

Term	Words	Bytes	Bits	Other
Byte	1	—	8	—
Word	1	2	16	—
Tribyte	—	3	24	—
Longword	2	4	32	—
Quadword	4	8	64	—
Octaword	8	16	128	Single read fill; that is, the cache space that can be filled in a single read access. It takes two read accesses to fill one Bcache line (see Hexword).
Hexword	16	32	256	Cache block, cache line. The space allocated to a single cache block.

Signal Name References

Signal names in text are printed in boldface lowercase. Mixed-case and uppercase signal naming conventions are ignored. These two examples illustrate the conventions used in this document:

- **MEM_WE_L[1]** is shown as **mem_we_l<1>**.
- **TEST_MODE[1]** is shown as **test_mode<1>**.

21172 Core Logic Chipset Overview

The Digital Semiconductor 21172-AA core logic chipset, hereafter referred to as the 21172 chipset, is used to support the Digital Semiconductor 21164 Alpha microprocessor (21164) in high-performance uniprocessor systems. The chipset includes an interface to the 64-bit peripheral component interconnect (PCI) bus, and associated control and data paths for the 21164, memory, and L3 backup cache (Bcache). Little discrete logic is needed on the module.

Note

All references to the 21164 microprocessor specifically refer to the version of the Digital Semiconductor 21164 Alpha microprocessor in 0.35- μ m CMOS. For the original version of the 21164 (0.5- μ m CMOS) please refer to the companion Digital Semiconductor 21171 core logic chipset documentation.

1.1 Features

The 21172 core logic chipset has the following features.

- Central processing unit—Supports the Digital Semiconductor 21164 Alpha microprocessor
- Components
 - One 21172-CA chip—The control, I/O interface, and address chip (CIA) has 388 pins and is housed in a plastic ball grid array (PBGA) package.
 - Four 21172-BA chips—The data switch chip (DSW) has 208 pins and is housed in a plastic quad flat pack (PQFP) package.
- Data paths
 - 64-bit, ECC-protected data path between CIA and DSW
 - 128-bit ECC-protected data path between the 21164 and DSW

21172 Core Logic Chipset Overview

1.1 Features

- 256-bit ECC-protected memory data path between DSW and memory
- Timing
 - Supports all 21164 clock frequencies
 - Supports system clock frequencies up to 33 MHz (PCI clock)
 - PCI clock, cache, and memory clock period are an integer multiple of the 21164 clock period
- Bcache
 - Third-level cache for the 21164
 - Cache sizes include 0MB, 1MB, and 2MB to 64MB
 - Write-back, ECC-protected cache
- Memory
 - Supports up to 8GB of ECC-protected physical memory
 - Industry-standard memory array configurations
- High-performance PCI
 - 64-bit multiplexed address and data
 - 64-bit PCI address handling
 - Scatter-gather map support
 - No glue logic required to connect PCI-compliant devices

1.2 System Overview

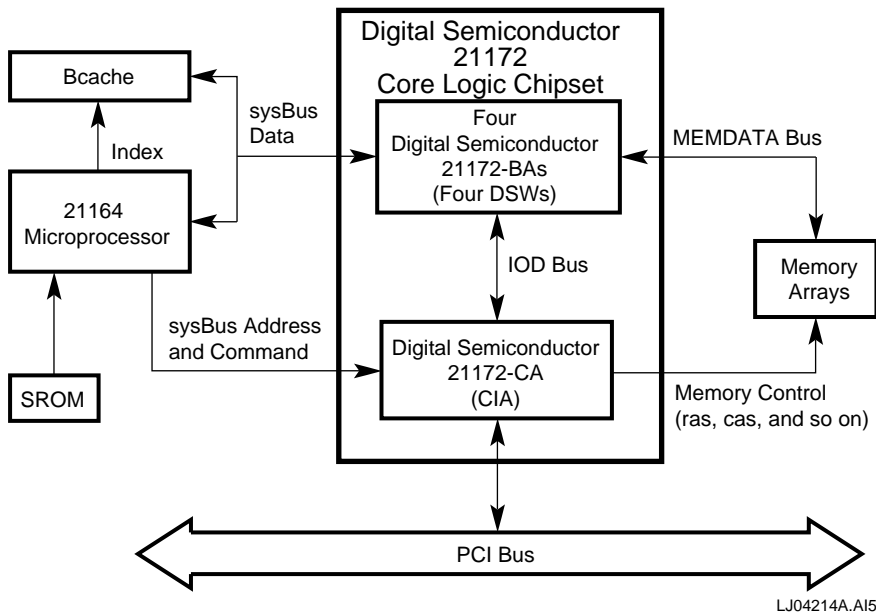
To build a system, such as that shown in Figure 1–1, with the 21172 core logic chipset, the following additional components are required:

- Digital Semiconductor 21164 Alpha microprocessor
- Bcache
- Memory arrays
- Serial ROM (SROM) (for 21164)
- Circuit to generate input clock for 21164
- Circuit to generate input clock for 21172
- Support logic
 - PCI interrupt controller

21172 Core Logic Chipset Overview 1.2 System Overview

- PCI arbiter
- Reset logic
- Console ROM logic and flash ROM
- PCI peripheral devices

Figure 1-1 21172 Core Logic Chipset System Block Diagram



The MEMDATA bus provides a 128-bit/256-bit, ECC-protected data path between the four data path slices (DSWs) and the memory arrays. Each bit slice (DSW) connects to 64 bits of the MEMDATA bus.

The input/output data (IOD) bus provides a 64-bit, ECC-protected data path between the CIA and DSW chips.

The 21164 is connected to the CIA and DSW by the system bus. The system bus contains address, data, and command signal lines. The system bus is the 21164 interface bus with several additional signals to handle direct memory access (DMA) transactions through the CIA and DSW.

21172 Core Logic Chipset Overview

1.2 System Overview

Systems built with the 21172 core logic chipset have the capability to support up to 8GB of memory. The chipset supports the following industry standard DRAM SIMMs in the memory arrays:

1MB × 36
2MB × 36
4MB × 36
8MB × 36
16MB × 36

1.3 21172 Core Logic Chipset Microarchitecture

As shown in Figure 1–1, the 21172 core logic chipset consists of the CIA and four DSW application-specific integrated circuits (ASICs).

- 21172–CA (CIA), 388-pin PBGA: Provides control functions to main memory, and a bridge to the 64-bit PCI bus. In addition, the CIA provides all control functions for the DSW and part of the I/O data path.
- 21172–BA (DSW), 208-pin PQFP: A bit slice ASIC (four required), provides the data path between the 21164, main memory, Bcache, the CIA, and part of the I/O data path.

1.3.1 CIA Microarchitecture

The CIA provides control functions to main memory and the PCI interface. As shown in Figure 1–2, the CIA is divided into the following major functional logic blocks.

The instruction and address logic accepts commands from the 21164 and directs instructions to either the memory port or the I/O port. A 3-entry 21164 instruction queue provides buffering in the event that the memory and I/O ports are busy. A flush address register is also provided to allow DMA read and write operations to interrogate a supported Bcache for the most recent data.

The memory logic provides the row and column addresses for the supported memory banks and all memory control functions (row address select [**ras**], column address select [**cas**], write-enable, and so on). In addition, the memory logic provides control signals to the DSW to initiate data transactions to and from memory.

The I/O address logic handles I/O read and write addresses. The decode logic extracts the PCI address from the dense or sparse space 21164 address encoding, generates PCI I/O read or write addresses, and passes them to the PCI data path logic. The I/O logic can also increment the current address each

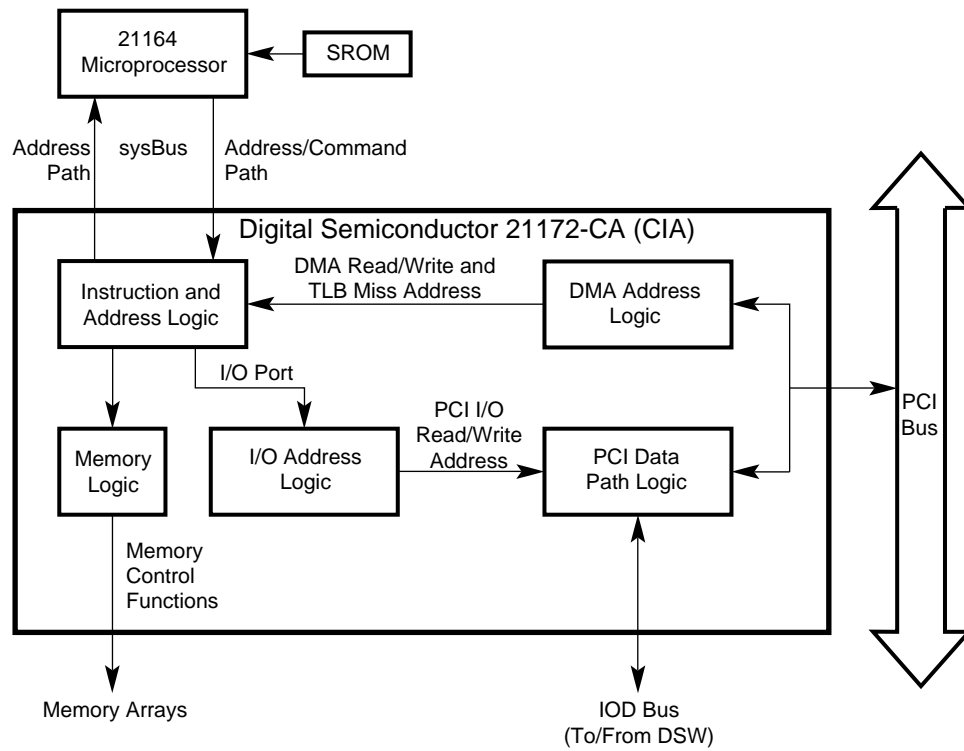
21172 Core Logic Chipset Overview

1.3 21172 Core Logic Chipset Microarchitecture

data cycle. The incremented address provides a pointer to the next data item to be transferred to or from the I/O read/write buffers.

A 3-entry I/O address queue is provided together with a single-entry current address register. These components allow four I/O write operations to be outstanding. The DSW chip provides four corresponding I/O write data buffers.

Figure 1–2 CIA Block Diagram



In addition, a bypass path is provided for the special case of a dense space I/O write transaction with all 32 data bytes valid.

The PCI data path logic provides the interface to the 64-bit PCI, and contains a portion of the data path control logic. Its major functions are to provide:

- ECC code generation on data transactions to the DSW
- ECC code check and correction on data transactions from the DSW

21172 Core Logic Chipset Overview

1.3 21172 Core Logic Chipset Microarchitecture

- I/O read/write addresses to the PCI
- Buffers for all I/O and DMA read and write data transactions

The DMA address logic converts PCI addresses to 21164 memory address space. Two address conversion methods are provided:

- A direct path where a base offset is concatenated with the PCI address
- A scatter-gather map, which maps any 8KB PCI page to any 8KB memory space page

The scatter-gather map translation lookaside buffer (TLB) has eight entries, with each entry containing four consecutive page table entries (PTEs).

1.3.2 DSW Microarchitecture

The DSW provides bidirectional data paths between the 21164, the memory arrays, and the CIA, as shown in Figure 1–3. Four DSW chips, using a bit slice design, provide the complete data path.

The majority of the DSW logic is comprised of data buffers and multiplexers. The CIA, using two encoded control fields, directs data flow to and from the DSW.

The victim buffer has a 64-byte single-entry buffer to contain victim data during a 21164 memory read transaction that generates a victim. Victim data is saved in the buffer, allowing the read transaction to fetch read data from memory and send it to the 21164. Later the victim data is written to memory.

The memory interface contains the data-out register and the data-in register. The data-out register buffers either 21164 victim data to memory, or DMA write data to memory. The data-in register accepts 21164 memory read data. The DSW clocks the register on optimum 15-ns clocks to minimize memory latency.

The I/O write buffer consists of four 32-byte buffers. The four buffers allow the chipset to sustain maximum bandwidth on a large copy operation from memory, through the 21164, to I/O space.

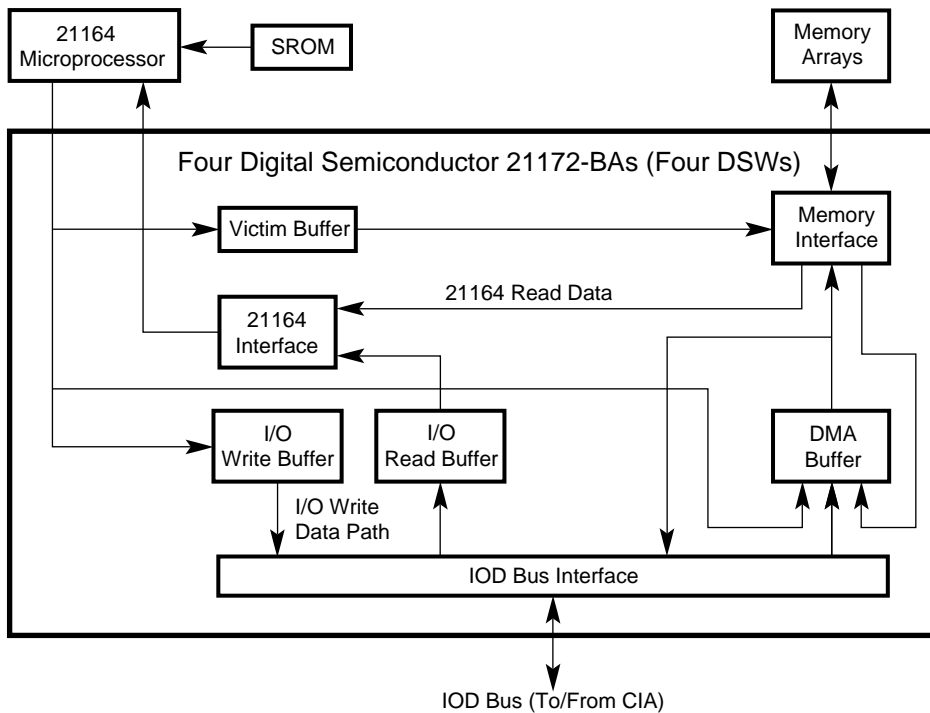
The 32-byte I/O read buffer is provided to assemble the 128-bit data required by the 21164 from the 64-bit data transfers from the IOD bus (CIA-DSW interface).

The DMA read/write buffers consist of two identical buffers (DMA buffer 0 and DMA buffer 1). Each buffer consists of three, 64-byte, single-entry data buffers used to hold data from memory, Bcache, and the PCI. When a buffer is written, its data valid bits are asserted, along with data, to the buffer. The valid bits are used to select appropriate read/write data.

21172 Core Logic Chipset Overview

1.3 21172 Core Logic Chipset Microarchitecture

Figure 1-3 DSW Block Diagram



LJ04216A.A15

1.4 System Clocks

The 21164 internal clock is divided by a programmable value to create an output clock, **sys_clk_out1_h**, which is used as the system clock. The divisor must produce a clock between 25 MHz and 33 MHz.

The system designer must supply a circuit to maintain stability of the system clock and distribute it to PCI devices.

The system clock can be used as a source for a circuit that generates a clock for (E)ISA devices.

21172 Core Logic Chipset Overview

1.5 21172 Core Logic Chipset Characteristics

1.5 21172 Core Logic Chipset Characteristics

Table 1–1 lists the characteristics of the CIA and DSW chips.

Table 1–1 21172–CA (CIA) and 21172–BA (DSW) Chip Characteristics

Characteristic	Specification
21172–CA (CIA) Chip Characteristics	
Operating temperature	Tj maximum = 125°C (257°F)
Power dissipation @Vdd max = 3.465 V, Freq = 33 MHz	3.5 W maximum
21172–BA (DSW) Chip Characteristics	
Operating temperature	Tj maximum = 125°C (257°F)
Power dissipation @Vdd max = 3.465 V, Freq = 33 MHz	1.5 W maximum
Absolute Maximum Ratings	
dc supply voltage	Vss 0.0 V Vdd 3.3 V ±5%
Vin min (3.3-V pins)	–1.0 V
Vin max (3.3-V pins)	Vdd+0.3 V
Vin min (5V-tolerant pins)	–1.0 V
Vin max (5V-tolerant pins)	6.5 V
Storage temperature range	–40°C (–40°F) to 125°C (257°F)
dc Characteristics	
Vil max (except pci_clk)	0.8 V
Vih min (except pci_clk)	2.0 V
Vil max (pci_clk)	0.3*Vdd
Vih min (pci_clk)	0.7*Vdd
Vol max	0.4 V
Voh min	2.4 V

21172 Core Logic Chipset Overview

1.5 21172 Core Logic Chipset Characteristics

Note

To prevent latchup due to the power differences between the gate array's PLL block and its core logic, the following rule must be strictly observed for both the CIA and DSW chips:

*Board **Vdd** on and off ramp time must be slower than 100 ms as measured from 10% to 90% points.*

Part I

21172–CA (CIA) Information

Part I contains five chapters that provide information about the 21172–CA (CIA) chip. The following table lists the topics described in each chapter:

Chapter	Description
2	Pin signals
3	Architecture
4	Control and status registers
5	Power-up and initialization
6	System address mapping

2

21172–CA Pin Descriptions

This chapter provides a description of the 21172–CA pin signals.

2.1 21172–CA Pin List

Table 2–1 lists the pin signals grouped by function. The following abbreviations are used in the Type column of the table:

- I = Input
- O = Output
- B = Bidirectional
- A = Analog
- P = Power

Table 2–1 21172–CA Pin List

Signal Name	Quantity	Type	Signal Name	Quantity	Type
System Bus Signals (59 Total)					
addr<39:4>	36	B	addr_bus_req	1	B
addr_cmd_par	1	B	cack	1	O
cmd<3:0>	4	B	dack	1	O
error¹	1	O	fill	1	O
fill_error	1	O	fill_id	1	O
idle_bc	1	O	int¹	1	O
int4_valid<3:0>	4	I	res<1:0>	2	I
tag_ctl_par	1	O	tag_dirty	1	O
victim_pending	1	I	—	—	—

¹See Section 2.2.1.

(continued on next page)

21172–CA Pin Descriptions

2.1 21172–CA Pin List

Table 2–1 (Cont.) 21172–CA Pin List

Signal Name	Quantity	Type	Signal Name	Quantity	Type
IOD Bus Signals (88 Total)					
cmc<8:0>	9	O	ioc<6:0>	7	O
iod<63:0>	64	B	iod_e<7:0>	8	B
Memory Control Signals (43 Total)					
cas<3:0> ² or cas_l<3:0> ³	4	O	mem_ack_l	1	O
mem_addr<12:0>	13	O	mem_b0	1	O
mem_en	1	O	mem_req_l	1	I
mem_we_l<1:0>	2	O	—	—	—
ras<3:0> ² or cas_l<7:4> ³	4	O	set_sel<15:0> ² or rasx_l<7:0> ³	16	O
PCI Local Bus Signals (88 Total)					
ack64_l	1	B	ad<63:0>	64	B
cbe_l<7:0>	8	B	devsel_l	1	B
frame_l	1	B	gnt_l	1	I
irdy_l	1	B	lock_l	1	I
mem_cs_l	1	I	par	1	B
par64	1	B	perr_l	1	B
req_l	1	O	req64_l	1	B
rst_l	1	O	serr_l	1	I
stop_l	1	B	trdy_l	1	B
Miscellaneous Signals (8 Total)					
pll_test_h<3:0>	4	I	sys_rst_l	1	I
test_mode<1:0>	2	I	test_out	1	O
Phase-Locked Loop Signals (3 Total)					
pll_agnd	1	A	pll_clk	1	I
pll_lp2	1	A	—	—	—

²Signal name in set select memory mode (see Section 2.2.4.1).

³Signal name in alternate memory mode (see Section 2.2.4.1).

(continued on next page)

21172–CA Pin Descriptions
2.1 21172–CA Pin List

Table 2–1 (Cont.) 21172–CA Pin List

Signal Name	Quantity	Type	Signal Name	Quantity	Type
Miscellaneous Power Signals (4 Total)					
aux_vdd	1	P	aux_vss	1	P
pll_vdd	1	P	pll_vss	1	P
Reserved Signals (5 Total)					
debug	1	O	gru_ack	1	I
gru_sel	1	O	iddt	1	I
td	1	I	—	—	—

Pin Totals

Total input pins:	23
Total output pins:	73
Total bidirectional pins:	196
Total analog pins:	2
	—
Total signal pins:	294
Total signal pins:	294
Total Vdd pins:	23
Total Ground pins:	65
Total miscellaneous power pins:	4
Total spare pins:	2
	—
Total pins:	388

21172–CA Pin Descriptions

2.2 21172–CA Signal Descriptions

2.2 21172–CA Signal Descriptions

This section provides a description of the pin function. Signal descriptions are grouped by function and correspond to the pin list (see Table 2–1).

Notes

- The 21164 microprocessor uses the rising edge of **sys_clk_out1_h** as a reference when generating and sampling the system interface signals.
 - In all cases, the term **Vdd** refers to 3.3-V CIA chip power rather than other board power levels such as 5.0 V.
-

2.2.1 System Bus Signals

The *Digital Semiconductor 21164 (366 MHz Through 433 MHz) Alpha Microprocessor Hardware Reference Manual* defines and describes all system bus signals in detail, except for **error** and **int**. Those two signals are described here.

error

The CIA asserts **error** to notify the 21164 that the CIA has detected an error other than a corrected ECC error. The **error** pin should be connected to one of the 21164 **irq_h<n>** pins.

int

The CIA asserts **int** to notify the 21164 that the CIA has detected a corrected ECC error. The **int** pin should be connected to one of the 21164 **irq_h<n>** pins.

2.2.2 PCI Local Bus Signals

This section provides information on the PCI local bus signals.

ack64_l

Acknowledge 64-bit transfer—Asserted by the device that has decoded its address as the target of the current access, indicating that the target device is willing to transfer 64-bit data.

ad<63:0>

The 64-bit address and 64 bits of data are multiplexed on these pins. The upper 32 bits (<63:32>) are reserved except for those transfers where **req_64_l** is asserted.

21172–CA Pin Descriptions

2.2 21172–CA Signal Descriptions

cbe_1<7:0>

Bus command and byte enable are multiplexed on the same pins. During an address phase (when using the DAC command and when **req_64_1** is asserted) the bus command is transferred on **cbe_1<7:0>**. During a data phase these signals carry the byte enable bits, indicating which bytes have meaningful data.

devsel_1

When asserted, indicates that the asserting device has decoded its address as the target of the current access. As an input, **devsel_1** indicates whether any device on the bus has been selected.

frame_1

The cycle frame signal is asserted by the current master to indicate the beginning and duration of an access. When **frame_1** is deasserted, the final phase of the access has started.

gnt_1

Assertion of **gnt_1** indicates that access to the bus has been granted. This is a point-to-point signal. Every master has its own grant line.

irdy_1

Assertion of **irdy_1** (initiator ready) indicates that the bus master is ready to complete the current phase of the transaction. Signal **irdy_1** is used in conjunction with **trdy_1**. A data phase is completed on any clock where both signals are sampled asserted.

lock_1

Assertion of **lock_1** indicates an atomic operation that may require multiple transactions to complete. When **lock_1** is asserted, nonexclusive transactions may proceed to an address that is not currently locked.

mem_cs_1

The *PCI Local Bus Specification* does not preclude product-specific function or performance enhancements by way of sideband signals. A sideband signal is loosely defined as any signal that is not part of the PCI specification and connects two or more PCI nodes and has meaning only to these nodes. Sideband signals must not violate PCI local bus specifications.

Signal **mem_cs_1** is a sideband signal. It is used by the PCI-to-EISA bridge to help manage the EISA memory map and peripheral component architecture (PCA) compatibility holes (see Section 6.4.4.2).

21172–CA Pin Descriptions

2.2 21172–CA Signal Descriptions

par

The **par** signal is used to create even parity across **ad<31:0>** and **cbe_1<3:0>**. Signal **par** is stable and valid one clock cycle after the address phase. Signal **par** remains valid until one clock cycle after the completion of the current data phase. The master asserts **par** for address and write data phases; the target asserts **par** for read data phases.

par64

The **par64** signal is used to create even parity across **ad<63:32>** and **cbe_1<7:4>**. Signal **par64** is stable and valid one clock cycle after the address phase. Signal **par64** remains valid until one clock cycle after the completion of the current data phase. The master asserts **par64** for address and write data phases; the target asserts **par64** for read data phases.

perr_1

The **perr_1** signal (parity error) is used to report data parity errors that occur during all PCI transactions except a special cycle. Signal **perr_1** is a tristate signal that must be asserted by the data-receiving node two cycles after data with bad parity is detected. Signal **perr_1** must be asserted for a minimum of one clock cycle for each phase during which bad data parity is detected. Signal **perr_1**, like all sustained tristate signals, must be asserted high for at least one cycle before being returned from asserted to the tristate condition.

req_1

Assertion of the **req_1** signal (request) indicates to the arbiter that this node wants to use the PCI bus. Every master node has its own **req_1** pin.

req64_1

The current bus master asserts **req64_1** (request 64-bit transfer), indicating that it wants to transfer data by using 64 bits. Signal **req64_1** has the same timing as **frame_1**. Signal **req64_1** is asserted low while signal **rst_1** is asserted. Nodes connected to the 64-bit PCI bus extension will see the signal asserted low. Those nodes not connected to the 64-bit PCI bus extension will see the signal asserted high by a pull-up device.

rst_1

The **rst_1** signal (reset) is used to bring PCI-specific registers, sequencers, and signals to a consistent state. When **rst_1** is asserted, all PCI output pins must be driven to their benign state, usually tristated. Although the signal is asynchronous, deassertion must be a clean, bounce-free edge.

21172–CA Pin Descriptions

2.2 21172–CA Signal Descriptions

serr_1

The **serr_1** pin (system error) is used to report address parity errors, data parity errors on the special cycles command, or any other system error where the result will be catastrophic. If a node does not want a nonexistent memory error to be generated, a different reporting mechanism is required. Signal **serr_1** is an open-drain output and is asserted for a minimum of one clock cycle by the node reporting the error.

stop_1

The **stop_1** pin indicates that the current target is requesting that the master stop the current transaction.

trdy_1

When asserted, the **trdy_1** (target ready) pin indicates the target node's ability to complete the current phase of the transaction. Signal **trdy_1** is used in conjunction with **irdy_1**. A data phase is completed on any clock assertion where both **trdy_1** and **irdy_1** are sampled asserted. During a read transaction, **trdy_1** indicates that valid data is present on **ad<31:0>**. During a write transaction, **trdy_1** indicates that the target is ready to accept data. Wait cycles are inserted until both **trdy_1** and **irdy_1** are asserted together.

21172–CA Pin Descriptions
2.2 21172–CA Signal Descriptions

2.2.3 IOD Bus Signals

This section provides a description of the IOD bus signals.

iod<63:0>

The **iod<63:0>** signals carry data between the CIA and DSW on the IOD bus.

iod_e<7:0>

The **iod_e<7:0>** signals carry ECC between the CIA and DSW on the IOD bus. All ECC generation and detection is performed in the 21164 and CIA.

ioc<6:0>

The **ioc<6:0>** signals allow the CIA to control data flow on the IOD bus. The data transfers are between memory and the PCI bus. CIA control is encoded as shown in Table 2–2.

Table 2–2 ioc<6:0> Control Functions

ioc<6:4>	ioc<3>	ioc<2>	ioc<1>	ioc<0>	Function
000	Buffer <i>n</i>	x	x	x	Clear valid bits in PCI buffer <i>n</i> .
001	x	x	x	x	Load the LW register (Reserved).
010	Buffer <i>n</i>	Adr <i>m2</i>	Adr <i>m1</i>	Adr <i>m0</i>	Write PCI buffer <i>n</i> , adr <i>m</i> .
011	Mask3	Mask2	Mask1	Mask0	Write IOR (QW mask).
100	Buffer <i>n</i>	Adr <i>m2</i>	Adr <i>m1</i>	Adr <i>m0</i>	Read DMA buffer <i>n</i> , adr <i>m</i> .
101	Buffer <i>n1</i>	Buffer <i>n0</i>	Adr <i>m1</i>	Adr <i>m0</i>	Read IOW buffer <i>n</i> , adr <i>m</i> .
110	Buffer <i>n1</i> Buffer <i>n</i>	Buffer <i>n0</i> Adr <i>m2</i>	Adr <i>m1</i>	Adr <i>m0</i>	Read buffer (IOW or DMA) low longword.
111	Mask3	Mask2	Mask1	Mask0	Start IOW (buffer mask).

21172–CA Pin Descriptions
2.2 21172–CA Signal Descriptions

cmc<8:0>

The **cmc<8:0>** signals control flow of data between the 21164 and the memory arrays. They also start and stop loading of the free-running buffers (victim, flush, IOW buffers) in the DSW. In Table 2–3, A5 and A4 are address bits. The Var field encodes the transaction being performed as shown in Table 2–4, Table 2–5, and Table 2–6.

Table 2–3 cmc<8:0> Control Functions

cmc<8:5>	cmc<4>	cmc<3>	cmc<2>	cmc<1>	cmc<0>	Function
0000	Var 4 ¹	Var 3	Var 2	Var 1	Var 0	Start/stop the buffer.
0010	x	A4 ²	Delay	Var 1	Var 0	Move fill data from memory to 21164.
0011	x	A4	x	Var 1	Var 0	Move IOR buffer data to 21164.
0100	A5 ²	A4	Delay	Var 1	Var 0	Move memory data to memory buffer 0.
0101	A5	A4	Delay	Var 1	Var 0	Move memory data to memory buffer 1.
0110	A5	A4	Delay	Stop IOW _{n1}	Stop IOW _{n0}	Move memory data to memory buffer 0, stop IOW buffer <i>n</i> .
0111	A5	A4	Delay	Stop IOW _{n1}	Stop IOW _{n0}	Move memory data to memory buffer 1, stop IOW buffer <i>n</i> .
1000	A5	A4	Var 2	Var 1	Var 0	Write victim buffer to memory.
1010	A5	A4	Var 2	Var 1	Var 0	Write DMA0 buffer to memory.
1011	A5	A4	Var 2	Var 1	Var 0	Write DMA1 buffer to memory.

¹The variable field encodes the transaction being performed. See Table 2–4, Table 2–5, and Table 2–6.

²A5 and A4 are address bits.

21172–CA Pin Descriptions
2.2 21172–CA Signal Descriptions

Table 2–4 Var Encodings for CMC Command 0

Var<4:3>	Function
00	None.
01	Start victim.
10	Start flush buffer 0 and clear PCI 0 valid bits.
11	Start flush buffer 1 and clear PCI 0 valid bits.

Table 2–5 Var Encodings for CMC Commands 0, 8, A, and B

Var<2:0>	Function
000	None.
001	Stop victim.
010	Stop flush buffer 0.
011	Stop flush buffer 1.
100	Stop IOW 0.
101	Stop IOW 1.
110	Stop IOW 2.
111	Stop IOW 3.

Table 2–6 Var Encodings for CMC Commands 2, 3, 4, 5, 6, and 7

Var<1:0>	Function
00	None.
01	Stop victim.
10	Stop flush buffer 0.
11	Stop flush buffer 1.

2.2.4 Memory Control Signals

This section provides a description of the memory control signals.

2.2.4.1 Set Select/Alternate Memory Modes

Certain memory control signals can be defined in one of two modes, set select mode or alternate mode. The following paragraphs define the two modes.

Set Select Memory Mode

In set select mode, the **ras**, **cas**, and **set_sel** pins are defined normally and are high-asserted. In this mode, the board must “AND” the **set_sel** and **ras** signals external to the CIA.

When signal **pll_test<0>** is deasserted (0) during reset, the CIA is put into set select memory mode. Status bit ALT_MEM reset in the CIA revision register (CIA_REV<8>=0) indicates this mode.

Alternate Memory Mode

In alternate mode, the **ras**, **cas**, and **set_sel** pins are redefined and are low-asserted. In this mode, there are no **set_sel** signals, only **ras** and **cas**.

When signal **pll_test<0>** is asserted (1) during reset, the CIA is put into alternate memory mode. Status bit ALT_MEM set in the CIA revision register (CIA_REV<8>=1) indicates this mode. In alternate memory mode, bits MBAx<15> (S1_VALID) must be zero.

Alternate memory mode redefines the memory control signals as listed in Table 2–7.

Table 2–7 Alternate Memory Mode Memory Control Signal Redefinition

Set Select Mode Signal	Is redefined as . . .	Alternate Mode Signal
cas<0>		cas_l<0>
cas<1>		cas_l<1>
cas<2>		cas_l<2>
cas<3>		cas_l<3>
ras<0>		cas_l<4>
ras<1>		cas_l<5>
ras<2>		cas_l<6>
ras<3>		cas_l<7>
set_sel<0>		rasa_l<0>
set_sel<1>		rasa_l<1>

(continued on next page)

21172–CA Pin Descriptions
 2.2 21172–CA Signal Descriptions

Table 2–7 (Cont.) Alternate Memory Mode Memory Control Signal Redefinition

Set Select Mode Signal	Is redefined as . . .	Alternate Mode Signal
set_sel<2>		rasa_l<2>
set_sel<3>		rasa_l<3>
set_sel<4>		rasa_l<4>
set_sel<5>		rasa_l<5>
set_sel<6>		rasa_l<6>
set_sel<7>		rasa_l<7>
set_sel<8>		rasb_l<0>
set_sel<9>		rasb_l<1>
set_sel<10>		rasb_l<2>
set_sel<11>		rasb_l<3>
set_sel<12>		rasb_l<4>
set_sel<13>		rasb_l<5>
set_sel<14>		rasb_l<6>
set_sel<15>		rasb_l<7>

Signals **rasa_l** and **rasb_l** are actually two copies of eight unique **ras** signals. Signals **rasa_l<x>** = signals **rasb_l<x>**. Each **ras_x_l** pair is controlled by a memory base address register (see Section 4.6.2).

- Memory base address register 0 (MBA0) controls **ras_x_l<0>**.
- Memory base address register 2 (MBA2) controls **ras_x_l<1>**.
- Memory base address register 4 (MBA4) controls **ras_x_l<2>**.
- Memory base address register 6 (MBA6) controls **ras_x_l<3>**.
- Memory base address register 8 (MBA8) controls **ras_x_l<4>**.
- Memory base address register 10 (MBAA) controls **ras_x_l<5>**.
- Memory base address register 12 (MBAC) controls **ras_x_l<6>**.
- Memory base address register 14 (MBAE) controls **ras_x_l<7>**.

MBAx registers can be configured to the same address ranges. This produces additional copies of **ras**. A system could have two copies of eight **ras** signals, four copies of four **ras** signals, eight copies of two **ras** signals, and so forth. Signals **cas_l<x>** are eight copies of a single **cas**.

2.2.4.2 Memory Control Signal Definitions

This section describes the memory control signals.

cas<3:0> or **cas_l<7:0>**

The **cas<3:0>** (set select mode) or **cas_l<7:0>** (alternate mode) signals select the column address in the memory arrays.

21172–CA Pin Descriptions

2.2 21172–CA Signal Descriptions

mem_ack_l

The **mem_ack_l** PCI sideband signal is asserted by the CIA when it has cleared the transaction path in preparation to receive a transaction from the PCI-to-EISA bridge. See **mem_req_l**.

mem_addr<12:0>

The **mem_addr<12:0>** signals carry the memory array address.

mem_b0

Signal **mem_b0** is a copy of signal **mem_addr<0>**. It is provided for use by some DIMMs that require extra drive due to loading. Signal **mem_b0** can be used to equalize the load, and unlike **mem_addr<0>** is output only.

mem_en

The **mem_en** signal enables operation of the memory arrays.

mem_req_l

Signal **mem_req_l** is asserted by the PCI-to-EISA bridge to notify the CIA that it has a transaction pending. The CIA completes all transactions in progress and then asserts **mem_ack_l** to the PCI-to-EISA bridge. The PCI-to-EISA bridge then starts the transaction. See **mem_ack_l**.

The PCI-to-EISA bridge set can be made using these two chips:

- Intel 82374EB EISA System Component (ESC)
- Intel 82375EB PCI-to-EISA Bridge (PCEB)

The PCEB allows a certain time period for a PCI transaction to complete. The **mem_req_l** and **mem_ack_l** PCI sideband signals are used to ensure that the CIA performs transactions from the PCI-to-EISA bridge immediately.

mem_we_l<1:0>

The **mem_we_l<1:0>** signals carry the write-enable signal to the memory arrays.

ras<3:0>

The **ras<3:0>** (set select mode) signals select the row address in the memory arrays.

set_sel<15:0> or rasx_l<7:0>

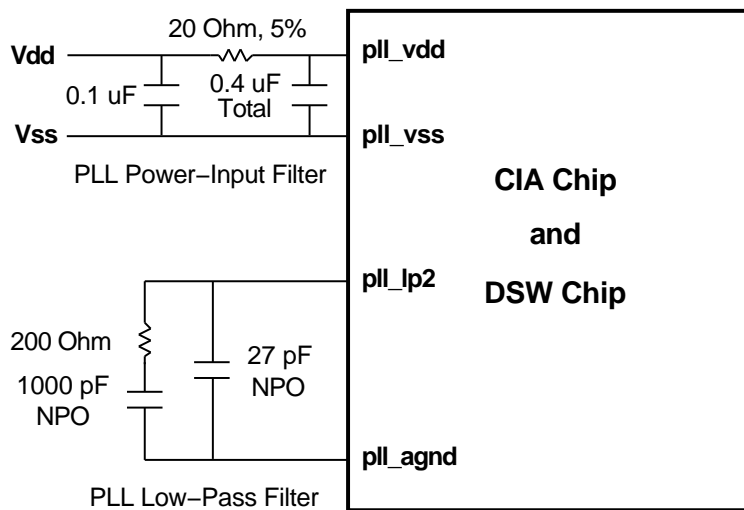
In set select memory mode, the **set_sel<15:0>** signals select a memory array. In alternate memory mode, the **rasx_l<7:0>** signals select the row address in the memory arrays.

21172-CA Pin Descriptions
 2.2 21172-CA Signal Descriptions

2.2.5 Phase-Locked Loop Signals

This section describes the phase-locked loop signals. See Figure 2-1 for required external filter components.

Figure 2-1 External PLL Filter Components



MK-1455-28

pll_agnd

Signal **pll_agnd** is an analog ground output from the PLL used for the low-pass filter connected to **pll_lp2**. It *must not* be connected to system ground (**Vss**).

pll_clk

The **pll_clk** signal is the system clock (or reference signal) input to the phase-locked loop (PLL) circuit. The **pll_clk** signal is the source for CIA timing. It supplies a clock with the same frequency as that supplied to the PCI.

pll_lp2

Signal **pll_lp2** is the PLL VCO loop filter input.

2.2.6 Miscellaneous Power Signals

This section describes the miscellaneous power signals. See Figure 2–1 for required external filter components.

aux_vdd

Signal **aux_vdd** is an isolated **Vdd** pin for the PLL section of the CIA.

aux_vss

Signal **aux_vss** is an isolated **Vss** pin for the PLL section of the CIA.

pll_vdd

Signal **pll_vdd** is the power input to the PLL. It is connected to **Vdd**.

pll_vss

Signal **pll_vss** is the ground input to the PLL. It is connected to the system ground (**Vss**).

2.2.7 Miscellaneous Signals

This section provides a description of the miscellaneous signals.

pll_test<3:0>

Signals **pll_test<3:0>** select which internal chip information is presented at output pin **debug**. Signals **pll_test_h<3:1>** are test pins reserved for use by Digital. They should be tied to ground. Signal **pll_test_h<0>** selects the memory mode: 0=set select mode and 1=alternate mode (see Section 2.2.4.1).

sys_rst_l

Signal **sys_rst_l** is the system reset input signal to the CIA. When **sys_rst_l** is asserted, the CIA resets its internal state.

test_mode<1:0>

Signals **test_mode<1:0>** define the four operating modes of the CIA as listed in Table 2–8.

Table 2–8 DSW Operating Modes

test_mode<1:0>	DSW Operating Mode
00	Tristate mode
01	Normal mode
10	Scan mode
11	PLL test mode

test_out

Signal **test_out** is the end output of the parametric NAND tree.

21172–CA Pin Descriptions

2.2 21172–CA Signal Descriptions

2.2.8 Reserved Signals

This section provides a description of the reserved signals.

debug

Signal **debug** presents internal CIA information selected by signals **pll_test_h<3:0>**. This signal is reserved to Digital.

gru_ack

Signal **gru_ack** is reserved for use by Digital.

gru_sel

Signal **gru_sel** is reserved for use by Digital.

iddt

Signal **iddt** is a test pin reserved for use by Digital. It should be tied to ground.

td

Signal **td** is a test pin reserved for use by Digital. It should be tied to ground.

21172–CA Pin Descriptions
2.2 21172–CA Signal Descriptions

2.2.9 21172–CA Pin Name List (Alphabetic)

Table 2–9 lists the 21172–CA (CIA) pins by name in alphabetical order. The following abbreviations are used in the Type column of the table:

- I = Input
- O = Output
- B = Bidirectional
- A = Analog
- P = Power

Table 2–9 21172–CA Pin Names (Alphabetic)

Pin Name	Pin Number	Type	5-V Tolerant	Pin Name	Pin Number	Type	5-V Tolerant
ack64_1	C10	B	Yes	ad<0>	AA4	B	Yes
ad<1>	Y3	B	Yes	ad<2>	V2	B	Yes
ad<3>	U2	B	Yes	ad<4>	T2	B	Yes
ad<5>	R2	B	Yes	ad<6>	M1	B	Yes
ad<7>	M4	B	Yes	ad<8>	L4	B	Yes
ad<9>	K4	B	Yes	ad<10>	J4	B	Yes
ad<11>	H4	B	Yes	ad<12>	E2	B	Yes
ad<13>	E3	B	Yes	ad<14>	E4	B	Yes
ad<15>	C4	B	Yes	ad<16>	C18	B	Yes
ad<17>	D19	B	Yes	ad<18>	C21	B	Yes
ad<19>	B23	B	Yes	ad<20>	C25	B	Yes
ad<21>	D26	B	Yes	ad<22>	F24	B	Yes
ad<23>	G25	B	Yes	ad<24>	J26	B	Yes
ad<25>	K23	B	Yes	ad<26>	L23	B	Yes
ad<27>	M24	B	Yes	ad<28>	T26	B	Yes
ad<29>	U26	B	Yes	ad<30>	U25	B	Yes
ad<31>	V24	B	Yes	ad<32>	AA1	B	Yes
ad<33>	W4	B	Yes	ad<34>	V4	B	Yes
ad<35>	T1	B	Yes	ad<36>	T4	B	Yes
ad<37>	R3	B	Yes	ad<38>	L1	B	Yes
ad<39>	L2	B	Yes	ad<40>	K3	B	Yes

(continued on next page)

21172–CA Pin Descriptions

2.2 21172–CA Signal Descriptions

Table 2–9 (Cont.) 21172–CA Pin Names (Alphabetic)

Pin Name	Pin Number	Type	5-V Tolerant	Pin Name	Pin Number	Type	5-V Tolerant
ad<41>	J2	B	Yes	ad<42>	F1	B	Yes
ad<43>	E1	B	Yes	ad<44>	D1	B	Yes
ad<45>	C1	B	Yes	ad<46>	B3	B	Yes
ad<47>	B2	B	Yes	ad<48>	C19	B	Yes
ad<49>	A22	B	Yes	ad<50>	D21	B	Yes
ad<51>	D22	B	Yes	ad<52>	D24	B	Yes
ad<53>	E25	B	Yes	ad<54>	F26	B	Yes
ad<55>	H23	B	Yes	ad<56>	J24	B	Yes
ad<57>	L26	B	Yes	ad<58>	M25	B	Yes
ad<59>	R26	B	Yes	ad<60>	R25	B	Yes
ad<61>	T23	B	Yes	ad<62>	U23	B	Yes
ad<63>	Y25	B	Yes	addr<4>	K2	B	No
addr<5>	J3	B	No	addr<6>	H3	B	No
addr<7>	F2	B	No	addr<8>	G4	B	No
addr<9>	F3	B	No	addr<10>	D2	B	No
addr<11>	C2	B	No	addr<12>	D3	B	No
addr<13>	F4	B	No	addr<14>	A3	B	No
addr<15>	B4	B	No	addr<16>	A4	B	No
addr<17>	D5	B	No	addr<18>	B5	B	No
addr<19>	D6	B	No	addr<20>	C6	B	No
addr<21>	A6	B	No	addr<22>	D7	B	No
addr<23>	B7	B	No	addr<24>	A18	B	No
addr<25>	C17	B	No	addr<26>	D17	B	No
addr<27>	B19	B	No	addr<28>	D18	B	No
addr<29>	B20	B	No	addr<30>	A21	B	No
addr<31>	C20	B	No	addr<32>	B21	B	No
addr<33>	D20	B	No	addr<34>	B22	B	No
addr<35>	A23	B	No	addr<36>	C22	B	No
addr<37>	A24	B	No	addr<38>	C23	B	No
addr<39>	B24	B	No	addr_bus_req	C12	B	No
addr_cmd_par	A16	B	No	aux_vdd	AD25	P	No
aux_vss	AC24	P	No	cack	K1	O	No

(continued on next page)

21172–CA Pin Descriptions
2.2 21172–CA Signal Descriptions

Table 2–9 (Cont.) 21172–CA Pin Names (Alphabetic)

Pin Name	Pin Number	Type	5-V Tolerant	Pin Name	Pin Number	Type	5-V Tolerant
cas<0> ¹	AD2	O	No	cas<1> ¹	AD1	O	No
cas<2> ¹	AC3	O	No	cas<3> ¹	AC2	O	No
cas_l<0> ²	AD2	O	No	cas_l<1> ²	AD1	O	No
cas_l<2> ²	AC3	O	No	cas_l<3> ²	AC2	O	No
cas_l<4> ²	AC1	O	No	cas_l<5> ²	AB4	O	No
cas_l<6> ²	AB2	O	No	cas_l<7> ²	AB1	O	No
cbe_l<0>	D12	B	Yes	cbe_l<1>	C14	B	Yes
cbe_l<2>	B16	B	Yes	cbe_l<3>	B17	B	Yes
cbe_l<4>	A15	B	Yes	cbe_l<5>	C15	B	Yes
cbe_l<6>	C16	B	Yes	cbe_l<7>	B18	B	Yes
cmc<0>	AE10	O	No	cmc<1>	AD10	O	No
cmc<2>	AC10	O	No	cmc<3>	AE9	O	No
cmc<4>	AE8	O	No	cmc<5>	AD9	O	No
cmc<6>	AC9	O	No	cmc<7>	AE7	O	No
cmc<8>	AD8	O	No	cmd<0>	B15	B	No
cmd<1>	D15	B	No	cmd<2>	A17	B	No
cmd<3>	D16	B	No	dack	L3	O	No
debug	W23	O	No	devsel_1	D8	B	Yes
error	AF9	O	Yes	fill	A12	O	No
fill_error	B11	O	No	fill_id	A11	O	No
frame_1	B6	B	Yes	gnt_1	C11	I	Yes
gru_ack	AC11	I	Yes	gru_sel	AD11	O	Yes
iddt	AE25	I	No	idle_bc	D10	O	No
int	AE11	O	Yes	int4_valid<0>	J1	I	No
int4_valid<1>	H2	I	No	int4_valid<2>	G2	I	No
int4_valid<3>	G3	I	No	ioc<0>	AF10	O	No
ioc<1>	AC12	O	No	ioc<2>	AD12	O	No
ioc<3>	AE12	O	No	ioc<4>	AF11	O	No
ioc<5>	AD13	O	No	ioc<6>	AE13	O	No
iod<0>	AF12	B	Yes	iod<1>	AD14	B	Yes
iod<2>	AE14	B	Yes	iod<3>	AD15	B	Yes

¹Signal name in set select memory mode.

²Signal name in alternate memory mode.

(continued on next page)

21172–CA Pin Descriptions
2.2 21172–CA Signal Descriptions

Table 2–9 (Cont.) 21172–CA Pin Names (Alphabetic)

Pin Name	Pin Number	Type	5-V Tolerant	Pin Name	Pin Number	Type	5-V Tolerant
iod<4>	AC15	B	Yes	iod<5>	AE15	B	Yes
iod<6>	AF15	B	Yes	iod<7>	AD16	B	Yes
iod<8>	AC16	B	Yes	iod<9>	AE16	B	Yes
iod<10>	AF16	B	Yes	iod<11>	AD17	B	Yes
iod<12>	AC17	B	Yes	iod<13>	AE17	B	Yes
iod<14>	AF17	B	Yes	iod<15>	AD18	B	Yes
iod<16>	AC18	B	Yes	iod<17>	AE18	B	Yes
iod<18>	AF18	B	Yes	iod<19>	AD19	B	Yes
iod<20>	AC19	B	Yes	iod<21>	AE19	B	Yes
iod<22>	AE20	B	Yes	iod<23>	AD20	B	Yes
iod<24>	AC20	B	Yes	iod<25>	AF21	B	Yes
iod<26>	AE21	B	Yes	iod<27>	AD21	B	Yes
iod<28>	AC21	B	Yes	iod<29>	AF22	B	Yes
iod<30>	AE22	B	Yes	iod<31>	AD22	B	Yes
iod<32>	AA26	B	Yes	iod<33>	W24	B	Yes
iod<34>	V23	B	Yes	iod<35>	W25	B	Yes
iod<36>	V25	B	Yes	iod<37>	U24	B	Yes
iod<38>	V26	B	Yes	iod<39>	T24	B	Yes
iod<40>	T25	B	Yes	iod<41>	R23	B	Yes
iod<42>	R24	B	Yes	iod<43>	P24	B	Yes
iod<44>	P25	B	Yes	iod<45>	N24	B	Yes
iod<46>	N25	B	Yes	iod<47>	M23	B	Yes
iod<48>	M26	B	Yes	iod<49>	L24	B	Yes
iod<50>	L25	B	Yes	iod<51>	K24	B	Yes
iod<52>	K25	B	Yes	iod<53>	K26	B	Yes
iod<54>	J23	B	Yes	iod<55>	J25	B	Yes
iod<56>	H24	B	Yes	iod<57>	H25	B	Yes
iod<58>	G24	B	Yes	iod<59>	G23	B	Yes
iod<60>	F25	B	Yes	iod<61>	F23	B	Yes
iod<62>	E26	B	Yes	iod<63>	E24	B	Yes
iod_e<0>	AC22	B	Yes	iod_e<1>	AF23	B	Yes
iod_e<2>	AE23	B	Yes	iod_e<3>	AD23	B	Yes

(continued on next page)

21172–CA Pin Descriptions 2.2 21172–CA Signal Descriptions

Table 2–9 (Cont.) 21172–CA Pin Names (Alphabetic)

Pin Name	Pin Number	Type	5-V Tolerant	Pin Name	Pin Number	Type	5-V Tolerant
iod_e<4>	E23	B	Yes	iod_e<5>	D25	B	Yes
iod_e<6>	C26	B	Yes	iod_e<7>	B25	B	Yes
irdy_1	B10	B	Yes	lock_1	C9	I	Yes
mem_ack_1	W2	O	Yes	mem_addr<0>	V1	O	No
mem_addr<1>	V3	O	No	mem_addr<2>	U1	O	No
mem_addr<3>	U4	O	No	mem_addr<4>	U3	O	No
mem_addr<5>	T3	O	No	mem_addr<6>	R1	O	No
mem_addr<7>	R4	O	No	mem_addr<8>	P2	O	No
mem_addr<9>	N2	O	No	mem_addr<10>	N3	O	No
mem_addr<11>	M2	O	No	mem_addr<12>	M3	O	No
mem_b0	W3	O	No	mem_cs_1	AA3	I	Yes
mem_en	AE2	O	No	mem_req_1	Y2	I	Yes
mem_we_1<0>	AA2	O	No	mem_we_1<1>	Y4	O	No
par	A5	B	Yes	par64	C5	B	Yes
perr_1	C7	B	Yes	pll_agnd	AD26	A	No
pll_clk	AC26	I	No	pll_lp2	AB23	A	No
pll_test<0>	Y23	I	Yes	pll_test<1>	AB26	I	Yes
pll_test<2>	AA25	I	Yes	pll_test<3>	Y24	I	Yes
pll_vdd	AB24	P	No	pll_vss	AC25	P	No
ras<0>¹	AC1	O	No	ras<1>¹	AB4	O	No
ras<2>¹	AB2	O	No	ras<3>¹	AB1	O	No
rasa_1<0>²	AF6	O	No	rasa_1<1>²	AC8	O	No
rasa_1<2>²	AD7	O	No	rasa_1<3>²	AE6	O	No
rasa_1<4>²	AF5	O	No	rasa_1<5>²	AC7	O	No
rasa_1<6>²	AD6	O	No	rasa_1<7>²	AE5	O	No
rasb_1<0>²	AF4	O	No	rasb_1<1>²	AC6	O	No
rasb_1<2>²	AD5	O	No	rasb_1<3>²	AE4	O	No
rasb_1<4>²	AF3	O	No	rasb_1<5>²	AC5	O	No
rasb_1<6>²	AD4	O	No	rasb_1<7>²	AE3	O	No
req_1	B12	O	Yes	req64_1	D11	B	Yes
res<0>	B13	I	No	res<1>	B14	I	No

¹Signal name in set select memory mode.

²Signal name in alternate memory mode.

(continued on next page)

21172–CA Pin Descriptions

2.2 21172–CA Signal Descriptions

Table 2–9 (Cont.) 21172–CA Pin Names (Alphabetic)

Pin Name	Pin Number	Type	5-V Tolerant	Pin Name	Pin Number	Type	5-V Tolerant
rst_l	AB3	O	Yes	serr_l	A10	I	Yes
set_sel<0> ¹	AF6	O	No	set_sel<1> ¹	AC8	O	No
set_sel<2> ¹	AD7	O	No	set_sel<3> ¹	AE6	O	No
set_sel<4> ¹	AF5	O	No	set_sel<5> ¹	AC7	O	No
set_sel<6> ¹	AD6	O	No	set_sel<7> ¹	AE5	O	No
set_sel<8> ¹	AF4	O	No	set_sel<9> ¹	AC6	O	No
set_sel<10> ¹	AD5	O	No	set_sel<11> ¹	AE4	O	No
set_sel<12> ¹	AF3	O	No	set_sel<13> ¹	AC5	O	No
set_sel<14> ¹	AD4	O	No	set_sel<15> ¹	AE3	O	No
stop_l	C8	B	Yes	sys_rst_l	AE24	I	Yes
tag_ctl_par	B9	O	Yes	tag_dirty	A9	O	Yes
td	AA24	I	Yes	test_mode<0>	AB25	I	Yes
test_mode<1>	AA23	I	Yes	test_out	AF24	O	No
trdy_l	D9	B	Yes	victim_pending	B8	I	No

Pin Name	Pin Number	Type	5-V Tolerant
Vdd	A7, A14, A20, A26, D4, D13, D23, G1, G26, N1, N23, P4, P26, Y1, Y26, AC4, AC14, AC23, AF1, AF7, AF13, AF20, AF26	P	—
Ground	A1, A2, A8, A13, A19, A25, B1, B26, C3, C24, D14, H1, H26, L11, L12, L13, L14, L15, L16, M11, M12, M13, M14, M15, M16, N4, N11, N12, N13, N14, N15, N16, N26, P1, P11, P12, P13, P14, P15, P16, P23, R11, R12, R13, R14, R15, R16, T11, T12, T13, T14, T15, T16, W1, W26, AC13, AD3, AD24, AE1, AE26, AF2, AF8, AF14, AF19, AF25	P	—
Spare	C13, P3	—	—

¹Signal name in set select memory mode.

21172–CA Pin Descriptions
2.2 21172–CA Signal Descriptions

2.2.10 21172–CA Pin Assignment List (Alphanumeric)

Table 2–10 lists the 21172–CA pins in alphanumeric order. The following abbreviations are used in the Type column of the table:

- I = Input
- O = Output
- B = Bidirectional
- A = Analog
- P = Power

Table 2–10 21172–CA Pin Assignment (Alphanumeric)

Pin Number	Pin Name	Type	5-V Tolerant	Pin Number	Pin Name	Type	5-V Tolerant
A1	Ground	P	—	A2	Ground	P	—
A3	addr<14>	B	No	A4	addr<16>	B	No
A5	par	B	Yes	A6	addr<21>	B	No
A7	Vdd	P	—	A8	Ground	P	—
A9	tag_dirty	O	Yes	A10	serr_1	I	Yes
A11	fill_id	O	No	A12	fill	O	No
A13	Ground	P	—	A14	Vdd	P	—
A15	cbe_l<4>	B	Yes	A16	addr_cmd_par	B	No
A17	cmd<2>	B	No	A18	addr<24>	B	No
A19	Ground	P	—	A20	Vdd	P	—
A21	addr<30>	B	No	A22	ad<49>	B	Yes
A23	addr<35>	B	No	A24	addr<37>	B	No
A25	Ground	P	—	A26	Vdd	P	—
B1	Ground	P	—	B2	ad<47>	B	Yes
B3	ad<46>	B	Yes	B4	addr<15>	B	No
B5	addr<18>	B	No	B6	frame_1	B	Yes
B7	addr<23>	B	No	B8	victim_pending	I	No
B9	tag_ctl_par	O	Yes	B10	irdy_1	B	Yes
B11	fill_error	O	No	B12	req_1	O	Yes
B13	res<0>	I	No	B14	res<1>	I	No
B15	cmd<0>	B	No	B16	cbe_l<2>	B	Yes

(continued on next page)

21172–CA Pin Descriptions

2.2 21172–CA Signal Descriptions

Table 2–10 (Cont.) 21172–CA Pin Assignment (Alphanumeric)

Pin Number	Pin Name	Type	5-V Tolerant	Pin Number	Pin Name	Type	5-V Tolerant
B17	cbe_l<3>	B	Yes	B18	cbe_l<7>	B	Yes
B19	addr<27>	B	No	B20	addr<29>	B	No
B21	addr<32>	B	No	B22	addr<34>	B	No
B23	ad<19>	B	Yes	B24	addr<39>	B	No
B25	iod_e<7>	B	Yes	B26	Ground	P	—
C1	ad<45>	B	Yes	C2	addr<11>	B	No
C3	Ground	P	—	C4	ad<15>	B	Yes
C5	par64	B	Yes	C6	addr<20>	B	No
C7	perr_l	B	Yes	C8	stop_l	B	Yes
C9	lock_l	I	Yes	C10	ack64_l	B	Yes
C11	gnt_l	I	Yes	C12	addr_bus_req	B	No
C13	Spare	—	—	C14	cbe_l<1>	B	Yes
C15	cbe_l<5>	B	Yes	C16	cbe_l<6>	B	Yes
C17	addr<25>	B	No	C18	ad<16>	B	Yes
C19	ad<48>	B	Yes	C20	addr<31>	B	No
C21	ad<18>	B	Yes	C22	addr<36>	B	No
C23	addr<38>	B	No	C24	Ground	P	—
C25	ad<20>	B	Yes	C26	iod_e<6>	B	Yes
D1	ad<44>	B	Yes	D2	addr<10>	B	No
D3	addr<12>	B	No	D4	Vdd	P	—
D5	addr<17>	B	No	D6	addr<19>	B	No
D7	addr<22>	B	No	D8	devsel_l	B	Yes
D9	trdy_l	B	Yes	D10	idle_bc	O	No
D11	req64_l	B	Yes	D12	cbe_l<0>	B	Yes
D13	Vdd	P	—	D14	Ground	P	—
D15	cmd<1>	B	No	D16	cmd<3>	B	No
D17	addr<26>	B	No	D18	addr<28>	B	No
D19	ad<17>	B	Yes	D20	addr<33>	B	No
D21	ad<50>	B	Yes	D22	ad<51>	B	Yes
D23	Vdd	P	—	D24	ad<52>	B	Yes
D25	iod_e<5>	B	Yes	D26	ad<21>	B	Yes
E1	ad<43>	B	Yes	E2	ad<12>	B	Yes

(continued on next page)

21172–CA Pin Descriptions 2.2 21172–CA Signal Descriptions

Table 2–10 (Cont.) 21172–CA Pin Assignment (Alphanumeric)

Pin Number	Pin Name	Type	5-V Tolerant	Pin Number	Pin Name	Type	5-V Tolerant
E3	ad<13>	B	Yes	E4	ad<14>	B	Yes
E23	iod_e<4>	B	Yes	E24	iod<63>	B	Yes
E25	ad<53>	B	Yes	E26	iod<62>	B	Yes
F1	ad<42>	B	Yes	F2	addr<7>	B	No
F3	addr<9>	B	No	F4	addr<13>	B	No
F23	iod<61>	B	Yes	F24	ad<22>	B	Yes
F25	iod<60>	B	Yes	F26	ad<54>	B	Yes
G1	Vdd	P	—	G2	int4_valid<2>	I	No
G3	int4_valid<3>	I	No	G4	addr<8>	B	No
G23	iod<59>	B	Yes	G24	iod<58>	B	Yes
G25	ad<23>	B	Yes	G26	Vdd	P	—
H1	Ground	P	—	H2	int4_valid<1>	I	No
H3	addr<6>	B	No	H4	ad<11>	B	Yes
H23	ad<55>	B	Yes	H24	iod<56>	B	Yes
H25	iod<57>	B	Yes	H26	Ground	P	—
J1	int4_valid<0>	I	No	J2	ad<41>	B	Yes
J3	addr<5>	B	No	J4	ad<10>	B	Yes
J23	iod<54>	B	Yes	J24	ad<56>	B	Yes
J25	iod<55>	B	Yes	J26	ad<24>	B	Yes
K1	cack	O	No	K2	addr<4>	B	No
K3	ad<40>	B	Yes	K4	ad<9>	B	Yes
K23	ad<25>	B	Yes	K24	iod<51>	B	Yes
K25	iod<52>	B	Yes	K26	iod<53>	B	Yes
L1	ad<38>	B	Yes	L2	ad<39>	B	Yes
L3	dack	O	No	L4	ad<8>	B	Yes
L11	Ground	P	—	L12	Ground	P	—
L13	Ground	P	—	L14	Ground	P	—
L15	Ground	P	—	L16	Ground	P	—
L23	ad<26>	B	Yes	L24	iod<49>	B	Yes
L25	iod<50>	B	Yes	L26	ad<57>	B	Yes
M1	ad<6>	B	Yes	M2	mem_addr<11>	O	No
M3	mem_addr<12>	O	No	M4	ad<7>	B	Yes

(continued on next page)

21172–CA Pin Descriptions
2.2 21172–CA Signal Descriptions

Table 2–10 (Cont.) 21172–CA Pin Assignment (Alphanumeric)

Pin Number	Pin Name	Type	5-V Tolerant	Pin Number	Pin Name	Type	5-V Tolerant
M11	Ground	P	—	M12	Ground	P	—
M13	Ground	P	—	M14	Ground	P	—
M15	Ground	P	—	M16	Ground	P	—
M23	iod<47>	B	Yes	M24	ad<27>	B	Yes
M25	ad<58>	B	Yes	M26	iod<48>	B	Yes
N1	Vdd	P	—	N2	mem_addr<9>	O	No
N3	mem_addr<10>	O	No	N4	Ground	P	—
N11	Ground	P	—	N12	Ground	P	—
N13	Ground	P	—	N14	Ground	P	—
N15	Ground	P	—	N16	Ground	P	—
N23	Vdd	P	—	N24	iod<45>	B	Yes
N25	iod<46>	B	Yes	N26	Ground	P	—
P1	Ground	P	—	P2	mem_addr<8>	O	No
P3	Spare	—	—	P4	Vdd	P	—
P11	Ground	P	—	P12	Ground	P	—
P13	Ground	P	—	P14	Ground	P	—
P15	Ground	P	—	P16	Ground	P	—
P23	Ground	P	—	P24	iod<43>	B	Yes
P25	iod<44>	B	Yes	P26	Vdd	P	—
R1	mem_addr<6>	O	No	R2	ad<5>	B	Yes
R3	ad<37>	B	Yes	R4	mem_addr<7>	O	No
R11	Ground	P	—	R12	Ground	P	—
R13	Ground	P	—	R14	Ground	P	—
R15	Ground	P	—	R16	Ground	P	—
R23	iod<41>	B	Yes	R24	iod<42>	B	Yes
R25	ad<60>	B	Yes	R26	ad<59>	B	Yes
T1	ad<35>	B	Yes	T2	ad<4>	B	Yes
T3	mem_addr<5>	O	No	T4	ad<36>	B	Yes
T11	Ground	P	—	T12	Ground	P	—
T13	Ground	P	—	T14	Ground	P	—
T15	Ground	P	—	T16	Ground	P	—

(continued on next page)

21172–CA Pin Descriptions 2.2 21172–CA Signal Descriptions

Table 2–10 (Cont.) 21172–CA Pin Assignment (Alphanumeric)

Pin Number	Pin Name	Type	5-V Tolerant	Pin Number	Pin Name	Type	5-V Tolerant
T23	ad<61>	B	Yes	T24	iod<39>	B	Yes
T25	iod<40>	B	Yes	T26	ad<28>	B	Yes
U1	mem_addr<2>	O	No	U2	ad<3>	B	Yes
U3	mem_addr<4>	O	No	U4	mem_addr<3>	O	No
U23	ad<62>	B	Yes	U24	iod<37>	B	Yes
U25	ad<30>	B	Yes	U26	ad<29>	B	Yes
V1	mem_addr<0>	O	No	V2	ad<2>	B	Yes
V3	mem_addr<1>	O	No	V4	ad<34>	B	Yes
V23	iod<34>	B	Yes	V24	ad<31>	B	Yes
V25	iod<36>	B	Yes	V26	iod<38>	B	Yes
W1	Ground	P	—	W2	mem_ack_1	O	Yes
W3	mem_b0	O	No	W4	ad<33>	B	Yes
W23	debug	O	No	W24	iod<33>	B	Yes
W25	iod<35>	B	Yes	W26	Ground	P	—
Y1	Vdd	P	—	Y2	mem_req_1	I	Yes
Y3	ad<1>	B	Yes	Y4	mem_we_1<1>	O	No
Y23	pll_test<0>	I	Yes	Y24	pll_test<3>	I	Yes
Y25	ad<63>	B	Yes	Y26	Vdd	P	—
AA1	ad<32>	B	Yes	AA2	mem_we_1<0>	O	No
AA3	mem_cs_1	I	Yes	AA4	ad<0>	B	Yes
AA23	test_mode<1>	I	Yes	AA24	td	I	Yes
AA25	pll_test<2>	I	Yes	AA26	iod<32>	B	Yes
AB1	ras<3> ¹ or cas_1<7> ²	O	No	AB2	ras<2> ¹ or cas_1<6> ²	O	No
AB3	rst_1	O	Yes	AB4	ras<1> ¹ or cas_1<5> ²	O	No
AB23	pll_lp2	A	No	AB24	pll_vdd	P	No
AB25	test_mode<0>	I	Yes	AB26	pll_test<1>	I	Yes
AC1	ras<0> ¹ or cas_1<4> ²	O	No	AC2	cas<3> ¹ or cas_1<3> ²	O	No
AC3	cas<2> ¹ or cas_1<2> ²	O	No	AC4	Vdd	P	—

¹Signal name in set select memory mode.

²Signal name in alternate memory mode.

(continued on next page)

21172–CA Pin Descriptions

2.2 21172–CA Signal Descriptions

Table 2–10 (Cont.) 21172–CA Pin Assignment (Alphanumeric)

Pin Number	Pin Name	Type	5-V Tolerant	Pin Number	Pin Name	Type	5-V Tolerant
AC5	set_sel<13> ¹ or rasb_l<5> ²	O	No	AC6	set_sel<9> ¹ or rasb_l<1> ²	O	No
AC7	set_sel<5> ¹ or rasa_l<5> ²	O	No	AC8	set_sel<1> ¹ or rasa_l<1> ²	O	No
AC9	cmc<6>	O	No	AC10	cmc<2>	O	No
AC11	gru_ack	I	Yes	AC12	ioc<1>	O	No
AC13	Ground	P	—	AC14	Vdd	P	—
AC15	iod<4>	B	Yes	AC16	iod<8>	B	Yes
AC17	iod<12>	B	Yes	AC18	iod<16>	B	Yes
AC19	iod<20>	B	Yes	AC20	iod<24>	B	Yes
AC21	iod<28>	B	Yes	AC22	iod_e<0>	B	Yes
AC23	Vdd	P	—	AC24	aux_vdd	P	—
AC25	pll_vss	P	No	AC26	pll_clk	I	No
AD1	cas<1> ¹ or cas_l<1> ²	O	No	AD2	cas<0> ¹ or cas_l<0> ²	O	No
AD3	Ground	P	—	AD4	set_sel<14> ¹ or rasb_l<6> ²	O	No
AD5	set_sel<10> ¹ or rasb_l<2> ²	O	No	AD6	set_sel<6> ¹ or rasa_l<6> ²	O	No
AD7	set_sel<2> ¹ or rasa_l<2> ²	O	No	AD8	cmc<8>	O	No
AD9	cmc<5>	O	No	AD10	cmc<1>	O	No
AD11	gru_sel	O	Yes	AD12	ioc<2>	O	No
AD13	ioc<5>	O	No	AD14	iod<1>	B	Yes
AD15	iod<3>	B	Yes	AD16	iod<7>	B	Yes
AD17	iod<11>	B	Yes	AD18	iod<15>	B	Yes
AD19	iod<19>	B	Yes	AD20	iod<23>	B	Yes
AD21	iod<27>	B	Yes	AD22	iod<31>	B	Yes
AD23	iod_e<3>	B	Yes	AD24	Ground	P	—
AD25	aux_vss	P	—	AD26	pll_agnd	A	No
AE1	Ground	P	—	AE2	mem_en	O	No
AE3	set_sel<15> ¹ or rasb_l<7> ²	O	No	AE4	set_sel<11> ¹ or rasb_l<3> ²	O	No

¹Signal name in set select memory mode.

²Signal name in alternate memory mode.

(continued on next page)

21172–CA Pin Descriptions
2.2 21172–CA Signal Descriptions

Table 2–10 (Cont.) 21172–CA Pin Assignment (Alphanumeric)

Pin Number	Pin Name	Type	5-V Tolerant	Pin Number	Pin Name	Type	5-V Tolerant
AE5	set_sel<7> ¹ or rasa_l<7> ²	O	No	AE6	set_sel<3> ¹ or rasa_l<3> ²	O	No
AE7	cmc<7>	O	No	AE8	cmc<4>	O	No
AE9	cmc<3>	O	No	AE10	cmc<0>	O	No
AE11	int	O	Yes	AE12	ioc<3>	O	No
AE13	ioc<6>	O	No	AE14	iod<2>	B	Yes
AE15	iod<5>	B	Yes	AE16	iod<9>	B	Yes
AE17	iod<13>	B	Yes	AE18	iod<17>	B	Yes
AE19	iod<21>	B	Yes	AE20	iod<22>	B	Yes
AE21	iod<26>	B	Yes	AE22	iod<30>	B	Yes
AE23	iod_e<2>	B	Yes	AE24	sys_rst_l	I	Yes
AE25	iddt	I	No	AE26	Ground	P	—
AF1	Vdd	P	—	AF2	Ground	P	—
AF3	set_sel<12> ¹ or rasb_l<4> ²	O	No	AF4	set_sel<8> ¹ or rasb_l<0> ²	O	No
AF5	set_sel<4> ¹ or rasa_l<4> ²	O	No	AF6	set_sel<0> ¹ or rasa_l<0> ²	O	No
AF7	Vdd	P	—	AF8	Ground	P	—
AF9	error	O	Yes	AF10	ioc<0>	O	No
AF11	ioc<4>	O	No	AF12	iod<0>	B	Yes
AF13	Vdd	P	—	AF14	Ground	P	—
AF15	iod<6>	B	Yes	AF16	iod<10>	B	Yes
AF17	iod<14>	B	Yes	AF18	iod<18>	B	Yes
AF19	Ground	P	—	AF20	Vdd	P	—
AF21	iod<25>	B	Yes	AF22	iod<29>	B	Yes
AF23	iod_e<1>	B	Yes	AF24	test_out	O	No
AF25	Ground	P	—	AF26	Vdd	P	—

¹Signal name in set select memory mode.

²Signal name in alternate memory mode.

21172-CA Pin Descriptions

2.3 21172-CA Mechanical Specifications


2.3 21172-CA Mechanical Specifications

The 21172-CA is contained in a custom 388-pin encapsulated plastic ball grid array (EPBGA) package, shown in Figure 2-2.

Table 2-11 lists the 388-pin 21172-CA package dimensions in millimeters.

21172-CA Pin Descriptions
 2.3 21172-CA Mechanical Specifications

Figure 2-2 21172-CA 388-Pin EPBGA Package Dimensions



21172–CA Pin Descriptions
2.3 21172–CA Mechanical Specifications

Table 2–11 21172–CA 388-Pin EPBGA Package Dimensions

Symbol	Dimension	Value (mm)
A	Package overall height	2.30 minimum to 2.90 maximum (2.64 nominal)
A1	Package standoff height	0.50 minimum to 0.70 maximum (0.60 nominal)
A2	Encapsulated package thickness	1.00 minimum to 1.40 maximum (1.20 nominal)
A3	Printed circuit board thickness	0.54 ref ¹
b	Solder ball width	0.6 minimum to 0.9 maximum (0.75 nominal)
D	Package overall width	34.80 minimum to 35.20 maximum (35.00 nominal)
D1	Package width	31.75 BSC ²
D2	Encapsulation width	30.00 BSC ²
D3	Thermal ball/Vss location	6.35 BSC ²
E	Package overall length	34.80 minimum to 35.20 maximum (35.00 nominal)
E1	Package length	31.75 BSC ²
E2	Encapsulation length	30.00 BSC ²
E3	Thermal ball/Vss location	6.35 BSC ²
e	Solder ball pitch	1.27 BSC ²
aaa	Surface planarity	0.15 maximum
ddd	Encapsulation positional tolerance	0.30 maximum

¹The value for this measurement is for reference only.

²ANSI Y14.5M-1982 American National Standard Dimensioning and Tolerancing, Section 1.3.2, defines Basic Dimension (BSC) as: A numerical value used to describe the theoretically exact size, profile, orientation, or location of a feature or datum target. It is the basis from which permissible variations are established by tolerances on other dimensions, in notes, or in feature control frames.

Note

The drawing/dimensions are for reference only. Examples of board layout, including detailed engineering drawings and plot files, are available from Digital Semiconductor.

21172–CA Architecture Overview

This chapter describes the 21172–CA (CIA) data paths with functional units, transactions, and architecture.

3.1 CIA Data Paths with Functional Units Description

This section shows the CIA data path functional units and describes operation of the functional units.

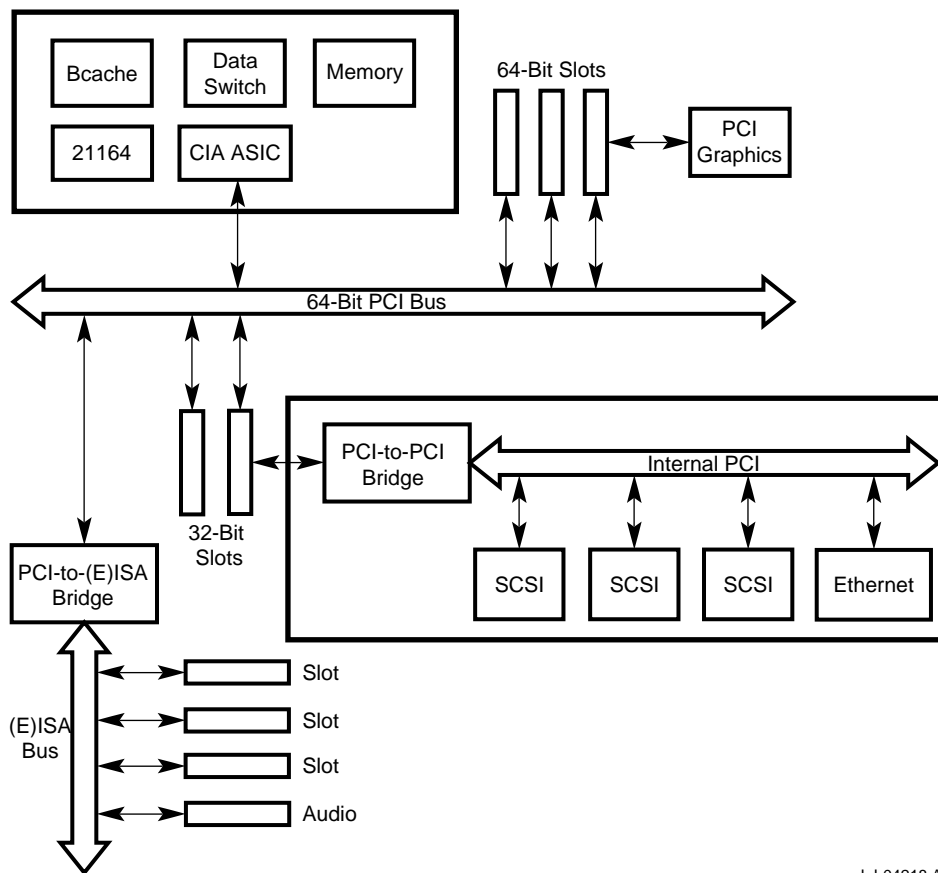
3.1.1 System Functional Units

The CIA operates as the host bridge on a PCI-based system along with the 21164, four DSWs, Bcache, memory arrays, and support logic. A simplified description of the system data paths and functional units is presented here to show where the CIA operates within the system.

The CIA performs control, I/O, and address tasks within the 21172 Core Logic Chipset, as shown in Figure 3–1. The host bridge module/chip is shown connected to the 64-bit PCI bus.

21172-CA Architecture Overview
3.1 CIA Data Paths with Functional Units Description

Figure 3-1 System Functional Block Diagram



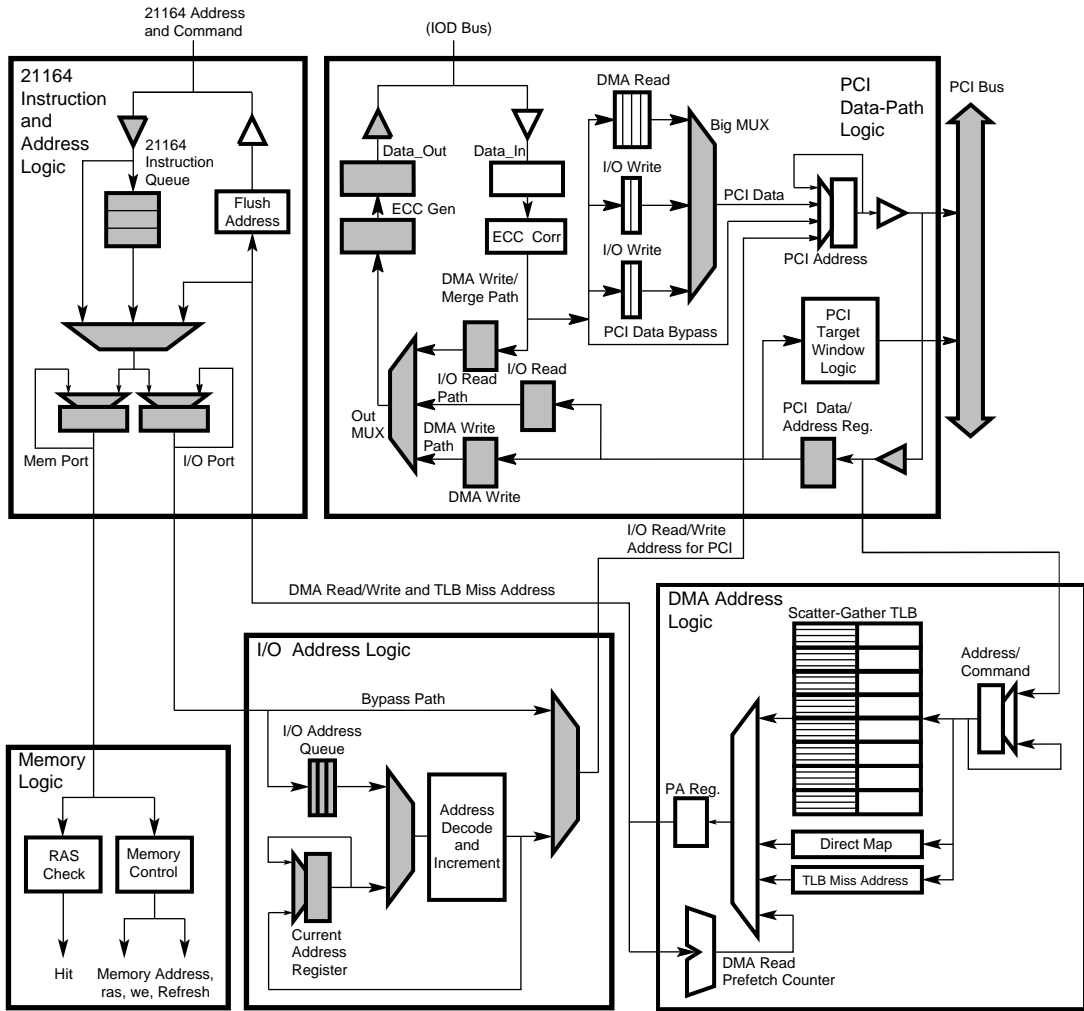
LJ-04218.AI

21172-CA Architecture Overview
3.1 CIA Data Paths with Functional Units Description

3.1.2 CIA Functional Units

Figure 3-2 shows the CIA data paths and functional units.

Figure 3-2 CIA Functional Block Diagram



LJ-04219.AI

21172–CA Architecture Overview

3.1 CIA Data Paths with Functional Units Description

The CIA functional units are listed here and are described in Section 3.1.2.1 through Section 3.1.2.5:

- 21164 instruction and address logic
- PCI data path logic
- Memory logic
- I/O address logic
- DMA address logic

3.1.2.1 21164 Instruction and Address Logic

The 21164 instruction and address logic accepts commands from the 21164 and directs the instruction to the memory port or the I/O port. The DMA read/write address (or the scatter-gather TLB miss servicing address) also has access to the memory port through a multiplexer.

The 21164 instruction queue is provided to capture 21164 commands should the memory or I/O port be busy. The 21164 can issue, at most, two read miss transactions—with one having a victim. This maximum condition would require holding three addresses and commands, so a 3-entry buffer is needed.

The flush address register is used during DMA read or write transactions to the Bcache for the latest data. The Bcache is implemented as a write invalidate cache and so might have the only valid copy of the required data.

3.1.2.2 PCI Data Path Logic

Because the CIA is the PCI host bridge it generates and decodes PCI addresses and also supplies and receives data. The data path to and from the PCI goes through the CIA and the DSW. The ECC generation and check logic is excluded from the DSW because of the bit slice nature of the DSW and so is performed by the CIA.

Separate buffers are provided in the CIA for DMA and I/O read/write transactions. These buffers are either 32 bytes or 64 bytes in size, and are shown partitioned into 16-byte units, which matches the width of the 21164 data bus.

Later, when the DSW is examined, it will be seen that the DSW also has buffers for the DMA and I/O paths. Control logic simplification is the reason for this duplication.

21172–CA Architecture Overview
3.1 CIA Data Paths with Functional Units Description

Table 3–1 lists the PCI commands that the CIA responds to and sends.

Table 3–1 PCI Commands Recognized by the CIA

PCI cbe_l<3:0>	Command	CIA Slave	CIA Master
0000	Interrupt acknowledge	No	Yes
0001	Special cycle	No	Yes
0010	I/O read	No	Yes
0011	I/O write	No	Yes
0100	Reserved	—	—
0101	Reserved	—	—
0110	Memory read	Yes	Yes
0111	Memory write	Yes	Yes
1000	Reserved	—	—
1001	Reserved	—	—
1010	Configuration read	No	Yes
1011	Configuration write	No	Yes
1100	Memory read multiple	Yes	No
1101	Dual address cycle	Yes	No
1110	Memory read line	Yes	No
1111	Memory write and invalidate	Yes ¹	No

¹Accepted as memory write

The CIA has these PCI features:

- Supports 64-bit PCI bus width
- Supports 64-bit PCI addressing (DAC cycles)
- Accepts PCI fast back-to-back cycles
- Issues PCI fast back-to-back cycles in dense address space

21172–CA Architecture Overview

3.1 CIA Data Paths with Functional Units Description

3.1.2.3 Memory Logic

Memory logic provides the row and column address for the memory banks and all the control signals (**cas**<3:0>, **cas_l**<7:4>, **ras**<3:0>, **rasx_l**<7:0>, **mem_we_l**<1:0>, and **mem_en**). The memory logic provides control signals to the DSW to inform it when to send/strobe the data to/from memory. The CIA also controls memory refresh.

3.1.2.4 I/O Address Logic

The I/O address logic handles the I/O read and write addresses. The address decode logic extracts the PCI address from the dense/sparse space 21164 address encodings (see Chapter 6 for more details).

This logic also increments the current address of each data cycle for two reasons:

- In case of a PCI retry, which will require the transaction to be resumed later, we need to start with the address of the aborted data.
- We need to provide a pointer to the next data item to be loaded/sent from the I/O read/write buffers.

The CIA can queue up to six I/O write transactions. Two of the I/O write transactions can be waiting in the 21164 queue and four I/O write transactions can be sitting in the I/O address queue (a 3-entry I/O address queue is provided together with a single-entry current address register). This allows six I/O write transactions to be outstanding (the DSW provides four corresponding I/O write data buffers). The bypass path is enabled/disabled by CIA_CTRL[CSR_IOA_BYPASS] (CSR_CTRL<15>) being equal to 0/1 respectively.

21172–CA Architecture Overview

3.1 CIA Data Paths with Functional Units Description

The CIA provides two more buffers that are needed to sustain maximum bandwidth for memory copy to I/O space. A bypass path is provided for certain dense-space I/O write transactions; that is, the valid bits for the first 16 bytes of data are *xxx1*.

3.1.2.5 DMA Address Logic

The DMA address logic converts PCI addresses to CPU memory addresses. Two conversion methods are provided:

- A direct path where a base offset is concatenated with the PCI address
- A scatter-gather map, which maps any 8KB PCI page to any 8KB memory space page

The scatter-gather TLB is eight entries deep with each entry holding four consecutive PTEs. A TLB miss is handled by hardware but software is required to invalidate stale entries by writing to the translation buffer invalidate register (TBIA). See Chapter 6 for more information on address mapping.

The output of the physical address register (PA) has a counter that generates the prefetch address for certain DMA read miss transactions. An 8KB detector prevents prefetching across page boundaries.

3.2 CIA Transactions

The CIA performs the transactions listed here:

- 21164 memory read miss
- 21164 memory read miss with victim
- 21164 I/O read
- 21164 I/O write
- DMA read
- DMA read (prefetch)
- DMA write

These transactions are described in detail in Section 3.2.1 through Section 3.2.8.

21172–CA Architecture Overview

3.2 CIA Transactions

3.2.1 21164 Read Miss Transaction

The 21164 read miss command and address are sent to the CIA. If the CIA is idle, the command will be stored directly in the memory port register; otherwise, the command enters the 3-deep CPU instruction queue and remains there until the memory port is free. The CIA can accept another subsequent read miss command from the 21164 and store it in the 21164 instruction queue.

The three data paths to the memory port have to arbitrate for use of the memory port: (1) the direct path and (2) the 21164 instruction queue previously mentioned and (3) the DMA read/write address path. The DMA read/write address path is also the path for scatter-gather TLB miss addresses. Upstream of the multiplexer all instruction ordering from the 21164 is preserved. Downstream of the multiplexer ordering between I/O write transactions and 21164 memory transactions is lost. This “post and run” I/O write coherency issue is discussed in Section 3.3.2.

When the 21164 read miss command enters the memory logic, the memory controller generates the appropriate address signals for the memory array. The memory controller then waits for the memory access delay before instructing the DSW to accept the memory data. The MEMDATA bus width is 128 or 256 bits (16 or 32 bytes), and so two or four memory cycles are required to access the 64-byte block.

The DSW synchronizes the data to the 21164 clock. The fixed 64-byte block size read data is returned to the 21164 in wrapped-order.

3.2.2 21164 Read Miss with Victim Transaction

The read miss with victim transaction is similar to the read miss transaction described in Section 3.2.1, except the 21164 will also use a Bcache victim command to send out the victim block. The victim block is the modified (dirty) block that is to be replaced in the Bcache by the read miss transaction data.

The read miss command is followed by the Bcache victim command and address. The read miss command will get to the memory port first and so the Bcache victim command and address are always written into the 21164 instruction queue. The victim block data is saved in the victim buffer in the DSW.

The CIA arbiter ensures that a read miss transaction and a victim write transaction are an atomic operation. The victim data is written to memory after the read data has been fetched from memory. The row address portion of the victim address is compared to the row address of the current read miss transaction. If they match, then no memory **ras** cycle is required; instead, only a **cas** cycle is needed. The address bits for the memory have been carefully interspersed to maximize the chance of a victim row address “hitting” the read

21172—CA Architecture Overview

3.2 CIA Transactions

address—this performance feature is transparent to software. The memory controller in the CIA generates the memory write pulses and instructs the DSW to send the victim data to memory.

The DSW victim buffer is invalidated if a DMA write transaction (or DMA read with lock transaction) “hits” in the victim buffer. Until the 21164 has its read miss transaction data returned, the victim block is still in the Bcache and is still “owned” by the 21164. It is possible for the read miss with victim transaction from the 21164 to be stalled behind a DMA write transaction. The DMA write transaction could “hit” the victim block – in this case, the DMA write transaction will cause a flush command to be issued to the 21164, which will result in the 21164 providing the victim data to the DSW and then the 21164 will invalidate the victim block in Bcache. Thus, the victim data waiting in the victim buffer is no longer valid and is marked invalid by logic in the CIA.

3.2.3 21164 Read Transaction to Noncacheable Space

A read transaction by the 21164 to noncacheable space can be to one of four address spaces:

- PCI memory dense or sparse space
- PCI I/O dense or sparse space
- PCI configuration space
- CIA control and status register (CSR) space

The read command to noncacheable space is accepted by the CIA like the read miss command, except the instruction goes to the I/O port. All read commands to noncacheable space go to the I/O address queue. There is no bypass path for dense space I/O read transactions.

Decode logic is provided on the output of the I/O address queue to extract the byte address for the PCI from the 21164’s sparse space encoding, and to decode the address for the correct region (PCI memory, I/O, configuration, or CSR).

A read command to noncacheable space destined for the PCI is described here because it is the more interesting case. Logic increments the current longword/quadword address stored in the current address register each data cycle. This is needed in case of a PCI retry and is also used to index the next data item in the I/O write/read data buffers. The value in the current address register corresponds to the address of the data item on the PCI bus (it has the same canonical time as the PCI data-out register).

21172–CA Architecture Overview

3.2 CIA Transactions

The I/O read command address is sent to the PCI after any prior I/O write transactions have completed (strict ordering is maintained). The PCI returns the requested data and places it in the I/O read buffer. The contents of this I/O read buffer are next copied to the I/O read buffer in the DSW and then are sent to the 21164. The I/O read buffers are replicated in the CIA for the following reasons:

- The path to the DSW might be busy with the end of a prior DMA write transaction (especially if the data has to be merged with memory data to build it up to the ECC width).
- At least 64 bits of storage are required for ECC because the PCI can return 32 bits of data at a time.
- It is convenient to handle transaction retries forced by the PCI protocol. The I/O read buffer in the DSW was provided to build the data up from 64 bits (IOD bus data lines between the CIA and DSW) to the 128 bits required by the 21164.

I/O read transactions from the 21164 have 8-byte resolution and the CIA always returns 32 bytes. If smaller resolution is required, sparse space addressing must be used. Sparse space addressing is described in Chapter 6.

Data is returned in the appropriate byte lanes.

3.2.4 21164 Write Transaction to Noncacheable Space

The 21164 issues write transactions to noncacheable space with longword resolution. For smaller granularity, sparse space addressing must be used (see Chapter 6).

Data for I/O write transactions is captured in the I/O write buffer in the DSW. Four 32-byte buffers are available in the DSW and a further two in the CIA. This number of entries allows the CIA to sustain maximum bandwidth on a large copy operation from memory through the 21164-to-I/O space.

The data from these I/O write buffers is sent to the two I/O write buffers in the CIA: two buffers are provided, allowing one to be emptied to the PCI bus while the other is being filled. Each 32-byte buffer in the CIA requires a separate PCI transaction (no merging of the write buffers occurs).

Another benefit of the CIA's I/O write buffers is simpler data-flow management when a target PCI device stalls the I/O write transactions.

The address for an I/O write transaction is sent to the CIA I/O port. For a 32-byte aligned dense space write transaction, the I/O write command is either sent directly to the PCI bus via the bypass path, or queued up in the I/O address queue. The direct (bypass) path is used if:

21172–CA Architecture Overview

3.2 CIA Transactions

- There are no outstanding I/O commands.
- The command is an I/O write command to dense space and the first 16 bytes are *xxx1*.

This command also goes into the I/O address queue, but only as a convenient path in the I/O addressing logic section, to get the address into the current address register.

A total of six I/O write addresses can be queued: the first in the current address register, two in the 21164 queue, and three in the I/O address queue. The I/O address queue maintains strict ordering for I/O transactions but does not maintain strict ordering of I/O write transactions relative to any memory read or write transactions (see Section 3.3).

3.2.5 DMA Transactions

The PCI address and command are captured in the address/command register and the data/address register. The address is compared against four address windows to determine if this PCI command should be accepted or ignored by the CIA.

Address windows are a requirement of the PCI specification and are software programmable—they are described in detail in Chapter 6. All PCI commands destined for memory are accepted by the CIA.

Note

Two registers capture incoming samples from the PCI bus.

- The address register captures the address and command, holding them for the duration of the transaction.
- The data-in register captures the data, but also captures the address.

The target window logic is attached to the data-in register. It is located there to capture a buffer of data (64 bytes) for use with a DMA write transaction scatter-gather TLB “miss.” The target window logic retries the master PCI device in case it has more data.

21172—CA Architecture Overview

3.2 CIA Transactions

The CIA has released the PCI for further transactions but is busy servicing the TLB “miss” so must hold the virtual address in the address register. Should another PCI transaction occur while the CIA is servicing the TLB “miss,” the CIA will accept that address and compare it to the target window.

There are three registers associated with each of the four PCI windows:

- Window base register 0–3 (Wn_BASE)—defines the start of the target window. $Wn_BASE[Wn_BASE_SG]$ determines if the scatter-gather map is used for the translation.
- Window mask register 0–3 (Wn_MASK)—defines the size of the window.
- Translated base register (Tn_BASE)—holds the base address used to relocate the PCI address in the 21164’s memory space (for direct mapping) and is also used to hold the scatter-gather map base address for scatter-gather mapping.

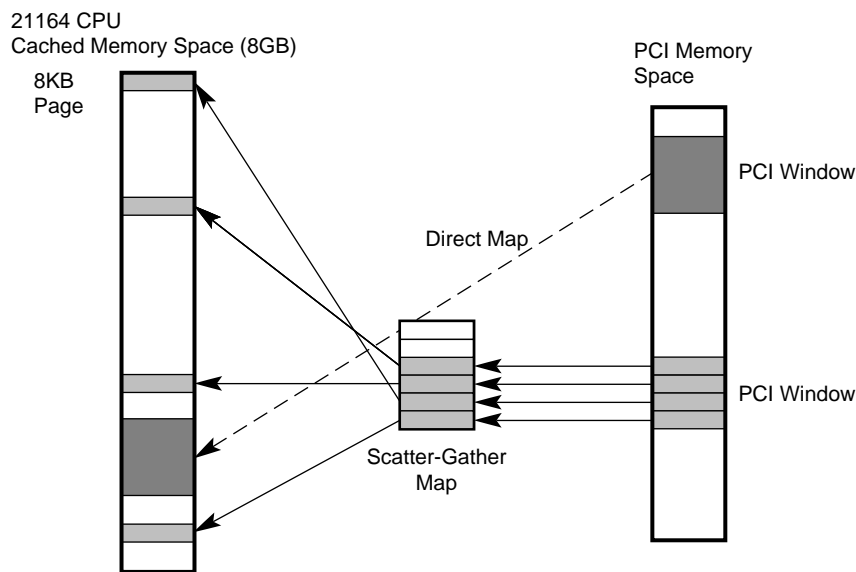
$Wn_BASE[Wn_BASE_SG]$ determines how the PCI address is translated as follows:

- If $Wn_BASE[Wn_BASE_SG]$ is clear, the address is mapped directly.
- If $Wn_BASE[Wn_BASE_SG]$ is set, the address is mapped through the scatter-gather table, allowing any 8KB PCI address to map to any 8KB memory address.

Figure 3–3 illustrates the mapping. This scatter-gather table is located in memory, but the CIA provides an 8-entry TLB that caches the most recent scatter-gather table entries. Each TLB entry holds four consecutive scatter-gather table entries, mapping a contiguous 32KB of virtual PCI addresses to any four 8KB memory pages.

21172-CA Architecture Overview 3.2 CIA Transactions

Figure 3-3 Address Space Overview



LJ04220A.AI5

3.2.6 DMA Read Transaction

The translated address stored in the PA register is sent to the 21164 through the flush/read address register and to memory through the memory port. If the 21164 data cache contains valid modified data, then a copy is sent to the Bcache-buffer part of one of the two DMA buffers in the DSW and the valid bit is set. Memory data is always fetched and put in the memory buffer of the DMA buffer. The PCI-buffer part of the DMA buffer is not used for DMA read transactions.

When the first valid QW is available in the DSW's DMA buffer, it is sent to the DMA read buffer in the CIA through the multiplexers. Any ECC correction is done in the CIA. From the DMA buffer in the CIA, the data goes from a register onto the PCI a cycle later. The DMA read data also takes the bypass path around the multiplexer for as long as the PCI can accept this stream of data. Once the PCI stalls, then the buffers are used.

21172–CA Architecture Overview

3.2 CIA Transactions

3.2.7 DMA Read Transaction (Prefetch)

The PCI supports three types of memory read commands, which are used as specified in the *PCI Local Bus Specification*:

- Read—used for small transfers (up to 1/2 a cache line)
- Read line—used for medium transfers (1/2 to 3 cache lines)
- Read multiple—used for large transfers (more than 3 cache lines)

In the PCI environment, most devices assume a processor cache of 32 bytes (compared to the CIA cache line of 64 bytes). The CIA prefetch strategy is shown in Table 3–2.

Table 3–2 CIA Prefetch Strategy

PCI Memory Read Command	Prefetch Operation
Read	None.
Read line	Prefetch one block.
Read multiple	Prefetch until end of transaction.

A counter is used to increment the memory block address for prefetching. The output of this counter goes to the Bcache and memory like a normal DMA read command address. The returned data is sent to the next free DMA buffer in the DSW. So as one buffer is being copied down to the CIA, the other is free to accept DMA prefetched data.

At the end of a transaction, all prefetched data in the DMA read buffers is ignored (that is, no prefetch caching is performed). Prefetching does not continue over 8KB page boundaries.

The 64-byte DMA read buffer in the CIA acts like two 32-byte halves during DMA commands—as one half is emptying to the PCI bus, the other half is being filled.

3.2.8 DMA Write Transaction

The translated address stored in the PA register is sent to the Bcache through the flush/read address register and also to memory via the memory port. If the Bcache contains valid (modified) data, the data is sent to the Bcache-buffer portion of the DMA buffer in the DSW and the Bcache data is invalidated. Memory data is always fetched and is put in the memory-buffer part of the DMA buffer.

21172–CA Architecture Overview

3.2 CIA Transactions

The DMA write data is taken from the PCI bus and is placed in the DMA write buffer. If the data is a complete quadword, then ECC is generated and the data is sent to the PCI-write part of the DMA buffer in the DSW. If the data is an incomplete quadword, then a merge operation has to be performed. The valid Bcache or memory quadword is sent to the CIA from the DMA buffer in the DSW, and the valid bytes of the DMA write data are merged. This merged quadword then has ECC generated and is sent to the PCI-write buffer portion of the DMA buffer in the DSW.

The DMA buffer in the DSW builds the data up to 64 bytes (the block size) before sending the data to memory. The memory logic generates the memory write pulses. The memory address is read from the PA register by way of a multiplexer.

Note

If the I/O port is busy with an I/O write/read transaction, the memory port will be available for DMA write commands.

One reason for providing a copy of the DMA write buffer in the CIA is that the data path to the DSW may not be available at the start of a DMA write transaction. It could be busy transferring I/O write transaction data from the DSW to the I/O write buffers if a 21164 I/O write command had been received concurrently with a PCI DMA write command.

DMA write transaction data is always written to memory, and is never stored (cached) in the DMA write buffers across transactions.

3.3 System Data Coherency Issues

The 21164 is the microprocessor in systems using the 21172 core logic chipset. The data coherency requirements of the 21164 dominate CIA data coherency issues. The CIA supports the flush-based data coherency protocol with a Bcache but with no duplicate tag stores for either the Scache or the Bcache. Therefore, the CIA uses the read and flush transactions to probe the 21164 cache.

21172–CA Architecture Overview

3.3 System Data Coherency Issues

3.3.1 Basic Properties of the System

The system has these basic properties:

1. The memory arrays do only one thing at a time.
2. A read or write transaction to memory is atomic. No operation of the memory arrays can occur between the read and write parts of a read-modify-write sequence.
3. The memory array sequence for a 21164 cache “miss” with victim transaction is atomic in the memory array. A cache miss with victim transaction causes the 21164 to perform a read miss transaction for the block required for the cache miss transaction and then a Bcache victim transaction to write the victim data to memory. No other memory access can come between the read miss transaction and the Bcache victim transaction.
4. When “transaction A” gains access to memory it is said to “own” memory. Once “transaction A” “owns” memory there will be no access for any other transaction until “transaction A” has completed all of its memory activity.
5. When an I/O TLB “miss” occurs, the CIA issues a read command to the 21164 for the TLB entry and a memory access is also made for this data. The read command is not issued to the 21164 until memory is owned for this access.
6. The CIA processes a PCI device read transaction to memory by issuing a read command to the 21164 for the data and making a memory access for the required data. The read command is not issued until memory is “owned” for the memory cycle.
7. A device write transaction to memory causes the CIA to issue a flush transaction to the 21164 for the data and to start a memory access for the data. The memory access serves to write data obtained from the 21164 and also serves to write the data from the PCI device to memory. The CIA does not issue the flush command until memory is “owned” for the memory write transaction.
8. To conclude, the 2-step sequences described in 5, 6, and 7 are all mutually atomic. In each case, the 2-step sequence is a probe of 21164 and a memory cycle. For any pair of sequences of the types describe in 5, 6, and 7, both steps of one of the sequences happened entirely before either of the two steps of the other sequences.

21172–CA Architecture Overview
3.3 System Data Coherency Issues

3.3.2 Post and Run I/O Write Data Coherency

The Alpha architecture allows “posted” or “buffered” write transactions to I/O devices, allowing the CIA to buffer write transactions from the 21164 to a PCI device. Several read transactions from a PCI device to main memory might be completed while write transactions to the same PCI device are buffered. The usual way for a software programmer to ensure that the write data has reached the device is to read a register on the same device.

The CIA performs the following operations to implement the Alpha architecture post and run feature.

- Before a read transaction to an I/O location is processed, all preceding 21164 write transactions to devices are flushed out of the buffers.
- Before a read transaction from a PCI device to main memory is processed, all preceding write transactions by the PCI device to main memory are flushed out of the write buffers.

3.4 CIA Memory *ras*, *cas*, and Address Operation

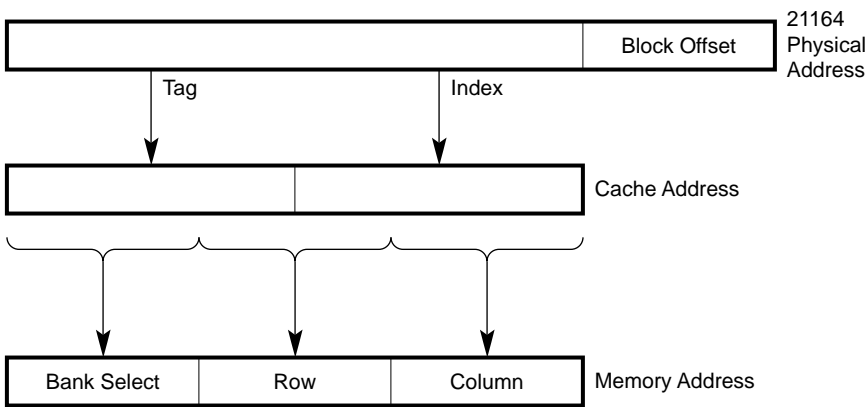
The technique used by the CIA to eliminate **ras** cycles for victim write transactions is described here. The method is transparent to software.

3.4.1 Traditional Mapping of 21164 Address to Memory Address

Figure 3–4 shows the traditional approach of mapping the 21164 address bit for bit to the memory address.

21172-CA Architecture Overview
3.4 CIA Memory ras, cas, and Address Operation

Figure 3-4 Traditional Mapping to a Memory Address

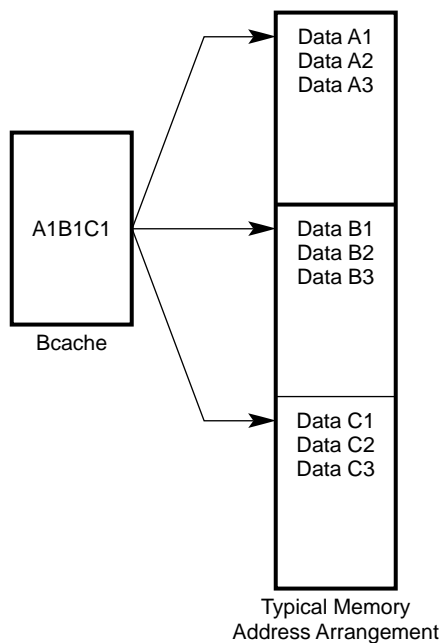


LJ-04221.AI

Because a cache is much smaller than available memory, multiple memory locations will correspond to the same cache location. This correspondence is shown in the left-hand side of Figure 3-5 where data A1, data B1, and data C1 all map to the same cache location (A1B1C1).

21172–CA Architecture Overview
3.4 CIA Memory ras, cas, and Address Operation

Figure 3–5 Traditional Victim Cache to Memory Correspondence



LJ-04222.AI

Suppose that data A1 currently resides in the Bcache, but the 21164 wants to access data C1. A cache miss will occur that will fetch data C1 and overwrite A1 in the cache. If A1 is the only valid copy of the data (for example, a 21164 write transaction had previously updated A1), then A1 has to be written to memory before data C1 can be written over data A1. The displaced data A1 is referred to as a victim.

Look again at Figure 3–5 and notice that in the traditional memory addressing scheme, data C1 and the potential victim (data A1) are well separated in memory. They are separated by a multiple of n times the cache size. This sparse distribution of potential victim locations means a read fill block and its victim may not reside in the same memory page.

21172-CA Architecture Overview

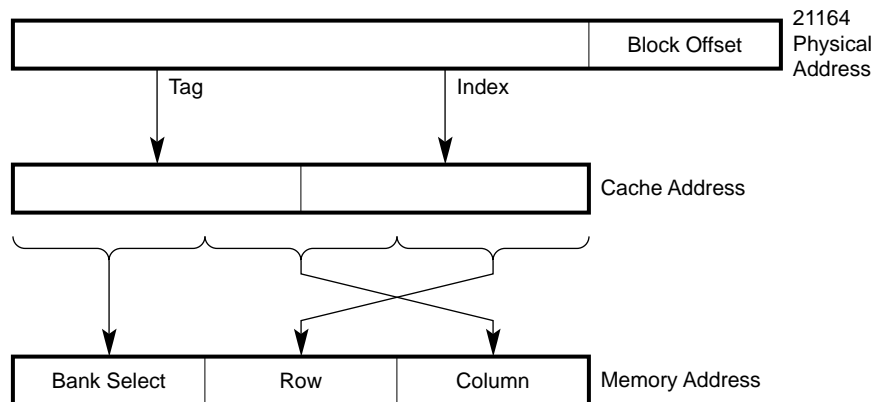
3.4 CIA Memory ras, cas, and Address Operation

3.4.2 CIA Mapping of 21164 Address to Memory Address

The highest tag bits are used for memory bank selection in both the traditional and CIA address mapping schemes.

The CIA shuffles the memory address bits so that some of the high-order 21164 address bits (part of the Bcache tag) become the memory column address and the low-order CPU address bits (part of the cache index) become the memory row address, as shown in Figure 3-6.

Figure 3-6 Modified Mapping to a Memory Address

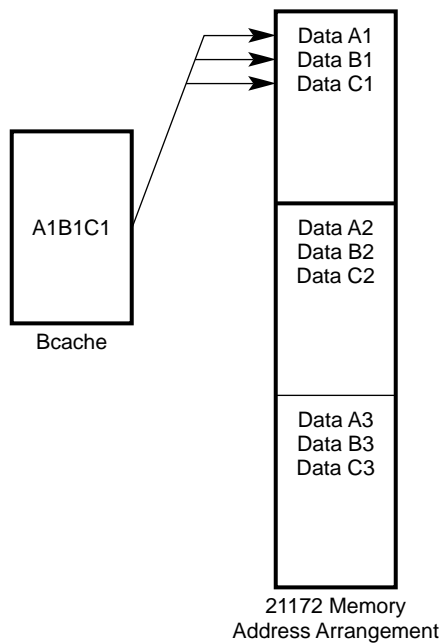


LJ-04223.AI

Interchanging the memory row and column address bits, the write transaction target and the potential victim locations reside in the same SIMMs, even within the same row within the SIMMs, as shown in Figure 3-7. The chance of a cache block and its victim residing in the same memory page is greatly increased. This technique will remove all **ras** cycles from victim write transactions. In fact, all victims will be able to share a **ras** signal with the read miss transaction.

21172-CA Architecture Overview
3.4 CIA Memory ras, cas, and Address Operation

Figure 3-7 CIA Victim Cache to Memory Correspondence



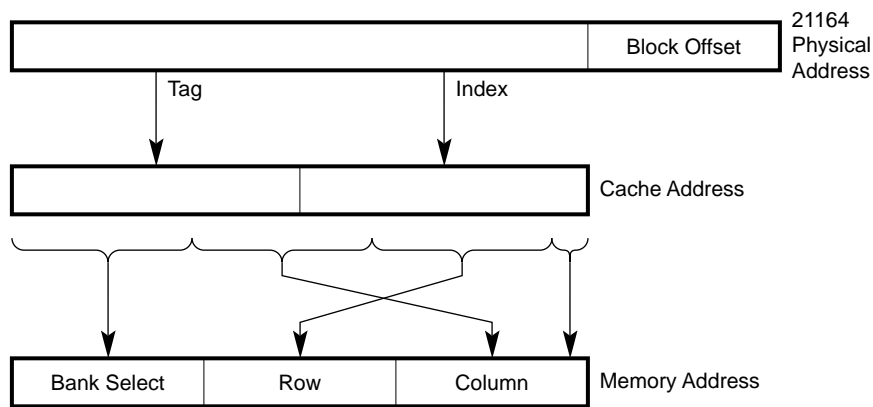
LJ-04224.AI

However, this approach hinders DMA transactions. Consecutive blocks (data A1, data A2, and data A3), as shown in Figure 3-7, are scattered a page apart by using this memory addressing scheme. This implies that consecutive DMA read/write transaction blocks will require a **ras** cycle, reducing possible bandwidth. The traditional scheme does not suffer this problem.

A compromise between the two schemes, as shown in Figure 3-8, is used by the CIA. The 4 low-order bits of the index map straight to the low-order bits of the column address (just like in a traditional scheme). The remaining high-order index bits go to the memory row-address. The remainder of the column address uses the tag portion of the physical address.

21172-CA Architecture Overview
 3.4 CIA Memory ras, cas, and Address Operation

Figure 3-8 CIA Mapping to a Memory Address



LJ-04225.AI

Using this compromise scheme, a DMA transaction will march through four 64-byte blocks (64 longwords), on average, before requiring a **ras** cycle. This constitutes a large DMA transfer, minimizing the **ras** penalty.

The logic that determines if a **ras** cycle is required is located in the memory logic block of Figure 3-2.

Figures 3-9 and 3-10 show the relationship between CPU/DMA address bits <28:04> and column and row address bits to memory. The memory address maps for the 128-bit data path and the 256-bit data path are shown in Tables 3-9 and 3-10, respectively. The user specifies row type in $MBA_n[ROW_TYPE]$ ($MBA_n<3:1>$) as described in Section 4.6.2.

Table 3-3 lists the row and column map for 128-bit and 256-bit banks.

21172–CA Architecture Overview
3.4 CIA Memory ras, cas, and Address Operation

Table 3–3 Row/Column Map

ROW_TYPE<2:0>	Row × Column	Bank Size	
		128-Bit	256-Bit
0 0 0	10 × 10	16MB	32MB
0 0 1	11 × 11 or 12 × 10	64MB	128MB
0 1 0	12 × 12 or 13 × 11	256MB	512MB
0 1 1	10 × 10, 11 × 9, or 12 × 8	16MB	32MB
1 0 0	10 × 9	NA	16MB
1 0 1	11 × 10	32MB	64MB
1 1 0	12 × 11	128MB	256MB
1 1 1	13 × 12	512MB	NA

21172-CA Architecture Overview
3.4 CIA Memory ras, cas, and Address Operation

Figure 3-10 Memory Address Maps for 256-Bit Data Path

Row Type	DRAM Size	DRAM Configuration	CPU/DMA Address Bits <28:4>															CPU Tag Index of Minimum Cache Size of 2MB (15 Address Lines Required)														
			28	27	26	25	24	23	22	21	20	19	18	17	16	15	14	13	12	11	10	9	8	7	6	5	4					
000	1M X 4 X 16	10R X 10C	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-					
001	4M X 4 X 16	11R X 11C 12R X 10C	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-					
010	16M X 4 X 16	12R X 12C 13R X 11C	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-					
011	1M X 4 X 16	10R X 10C 12R X 8C	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-					
100	512K X 8	10R X 9C	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-					
101	2M X 8	11R X 10C	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-					
110	8M X 8	12R X 11C	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-					
111	32M X 8	13R X 12C	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-					

MK11455-23

3.5 MB and LOCK Instructions

This section describes the MB (memory barrier) and LOCK instructions.

3.5.1 MB Instruction

The MB instruction may or may not be disabled from leaving the 21164 by clearing BC_CONTROL[EI_CMD_GRP3]. If BC_CONTROL[EI_CMD_GRP3] is clear, the CIA should never acknowledge receipt of an MB instruction.

CACK_EN[MB_ENABLE], in the CIA, enables or disables the CIA from acknowledging receipt of an MB instruction from the 21164. If CACK_EN[MB_ENABLE] is clear, the CIA will treat an MB instruction as a NOP. See Section 3.3.2.

3.5.2 LOCK Instruction

The CIA has to contend with locks originating on either the 21164 or the PCI. There are two cases that require that the 21164 internal lock state on a block be cleared (from a system point of view). The first is during a DMA write transaction to the locked block and the second is when a PCI device has established a lock.

The things to consider for system lock behavior are:

- DMA write transaction case—The 21164 maintains its lock flag and lock address in its bus interface unit (BIU). This lock status is correctly maintained, even if the locked block is displaced from the Scache, as long as the system sends all DMA write FLUSH commands to the 21164 (no duplicate cache tag).
- PCI lock established—A PCI device can only establish a lock on a 16-byte block of data during a DMA read transaction. Once the lock is established by the successful DMA read transaction, the PCI device has exclusive access to that 16-byte block (the PCI target may lock a larger block). This PCI lock allows the device to do an atomic read-modify-write on a 16-byte data unit.

The Alpha and PCI architectures have different mechanisms for establishing a lock:

- For PCI architecture, it is a successful read transaction (with the PCI LOCK signal asserted) that establishes a lock.
- For Alpha architecture, it is a successful write transaction (ST_n_C) that establishes a lock.

21172–CA Architecture Overview 3.5 MB and LOCK Instructions

So, a PCI read transaction with the PCI LOCK signal asserted must be treated by the 21164 as a write transaction (that is, for the Alpha architecture, it must have the same effect as a write transaction from some other processor). Example 3–1 shows an example of the 21164 and PCI competing for a memory flag.

Example 3–1 21164 and PCI Device Lock Contention

21164 Atomic Update -----	PCI Device Atomic Update -----
<pre>try_again: LDQ_L R1, mem_flag { modify memory flag } STQ_C R1, mem_flag BEQ R,no_store ----- -----</pre>	<pre>READ/LOCK mem_flag { modify mem_flag } WRITE/UNLOCK mem_flag</pre>
<pre>no_store: { check for excessive iterations } BR try_again</pre>	

Both the 21164 and the PCI device are attempting to perform an atomic read-modify-write operation on a memory flag (a semaphore). The 21164 begins with a LDQ_L instruction but will not know if the update is successful until the companion STQ_C instruction completes successfully.

Assume that the PCI device performs a PCI read transaction with the PCI LOCK signal asserted, and obtains the lock to the memory flag before the 21164 STQ_C instruction is executed. The 21164's STQ_C instruction must fail, otherwise, the 21164 and the PCI device will both believe they have successfully modified the flag.

Before a PCI device establishes the PCI lock, the 21164 must clear its lock flag. The CIA must perform a flush transaction to the locked block, causing the 21164 to clear its lock flag. Unfortunately, the system must not only perform a DMA read transaction, but must also write the flushed block to memory. Once the CIA successfully sends the data to the PCI device, the PCI lock is established.

21172—CA Architecture Overview

3.5 MB and LOCK Instructions

While the PCI lock is in effect, the 21164 may attempt to refill the flushed block. This will happen if a loop, as shown in Example 3–1, is repeatedly executing an `LD x _L` instruction. The requested block cannot be provided until the PCI lock has completed because once the block is in the 21164 internal cache, the CIA loses control of the block.

Initially, the CIA could prevent the 21164 from setting its internal lock flag for the block by using the 21164 input signal line **system_lock_flag_h**. The next time the `LD x _L/ST n _C` loop, the 21164 will have a cache hit on the block and will be unaware of the system lock state, so the loop will succeed this second time.

The CIA lock rules are:

- The 21164 **system_lock_flag_h** input signal line is not connected to the CIA and is always tied true.
- All DMA write transactions will result in the CIA performing a flush transaction to the 21164.
- All DMA read transactions with the PCI LOCK signal asserted will result in the CIA performing a flush transaction to the 21164. The CIA must write the flushed block to memory. The internal CIA lock address register will be set with the locking address.
- If the 21164 requests a fill that hits the CIA lock address register, then the fill is stalled until the PCI lock state is relinquished by the PCI device.
- The 21164 lock transaction will be treated as a NOP for systems with duplicate tag stores.
- The 21164 write transaction with lock will be treated as a write transaction.

3.5.3 Locks to Uncached Space

The CIA does not support `LD x _L` and `ST n _C` to noncacheable space. They are converted to `LD x` and `ST n` .

4

21172–CA Control and Status Registers

This chapter describes the 21172–CA (CIA) control and status registers (CSRs). It also provides information about programming memory timing, configuring memory, and initializing the backup cache (Bcache).

Note

Programmers must write only zeros to register fields defined as reserved.

4.1 21172–CA Registers

The 21172–CA CSRs are 32 bits wide and are located on NATURALLY ALIGNED 64-byte addresses. Write transactions to read-only registers could result in UNPREDICTABLE behavior while read transactions are nondestructive. Only zeros should be written to reserved bits. The register descriptions contain the initialized states.

The 21172–CA CSRs are presented here in seven groups:

- General registers
- Diagnostic registers
- Performance monitor registers
- Error registers
- System configuration registers
- PCI address and scatter-gather registers
- Address translation registers

21172–CA Control and Status Registers

4.1 21172–CA Registers

The 64-bit wide 21164 Alpha microprocessor CSRs are described in the *Digital Semiconductor 21164 (366 MHz Through 433 MHz) Alpha Microprocessor Hardware Reference Manual*. The CSRs of most interest to uniprocessor system designers are:

- Scache control register (SC_CTL)
- Bcache control register (BC_CONTROL)
- Bcache configuration register (BC_CONFIG)

4.1.1 21172–CA General Registers

Table 4–1 lists the 21172–CA (CIA) general registers and Section 4.2 shows and describes the individual registers.

Table 4–1 21172–CA General Registers

Address	Mnemonic	Name
87.4000.0080	CIA_REV	CIA revision register
87.4000.00C0	PCI_LAT	PCI latency register
87.4000.0100	CIA_CTRL	CIA control register
87.4000.0140	CIA_CNFG	CIA configuration register
87.4000.0400	HAE_MEM	Hardware address extension register
87.4000.0440	HAE_IO	Hardware address extension I/O register
87.4000.0480	CFG	Configuration register
87.4000.0600	CACK_EN	CIA acknowledgment enable register

4.1.2 21172–CA Diagnostic Registers

Table 4–2 lists the 21172–CA (CIA) diagnostic registers and Section 4.3 shows and describes the individual registers.

Table 4–2 21172–CA Diagnostic Registers

Address	Mnemonic	Name
87.4000.2000	CIA_DIAG	CIA diagnostic control register
87.4000.3000	DIAG_CHECK	Diagnostic check register

21172–CA Control and Status Registers

4.1 21172–CA Registers

4.1.3 21172–CA Performance Monitor Registers

Table 4–3 lists the 21172–CA (CIA) performance monitor registers and Section 4.4 shows and describes the individual registers.

Table 4–3 21172–CA Performance Monitor Registers

Address	Mnemonic	Name
87.4000.4000	PERF_MONITOR	Performance monitor register
87.4000.4040	PERF_CONTROL	Performance control register

4.1.4 21172–CA Error Registers

Table 4–4 lists the 21172–CA (CIA) error registers and Section 4.5 shows and describes the individual registers.

Table 4–4 21172–CA Error Registers

Address	Mnemonic	Name
87.4000.8000	CPU_ERR0	CPU error information register 0
87.4000.8040	CPU_ERR1	CPU error information register 1
87.4000.8200	CIA_ERR	CIA error register
87.4000.8240	CIA_STAT	CIA status register
87.4000.8280	ERR_MASK	CIA error mask register
87.4000.8300	CIA_SYN	CIA syndrome register
87.4000.8400	MEM_ERR0	CIA memory port status register 0
87.4000.8440	MEM_ERR1	CIA memory port status register 1
87.4000.8800	PCI_ERR0	PCI error register 0
87.4000.8840	PCI_ERR1	PCI error register 1
87.4000.8880	PCI_ERR2	PCI error register 2

21172–CA Control and Status Registers

4.1 21172–CA Registers

4.1.5 21172–CA System Configuration Registers

Table 4–5 lists the 21172–CA (CIA) system configuration registers and Section 4.6 shows and describes the individual registers.

Table 4–5 21172–CA System Configuration Registers

Address	Mnemonic	Name
87.5000.0000	MCR	Memory configuration register
87.5000.0600	MBA0	Memory base address register 0
87.5000.0680	MBA2	Memory base address register 2
87.5000.0700	MBA4	Memory base address register 4
87.5000.0780	MBA6	Memory base address register 6
87.5000.0800	MBA8	Memory base address register 8
87.5000.0880	MBAA	Memory base address register 10
87.5000.0900	MBAC	Memory base address register 12
87.5000.0980	MBAE	Memory base address register 14
87.5000.0B00	TMG0	Memory timing information register 0
87.5000.0B40	TMG1	Memory timing information register 1
87.5000.0B80	TMG2	Memory timing information register 2

21172–CA Control and Status Registers
4.1 21172–CA Registers

4.1.6 21172–CA PCI Address and Scatter-Gather Registers

Table 4–6 lists the 21172–CA (CIA) PCI address and scatter-gather registers, and Section 4.7 describes the individual registers.

Table 4–6 21172–CA PCI Address and Scatter-Gather Registers

Address	Mnemonic	Name
87.6000.0100	TBIA	Scatter-gather translation buffer invalidate register
87.6000.0400	W0_BASE	Window base register 0
87.6000.0440	W0_MASK	Window mask register 0
87.6000.0480	T0_BASE	Translated base register 0
87.6000.0500	W1_BASE	Window base register 1
87.6000.0540	W1_MASK	Window mask register 1
87.6000.0580	T1_BASE	Translated base register 1
87.6000.0600	W2_BASE	Window base register 2
87.6000.0640	W2_MASK	Window mask register 2
87.6000.0680	T2_BASE	Translated base register 2
87.6000.0700	W3_BASE	Window base register 3
87.6000.0740	W3_MASK	Window mask register 3
87.6000.0780	T3_BASE	Translated base register 3
87.6000.07C0	W_DAC	Window base dual address control register

21172–CA Control and Status Registers

4.1 21172–CA Registers

4.1.7 21172–CA Address Translation Registers

Table 4–7 lists the 21172–CA (CIA) address translation registers and Section 4.8 describes the individual registers.

Table 4–7 21172–CA Address Translation Registers

Address	Mnemonic	Name
87.6000.0800	LTB_TAG0	Lockable translation buffer tag 0
87.6000.0840	LTB_TAG1	Lockable translation buffer tag 1
87.6000.0880	LTB_TAG2	Lockable translation buffer tag 2
87.6000.08C0	LTB_TAG3	Lockable translation buffer tag 3
87.6000.0900	TB_TAG0	Translation buffer tag 0
87.6000.0940	TB_TAG1	Translation buffer tag 1
87.6000.0980	TB_TAG2	Translation buffer tag 2
87.6000.09C0	TB_TAG3	Translation buffer tag 3
87.6000.1000	TB0_PAGE0	Translation buffer 0 page 0
87.6000.1040	TB0_PAGE1	Translation buffer 0 page 1
87.6000.1080	TB0_PAGE2	Translation buffer 0 page 2
87.6000.10C0	TB0_PAGE3	Translation buffer 0 page 3
87.6000.1100	TB1_PAGE0	Translation buffer 1 page 0
87.6000.1140	TB1_PAGE1	Translation buffer 1 page 1
87.6000.1180	TB1_PAGE2	Translation buffer 1 page 2
87.6000.11C0	TB1_PAGE3	Translation buffer 1 page 3
87.6000.1200	TB2_PAGE0	Translation buffer 2 page 0
87.6000.1240	TB2_PAGE1	Translation buffer 2 page 1
87.6000.1280	TB2_PAGE2	Translation buffer 2 page 2
87.6000.12C0	TB2_PAGE3	Translation buffer 2 page 3

(continued on next page)

21172–CA Control and Status Registers
4.1 21172–CA Registers

Table 4–7 (Cont.) 21172–CA Address Translation Registers

Address	Mnemonic	Name
87.6000.1300	TB3_PAGE0	Translation buffer 3 page 0
87.6000.1340	TB3_PAGE1	Translation buffer 3 page 1
87.6000.1380	TB3_PAGE2	Translation buffer 3 page 2
87.6000.13C0	TB3_PAGE3	Translation buffer 3 page 3
87.6000.1400	TB4_PAGE0	Translation buffer 4 page 0
87.6000.1440	TB4_PAGE1	Translation buffer 4 page 1
87.6000.1480	TB4_PAGE2	Translation buffer 4 page 2
87.6000.14C0	TB4_PAGE3	Translation buffer 4 page 3
87.6000.1500	TB5_PAGE0	Translation buffer 5 page 0
87.6000.1540	TB5_PAGE1	Translation buffer 5 page 1
87.6000.1580	TB5_PAGE2	Translation buffer 5 page 2
87.6000.15C0	TB5_PAGE3	Translation buffer 5 page 3
87.6000.1600	TB6_PAGE0	Translation buffer 6 page 0
87.6000.1640	TB6_PAGE1	Translation buffer 6 page 1
87.6000.1680	TB6_PAGE2	Translation buffer 6 page 2
87.6000.16C0	TB6_PAGE3	Translation buffer 6 page 3
87.6000.1700	TB7_PAGE0	Translation buffer 7 page 0
87.6000.1740	TB7_PAGE1	Translation buffer 7 page 1
87.6000.1780	TB7_PAGE2	Translation buffer 7 page 2
87.6000.17C0	TB7_PAGE3	Translation buffer 7 page 3


21172-CA Control and Status Registers
4.2 21172-CA General Registers

4.2 21172-CA General Registers

Sections 4.2.1 through 4.2.8 define and describe the 21172-CA general registers.

4.2.1 CIA Revision Register

Register mnemonic: CIA_REV
 Register address: 87.4000.0080




LJ04226A.AI5

Field	Name	Type	Description
<07:00>	CIA_REV	RO,CIA Rev	CIA revision field. 2—Pass 3 CIA chip present. Note: Values 0 and 1, for Pass 1 and 2 CIA chips respectively, do not apply to the 21172-CA.
<08>	ALT_MEM	RO,0	Alternate memory mode 0—High-asserted ras and cas , SET SELECTS. 1—Low-asserted ras and cas .
<31:09>	Reserved	RO,0	—

21172–CA Control and Status Registers
4.2 21172–CA General Registers

4.2.2 PCI Latency Register

Register mnemonic: PCI_LAT
 Register address: 87.4000.00C0



LJ04227A.AI5

Field	Name	Type	Description
<03:00>	TRGT_RET	RW,0	PCI target retry count in PCI clock cycles. This is the number of cycles that the CIA will wait when a resource is busy until it will stop. The value should be tuned for performance.
			TRGT_RET<3:0> Cycles
			0000 0
			0001 1
			. .
			. .
			. .
			1110 15
			1111 Indefinite


21172–CA Control and Status Registers
4.2 21172–CA General Registers

Field	Name	Type	Description														
<07:04>	MSTR_RET	RW,0	PCI master retry count in multiples of 4 PCI clock cycles. This is the number of cycles that the CIA will wait until it retries when it has been stopped. Recommended value = 0.														
			<table border="1"> <thead> <tr> <th>MSTR_RET<3:0></th> <th>Cycles</th> </tr> </thead> <tbody> <tr> <td>0000</td> <td>2</td> </tr> <tr> <td>0001</td> <td>6</td> </tr> <tr> <td>.</td> <td>.</td> </tr> <tr> <td>.</td> <td>.</td> </tr> <tr> <td>1110</td> <td>62</td> </tr> <tr> <td>1111</td> <td>2–66 (random)</td> </tr> </tbody> </table>	MSTR_RET<3:0>	Cycles	0000	2	0001	6	1110	62	1111	2–66 (random)
MSTR_RET<3:0>	Cycles																
0000	2																
0001	6																
.	.																
.	.																
1110	62																
1111	2–66 (random)																
<15:08>	MSTR_LAT	RW,0	PCI master latency timer in PCI clock cycles.														
<31:16>	Reserved	RO,0	—														

21172–CA Control and Status Registers
4.2 21172–CA General Registers

4.2.3 CIA Control Register

Register mnemonic: CIA_CTRL
 Register address: 87.4000.0100



LJ-04228.AI

21172–CA Control and Status Registers

4.2 21172–CA General Registers

Field	Name	Type	Description
<00>	PCI_EN	RW,0	0—CIA asserts reset to the PCI. 1—CIA does not assert reset to the PCI.
<01>	PCI_LOCK_EN	RW,0	0—CIA will not lock when the PCI tries to lock it. 1—CIA will lock when the PCI tries to lock it.
<02>	PCI_LOOP_EN	RW,0	0—CIA will not respond as a target when it is the master. 1—CIA will respond as a target when it is the master.
<03>	FST_BB_EN	RW,0	0—CIA will not initiate fast back-to-back PCI transactions. 1—CIA will initiate fast back-to-back PCI transactions.
<04>	PCI_MST_EN	RW,0	0—CIA will not initiate PCI transactions. 1—CIA will initiate PCI transactions.
<05>	PCI_MEM_EN	RW,0	0—CIA will not respond to PCI transactions. 1—CIA will respond to PCI transactions.
<06>	PCI_REQ64_EN	RW,0	0—CIA will not request 64-bit PCI data transactions. 1—CIA will request 64-bit PCI data transactions.
<07>	PCI_ACK64_EN	RW,0	0—CIA will not accept 64-bit PCI data transactions. 1—CIA will accept 64-bit PCI data transactions.
<08>	ADDR_PE_EN	RW,0	0—CIA will not check PCI address parity errors. 1—CIA will check PCI address parity errors.
<09>	PERR_EN	RW,0	0—CIA will not check PCI data parity errors. 1—CIA will check PCI data parity errors.
<10>	FILL_ERR_EN	RW,0	0—CIA will not assert fill_error on the system bus. 1—CIA will assert fill_error if an error occurs during a 21164 read miss transaction.
<11>	MCHK_ERR_EN	RW,0	0—CIA will not assert error on the system bus. 1—CIA will assert error to report system machine check conditions.
<12>	ECC_CHK_EN	RW,0	0—CIA will not check IOD data for ECC errors. 1—CIA will check IOD data for ECC errors.
<13>	ASSERT_IDLE_BC	RW,0	0—CIA will not assert idle_bc on the system bus when asserting addr_bus_req . 1—CIA will assert idle_bc when asserting addr_bus_req .
<14>	CON_IDLE_BC	RW,0	0—CIA may generate a noncontiguous idle_bc signal. 1—CIA will guarantee that the idle_bc signal is contiguous.

21172–CA Control and Status Registers
4.2 21172–CA General Registers

Field	Name	Type	Description
<15>	CSR_IOA_BYPASS	RW,0	0—CIA will not bypass the I/O address queue. 1—CIA will bypass the I/O address queue. See Table 4–8.
<16>	IO_FLUSHREQ_EN	RW,0	Used in combination with CPU_FLUSHREQ_EN. The 2 bits control the CIA’s response to a PCI master, causing a flush request to the 21164.
<17>	CPU_FLUSHREQ_EN	RW,0	See IO_FLUSHREQ_EN and Table 4–8.
<18>	ARB_CPU_EN	RW,0	0—Disable the bypass path from the 21164 into the memory and IOA queue (for a cache less than 2MB). 1—Enable the bypass path from the 21164 into the memory and IOA queue.
<19>	EN_ARB_LINK	RW,0	0—Disable CPU memory read transactions from creating an arbitration link (sharing a common ras cycle). 1—Enable CPU memory read transactions to create an arbitration link.
<21:20>	RD_TYPE	RW,0	This field controls the prefetch algorithm used for PCI memory read commands. See Table 4–9.
<23:22>	Reserved	RO,0	—
<25:24>	RL_TYPE	RW,0	This field controls the prefetch algorithm used for PCI memory read line commands. See Table 4–9.
<27:26>	Reserved	RO,0	—
<29:28>	RM_TYPE	RW,0	This field controls the prefetch algorithm used for PCI memory read multiple commands. See Table 4–9.
<30>	Reserved	RO,0	—
<31>	EN_DMA_RD_PERF	RW,1	0—Disable the DMA read performance logic. 1—Enable the DMA read performance logic.

21172–CA Control and Status Registers

4.2 21172–CA General Registers

Table 4–8 CPU and I/O Flush Request

IO_FLUSH_REQ_EN	CPU_FLUSH_REQ_EN	Response
0	0	Assert mem_ack_1 immediately.
0	1	Assert mem_ack_1 immediately and do not allow new CPU commands to proceed.
1	0	Illegal.
1	1	Wait until IOA queue is empty, then assert mem_ack_1 but do not allow new CPU commands to proceed.


Table 4–9 PCI READ Prefetch Algorithm

RL_TYPE	Type	Description
0 0	Short	Fetch requested longword or quadword and the remainder of the cache block.
0 1	Medium	Prefetch next block and no more. Will not cross 8KB page boundary.
1 0	Long	Keep prefetching until PCI transaction completes. Will not cross 8KB page.
1 1	Reserved	—

21172–CA Control and Status Registers
4.2 21172–CA General Registers

4.2.4 CIA Configuration Register

Register mnemonic: CIA_CNFG
 Register address: 87.4000.0140



LJ-04867.A15

Field	Name	Type	Description
<00>	IOA_BWEN	RW,0	Byte/word enable. 0—21164 byte/word support disabled. 1—21164 byte/word support enabled.
<03:01>	Reserved	RO,0	—
<04>	PCI_MWEN	RW,0	Monster window enable. 0—PCI target monster window disabled. 1—PCI target monster window enabled.
<05>	PCI_DWEN	RW,0	0—Second PCI target DMA write buffer disabled. 1—Second PCI target DMA write buffer enabled.
<07:06>	Reserved	RO,0	—
<08>	PCI_WLEN	RW,0	0—CIA will not link consecutive write operations into a single PCI transaction. 1—CIA will link consecutive write operations into a single PCI transaction.
<31:09>	Reserved	RO,0	—

21172-CA Control and Status Registers


4.2 21172-CA General Registers

4.2.5 Hardware Address Extension Register

Register mnemonic: HAE_MEM
Register address: 87.4000.0400

This hardware address extension register is used to extend a PCI sparse-space memory address up to the full 32-bit PCI address.

In sparse-addressing mode, the 21164 address provides the low-order PCI address bits. The high-order PCI address bits <31:26> are obtained from either the hardware address extension register or the 21164 address, depending on sparse regions, as shown in Table 4-10.



LJ-04229.AI

21172–CA Control and Status Registers
4.2 21172–CA General Registers

Table 4–10 HAE_MEM High-Order Sparse-Space Bits

Region ¹	PCI Address					
	<31>	<30>	<29>	<28>	<27>	<26>
1	HAE_MEM <31>	HAE_MEM <30>	HAE_MEM <29>	21164 <33>	21164 <32>	21164 <31>
2	HAE_MEM <15>	HAE_MEM <14>	HAE_MEM <13>	HAE_MEM <12>	HAE_MEM <11>	21164 <31>
3	HAE_MEM <7>	HAE_MEM <6>	HAE_MEM <5>	HAE_MEM <4>	HAE_MEM <3>	HAE_MEM <2>

¹Region 1 is 80.0000.0000 to 83.FFFF.FFFF.
Region 2 is 84.0000.0000 to 84.FFFF.FFFF.
Region 3 is 85.0000.0000 to 85.7FFF.FFFF.

Table 4–11 lists the address range of each of the three regions.

Table 4–11 Hardware Address Extension Memory Register

Field	Name	Type	21164 Address Range
<31:29>	Region 1	RW,0	80.0000.0000–83.FFFF.FFFF
<28:16>	Reserved	RW,0	—
<15:11>	Region 2	RW,0	84.0000.0000–84.FFFF.FFFF
<10:8>	Reserved	RW,0	—
<7:2>	Region 3	RW,0	85.0000.0000–85.7FFF.FFFF
<1:0>	Reserved	RW,0	—

Initializing this register to 0000.2028₁₆ will make all three regions contiguous starting at PCI address 0.


21172–CA Control and Status Registers
4.2 21172–CA General Registers

4.2.6 Hardware Address Extension I/O Register

Register mnemonic: HAE_IO
 Register address: 87.4000.0440

This hardware address extension register is used to extend a PCI sparse-space I/O address up to the full 32-bit PCI address. In sparse-addressing mode, the 21164 address provides the PCI address <24:00> and HAE_IO provides <31:25>.

At power-up this register is set to zero. In this case, sparse I/O region A and region B both map to the lower 32MB of sparse I/O space. Setting HAE_IO to 0200 0000 will make regions A and B consecutive in the lower 64MB of PCI I/O space.



LJ-04230.AI


Field	Name	Type	Description
<24:00>	Reserved	RO,0	—
<31:25>	HAE_IO	RW,0	In sparse-address mode, this field provides address bits <31:25>.

21172–CA Control and Status Registers
4.2 21172–CA General Registers

4.2.7 Configuration Register

Register mnemonic: CFG
 Register address: 87.4000.0480

The CFG field bits are used as the low 2 address bits during an access to PCI configuration space.



LJ-04231.AI

Field	Name	Type	Description								
<01:00>	CFG	RW,0	PCI configuration space access type: <table border="1" style="margin-left: 20px;"> <thead> <tr> <th>CFG<1:0></th> <th>Meaning</th> </tr> </thead> <tbody> <tr> <td>0 0</td> <td>Type 0 configuration cycle</td> </tr> <tr> <td>0 1</td> <td>Type 1 configuration cycle</td> </tr> <tr> <td>1 x</td> <td>Reserved</td> </tr> </tbody> </table>	CFG<1:0>	Meaning	0 0	Type 0 configuration cycle	0 1	Type 1 configuration cycle	1 x	Reserved
CFG<1:0>	Meaning										
0 0	Type 0 configuration cycle										
0 1	Type 1 configuration cycle										
1 x	Reserved										
<31:02>	Reserved	RO,0	—								


21172–CA Control and Status Registers

4.2 21172–CA General Registers

4.2.8 CIA Acknowledgment Enable Register

Register mnemonic: CACK_EN
 Register address: 87.4000.0600

The CIA acknowledgment register (CACK_EN) allows software to decide if the CIA will acknowledge receipt any of four particular commands from the 21164 or to ignore any of these commands (NOP response).



LJ-04232.AI

Field	Name	Type	Description
<00>	LOCK_EN	RW,1	1—Enables CIA to acknowledge receipt of a LOCK command from the 21164.
<01>	MB_EN	RW,1	1—Enables CIA to acknowledge receipt of an MB command from 21164.
<02>	SET_DIRTY_EN	RW,1	1—Enables CIA to acknowledge receipt of a SET DIRTY command from the 21164.
<03>	BC_VICTIM_EN	RW,1	1—Enables CIA to acknowledge receipt of a BC_VICTIM command from the 21164.
<31:04>	Reserved	RO,0	—

21172–CA Control and Status Registers
4.3 21172–CA Diagnostic Registers


4.3 21172–CA Diagnostic Registers

Sections 4.3.1 and 4.3.2 define and describe the 21172–CA diagnostic registers.

4.3.1 CIA Diagnostic Control Register

Register mnemonic: CIA_DIAG
 Register address: 87.4000.2000

The CIA_DIAG is used as a diagnostic/debug tool, allowing various errors to be tested.



LJ-04233.AI

Field	Name	Type	Description
<00>	FROM_WRT_EN	RW,0	0—Flash ROM cannot be programmed. A write transaction to flash ROM results in a system machine check. 1—The flash ROM can be programmed.
<01>	USE_CHECK	RW,0	1—Use DIAG_CHECK<7:0> for DMA write cycle ECC sent on the IOD bus.
<27:02>	Reserved	RO,0	—

21172–CA Control and Status Registers

4.3 21172–CA Diagnostic Registers


Field	Name	Type	Description												
<29:28>	FPE_PCI	RW,0	This field allows bad parity to be forced on the PCI for diagnostic purposes:												
			<table border="1"> <thead> <tr> <th colspan="2">FPE_PCI</th> </tr> <tr> <th><1:0></th> <th>Meaning</th> </tr> </thead> <tbody> <tr> <td>0 0</td> <td>Normal parity asserted on PCI.</td> </tr> <tr> <td>0 1</td> <td>Bad parity is forced onto ad<31:00> during data cycles.</td> </tr> <tr> <td>1 0</td> <td>Bad parity is forced onto ad<63:32> during data cycles.</td> </tr> <tr> <td>1 1</td> <td>Bad parity is forced onto ad<31:00> and ad<63:32> during address and data cycles.</td> </tr> </tbody> </table>	FPE_PCI		<1:0>	Meaning	0 0	Normal parity asserted on PCI.	0 1	Bad parity is forced onto ad<31:00> during data cycles.	1 0	Bad parity is forced onto ad<63:32> during data cycles.	1 1	Bad parity is forced onto ad<31:00> and ad<63:32> during address and data cycles.
FPE_PCI															
<1:0>	Meaning														
0 0	Normal parity asserted on PCI.														
0 1	Bad parity is forced onto ad<31:00> during data cycles.														
1 0	Bad parity is forced onto ad<63:32> during data cycles.														
1 1	Bad parity is forced onto ad<31:00> and ad<63:32> during address and data cycles.														
<30>	Reserved	RO,0	—												
<31>	FPE_TO_CPU	RW,0	1—A parity error is forced on the 21164 address/command bus when the CIA is the bus master.												

21172–CA Control and Status Registers
4.3 21172–CA Diagnostic Registers

4.3.2 Diagnostic Check Register

Register mnemonic: DIAG_CHECK
 Register address: 87.4000.3000

The DIAG_CHECK causes diagnostic DMA write transactions to force a known ECC pattern into memory. This register provides the ECC value that is written to memory if CIA_DIAG[USE_CHECK] is set.



LJ-04234.AI

Field	Name	Type	Description
<07:00>	DIAG_CHECK	RW,UNDEFINED	This register provides the quadword ECC for diagnostic DMA write transactions.
<31:08>	Reserved	RO,0	—

21172–CA Control and Status Registers
4.4 21172–CA Performance Monitor Registers


4.4 21172–CA Performance Monitor Registers

Sections 4.4.1 and 4.4.2 define and describe the 21172–CA (CIA) performance monitor registers.

4.4.1 Performance Monitor Register

Register mnemonic: **PERF_MONITOR**
 Register address: **87.4000.4000**

The performance monitor register is two 16-bit counters that can be programmed to count a variety of events. The performance control register is used to set up the counters. Each counter can be programmed to count events such as 21164 read miss transactions received by the CIA or DMA write transactions. This register can also be set up as a single 32-bit counter by directing the high counter to count the low counter overflow.




LJ-04235.AI

Field	Name	Type	Description
<15:00>	LOW_COUNT	RO,0	Value of the low counter.
<31:16>	HIGH_COUNT	RO,0	Value of the high counter.

21172–CA Control and Status Registers
4.4 21172–CA Performance Monitor Registers

4.4.2 Performance Control Register

Register mnemonic: PERF_CONTROL
Register address: 87.4000.4040



LJ-04236.AI

Field	Name	Type	Description
<07:00>	LOW_SELECT	RW,0	See Table 4–12.
<12:08>	Reserved	RO,0	—
<13>	LOW_COUNT_CLR	WO,0	Write a one to clear the low counter.
<14>	LOW_ERR_STOP	RW,0	If the CIA encounters an error and this bit is set, then stop the low error counter.
<15>	LOW_COUNT_START	RW,0	0—Do not count. Keep current values. 1—Start counting.
<23:16>	HIGH_SELECT	RW,0	See Table 4–12.
<28:24>	Reserved	RO,0	—
<29>	HIGH_COUNT_CLR	WO,0	Write a one to clear the high counter.
<30>	HIGH_ERR_STOP	RW,0	If the CIA encounters an error and this bit is set, then stop the high error counter.
<31>	HIGH_COUNT_START	RW,0	0—Do not count. Keep current values. 1—Start counting.

21172–CA Control and Status Registers
4.4 21172–CA Performance Monitor Registers

Table 4–12 PERF_CONTROL Low/High Select Encodings

Select ¹	Description
0000 0000	LOW_SELECT—The value is reserved. HIGH_SELECT—Make 32-bit counter.
0000 0001	Counting clock cycles—always increment.
0000 0010	Counting refresh cycles.
0001 0000	Counting number of 21164 commands acknowledged.
0001 0001	Counting number of 21164 read commands (modify or not).
0001 0010	Counting number of 21164 read miss commands (not modify).
0001 0011	Counting number of 21164 read miss modify commands.
0001 0100	Counting number of 21164 Bcache victim commands that are acknowledged by the CIA.
0001 0101	Counting number of 21164 lock commands that are acknowledged by the CIA.
0001 0110	Counting number of 21164 memory barrier commands that are acknowledged by the CIA.
0001 0111	Counting number of 21164 FETCH or FETCH_M commands.
0001 1000	Counting number of 21164 write block commands (lock or not).
0010 0000	Counting number of 21164 memory commands.
0010 0001	Counting number of 21164 I/O commands.
0010 0010	Counting number of 21164 I/O read commands.
0010 0011	Counting number of 21164 I/O write commands.
0010 0100	Counting number of CIA system commands issued (read/flush).
0010 0101	Counting number of 21164 system read commands issued.
0010 0110	Counting number of 21164 system flush commands issued.
0010 0111	Counting number of times CIA received NOACK as a response.
0010 1000	Counting number of times CIA received Scache/ACK as a response.
0010 1001	Counting number of times CIA received Bcache/ACK as a response.
0011 0000	Counting number of DMA read commands (total).
0011 0001	Counting number of DMA read commands.

¹Select = HIGH_SELECT <23:16> and LOW_SELECT <07:00>.

(continued on next page)

21172–CA Control and Status Registers
4.4 21172–CA Performance Monitor Registers

Table 4–12 (Cont.) PERF_CONTROL Low/High Select Encodings

Select ¹	Description
0011 0010	Counting number of DMA read commands (read line).
0011 0011	Counting number of DMA read commands (read multiple).
0011 0100	Counting number of DMA write commands (total).
0011 0101	Counting number of DMA write commands.
0011 0110	Counting number of DMA write commands (write and invalidate).
0011 0111	Counting number of DMA dual address cycles.
0011 1000	Counting number of DMA cycles that CIA responded to by issuing a retry.
0011 1001	Counting number of I/O cycles on which CIA received a retry response.
0100 0000	Counting number of times PCI bus lock was established.
0100 0001	Counting number of times 21164 tried to access block that was locked.
0100 0010	Counting number of times DMA commands (that caused a flush) hit on the victim address.
0101 0000	Counting number of times CIA had to refill the translation lookaside buffer (TLB).
0110 0000	Counting number of single-bit ECC errors detected.
All others	Reserved, unused, not counting.

¹Select = HIGH_SELECT <23:16> and LOW_SELECT <07:00>.

21172–CA Control and Status Registers
4.5 21172–CA Error Registers

4.5 21172–CA Error Registers


Sections 4.5.1 through 4.5.11 define and describe the 21172–CA (CIA) error registers.

4.5.1 CPU Error Information Register 0

Register mnemonic: CPU_ERR0
 Register address: 87.4000.8000

The low-order address bits of the system bus are locked into this register when the CIA detects an error event. Clearing all the error bits in CIA_ERR unlocks this register. When CPU_ERR0 is not locked, its contents are UNDEFINED.

The information in CPU_ERR0 and CPU_ERR1 is only related to system bus parity errors detected by the CIA.



LJ-04237.AI

Field	Name	Type	Description
<03:00>	Reserved	R0,0	—
<31:04>	addr_h<31:4>	RO,UNDEFINED	Contains address bits <31:4> of the current address on the system bus when a system bus error occurs.


21172-CA Control and Status Registers
4.5 21172-CA Error Registers

4.5.2 CPU Error Information Register 1

Register mnemonic: CPU_ERR1
 Register address: 87.4000.8040

The mask and command field and the remaining address field on the system bus are locked into CPU_ERR1 when the CIA detects a event error. Clearing all the error bits in CIA_ERR unlocks CPU_ERR1. When CPU_ERR1 is not locked, its contents are UNDEFINED.

The information in CPU_ERR0 and CPU_ERR1 is only related to system bus parity errors detected by the CIA (CPU_ERR<02 = 1>).



LJ-04238.AI

21172–CA Control and Status Registers

4.5 21172–CA Error Registers

Table 4–13 CPU Error Information Register 1

Field	Name	Type	Description
<02:00>	ADDR<34:32>	RO, UNDEFINED	Contains addr<34:32> from the system bus.
<06:03>	Reserved	RO,0	—
<07>	ADDR<39>	RO, UNDEFINED	Contains addr<39> from the system bus.
<11:08>	CMD	RO, UNDEFINED	Contains the command from the system bus.
<15:12>	INT4_VALID	RO, UNDEFINED	Contains int4_valid<3:0> from the system bus.
<20:16>	Reserved	RO,0	—
<21>	ADDR_CMD_PAR	RO, UNDEFINED	Contains the parity bit from the system bus.
<29:22>	Reserved	RO,0	—
<30>	FPE_2_CPU	RO, UNDEFINED	Contains a copy of the CSR bit to force bad parity on the system bus.
<31>	CPU_PE	RO, UNDEFINED	1—Indicates that the CIA has logged a system bus error.

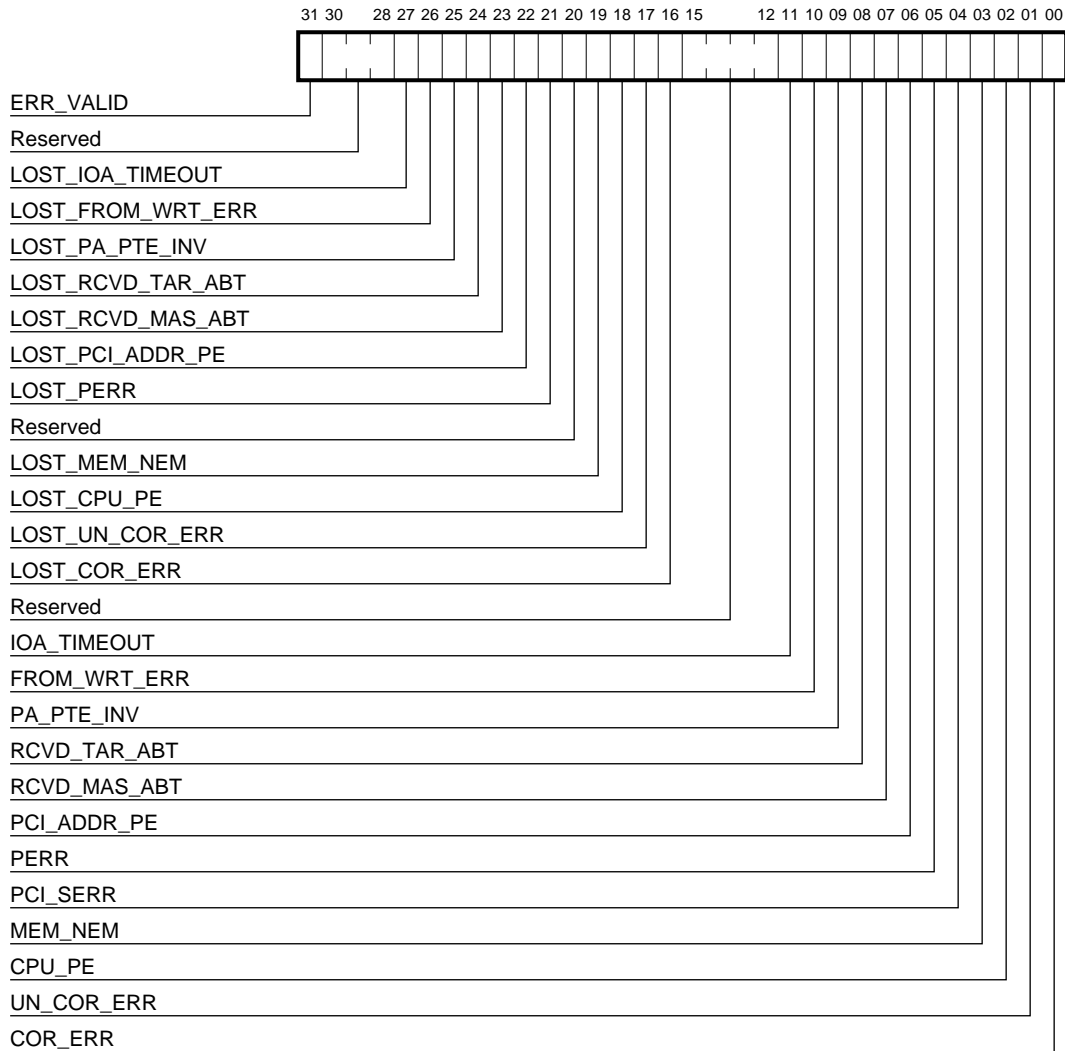
4.5.3 CIA Error Register

Register mnemonic: CIA_ERR
 Register address: 87.4000.8200

The CIA error register is used by the CIA to log information about an error condition detected in the CIA. All bits of CIA_ERR, except the LOST bit, will be locked until the CIA_ERR register is cleared by a write transaction. The LOST bit will be set whenever an error is detected and CIA_ERR is already locked.

21172-CA Control and Status Registers

4.5 21172-CA Error Registers



LJ-04239.AI

21172–CA Control and Status Registers

4.5 21172–CA Error Registers

Field	Name	Type	Description
<00>	COR_ERR	RW1C,0	<p>Corrected single-bit ECC error detected. This error cannot occur for a 21164 to memory read/write:</p> <ul style="list-style-type: none"> • 21164/memory read ECC errors are detected by the 21164. • 21164/memory write transactions are not checked. <p>This error is applicable to a DMA transaction, a scatter-gather TLB miss, or an I/O write transaction from the 21164.</p>
<01>	UN_COR_ERR	RW1C,0	<p>Uncorrectable ECC error detected (double-bit error or single-bit error if correction disabled). This error cannot occur for a 21164-to-memory read/write transaction because:</p> <ul style="list-style-type: none"> • 21164/memory read transaction ECC errors are detected by the 21164. • 21164/memory write transactions are not checked. <p>This error is applicable to a DMA transaction, a scatter-gather TLB miss, or an I/O write transaction from the 21164.</p>
<02>	CPU_PE	RW1C,0	system bus parity error detected.
<03>	MEM_NEM	RW1C,0	Access to nonexistent memory detected.
<04>	PCI_SERR	RW1C,0	PCI bus system error (SERR) detected.
<05>	PERR	RW1C,0	PCI bus data parity error (PERR) detected.
<06>	PCI_ADDR_PE	RW1C,0	PCI bus address parity error detected.
<07>	RCVD_MAS_ABT	RW1C,0	PCI master state machine generated master abort.
<08>	RCVD_TAR_ABT	RW1C,0	PCI master state machine received target abort.
<09>	PA_PTE_INV	RW1C,0	Invalid page table entry (PTE) on scatter-gather access.
<10>	FROM_WRT_ERR	RW1C,0	Write to flash ROM attempted without setting CIA_DIAG[FROM_WRT_EN].
<11>	IOA_TIMEOUT	RW1C,0	I/O timeout occurred. I/O read/write transaction failed to complete in 1 second.

21172–CA Control and Status Registers

4.5 21172–CA Error Registers

Field	Name	Type	Description
<15:12>	Reserved	RO,0	—
<16>	LOST_COR_ERR	RO,0	A correctable error was detected while CIA_ERR was locked.
<17>	LOST_UN_COR_ERR	RO,0	An uncorrectable error was detected while CIA_ERR was locked.
<18>	LOST_CPU_PE	RO,0	A system bus parity error was detected while CIA_ERR was locked.
<19>	LOST_MEM_NEM	RO,0	An access to nonexistent memory was detected while CIA_ERR was locked.
<20>	Reserved	RO,0	—
<21>	LOST_PERR	RO,0	A PCI data parity error was detected while CIA_ERR was locked.
<22>	LOST_PCI_ADDR_PE	RO,0	A PCI address parity error was detected while CIA_ERR was locked.
<23>	LOST_RCVD_MAS_ABT	RO,0	The PCI master state machine generated a master abort while CIA_ERR was locked.
<24>	LOST_RCVD_TAR_ABT	RO,0	The PCI master state machine received a target abort while CIA_ERR was locked.
<25>	LOST_PA_PTE_INV	RO,0	An invalid page table entry on scatter-gather access occurred while CIA_ERR was locked.
<26>	LOST_FROM_WRT_ERR	RO,0	A write transaction to flash ROM was attempted, with CIA_DIAG[FROM_WRT_EN] clear, while CIA_ERR was locked.
<27>	LOST_IOA_TIMEOUT	RO,0	An I/O timeout occurred while CIA_ERR was locked. An I/O read/write transaction failed to complete in 1 second.
<30:28>	Reserved	RO,0	—
<31>	ERR_VALID	RO,0	An error has been detected and the CIA error registers are all locked.


21172—CA Control and Status Registers

4.5 21172—CA Error Registers

4.5.4 CIA Status Register

Register mnemonic: CIA_STAT
 Register address: 87.4000.8240

This register contains status information about the error stored in CIA_ERR. The register provides the state of the CIA when the error was detected.



LJ-04240.AI

Field	Name	Type	Description
<00>	PCI_STATUS<0>	RO,0	1—The PCI target state machine is active.
<01>	PCI_STATUS<1>	RO,0	1—The PCI master state machine is active.
<02>	Reserved	RO,0	—
<03>	MEM_SOURCE	RO,0	0—21164 is the source of the memory cycle. 1—PCI is the source of the memory cycle.
<07:04>	IOA_VALID<3:0>	RO,0	Valid bits for the I/O command/address queue.
<10:08>	CPU_QUEUE<2:0>	RO,0	Valid bits for the 21164 command/address queue.
<11>	TLB_MISS	RO,0	1—A TLB miss refill was in progress when the error occurred.

21172–CA Control and Status Registers
4.5 21172–CA Error Registers

Field	Name	Type	Description																		
<15:12>	DM_ST<3:0>	RO, UNDEFINED	This field represents the state of the logic that moves data between the CIA and the DSW:																		
			<table border="1" style="width: 100%; border-collapse: collapse;"> <thead> <tr> <th style="text-align: left;">DM_ST<3:0></th> <th style="text-align: left;">Definition</th> </tr> </thead> <tbody> <tr> <td>0000</td> <td>Idle</td> </tr> <tr> <td>0001</td> <td>Restarting DSW IOW buffers</td> </tr> <tr> <td>0010</td> <td>I/O write transaction to the 64-bit PCI bus</td> </tr> <tr> <td>0110</td> <td>DMA read transaction or TLB miss</td> </tr> <tr> <td>0111</td> <td>DMA write transaction</td> </tr> <tr> <td>1010</td> <td>I/O read transaction to an internal CIA CSR</td> </tr> <tr> <td>1011</td> <td>I/O read transaction to the PCI bus (32-bit/64-bit).</td> </tr> <tr> <td>All others</td> <td>Reserved</td> </tr> </tbody> </table>	DM_ST<3:0>	Definition	0000	Idle	0001	Restarting DSW IOW buffers	0010	I/O write transaction to the 64-bit PCI bus	0110	DMA read transaction or TLB miss	0111	DMA write transaction	1010	I/O read transaction to an internal CIA CSR	1011	I/O read transaction to the PCI bus (32-bit/64-bit).	All others	Reserved
DM_ST<3:0>	Definition																				
0000	Idle																				
0001	Restarting DSW IOW buffers																				
0010	I/O write transaction to the 64-bit PCI bus																				
0110	DMA read transaction or TLB miss																				
0111	DMA write transaction																				
1010	I/O read transaction to an internal CIA CSR																				
1011	I/O read transaction to the PCI bus (32-bit/64-bit).																				
All others	Reserved																				
<17:16>	PA_CPU_RES<1:0>	RO,0	21164 response for the DMA: 00 = No response 01 = NOACK 10 = Scache hit 11 = Bcache hit																		
<31:18>	Reserved	RO,0	—																		


21172-CA Control and Status Registers

4.5 21172-CA Error Registers

4.5.5 CIA Error Mask Register

Register mnemonic: ERR_MASK
Register address: 87.4000.8280

ERR_MASK is used to disable the logging and reporting of errors. A zero disables logging/reporting of errors while a one enables logging/reporting of errors.



LJ-04241.AI

21172–CA Control and Status Registers
4.5 21172–CA Error Registers

Field	Name	Type	Description
<00>	CORR_ERR	RW,0	Disable/enable error logging/reporting for correctable ECC errors.
<01>	UN_COR_ERR	RW,0	Disable/enable error logging/reporting for uncorrectable ECC errors.
<02>	CPU_PE	RW,0	Disable/enable error logging/reporting for 21164 parity errors.
<03>	MEM_NEM	RW,0	Disable/enable error logging/reporting for non-existent memory access errors.
<04>	PCI_SERR	RW,0	Disable/enable error logging/reporting for PCI system errors.
<05>	PERR	RW,0	Disable/enable error logging/reporting for PCI data parity errors.
<06>	PCI_ADDR_PE	RW,0	Disable/enable error logging/reporting for PCI address parity errors.
<07>	RCVD_MAS_ABT	RW,0	Disable/enable error logging/reporting for PCI master abort errors.
<08>	RCVD_TAR_ABT	RW,0	Disable/enable error logging/reporting for PCI target abort errors.
<09>	PA_PTE_INV	RW,0	Disable/enable error logging/reporting for invalid PTE errors.
<10>	FROM_WRT_ERR	RW,0	Disable/enable error logging/reporting for flash ROM write errors.
<11>	IOA_TIMEOUT	RW,0	Disable/enable error logging/reporting for I/O timeout errors.
<31:12>	Reserved	RW,0	—


21172–CA Control and Status Registers

4.5 21172–CA Error Registers

4.5.6 CIA Syndrome Register

Register mnemonic: CIA_SYN
Register address: 87.4000.8300

The CIA uses CIA_SYN to log syndrome information about an error detected by the ECC checkers. The syndrome information is locked into CIA_SYN when a CIA error occurs. Clearing all the error bits in CIA_ERR unlocks this register. When CIA_SYN is not locked its contents are UNDEFINED.



LJ-04242.AI

Field	Name	Type	Description
<07:00>	ECC_SYNDROME	RO, UNDEFINED	The syndrome is locked in ECC_SYNDROME <07:00> when an error is detected by the CIA.
<31:08>	Reserved	RO,0	—

21172–CA Control and Status Registers
4.5 21172–CA Error Registers

Table 4–14 lists the ECC syndromes for single-bit errors.

Table 4–14 ECC Syndromes for Single-Bit Errors


Bit ₁₀	Syndrome ₁₆	Bit ₁₀	Syndrome ₁₆	Bit ₁₀	Syndrome ₁₆	Bit ₁₀	Syndrome ₁₆
Syndromes for Data Bits							
00	CE	01	CB	02	D3	03	D5
04	D6	05	D9	06	DA	07	DC
08	23	09	25	10	26	11	29
12	2A	13	2C	14	31	15	34
16	0E	17	0B	18	13	19	15
20	16	21	19	22	1A	23	1C
24	E3	25	E5	26	E6	27	E9
28	EA	29	EC	30	F1	31	F4
32	4F	33	4A	34	52	35	54
36	57	37	58	38	5B	39	5D
40	A2	41	A4	42	A7	43	A8
44	AB	45	AD	46	B0	47	B5
48	8F	49	8A	50	92	51	94
52	97	53	98	54	9B	55	9D
56	62	57	64	58	67	59	68
60	6B	61	6D	62	70	63	75
Syndromes for Check Bits							
00	01	01	02	02	04	03	08
04	10	05	20	06	40	07	80

21172–CA Control and Status Registers
4.5 21172–CA Error Registers

4.5.7 CIA Memory Port Status Register 0

Register mnemonic: MEM_ERR0
 Register address: 87.4000.8400

The low-order address bits of the memory port address bus are locked into MEM_ERR0 upon a CIA-detected error. Clearing all the error bits in CIA_ERR unlocks this register. If the register is not locked, the contents are UNDEFINED.



LJ-04243.AI


Field	Name	Type	Description
<03:00>	Reserved	RO, UNDEFINED	—
<31:04>	MEM_PORT_ADDR<31:4>	RO, UNDEFINED	Contains address bits <31:4> of the current address in the memory port when the CIA detects an error.

21172–CA Control and Status Registers
4.5 21172–CA Error Registers

4.5.8 CIA Memory Port Status Register 1

Register mnemonic: MEM_ERR1
 Register address: 87.4000.8440

The command memory mask (**int4_valid<3:0>** from the CPU), memory sequencer state, the source of the command, and encoded set select field, and the remaining address field are locked into MEM_ERR1 upon a CIA error. Clearing all error bits in CIA_ERR unlocks this register. If the register is not locked, the contents are UNDEFINED.



LJ-04244.AI

21172–CA Control and Status Registers

4.5 21172–CA Error Registers

Field	Name	Type	Description
<01:00>	MEM_PORT_ADDR<33:32>	RO, UNDEFINED	Captures memory port address bits <33:32> when the CIA detects an error.
<06:02>	Reserved	RO,0	—
<07>	MEM_PORT_ADDR<39>	RW, UNDEFINED	Captures memory port address bit <39> when the CIA detects an error.
<11:08>	MEM_PORT_CMD	RO, UNDEFINED	Contains command field of the memory port when a CIA error is detected, as listed in Table 4–15.
<15:12>	MEM_PORT_MASK	RW, UNDEFINED	The mask bits when the error occurred.
<19:16>	SEQ_ST	RW, UNDEFINED	The memory sequencer state when the error occurred, as listed in Table 4–16.
<20>	MEM_PORT_SOURCE	RW, UNDEFINED	Source of the memory command: 0—21164 is source of memory command. 1—DMA is source of memory command.
<23:21>	Reserved	RO,0	—
<28:24>	SET_SEL_ENC	RW, UNDEFINED	Encoded set select indicates which memory set was active when the error occurred, as listed in Table 4–17.
<31:29>	Reserved	RO,0	—

Table 4–15 Memory Port Command Field (MEM_PORT_CMD)

MEM_PORT_SOURCE	MEM_PORT_CMD	Description
0	011X	21164 WRITE BLOCK or WRITE BLOCK LOCK
0	10XX	21164 READ MISS, READ MISS MODIFY
0	1100	21164 BC_VICTIM
0	111X	21164 READ MISS MODIFY
1	10XX	DMA READ, DMA READ MODIFY
1	001X	DMA WRITE

21172–CA Control and Status Registers
4.5 21172–CA Error Registers

Table 4–16 Memory Sequencer State Field (SEQ_ST)

SEQ_ST	Description
0000	Idle.
0001	DMA read or write.
0010, 0011	21164 READ MISS (or READ MISS MDOIFY) with victim.
0100, 0101, 0110	21164 READ MISS (or READ MISS MODIFY) with no victim.
01F11, 1000, 1001	Refresh.
1100	Idle. Waiting for DMA pending read.
1110, 1111	Idle. ras precharge

Table 4–17 Encoded Set Select Field (SET_SEL_ENC)


SET_SEL_ENC	Description
00000	Set 0 selected
00001	Set 1 selected
00010	Set 2 selected
00011	Set 3 selected
00100	Set 4 selected
00101	Set 5 selected
00110	Set 6 selected
00111	Set 7 selected
01000	Set 8 selected
01001	Set 9 selected
01010	Set A selected
01011	Set B selected
01100	Set C selected
01101	Set D selected
01110	Set E selected
01111	Set F selected
10000	No set selected
11111	Refresh cycle.
—	All other codes reserved.

21172–CA Control and Status Registers
4.5 21172–CA Error Registers

4.5.9 PCI Error Register 0

Register mnemonic: PCI_ERR0
 Register address: 87.4000.8800

PCI_ERR0 is used by the CIA to log information about the state of the PCI interface when an error is detected by the CIA. PCI_ERR0 is locked, as are all CIA error registers, when the CIA detects an error. PCI_ERR0 is unlocked when CIA_ERR is cleared. When the register is not locked, the contents are UNPREDICTABLE.



LJ04245A.AI5

The data in the WINDOW, DMA_DAC, and DMA_CMD files is associated with the address stored in PCI_ERR1. This group and PCI_ERR1 hold information related to the following memory errors that occur while the CIA is handling a DMA transaction:

- Correctable errors (CIA_ERR<00>)
- Uncorrectable errors (CIA_ERR<01>)
- Access to nonexistent memory (CIA_ERR<03>)
- Invalid page table entry (CIA_ERR<09>)

21172–CA Control and Status Registers

4.5 21172–CA Error Registers

The data in the PCI_DAC, PCI_CMD, TARGET_STATE, and MASTER_STATE fields is associated with the address stored in PCI_ERR2. This group and PCI_ERR2 hold information about the following PCI bus errors:

- PCI data parity error (CIA_ERR<05>)
- PCI address parity error (CIA_ERR<06>)
- PCI master abort (CIA_ERR<07>)
- PCI target abort (CIA_ERR<08>)
- IOA timeout (CIA_ERR<11>)

The LOCK_STATE field has general information about the current state of the CIA not specifically associated with either PCI_ERR1 or PCI_ERR2.

Field	Name	Type	Description														
<03:00>	DMA_CMD	RO, UNDEFINED	This field holds the current DMA command on the PCI.														
<04>	LOCK_STATE	RO, UNDEFINED	1—The CIA was locked when the error was detected.														
<05>	DMA_DAC	RO, UNDEFINED	1—The error occurred during a PCI dual address cycle (DAC).														
<07:06>	Reserved	RO,0	—														
<12:08>	WINDOW	RO, UNDEFINED	Indicates which window (if any) was selected by the PCI address.														
			<table style="width: 100%; border-collapse: collapse;"> <thead> <tr> <th style="text-align: left;">WINDOW<4:0></th> <th style="text-align: left;">Window Selected</th> </tr> </thead> <tbody> <tr> <td>00000</td> <td>No window active</td> </tr> <tr> <td>00001</td> <td>Window 0 hit</td> </tr> <tr> <td>00010</td> <td>Window 1 hit</td> </tr> <tr> <td>00100</td> <td>Window 2 hit</td> </tr> <tr> <td>01000</td> <td>Window 3 hit</td> </tr> <tr> <td>10000</td> <td>Window 4 hit Monster window</td> </tr> </tbody> </table>	WINDOW<4:0>	Window Selected	00000	No window active	00001	Window 0 hit	00010	Window 1 hit	00100	Window 2 hit	01000	Window 3 hit	10000	Window 4 hit Monster window
WINDOW<4:0>	Window Selected																
00000	No window active																
00001	Window 0 hit																
00010	Window 1 hit																
00100	Window 2 hit																
01000	Window 3 hit																
10000	Window 4 hit Monster window																
<15:13>	Reserved	RO,0	—														

21172–CA Control and Status Registers
4.5 21172–CA Error Registers


Field	Name	Type	Description
<19:16>	MASTER_STATE	RO,0	Master state indicated here: 0 = Idle 1 = Drive bus 2 = Address step cycle 3 = Address cycle 4 = Data cycle 5 = Last read data cycle 6 = Last write data cycle 7 = Read stop cycle 8 = Write stop cycle 9 = Read turnaround cycle A = Write turnaround cycle B = Reserved C = Reserved D = Reserved E = Reserved F = Unknown state
<23:20>	TARGET_STATE	RO,0	Target state indicated here: 0 = Idle 1 = Busy 2 = Read data cycle 3 = Write data cycle 4 = Read stop cycle 5 = Write stop cycle 6 = Read turnaround cycle 7 = Write turnaround cycle 8 = Read wait cycle 9 = Write wait cycle A = Reserved B = Reserved C = Reserved D = Reserved E = Reserved F = Unknown state
<27:24>	PCI_CMD	RO, UNDEFINED	The current PCI command.
<28>	PCI_DAC	RO, UNDEFINED	1—The current PCI command is a dual address cycle command.
<31:29>	Reserved	RO,0	—

21172–CA Control and Status Registers
4.5 21172–CA Error Registers

4.5.10 PCI Error Register 1

Register mnemonic: PCI_ERR1
 Register address: 87.4000.8840

The CIA uses PCI_ERR1 to log DMA **ad<31:0>** when an error condition is logged into PCI_ERR0. This register is locked whenever the CIA detects an error. This register always captures **ad<31:0>**—even for a DMA DAC cycle.



LJ-04246.AI

Signal **ad<39:32>** can be obtained from W_DAC; **ad<63:40>** must be zero for the CIA to hit on the DAC cycle.

PCI_ERR1 is unlocked when the error bits in CIA_ERR have all been cleared. The contents of PCI_ERR1 are UNPREDICTABLE when it is not locked.

PCI_ERR1 and some fields in PCI_ERR0 (WINDOW, DMA_DAC, DMA_CMD) hold information related to the following errors, which are associated with the memory while the CIA is handling a DMA transaction:

- Correctable ECC error (CIA_ERR<00>)
- Uncorrectable ECC error (CIA_ERR<01>)
- Access to nonexistent memory (CIA_ERR<03>)
- Invalid page table entry (CIA_ERR<09>)


Field	Name	Type	Description
<31:00>	DMA_ADDRESS<31:0>	RO, UNDEFINED	Contains DMA ad<31:0> .

21172–CA Control and Status Registers
4.5 21172–CA Error Registers

4.5.11 PCI Error Register 2

Register mnemonic: PCI_ERR2
 Register address: 87.4000.8880

The CIA uses PCI_ERR2 to log **ad<31:0>** when an error condition is logged into PCI_ERR0. This register is locked whenever the CIA detects an error. This register always captures **ad<31:0>** — even for a DMA DAC cycle.



LJ-04247.AI

Signal **ad<39:32>** can be obtained from W_DAC; **ad<63:40>** must be zero for the CIA to hit on the DAC cycle.

PCI_ERR2 is unlocked when the error bits in CIA_ERR have all been cleared. The contents of PCI_ERR1 are UNPREDICTABLE when it is not locked.

PCI_ERR2 and some fields in PCI_ERR0 (PCI_DAC, PCI_CMDK, TARGET_STATE, MASTER_STATE) hold information related to the following errors, which are associated with the PCI bus:

- PCI data parity error (CIA_ERR<05>)
- PCI address parity error (CIA_ERR<06>)
- PCI master abort (CIA_ERR<07>)
- PCI target abort (CIA_ERR<08>)
- IOA timeout (CIA_ERR<11>)

Field	Name	Type	Description
<31:00>	PCI_ADDRESS<31:0>	RO, UNDEFINED	Contains PCI ad<31:0> .

21172–CA Control and Status Registers
4.6 21172–CA System Configuration Registers


4.6 21172–CA System Configuration Registers

Sections 4.6.1 through 4.6.3 define and describe the 21172–CA system configuration registers.

4.6.1 Memory Configuration Register

Register mnemonic: MCR
 Register address: 87.5000.0000

MCR defines the system Bcache and memory configuration. The register contents are used to configure the CIA memory controller.



LJ-04248.AI

Field	Name	Type	Description
<00>	MEM_SIZE	RW,0	This bit must be written by firmware. 0—The memory port is configured to a 128-bit data path. See Figure 3–9. 1—The memory port is configured to a 256-bit data path. See Figure 3–10. MMB0 and MMB1 must be populated identically or an illegal configuration is flagged.

21172–CA Control and Status Registers

4.6 21172–CA System Configuration Registers

Field	Name	Type	Description
<03:01>	Reserved	RO,0	—
<06:04>	CACHE_SIZE	RW,0	SROM code determines the cache size and writes these bits before any memory cycles are started.
CACHE_SIZE			
		<2:0>	Cache RAM Size
		0 0 0	No cache present
		0 0 1	Reserved
		0 1 0	128K * X, 2MB—complete cache
		0 1 1	256K * X, 4MB or larger—complete cache
		100–111	Reserved
<07>	Reserved	RO,0	—
<17:08>	REF_RATE	RW,0	This field controls the memory refresh rate. A 10-bit, free-running counter, starting at zero, counts up to the REF_RATE and resets to zero. Memory is refreshed every time the REF_RATE value is reached. Setting this field to zero will disable refresh cycles.
<19:18>	REF_BURST	RW,0	The refresh state machine can be set up to perform all or half the single inline memory modules (SIMM) at once and also to do one or two refreshes when the REF_RATE counter rolls over.
		REF_BURST<1:0>	Refresh Burst Size
		0 0	ras to all SIMMs at once, single refresh
		0 1	ras to all SIMMs at once, double refresh (burst)
		1 0	ras to half of SIMMs at once, single refresh
		1 1	ras to half of SIMMs at once, double refresh (burst)

21172–CA Control and Status Registers
4.6 21172–CA System Configuration Registers

Field	Name	Type	Description	
<21:20>	TMG_R0	RW,0	Controls the row address nominal setup time.	
			TMG_R0<1:0>	Row Address Setup (Nominal)
			00	15 ns
			01	30 ns
			10	45 ns
11	60 ns			
<22>	LONG_CBR_CAS	RW,1	Refresh (cas before ras) cas pulse width. 0—Pulse width = 60 ns. 1—Pulse width = 90 ns.	
<25:23>	Reserved	RO,0	—	
<27:26>	DLY_IDLE_BC	RW,00	Write only 00₂ to this field. Writing any other value may cause UNPREDICTABLE behavior.	
<28>	Reserved	RO,0	—	
<29>	EARLY_IDLE_BC	RW,1	Write only 1₂ to this bit. Writing a 0 ₂ may cause UNPREDICTABLE behavior.	
<31:30>	Reserved	RO,0	—	

21172—CA Control and Status Registers
4.6 21172—CA System Configuration Registers


4.6.2 Memory Base Address Register *n*

Register mnemonic: *MBA_n*
 Register address: See below

The eight memory base address registers correspond to the 16 possible banks of memory that the CIA can support. The bits within the MBA register provide a pattern that is compared with the incoming address to determine which bank is being accessed. The minimum bank size for an MBA register is 16MB.

The addresses for the eight memory base address registers follow:

- | | |
|-------------------|-------------------|
| MBA0—87.5000.0600 | MBA2—87.5000.0680 |
| MBA4—87.5000.0700 | MBA6—87.5000.0780 |
| MBA8—87.5000.0800 | MBAA—87.5000.0880 |
| MBAC—87.5000.0900 | MBAE—87.5000.0980 |



LJ04249A.AI5

21172–CA Control and Status Registers

4.6 21172–CA System Configuration Registers

Field	Name	Type	Description																		
<00>	S0_VALID	RW,0	1—Indicates that side 0 for the bank is valid.																		
<03:01>	ROW_TYPE	RW,0	Specifies the row and column configuration of the physical bank being referenced. ROW_TYPE is used for generating the memory address map. See Figures 3–9 and 3–10.																		
			<table border="1"> <thead> <tr> <th>ROW_TYPE<2:0></th> <th>Row × Column</th> </tr> </thead> <tbody> <tr> <td>0 0 0</td> <td>10 × 10</td> </tr> <tr> <td>0 0 1</td> <td>11 × 11 or 12 × 10</td> </tr> <tr> <td>0 1 0</td> <td>12 × 12 or 13 × 11</td> </tr> <tr> <td>0 1 1</td> <td>10 × 10, 11 × 9, or 12 × 8</td> </tr> <tr> <td>1 0 0</td> <td>10 × 9</td> </tr> <tr> <td>1 0 1</td> <td>11 × 10</td> </tr> <tr> <td>1 1 0</td> <td>12 × 11</td> </tr> <tr> <td>1 1 1</td> <td>13 × 12</td> </tr> </tbody> </table>	ROW_TYPE<2:0>	Row × Column	0 0 0	10 × 10	0 0 1	11 × 11 or 12 × 10	0 1 0	12 × 12 or 13 × 11	0 1 1	10 × 10, 11 × 9, or 12 × 8	1 0 0	10 × 9	1 0 1	11 × 10	1 1 0	12 × 11	1 1 1	13 × 12
ROW_TYPE<2:0>	Row × Column																				
0 0 0	10 × 10																				
0 0 1	11 × 11 or 12 × 10																				
0 1 0	12 × 12 or 13 × 11																				
0 1 1	10 × 10, 11 × 9, or 12 × 8																				
1 0 0	10 × 9																				
1 0 1	11 × 10																				
1 1 0	12 × 11																				
1 1 1	13 × 12																				
<08:04>	MASK	RW,0	Indicates bits to use when comparing MBA _n [PATTERN] with the physical address. The MASK field implicitly indicates the size of the memory SIMMs in the bank corresponding to the MBA _n register. For a 256-bit memory data bus the valid MASK fields are:																		
			<table border="1"> <thead> <tr> <th>SIMM Size</th> <th>MASK</th> <th>Comparison Range</th> </tr> </thead> <tbody> <tr> <td>4MB or 8MB</td> <td>00001</td> <td>addr<24> is not used in the address comparison.</td> </tr> <tr> <td>8MB or 16MB</td> <td>00011</td> <td>addr<25:24> are not used in the address comparison.</td> </tr> <tr> <td>16MB or 32MB</td> <td>00111</td> <td>addr<26:24> are not used in the address comparison.</td> </tr> <tr> <td>32MB or 64MB</td> <td>01111</td> <td>addr<27:24> are not used in the address comparison.</td> </tr> <tr> <td>64MB or 128MB</td> <td>11111</td> <td>addr<28:24> are not used in the address comparison.</td> </tr> </tbody> </table>	SIMM Size	MASK	Comparison Range	4MB or 8MB	00001	addr<24> is not used in the address comparison.	8MB or 16MB	00011	addr<25:24> are not used in the address comparison.	16MB or 32MB	00111	addr<26:24> are not used in the address comparison.	32MB or 64MB	01111	addr<27:24> are not used in the address comparison.	64MB or 128MB	11111	addr<28:24> are not used in the address comparison.
SIMM Size	MASK	Comparison Range																			
4MB or 8MB	00001	addr<24> is not used in the address comparison.																			
8MB or 16MB	00011	addr<25:24> are not used in the address comparison.																			
16MB or 32MB	00111	addr<26:24> are not used in the address comparison.																			
32MB or 64MB	01111	addr<27:24> are not used in the address comparison.																			
64MB or 128MB	11111	addr<28:24> are not used in the address comparison.																			
<14:09>	Reserved	RO,0	—																		
<15>	S1_VALID	RW,0	1—Indicates that side 1 for the bank is valid. Side 1 cannot be valid unless side 0 is also valid.																		

21172–CA Control and Status Registers

4.6 21172–CA System Configuration Registers

Field	Name	Type	Description															
<25:16>	PATTERN	RW,0	MBA _n [PATTERN] is compared with the incoming system bus address addr <33:24> and then ORed with MASK<8:4> to determine if the bank corresponding to this MBA is being accessed. PATTERN indicates the base address of the memory bank while MASK<8:4> indicates the size of the bank.															
<27:26>	Reserved	RO,0	—															
<29:28>	TMG_SEL	RW,0	Timing register select:															
<table border="1"> <thead> <tr> <th>TMG_SEL</th> <th>TMG_n Selected</th> <th>Description</th> </tr> </thead> <tbody> <tr> <td>0 0</td> <td>TMG0</td> <td>Used to control timing of 50-ns and 60-ns SIMMs</td> </tr> <tr> <td>0 1</td> <td>TMG1</td> <td>Used to control timing of 70-ns SIMMs</td> </tr> <tr> <td>1 0</td> <td>TMG2</td> <td>Used to control timing of 80-ns SIMMs</td> </tr> <tr> <td>1 1</td> <td>Reserved</td> <td>—</td> </tr> </tbody> </table>				TMG_SEL	TMG _n Selected	Description	0 0	TMG0	Used to control timing of 50-ns and 60-ns SIMMs	0 1	TMG1	Used to control timing of 70-ns SIMMs	1 0	TMG2	Used to control timing of 80-ns SIMMs	1 1	Reserved	—
TMG_SEL	TMG _n Selected	Description																
0 0	TMG0	Used to control timing of 50-ns and 60-ns SIMMs																
0 1	TMG1	Used to control timing of 70-ns SIMMs																
1 0	TMG2	Used to control timing of 80-ns SIMMs																
1 1	Reserved	—																
<31:30>	Reserved	RO,0	—															

21172–CA Control and Status Registers
4.6 21172–CA System Configuration Registers


4.6.3 Memory Timing Information Register *n*

Register mnemonic: TMG_n
 Register address: See below

The addresses for the three memory timing information registers follow:

- TMG0—87.5000.0B00
- TMG1—87.5000.0B40
- TMG2—87.5000.0B80

The memory timing registers contain the system memory parameters that control operation of the CIA memory sequencer. See $MBA_n<29:28>$ in Section 4.6.2 for TMG_n register selection.



LJ-04250.AI

21172–CA Control and Status Registers

4.6 21172–CA System Configuration Registers

Field	Name	Type	Description
<1:0>	TMG_R1	RW,0	Read starting data delay.
<3:2>	TMG_R2	RW,0	Row address hold.
<5:4>	TMG_R3	RW,0	Read cycle time.
<7:6>	TMG_R4	RW,0	Read cas assertion delay.
<9:8>	TMG_R5	RW,0	Read cas pulse width.
<11:10>	TMG_R6	RW,0	Read column address hold.
<13:12>	TMG_W1	RW,0	Write data delay.
<16:14>	TMG_W4	RW,0	Write cas assertion delay.
<17>	TMG_PRE	RW,0	ras precharge delay.
<19:18>	TMG_V3	RW,0	Write cycle time.
<22:20>	TMG_V4	RW,0	Linked victim, cas assertion delay.
<23>	Reserved	RO,0	—
<25:24>	TMG_V5	RW,0	Victim/write, cas pulse width.
<27:26>	TMG_V6	RW,0	Victim/write, column address hold.
<29:28>	TMG_RV	RW,0	Read-to-victim start delay.
<31:30>	TMG_RD_DLY	RW,0	Read data delay (affects DSW).

21172–CA Control and Status Registers
4.6 21172–CA System Configuration Registers

Table 4–18 lists the possible values of the encoded memory parameters.

Table 4–18 Memory Timing Parameters, Encoded Values

Parameter	Encoded Values							
	000	001	010	011	100	101	110	111
TMG_R1	30	— ¹	60	—	—	—	—	—
TMG_R2	30	45	60	75	—	—	—	—
TMG_R3	60	—	90	—	—	—	—	—
TMG_R4	30	45	60	75	—	—	—	—
TMG_R5	30	45	60	75	—	—	—	—
TMG_R6	30	45	60	75	—	—	—	—
TMG_W1	30	—	60	—	—	—	—	—
TMG_W4	60	75	90	105	120	135	—	—
TMG_PRE	0	30	—	—	—	—	—	—
TMG_V3	60	—	90	—	—	—	—	—
TMG_V4	60	75	90	105	120	135	—	—
TMG_V5	30	45	60	75	—	—	—	—
TMG_V6	30	45	60	75	—	—	—	—
TMG_RV ²	0	0	30	30	—	—	—	—
TMG_RV ³	15	30	45	60	—	—	—	—
TMG_RD_DLY	0	15	30	45	—	—	—	—

¹Illegal parameters are indicated by —.

²Delay between a read and a linked victim.


³Delay after the read to assertion of **mem_en**.

21172-CA Control and Status Registers

4.6 21172-CA System Configuration Registers

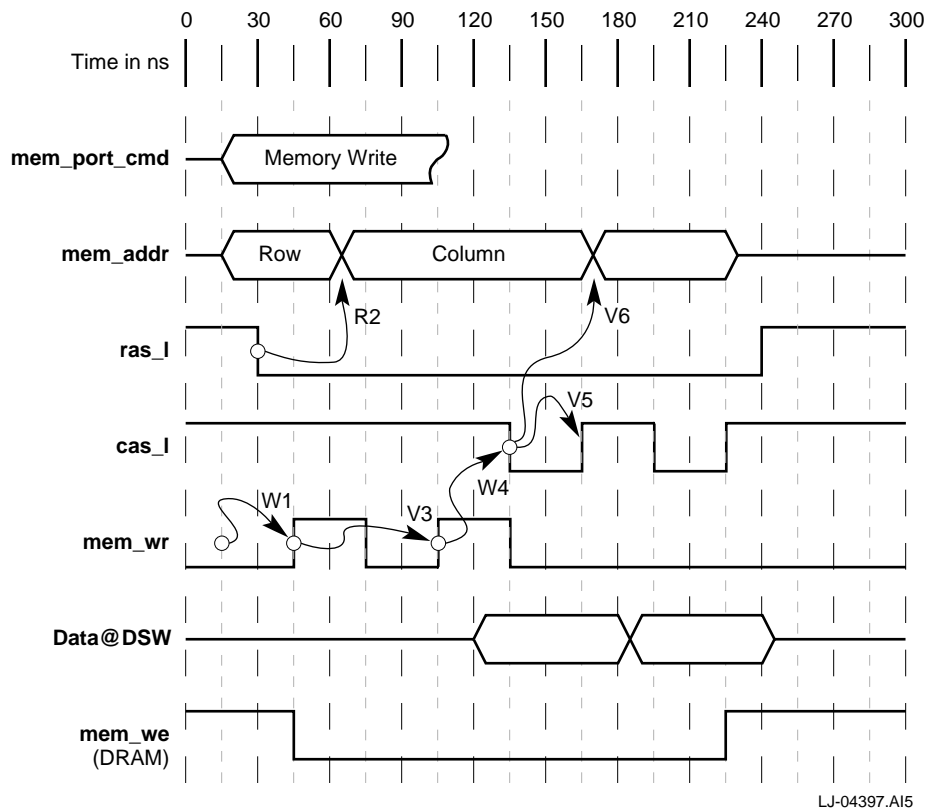
Figures 4-1 and 4-2 demonstrate TMG_n parameters applied to 21172 timing for read and write transactions. In these examples of 21172 timing, the 21172 accepts and uses a 33-MHz clock input. The 21172 doubles that frequency, generating a 66-MHz clock for its internal use.

Figure 4-1 Memory Read Timing Sample



**21172-CA Control and Status Registers
4.6 21172-CA System Configuration Registers**

Figure 4-2 Memory Write Timing Sample



21172–CA Control and Status Registers
4.6 21172–CA System Configuration Registers

Figures 4–3 and 4–4 list the minimum and maximum timing limits for memory and other 21172 chipset signals.

Figure 4–3 CIA Signal Timing Limits

CIA Signal	Timing	Reference	Setup	Hold	Output Delay							
					Max 15 pF		Min 15 pF		Max 50 pF		Min 50 pF	
					Rise	Fall	Rise	Fall	Rise	Fall	Rise	Fall
ack64_l	CLK	7.0	–0.4	6.3	5.0	1.4	1.1	7.2	5.8	1.8	1.4	
ad	CLK	5.1	0.2	8.0	7.1	1.5	1.4	8.9	7.9	1.9	1.7	
addr	CLK2X	5.0	0.7	6.4	6.7	0.2	0.2	8.3	8.3	1.0	0.9	
addr<39>	CLK2X	7.3	0.0	6.4	6.7	1.0	1.0	8.3	8.3	1.8	1.7	
mem_addr (row)	addr	—	—	5.3	5.5	1.7	1.8	6.5	6.5	2.2	2.2	
mem_b0 (row)	addr	—	—	5.3	5.5	2.0	2.1	6.5	6.5	2.5	2.5	
addr_bus_req	CLK2X	—	—	5.5	5.3	0.5	0.5	7.3	6.9	1.3	1.1	
addr_cmd_par	CLK2X	4.4	0.5	6.5	6.5	0.6	0.8	8.3	8.1	1.4	1.5	
cack	CLK2X	—	—	5.1	5.1	0.3	0.3	7.0	6.7	1.1	1.0	
cas ¹	CLK2X	—	—	5.1	5.1	0.3	0.3	6.3	6.1	0.8	0.7	
cas_l ² (cas & ras)	CLK2X	—	—	5.3	5.3	0.3	0.3	6.5	6.3	0.8	0.7	
cbe_l	CLK	4.5	0.1	6.4	5.9	1.6	1.1	7.3	6.7	2.0	1.4	
mem_addr (column)	CLK2X	—	—	6.8	7.1	0.2	0.3	8.0	8.0	0.7	0.7	
mem_b0 (column)	CLK2X	—	—	6.2	6.4	0.4	0.5	7.4	7.3	0.9	0.9	
cmc	CLK	—	—	4.7	4.6	1.1	1.1	6.6	6.2	1.9	1.7	
cmd	CLK2X	7.1	0.3	5.7	6.1	0.4	0.4	7.6	7.6	1.2	1.1	
dack	CLK2X	—	—	5.1	5.1	0.3	0.3	6.9	6.6	1.1	1.0	
devsel_l	CLK	3.5	0.0	6.3	4.6	1.3	1.1	7.3	5.3	1.7	1.4	
error	CLK	—	—	8.5	8.8	2.2	2.1	10.4	10.6	3.0	2.9	
fill	CLK2X	—	—	5.6	5.4	0.5	0.5	7.5	7.0	1.3	1.2	
fill_error	CLK2X	—	—	5.5	5.3	0.5	0.5	7.4	6.9	1.3	1.1	
fill_id	CLK2X	—	—	5.4	5.3	0.5	0.4	7.3	6.9	1.3	1.1	
frame_l	CLK	6.4	–0.4	6.5	5.0	1.6	1.2	7.5	5.8	2.0	1.5	
gnt_l	CLK	6.4	–0.5	—	—	—	—	—	—	—	—	
idle_bc	CLK2X	—	—	7.4	7.8	0.9	0.8	9.3	9.4	1.7	1.5	
int	CLK	—	—	6.2	6.1	2.0	2.0	8.1	7.9	2.8	2.7	
int4_valid	CLK2X	5.5	0.8	—	—	—	—	—	—	—	—	

Notes:

- ¹ Set select memory mode
- ² Alternate memory mode

MK1455–25

(continued on next page)

21172–CA Control and Status Registers
4.6 21172–CA System Configuration Registers

Figure 4–3 (Cont.) CIA Signal Timing Limits

CIA Signal	Timing Reference	Setup	Hold	Output Delay							
				Max 15 pF		Min 15 pF		Max 50 pF		Min 50 pF	
				Rise	Fall	Rise	Fall	Rise	Fall	Rise	Fall
ioc	CLK	—	—	5.3	5.1	1.2	1.2	7.2	6.6	2.0	1.8
iod	CLK	4.1	0.3	8.8	8.8	1.5	1.5	10.7	10.6	2.3	2.3
iod_e	CLK	3.0	0.1	9.1	9.2	1.6	1.6	11.0	11.0	2.4	2.4
irdy_l	CLK	6.0	-0.2	6.4	4.7	1.4	1.1	7.3	5.5	1.8	1.4
lock_l	CLK	3.7	-0.2	—	—	—	—	—	—	—	—
mem_ack_l	CLK	—	—	5.7	6.0	1.6	1.8	7.6	7.8	2.5	2.6
mem_cs_l	CLK	4.5	-0.8	—	—	—	—	—	—	—	—
mem_en	CLK2X	—	—	4.7	4.8	0.1	0.2	5.9	5.8	0.7	0.6
mem_req_l	CLK	4.0	-0.2	—	—	—	—	—	—	—	—
mem_we_l	CLK2X	—	—	4.7	4.9	0.1	0.2	5.9	5.8	0.6	0.6
par	CLK	5.1	-0.7	6.9	5.1	1.8	1.3	7.8	5.9	2.2	1.6
par64	CLK	5.1	-0.6	7.5	5.5	1.7	1.0	8.4	6.3	2.1	1.3
perr_l	CLK	7.2	0.0	7.7	5.7	1.7	1.1	8.7	6.5	2.1	1.5
ras ¹	CLK2X	—	—	5.3	5.3	0.3	0.3	6.5	6.3	0.8	0.7
rasx_l ² (set_sel)	CLK2X	—	—	5.3	5.5	0.0	0.0	6.5	6.5	0.4	0.4
req_l	CLK	—	—	4.9	4.6	1.3	1.2	5.9	5.3	1.7	1.5
req64_l	CLK	3.8	0.0	6.2	4.7	1.4	1.0	7.2	5.5	1.8	1.3
res	CLK2X	4.5	0.7	—	—	—	—	—	—	—	—
rst_l	CLK	—	—	6.7	5.8	2.3	2.0	7.6	6.5	2.7	2.3
serr_l	CLK	3.4	-0.2	—	—	—	—	—	—	—	—
set_sel ¹	CLK2X	—	—	10.8	9.7	1.9	1.6	12.0	10.6	2.4	2.0
stop_l	CLK	5.6	-0.4	6.3	4.7	1.3	1.1	7.3	5.4	1.7	1.4
tag_ctl_par	CLK	—	—	6.9	7.6	1.8	1.9	8.3	8.9	2.4	2.5
tag_dirty	CLK	—	—	6.9	7.6	1.8	1.9	8.3	8.9	2.4	2.5
trdy_l	CLK	6.2	-0.4	6.3	4.6	1.3	1.1	7.3	5.4	1.7	1.4
victim_pending	CLK2X	6.1	0.5	—	—	—	—	—	—	—	—

Notes:

- ¹ Set select memory mode
- ² Alternate memory mode

MK1455–26

21172–CA Control and Status Registers
4.6 21172–CA System Configuration Registers

Figure 4–4 DSW Signal Timing Limits

DSW Signal	Timing			Output Delay							
	Reference	Setup	Hold	Max 15pF		Min 15pF		Max 30pF		Min 30pF	
				Rise	Fall	Rise	Fall	Rise	Fall	Rise	Fall
iod	clk	4.7	0.0	13.0	12.9	2.3	2.5	14.9	14.7	3.1	3.2
mem_dat	clk	3.8	0.2	6.1	6.9	1.5	1.9	7.5	8.2	2.1	2.4
mem_dat (delayed)	clk	—	—	15.2	16.6	4.1	4.4	16.5	17.9	4.7	4.9
cmc	—	4.6	0.2	—	—	—	—	—	—	—	—
cpu_dat	—	3.5	0.1	6.3	5.6	0.9	1.0	8.2	7.1	1.7	1.6
ioc	—	6.6	-0.2	—	—	—	—	—	—	—	—
iod	ioc<6>	—	—	6.9	7.6	2.5	2.7	8.8	9.4	3.3	3.5
mem_dat	mem_en	—	—	6.9	7.9	2.5	2.9	8.3	9.2	3.1	3.4

MK1455–27

21172–CA Control and Status Registers
4.7 21172–CA PCI Address and Scatter-Gather Registers


4.7 21172–CA PCI Address and Scatter-Gather Registers

Sections 4.7.1 through 4.7.5 define and describe the 21172–CA system configuration registers.

4.7.1 Scatter-Gather Translation Buffer Invalidate Register

Register mnemonic: TBIA
 Register address: 87.6000.0100

Writing to the TBIA causes the specified group of scatter-gather TLB tags to be unlocked and marked invalid.



Field	Name	Type	Description
<01:00>	TBIA<1:0>	WO,0	A write transaction to this register will invalidate the specified scatter-gather translation buffers.
		TBIA<1:0>	Operation
		0 0	No operation.
		0 1	Invalidate and unlock the TLB tags that are currently locked.
		1 0	Invalidate the TLB tags that currently unlocked.
		1 1	Invalidate and unlock all TLB tags.
<31:02>	Reserved	RO,0	—

21172–CA Control and Status Registers
4.7 21172–CA PCI Address and Scatter-Gather Registers

4.7.2 Window Base Register

Register mnemonic: *W_n_BASE*
 Register address: See below

There are four window base registers, each providing the base address for a particular target window. The window base registers should not be modified unless software ensures that no PCI traffic is targeted for the window being modified.


The incoming PCI address bits <31:20> are compared with each of the four window base registers to determine if a hit occurs in the target window. The window mask register contents determine which bits are involved in the comparison (see Sections 6.4.4.1 and 6.4.4.2).

The target window is hit when the masked addresses match a valid window base register. If MEMCS_ENABLE is set, the hit is further qualified by the MEMCS input signal—this is used if PC compatibility holes are required in the CIA.

When W3_BASE[DAC_ENABLE] is set, W_DAC is used to compare PCI address bits <39:32> during a DAC cycle.

The addresses for the four window base address registers follow:

W0_BASE–87.6000.0400	W1_BASE–87.6000.0500
W2_BASE–87.6000.0600	W3_BASE–87.6000.0700



¹ W0_BASE only.
² W3_BASE only.

21172–CA Control and Status Registers
4.7 21172–CA PCI Address and Scatter-Gather Registers

Field	Name	Type	Description
<00>	W_EN	RW, UNDEFINED	0—The PCI target window is disabled and will not be used to respond to PCI-initiated transfers. 1—The PCI target window is enabled and is used to respond to PCI-initiated transfers that hit in the address range of the target window.
<01>	Wn_BASE_SG	RW, UNDEFINED	0—The PCI target window uses direct mapping to translate a PCI address to a 21164 address. See Table 4–20. 1—The PCI target window uses scatter-gather mapping to translate a PCI address to a physical memory address. See Table 4–19.
<02>	MEMCS_ENABLE ¹	RW, UNDEFINED	1—The MEMCS signal from the PCI-to-EISA bridge is ANDed with the normal window hit.
<03>	DAC_ENABLE ²	RW, UNDEFINED	1—W_DAC is compared with PCI address bits <39:32> for a PCI DAC cycle. If the compare is a hit, and the 32-bit portion of the PCI address is also hit, then a DAC cycle hit occurs.
<19:04>	Reserved	RO,0	—
<31:20>	W_BASE	RW, UNDEFINED	Specifies the PCI base address of the PCI target window and is used to determine a hit in the target window. See MEMCS_ENABLE and DAC_ENABLE.

¹Only in W0_BASE.
²Only in W3_BASE.

21172–CA Control and Status Registers
4.7 21172–CA PCI Address and Scatter-Gather Registers

4.7.3 Window Mask Register *n*

Register mnemonic: *W_n_MASK*
 Register address: See below


The window mask registers (*W_n_MASK*) provide a mask corresponding to **ad<31:20>** of an incoming PCI address. The size of each window can be programmed to be from 1MB to 4GB in powers of 2 by masking bits of an incoming PCI address with the window mask register bits <31:20>.

There are four window mask registers. The window mask registers should not be modified unless software ensures that no PCI traffic is targeted for the window being modified.

The incoming PCI address bits <31:20> are compared with each of the four window base registers to determine a hit in the target window. The window mask register contents determine which bits are involved in the comparison. The target window is hit when the masked addresses match a valid window base register.

The addresses of the window mask registers follow:

<i>W0_MASK</i> –87.6000.0440	<i>W1_MASK</i> –87.6000.0540
<i>W2_MASK</i> –87.6000.0640	<i>W3_MASK</i> –87.6000.0740



LJ-04253.AI

21172–CA Control and Status Registers
4.7 21172–CA PCI Address and Scatter-Gather Registers

Field	Name	Type	Description																														
<19:00>	Reserved	RO,0	—																														
<31:20>	W_MASK	RW, UNDEFINED	This field specifies the size of the PCI target window. It is also used to mask out address bits not used when determining a PCI target window hit.																														
			<table border="1"> <thead> <tr> <th>W_MASK<31:20></th> <th>Size of Window</th> </tr> </thead> <tbody> <tr><td>0000 0000 0000</td><td>1MB</td></tr> <tr><td>0000 0000 0001</td><td>2MB</td></tr> <tr><td>0000 0000 0011</td><td>4MB</td></tr> <tr><td>0000 0000 0111</td><td>8MB</td></tr> <tr><td>0000 0000 1111</td><td>16MB</td></tr> <tr><td>0000 0001 1111</td><td>32MB</td></tr> <tr><td>0000 0011 1111</td><td>64MB</td></tr> <tr><td>0000 0111 1111</td><td>128MB</td></tr> <tr><td>0000 1111 1111</td><td>256MB</td></tr> <tr><td>0001 1111 1111</td><td>512MB</td></tr> <tr><td>0011 1111 1111</td><td>1GB</td></tr> <tr><td>0111 1111 1111</td><td>2GB</td></tr> <tr><td>1111 1111 1111</td><td>4GB</td></tr> <tr><td>Other combinations</td><td>Not supported</td></tr> </tbody> </table>	W_MASK<31:20>	Size of Window	0000 0000 0000	1MB	0000 0000 0001	2MB	0000 0000 0011	4MB	0000 0000 0111	8MB	0000 0000 1111	16MB	0000 0001 1111	32MB	0000 0011 1111	64MB	0000 0111 1111	128MB	0000 1111 1111	256MB	0001 1111 1111	512MB	0011 1111 1111	1GB	0111 1111 1111	2GB	1111 1111 1111	4GB	Other combinations	Not supported
W_MASK<31:20>	Size of Window																																
0000 0000 0000	1MB																																
0000 0000 0001	2MB																																
0000 0000 0011	4MB																																
0000 0000 0111	8MB																																
0000 0000 1111	16MB																																
0000 0001 1111	32MB																																
0000 0011 1111	64MB																																
0000 0111 1111	128MB																																
0000 1111 1111	256MB																																
0001 1111 1111	512MB																																
0011 1111 1111	1GB																																
0111 1111 1111	2GB																																
1111 1111 1111	4GB																																
Other combinations	Not supported																																

21172–CA Control and Status Registers
4.7 21172–CA PCI Address and Scatter-Gather Registers

4.7.4 Translated Base Register *n*


Register mnemonic: *T_n_BASE*
 Register address: See below

There are four translated base registers (*T_n_BASE*), one for each window. They are used to map PCI addresses into memory addresses. The addresses of the four translated base registers follow:

T0_BASE–87.6000.0480 *T1_BASE*–87.6000.0580
T2_BASE–87.6000.0680 *T3_BASE*–87.6000.0780

If the scatter-gather bit of the window base register is set, the *T_n_BASE* provides the base address of the scatter-gather map for this window. If the scatter-gather bit is clear, the *T_n_BASE* provides the base physical address of this window.

The *T_n_BASE* register should not be modified unless software can ensure that no PCI traffic is targeted for the window being modified.



Field	Name	Type	Description
<07:00>	Reserved	RO,0	—
<31:08>	<i>T_BASE</i> <33:10>	RW, UNDEFINED	If scatter-gather mapping is enabled, this field specifies the base 21164 address of the scatter-gather map table for the PCI target window. See Table 4–19. If scatter-gather mapping is disabled, this field specifies the base 21164 address of the translated PCI address for the PCI target window. See Table 4–20.

21172–CA Control and Status Registers

4.7 21172–CA PCI Address and Scatter-Gather Registers

W_MASK<31:20> sets the size of the PCI target window and the number of 8KB pages that fall into the window. Every 8KB page requires one 8-byte scatter-gather map entry. Table 4–19 shows the relationship of the W_MASK field to the size of the scatter-gather map in memory.

$$\frac{\text{Size of the Window in Bytes}}{8KB} = \text{Number of Entries Required}$$

$$\text{Number of Entries} * 8 \text{ bytes} = \text{Size of the Scatter – Gather Table}$$

Table 4–19 PCI Address Translation—Scatter-Gather Mapping Enabled

W_MASK<31:20>	S-G Map Table Size	Scatter-Gather Map Address addr_h<33:0> (Used to Index S-G Table in Memory)
0000 0000 0000	1KB	T_BASE<33:10> : ad_h<19:13> : 000
0000 0000 0001	2KB	T_BASE<33:11> : ad_h<20:13> : 000
0000 0000 0011	4KB	T_BASE<33:12> : ad_h<21:13> : 000
0000 0000 0111	8KB	T_BASE<33:13> : ad_h<22:13> : 000
0000 0000 1111	16KB	T_BASE<33:14> : ad_h<23:13> : 000
0000 0001 1111	32KB	T_BASE<33:15> : ad_h<24:13> : 000
0000 0011 1111	64KB	T_BASE<33:16> : ad_h<25:13> : 000
0000 0111 1111	128KB	T_BASE<33:17> : ad_h<26:13> : 000
0000 1111 1111	256KB	T_BASE<33:18> : ad_h<27:13> : 000
0001 1111 1111	512KB	T_BASE<33:19> : ad_h<28:13> : 000
0011 1111 1111	1MB	T_BASE<33:20> : ad_h<29:13> : 000
0111 1111 1111	2MB	T_BASE<33:21> : ad_h<30:13> : 000
1111 1111 1111	4MB	T_BASE<33:22> : ad_h<31:13> : 000

Note

Unused translated base must be zero for correct operation.

21172–CA Control and Status Registers
4.7 21172–CA PCI Address and Scatter-Gather Registers

The quadword address used to index into the table is formed from concatenating the appropriate T_BASE and PCI address bits based on the size of the scatter-gather map. The PCI address forms the index into the table, while T_BASE forms the NATURALLY ALIGNED base of the table.

For example, there are 128 entries in the scatter-gather table, with a table size of 1KB, for a mask of 0000 0000 0000. Entries are quadwords, so the lower 3 bits, <2:0>, of the address are always zero. Now, mask off **ad<31:20>** because of W_MASK. Then use **ad<19:13>** (7 bits, $2^7 = 128$ entries in the table) as the table index. Use T_BASE <31:08> to get the other bits of the 34-bit address.

Table 4–20 PCI Address Translation—Scatter-Gather Mapping Disabled

W_MASK<31:20>	Translated Address addr_h<33:0>	Unused Translated Base Register Bits
0000 0000 0000	T_BASE<33:20> : ad_h<19:0>	T_BASE<19:10>
0000 0000 0001	T_BASE<33:21> : ad_h<20:0>	T_BASE<20:10>
0000 0000 0011	T_BASE<33:22> : ad_h<21:0>	T_BASE<21:10>
0000 0000 0111	T_BASE<33:23> : ad_h<22:0>	T_BASE<22:10>
0000 0000 1111	T_BASE<33:24> : ad_h<23:0>	T_BASE<23:10>
0000 0001 1111	T_BASE<33:25> : ad_h<24:0>	T_BASE<24:10>
0000 0011 1111	T_BASE<33:26> : ad_h<25:0>	T_BASE<25:10>
0000 0111 1111	T_BASE<33:27> : ad_h<26:0>	T_BASE<26:10>
0000 1111 1111	T_BASE<33:28> : ad_h<27:0>	T_BASE<27:10>
0001 1111 1111	T_BASE<33:29> : ad_h<28:0>	T_BASE<28:10>
0011 1111 1111	T_BASE<33:30> : ad_h<29:0>	T_BASE<29:10>
0111 1111 1111	T_BASE<33:31> : ad_h<30:0>	T_BASE<30:10>
1111 1111 1111	T_BASE<33:32> : ad_h<31:0>	T_BASE<31:10>

21172–CA Control and Status Registers
4.7 21172–CA PCI Address and Scatter-Gather Registers

4.7.5 Window Base DAC Register

Register mnemonic: W_DAC
 Register address: 87.6000.07C0


W_DAC provides DAC_BASE<7:0> for comparison against PCI address bits <39:32> during a DAC cycle. PCI address bits <63:40> must be zero for a PCI window hit. This register is used with W3_BASE.

W_DAC is only applicable to window 3 and only if enabled by W3_BASE[DAC_ENABLE].

Determining a Hit in the Target Window

The target window is hit when the following is satisfied:

- The incoming PCI address bits <31:20> match one of the four window base registers. Wn_MASK[W_MASK]<31:20> determines which bits are involved in the comparison.
- PCI address bits <63:40> are zero.
- PCI address bits <39:32> match DAC_BASE<7:0>.



LJ-04255.AI

Field	Name	Type	Description
<07:00>	DAC_BASE<7:0>	RW, UNDEFINED	Specifies bits <39:32> of the PCI base address used to determine a hit in the target window for a DAC cycle.
<31:08>	Reserved	RO,0	—

21172-CA Control and Status Registers
4.8 21172-CA Address Translation Registers

4.8 21172-CA Address Translation Registers

Sections 4.8.1 through 4.8.3 define and describe the 21172-CA address translation registers.

4.8.1 Lockable Translation Buffer Tag *n* Registers

Register mnemonic: LTB_TAG n
 Register address: See below

Software can write to these four lockable translation buffer tag n registers. Hardware cannot evict the register contents when a scatter-gather TLB miss occurs if the LOCKED bit is set. The addresses of the LTB_TAG n registers follow:


LTB_TAG0-87.6000.0800 LTB_TAG1-87.6000.0840
 LTB_TAG2-87.6000.0880 LTB_TAG3-87.6000.08C0

Determining a Hit in the Translation Buffer

If $Wn_BASE[Wn_BASE_SG]$ is true when a PCI address hits one of the Wn_BASE registers, the incoming PCI address bits <31:15> are compared with each of the eight translation buffer tag registers. If there is a match, the corresponding translation buffer page register group is indexed by PCI address bits <14:13> and if it is valid, there is a translation buffer hit.

Operation on an Scatter-Gather TLB Miss

Hardware uses a round-robin algorithm to handle a scatter-gather TLB miss. A scatter-gather TLB entry is overwritten if it is not locked, but will not be evicted if it is locked. The hardware will write all four PTEs on a miss.



LJ-04256.AI

21172–CA Control and Status Registers
4.8 21172–CA Address Translation Registers

Field	Name	Type	Description
<00>	VALID	RW,0	If both VALID and the corresponding $Wn_BASE[Wn_BASE_SG]$ are set, this entry will be used for address translation.
<01>	LOCKED	RW,0	1—The hardware will not evict this entry.
<02>	DAC	RW,0	0—This TAG entry belongs to a 32-bit PCI address. 1—This tag entry corresponds to a 64-bit PCI address.
<14:03>	Reserved	RO,0	—
<31:15>	TB_TAG	RW, UNDEFINED	The TAG for each translation buffer entry.

21172–CA Control and Status Registers
4.8 21172–CA Address Translation Registers

4.8.2 Translation Buffer Tag *n* Register

Register mnemonic: TB_TAG*n*
 Register address: See below


The addresses for the four TB_TAG*n* registers follow:

TB_TAG0–87.6000.0900 TB_TAG1–87.6000.0940
 TB_TAG2–87.6000.0980 TB_TAG3–87.6000.09C0

Software can write to these TLB tag entries, but the entries cannot be locked and so may be evicted by the hardware on a scatter-gather TLB miss.

The incoming PCI address bits <31:15> are compared with each of the eight translation buffer tag registers. If there is a match, the corresponding translation buffer page register group is indexed by PCI address bits <14:13>, and if it is valid, then there is a translation buffer hit.

A scatter-gather TLB miss is handled by hardware using a round-robin algorithm. An entry is overwritten if it is not locked. The hardware will write all four PTEs in the event of a scatter-gather TLB miss.



LJ-04257.AI

Field	Name	Type	Description
<00>	VALID	RW,0	If both VALID and the corresponding W <i>n</i> _BASE[W <i>n</i> _BASE_SG] are set, this entry is used for address translation.
<01>	Reserved	RO,0	—
<02>	DAC	RW,0	0—This TAG entry belongs to a 32-bit PCI address. 1—This TAG entry corresponds to a 64-bit PCI address.
<14:03>	Reserved	RO,0	—

21172–CA Control and Status Registers
4.8 21172–CA Address Translation Registers

Field	Name	Type	Description
<31:15>	TB_TAG	RW, UNDEFINED	The TAG for each translation buffer entry.

21172–CA Control and Status Registers
4.8 21172–CA Address Translation Registers

4.8.3 Translation Buffer *m* Page *n* Registers

Register mnemonic: $TB_m_PAGE_n$
Register address: See below


There are 32 translation buffer data registers: a group of four for each of the eight translation buffer entries. The $TB_m_PAGE_n$ registers are automatically updated on a TLB miss (a group of four at a time) by the CIA hardware.

The addresses for the 32 $TB_m_PAGE_n$ registers follow:

TB0_PAGE0–87.6000.1000	TB0_PAGE1–87.6000.1040
TB0_PAGE2–87.6000.1080	TB0_PAGE3–87.6000.10C0
TB1_PAGE0–87.6000.1100	TB1_PAGE1–87.6000.1140
TB1_PAGE2–87.6000.1180	TB1_PAGE3–87.6000.11C0
TB2_PAGE0–87.6000.1200	TB2_PAGE1–87.6000.1240
TB2_PAGE2–87.6000.1280	TB2_PAGE3–87.6000.12C0
TB3_PAGE0–87.6000.1300	TB3_PAGE1–87.6000.1340
TB3_PAGE2–87.6000.1380	TB3_PAGE3–87.6000.13C0
TB4_PAGE0–87.6000.1400	TB4_PAGE1–87.6000.1440
TB4_PAGE2–87.6000.1480	TB4_PAGE3–87.6000.14C0
TB5_PAGE0–87.6000.1500	TB5_PAGE1–87.6000.1540
TB5_PAGE2–87.6000.1580	TB5_PAGE3–87.6000.15C0
TB6_PAGE0–87.6000.1600	TB6_PAGE1–87.6000.1640
TB6_PAGE2–87.6000.1680	TB6_PAGE3–87.6000.16C0
TB7_PAGE0–87.6000.1700	TB7_PAGE1–87.6000.1740
TB7_PAGE2–87.6000.1780	TB7_PAGE3–87.6000.17C0

21172–CA Control and Status Registers

4.8 21172–CA Address Translation Registers



LJ-04258.AI

Field	Name	Type	Description
<00>	VALID	RW, UNDEFINED	1—This entry can be used for address translation.
<21:01>	PAGE_ADDRESS	RW, UNDEFINED	This field forms physical address<33:13>. PCI ad <12:0> forms physical address<12:0>.
<31:22>	Reserved	RO,0	—

Incoming PCI address bits <31:15> are compared with each of the eight translation buffer tag registers. If there is a match, the corresponding translation buffer page register group is indexed using PCI address bits <14:13>, and if it is valid, there is a translation buffer hit.

If the address bits do not match the tag, or the page entry is invalid, a TLB miss occurs. If the PTE, fetched by the hardware TLB-miss handler, is still invalid, the CIA_ERROR interrupt is asserted for a DMA write transaction.

21172–CA Power-Up and Initialization

This chapter describes the behavior of the 21172–CA (CIA) chip on power-up and assertion of **sys_rst_1**. It also describes the system level requirements and the various registers that must be initialized after **sys_rst_1** is deasserted.

5.1 Power-Up

On power-up, the **sys_rst_1** input of the CIA should be asserted. It should be kept asserted until the system clocks are up and running for eight cycles.

5.2 Internal Reset

The assertion and deassertion of the **sys_rst_1** pin on the module is asynchronous to the CIA. The CIA generates an internal reset signal from **sys_rst_1**. The internal reset signal is asserted asynchronously, as soon as **sys_rst_1** is asserted, but is deasserted synchronously. Because of the synchronous deassertion of the internal reset, the CIA requires that no external transaction start until eight system clock cycles after the deassertion of **sys_rst_1**.

21172—CA Power-Up and Initialization
 5.3 State of Pins on Reset Assertion

5.3 State of Pins on Reset Assertion

The following are general rules and requirements for the behavior of the CIA pins during reset:

- All input control signals (except the clocks and **sys_rst_l**) must be in the deasserted state as long as **sys_rst_l** is asserted.
- All output and bidirectional signals are in the states listed in Table 5–1.

Table 5–1 lists the state of CIA pins while **sys_rst_l** is asserted.

Table 5–1 CIA Pin State During Reset

Signal	Reset State	Signal	Reset State
ack64_l	Z	ad<63:0>	Z
addr<39:4>	Z	addr_bus_req	Z
addr_cmd_par	Z	cack	Z
cas<3:0>	0	cbe_l<7:0>	Z
cmc<8:0>	0	cmd<3:0>	Z
dack	Z	devsel_l	Z
error	0	fill	Z
fill_error	Z	fill_id	Z
frame_l	Z	idle_bc	Z
int	0	ioc<6:0>	5 ₂
iod<63:0>	Z	iod_e<7:0>	Z
irdy_l	Z	mem_ack_l	x ¹
mem_addr<12:0>	x	mem_b0	x
mem_en	1	mem_we_l<1:0>	11 ₂
par	Z	par64	Z
perr_l	Z	ras<3:0>	0
req_l	1	req64_l	Z
rst_l	0	serr_l	Z
set_sel<15:0>	0	stop_l	Z
tag_ctl_par	Z	tag_dirty	Z
test_out	0	trdy_l	Z

¹Dependent upon the state of input pin **mem_req_l**.
 During reset **mem_ack_l** = **mem_req_l**.

21172–CA Power-Up and Initialization
5.4 Configuration After Reset Deassertion

5.4 Configuration After Reset Deassertion

Table 5–2 lists the state of CIA pins after **sys_rst_l** has been deasserted.

Table 5–2 CIA Pin State After Reset

Signal	Reset State	Signal	Reset State
ack64_l	Z	ad<63:32>	Z
ad<31:0>	x or Z ¹	addr<39:4>	Z
addr_bus_req	0	addr_cmd_par	Z
cack	0	cas<3:0>	0
cbe_l<7:4>	Z	cbe_l<3:0>	x or Z ¹
cmc<8:0>	0	cmd<3:0>	Z
dack	0	devsel_l	Z
error	0	fill	0
fill_error	0	fill_id	0
frame_l	Z	idle_bc	0
int	0	ioc<6:0>	52
iod<63:0>	Z	iod_e<7:0>	Z
irdy_l	Z	mem_ack_l	x ²
mem_addr<12:0>	x	mem_b0	x
mem_en	1	mem_we_l<1:0>	11 ₂
par	x or Z ¹	par64	>
perr_l	Z	ras<3:0>	0
req_l	1	req64_l	0
rst_l	0	serr_l	Z
set_sel<15:0>	0	stop_l	Z
tag_ctl_par	Z	tag_dirty	Z
test_out	0	trdy_l	Z

¹Dependent upon the state of PCI **gnt_l**.
If **gnt_l** is low, (CIA granted use of the PCI) the state of these signal lines is x.
If **gnt_l** is high, (CIA not granted use of the PCI) the state of these signal lines is Z.

²Dependent upon the state of input pin **mem_req_l**.
After reset **mem_ack_l** = **mem_req_l**.

21172–CA Power-Up and Initialization

5.4 Configuration After Reset Deassertion

After `sys_rst_1` has been deasserted, the 21164 reads instructions from the serial ROM (SROM) into its internal Icache and then executes those instructions. The instructions from the SROM initialize the registers listed in Table 5–3 to the values shown.

Table 5–3 Registers Initialized by SROM Code

Register	Acronym	Initialized Value (Hex)
Memory configuration register	MCR	Depends upon memory configuration.
Memory base address registers 0,2,4,6,8,A,C,E	MBA n	Depends upon memory configuration.
Memory timing registers 0–2	TMG0, TMG1, TMG2	Depends upon memory configuration.

6

21172–CA Address Mapping

This chapter describes 21172–CA (CIA) mapping of 40-bit 21164 physical addresses to memory addresses and I/O space addresses. It also describes translation of a 21164-initiated address into a PCI address and translation of a PCI-initiated address into a physical memory address.

Topics include dense and sparse address space¹, PCI addressing, scatter-gather address translation for DMA operations, and Extended Industry Standard Architecture (EISA) requirements.

6.1 Address Mapping Overview

The 21164 address space is divided into two regions, as shown in Figure 6–1, using physical address bit **addr<39>**. When clear, the 21164 access is to cacheable memory space (partly reserved). When set, the 21164 access is to noncacheable address space. The noncached address space is used for the CSRs, uncached diagnostic memory access, and to access the memory-mapped PCI I/O address space.

The PCI defines three physical address spaces:

- A 64-bit PCI memory space
- A 4GB PCI I/O space
- A 256-byte-per-device PCI configuration space

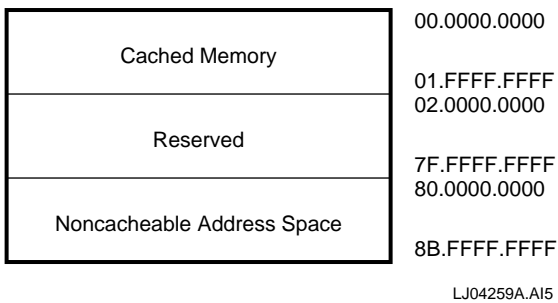
In addition to these three PCI address spaces, the 21164's noncached space is also used to generate PCI interrupt acknowledge and special cycles.

¹ Dense and sparse address spaces are defined in Section 6.3.1 and Section 6.3.2, respectively.

21172–CA Address Mapping

6.1 Address Mapping Overview

Figure 6–1 21164 Address Space



6.2 21164 Address Space Configuration Supported by the CIA

As shown in Figure 6–2, the CIA supports only the first 8GB of cacheable memory space; the remainder is reserved. The cacheable memory space block size is fixed at 64 bytes. The CIA will send READ and FLUSH commands to the 21164 caches for DMA traffic to the 8GB memory address area.


The CIA supports 21164 access to memory-mapped I/O devices in noncacheable address space. The CIA defines five address spaces within the noncacheable space.

- 22GB PCI memory sparse space
- 2GB PCI I/O space
- 4GB PCI memory dense space
- 4GB that includes address space for:
 - PCI configuration
 - Special/interrupt acknowledge cycles
 - CIA control and status registers (CSR)
 - Flash ROM and support logic registers
- 16GB PCI byte/word I/O space

21172-CA Address Mapping

6.2 21164 Address Space Configuration Supported by the CIA

Figure 6-2 21164 Address Space Configuration



LJ-04868.A15

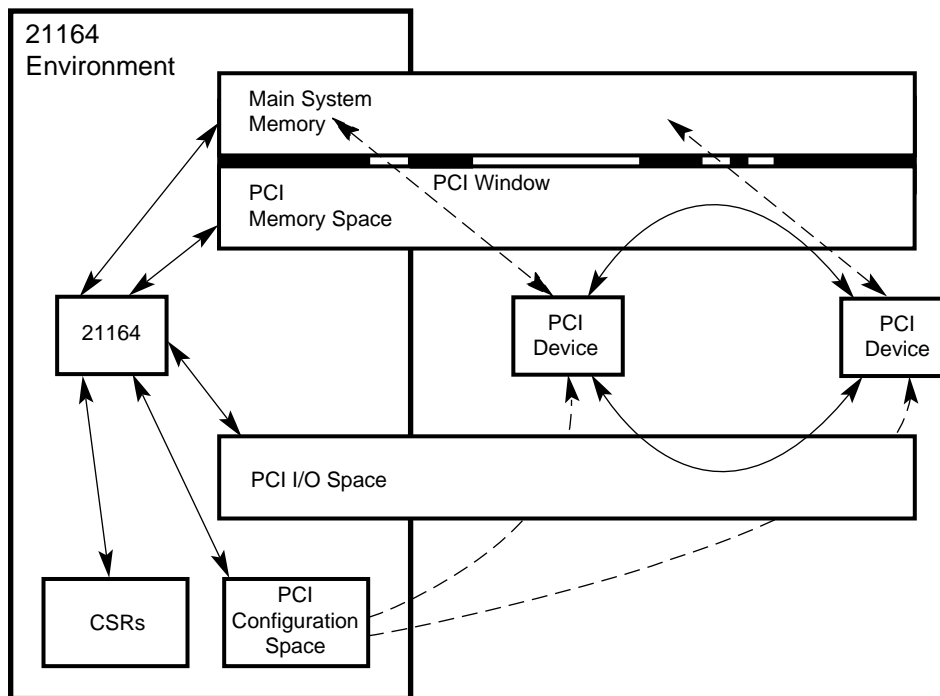
21172-CA Address Mapping

6.2 21164 Address Space Configuration Supported by the CIA

6.2.1 21164 Access to Address Space

The 21164 has access to the complete address space: it can access cached memory and CSRs, as well as all the PCI memory, I/O, and configuration regions, as shown in Figure 6-3.

Figure 6-3 Address Space Overview



LJ-04261.AI

6.2.2 PCI Access to Address Space

PCI devices have a restricted view of the address space. They can access any PCI device through the PCI memory or PCI I/O space, but they have no access to PCI configuration space.

Furthermore, the CIA restricts PCI access to the system memory (for DMA operations) through five programmable address windows in PCI memory space, as shown in Figure 6-3.


21172-CA Address Mapping

6.2 21164 Address Space Configuration Supported by the CIA

DMA access to system memory is achieved by means of windows through one of three access methods:

- The direct method uses the “Monster Window” with dual address cycles where PCI address <33:0> equals memory address <33:0>.
- The directly mapped method concatenates an offset to a portion of the PCI address.
- The virtually mapped method uses a scatter-gather translation map. The scatter-gather map allows any 8KB page of PCI memory address region (page) to be redirected to any 8KB cached memory page, as shown in Figure 6-4.

Figure 6-4 Address Mapping Overview



LJ04220A.A15

The CIA generates 32-bit PCI addresses but accepts both 64-bit address cycles (DAC²) and 32-bit PCI address cycles (SAC³). The following window restrictions apply to PCI main memory accesses:

² Dual address cycle (PCI 64-bit address transfer)—used only if address bits **ad<63:32>** are nonzero.

³ Single address cycle (PCI 32-bit address transfer).

21172–CA Address Mapping

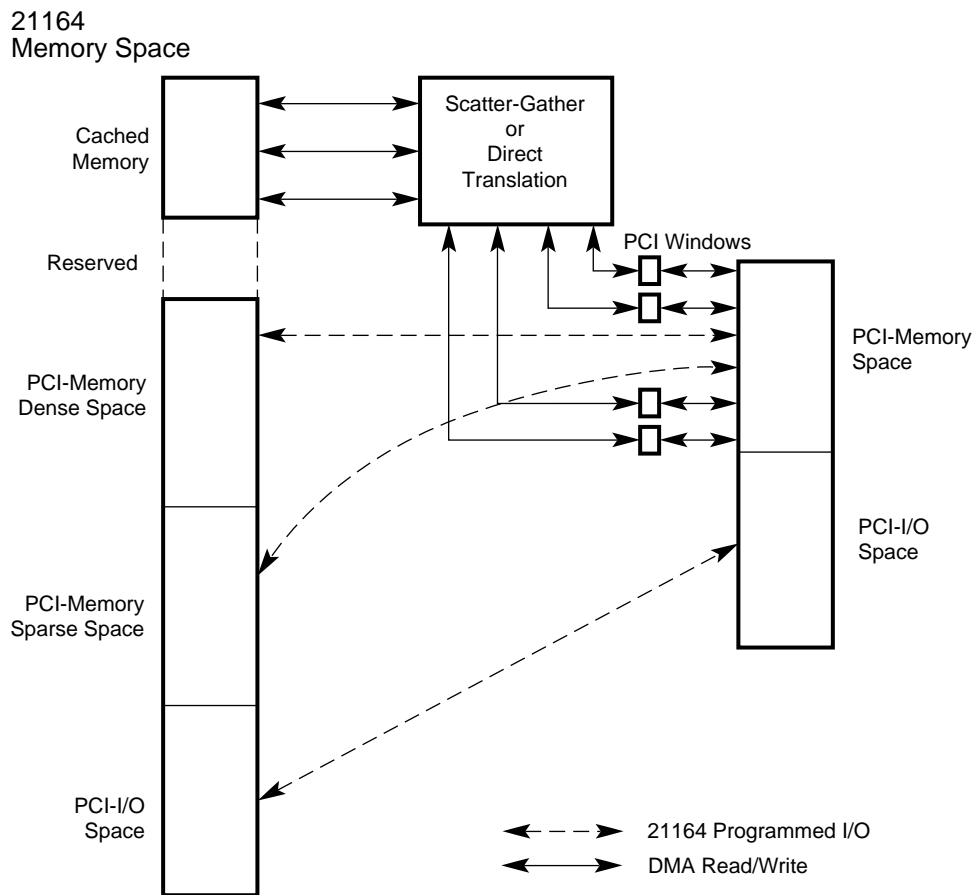
6.2 21164 Address Space Configuration Supported by the CIA

- Window 4, the “Monster Window,” provides full access to main memory. It is accessed by DAC cycles only with PCI address bit <40>=1. Memory address bits <33:0> equal PCI address bits <33:0>.
- Window 3 can be accessed either by DAC or SAC cycles, but not both. If using DAC cycles, the following three restrictions apply:
 1. PCI address bits <63:40> must be zero
 2. PCI address bits <39:32> must match the DAC register
 3. PCI address bits <31:0> must hit in window 3.
- Windows 0, 1, and 2 accept only SAC cycles.

Figure 6–5 shows how the 21164 address map translates to the PCI address space. It also shows how PCI devices access 21164 memory space via DMAs. Notice that the PCI memory space is double mapped using dense and sparse space.

21172-CA Address Mapping
6.2 21164 Address Space Configuration Supported by the CIA

Figure 6-5 21164 and DMA Read and Write Transactions



LJ04263A.A15

21172–CA Address Mapping

6.3 21164 Address Space

6.3 21164 Address Space

This section lists and describes the 21164 address space regions. The requirements for using the address regions are also shown and described. Table 6–1 lists the 21164 address regions.

Table 6–1 21164 Physical Address Space

21164 Address	Size (GB)	Selection
00.0000.0000–01.FFFF.FFFF	8	Main memory
80.0000.0000–83.FFFF.FFFF	16	PCI memory—512MB Sparse Space—Region 0
84.0000.0000–84.FFFF.FFFF	4	PCI memory—128MB Sparse Space—Region 1
85.0000.0000–85.7FFF.FFFF	2	PCI memory—64MB Sparse Space—Region 2
85.8000.0000–85.BFFF.FFFF	1	PCI I/O space—32MB Sparse Space—Region A
85.C000.0000–85.FFFF.FFFF	1	PCI I/O space—32MB Sparse Space—Region B
86.0000.0000–86.FFFF.FFFF	4	PCI memory, 4GB—Dense space
87.0000.0000–87.1FFF.FFFF	0.5	PCI configuration Sparse space
87.2000.0000–87.3FFF.FFFF	0.5	PCI special/interrupt acknowledge Sparse space
87.4000.0000–87.4FFF.FFFF	0.25	CIA main CSRs Pseudospase ¹
87.5000.0000–87.5FFF.FFFF	0.25	CIA PCI memory control CSRs Pseudospase ¹
87.6000.0000–87.6FFF.FFFF	0.25	CIA PCI address translation Pseudospase ¹
87.7000.0000–87.7FFF.FFFF	0.25	Reserved
88.0000.0000–88.FFFF.FFFF	4	21164 byte/word PCI memory space
89.0000.0000–89.FFFF.FFFF	4	21164 byte/word PCI I/O space
8A.0000.0000–8A.FFFF.FFFF	4	21164 byte/word PCI configuration space—type 0
8B.0000.0000–8B.FFFF.FFFF	4	21164 byte/word PCI configuration space—type 1

¹Pseudospase space is a hardware-specific restricted version of sparse space.

21172–CA Address Mapping

6.3 21164 Address Space

The reasons for using the 21164 I/O space address map are as follows:

- Provides 4GB of dense space to completely map the 32-bit PCI memory space.
- Provides abundant PCI sparse memory space because sparse memory space has byte granularity and is the safest memory space to use (no prefetching). Furthermore, the larger the space, the less likely software will need to dynamically relocate the sparse space segments. The main problem with sparse space is that it is wasteful of 21164 address space. For instance, 16GB of 21164 address space maps to only 512MB of PCI sparse space.

The CIA supports three PCI sparse space memory regions, allowing 704MB of total sparse memory space. The three regions can be relocated using the HAE_MEM CSR, and the simplest configuration allows for 704MB of contiguous memory space:

- 512MB region, which may be located in any NATURALLY ALIGNED 512MB segment of the PCI memory space. Software developers may find this region sufficient for their needs and can ignore the remaining two regions.
- 128MB region, which may be located on any NATURALLY ALIGNED 128MB segment of the PCI memory space.
- 64MB region, which may be located on any NATURALLY ALIGNED 64MB segment of the PCI memory space.
- Limits the PCI I/O space to sparse space. Although the PCI I/O space can handle 4GB, some chips can access only 64KB. So most, if not all, PCI devices will not exceed 64KB for the foreseeable future. Therefore, the CIA supports 64MB of sparse I/O space.
- Supports two CIA PCI sparse space I/O regions. Region A contains 32MB and is fixed in PCI segment 0–32 MB. Region B also contains 32MB, but can be relocated using the HAE_IO register.

6.3.1 PCI Dense Memory Space

PCI dense memory space is located in the range 86.0000.0000 to 86.FFFF.FFFF. It is typically used for memory-like data buffers such as video frame buffers or nonvolatile RAM (NVRAM). Dense space does not allow byte or word access but has the following advantages over sparse space:

- Contiguous memory locations—Some software, like the default graphics routines of Windows NT, require memory-like access. These routines cannot use sparse space addresses, because sparse space addresses require

21172–CA Address Mapping

6.3 21164 Address Space

PCI transactions to be at adjacent Alpha addresses, instead of being widely separated as in sparse space. As a result, if the user-mode driver uses sparse space for its frame-buffer manipulation, it cannot “hand over” the buffer to the common Windows NT graphics code.

- Higher bus bandwidth—PCI bus burst transfers are not usable in sparse space except for a 2-longword burst for quadword write transactions. Dense space is defined to allow both burst read and write transactions.
- Efficient read/write buffering—In sparse space, separate accesses use separate read or write buffer entries. Dense space allows separate accesses to be collapsed in read and write buffers (this is exactly what the 21164 does).
- Few memory barriers—In general, sparse space accesses are separated by memory barriers to avoid read/write buffer collapsing. Dense space accesses only require barriers when explicit ordering is required by the software.

Dense space is provided for the 21164 to access PCI memory space, but not for accessing PCI I/O space. Dense space has the following characteristics:

- There is a one-to-one mapping between 21164 addresses and PCI addresses. A longword address from the 21164 will map to a longword on the PCI with no shifting of the address field.
- The concept of dense space (and sparse space) is applicable only to 21164-generated addresses. There is no such thing as dense space (or sparse space) for PCI-generated address.
- Byte or word accesses are *not* possible in this space. The minimum access granularity is a longword on write transactions and a quadword on read transactions. The maximum transfer length is 32 bytes (performed as a burst of 8 longwords on the PCI). Any combination of longwords may be valid on write transactions. Valid longwords surrounding invalid longwords (called a “hole”) are required to be handled correctly by all PCI devices. The CIA will allow such “holes” to be issued.
- Read transactions will always be performed as a burst of two or more longwords on the PCI because the minimum granularity is a quadword. The processor can request a longword but the CIA will always fetch a quadword, thus prefetching a second longword. Therefore, this space cannot be used for devices that have read side effects. Although a longword may be prefetched, the prefetch buffer is not treated as a cache and thus coherency is not an issue. A quadword read transaction is not atomic on the PCI, that is, the target device is at liberty to force a retry after the first longword of data is sent, and then allow another device to take control

21172–CA Address Mapping

6.3 21164 Address Space

- For a write transaction, **ad<2>** equals **addr<2>**.

6.3.2 PCI Sparse Memory Space

The CIA supports three regions of contiguous 21164 address space that maps to PCI sparse memory space. The total 21164 range is from 80.0000.0000 to 85.7FFF.FFFF. Sparse address space maps a large piece of 21164 memory address space to a small PCI address space. For example, a 32-byte memory address might map to a single-byte PCI address.

A problem arises because the Alpha instruction set can express only aligned longword and quadword data references. The PCI bus requires the ability to express byte, word, tribyte, longword, and quadword references. The CIA must also be able to emulate PCI transactions for PCI devices designed for systems that are capable of generating the UNALIGNED references.

The CIA accomplishes UNALIGNED PCI references by encoding the size of the data transfer (byte, word, and so on) and the byte enable information in the 21164 address. Address bits **addr<6:3>** are used for this purpose. The PCI longword address **ad<26:3>** is generated using remaining address bits **addr<31:7>**.

Quadword address encoding is provided by **addr<6:3>** with **addr<7>** assumed to be zero by the CIA hardware. See Table 6–3.

The loss of address bits **addr<6:3>** has resulted in a 22GB “sparse” address space that maps to only 704MB of address space on the PCI.

The rules for accessing sparse space follow:

- Sparse space supports all the byte encodings that are expected to be generated in a system to ensure compatibility with existing and expected PCI devices/drivers. The results of some references are not explicitly defined (these are the missing entries in Table 6–3, such as word size with address **addr<6:5>** = 11). The hardware will complete the reference, but the reference is not required to produce any particular result nor will the CIA report an error.
- Software must use longword load or store instructions (LDL/STL) to perform a reference that is of longword length or less on the PCI bus.
 - The bytes to be transferred must be positioned within the longword in the correct byte lanes as indicated by the PCI byte enable.
 - The hardware will do no byte shifting within the longword.
- Quadword load and store instructions must be used only to perform quadword transfers. Use of STQ/LDQ instructions for any other references will produce UNPREDICTABLE results.

21172–CA Address Mapping
6.3 21164 Address Space

- Read-ahead (prefetch) is not performed in sparse space by the CIA hardware because the read-ahead might have unwanted side effects.
- Programmers are required to insert memory barrier (MB) instructions between sparse space accesses to prevent collapsing in the 21164 write buffer. However, this is not always necessary. For instance, consecutive sparse space addresses will be separated by 32 bytes (and will not be collapsed by the 21164).
- Programmers are required to insert MB instructions if the sparse space address ordering/coherency to a dense space address is to be maintained.
- On read transactions, the 21164 sends out **addr<4:3>** indirectly on the **int4_valid_h<3:0>**.
- Accesses with **addr<2:0>** nonzero will produce UNPREDICTABLE results.
- The relationship between **int4_valid_h<3:0>** and 21164 **addr<4:3>** for a sparse space write transaction is shown in Table 6–2. The important point is that all other **int4_valid_h<3:0>** patterns will produce UNPREDICTABLE results such as the result of collapsing in the 21164 write buffer or issuing an STQ instruction when an STL instruction was required.

Table 6–2 int4_valid_h<3:0> and addr<4:3> for Sparse Space Write Transactions

21164 Data Cycle	int4_valid_h<3:0> ¹	addr_h<4:3>
First	0 0 0 1	0 0
	0 0 1 0	0 0
	0 1 0 0	0 1
	1 0 0 0	0 1
Second	0 0 0 1	1 0
	0 0 1 0	1 0
	0 1 0 0	1 1
	1 0 0 0	1 1
	1 1 0 0 (STQ) ²	1 1

¹All other **int4_valid_h<3:0>** patterns cause UNPREDICTABLE results.

²Only one STQ case is allowed.

Table 6–3 defines the low-order PCI sparse memory address bits. Address bits **addr<7:3>** are used to generate the length of the PCI transaction in bytes, the byte enables, and **ad<2:0>**. Address bits **addr<30:8>** correspond to the quadword PCI address and are sent out on the PCI as **ad<25:3>**.

21172–CA Address Mapping
6.3 21164 Address Space

Table 6–3 PCI Memory Sparse Space Read/Write Encodings

Size	Byte Offset	21164	PCI	Data-In Register	
addr<4:3>	addr<6:5> ¹	Instruction	ad<2:0> ²	Byte ³ Enable	
Byte					
00	00	LDL, STL	addr<7>,00	1110	<0>
	01			1101	<1>
	10			1011	<2>
	11			0111	<3>
Word					
01	00	LDL, STL	addr<7>,00	1100	<1:0>
	01			1001	<2:1>
	10			0011	<3:2>
Tribyte					
10	00	LDL, STL	addr<7>,00	1000	<2:0>
	01			0001	<3:1>
Longword					
11	00	LDL, STL	addr<7>,00	0000	<3:0>
Quadword					
11	11	LDQ, STQ	000	0000	<7:0>

¹Missing entries (such as word size with **addr<6:5>** = 11₂) cause UNPREDICTABLE results.

²In PCI sparse memory space, **ad<1:0>** is always equal to zero.

³Byte enable set to zero indicates that byte lane carries meaningful data.

The high-order PCI address bits **ad<31:26>** are obtained from either the hardware address extension register (HAE_MEM) or the 21164 address, depending on sparse space regions, as shown in Table 6–4.

21172-CA Address Mapping
6.3 21164 Address Space

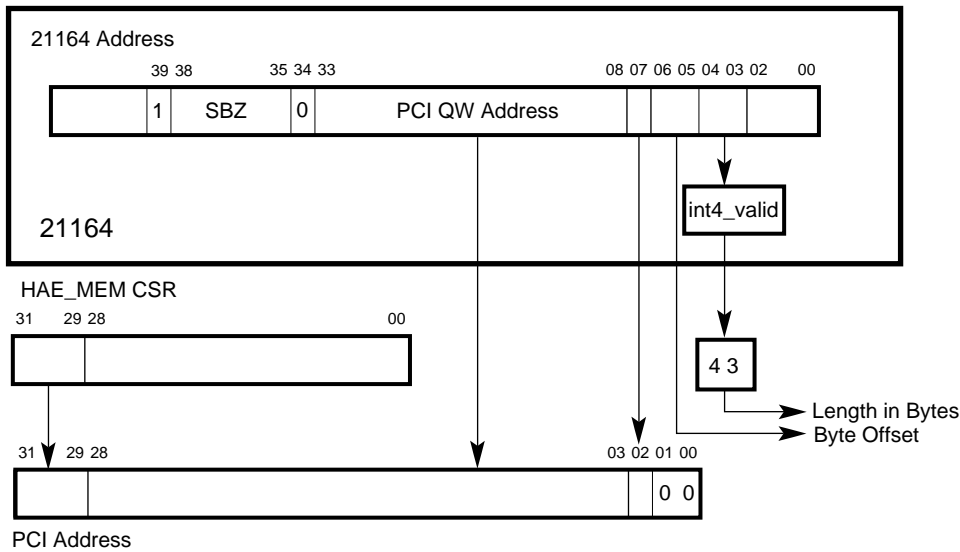
Table 6-4 HAE_MEM High Order Sparse Space Bits

Region ¹	PCI Address					
	<31>	<30>	<29>	<28>	<27>	<26>
1	HAE_MEM <31>	HAE_MEM <30>	HAE_MEM <29>	21164 <33>	21164 <32>	21164 <31>
2	HAE_MEM <15>	HAE_MEM <14>	HAE_MEM <13>	HAE_MEM <12>	HAE_MEM <11>	21164 <31>
3	HAE_MEM <7>	HAE_MEM <6>	HAE_MEM <5>	HAE_MEM <4>	HAE_MEM <3>	HAE_MEM <2>

¹Region 1 is 80.0000.0000 to 83.FFFF.FFFF.
Region 2 is 84.0000.0000 to 84.FFFF.FFFF.
Region 3 is 85.0000.0000 to 85.7FFF.FFFF.

HAE_MEM is located in the CIA and is described in Section 4.2.5. Figures 6-7 through 6-9 show mapping for the three regions.

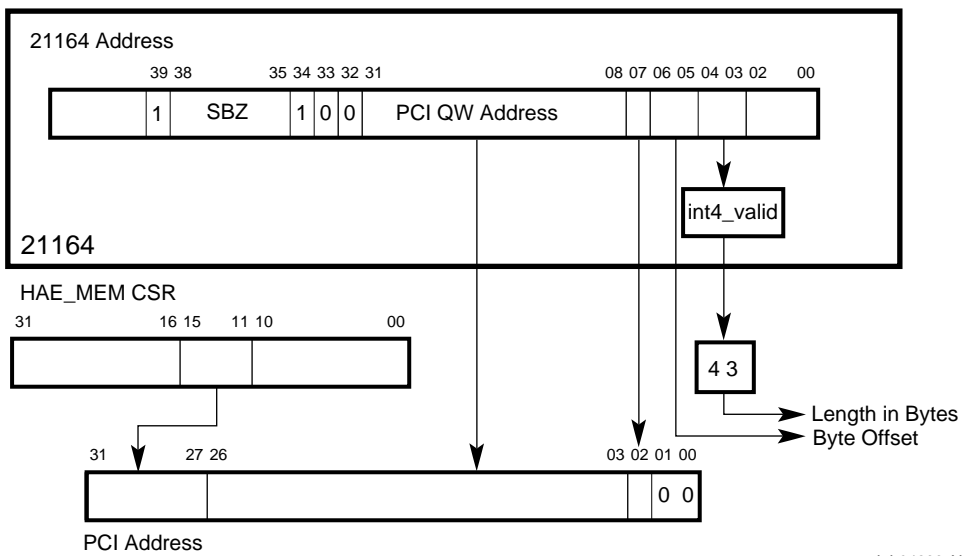
Figure 6-7 PCI Memory Sparse Space Address Generation (Region 1)



LJ04265A.A15

21172-CA Address Mapping
6.3 21164 Address Space

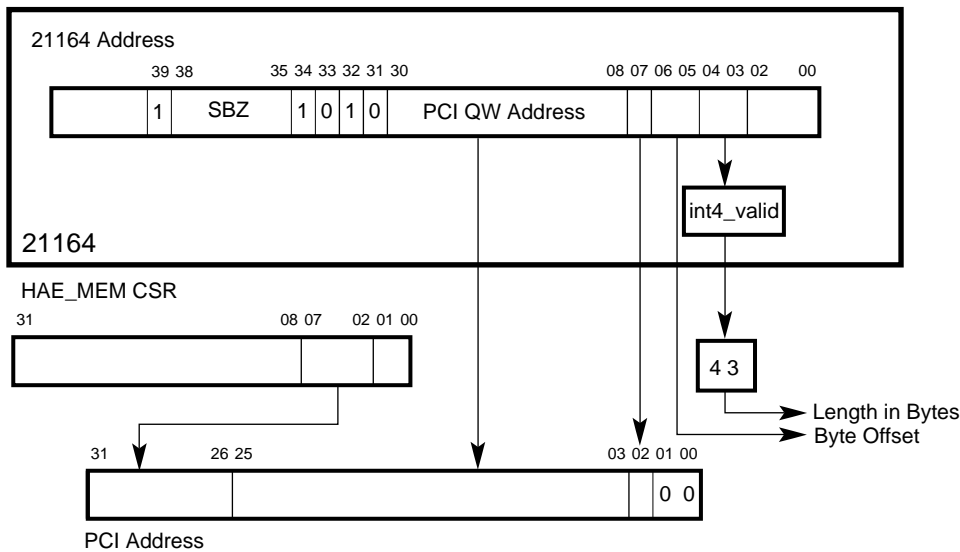
Figure 6-8 PCI Memory Sparse Space Address Generation (Region 2)



LJ-04266.AI

21172-CA Address Mapping
6.3 21164 Address Space

Figure 6-9 PCI Memory Sparse Space Address Generation (Region 3)



LJ-04267.AI

The 21164 provides six physical address bits <39:34> that can be used to backfill the “lost” sparse space bits. However, other 21164 platforms use these high-order bits in different ways, encoding multiple PCI ports for instance, and so for easier software portability, these bits are not used.

6.3.3 PCI Sparse I/O Space

PCI sparse I/O space has characteristics similar to the PCI sparse memory space. PCI sparse I/O space is located in the address range 85.8000.0000 to 85.FFFF.FFFF. This 2GB 21164 address segment maps to two 32MB regions of PCI I/O address space. A read or write transaction to this space causes a PCI I/O read or write transaction.


The high-order PCI address bits for region A are handled as follows:

- This region has **addr<34:30>** equal to 10110₂, and addresses the lower 32MB of PCI sparse I/O space; thus, **ad<31:25>** are set to zero by the hardware. See Figure 6-10.
- **ad<24:3>** are derived from **addr<29:8>**.
- This region is used for (E)ISA addressing (the EISA 64KB I/O space cannot be relocated).

21172-CA Address Mapping
6.3 21164 Address Space

- **ad<2:0>** are defined in Table 6-5.

Figure 6-10 PCI Sparse I/O Space Address Translation (Region A)




The high-order PCI address bits for region B are handled as follows:

- This region has **addr<34:30>** equal to 10111₂, and addresses 32MB of PCI sparse I/O space that can be relocated.
- This 32MB segment is relocated by assigning **ad<31:25>** equal to HAE_IO<31:25>, as shown in Figure 6-11.
HAE_IO is defined in Section 4.2.6. The power-on self-test (POST), running POST11 software, should initialize the contents of this register and the register should then remain unchanged.
- **ad<24:3>** are derived from **addr<29:8>**.
- **ad<2:0>** are defined in Table 6-5.

21172-CA Address Mapping
 6.3 21164 Address Space

Figure 6-11 PCI Sparse I/O Space Address Translation (Region B)



LJ04269A.AI5

21172–CA Address Mapping
6.3 21164 Address Space

Table 6–5 PCI I/O Sparse Space Read/Write Encodings

Size	Byte Offset	21164	PCI		Data-In Register
addr<4:3>	addr<6:5> ¹	Instruction	ad<2:0> ¹	Byte ² Enable	Byte Lanes [7:0]
Byte					
00	00	LDL, STL	addr<7>,00	1110	<0>
	01		addr<7>,01	1101	<1>
	10		addr<7>,10	1011	<2>
	11		addr<7>,11	0111	<3>
Word					
01	00	LDL, STL	addr<7>,00	1100	<1:0>
	01		addr<7>,01	1001	<2:1>
	10		addr<7>,10	0011	<3:2>
Tribyte					
10	00	LDL, STL	addr<7>,00	1000	<2:0>
	01		addr<7>,01	0001	<3:1>
Longword					
11	00	LDL, STL	addr<7>,00	0000	<3:0>
Quadword					
11	11	LDQ, STQ	000	0000	<7:0>

¹Missing entries (such as word size with **addr<6:5>** = 11₂) cause UNPREDICTABLE results.

²Byte enable set to zero indicates that byte lane carries meaningful data.

The (E)ISA devices have reserved the lower 64KB of PCI I/O space (85.8000.0000 to 85.801F.FFFF). Hence, all PCI devices should be relocated above this region. Table 6–6 shows the PCI and (E)ISA I/O map for a 21164 system. Four (E)ISA slots are hardwired, allocating 4KB per slot (as per EISA standard). The first slot is reserved for the EISA system board (X-bus addressing, interrupt controller, and so forth).

21172-CA Address Mapping
6.3 21164 Address Space

Table 6-6 CIA PCI and (E)ISA I/O Map

21164 Address	PCI Address	Range (KB)	Selection
85.8000.0000	0000.0000	0-4	EISA system board
85.8002.0000	0000.1000	4-8	EISA slot 1
85.8004.0000	0000.2000	8-12	EISA slot 2
85.8006.0000	0000.3000	12-16	EISA slot 3
85.8008.0000	0000.4000	16-20	EISA slot 4
85.800A.0000	0000.5000	20-24	Reserved
85.800C.0000	0000.6000	24-28	Reserved
85.800E.0000	0000.7000	28-32	Reserved
85.8010.0000	0000.8000	32-36	Reserved
85.8012.0000- 85.801F.FFFF	0000.9000- 0000.FFFF	36-64	Reserved
85.8020.0000- 85.BFFF.FFFF	0001.0000- 01FF.FFFF	64KB- 32MB	PCI I/O area—fixed
85.C000.0000- 85.FFFF.FFFF	0200.0000- 03FF.FFFF	32MB	PCI I/O area—can be relocated

Note

A quadword access to the PCI sparse I/O space will result in a 2-longword burst on the PCI. However, PCI devices might not support bursting in I/O space.

6.3.4 PCI Configuration Space

PCI configuration space is located in the address range: 87.0000.0000 to 87.1FFF.FFFF. Software designers are advised to clear CIA_CTRL [FILL_ERR_EN] when probing for PCI devices using configuration space read transactions. This will prevent the CIA from generating an ECC error if no device responds to the configuration cycle and UNPREDICTABLE data is read from the PCI bus.

A 21164 read or write access to this address space causes a configuration read or write cycle on the PCI. The two types of targets selected, depending upon the value of the configuration register (CFG), are listed here and are shown in Figure 6-12.

21172-CA Address Mapping


6.3 21164 Address Space

- Type 0—These are targets on the primary 64-bit 21164 system PCI bus. These are selected by making CFG<1:0> equal to 00₂.
- Type 1—These are targets on the secondary 32-bit 21164 system PCI bus (that is, behind a PCI-PCI bridge). These are selected by making CFG<1:0> equal to 01₂.

Note

CFG<1:0> equal to 10₂ and 11₂ are reserved by the PCI specification.

Figure 6-12 PCI Configuration Space Definition



LJ04270A.A15

Software must program CFG before running a configuration cycle. Sparse address decoding is used.

Note

The CIA uses CFG<1:0> instead of unused **addr<38:35>** to be compatible with the Digital Semiconductor 21071 core logic chipset. The Digital Semiconductor 21071 core logic chipset is used with Alpha 21064 series microprocessors.

- The high-order PCI address bits **ad<31:24>** are always zero.

21172–CA Address Mapping
6.3 21164 Address Space

- Address bits **addr<28:7>** correspond to PCI **ad<23:2>** and provide the configuration command information (which device to select).
- Address bits **addr<6:3>** are used to generate both the length of the PCI transaction in bytes and the byte enables, as shown in Table 6–7.
- Address bits **ad<1:0>** are obtained from CFG<1:0>.

Table 6–7 PCI Configuration Space Read/Write Encodings

Size	Byte Offset	21164	PCI		Data-In Register
addr<4:3>	addr<6:5> ¹	Instruction	ad<1:0> ¹	Byte ² Enable	Byte Lanes [7:0]
Byte					
00	00	LDL, STL	CFG<1:0>	1110	<0>
	01		CFG<1:0>	1101	<1>
	10		CFG<1:0>	1011	<2>
	11		CFG<1:0>	0111	<3>
Word					
01	00	LDL, STL	CFG<1:0>	1100	<1:0>
	01		CFG<1:0>	1001	<2:1>
	10		CFG<1:0>	0011	<3:2>
Tribyte					
10	00	LDL, STL	CFG<1:0>	1000	<2:0>
	01		CFG<1:0>	0001	<3:1>
Longword					
11	00	LDL, STL	CFG<1:0>	0000	<3:0>
Quadword					
11	11	LDQ, STQ	CFG<1:0>	0000	<7:0>

¹Missing entries (such as word size with **addr<6:5>** = 11₂) cause UNPREDICTABLE results.

²Byte enable set to zero indicates that byte lane carries meaningful data.

21172–CA Address Mapping
6.3 21164 Address Space

6.3.4.1 Device Select (IDSEL)

Peripherals are selected during a PCI configuration cycle if these three statements are true:

- Their IDSEL pin is asserted.
- The PCI bus command indicates a configuration read or write transaction.
- Address bits <1:0> are 00.

Address bits <7:2> select a longword register in the peripheral's 256-byte configuration address space. Transactions can use byte masks.

Peripherals that integrate multiple functional units (like SCSI and Ethernet) can provide configuration space for each function. Address bits **ad<10:8>** can be decoded by the peripheral to select one of eight functional units. Address bits **ad<31:11>** are available to generate the IDSELS. (Note that IDSELS behind a PCI-PCI bridge are determined from the device field encoding of a type 1 access.)

The IDSEL pin of each device corresponds to a unique PCI address bit from **ad<31:11>**. The binary value of **addr<20:16>** is used to select an address that is asserted on **ad<31:11>**, as listed in Table 6–8.

Table 6–8 Generating IDSEL Pin Signals

addr<20:16>	ad<31:11>–IDSEL
0000	0000 0000 0000 0000 0000 1
0001	0000 0000 0000 0000 0001 0
00010	0000 0000 0000 0000 0010 0
00011	0000 0000 0000 0000 0100 0
.....
10011	0100 0000 0000 0000 0000 0
10100	1000 0000 0000 0000 0000 0
10101	0000 0000 0000 0000 0000 0 ¹
.....
11111	0000 0000 0000 0000 0000 0 ¹

¹No device selected.

21172–CA Address Mapping

6.3 21164 Address Space

Note

If a quadword access is specified for the configuration cycle, then the least significant bit (LSB) of the register number field, **ad<2>**, must be zero. A quadword read/write transaction must access quadword-aligned registers.

If the PCI cycle is a configuration read or write cycle, but **ad<1:0>** equals 01_2 (like a type 1 transfer), then a device on a hierarchical bus is being selected using a PCI–PCI bridge. This cycle is accepted by the PCI–PCI bridge for propagation to its secondary PCI bus. During this cycle, **ad<23:16>** selects a unique bus number; **ad<15:8>** selects a device on that bus (typically decoded by the PCI–PCI bridge to generate the secondary PCI address pattern for IDSEL); and **ad<7:2>** selects a longword in the device's configuration space.

Each PCI–PCI bridge can be configured by PCI configuration cycles on its primary PCI interface. Configuration parameters in the PCI–PCI bridge will identify the bus number for its secondary PCI interface, and a range of bus numbers that can exist hierarchically behind it.

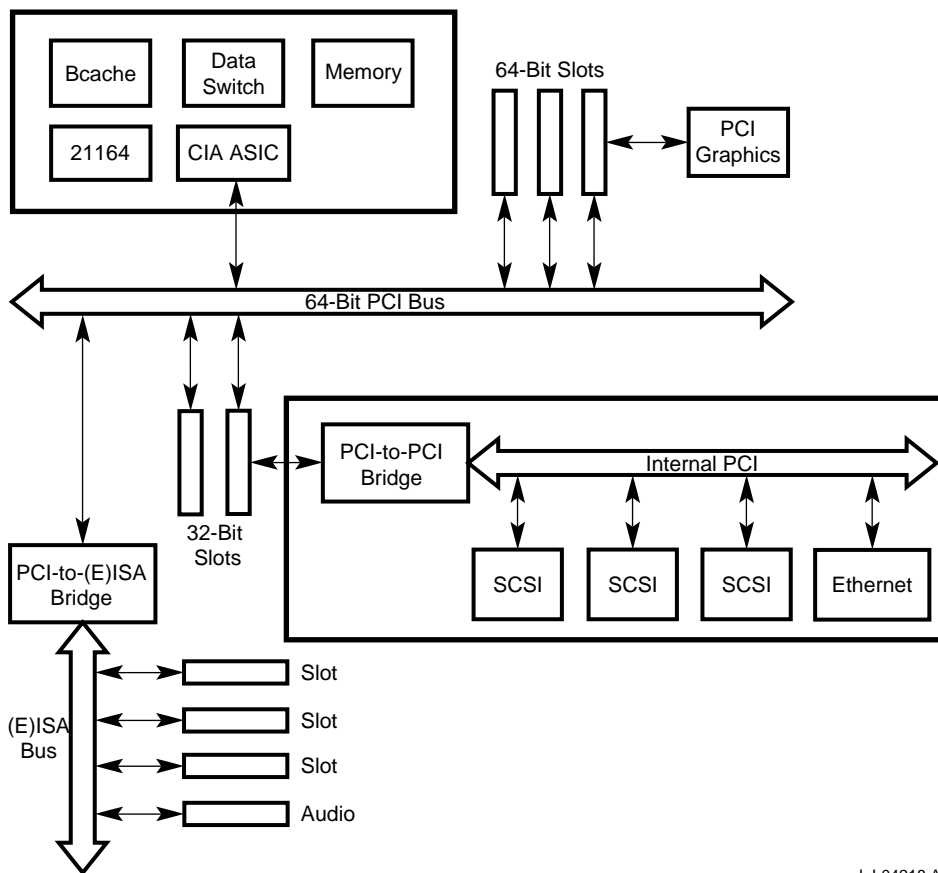
If the bus number of the configuration cycle matches the bus number of the PCI–PCI bridge secondary PCI interface, it will accept the configuration cycle, decode it, and generate a PCI configuration cycle with address bits $\langle 1:0 \rangle = 00$ on its secondary PCI interface.

If the bus number is within the range of bus numbers that may exist hierarchically behind its secondary PCI interface, the PCI–PCI bridge passes the PCI configuration cycle on unmodified (address bits $\langle 1:0 \rangle = 01$). It will be accepted by a bridge further downstream.

Figure 6–13 shows the PCI hierarchy for systems using the CIA. The IDSEL lines are significant in type 0 configuration cycles.

21172-CA Address Mapping
6.3 21164 Address Space

Figure 6-13 21164 System PCI Bus Hierarchy Example



LJ-04218.AI

6.3.4.2 PCI Special/Interrupt Acknowledge Cycles

PCI special/interrupt acknowledge cycle addresses are located in the range 87.2000.0000 to 87.3FFF.FFFF.

The special cycle command provides a simple message broadcasting mechanism on the PCI.

The special cycle contains no explicit destination address, but is broadcast to all devices. The CIA will drive all zeros as the special cycle address. Each receiving device must determine if the message contained in the data field is applicable to it.

21172–CA Address Mapping

6.3 21164 Address Space

A write transaction to address range 87.2000.0000 to 87.3FFF.FFFF causes a special cycle write transaction on the PCI. The 21164 write data will be passed unmodified to the PCI. Software must write the data in longword 0 of the hexword within the following fields:

- Bytes 0 and 1 contain the encoded message.
- Bytes 2 and 3 contain a message-dependent (optional) data field.

A read to address range 87.2000.0000 to 87.3FFF.FFFF will result in an interrupt acknowledge cycle on the PCI returning the vector data, which is provided by the PCI-to-EISA bridge, to the 21164.

6.3.4.3 Hardware-Specific and Miscellaneous Register Space

Hardware-specific and miscellaneous register space is located in the range 87.4000.0000 to 87.FFFF.FFFF. Table 6–9 lists the regions, with hardware registers, within this space.

Table 6–9 Hardware-Specific Register Address Space

addr<39:28>	Selected Region	addr<27:6>	addr<5:0>
1000.0111.0100 ¹	CIA control, diagnostic, error registers	LW address	000000
1000.0111.0101 ¹	CIA memory control registers	LW address	000000
1000.0111.0110 ¹	CIA PCI address translation (S/G, windows, and so on)	LW address	000000
1000.0111.0111 to 1000.0111.1111	Reserved	—	—

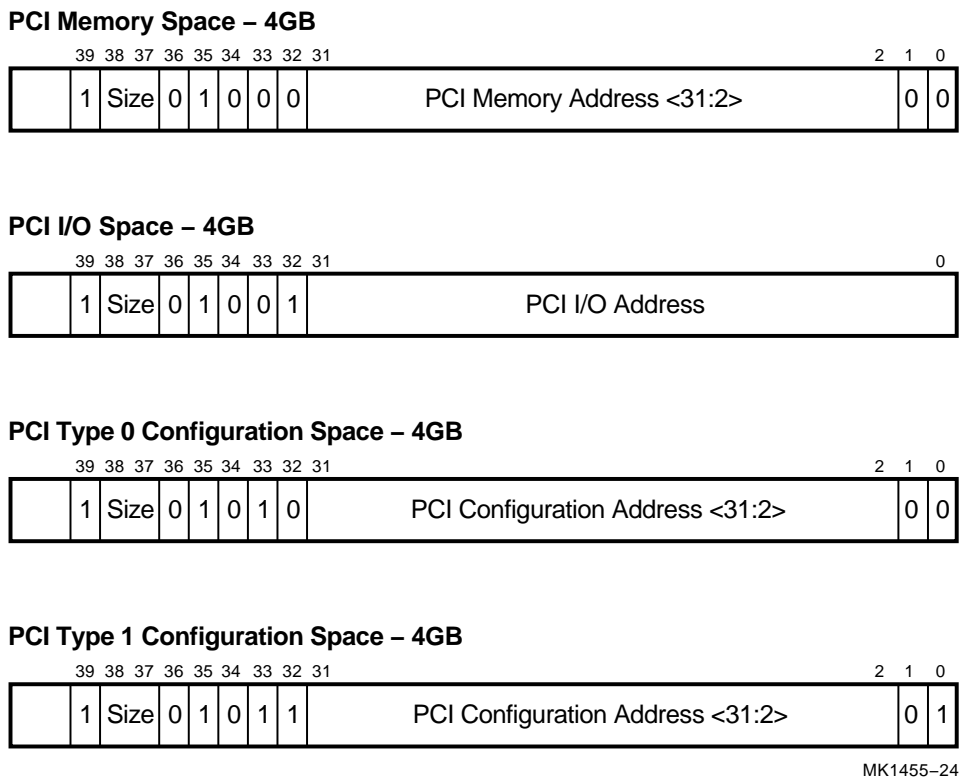
¹This address space is a hardware-specific variant of sparse space encoding. For the CSRs, **addr<27:6>** specify a longword address where **addr<5:0>** must be zero. All CIA registers are accessed with LW granularity. For more specific details on the CIA CSRs, refer to Chapter 4.

6.3.5 Byte/Word PCI Space

The 21164 supports byte/word instructions that allow software to access I/O space with byte granularity without using sparse space. Byte/word space is divided into four regions as shown in Figure 6–14.

21172-CA Address Mapping
6.3 21164 Address Space

Figure 6-14 Byte/Word PCI Space



Each of these regions operates identically as follows:

- Byte and word instructions issue a single byte/word read or write PCI transaction.
- Longword instructions issue a single longword read transaction, and up to eight longword write transactions.
- Quadword instructions issue up to four quadword read or write transactions.

21172–CA Address Mapping

6.3 21164 Address Space

The size field (address bits <38:37>) is added by the 21164 hardware as shown in the following list. The software value is zero.

Size	Data size
00	INT8
01	INT4
10	INT2
11	INT1

The following operations have single data transfers on the PCI:

- INT1 read/write operations
- INT2 read/write operations
- INT4 read operations

The following operations may have multiple data transfers on the PCI:

- INT4 write operations
- INT8 read/write operations

Byte/word support is enabled when 21164 CSR ICSR<17>=1 and when CIA CSR CIA_CNFG<IOA_BWEN>=1. Table 6–10 shows noncached 21164 addresses when byte/word support is enabled.

Table 6–10 21164 Byte/Word Addressing

Instruction	addr<38:37>	int4_valid<3:0>
LDQ	00	INT8 mask
LDL	01	addr<3:2> ,<1:0>undefined
LDWU	10	addr<3:1> ,<0>undefined
LDBU	11	addr<3:0>
STQ	00	INT4 mask
STL	01	INT4 mask
STW	10	addr<3:1> ,<0>undefined
STB	11	addr<3:0>

21172–CA Address Mapping

6.4 PCI to Physical Memory Addressing

6.4 PCI to Physical Memory Addressing

This section describes direct and scatter-gather mapping through the use of windows.

6.4.1 Address Mapping Windows

PCI addresses coming into the CIA (32-bit or 64-bit) are mapped to the 21164 cached memory space (8GB). The CIA provides five programmable address windows that control access of PCI peripherals to system memory. Each window location is defined by its base register (Wn_BASE) and its size is defined by its mask register (Wn_MASK). The five PCI address windows are also referred to as the PCI target windows.

Mapping from the PCI address to a physical memory address can be direct (physical mapping with an address offset) or scatter-gather (virtual mapping).

Windows [3:0]

Windows [3:0] have three registers associated with them. They are:

- PCI base register (Wn_BASE)
- PCI mask register (Wn_MASK)
- Translation base register (Tn_BASE)

In addition, there is another register that is associated with window 3 only. It is the PCI window DAC base register (W_DAC). It is used for PCI 64-bit addresses (DAC).

Wn_MASK provides a mask corresponding to bits <31:20> of an incoming PCI address. The size of each window can be programmed to be from 1MB to 4GB in powers of two, by masking bits of the incoming PCI address using Wn_MASK as shown in Table 6–11.

21172–CA Address Mapping
6.4 PCI to Physical Memory Addressing

Table 6–11 PCI Target Window MASK Register (W_n_MASK)

$W_MASK_{<31:20>}$	Size of Window	Value of n^1
0000 0000 0000	1MB	20
0000 0000 0001	2MB	21
0000 0000 0011	4MB	22
0000 0000 0111	8MB	23
0000 0000 1111	16MB	24
0000 0001 1111	32MB	25
0000 0011 1111	64MB	26
0000 0111 1111	128MB	27
0000 1111 1111	256MB	28
0001 1111 1111	512MB	29
0011 1111 1111	1GB	30
0111 1111 1111	2GB	31
1111 1111 1111	4GB	32
All others	UNPREDICTABLE	—

¹Only incoming $ad_{<31:n>}$ are compared with PCI base register $<31:n>$ as shown in Figure 6–16. If $n=32$, no comparison is performed. Windows are not allowed to overlap.

Based on the value of the window mask register, the unmasked bits of the incoming PCI address are compared with the corresponding bits of each window's W_n_BASE . If the address in one of the W_n_BASE registers and the incoming PCI address match, then the PCI address has hit that PCI target window. Otherwise, it has missed the window.

A window enable bit, W_n_BASE [W_EN], lets windows be independently enabled ($[W_EN]=1$) or disabled ($[W_EN]=0$). If a hit occurs in any of the windows that are enabled, then the CIA will respond to the PCI cycle by asserting **devsel₁**. The PCI target windows must be programmed so that their address ranges do not overlap. If they overlap, the compare results are UNDEFINED.

The window base address must be on a NATURALLY ALIGNED address boundary corresponding to the size of the window. For example, a 4MB window cannot start at address 1MB; it must start at address 4MB, or 8MB, or 12MB, and so on.

21172—CA Address Mapping

6.4 PCI to Physical Memory Addressing

6.4.1.1 PCI Device Address Space

A PCI device specifies the amount of memory space it requires by using base registers in its configuration space. The registers are implemented such that the address space consumed by the device is a power of two in size, and is NATURALLY ALIGNED on the size of the space consumed.

A PCI device need not use all of the address range it consumes, that is, the size of the PCI address window defined by the base address. Also, a PCI device need not respond to unused portions of the address space. The one exception to this is a PCI bridge that requires two additional registers (the base and limit address registers). These registers specify the address space that the PCI bridge will respond to in transactions.

A PCI bridge responds to all addresses in the range: $\text{base} \leq \text{address} < \text{limit}$.


These base and limit address registers are initialized by POST code at power-up. The CIA, as a PCI host-bridge device, does not have base and limit registers. Host bridges, because they are under system control, do not have to operate within the rules for other PCI devices. The CIA does respond to all the addresses within the four windows.

6.4.1.2 Address Mapping Example

Figure 6-15 shows how the DMA address ranges of a number of PCI devices are within the PCI window ranges. PCI devices are allowed to have multiple DMA address ranges, like device 2.

21172-CA Address Mapping
6.4 PCI to Physical Memory Addressing

Figure 6-15 PCI DMA Addressing Example



LJ-04272.AI

Figure 6-15 also shows that the CIA window can be larger than the corresponding device's DMA address range, as with device 0. Devices 1 and 2 have address ranges that are accepted by one CIA window. Wn_BASE [Wn_BASE_SG] for each window determines whether direct mapping or scatter-gather mapping is used to access physical memory.

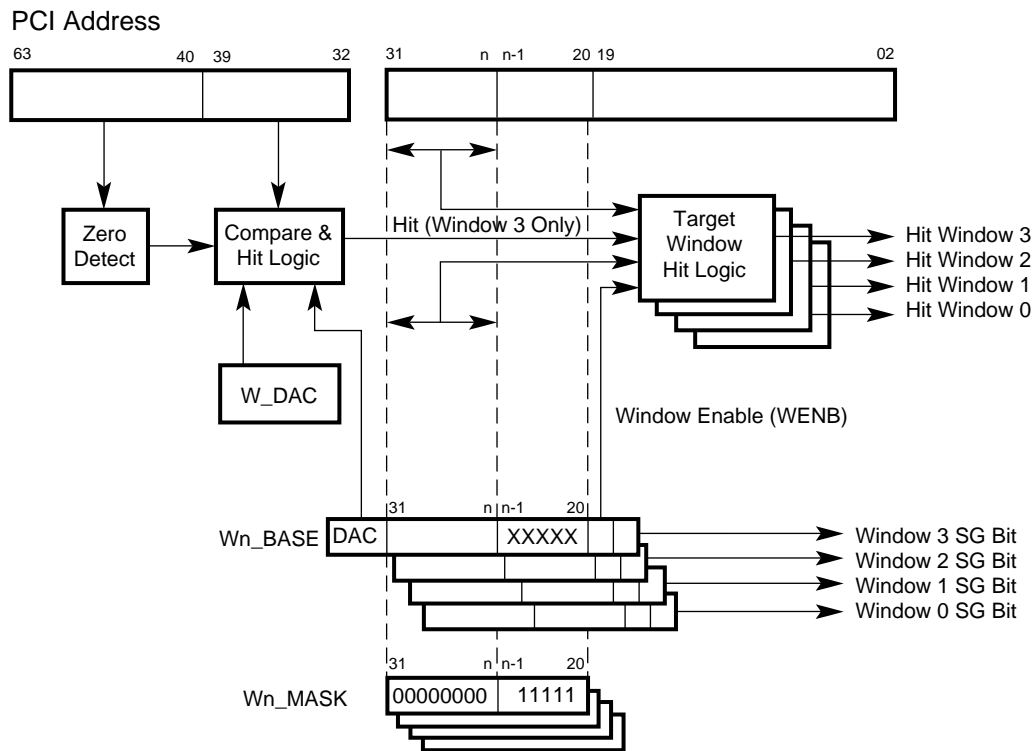
PCI single and dual address cycles have the following characteristics:

- Dual address cycle (DAC)—issued only if $\langle 63:32 \rangle$ are nonzero in 64-bit PCI address mode, and only if enabled by $W3_BASE[DAC_ENABLE]$.
- Single address cycle (SAC)—all 32-bit addresses. A PCI device must use SAC if bits $\langle 63:32 \rangle$ of a 64-bit address equal zero.

Figure 6-16 shows the PCI window comparison logic.

21172–CA Address Mapping
6.4 PCI to Physical Memory Addressing

Figure 6–16 PCI Target Window Compare



LJ04273A.A15

The comparison logic associated with **ad<63:32>** is used only for DAC with window 3. The other windows recognize only 32-bit PCI addresses (SAC).

For a hit to occur on a DAC address, address bits **<63:40>** must be zero and **ad<39:32>** must match the **W_DAC[DAC_BASE<7:0>]**. The low-order address bits **ad<31:20>** must also hit. This scheme allows a NATURALLY ALIGNED 1MB to 4GB PCI window to be placed anywhere in the first 1TB of a 64-bit PCI address space.


When an address match occurs with a PCI target window, the CIA translates the 32-bit PCI address **ad<31:0>** to a memory address **<33:0>**. The translated address is generated in one of two ways as determined by **Wn_BASE[Wn_BASE_SG]**.

21172–CA Address Mapping
6.4 PCI to Physical Memory Addressing

6.4.2 Direct Mapped Addressing

If W_n_BASE [$W_n_BASE_SG$] is clear, the DMA address is direct mapped. The translated address is generated by concatenating bits from the matching window's T_n_BASE with bits from the incoming PCI address (**ad**<**31:0**>). This process is shown in Figure 6–17 with n being the LSB from the T_n_BASE column of Table 6–12.

Figure 6–17 Direct Mapped Translation



LJ04274A.A15

The bits involved in the concatenation are defined by the window's W_n_MASK , as shown in Table 6–12. Because memory is located in the lower 8GB of the 21164 address space, the CIA implicitly ensures that **addr**<**39:33**> are always zero. Because T_n_BASE is simply concatenated to the PCI address, direct mapping is to a NATURALLY ALIGNED memory region. For example, a 4MB direct-mapped window will map to any 4MB region in main memory that falls on a 4MB boundary.

21172–CA Address Mapping

6.4 PCI to Physical Memory Addressing

Table 6–12 Direct-Mapped PCI Target Address Translation

W_MASK <31:20>	Size of Window	Translated Address Source	
		Tn_BASE ¹	PCI Address
0000 0000 0000	1MB	addr<32:20>	addr<19:2>
0000 0000 0001	2MB	addr<32:21>	addr<20:2>
0000 0000 0011	4MB	addr<32:22>	addr<21:2>
0000 0000 0111	8MB	addr<32:23>	addr<22:2>
0000 0000 1111	16MB	addr<32:24>	addr<23:2>
0000 0001 1111	32MB	addr<32:25>	addr<24:2>
0000 0011 1111	64MB	addr<32:26>	addr<25:2>
0000 0111 1111	128MB	addr<32:27>	addr<26:2>
0000 1111 1111	256MB	addr<32:28>	addr<27:2>
0001 1111 1111	512MB	addr<32:29>	addr<28:2>
0011 1111 1111	1GB	addr<32:30>	addr<29:2>
0111 1111 1111	2GB	addr<32:31>	addr<30:2>
1111 1111 1111	4GB	addr<32>	addr<31:2>

¹Unused bits of Tn_BASE must be cleared because the hardware performs an OR operation for the concatenation.

6.4.3 Scatter-Gather Addressing

When Wn_BASE[Wn_BASE_SG] is set, the translated address is generated using a table lookup. The table is referred to as a scatter-gather map.

The incoming PCI address is compared to the PCI window addresses for a hit. The Tn_BASE of the window that was hit is used to specify the starting address of the scatter-gather map table in memory.

Part of the incoming PCI address is used as an offset from this starting address to access the scatter-gather page table entry (PTE). This PTE, together with the remaining least-significant PCI address bits, forms the 21164 memory address.

Each scatter-gather map entry maps an 8KB page of PCI address space into an 8KB page of 21164 address space. This offers a number of advantages to software:

- Performance—ISA devices map to the lower 16MB of memory. The Windows NT operating system currently copies data from here to user space. The scatter-gather map avoids this copy operation.


21172–CA Address Mapping 6.4 PCI to Physical Memory Addressing

- Address management—User I/O buffers might not be physically contiguous or contained within a page. Without scatter-gather maps, software has to manage the scattered nature of the user buffer.

In peripheral components architecture, the term scatter-gather is not an address translation scheme but instead is used to signify a DMA transfer list. An element of this transfer list contains the DMA address and the number of data items to transfer. The DMA device fetches each item of the list until the list is empty. Many of the PCI devices, such as the PCI-to-EISA bridge, support this form of scatter-gather process.

Each scatter-gather PTE is a quadword and has a valid bit in PTE<0> as shown in Figure 6–18. Page address bit <13> is at PTE<1>.

Figure 6–18 Scatter-Gather PTE Format



LJ-04275.AI

Because the CIA implements only valid memory addresses up to 8GB, scatter-gather PTE bits <63:21> must be set to zero. PTE bits <20:1> are used to generate the physical page address. This address is appended to **ad<12:5>** of the incoming PCI address to generate the memory address.

The size of the scatter-gather map table is determined by the size of the PCI target window defined by Wn_MASK , as shown in Table 6–13. The number of entries is the window size divided by the page size (8KB). The size of the table is simply the number of entries multiplied by 8 bytes.

The scatter-gather map table address is obtained from Tn_BASE and the PCI address, as shown in Table 6–13.

21172–CA Address Mapping
6.4 PCI to Physical Memory Addressing

Table 6–13 Scatter-Gather Mapped PCI Target Address Translation

W_MASK <31:20>	Size of Window	S-G Map Table Size	Scatter-Gather Map Address <33:3>	
			T _n _Base ¹	PCI Address
0000 0000 0000	1MB	1KB	<32:10 >	<19:13>
0000 0000 0001	2MB	2KB	<32:11>	<20:13>
0000 0000 0011	4MB	4KB	<32:12>	<21:13>
0000 0000 0111	8MB	8KB	<32:13>	<22:13>
0000 0000 1111	16MB	16KB	<32:14>	<23:13>
0000 0001 1111	32MB	32KB	<32:15>	<24:13>
0000 0011 1111	64MB	64KB	<32:16>	<25:13>
0000 0111 1111	128MB	128KB	<32:17>	<26:13>
0000 1111 1111	256MB	256KB	<32:18>	<27:13>
0001 1111 1111	512MB	512KB	<32:19>	<28:13>
0011 1111 1111	1GB	1MB	<32:20>	<29:13>
0111 1111 1111	2GB	2MB	<32:21>	<30:13>
1111 1111 1111	4GB	4MB	<32:22>	<31:13>

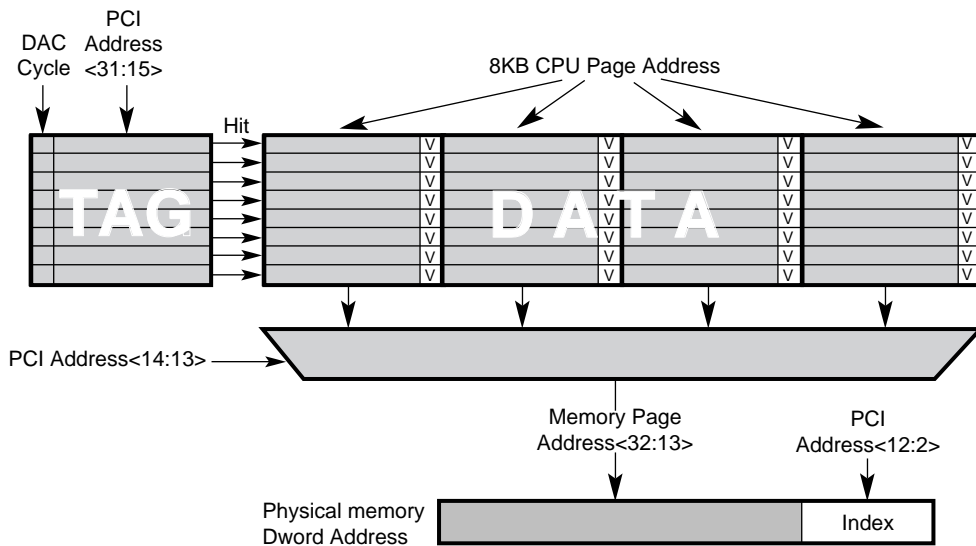
¹Unused bits of T_n_BASE must be zero for correct operation.

6.4.3.1 Scatter-Gather Translation Lookaside Buffer (TLB)

An 8-entry translation lookaside buffer (TLB) is provided in the CIA for scatter-gather PTEs. The TLB is a fully associative cache and holds the eight most recent scatter-gather map lookups. Four of these entries can be “locked,” preventing their displacement by the hardware TLB-miss handler. Each of the eight TLB entries holds a PCI address for the tag, and four consecutive 8KB page addresses as the TLB data, as shown in Figure 6–19.

21172-CA Address Mapping 6.4 PCI to Physical Memory Addressing

Figure 6-19 Scatter-Gather Associative TLB



LJ04276A.A15

Each time an incoming PCI address hits in a PCI target window that has scatter-gather enabled, **ad<31:15>** are compared with the 32KB PCI page address in the TLB tag. If a match is found, the required 21164 page address is one of the four items provided by the data in the matching TLB entry. Address bits **ad<14:13>** select the correct 8KB page address from the four addresses fetched.

With a TLB hit, the scatter-gather map table lookup in memory is avoided, resulting in enhanced performance. If no match is found in the TLB, the scatter-gather map lookup is performed and four PTE entries are fetched and written over an existing entry in the TLB. The TLB entry to be replaced is determined by a round-robin algorithm on the “unlocked” entries. Coherency of the TLB is maintained by software write transactions (invalidates) to the TBIA. See Section 4.7.1.

The TAG portion of the TLB entry contains a DAC flag to indicate that PCI tag address bits <31:15> correspond to a 64-bit DAC address. Only 1 bit is required instead of the high-order PCI address bits **ad<39:32>** because only window 3 is assigned to a DAC cycle, and the window hit logic has already performed a comparison of the high-order address bits against W_DAC.

21172–CA Address Mapping

6.4 PCI to Physical Memory Addressing

Figure 6–20 shows the entire translation process, from PCI address to physical address, on a window that implements scatter-gather. The MSB from the PCI address column of Table 6–13 equals $n-1$. Both paths are indicated: the path for a TLB hit is to the right, and the path for a TLB miss is to the left. The scatter-gather TLB is shown in a slightly simplified but functionally equivalent form.

6.4.3.1.1 Scatter-Gather TLB Hit Process The process for a scatter-gather TLB hit is as follows:

- The window compare logic determines if the PCI address has hit in one of the four windows and $Wn_BASE[Wn_BASE_SG]$ determines if the scatter-gather path should be taken. If window 3 has DAC mode enabled, and the PCI cycle is a DAC cycle, then a further comparison is made between the high-order PCI bits and W_DAC .
- Address bits **ad<31:13>** are sent to the TLB associative tag together with the DAC hit indication. If the address and DAC bits match in the TLB, the corresponding 8KB page, 21164 memory address is read out of the TLB. If this entry is valid, then a TLB hit has occurred and this page address is concatenated with **ad<12:2>** to form the physical memory address. If the data entry is invalid or if the TAG compare failed, a TLB miss occurs.

6.4.3.1.2 Scatter-Gather TLB Miss Process The process for a scatter-gather TLB miss is as follows:

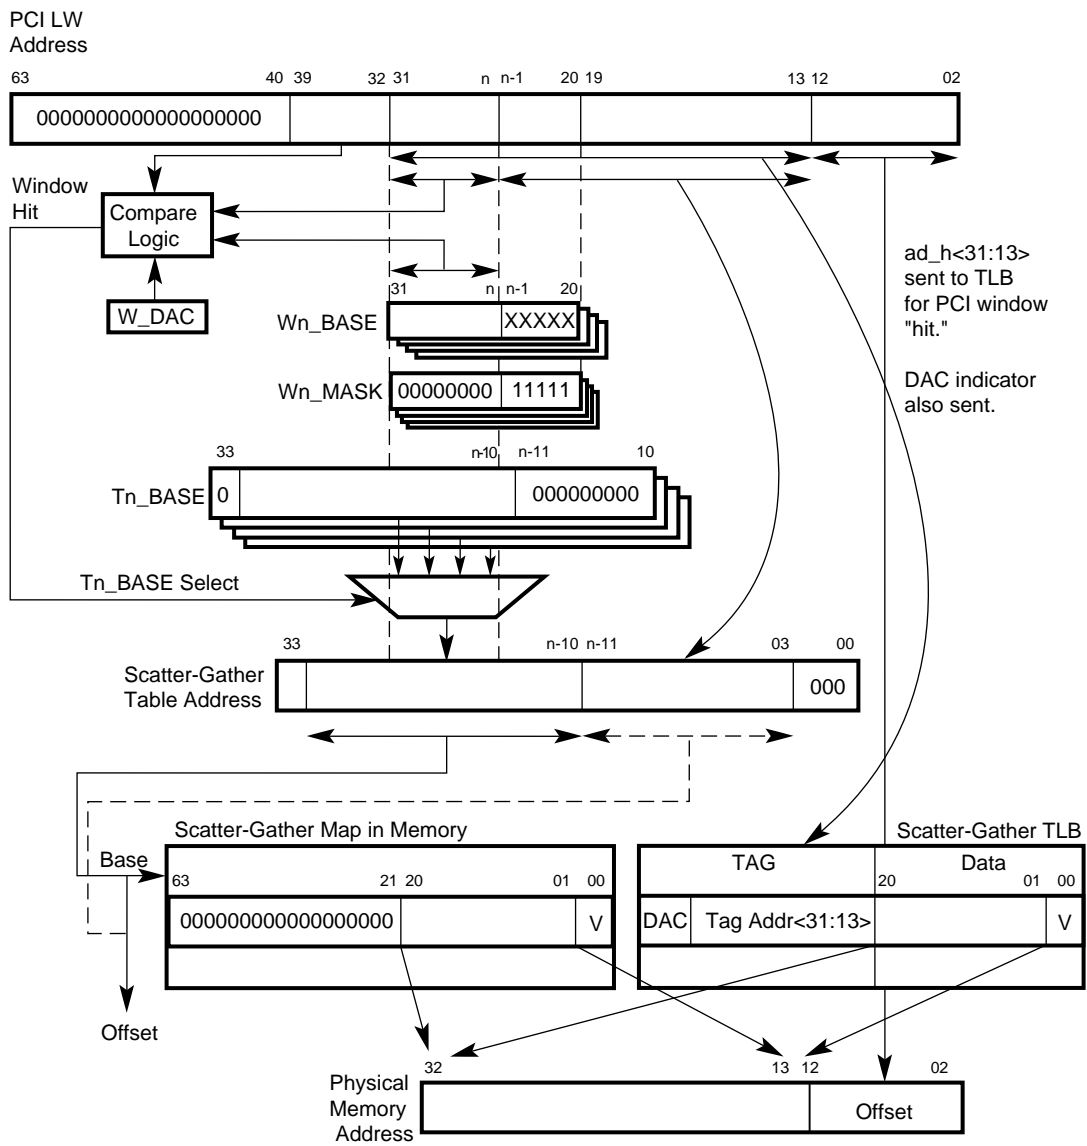
- The relevant bits of the PCI address (as determined by Wn_MASK) are concatenated with the relevant Tn_BASE bits to form the address used to access the scatter-gather PTE from a table located in main memory.
- Scatter-gather PTE<20:1> are used to generate the page address that is appended to the page offset to generate the physical memory address.

At this point, the TLB is also updated (round-robin algorithm) with the four PTE entries that correspond to the 32KB PCI page 21164 memory address that first missed the TLB. The tag portion of the TLB is loaded with this PCI page address, and the DAC bit is set if this PCI cycle is a DAC cycle.

- If the requested PTE is marked invalid (bit 0 clear), then a TLB invalid entry exception is taken.

21172-CA Address Mapping 6.4 PCI to Physical Memory Addressing

Figure 6-20 Scatter-Gather Map Translation



LJ-04277.A15


21172-CA Address Mapping

6.4 PCI to Physical Memory Addressing

6.4.4 Suggested Use of a PCI Window

Figure 6-21 shows a power-up PCI window assignment (as configured by firmware) and Table 6-14 lists the details. PCI window 0 was chosen for the 8MB to 16MB EISA region because this window incorporates the **mem_cs_1** logic. PCI window 3 was not used as it incorporates the DAC cycle logic. PCI window 1 was chosen arbitrarily for the 1GB, direct-mapped region, and PCI window 2 is not assigned.

Figure 6-21 Default PCI Window Allocation



LJ-04278.AI

21172–CA Address Mapping
6.4 PCI to Physical Memory Addressing

Table 6–14 PCI Window Power-Up Configuration

PCI Window	Assignment	Size	Comments
0	Scatter-Gather	8MB	Not used by firmware. mem_cs_1 disabled.
1	Direct Mapped	1GB	Mapped to 0GB to 1GB of main memory.
2	Disabled	—	—
3	Disabled	—	—

6.4.4.1 Peripheral Component Architecture Compatibility Addressing and Holes

The peripheral component architecture allows certain (E)ISA devices to respond to hardwired memory addresses. An example is a VGA graphics device that has its frame buffer located in memory address region A0000–BFFFF. Such devices “pepper” memory space with holes that are collectively known as peripheral component architecture compatibility holes.

The PCI-to-EISA bridge decodes PCI addresses and generates a signal, **mem_cs_1**, that takes into account the various compatibility holes.

6.4.4.2 Memory Chip Select Signal mem_cs_1

The PCI-to-EISA bridge set can be made using these two chips:

- Intel 82374EB EISA System Component (ESC)
- Intel 82375EB PCI-to-EISA Bridge (PCEB)

The PCI-to-EISA bridge provides address decode logic with attributes (such as read only, write only, VGA frame buffer, memory holes, and BIOS shadowing) to help manage the EISA memory map and peripheral component architecture compatibility holes.

This is known as main memory decoding in the PCEB chip, and results in the generation of the memory chip select (**mem_cs_1**) signal. One exception is the VGA memory hole region that never asserts **mem_cs_1**. If enabled, the CIA uses **mem_cs_1** with W0_BASE.

21172–CA Address Mapping

6.4 PCI to Physical Memory Addressing

In Figure 6–22, the two main holes are lightly shaded, while the **mem_cs_1** range is darkly shaded.

The **mem_cs_1** range of Figure 6–22 is subdivided into several portions (such as the BIOS area) that are individually enabled/disabled using CSRs as listed here:


- The MCSTOM (top of memory) register has 2MB granularity and can be programmed to select the regions from 1MB up to 512MB.
- The MCSTOH (top of hole) and MCSBOH (bottom of hole) registers define a memory hole region where **mem_cs_1** is not selected. The granularity of the hole is 64KB.
- The MAR1, 2, and 3 registers enable various BIOS regions.
- The MCSCON (control) register enables the **mem_cs_1** decode logic, and in addition selects a number of regions (0KB to 512KB).

Note

For more detail, please refer to the Intel 82375EB specification.

21172-CA Address Mapping
6.4 PCI to Physical Memory Addressing

Figure 6-22 Memory Chip Select Signal (mem_cs_l) Decode Area



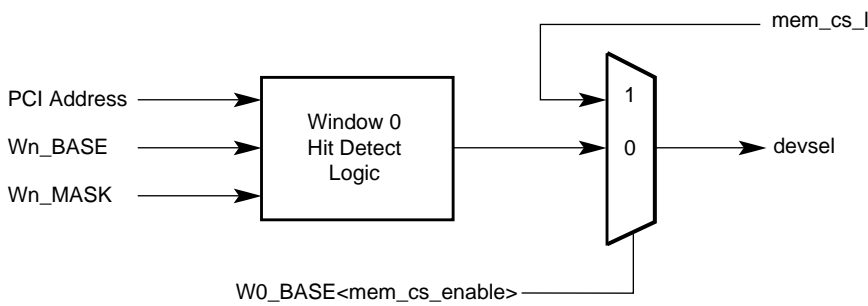
LJ-04279.AI5

21172–CA Address Mapping

6.4 PCI to Physical Memory Addressing

As shown in Figure 6–23, PCI window 0 in the CIA can be enabled to accept the **mem_cs_1** signal as the PCI memory decode signal. With this path enabled, the PCI window hit logic simply uses the **mem_cs_1** signal. For example, if **mem_cs_1** is asserted, then a PCI window 0 hit occurs and the **dev_sel_1** signal is asserted on the PCI.

Figure 6–23 Memory Chip Select Signal (mem_cs_1) Logic



LJ-04280.AI

Consequently, the window address area must be large enough to encompass the **mem_cs_1** region programmed into the PCI-to-EISA bridge. The remaining window attributes are still applicable and/or required.

- W0_BASE [Wn_BASE_SG] determines if scatter-gather or direct mapping is applicable.
- W0_MASK size information must match the **mem_cs_1** size for the scatter-gather and direct-mapping algorithms to correctly use the translated base register.
- The **mem_cs_1** enable bit, W0_BASE [MEMCS_ENABLE], takes precedence over W0_BASE [W_EN].

Part II

21172–BA (DSW) Information

Part II contains two chapters that provide information about the 21172–BA (DSW) chip. The following table lists the topics described in each chapter:

Chapter	Description
7	Pin signals
8	Architecture

21172–BA Pin Descriptions

This chapter provides a description of the 21172–BA (DSW) pin signals.

7.1 21172–BA Pin List

Table 7–1 lists the 21172–BA (DSW) pins grouped by function. The following abbreviations are used in the Type column of the table:

- I = Input
- O = Output
- B = Bidirectional
- A = Analog
- P = Power

Table 7–1 21172–BA Pin List

Signal Name	Quantity	Type	Signal Name	Quantity	Type
System Bus Signals (36 Total)					
cpu_dat<35:0>	36	B	—	—	—
IOD Bus Signals (34 Total)					
cmc<8:0>	9	I	ioc<6:0>	7	O
iod<17:0>	18	B	—	—	—
MEMDATA Bus Signals (73 Total)					
mem_dat<71:0>	72	B	mem_en	1	I

(continued on next page)

21172–BA Pin Descriptions
7.1 21172–BA Pin List

Table 7–1 (Cont.) 21172–BA Pin List

Signal Name	Quantity	Type	Signal Name	Quantity	Type
Miscellaneous Signals (6 Total)					
config<1:0>	2	I	reset_1	1	I
scan_out	1	O	test_mode<1:0>	2	I
Phase-Locked Loop Signals (3 Total)					
pll_agnd	1	A	pll_clk	1	I
pll_lp2	1	A	—	—	—
Miscellaneous Power Signals (4 Total)					
aux_vdd	1	P	aux_vss	1	P
pll_vdd	1	P	pll_vss	1	P
Reserved Signals (6 Total)					
iddt	1	I	procmon_out	1	O
Miscellaneous	4	I	—	—	—

Pin Totals

Total input pins:	21
Total output pins:	9
Total bidirectional pins:	126
Total analog pins:	2
	—
Total signal pins:	158
	—
Total signal pins:	158
Total Vdd pins:	22
Total Ground pins:	24
Total miscellaneous power pins:	4
	—
Total pins:	208

7.2 21172–BA Signal Descriptions

This section provides pin signal information, including a description of the signal; the clock edge on which the signal changes; and rules about signal usage during various system bus transactions.

For simplicity, the rising edge of **sys_clk_out1_h** will be indicated by the name **clk1R** and the falling edge of **sys_clk_out1_h** will be indicated by the name **clk1F**. See Section 7.2.5 for more information about clock use on the DSW.

Signal descriptions are grouped by function and correspond to the pin list (see Table 7–1).

Note

In all cases, the term **Vdd** refers to 3.3-V DSW chip power rather than other board power levels such as 5.0 V.

7.2.1 System Bus Signals

This section describes the system bus data signal lines.

cpu_dat<35:0>

The **cpu_dat<35:0>** signals carry data between the 21164, Bcache, and the DSW. Using four DSW chips in parallel results in 144 signal lines that carry 128 data bits and 16 ECC bits. There are no specific data and ECC bits. All bits are treated the same in the DSW. ECC generation and checking is performed in the 21164 and CIA.

By default, **cpu_dat<35:0>** are in the tristate condition as a DSW internal signal; **Drive_CPU_L**, is in the high state. This allows data to be read into the DSW from the 21164 or Bcache. The data source is controlled by the 21164. Data is clocked into the DSW every CLKx cycle.

When **Drive_CPU_L** is in the low state, the DSW drives data onto **cpu_dat<35:0>** during each CLKx cycle.

7.2.2 IOD Bus Signals

This section describes the IOD bus signals.

21172–BA Pin Descriptions

7.2 21172–BA Signal Descriptions

iod<17:0>

The **iod<17:0>** signals carry data and ECC between the CIA and DSW on the IOD bus. Four DSW slices together drive and receive 72 data/ECC signal lines on the IOD bus between the DSW and the CIA.

The DSW slices treat all bits in the same manner and are unaware of the concept of ECC. No specific bits are dedicated to ECC nor data. All ECC generation and detection is performed in the 21164 and CIA.

The DSW puts these signal lines in tristate mode by default by deasserting an internal DSW signal, Drive_IOD_L. The DSW asserts Drive_IOD_L and so drives data to the CIA every CLKx cycle.

ioc<6:0>

The **ioc<6:0>** signals allow the CIA to control data flow on the IOD bus. The data transfers are between memory and the PCI bus. CIA control is encoded as shown in Table 7–2.

Table 7–2 ioc<6:0> Control Functions

ioc <6:4>	ioc <3>	ioc <2>	ioc <1>	ioc <0>	Function
000	Buffer <i>n</i>	x	x	x	Clear valid bits in PCI buffer <i>n</i> .
001	x	x	x	x	Load the LW register (Reserved).
010	Buffer <i>n</i>	Adr m2	Adr m1	Adr m0	Write PCI buffer <i>n</i> , adr m.
011	Mask3	Mask2	Mask1	Mask0	Write IOR (QW mask).
100	Buffer <i>n</i>	Adr m2	Adr m1	Adr m0	Read DMA buffer <i>n</i> , adr m.
101	Buffer <i>n1</i>	Buffer <i>n0</i>	Adr m1	Adr m0	Read IOW buffer <i>n</i> , adr m.
110	Buffer <i>n1</i> Buffer <i>n</i>	Buffer <i>n0</i> Adr m2	Adr m1	Adr m0	Read buffer (IOW or DMA) low longword.
111	Mask3	Mask2	Mask1	Mask0	Start IOW (buffer mask).

cmc<8:0>

The **cmc<8:0>** signal lines control flow of data between the 21164 and the memory arrays. They also start and stop loading of the free-running buffers (victim, flush, IOW buffers) in the DSW. In Table 7–3, A5 and A4 are address bits. The Var field encodes the transaction being performed as shown in Table 7–4, Table 7–5, and Table 7–6.

21172–BA Pin Descriptions
7.2 21172–BA Signal Descriptions

Table 7–3 cmc<8:0> Control Functions

cmc<8:5>	cmc<4>	cmc<3>	cmc<2>	cmc<1>	cmc<0>	Function
0000	Var 4 ¹	Var 3	Var 2	Var 1	Var 0	Start/stop the buffer.
0010	x	A4 ²	Delay	Var 1	Var 0	Move fill data from memory to 21164.
0011	x	A4	x	Var 1	Var 0	Move IOR buffer data to 21164.
0100	A5 ²	A4	Delay	Var 1	Var 0	Move memory data to memory buffer 0.
0101	A5	A4	Delay	Var 1	Var 0	Move memory data to memory buffer 1.
0110	A5	A4	Delay	Stop IOW _{n1}	Stop IOW _{n0}	Move memory data to memory buffer 0, stop IOW buffer <i>n</i> .
0111	A5	A4	Delay	Stop IOW _{n1}	Stop IOW _{n0}	Move memory data to memory buffer 1, stop IOW buffer <i>n</i> .
1000	A5	A4	Var 2	Var 1	Var 0	Write victim buffer to memory.
1010	A5	A4	Var 2	Var 1	Var 0	Write DMA0 buffer to memory.
1011	A5	A4	Var 2	Var 1	Var 0	Write DMA1 buffer to memory.

¹The variable field encodes the transaction being performed. See Table 7–4, Table 2–5, and Table 2–6.

²A5 and A4 are address bits.

Table 7–4 Var Encodings for CMC Command 0

Var<4:3>	Function
00	None.
01	Start victim.
10	Start flush buffer 0 and clear PCI 0 valid bits.
11	Start flush buffer 1 and clear PCI 0 valid bits.

Table 7–5 Var Encodings for CMC Commands 0, 8, A, and B

Var<2:0>	Function
000	None.

(continued on next page)

21172–BA Pin Descriptions

7.2 21172–BA Signal Descriptions

Table 7–5 (Cont.) Var Encodings for CMC Commands 0, 8, A, and B

Var<2:0>	Function
001	Stop victim.
010	Stop flush buffer 0.
011	Stop flush buffer 1.
100	Stop IOW 0.
101	Stop IOW 1.
110	Stop IOW 2.
111	Stop IOW 3.

Table 7–6 Var Encodings for CMC Commands 2, 3, 4, 5, 6, and 7

Var<1:0>	Function
00	None.
01	Stop victim.
10	Stop flush buffer 0.
11	Stop flush buffer 1.

7.2.3 MEMDATA Bus Signals

This section describes the MEMDATA bus signals.

mem_dat<71:0>

The **mem_dat<71:0>** signals carry data and ECC. Four DSW slices in parallel drive and receive 288 data/ECC signal lines on the MEMDATA bus between the DSW and memory arrays.

The DSW slices treat all bits in the same manner and are unaware of the concept of ECC. No specific bits are dedicated to ECC nor data. All ECC generation and detection is performed in the 21164 and CIA. By default, **mem_dat<35:0>** are in the active condition as a DSW internal signal; Drive_MEM_L is in the low state. This allows data to be driven onto the MEMDATA bus every CLKx cycle.

When Drive_MEM_L is in the high state, the DSW is in the tristate mode and data is read into the DSW from the memory arrays on the edge of each CLK2x cycle.

mem_en

When the CIA drives **mem_en** high, the DSW asserts Drive_MEM_L, enabling the DSW to drive the MEMDATA bus. When the CIA does not assert **mem_en**, the DSW deasserts Drive_MEM_L, enabling the memory arrays to drive the MEMDATA bus.

21172–BA Pin Descriptions
7.2 21172–BA Signal Descriptions

7.2.4 Miscellaneous Signals

This section describes the miscellaneous DSW signals.

config<1:0>

Signals **config<1:0>** configure the data width of the IOD and MEMDATA buses as listed in Table 7–7.

Table 7–7 IOD Bus and MEMDATA Bus Data Width Selection

config<1:0>	IOD Bus Data Width	MEMDATA Bus Data Width
00 ¹	32 bits	128 bits
01 ¹	32 bits	256 bits
10	64 bits	128 bits
11	64 bits	256 bits

¹Not implemented.

reset_l

Signal **reset_l** is the system reset input signal to the DSW. When **reset_l** is asserted, the DSW resets its internal state.

scan_out

Signal **scan_out** is the end output of the parametric NAND tree.

test_mode<1:0>

Signals **test_mode<1:0>** define the four operating modes of the DSW as listed in Table 7–8.

Table 7–8 DSW Operating Modes

test_mode<1:0>	DSW Operating Mode
00	Tristate mode
01	Normal mode
10	Scan mode
11	PLL test mode

7.2.5 Phase-Locked Loop Signals

This section describes the DSW phase-locked loop signals. See Figure 2–1 for required external filter components.

21172–BA Pin Descriptions

7.2 21172–BA Signal Descriptions

pll_agnd

Signal **pll_agnd** is an analog ground output from the PLL used for the low-pass filter connected to **pll_lp2**. It *must not* be connected to system ground (**Vss**).

pll_clk

The **pll_clk** signal is the system clock (or reference signal) input to the phase-locked loop (PLL) circuit. It supplies a clock with the same frequency as that supplied to the PCI.

pll_lp2

Signal **pll_lp2** is the PLL VCO loop filter input.

7.2.6 Miscellaneous Power Signals

This section describes the DSW miscellaneous power signals. See Figure 2–1 for required external filter components.

aux_vdd

Signal **aux_vdd** is an input to the PLL. It is connected to the system board power.

aux_vss

Signal **aux_vss** is an input to the PLL. It is connected to the system board ground.

pll_vdd

Signal **pll_vdd** is an input to the PLL. It is connected to the system board power.

pll_vss

Signal **pll_vss** is an input to the PLL. It is connected to the system board ground.

7.2.7 Reserved Signals

This section describes the DSW reserved signals.

iddt

Signal **iddt** is a test pin reserved for use by Digital. It should be tied to ground.

procmon_out

Signal **procmon_out** is a test pin reserved for use by Digital.

21172–BA Pin Descriptions
7.2 21172–BA Signal Descriptions

Test pins 1 through 4

The four test pins (Pins 204 through 207) are reserved for manufacturing test purposes. These pins must be pulled up to **Vss** on the module.

21172–BA Pin Descriptions
7.2 21172–BA Signal Descriptions

7.2.8 21172–BA Pin Name List (Alphabetic)

Table 7–9 lists the 21172–BA (DSW) pins in alphabetical order. The following abbreviations are used in the Type column of the table:

- I = Input
- O = Output
- B = Bidirectional
- A = Analog
- P = Power

Table 7–9 21172–BA Pin Assignment List (Alphabetic)

Pin Name	Pin Number	Type	5-V Tolerant	Pin Name	Pin Number	Type	5-V Tolerant
aux_vdd	105	P	No	aux_vss	106	P	No
cmc<0>	102	I	No	cmc<1>	101	I	No
cmc<2>	100	I	No	cmc<3>	99	I	No
cmc<4>	98	I	No	cmc<5>	116	I	No
cmc<6>	117	I	No	cmc<7>	118	I	No
cmc<8>	119	I	No	config<0>	49	I	Yes
config<1>	50	I	Yes	cpu_dat<0>	3	B	No
cpu_dat<1>	4	B	No	cpu_dat<2>	12	B	No
cpu_dat<3>	13	B	No	cpu_dat<4>	21	B	No
cpu_dat<5>	22	B	No	cpu_dat<6>	30	B	No
cpu_dat<7>	31	B	No	cpu_dat<8>	55	B	No
cpu_dat<9>	56	B	No	cpu_dat<10>	64	B	No
cpu_dat<11>	65	B	No	cpu_dat<12>	73	B	No
cpu_dat<13>	74	B	No	cpu_dat<14>	82	B	No
cpu_dat<15>	83	B	No	cpu_dat<16>	91	B	No
cpu_dat<17>	92	B	No	cpu_dat<18>	127	B	No
cpu_dat<19>	126	B	No	cpu_dat<20>	136	B	No
cpu_dat<21>	135	B	No	cpu_dat<22>	145	B	No
cpu_dat<23>	144	B	No	cpu_dat<24>	154	B	No
cpu_dat<25>	153	B	No	cpu_dat<26>	159	B	No
cpu_dat<27>	160	B	No	cpu_dat<28>	168	B	No

(continued on next page)

21172–BA Pin Descriptions
7.2 21172–BA Signal Descriptions

Table 7–9 (Cont.) 21172–BA Pin Assignment List (Alphabetic)

Pin Name	Pin Number	Type	5-V Tolerant	Pin Name	Pin Number	Type	5-V Tolerant
cpu_dat<29>	169	B	No	cpu_dat<30>	177	B	No
cpu_dat<31>	178	B	No	cpu_dat<32>	186	B	No
cpu_dat<33>	187	B	No	cpu_dat<34>	195	B	No
cpu_dat<35>	196	B	No	iddt	113	I	No
ioc<0>	45	I	No	ioc<1>	43	I	No
ioc<2>	42	I	No	ioc<3>	41	I	No
ioc<4>	40	I	No	ioc<5>	39	I	No
ioc<6>	38	I	No	iod<0>	2	B	No
iod<1>	11	B	Yes	iod<2>	20	B	Yes
iod<3>	29	B	Yes	iod<4>	54	B	Yes
iod<5>	63	B	Yes	iod<6>	72	B	Yes
iod<7>	81	B	Yes	iod<8>	90	B	Yes
iod<9>	128	B	Yes	iod<10>	137	B	Yes
iod<11>	146	B	Yes	iod<12>	155	B	Yes
iod<13>	158	B	Yes	iod<14>	167	B	Yes
iod<15>	176	B	Yes	iod<16>	184	B	Yes
iod<17>	194	B	Yes	mem_dat<0>	6	B	Yes
mem_dat<1>	7	B	Yes	mem_dat<2>	8	B	Yes
mem_dat<3>	9	B	Yes	mem_dat<4>	15	B	Yes
mem_dat<5>	16	B	Yes	mem_dat<6>	17	B	Yes
mem_dat<7>	18	B	Yes	mem_dat<8>	24	B	Yes
mem_dat<9>	25	B	Yes	mem_dat<10>	26	B	Yes
mem_dat<11>	27	B	Yes	mem_dat<12>	33	B	Yes
mem_dat<13>	34	B	Yes	mem_dat<14>	35	B	Yes
mem_dat<15>	36	B	Yes	mem_dat<16>	58	B	Yes
mem_dat<17>	59	B	Yes	mem_dat<18>	60	B	Yes
mem_dat<19>	61	B	Yes	mem_dat<20>	67	B	Yes
mem_dat<21>	68	B	Yes	mem_dat<22>	69	B	Yes
mem_dat<23>	70	B	Yes	mem_dat<24>	76	B	Yes
mem_dat<25>	77	B	Yes	mem_dat<26>	78	B	Yes
mem_dat<27>	79	B	Yes	mem_dat<28>	85	B	Yes
mem_dat<29>	86	B	Yes	mem_dat<30>	87	B	Yes

(continued on next page)

21172–BA Pin Descriptions
7.2 21172–BA Signal Descriptions

Table 7–9 (Cont.) 21172–BA Pin Assignment List (Alphabetic)

Pin Name	Pin Number	Type	5-V Tolerant	Pin Name	Pin Number	Type	5-V Tolerant
mem_dat<31>	88	B	Yes	mem_dat<32>	94	B	Yes
mem_dat<33>	95	B	Yes	mem_dat<34>	96	B	Yes
mem_dat<35>	97	B	Yes	mem_dat<36>	124	B	Yes
mem_dat<37>	123	B	Yes	mem_dat<38>	122	B	Yes
mem_dat<39>	121	B	Yes	mem_dat<40>	133	B	Yes
mem_dat<41>	132	B	Yes	mem_dat<42>	131	B	Yes
mem_dat<43>	130	B	Yes	mem_dat<44>	142	B	Yes
mem_dat<45>	141	B	Yes	mem_dat<46>	140	B	Yes
mem_dat<47>	139	B	Yes	mem_dat<48>	151	B	Yes
mem_dat<49>	150	B	Yes	mem_dat<50>	149	B	Yes
mem_dat<51>	148	B	Yes	mem_dat<52>	162	B	Yes
mem_dat<53>	163	B	Yes	mem_dat<54>	164	B	Yes
mem_dat<55>	165	B	Yes	mem_dat<56>	171	B	Yes
mem_dat<57>	172	B	Yes	mem_dat<58>	173	B	Yes
mem_dat<59>	174	B	Yes	mem_dat<60>	180	B	Yes
mem_dat<61>	181	B	Yes	mem_dat<62>	182	B	Yes
mem_dat<63>	183	B	Yes	mem_dat<64>	189	B	Yes
mem_dat<65>	190	B	Yes	mem_dat<66>	191	B	Yes
mem_dat<67>	192	B	Yes	mem_dat<68>	198	B	Yes
mem_dat<69>	199	B	Yes	mem_dat<70>	200	B	Yes
mem_dat<71>	201	B	Yes	mem_en	46	I	No
pll_agnd	108	A	No	pll_clk	114	I	No
pll_lp2	107	A	No	pll_vdd	110	P	No
pll_vss	109	P	No	procmon_out	51	O	No
reset_l	115	I	Yes	scan_out	203	O	No
test_mode<0>	47	I	Yes	test_mode<1>	48	I	Yes
Test pin 1	204	I	Yes	Test pin 2	205	I	Yes
Test pin 3	206	I	Yes	Test pin 4	207	I	Yes

(continued on next page)

21172–BA Pin Descriptions 7.2 21172–BA Signal Descriptions

Table 7–9 (Cont.) 21172–BA Pin Assignment List (Alphabetic)

Pin Name	Pin Number	Type	5-V Tolerant
Vdd	1, 10, 19, 28, 37, 52, 57, 71, 80, 84, 104, 112, 125, 134, 143, 156, 157, 166, 175, 179, 193, 202	P	—
Ground	5, 14, 23, 32, 44, 53, 62, 66, 75, 89, 93, 103, 111, 120, 129, 138, 147, 152, 161, 170, 185, 188, 197, 208	P	—

7.2.9 21172–BA Pin Assignment List (Numeric)

Table 7–10 lists the 21172–BA (DSW) pins in numeric order. The following abbreviations are used in the Type column of the table:

- I = Input
- O = Output
- B = Bidirectional
- A = Analog
- P = Power

Table 7–10 21172–BA Pin Assignment (Numeric)

Pin Number	Pin Name	Type	5-V Tolerant	Pin Number	Pin Name	Type	5-V Tolerant
1	Vdd	P	—	2	iod<0>	B	Yes
3	cpu_dat<0>	B	No	4	cpu_dat<1>	B	No
5	Ground	P	—	6	mem_dat<0>	B	Yes
7	mem_dat<1>	B	Yes	8	mem_dat<2>	B	Yes
9	mem_dat<3>	B	Yes	10	Vdd	P	—
11	iod<1>	B	Yes	12	cpu_dat<2>	B	No
13	cpu_dat<3>	B	No	14	Ground	P	—
15	mem_dat<4>	B	Yes	16	mem_dat<5>	B	Yes
17	mem_dat<6>	B	Yes	18	mem_dat<7>	B	Yes
19	Vdd	P	—	20	iod<2>	B	Yes
21	cpu_dat<4>	B	No	22	cpu_dat<5>	B	No

(continued on next page)

21172–BA Pin Descriptions
7.2 21172–BA Signal Descriptions

Table 7–10 (Cont.) 21172–BA Pin Assignment (Numeric)

Pin Number	Pin Name	Type	5-V Tolerant	Pin Number	Pin Name	Type	5-V Tolerant
23	Ground	P	—	24	mem_dat<8>	B	Yes
25	mem_dat<9>	B	Yes	26	mem_dat<10>	B	Yes
27	mem_dat<11>	B	Yes	28	Vdd	P	—
29	iod<3>	B	Yes	30	cpu_dat<6>	B	No
31	cpu_dat<7>	B	No	32	Ground	P	—
33	mem_dat<12>	B	Yes	34	mem_dat<13>	B	Yes
35	mem_dat<14>	B	Yes	36	mem_dat<15>	B	Yes
37	Vdd	P	—	38	ioc<6>	I	No
39	ioc<5>	I	No	40	ioc<4>	I	No
41	ioc<3>	I	No	42	ioc<2>	I	No
43	ioc<1>	I	No	44	Ground	P	—
45	ioc<0>	I	No	46	mem_en	I	No
47	test_mode<0>	I	Yes	48	test_mode<1>	I	Yes
49	config<0>	I	Yes	50	config<1>	I	Yes
51	procmon_out	O	—	52	Vdd	P	—
53	Ground	P	—	54	iod<4>	B	Yes
55	cpu_dat<8>	B	No	56	cpu_dat<9>	B	No
57	Vdd	P	—	58	mem_dat<16>	B	Yes
59	mem_dat<17>	B	Yes	60	mem_dat<18>	B	Yes
61	mem_dat<19>	B	Yes	62	Ground	P	—
63	iod<5>	B	Yes	64	cpu_dat<10>	B	No
65	cpu_dat<11>	B	No	66	Ground	P	—
67	mem_dat<20>	B	Yes	68	mem_dat<21>	B	Yes
69	mem_dat<22>	B	Yes	70	mem_dat<23>	B	Yes
71	Vdd	P	—	72	iod<6>	B	Yes
73	cpu_dat<12>	B	No	74	cpu_dat<13>	B	No
75	Ground	P	—	76	mem_dat<24>	B	Yes
77	mem_dat<25>	B	Yes	78	mem_dat<26>	B	Yes
79	mem_dat<27>	B	Yes	80	Vdd	P	—
81	iod<7>	B	Yes	82	cpu_dat<14>	B	No
83	cpu_dat<15>	B	No	84	Vdd	P	—
85	mem_dat<28>	B	Yes	86	mem_dat<29>	B	Yes

(continued on next page)

21172–BA Pin Descriptions
7.2 21172–BA Signal Descriptions

Table 7–10 (Cont.) 21172–BA Pin Assignment (Numeric)

Pin Number	Pin Name	Type	5-V Tolerant	Pin Number	Pin Name	Type	5-V Tolerant
87	mem_dat<30>	B	Yes	88	mem_dat<31>	B	Yes
89	Ground	P	—	90	iod<8>	B	Yes
91	cpu_dat<16>	B	No	92	cpu_dat<17>	B	No
93	Ground	P	—	94	mem_dat<32>	B	Yes
95	mem_dat<33>	B	Yes	96	mem_dat<34>	B	Yes
97	mem_dat<35>	B	Yes	98	cmc<4>	I	No
99	cmc<3>	I	No	100	cmc<2>	I	No
101	cmc<1>	I	No	102	cmc<0>	I	No
103	Ground	P	—	104	Vdd	P	—
105	aux_vdd	P	—	106	aux_vss	P	—
107	pll_lp2	A	—	108	pll_agnd	A	—
109	pll_vss	P	—	110	pll_vdd	P	—
111	Ground	P	—	112	Vdd	P	—
113	iddt	I	No	114	pll_clk	I	No
115	reset_l	I	Yes	116	cmc<5>	I	No
117	cmc<6>	I	No	118	cmc<7>	I	No
119	cmc<8>	I	No	120	Ground	P	—
121	mem_dat<39>	B	Yes	122	mem_dat<38>	B	Yes
123	mem_dat<37>	B	Yes	124	mem_dat<36>	B	Yes
125	Vdd	P	—	126	cpu_dat<19>	B	No
127	cpu_dat<18>	B	No	128	iod<9>	B	Yes
129	Ground	P	—	130	mem_dat<43>	B	Yes
131	mem_dat<42>	B	Yes	132	mem_dat<41>	B	Yes
133	mem_dat<40>	B	Yes	134	Vdd	P	—
135	cpu_dat<21>	B	No	136	cpu_dat<20>	B	No
137	iod<10>	B	Yes	138	Ground	P	—
139	mem_dat<47>	B	Yes	140	mem_dat<46>	B	Yes
141	mem_dat<45>	B	Yes	142	mem_dat<44>	B	Yes
143	Vdd	P	—	144	cpu_dat<23>	B	No
145	cpu_dat<22>	B	No	146	iod<11>	B	Yes
147	Ground	P	—	148	mem_dat<51>	B	Yes
149	mem_dat<50>	B	Yes	150	mem_dat<49>	B	Yes

(continued on next page)

21172–BA Pin Descriptions
7.2 21172–BA Signal Descriptions

Table 7–10 (Cont.) 21172–BA Pin Assignment (Numeric)


Pin Number	Pin Name	Type	5-V Tolerant	Pin Number	Pin Name	Type	5-V Tolerant
151	mem_dat<48>	B	Yes	152	Ground	P	—
153	cpu_dat<25>	B	No	154	cpu_dat<24>	B	No
155	iod<12>	B	Yes	156	Vdd	P	—
157	Vdd	P	—	158	iod<13>	B	Yes
159	cpu_dat<26>	B	No	160	cpu_dat<27>	B	No
161	Ground	P	—	162	mem_dat<52>	B	Yes
163	mem_dat<53>	B	Yes	164	mem_dat<54>	B	Yes
165	mem_dat<55>	B	Yes	166	Vdd	P	—
167	iod<14>	B	Yes	168	cpu_dat<28>	B	No
169	cpu_dat<29>	B	No	170	Ground	P	—
171	mem_dat<56>	B	Yes	172	mem_dat<57>	B	Yes
173	mem_dat<58>	B	Yes	174	mem_dat<59>	B	Yes
175	Vdd	P	—	176	iod<15>	B	Yes
177	cpu_dat<30>	B	No	178	cpu_dat<31>	B	No
179	Vdd	P	—	180	mem_dat<60>	B	Yes
181	mem_dat<61>	B	Yes	182	mem_dat<62>	B	Yes
183	mem_dat<63>	B	Yes	184	iod<16>	B	Yes
185	Ground	P	—	186	cpu_dat<32>	B	No
187	cpu_dat<33>	B	No	188	Ground	P	—
189	mem_dat<64>	B	Yes	190	mem_dat<65>	B	Yes
191	mem_dat<66>	B	Yes	192	mem_dat<67>	B	Yes
193	Vdd	P	—	194	iod<17>	B	Yes
195	cpu_dat<34>	B	No	196	cpu_dat<35>	B	No
197	Ground	P	—	198	mem_dat<68>	B	Yes
199	mem_dat<69>	B	Yes	200	mem_dat<70>	B	Yes
201	mem_dat<71>	B	Yes	202	Vdd	P	—
203	scan_out	O	—	204	Test pin 1	I	Yes
205	Test pin 2	I	Yes	206	Test pin 3	I	Yes
207	Test pin 4	I	Yes	208	Ground	P	—

21172-BA Pin Descriptions
7.3 21172-BA Mechanical Specifications

7.3 21172-BA Mechanical Specifications

The 21172-BA is contained in an industry-standard 208-pin plastic quad flat pack (PQFP) package, shown in Figure 7-1.

Figure 7-1 21172-BA 208-Pin PQFP Package Dimensions



21172–BA Pin Descriptions

7.3 21172–BA Mechanical Specifications

Table 7–11 lists the 208-pin package dimensions in millimeters.

Table 7–11 208-Pin PQFP Package Dimensions

Symbol	Dimension	Value (mm)
LL	Lead length	1.30 reference ¹
e	Lead pitch	0.50 BSC ²
L	Foot length	0.50 minimum to 0.75 maximum
A	Package overall height	4.20 ¹
A1	Package standoff height	0.25
A2	Package thickness	3.17 minimum to 4.25 maximum
b	Lead width	0.17 minimum to 0.27 maximum
c	Lead thickness	0.09 minimum to 0.20 maximum
ccc	Coplanarity	0.08
ddd	Lead skew	0.08
D	Package overall width	30.60 BSC ²
D1	Package width	28.00 BSC ²
E	Package overall length	30.60 BSC ²
E1	Package length	28.00 BSC ²
R	Ankle radius	0.08 minimum to 0.25 maximum

¹The value for this measurement is for reference only.

²ANSI Y14.5M-1982 American National Standard Dimensioning and Tolerancing, Section 1.3.2, defines Basic Dimension (BSC) as: A numerical value used to describe the theoretically exact size, profile, orientation, or location of a feature or datum target. It is the basis from which permissible variations are established by tolerances on other dimensions, in notes, or in feature control frames.

Note

The drawing/dimensions are for reference only. Examples of board layout, including detailed engineering drawings and plot files, are available from Digital Semiconductor.

21172–BA Architecture Overview

This chapter describes the 21172–BA (DSW) architecture.

8.1 DSW Functional Units

The DSW connects the data paths between the memory, the 21164/Bcache, and the CIA (for PCI data).


- The memory data path, on the MEMDATA bus, has 128 bits or 256 bits of data and 16 bits or 32 bits of ECC. The width of the data path, 128 bits or 256 bits, is selected using input pins **config<1:0>**.
- The 21164/Bcache data path, on the system bus, has 128 bits of data and 16 bits of ECC.
- The CIA data path, on the IOD bus, has 64 bits of data and 8 bits of ECC.

The data paths supported by the DSW require four DSW slices to accommodate the data path widths. Each DSW slice is identical and there are no restrictions on how the bits of any bus are assigned.

The DSW is composed of mainly data buffers and multiplexers, as shown in Figure 8–1. All control for the DSW is supplied by the CIA (though some encoding and simple sequencing is performed by the DSW).

21172-BA Architecture Overview
8.1 DSW Functional Units

Figure 8-1 DSW Functional Block Diagram



LJ-04283.AI

21172—BA Architecture Overview

8.1 DSW Functional Units

The DSW functional units are:

- Victim buffer—64 bytes
- I/O read buffer—32 bytes
- Four I/O write buffers—4 * 32 bytes
- Two DMA buffer sets are provided for read and write transactions. Each set contains:
 - Memory buffer
 - Flush buffer
 - PCI buffer

8.1.1 Victim Buffer—64 Bytes

The victim buffer in each DSW slice has four 36-bit entries. Using four DSW slices results in a victim buffer with four 128-bit data entries and four 16-bit ECC entries.

The victim buffer is a free-loading first-in/first-out buffer (FIFO). The values in **cmc<8:0>** “start” and “stop” loading of the victim buffer. When buffer loading is “started,” data is shifted into the buffer during every CLKx cycle. When buffer loading is “stopped,” the contents of the buffer are held and loading is inhibited. Normally buffer loading is “stopped” before data is read out to the MEMDATA bus.

The MEMDATA bus can be configured, using **config<1:0>**, for 128-bit or 256-bit width. The buffer can be addressed in two ways:

- 128-bit—The victim buffer data has addresses 0, 1, 2, or 3. The data will be sent to memory by using the low-order memory bus bits.
- 256-bit—The victim buffer data has addresses 0 or 2 and two octawords of data plus ECC are sent to the memory in one cycle.

Data is sent to memory from a flip-flop loaded at CLKx.

8.1.2 I/O Read Buffer (IOR)—32 Bytes

The IOR buffer in each DSW slice has two 36-bit entries. Using four DSW slices results in an IOR buffer with two 128-bit data entries and two associated 16-bit ECC entries.

The IOR buffer is loaded from the IOD bus by using a quadword mask. Any or all of the quadwords can be loaded in one cycle. Only one quadword of data comes from the IOD bus during one cycle so if more than a single quadword is selected, the same data will be written into each IOR buffer quadword.

21172–BA Architecture Overview

8.1 DSW Functional Units

One instruction to the DSW controls writing the entire IOR buffer. The instruction specifies which quadword to write first. Data is read from the IOR buffer and is written into the IOR buffer on the edge of CLKx.

8.1.3 I/O Write (IOW) Buffers—4 * 32 Bytes

The IOW buffer in each DSW slice has four banks, with each bank containing two 36-bit registers. Using four DSW slices results in an IOW buffer with four banks, with each bank containing two 144-bit registers. The 144 bits contain 128 bits of data and 16 bits of ECC.

The IOW buffer is free-loading. Signals **ioc<6:0>** and **cmc<8:0>** “start” and “stop” loading of the IOW buffer.

When the IOW buffer is “started,” new data is loaded every CLKx cycle from the system bus data path (the same data is loaded into each bank). When the IOW buffer is “stopped,” loading is inhibited and the current buffer contents are held.

Data is read from the IOW buffer to the IOD bus a quadword at each CLKx cycle by the CIA issuing an instruction on **ioc<6:0>**. The instruction specifies the bank number and the bank address of the quadword.

Quadword valid bits are provided for the IOW buffer. These valid bits are used to merge the appropriate IOW buffer quadword with the memory/Bcache data. Finer granularity merging (such as bytes) is performed in the CIA (because of ECC) by looping the appropriate memory/Bcache quadword through the CIA and merging the valid write bytes.

No data is preserved in the IOW buffers across transactions (no I/O read prefetched data or posted I/O write data). Consequently, there are no coherency issues associated with I/O read/write data lingering in the buffers.

8.1.4 DMA Buffer Sets

Two DMA buffer sets are provided for read and write transactions. Each buffer set consists of three buffers:

- Memory buffer for memory data
- Flush buffer for system bus data
- PCI buffer for PCI DMA write data (not used during DMA read transactions)

The system bus and memory data could have shared one buffer, because only one of these two will provide valid data, but two buffers simplified the control logic.

21172–BA Architecture Overview

8.1 DSW Functional Units

The memory buffer (memory buffers 0 and 1) in each DSW slice has four 36-bit entries. Using four DSW slices results in a 4-entry memory buffer with 144 bits per entry (includes 128 data bits and 16 ECC bits).

The memory buffer is loaded from the MEMDATA bus on a CLKx cycle. Only the entry addressed by the instruction on **cmc<8:0>** is loaded.

Data is read from the memory buffer by addressing an octaword or quadword location in the DMA buffer while the flush buffer and PCI buffer are *not* valid. Octawords are addressed for data to the memory arrays, and quadwords are addressed for data to the IOD bus.

8.1.4.1 Flush Buffer

The flush buffer (both flush 0 and flush 1) in each DSW slice has four 36-bit entries. Using four DSW slices results in a flush buffer with four 144-bit entries (128 bits of data and 16 bits of ECC).

The flush buffer is a free-loading FIFO. Signal **cmc<8:0>** “starts” and “stops” loading of the flush buffer. While it is “started,” it is loaded from the system bus every CLKx cycle. When the flush buffer is “stopped,” loading is inhibited and the current contents are held.

“Stopping” the flush buffer makes the buffer contents valid. When the flush buffer is valid, its contents will be selected over the contents of the memory buffer during a DMA read transaction.

Octawords are addressed for data sent to the memory arrays and quadwords are addressed for data sent to the IOD bus.

8.1.4.2 PCI Buffer

The PCI buffer (both PCI 0 and PCI 1) in each DSW slice has eight 18-bit entries. Using four DSW slices results in a PCI buffer of eight 72-bit entries (64 bits data and 8 bits ECC).

The PCI buffer is loaded from the IOD bus with a quadword address under the control of **ioc<6:0>**. Each quadword has a valid bit that is set when the quadword is written.

A DMA write transaction causes data to be read from the PCI buffer. When the write address points to valid data in the PCI buffer that data will be written to the memory arrays.

A

Technical Support and Ordering Information

A.1 Obtaining Technical Support

If you need technical support or help deciding which literature best meets your needs, call the Digital Semiconductor Information Line:

United States and Canada **1-800-332-2717**
Outside North America **+1-508-628-4760**

A.2 Ordering Digital Semiconductor Products

To order the Digital Semiconductor 21164 Alpha microprocessor, Digital Semiconductor 21172 chipset, and evaluation boards, contact your local distributor.

To obtain a *Digital Semiconductor Product Catalog*, contact the Digital Semiconductor Information Line. The following table lists some of the semiconductor products available from Digital:

Product	Order Number
Digital Semiconductor 21164 366-MHz Alpha Microprocessor	21164-EB
Digital Semiconductor 21164 400-MHz Alpha Microprocessor	21164-FB
Digital Semiconductor 21164 433-MHz Alpha Microprocessor	21164-HB
Digital Semiconductor 21164 300-MHz Alpha Microprocessor for PC Products	21164-P4
Digital Semiconductor 21164 366-MHz Alpha Microprocessor for PC Products	21164-P5
Digital Semiconductor 21172-AA Core Logic Chipset	21172-AA

Technical Support and Ordering Information

A.2 Ordering Digital Semiconductor Products

Product	Order Number
Digital Semiconductor 21172-BA Data Switch Chip (four required)	21172-BA
Digital Semiconductor 21172-CA Control, I/O Interface, and Address Chip	21172-AA

A.3 Ordering Digital Semiconductor Literature

The following table lists some of the available Digital Semiconductor literature. For a complete list contact the Digital Semiconductor Information Line.

Title	Order Number
Alpha AXP Architecture Reference Manual ¹	EY-T132E-DP
Digital Semiconductor 21172 Core Logic Chipset Product Brief	EC-QUQHA-TE
Digital Semiconductor 21164 (366 MHz Through 433 MHz) Alpha Microprocessor Product Brief	EC-QP97B-TE
Digital Semiconductor 21164 (366 MHz Through 433 MHz) Alpha Microprocessor Data Sheet	EC-QP98A-TE
Digital Semiconductor 21164 (366 MHz Through 433 MHz) Alpha Microprocessor Hardware Reference Manual	EC-QP99A-TE
Digital Semiconductor AlphaPC164 Microprocessor Evaluation Board Kit Product Brief	EC-QPFXA-TE
Digital Semiconductor AlphaPC164 Microprocessor Evaluation Board User's Guide	EC-QPFYA-TE
Digital Semiconductor AlphaPC164 Microprocessor Evaluation Board Read Me First	EC-QPFZA-TE
Digital Semiconductor AlphaPC164 Motherboard Product Brief	EC-QUQKA-TE
Digital Semiconductor AlphaPC164 Motherboard User's Manual	EC-QPG0A-TE

¹To purchase the *Alpha AXP Architecture Reference Manual*, call **1-800-DIGITAL** from the U.S. or Canada, contact your local Digital office, or call Digital Press at 1-800-366-2665.

Technical Support and Ordering Information
A.4 Ordering Associated Third-Party Literature

A.4 Ordering Associated Third-Party Literature

You can order the following third-party literature directly from the vendor:

Title	Vendor
PCI Local Bus Specification PCI System Design Guide	PCI Special Interest Group 1-800-433-5177 (US) 1-503-797-4207 (International) 1-503-234-6762 (FAX)

Index

21164 Alpha microprocessor
 control and status registers, 4-2
 instruction and address logic, 3-4
21172-BA
 See DSW
21172-BA architecture
 See DSW architecture
21172-CA
 See CIA
21172-CA architecture
 See CIA architecture
21172 definition, 1-1

A

Absolute maximum ratings, 1-8
ack64_1, 2-2, 2-4, 5-2, 5-3
ad<63:00>, 4-22
ad<63:0>, 2-2, 2-4, 2-6, 2-7, 4-47, 4-48,
 4-66, 4-70, 4-77, 5-2, 5-3, 6-5 to 6-40
addr<39:4>, 2-1, 4-30, 4-53, 4-54, 5-2,
 5-3, 6-1 to 6-35
Address mapping, 6-1
 address space configuration, 6-2
 overview, 6-1
Address space, 6-2, 6-8
 21164 access, 6-4
 PCI access, 6-4
 PCI configuration space, 6-21
 device select, 6-24
 hardware-specific register space,
 6-27
 miscellaneous register space, 6-27
 special/interrupt acknowledge cycles,
 6-26

Address space (cont'd)
 PCI dense memory space, 6-9
 PCI sparse I/O space, 6-17
 PCI sparse memory space, 6-12
addr_bus_req, 2-1, 4-12, 5-2, 5-3
addr_cmd_par, 2-1, 5-2, 5-3
addr_res
 See **res<1:0>**
Alpha documentation, A-2
Alternate memory mode, 2-11
Associated literature, A-2
aux_vdd, 2-3, 2-15, 7-2, 7-8
aux_vss, 2-3, 2-15, 7-2, 7-8

B

Byte/word space, 6-27

C

cack, 2-1, 5-2, 5-3
CACK_EN register, 4-20
cas<3:0>, 2-2, 2-12, 3-6, 4-51, 5-2, 5-3
cas_l<3:0>, 2-2
cas_l<7:0>, 2-12, 3-6
cas_l<7:4>, 2-2
cbe_l<7:0>, 2-2, 2-5, 2-6, 3-5, 5-2, 5-3
CFG register, 4-19
Characteristics
 absolute maximum ratings, 1-8
 dc, 1-8
Chipset overview, 1-1
 See also System overview
 chipset characteristics, 1-8
 CIA microarchitecture, 1-4

- Chipset overview (cont'd)
 - DSW microarchitecture, 1-6
 - features, 1-1
 - microarchitecture, 1-4
 - system clocks, 1-7
- CIA, 1-1
 - definition, 1-1
- CIA architecture
 - basic properties of system, 3-16
 - data path functional units, 3-1
 - functional units, 3-3
 - DMA address logic, 3-7
 - I/O address logic, 3-6
 - memory logic, 3-6
 - PCI data path logic, 3-4
 - system data coherency, 3-15
 - transactions, 3-7
 - DMA, 3-11
 - DMA read, 3-13
 - DMA read (prefetch), 3-14
 - DMA write, 3-14
 - read miss, 3-8
 - read miss to noncacheable space, 3-9
 - read miss with victim, 3-8
 - write to noncacheable space, 3-10
- CIA control and status registers, 4-1
 - address translation registers, 4-6, 4-72
 - diagnostic registers, 4-2, 4-21
 - error registers, 4-3, 4-28
 - general registers, 4-2, 4-8
 - PCI address and scatter-gather registers, 4-5, 4-63
 - performance monitor registers, 4-3, 4-24
 - system configuration registers, 4-4, 4-49
- CIA memory operation, 3-17
 - address mapping, 3-20
 - traditional address mapping, 3-17
- CIA package dimensions, 2-30
- CIA pin descriptions, 2-1, 2-4
 - IOD bus signals, 2-8
 - memory control signals, 2-11
 - miscellaneous power signals, 2-15
 - miscellaneous signals, 2-15
 - PCI local bus signals, 2-4
 - pin groups, 2-1
- CIA pin descriptions
 - pin groups (cont'd)
 - pin totals, 2-3
 - PLL signals, 2-14
 - reserved signals, 2-16
 - system bus signals, 2-4
- CIA pin name list (alphabetic), 2-17
- CIA pin name list (alphanumeric), 2-23
- CIA_CNFG register, 4-15
- CIA_CTRL register, 4-11
- CIA_DIAG register, 4-21
- CIA_ERR register, 4-30
- CIA_REV register, 4-8
- CIA_STAT register, 4-34
- CIA_SYN register, 4-38
- cmc<8:0>**, 2-2, 2-9, 5-2, 5-3, 7-1, 7-4, 8-3, 8-4, 8-5
- cmd<3:0>**, 2-1, 5-2, 5-3
- config<1:0>**, 7-2, 7-7, 8-1, 8-3
- Control, I/O interface, and address chip
 - See* CIA
- Control and status registers
 - See* CSRs
- cpu_dat<35:0>**, 7-1, 7-3
- CPU_ERR0 register, 4-28
- CPU_ERR1 register, 4-29
- CSRs, 4-1
 - See also* CIA control and status registers

D

- dack**, 2-1, 5-2, 5-3
- Data switch chip
 - See* DSW
- Data units, xvii
- DC characteristics, 1-8
- debug**, 2-3, 2-15, 2-16
- devsel_1**, 2-2, 2-5, 5-2, 5-3, 6-31, 6-46
- DIAG_CHECK register, 4-23
- Documentation, A-2
- Document conventions, xiv
- DSW, 1-1
 - definition, 1-1

- DSW architecture, 8-1
 - functional units, 8-1
 - DMA buffer sets, 8-4
 - I/O read buffer, 8-3
 - I/O write buffer, 8-4
 - victim buffer, 8-3
- DSW package dimensions, 7-17
- DSW pin descriptions, 7-1, 7-3
 - IOD bus signals, 7-3
 - MEMDATA bus signals, 7-6
 - miscellaneous power signals, 7-8
 - miscellaneous signals, 7-7
 - pin groups, 7-1
 - pin totals, 7-2
 - PLL signals, 7-7
 - reserved signals, 7-8
 - system bus signals, 7-3
- DSW pin name list (alphabetic), 7-10
- DSW pin name list (numeric), 7-13

E

- ECC
 - cache, 1-2
 - data paths, 1-1, 1-3
 - double-bit errors, 4-32, 4-37, 4-47
 - errors, 4-38
 - IOD bus, 4-12, 4-21
 - memory, 1-2, 4-23
 - single-bit errors, 4-27, 4-31, 4-37, 4-47
 - syndromes for single-bit errors, 4-38
- EISA System Component
 - See* ESC
- error**, 2-1, 2-4, 4-12, 5-2, 5-3
- Error correction code
 - See* ECC
- ERR_MASK register, 4-36
- ESC, 6-43

F

- fill**, 2-1, 5-2, 5-3
- fill_error**, 2-1, 4-12, 5-2, 5-3
- fill_id**, 2-1, 5-2, 5-3
- frame_l**, 2-2, 2-5, 2-6, 5-2, 5-3

G

- gnt_l**, 2-2, 2-5
- gru_ack**, 2-3, 2-16
- gru_sel**, 2-3, 2-16

H

- HAE_IO register, 4-18
- HAE_MEM register, 4-16

I

- iddt**, 2-3, 2-16, 7-2, 7-8
- idle_bc**, 2-1, 4-12, 5-2, 5-3
- Initialization
 - by SROM code, 5-4
 - pin state after **sys_rst_l** deassertion, 5-3
 - pin state with **sys_rst_l** asserted, 5-2
- Input/output data bus
 - See* IOD bus
- int**, 2-1, 2-4, 5-2, 5-3
- int4_valid<3:0>**, 2-1, 4-30, 4-41
- int4_valid_h<3:0>**, 6-11, 6-13
- ioc<6:0>**, 2-2, 2-8, 5-2, 5-3, 7-1, 7-4, 8-4, 8-5
- ioc<17:0>**, 7-1, 7-4
- ioc<63:0>**, 2-2, 2-8, 5-2, 5-3, 7-4
- IOD bus
 - signal description, 2-8
 - signal group, 2-2
- iod_e<7:0>**, 2-2, 2-8, 5-2, 5-3
- irdy_l**, 2-2, 2-5, 2-7, 5-2, 5-3

L

Literature, A-2
LOCK instruction, 3-26
 CIA lock rules, 3-28
 locks to uncached space, 3-28
lock_l, 2-2, 2-5
LTB_TAG*n* register, 4-72

M

MBA*n* register, 4-52
MB instruction, 3-26
MCR register, 4-49
Mechanical specifications
 CIA, 2-30
 DSW, 7-17
Memory addressing
 PCI to physical, 6-30
Memory control signals
 signal description, 2-11
 signal group, 2-2
Memory modes, 2-11
Memory timing
 maximum signal limits, 4-60
 memory read timing figure, 4-58
 memory timing parameters, 4-57
 memory timing register *n*, 4-55
 memory write timing figure, 4-59
 minimum signal limits, 4-60
mem_ack_l, 2-2, 2-13, 4-14, 5-2, 5-3
mem_addr<12:0>, 2-2, 2-13, 5-2, 5-3
mem_b0, 2-2, 2-13, 5-2, 5-3
mem_cs_l, 2-2, 2-5, 6-42, 6-43, 6-44, 6-46
mem_dat<71:0>, 7-1, 7-6
mem_en, 2-2, 2-13, 3-6, 4-57, 5-2, 5-3,
 7-1, 7-6
MEM_ERR0 register, 4-40
MEM_ERR1 register, 4-41
mem_req_l, 2-2, 2-13
mem_we_l<1:0>, 2-2, 2-13, 3-6, 5-2, 5-3

N

NMI, 2-7

O

Operating temperature (CIA), 1-8
Operating temperature (DSW), 1-8
Ordering associated third-party literature,
 A-3
Ordering products, A-1
Organization of document, xiii

P

par, 2-2, 2-6, 5-2, 5-3
par64, 2-2, 2-6, 5-2, 5-3
Parts
 ordering, A-1
PCEB, 6-43
PCI local bus signals
 signal description, 2-4
 signal group, 2-2
PCI-to-EISA bridge, 6-43
 See PCEB
PCI window
 suggested use, 6-42
PCI_ERR0 register, 4-44
PCI_ERR1 register, 4-47
PCI_ERR2 register, 4-48
PCI_LAT register, 4-9
PERF_CONTROL register, 4-25
PERF_MONITOR register, 4-24
perr_l, 2-2, 2-6, 5-2, 5-3
Phase-locked loop signals
 signal description, 2-14
 signal group, 2-2
Pin descriptions
 CIA, 2-1 to 2-16
 DSW, 7-1 to 7-9
pll_agnd, 2-2, 2-14, 7-2, 7-8
pll_clk, 2-2, 2-14, 7-2, 7-8

pll_lp2, 2-2, 2-14, 7-2, 7-8
pll_test_h<3:0>, 2-2, 2-15, 2-16
pll_vdd, 2-3, 2-15, 7-2, 7-8
pll_vss, 2-3, 2-15, 7-2, 7-8
Power
 dc supply, 1-8
 dissipation (CIA), 1-8
 dissipation (DSW), 1-8
Power-up, 5-1
 internal reset, 5-1
procmon_out, 7-2, 7-8

R

ras<3:0>, 2-2, 2-13, 3-6, 4-50, 4-51, 5-2, 5-3
rasx_l<7:0>, 2-2, 2-13

Registers

CACK_EN, 4-20
CFG, 4-19
CIA_CNFG, 4-15
CIA_CTRL, 4-11
CIA_DIAG, 4-21
CIA_ERR, 4-30
CIA_REV, 4-8
CIA_STAT, 4-34
CIA_SYN, 4-38
CPU_ERR0, 4-28
CPU_ERR1, 4-29
DIAG_CHECK, 4-23
ERR_MASK, 4-36
HAE_IO, 4-18
HAE_MEM, 4-16
LTB_TAG n , 4-72
MBA n , 4-52
MCR, 4-49
MEM_ERR0, 4-40
MEM_ERR1, 4-41
PCI_ERR0, 4-44
PCI_ERR1, 4-47
PCI_ERR2, 4-48
PCI_LAT, 4-9
PERF_CONTROL, 4-25
PERF_MONITOR, 4-24
TBIA, 4-63

Registers (cont'd)

TBm_PAGE n , 4-76
TB_TAG n , 4-74
TMG n , 4-55
Tn_BASE, 4-68
Wn_BASE, 4-64
Wn_MASK, 4-66
W_DAC, 4-71
Related documentation, A-2
req64_l, 2-2, 2-4, 2-5, 2-6, 5-2, 5-3
req_l, 2-2, 2-6, 5-2, 5-3
res<1:0>, 2-1
reset_l, 7-2, 7-7
rst_l, 2-2, 2-6, 5-2, 5-3

S

scan_out, 7-2, 7-7
Semiconductor Information Line, A-1
serr_l, 2-2, 2-7, 5-2, 5-3
Set select memory mode, 2-11
set_sel<15:0>, 2-2, 2-13, 5-2, 5-3
stop_l, 2-2, 2-7, 5-2, 5-3
System bus
 signal descriptions, 2-4
 signal group, 2-1
System clocks, 1-7
System data coherency, 3-15
 post and run I/O write data, 3-17
 system basic properties, 3-16
System functional units, 3-1
System overview, 1-2
 See also Chipset overview
sys_clk_out1_h, 1-7, 2-4, 7-3
sys_rst_l, 2-2, 2-15, 5-1

T

tag_ctl_par, 2-1, 5-2, 5-3
tag_dirty, 2-1, 5-2, 5-3
TBIA register, 4-63
TBm_PAGE n register, 4-76
TB_TAG n register, 4-74

td, 2-3, 2-16
Technical support, A-1
Temperature
 Operating (CIA), 1-8
 Operating (DSW), 1-8
 storage, 1-8
test_mode<1:0>, 2-2, 2-15, 7-2, 7-7
test_out, 2-2, 2-15, 5-2, 5-3
TMG n register, 4-55
T n _BASE register, 4-68
trdy_1, 2-2, 2-5, 2-7, 5-2, 5-3

V

Vdd, 2-4
victim_pending, 2-1
Vss, 7-3

W

W3_BASE
 dual address cycle (DAC) enable, 6-33
W3_BASE register
 dual address cycle (DAC), 4-64
 illustration and description, 4-64
W n _BASE register, 4-64
W n _MASK register, 4-66
W_DAC register, 4-71