

Universe II™ User Manual



Spring 1998

<http://www.tundra.com>

The information in this document is subject to change without notice and should not be construed as a commitment by Tundra Semiconductor Corporation. While reasonable precautions have been taken, Tundra Semiconductor Corporation assumes no responsibility for any errors that may appear in this document.

No part of this document may be copied or reproduced in any form or by any means without the prior written consent of Tundra Semiconductor Corporation.

Tundra® products are not designed, intended, or authorized for use as components in systems intended for surgical implant into the body, or other applications intended to support or sustain life, or for any other application in which the failure of the Tundra product could create a situation where personal injury or death may occur. Should Buyer purchase or use Tundra products for any such unintended or unauthorized application, Buyer shall indemnify and hold Tundra and its officers, employees, subsidiaries, affiliates and distributors harmless against all claims, costs, damages, and expenses, and reasonable attorney fees arising out of, directly or indirectly, any claim of personal injury or death associated with such unintended or unauthorized use, even if such claim alleges that Tundra was negligent regarding the design or manufacture of the part.

The acceptance of this document will be construed as an acceptance of the foregoing conditions.

Universe II™ User Manual

Copyright 1998, Tundra Semiconductor Corporation
All rights reserved.

Document: 8091142.MD300.01

Printed in Canada

Tundra and Tundra logo are registered trademarks of Tundra Semiconductor Corporation. Universe II, Universe, SCV64, Trooper II and QSpan are trademarks of Tundra Semiconductor Corporation. BI-Mode is a registered trademark of DY-4 Systems, Inc.

Overview

Chapter 1	Introduction
Chapter 2	Functional Description
Chapter 3	Description of Signals
Chapter 4	Signals and DC Characteristics
Appendix A	Registers
Appendix B	Performance
Appendix C	Typical Applications
Appendix D	Reliability Prediction
Appendix E	Cycle Mapping
Appendix F	Operating and Storage Conditions
Appendix G	Mechanical and Ordering Information

Table of Contents

1	Introduction	1-1
1.1	Features	1-1
1.2	Benefits of the Universe II	1-2
1.3	Past and Future of the Universe	1-3
1.4	About This Document	1-4
1.5	Universe II Technical Support	1-5
1.6	Universe II Documentation	1-5
1.7	Conventions	1-6
	1.7.1 Signals	1-6
	1.7.2 Symbols	1-6
	1.7.3 Terminology	1-7
2	Functional Description.....	2-1
2.1	Architectural Overview	2-2
2.1.1	VMEbus Interface	2-2
2.1.1.1	Universe II as VMEbus Slave	2-2
2.1.1.2	Universe II as VMEbus Master	2-3
2.1.2	PCI Bus Interface	2-4
2.1.2.1	Universe II as PCI Target	2-4
2.1.2.2	Universe II as PCI Master	2-4
2.1.3	Interrupter and Interrupt Handler	2-4
2.1.3.1	Interrupter	2-4
2.1.3.2	VMEbus Interrupt Handling	2-5
2.1.4	DMA Controller	2-5
2.2	VMEbus Interface	2-6
2.2.1	VMEbus Requester	2-6
2.2.1.1	Internal Arbitration for VMEbus Requests	2-6
2.2.1.2	Request Modes	2-7
2.2.1.3	VMEbus Release	2-8

2.2.2	Universe II as VMEbus Master	2-9
2.2.2.1	Addressing Capabilities	2-9
2.2.2.2	Data Transfer Capabilities	2-10
2.2.2.3	Cycle Terminations	2-13
2.2.3	Universe as VMEbus Slave	2-13
2.2.3.1	Coupled Transfers	2-14
2.2.3.2	Posted Writes	2-15
2.2.3.3	Prefetched Block Reads	2-16
2.2.3.4	VMEbus Lock Commands (ADOH Cycles).....	2-18
2.2.3.5	VMEbus Read-Modify-Write Cycles (RMW Cycles)..	2-19
2.2.3.6	Register Accesses.....	2-19
2.2.3.7	Location Monitors.....	2-19
2.2.3.8	Generating PCI Configuration Cycles	2-20
2.2.4	VMEbus Configuration	2-23
2.2.4.1	First Slot Detector	2-23
2.2.4.2	VMEbus Register Access at Power-up	2-23
2.2.5	Automatic Slot Identification	2-24
2.2.5.1	Auto Slot ID: VME64 Specified.....	2-24
2.2.5.2	Auto-ID: A Proprietary Tundra Method	2-25
2.2.6	System Controller Functions	2-26
2.2.6.1	System Clock Driver.....	2-26
2.2.6.2	VMEbus Arbiter.....	2-27
2.2.6.3	IACK Daisy-Chain Driver Module.....	2-27
2.2.6.4	VMEbus Time-out	2-28
2.2.7	BI-Mode	2-28
2.3	PCI Bus Interface	2-30
2.3.1	PCI Cycles—Overview	2-30
2.3.1.1	32-Bit Versus 64-Bit PCI	2-30
2.3.1.2	PCI Bus Request and Parking	2-31
2.3.1.3	Address Phase	2-31
2.3.1.4	Data Transfer	2-33

	2.3.1.5	Termination Phase	2-33
	2.3.1.6	Parity Checking.....	2-34
	2.3.2	Universe II as PCI Master	2-35
	2.3.2.1	PCI Burst Transfers	2-36
	2.3.2.2	Termination.....	2-37
	2.3.2.3	Parity	2-37
	2.3.3	Universe II as PCI Target.....	2-38
	2.3.3.1	Overview.....	2-38
	2.3.3.2	Data Transfer	2-39
	2.3.3.3	Coupled Transfers.....	2-42
	2.3.3.4	Posted Writes	2-44
	2.3.3.5	The Special Cycle Generator	2-45
	2.3.3.6	Using the VOWN bit	2-48
	2.3.3.7	Terminations	2-49
2.4		Slave Image Programming.....	2-50
	2.4.1	VME Slave Images	2-50
	2.4.1.1	VMEbus Fields	2-51
	2.4.1.2	PCI Bus Fields	2-51
	2.4.1.3	Control Fields	2-52
	2.4.2	PCI Bus Target Images	2-53
	2.4.2.1	PCI Bus Fields	2-54
	2.4.2.2	VMEbus Fields	2-54
	2.4.2.3	Control Fields	2-55
	2.4.3	Special PCI Target Image	2-55
2.5		Bus Error Handling	2-58
	2.5.1	Coupled Cycles	2-58
	2.5.2	Decoupled Transactions	2-58
	2.5.2.1	Posted Writes	2-58
	2.5.2.2	Prefetched Reads.....	2-60
	2.5.2.3	DMA Errors	2-60
	2.5.2.4	Parity Errors	2-60

2.6	Interrupt Generation	2-62
2.6.1	PCI Interrupt Generation	2-63
2.6.2	VMEbus Interrupt Generation.....	2-65
2.7	Interrupt Handling	2-68
2.7.1	PCI Interrupt Handling.....	2-68
2.7.2	VMEbus Interrupt Handling.....	2-68
2.7.2.1	Bus Error During VMEbus IACK Cycle	2-70
2.7.3	Internal Interrupt Handling.....	2-71
2.7.3.1	VMEbus and PCI Software Interrupts	2-73
2.7.3.2	Software IACK Interrupt	2-74
2.7.3.3	VMEbus Ownership Interrupt.....	2-75
2.7.3.4	DMA Interrupt	2-75
2.7.3.5	Mailbox Register Access Interrupts.....	2-75
2.7.3.6	Location Monitors.....	2-75
2.7.3.7	PCI and VMEbus Error Interrupts	2-76
2.7.4	VME64 Auto-ID	2-76
2.8	DMA Controller	2-77
2.8.1	DMA Registers Outline.....	2-77
2.8.1.1	Source and Destination Addresses.....	2-78
2.8.1.2	Transfer Size	2-79
2.8.1.3	Transfer Data Width	2-79
2.8.1.4	DMA Command Packet Pointer	2-80
2.8.1.5	DMA Control and Status	2-80
2.8.2	Direct Mode Operation.....	2-83
2.8.3	Linked-List Operation	2-86
2.8.3.1	Linked List Updating	2-91
2.8.4	FIFO Operation and Bus Ownership.....	2-92
2.8.4.1	PCI to VMEbus Transfers.....	2-92
2.8.4.2	VMEbus to PCI Transfers.....	2-94
2.8.5	DMA Interrupts	2-95
2.8.6	Interactions with Other Channels	2-96

2.8.7	DMA Error Handling	2-96
2.8.7.1	DMA Software Response to Error	2-97
2.8.7.2	DMA Hardware Response to Error	2-97
2.8.7.3	Resuming DMA Transfers	2-98
2.9	Registers	2-100
2.9.1	Overview of Universe II Registers	2-100
2.9.2	Register Access from the PCI Bus	2-101
2.9.2.1	PCI Configuration Access	2-102
2.9.2.2	Memory or I/O Access	2-103
2.9.2.3	Locking the Register Block from the PCI bus	2-103
2.9.3	Register Access from the VMEbus	2-104
2.9.3.1	VMEbus Register Access Image (VRAI)	2-104
2.9.3.2	CR/CSR Accesses	2-106
2.9.3.3	RMW and ADOH Register Access Cycles	2-106
2.9.4	Mailbox Registers	2-108
2.9.5	Semaphores	2-109
2.10	Utility Functions	2-110
2.10.1	Resets	2-110
2.10.1.1	Overview of Reset Support	2-110
2.10.1.2	Universe II Reset Circuitry	2-112
2.10.1.3	Reset Implementation Cautions	2-114
2.10.2	Power-Up Options	2-115
2.10.2.1	Power-up Option Descriptions	2-117
2.10.2.2	Power-Up Option Implementation	2-119
2.10.3	Hardware Initialization (Normal Operating Mode)	2-121
2.10.4	Test Modes	2-122
2.10.4.1	Auxiliary Test Modes	2-122
2.10.4.2	JTAG support	2-123
2.10.5	Clocks	2-123
3	Description of Signals	3-1
3.1	VMEbus Signals	3-1

3.2	PCI Bus Signals	3-5
4	Signals and DC Characteristics	4-1
4.1	Terminology	4-1
4.2	DC Characteristics and Pin Assignments.....	4-2
Appendix A – Registers		App A-1
Appendix B – Performance		App B-1
B.1	PCI Slave Channel	App B-2
B.1.1	Coupled Cycles	App B-2
B.1.1.1	Request of VMEbus.....	App B-2
B.1.1.2	Read Cycles	App B-2
B.1.1.3	Write Cycles.....	App B-4
B.1.2	Decoupled Cycles.....	App B-4
B.2	VME Slave Channel.....	App B-6
B.2.1	Coupled Cycles	App B-6
B.2.1.1	Block vs. non-Block Transfers	App B-6
B.2.1.2	Read Cycles	App B-7
B.2.1.3	Write Cycles.....	App B-8
B.2.2	Decoupled Cycles.....	App B-9
B.2.2.1	Write Cycles.....	App B-9
B.2.2.2	Prefetched Read Cycles	App B-12
B.3	DMA Channel	App B-14
B.3.1	Relative FIFO sizes	App B-14
B.3.2	VMEbus Ownership Modes	App B-14
B.3.3	VME Transfers	App B-14
B.3.3.1	Read Transfers	App B-15
B.3.3.2	Write Transfers	App B-15
B.3.4	PCI Transfers.....	App B-15
B.4	Summary	App B-17

Appendix C – Typical Applications	App C-1
C.1 VME Interface.....	App C-1
C.1.1 Transceivers	App C-1
C.1.1.1 Pull-down resistors	App C-5
C.1.2 Direction control	App C-5
C.1.3 Power-up Options.....	App C-5
C.2 PCI Bus Interface	App C-7
C.2.1 Resets	App C-8
C.2.2 Local Interrupts	App C-9
C.3 Manufacturing Test Pins	App C-10
C.4 Decoupling VDD and VSS on the Universe II	App C-10
Appendix D – Reliability Prediction	App D-1
D.1 Physical characteristics	App D-1
D.2 Thermal characteristics	App D-1
D.3 Universe II Ambient Operating Calculations.....	App D-2
D.4 Thermal vias.....	App D-3
Appendix E – Cycle Mapping	App E-1
E.1 Little-endian Mode.....	App E-1
Appendix F – Operating and Storage Conditions	App F-1
Appendix G – Mechanical and Ordering Information	App G-1
G.1 Mechanical Information.....	App G-1
G.2 Ordering Information	App G-5
Index	Index-1

List of Figures

Figure 2.1 :	Architectural Diagram for the Universe II	2-3
Figure 2.2 :	Influence of Transaction Data Width and Target Image Data Width on Data Packing/Unpacking	2-12
Figure 2.3 :	VMEbus Slave Channel Dataflow.....	2-14
Figure 2.4 :	Timing for Auto-ID Cycle.....	2-26
Figure 2.5 :	PCI Bus Target Channel Dataflow	2-40
Figure 2.6 :	Influence of Transaction Data Width and Target Image Data Width on Data Packing/Unpacking	2-41
Figure 2.7 :	Address Translation Mechanism for VMEbus to PCI Bus Transfers .	2-52
Figure 2.8 :	Address Translation Mechanism for PCI Bus to VMEbus Transfers .	2-54
Figure 2.9 :	Memory Mapping in the Special PCI Target Image	2-57
Figure 2.10 :	Universe Interrupt Circuitry	2-62
Figure 2.11 :	STATUS/ID Provided by Universe II.....	2-67
Figure 2.12 :	Sources of Internal Interrupts	2-72
Figure 2.13 :	Direct Mode DMA transfers.....	2-84
Figure 2.14 :	Command Packet Structure and Linked List Operation.....	2-87
Figure 2.15 :	DMA Linked List Operation	2-88
Figure 2.16 :	Universe II Control and Status Register Space	2-101
Figure 2.17 :	PCI Bus Access to UCSR as Memory or I/O Space	2-102
Figure 2.18 :	UCSR Access from the VMEbus Register Access Image.....	2-105
Figure 2.19 :	UCSR Access in VMEbus CR/CSR Space	2-107
Figure 2.20 :	Reset Circuitry	2-113
Figure 2.21 :	Resistor-Capacitor Circuit Ensuring Power-Up Reset Duration	2-114
Figure 2.22 :	Power-up Options Timing	2-120
Figure A.1 :	UCSR Access Mechanisms	App A-1
Figure B.1 :	Coupled Read Cycle - Universe II as VME Master	App B-3
Figure B.2 :	Several Coupled Read Cycles - Universe II as VME Master.....	App B-3
Figure B.3 :	Coupled Write Cycle - Universe II as VME Master	App B-4

Figure B.4 :	Several Non-Block Decoupled Writes - Universe II as VME Master	App B-6
Figure B.5 :	BLT Decoupled Write - Universe II as VME Master	App B-6
Figure B.6 :	Coupled Read Cycle - Universe II as VME Slave	App B-8
Figure B.7 :	Coupled Write Cycle - Universe II as VME Slave (bus parked at Universe II)	App B-9
Figure B.8 :	Non-Block Decoupled Write Cycle - Universe II as VME Slave	App B-10
Figure B.9 :	BLT Decoupled Write Cycle - Universe II as VME Slave	App B-11
Figure B.10 :	MBLT Decoupled Write Cycle - Universe II as VME Slave.....	App B-11
Figure B.11 :	BLT Pre-fetched Read Cycle - Universe II as VME Slave	App B-13
Figure B.12 :	PCI Read Transactions During DMA Operation.....	App B-16
Figure B.13 :	Multiple PCI Read Transactions During DMA Operation.....	App B-17
Figure C.1 :	Universe II Connections to the VMEbus Through TTL Buffers ..	App C-2
Figure C.2 :	Power-up Configuration Using Passive Pull-ups	App C-6
Figure C.3 :	Power-up Configuration Using Active Circuitry	App C-7
Figure C.4 :	Analog Isolation Scheme.....	App C-10
Figure C.5 :	Noise Filter Scheme	App C-11
Figure G.1 :	Mechanical Dimensions for the 324-Pin Ceramic BGA Package	App G-1
Figure G.2 :	313-PBGA (Generic).....	App G-3
Figure G.3 :	313-PBGA Top View (OMPAC and GTPAC Drawings).....	App G-4

List of Tables

Table 1.1 :	Suffixes for Active Low Signals	1-6
Table 2.2 :	PCI Address Line Asserted as a Function of VA[15:11]	2-21
Table 2.3 :	Command Type Encoding for Transfer Type	2-32
Table 2.4 :	Register Fields for the Special Cycle Generator	2-45
Table 2.5 :	VMEbus Fields for VMEbus Slave Image	2-50
Table 2.6 :	PCI Bus Fields for VMEbus Slave Image	2-50
Table 2.7 :	Control Fields for VMEbus Slave Image	2-50
Table 2.8 :	PCI Bus Fields for the PCI Bus Target Image	2-53
Table 2.9 :	VMEbus Fields for the PCI Bus Target Image	2-53
Table 2.10 :	Control Fields for PCI Bus Target Image	2-53
Table 2.11 :	PCI Bus Fields for the Special PCI Target Image	2-56
Table 2.12 :	VMEbus Fields for the Special PCI Bus Target Image	2-56
Table 2.13 :	Control Fields for the Special PCI Bus Target Image	2-56
Table 2.14 :	Source, Enabling, Mapping, and Status of PCI Interrupt Output	2-64
Table 2.15 :	Source, Enabling, Mapping, and Status of VMEbus Interrupt Outputs	2-66
Table 2.16 :	Internal Interrupt Routing	2-71
Table 2.17 :	DMA Interrupt Sources and Enable Bits	2-95
Table 2.18 :	Programming the VMEbus Register Access Image	2-104
Table 2.19 :	Hardware Reset Mechanism	2-110
Table 2.20 :	Software Reset Mechanism	2-111
Table 2.21 :	Functions Affected by Reset Initiations	2-112
Table 2.22 :	Power-Up Options	2-116
Table 2.23 :	VRAI Base Address Power-up Options	2-117
Table 2.24 :	Manufacturing Pin Requirements for Normal Operating Mode	2-121
Table 2.25 :	Test Mode Operation	2-122
Table 4.1 :	DC Electrical Characteristics ($V_{DD} = 5\text{ V} \pm 10\%$)	4-2
Table 4.2 :	Pin List and DC Characteristics for Universe II Signals	4-3

Table 4.3 :	PCI Bus Address/Data Pins	4-7
Table 4.4 :	VMEbus Address Pins.....	4-9
Table 4.5 :	VMEbus Data Pins	4-10
Table 4.6 :	Pin Assignments for Power and Ground	4-11
Table 4.7 :	Pinout for 313-pin Plastic BGA Package	4-12
Table 4.8 :	Pinout for 324-pin Ceramic BGA Package.....	4-13
Table A.1 :	Universe II Register Map	App A-2
Table A.2 :	PCI Configuration Space ID Register (PCI_ID)	App A-7
Table A.3 :	PCI Configuration Space Control and Status Register (PCI_CSR).....	App A-8
Table A.4 :	PCI Configuration Class Register (PCI_CLASS)	App A-10
Table A.5 :	PCI Configuration Miscellaneous 0 Register (PCI_MISC0)	App A-11
Table A.6 :	PCI Configuration Base Address Register (PCI_BS0)	App A-12
Table A.7 :	PCI Configuration Base Address 1 Register (PCI_BS1)	App A-13
Table A.8 :	PCI Configuration Miscellaneous 1 Register (PCI_MISC1)	App A-14
Table A.9 :	PCI Target Image 0 Control (LSI0_CTL)	App A-15
Table A.10 :	PCI Target Image 0 Base Address Register (LSI0_BS).....	App A-16
Table A.11 :	PCI Target Image 0 Bound Address Register (LSI0_BD)	App A-17
Table A.12 :	PCI Target Image 0 Translation Offset (LSI0_TO)	App A-18
Table A.13 :	PCI Target Image 1 Control (LSI1_CTL)	App A-19
Table A.14 :	PCI Target Image 1 Base Address Register (LSI1_BS).....	App A-20
Table A.15 :	PCI Target Image 1 Bound Address Register (LSI1_BD)	App A-21
Table A.16 :	PCI Target Image 1 Translation Offset (LSI1_TO)	App A-22
Table A.17 :	PCI Target Image 2 Control (LSI2_CTL)	App A-23
Table A.18 :	PCI Target Image 2 Base Address Register (LSI2_BS).....	App A-24
Table A.19 :	PCI Target Image 2 Bound Address Register (LSI2_BD)	App A-25
Table A.20 :	PCI Target Image 2 Translation Offset (LSI2_TO)	App A-26
Table A.21 :	PCI Target Image 3 Control (LSI3_CTL)	App A-27
Table A.22 :	PCI Target Image 3 Base Address Register (LSI3_BS).....	App A-28
Table A.23 :	PCI Target Image 3 Bound Address Register (LSI3_BD)	App A-29
Table A.24 :	PCI Target Image 3 Translation Offset (LSI3_TO)	App A-30

Table A.25 :	Special Cycle Control Register (SCYC_CTL).....	App A-31
Table A.26 :	Special Cycle PCI Bus Address Register (SCYC_ADDR).....	App A-32
Table A.27 :	Special Cycle Swap/Compare Enable Register (SCYC_EN)	App A-33
Table A.28 :	Special Cycle Compare Data Register (SCYC_CMP).....	App A-34
Table A.29 :	Special Cycle Swap Data Register (SCYC_SWP).....	App A-35
Table A.30 :	PCI Miscellaneous Register (LMISC).....	App A-36
Table A.31 :	Special PCI Target Image (SLSI).....	App A-37
Table A.32 :	PCI Command Error Log Register (L_CMDERR)	App A-39
Table A.33 :	PCI Address Error Log (LAERR)	App A-40
Table A.34 :	PCI Target Image 4 Control Register (LSI4_CTL).....	App A-41
Table A.35 :	PCI Target Image 4 Base Address Register (LSI4_BS).....	App A-42
Table A.36 :	PCI Target Image 4 Bound Address Register (LSI4_BD)	App A-43
Table A.37 :	PCI Target Image 4 Translation Offset (LSI4_TO)	App A-44
Table A.38 :	PCI Target Image 5 Control Register (LSI5_CTL).....	App A-45
Table A.39 :	PCI Target Image 5 Base Address Register (LSI5_BS).....	App A-46
Table A.40 :	PCI Target Image 5 Bound Address Register (LSI5_BD)	App A-47
Table A.41 :	PCI Target Image 5 Translation Offset (LSI5_TO)	App A-48
Table A.42 :	PCI Target Image 6 Control Register (LSI6_CTL).....	App A-49
Table A.43 :	PCI Target Image 6 Base Address Register (LSI6_BS).....	App A-50
Table A.44 :	PCI Target Image 6 Bound Address Register (LSI6_BD)	App A-51
Table A.45 :	PCI Target Image 6 Translation Offset (LSI6_TO)	App A-52
Table A.46 :	PCI Target Image 7 Control Register (LSI7_CTL).....	App A-53
Table A.47 :	PCI Target Image 7 Base Address Register (LSI7_BS).....	App A-54
Table A.48 :	PCI Target Image 7 Bound Address Register (LSI7_BD)	App A-55
Table A.49 :	PCI Target Image 7 Translation Offset (LSI7_TO)	App A-56
Table A.50 :	DMA Transfer Control Register (DCTL).....	App A-57
Table A.51 :	DMA Transfer Byte Count Register (DTBC)	App A-58
Table A.52 :	DMA PCI Bus Address Register (DLA)	App A-59
Table A.53 :	DMA VMEbus Address Register (DVA).....	App A-60
Table A.54 :	DMA Command Packet Pointer (DCPP)	App A-61
Table A.55 :	DMA General Control/Status Register (DGCS)	App A-62

Table A.56 :	DMA Linked List Update Enable Register (D_LLUE)	App A-64
Table A.57 :	PCI Interrupt Enable Register (LINT_EN)	App A-65
Table A.58 :	PCI Interrupt Status Register (LINT_STAT)	App A-67
Table A.59 :	PCI Interrupt Map 0 Register (LINT_MAP0)	App A-69
Table A.60 :	PCI Interrupt Map 1 Register (LINT_MAP1)	App A-70
Table A.61 :	VMEbus Interrupt Enable Register (VINT_EN)	App A-71
Table A.62 :	VMEbus Interrupt Status Register (VINT_STAT)	App A-73
Table A.63 :	VME Interrupt Map 0 Register (VINT_MAP0)	App A-75
Table A.64 :	VME Interrupt Map 1 Register (VINT_MAP1)	App A-76
Table A.65 :	Interrupt STATUS/ID Out Register (STATID)	App A-77
Table A.66 :	VIRQ1 STATUS/ID Register (V1_STATID)	App A-78
Table A.67 :	VIRQ2 STATUS/ID Register (V2_STATID)	App A-79
Table A.68 :	VIRQ3 STATUS/ID Register (V3_STATID)	App A-80
Table A.69 :	VIRQ4 STATUS/ID Register (V4_STATID)	App A-81
Table A.70 :	VIRQ5 STATUS/ID Register (V5_STATID)	App A-82
Table A.71 :	VIRQ6 STATUS/ID Register (V6_STATID)	App A-83
Table A.72 :	VIRQ7 STATUS/ID Register (V7_STATID)	App A-84
Table A.73 :	PCI Interrupt Map 2 Register (LINT_MAP2)	App A-85
Table A.74 :	VME Interrupt Map 2 Register (VINT_MAP2)	App A-86
Table A.75 :	Mailbox 0 Register (MBOX0)	App A-87
Table A.76 :	Mailbox 1 Register (MBOX1)	App A-88
Table A.77 :	Mailbox 2 Register (MBOX2)	App A-89
Table A.78 :	Mailbox 3 Register (MBOX3)	App A-90
Table A.79 :	Semaphore 0 Register (SEMA0)	App A-91
Table A.80 :	Semaphore 1 Register (SEMA1)	App A-92
Table A.81 :	Master Control Register (MAST_CTL)	App A-93
Table A.82 :	Miscellaneous Control Register (MISC_CTL)	App A-95
Table A.83 :	Miscellaneous Status Register (MISC_STAT)	App A-97
Table A.84 :	User AM Codes Register (USER_AM)	App A-98
Table A.85 :	VMEbus Slave Image 0 Control (VSI0_CTL)	App A-99
Table A.86 :	VMEbus Slave Image 0 Base Address Register (VSI0_BS)....	App A-100
Table A.87 :	VMEbus Slave Image 0 Bound Address Register (VSI0_BD)	App A-101
Table A.88 :	VMEbus Slave Image 0 Translation Offset (VSI0_TO)	App A-102

Table A.89 :	VMEbus Slave Image 1 Control (VSI1_CTL).....	App A-103
Table A.90 :	VMEbus Slave Image 1 Base Address Register (VSI1_BS)....	App A-104
Table A.91 :	VMEbus Slave Image 1 Bound Address Register (VSI1_BD)	App A-105
Table A.92 :	VMEbus Slave Image 1 Translation Offset (VSI1_TO)	App A-106
Table A.93 :	VMEbus Slave Image 2 Control (VSI2_CTL).....	App A-107
Table A.94 :	VMEbus Slave Image 2 Base Address Register (VSI2_BS)....	App A-108
Table A.95 :	VMEbus Slave Image 2 Bound Address Register (VSI2_BD)	App A-109
Table A.96 :	VMEbus Slave Image 2 Translation Offset (VSI2_TO)	App A-110
Table A.97 :	VMEbus Slave Image 3 Control (VSI3_CTL).....	App A-111
Table A.98 :	VMEbus Slave Image 3 Base Address Register (VSI3_BS)....	App A-112
Table A.99 :	VMEbus Slave Image 3 Bound Address Register (VSI3_BD)	App A-113
Table A.100 :	VMEbus Slave Image 3 Translation Offset (VSI3_TO)	App A-114
Table A.101 :	Location Monitor Control Register (LM_CTL)	App A-115
Table A.102 :	Location Monitor Base Address Register (LM_BS).....	App A-116
Table A.103 :	VMEbus Register Access Image Control Register (VRAI_CTL)	App A-117
Table A.104 :	VMEbus Register Access Image Base Address Register (VRAI_BS).....	App A-118
Table A.105 :	VMEbus CSR Control Register (VCSR_CTL).....	App A-119
Table A.106 :	VMEbus CSR Translation Offset (VCSR_TO)	App A-120
Table A.107 :	VMEbus AM Code Error Log (V_AMERR)	App A-121
Table A.108 :	VMEbus Address Error Log (VAERR)	App A-122
Table A.109 :	VMEbus Slave Image 4 Control (VSI4_CTL).....	App A-123
Table A.110 :	VMEbus Slave Image 4 Base Address Register (VSI4_BS)....	App A-124
Table A.111 :	VMEbus Slave Image 4 Bound Address Register (VSI4_BD)	App A-125
Table A.112 :	VMEbus Slave Image 4 Translation Offset (VSI4_TO)	App A-126
Table A.113 :	VMEbus Slave Image 5 Control (VSI5_CTL).....	App A-127
Table A.114 :	VMEbus Slave Image 5 Base Address Register (VSI5_BS)....	App A-128
Table A.115 :	VMEbus Slave Image 5 Bound Address Register (VSI5_BD)	App A-129
Table A.116 :	VMEbus Slave Image 5 Translation Offset (VSI5_TO)	App A-130
Table A.117 :	VMEbus Slave Image 6 Control (VSI6_CTL).....	App A-131
Table A.118 :	VMEbus Slave Image 6 Base Address Register (VSI6_BS)....	App A-132

Table A.119 :	VMEbus Slave Image 6 Bound Address Register (VSI6_BD)	App A-133
Table A.120 :	VMEbus Slave Image 6 Translation Offset (VSI6_TO)	App A-134
Table A.121 :	VMEbus Slave Image 7 Control (VSI7_CTL)	App A-135
Table A.122 :	VMEbus Slave Image 7 Base Address Register (VSI7_BS)....	App A-136
Table A.123 :	VMEbus Slave Image 7 Bound Address Register (VSI7_BD)	App A-137
Table A.124 :	VMEbus Slave Image 7 Translation Offset (VSI7_TO)	App A-138
Table A.125 :	VMEbus CSR Bit Clear Register (VCSR_CLR)	App A-139
Table A.126 :	VMEbus CSR Bit Set Register (VCSR_SET)	App A-140
Table A.127 :	VMEbus CSR Base Address Register (VCSR_BS)	App A-141
Table B.1 :	PCI Slave Channel Performance	App B-17
Table B.2 :	VME Slave Channel Performance	App B-18
Table B.3 :	DMA Channel Performance	App B-19
Table C.1 :	VMEbus Signal Drive Strength Requirements	App C-4
Table C.2 :	VMEbus Transceiver Requirements	App C-4
Table C.3 :	Reset Signals	App C-8
Table D.1 :	Ambient to Junction Thermal Impedance	App D-2
Table D.2 :	Maximum Universe II Junction Temperature	App D-2
Table E.1 :	Mapping of 32-bit Little-Endian PCI Bus to 32-bit VMEbus	App E-1
Table E.2 :	Mapping of 32-bit Little-Endian PCI Bus to 64-bit VMEbus	App E-2
Table F.1 :	Recommended Operating Conditions	App F-1
Table F.2 :	Absolute Maximum Ratings	App F-1
Table F.3 :	Power Dissipation	App F-2
Table G.1 :	Standard Ordering Information	App G-5

1 Introduction

1.1 Features

The Universe II (CA91C142) is the *de facto* industry standard PCI bus to VMEbus bridge, providing:

- 64-bit, 33 MHz PCI bus interface,
- fully compliant, high performance 64-bit VMEbus interface,
- integral FIFOs buffer multiple transactions in both directions,
- programmable DMA controller with linked-list support,
- industry leading performance,
- Wide range of VMEbus address and data transfer modes,
 - A32/A24/A16 master and slave, (not A64 or A40)
 - D64/D32/D16/D08 master and slave, (no MD32)
 - MBLT, BLT, ADOH, RMW, LOCK, location monitors,
- nine user programmable slave images on VMEbus and PCI bus ports,
- seven interrupt lines on either bus and flexible mapping of software and hardware sources of hardware interrupt,
- automatic initialization for slave-only applications,
- flexible register set, programmable from both the VMEbus and the PCI bus,
- four mailboxes and location monitor for message-oriented system,
- support for RMWs, lock cycles, and semaphores guarantee exclusive access,
- bus isolation mode for board maintenance, diagnostics, and live fault recovery,
- full VMEbus system controller functionality,
- several power-up options,
- IEEE 1149.1 JTAG testability support,
- commercial (0° to 70°C), industrial (-40° to 85°C) and extended temperature (-55° to 125°C) options, and
- available in 313-pin Plastic BGA and 324-pin Ceramic BGA.

1.2 Benefits of the Universe II

Interfacing the VMEbus with PCI presents a number of opportunities and challenges. The Universe II solves the problems and allows you to benefit from the opportunities.

The opportunities involve merging the best of the VMEbus and PCI bus worlds. The VMEbus is a proven standard specifically designed to support embedded systems. The distributed environment of the VMEbus supports multiprocessing and real-time intensive applications. A large number of off-the-shelf boards, software, and chassis components are available. VMEbus supports 21-slot systems without bridging and is continually evolving while providing backward compatibility. Hot swap solutions, and higher performance protocols have been defined and will be incorporated in future revisions of VMEbus.

Meanwhile, PCI has become a standard local bus. As a result, the leading semiconductor vendors have built PCI support into their newest processor and peripheral families.

The challenges involved in interfacing VMEbus to PCI include: address mapping, byte-lane swapping, and cycle mapping.

To allow VMEbus single-board computer vendors to benefit from PCI components, while overcoming the challenges involved in merging PCI with VMEbus, Tundra has developed a PCI to VMEbus interface controller, the Universe II. The Universe II is the industry proven, high performance 64-bit VMEbus to PCI interface, fully compliant with VME64 and tailored for the new generation of high performance PCI processors and peripherals.

The availability of the Universe II eases the development of multi-master, multi-processor architectures on VMEbus systems using PCI. The Universe II is ideally suited for CPU boards acting as both master and slave in the VMEbus system and that require access to PCI systems. With the Universe II, you know that as your system increases in complexity, you have silicon that continues to provide everything you need in a bridge. The elegant design of the Tundra Universe II, some of the best applications engineers in the industry, and this manual will make it as easy as possible for you to use the most sophisticated VMEbus interface.

1.3 Past and Future of the Universe

The Universe II (CA91C142) is a pin- and software-compatible revision of the Universe (CA91C042). The Universe was developed subsequently to the SCV64, Tundra's VME interface for non-PCI applications. The Universe II is the next generation of the Universe, and has been designed to exceed new customer expectations and to correct errata in the original Universe. The rich set of feature and performance enhancements are based on extensive consultation with our customers.

Given its history, the Universe II offers a low-risk, feature-rich, high performance solution for VMEbus-based PCI applications. Some of the performance enhancements offered by the Universe II include:

- improved PCI bus bandwidth utilization,
- improved register access performance,
- improved VMEbus slave and VME master performance,
- increased FIFO depth,
- DMA improved to optimize transfer rate on each bus,
- improved linked-list DMA performance,
- significantly improved coupled transfer performance,
- reduced power consumption.

Additional features include:

- four mailbox registers,
- eight semaphores,
- four location monitors,
- new software interrupts,
- more slave images.

The Universe II revision is another example of Tundra's commitment to supporting the VMEbus community. Tundra is actively participating in VMEbus, PCI bus, and CompactPCI bus standards and vendor associations, as well as related SIGs. Tundra will continue to propose and support enhancements to these specifications, while increasing both the range of options available to our customers and the compatibility between VME and PCI. Please visit our web site to keep abreast of these developments: <http://www.tundra.com>

1.4 About This Document

This manual is intended for users of the Tundra Universe II. Because of the differences between Universe and Universe II, this manual is not suitable for users of the original Tundra Universe. Universe users should continue to consult the previous VMEbus manual (document number 9000000.MD303.01).

The current manual is organized as follows:

- Chapter 1 - General Information,
- Chapter 2 - Functional Description,
- Chapter 3 - Signal Description,
- Chapter 4 - Signals and DC Characteristics,
- Appendix A - Registers,
- Appendix B- Performance,
- Appendix C- Typical Applications,
- Appendix D- Reliability,
- Appendix E- Cycle Mapping,
- Appendix F- Operating and Storage Conditions
- Appendix G- Mechanical and Ordering Information, and
- Index.

Chapter 1 introduces the Universe II and provides the reader with information about the necessary concepts and conventions required to use the manual.

Chapters 2 to 4 describe Universe II function, beginning with overall functionality in Chapter 2 to detailed signal descriptions in Chapters 3 and 4.

The Appendices are reference sources necessary for the implementation of the Universe II. In addition, the Appendices contain application information to aid the user in system design.

The Index provides a means to quickly access information on a keyword basis.

1.5 Universe II Technical Support

Tundra is dedicated to providing our customers with superior technical documentation and support. The following means of support are available:

1. **The Universe II User's Manual.** This is the main source of technical information. We strive hard to produce excellent documentation, and this manual contains the answers to most of our customers' questions.
2. **The Universe II Documentation Web Page.** This contains the latest manual, application notes, FAQ, articles, and any device errata and manual addenda. Please visit and bookmark <http://www.tundra.com>.
3. **The Designer's Resource Forum.** This is a public discussion forum at <http://www.tundra.com> which allows you to post questions and read threads pertaining specifically to the Universe II. Tundra technical support staff moderate this forum and promptly respond to customer inquiries.
4. **support@tundra.com.** You may also direct questions and feedback to Tundra using this e-mail address. Please include "Universe II" in the subject header of your message.
5. **Phone support.** Tundra's technical support staff may be reached at (613) 592-1320. Please ask for Universe II technical support.

1.6 Universe II Documentation

Tundra is dedicated to providing our customers with superior technical documentation and support.

1. **The Universe II User's Manual.** This is the main source of technical information.
2. **The Universe II Documentation Web Page.** This contains the latest manual, application notes, FAQ, articles, and any device errata and manual addenda. Please visit and bookmark <http://www.tundra.com>. Our documents are available as PDF files. These files are searchable and replete with hypertext links. You may read our documents on-line (with a PDF plug-in in your browser), or download them to your machine. The Acrobat indexes to these files allow you to perform very rapid searches through our documents.
3. **The Designer's Resource Center.** You can tailor how the Tundra web site is presented to you by using this web resource, available from <http://www.tundra.com>. There, you may also register to receive automatic e-

mail notification when the addendum, manual, or any other “resource” has changed. This is the best way to ensure that you always have the latest Universe II documentation.

1.7 Conventions

1.7.1 Signals

Signals on the VMEbus and PCI bus are either active high or active low. Active low signals are defined as true (asserted) when they are at logic low. Similarly, active high signals are defined as true at a logic high. Signals are considered asserted when active and negated when inactive, irrespective of voltage levels.

For voltage levels, the use of ‘0’ indicates a low voltage while a ‘1’ indicates a high voltage.

The names of Universe II signals on the VMEbus interface start with the letter ‘V’ (e.g., VRSYSRST#). Input signals on that interface start with ‘VR’ (e.g., VRSYSRST#), output signals start with ‘VX’ (e.g., VXBERR). Table 1.1 shows how the convention for denoting active low signals.

Table 1.1 : Suffixes for Active Low Signals

Suffix	Used for	Example
#	active low signals on PCI, and	RST#
	active low signals on VMEbus interface of the Universe II	VRSYSRST#
SIGNAL*	active low signals on the VMEbus backplane	SYSRESET*

1.7.2 Symbols



Caution: This symbol alerts the reader to procedures or operating levels which may result in misuse of or damage to the Universe II.



Note: This symbol directs the reader’s attention to useful information or suggestions.

1.7.3 Terminology

The term “cycle” refers to a single data beat, while a “transaction” is composed of one or more “cycles”.

To eliminate ambiguity, the expression “external master” is used to denote a master that is not the Universe II (or is a different Universe II). The capitalized expression “Master Interface” is used to denote the Universe II as bus master. The expression “external target” is used to denote a target that is not the Universe II (or is a different Universe II). The capitalized expression “Target Interface” or “Slave Interface” is used to denote the Universe II as target (or slave) of the bus. For example if the Universe II accesses a memory chip on the PCI bus, we might write: “The PCI Master Interface writes to the external PCI target.”

2 Functional Description

This chapter is organized as follows.

- Section 2.1, “Architectural Overview”, on page 2-2 summarizes the overall architecture of the Universe II,
- Section 2.2, “VMEbus Interface”, on page 2-6 presents the capabilities of the Universe II as VMEbus slave and VMEbus master,
- Section 2.3, “PCI Bus Interface”, on page 2-30 presents the capabilities of the Universe II as PCI bus target and PCI bus master,
- Section 2.4, “Slave Image Programming”, on page 2-50 explains how to program the mapping of cycles from one bus to the other,
- Section 2.5, “Bus Error Handling”, on page 2-58 explains how the Universe II handles and generates bus errors,
- Section 2.6, “Interrupt Generation”, on page 2-62 describes the interrupts which the Universe II can generate,
- Section 2.7, “Interrupt Handling”, on page 2-68 describes how the Universe II responds to interrupt conditions,
- Section 2.8, “DMA Controller”, on page 2-77 describes the operation of the Universe II’s Direct Memory Access Controller,
- Section 2.9, “Registers”, on page 2-100 gives an overview of the Universe II’s registers and how they can be accessed,
- Section 2.10, “Utility Functions”, on page 2-110 describes resets, power-up options, test modes, and clocks.

2.1 Architectural Overview

This section introduces the general architecture of the Universe II. This description makes frequent reference to the functional block diagram provided in Figure 2.1 on page 2-3. Notice that for each of the interfaces, VMEbus and PCI bus, there are three functionally distinct modules: master module, slave module, and interrupt module. These modules are connected to the different functional channels operating in the Universe II. These channels are:

- the VMEbus Slave Channel,
- the PCI Bus Target Channel,
- the DMA Channel,
- the Interrupt Channel, and
- the Register Channel.

The Architectural Overview is organized into the following sections:

- “VMEbus Interface”,
- “PCI Bus Interface”,
- “Interrupter and Interrupt Handler”, and
- “DMA Controller”.

These sections describe the operation of the Universe II in terms of the different modules and channels illustrated in Figure 2.1.

2.1.1 VMEbus Interface

2.1.1.1 Universe II as VMEbus Slave

The Universe II VMEbus Slave Channel accepts all of the addressing and data transfer modes documented in the VME64 specification (except A64 and those intended to augment 3U applications, i.e., A40 and MD32). Incoming write transactions from the VMEbus may be treated as either coupled or posted, depending upon the programming of the VMEbus slave image (see “VME Slave Images” on page 2-50). With posted write transactions, data is written to a Posted Write Receive FIFO (RXFIFO), and the VMEbus master receives data acknowledgment from the Universe II. Write data is transferred to the PCI resource from the RXFIFO without the involvement of the initiating VMEbus master (see “Posted Writes” on page 2-15 for a full explanation of this operation). With a coupled cycle, the VMEbus master only receives data acknowledgment when the transaction is complete on the PCI bus. This means that the VMEbus is unavailable to other masters while the PCI bus transaction is executed.

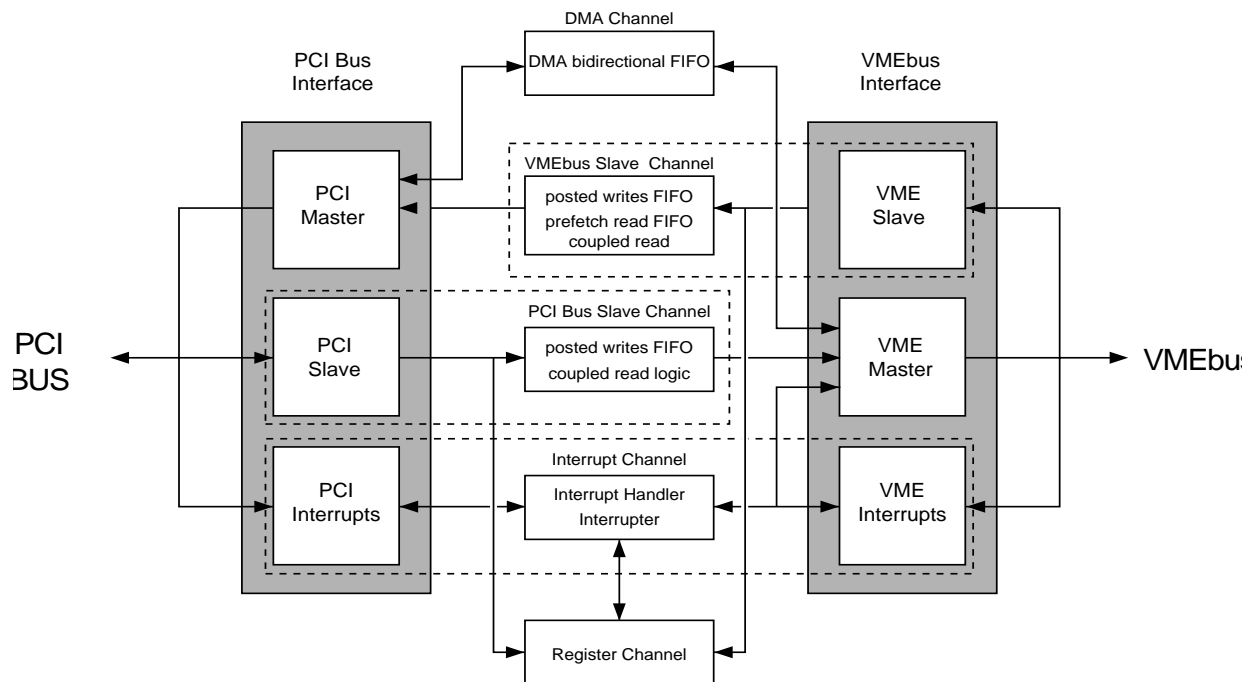


Figure 2.1 : Architectural Diagram for the Universe II

Read transactions may be either prefetched or coupled. If enabled by the user, a prefetched read is initiated when a VMEbus master requests a block read transaction (BLT or MBLT) and this mode is enabled. When the Universe II receives the block read request, it begins to fill its Read Data FIFO (RDFIFO) using burst transactions from the PCI resource. The initiating VMEbus master then acquires its block read data from the RDFIFO rather than from the PCI resources directly.

2.1.1.2 Universe II as VMEbus Master

The Universe II becomes VMEbus master when the VMEbus Master Interface is internally requested by the PCI Bus Target Channel, the DMA Channel, or the Interrupt Channel. The Interrupt Channel always has priority over the other two channels. Several mechanisms are available to configure the relative priority that the PCI Bus Target Channel and DMA Channel have over ownership of the VMEbus Master Interface.

The Universe II's VMEbus Master Interface generates all of the addressing and data transfer modes documented in the VME64 specification (except A64 and those intended to augment 3U applications, i.e. A40 and MD32). The Universe II is also compatible with all VMEbus modules conforming to pre-VME64 specifications. As VMEbus master, the Universe II supports Read-Modify-Write (RMW), and Address-Only-with-Handshake (ADOH) but does not accept RETRY* as a termination from the VMEbus slave. The ADOH cycle is used to implement the VMEbus Lock command allowing a PCI master to lock VMEbus resources.

2.1.2 PCI Bus Interface

2.1.2.1 Universe II as PCI Target

Read transactions from the PCI bus are always processed as coupled. Write transactions may be either coupled or posted, depending upon the setting of the PCI bus target image (see "PCI Bus Target Images" on page 2-53). With a posted write transaction, write data is written to a Posted Write Transmit FIFO (TXFIFO) and the PCI bus master receives data acknowledgment from the Universe II with zero wait-states. Meanwhile, the Universe II obtains the VMEbus and writes the data to the VMEbus resource independent of the initiating PCI master (see "Posted Writes" on page 2-44 for a full description of this operation).

To allow PCI masters to perform RMW and ADOH cycles, the Universe II provides a Special Cycle Generator. The Special Cycle Generator can be used in combination with a VMEbus ownership function to guarantee PCI masters exclusive access to VMEbus resources over several VMEbus transactions (see "The Special Cycle Generator" on page 2-45 and "Using the VOWN bit" on page 2-48 for a full description of this functionality).

2.1.2.2 Universe II as PCI Master

The Universe II becomes PCI master when the PCI Master Interface is internally requested by the VMEbus Slave Channel or the DMA Channel. There are mechanisms provided which allow the user to configure the relative priority of the VMEbus Slave Channel and the DMA Channel.

2.1.3 Interrupter and Interrupt Handler

2.1.3.1 Interrupter

The Universe II Interrupt Channel provides a flexible scheme to map interrupts to the PCI bus or VMEbus Interface. Interrupts are generated from hardware or software sources (see "Interrupt Generation" on page 2-62 and "Interrupt Handling" on page 2-68 for a full description of hardware and software sources). Interrupt sources can be mapped to any of the PCI bus or VMEbus interrupt output pins. Interrupt sources mapped to VMEbus interrupts are generated on the VMEbus interrupt output pins VIRQ# [7:1]. When a software and hardware source are assigned to the same VIRQn# pin, the software source always has higher priority.

Interrupt sources mapped to PCI bus interrupts are generated on one of the INT# [7:0] pins. To be fully PCI compliant, all interrupt sources must be routed to a single INT# pin.

For VMEbus interrupt outputs, the Universe II interrupter supplies an 8-bit STATUS/ID to a VMEbus interrupt handler during the IACK cycle, and optionally generates an internal interrupt to signal that the interrupt vector has been provided (see “VMEbus Interrupt Generation” on page 2-65).

Interrupts mapped to PCI bus outputs are serviced by the PCI interrupt controller. The CPU determines which interrupt sources are active by reading an interrupt status register in the Universe II. The source negates its interrupt when it has been serviced by the CPU (see “PCI Interrupt Generation” on page 2-63).

2.1.3.2 VMEbus Interrupt Handling

A VMEbus interrupt triggers the Universe II to generate a normal VMEbus IACK cycle and generate the specified interrupt output. When the IACK cycle is complete, the Universe II releases the VMEbus and the interrupt vector is read by the PCI resource servicing the interrupt output. Software interrupts are ROAK, while hardware, and internal interrupts are RORA.

2.1.4 DMA Controller

The Universe II provides an internal DMA controller for high performance data transfer between the PCI and VMEbus. DMA operations between the source and destination bus are decoupled through the use of a single bidirectional FIFO (DMAFIFO). Parameters for the DMA transfer are software configurable in the Universe II registers (see “DMA Controller” on page 2-77).

The principal mechanism for DMA transfers is the same for operations in either direction (PCI to VMEbus, or VMEbus to PCI), only the relative identity of the source and destination bus changes. In a DMA transfer, the Universe II gains control of the source bus and reads data into its DMAFIFO. Following specific rules of DMAFIFO operation (see “FIFO Operation and Bus Ownership” on page 2-92), it then acquires the destination bus and writes data from its DMAFIFO.

The DMA controller can be programmed to perform multiple blocks of transfers using entries in a linked-list. The DMA will work through the transfers in the linked-list following pointers at the end of each linked-list entry. Linked-list operation is initiated through a pointer in an internal Universe II register, but the linked-list itself resides in PCI bus memory.

2.2 VMEbus Interface

The VMEbus Interface incorporates all operations associated with the VMEbus. This includes master and slave functions, VMEbus configuration and system controller functions. These operations are covered as follows:

- “VMEbus Requester” below,
- “Universe II as VMEbus Master” on page 2-9,
- “Universe as VMEbus Slave” on page 2-13,
- “VMEbus Configuration” on page 2-23,
- “Automatic Slot Identification” on page 2-24, and
- “System Controller Functions” on page 2-26.
- “BI-Mode” on page 2-28.

For information concerning the Universe II VMEbus slave images, see “VME Slave Images” on page 2-50.

2.2.1 VMEbus Requester

2.2.1.1 Internal Arbitration for VMEbus Requests

Three different internal channels within the Universe II require use of the VMEbus: the Interrupt Channel, the PCI Target Channel, and the DMA Channel. These three channels do not directly request the VMEbus, instead they compete internally for ownership of the VMEbus Master Interface.

The Interrupt Channel (refer to Figure 2.1 on page 2-3) always has the highest priority for access to the VMEbus Master Interface. The DMA and PCI Target Channel requests are handled in a fair manner. The channel awarded VMEbus mastership maintains ownership of the VMEbus until it is ‘done’. The definition of ‘done’ for each channel is given below in “VMEbus Release” on page 2-8.

The Interrupt Channel requests the VMEbus master when it detects an enabled VMEbus interrupt line asserted and needs to run an interrupt acknowledge cycle to acquire the STATUS/ID.

The PCI Target Channel requests the VMEbus Master Interface to service the following conditions:

- the TXFIFO contains a complete transaction, or
- if there is a coupled cycle request.

The DMA Channel requests the VMEbus Master Interface if:

- the DMAFIFO has 64 bytes available (if it is reading from the VMEbus) or 64 bytes in its FIFO (if it is writing to the VMEbus), or
- the DMA block is complete (see “DMA Controller” on page 2-77).

In the case of the DMA Channel, the user can optionally use the DMA Channel VMEbus-off-timer to further qualify requests from this channel. The VMEbus-off-timer controls how long the DMA remains off the VMEbus before making another request (see “PCI to VMEbus Transfers” on page 2-92).

The Universe II provides a software mechanism for VMEbus acquisition through the VMEbus ownership bit (VOWN in the MAST_CTL register, Table A.81). When the VMEbus ownership bit is set, the Universe II acquires the VMEbus and sets an acknowledgment bit (VOWN_ACK in the MAST_CTL register, Table A.81) and optionally generates an interrupt to the PCI bus (see “VME Lock Cycles—Exclusive Access to VMEbus Resources” on page 2-47). The Universe II maintains VMEbus ownership until the ownership bit is cleared. During the VMEbus tenure initiated by setting the ownership bit, only the PCI Target Channel and Interrupt Channel can access the VMEbus Master Interface.

2.2.1.2 Request Modes

Request Levels

The Universe II is software configurable to request on all VMEbus request levels: BR3*, BR2*, BR1*, and BR0*. The default setting is for level 3 VMEbus request. The request level is a global programming option set through the VRL field in the MAST_CTL register (Table A.81). The programmed request level is used by the VMEbus Master Interface regardless of the channel (Interrupt Channel, DMA Channel, or PCI Target Channel) currently accessing the VMEbus Master Interface.

Fair and Demand

The Universe II requester may be programmed for either Fair or Demand mode. The request mode is a global programming option set through the VRM bits in the MAST_CTL register (Table A.81).

In Fair mode, the Universe II does not request the VMEbus until there are no other VMEbus requests pending at its programmed level. This mode ensures that every requester on an equal level has access to the bus.

In Demand mode (the default setting), the requester asserts its bus request regardless of the state of the BRn* line. By requesting the bus frequently, requesters far down the daisy chain may be prevented from ever obtaining bus ownership. This is referred to as “starving” those requesters. Note that in order to achieve fairness, all bus requesters in a VMEbus system must be set to Fair mode.

2.2.1.3 VMEbus Release

The Universe II VMEbus requester can be configured as either RWD (release when done) or ROR (release on request) using the VREL bit in the MAST_CTL register (Table A.81). The default setting is for RWD. ROR means the Universe II releases BBSY* only if a bus request is pending from another VMEbus master and once the channel that is the current owner of the VMEbus Master Interface is done. Ownership of the bus may be assumed by another channel without re-arbitration on the bus if there are no pending requests on any level on the VMEbus. When set for RWD, the VMEbus Master Interface releases BBSY* when the channel accessing the VMEbus Master Interface is done (see below). Note that the MYBBSY status bit in the MISC_STAT register (Table A.83) is set while the Universe II asserts the BBSY* output.

In RWD mode, the VMEbus is released when the channel (for example, the DMA Channel) is done, even if another channel has a request pending (for example, the PCI Target Channel). A re-arbitration of the VMEbus is required for any pending channel requests. Each channel has a set of rules that determine when it is ‘done’ with its VMEbus transaction.

The Interrupt Channel is done when a single interrupt acknowledge cycle is complete.

The PCI Target Channel is done under the following conditions:

- when the TXFIFO is empty (the TXFE bit is set by the Universe II in the MISC_STAT register, Table A.83),
- when the maximum number of bytes per PCI Target Channel tenure has been reached (as programmed with the PWON field in the MAST_CTL register, Table A.81)¹,
- after each posted write, if the PWON is equal to 0b1111, as programmed in the MAST_CTL register, Table A.81
- when the coupled cycle is complete and the Coupled Window Timer has expired,
- if the Coupled Request Timer (page 2-42) expires before a coupled cycle is retried by a PCI master, or
- when VMEbus ownership is acquired with the VOWN bit in the MAST_CTL register and then the VOWN bit is cleared (in other words, if the VMEbus is acquired through the use of the VOWN bit, the Universe II does not release BBSY* until the VOWN bit is cleared—see “VME Lock Cycles—Exclusive Access to VMEbus Resources” on page 2-47).

1. This setting is overridden if the VOWN mechanism is used.

The DMA Channel is done under the following conditions (see “FIFO Operation and Bus Ownership” on page 2-92 and “DMA Error Handling” on page 2-96):

- DMAFIFO full during VMEbus to PCI bus transfers,
- DMAFIFO empty during PCI bus to VMEbus transfers,
- if an error is encountered during the DMA operation,
- the DMA VMEbus Tenure Byte Counter has expired, or
- DMA block is complete.

The Universe II does not monitor BCLR* and so its ownership of the VMEbus is not affected by the assertion of BCLR*.

2.2.2 Universe II as VMEbus Master

The Universe II becomes VMEbus master as a result of the following chain of events:

1. a PCI master accesses a Universe II PCI target image (leading to VMEbus access) or the DMA Channel initiates a transaction,
2. either the Universe II PCI Target Channel or the DMA Channel wins access to the VMEbus Master Interface through internal arbitration, and
3. the Universe II Master Interface requests and obtains ownership of the VMEbus.

The Universe II will also become VMEbus master if the VMEbus ownership bit is set (see “VME Lock Cycles—Exclusive Access to VMEbus Resources” on page 2-47) and in its role in VMEbus interrupt handling (see “VMEbus Interrupt Handling” on page 2-68).

The following sections describe the function of the Universe II as a VMEbus master in terms of the different phases of a VMEbus transaction: addressing, data transfer, cycle termination, and bus release.

2.2.2.1 Addressing Capabilities

Depending upon the programming of the PCI target image (see “PCI Bus Target Images” on page 2-53), the Universe II generates A16, A24, A32, and CR/CSR address phases on the VMEbus. The address mode and type (supervisor/non-privileged and program/data) are also programmed through the PCI target image. Address pipelining is provided except during MBLT cycles, where the VMEbus specification does not permit it.

The address and AM codes that are generated by the Universe II are functions of the PCI address and PCI target image programming (see “PCI Bus Target Images” on page 2-53) or through DMA programming.

The Universe II generates Address-Only-with-Handshake (ADOH) cycles in support of lock commands for A16, A24, and A32 spaces. ADOH cycles must be generated through the Special Cycle Generator (see “The Special Cycle Generator” on page 2-45).

There are two User Defined AM codes that can be programmed through the USER_AM register (Table A.84). The USER_AM register can only be used to generate and accept AM codes 0x10 through 0x1F. These AM codes are designated as USERAM codes in the VMEbus specification. After power-up, the two values in the USER_AM register default to the same VME64 user-defined AM code.

If USER_AM codes are used with the VMEbus Slave Interface, ensure that the cycles use 32-bit addressing, and that only single cycle accesses are used. BLTs and MBLTs with USER_AM codes will lead to unpredictable behavior.

2.2.2.2 Data Transfer Capabilities

The data transfer between the PCI bus and VMEbus is depicted in Figure 2.2 on page 2-12. The Universe II can be seen as a funnel where the mouth of the funnel is the data width of the PCI transaction. The end of the funnel is the maximum VMEbus data width programmed into the PCI target image. For example, consider a 32-bit PCI transaction accessing a PCI target image with VDW set to 16 bits. A data beat with all byte lanes enabled will be broken into two 16-bit cycles on the VMEbus. If the PCI target image is also programmed with block transfers enabled, the 32-bit PCI data beat will result in a D16 block transfer on the VMEbus. Write data is unpacked to the VMEbus and read data is packed to the PCI bus data width.

If the data width of the PCI data beat is the same as the maximum data width of the PCI target image, then the Universe II maps the data beat to an equivalent VMEbus cycle. For example, consider a 32-bit PCI transaction accessing a PCI target image with VDW set to 32 bits. A data beat with all byte lanes enabled is translated to a single 32-bit cycle on the VMEbus.

As the general rule, if the PCI bus data width is less than the VMEbus data width then there is no packing or unpacking between the two buses. The only exception to this is during 32-bit PCI multi-data beat transactions to a PCI target image programmed with maximum VMEbus data width of 64 bits. In this case, packing/unpacking occurs to make maximum use of the full bandwidth on both buses.

Only aligned VMEbus transactions are generated, so if the requested PCI data beat has unaligned or non-contiguous byte enables, then it is broken into multiple aligned VMEbus transactions no wider than the programmed VMEbus data width. For example, consider a three-byte PCI data beat (on a 32-bit PCI bus) accessing a PCI target image with VDW set to 16 bits. The three-byte PCI data beat will be broken into two aligned VMEbus cycles: a single-byte cycle and a double-byte cycle (the ordering of the two cycles depends on the arrangement of the byte enables in the PCI data beat). If in the above example the PCI target image has a VDW set to 8 bits, then the three-byte PCI data beat will be broken into three single-byte VMEbus cycles.

BLT/MBLT cycles are initiated on the VMEbus if the PCI target image has been programmed with this capacity (see “PCI Bus Target Images” on page 2-53). The length of the BLT/MBLT transactions on the VMEbus will be determined by the initiating PCI transaction or the setting of the PWON field in the MAST_CTL register (Table A.81). For example, a single data beat PCI transaction queued in the TXFIFO results in a single data beat block transfer on the VMEbus. With the PWON field, the user can specify a transfer byte count that will be dequeued from the TXFIFO before the VMEbus Master Interface relinquishes the VMEbus. The PWON field specifies the minimum tenure of the Universe II on the VMEbus. However, tenure is extended if the VOWN bit in the MAST_CTL register is set (see “Using the VOWN bit” on page 2-48).

During DMA operations, the Universe II will attempt block transfers to the maximum length permitted by the VMEbus specification (256 bytes for BLT, 2 Kbytes for MBLT) and as limited by the VON counter (see “DMA VMEbus Ownership” on page 2-81).

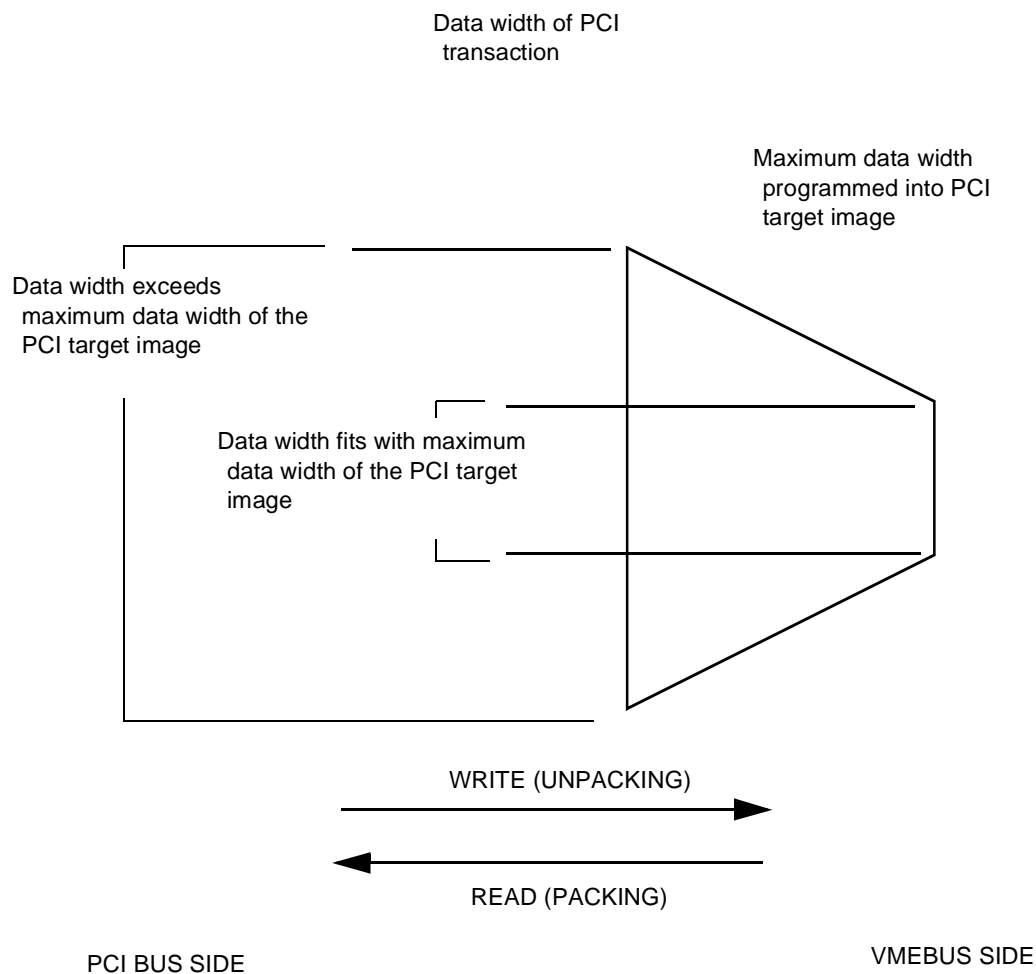


Figure 2.2 : Influence of Transaction Data Width and Target Image Data Width on Data Packing/Unpacking

The Universe II provides indivisible transactions with the VMEbus lock commands and the VMEbus ownership bit (see “VME Lock Cycles—Exclusive Access to VMEbus Resources” on page 2-47).

2.2.2.3 Cycle Terminations

The Universe II accepts BERR* or DTACK* as cycle terminations from the VMEbus slave. It does not support RETRY*. The assertion of BERR* indicates that some type of system error occurred and the transaction did not complete properly. A VMEbus BERR* received by the Universe II during a coupled transaction is communicated to the PCI master as a Target-Abort. No information is logged if the Universe II receives BERR* in a coupled transaction. If an error occurs during a posted write to the VMEbus, the Universe II uses the V_AMERR register (Table A.107) to log the AM code of the transaction (AMERR [5:0]), and the state of the IACK* signal (IACK bit, to indicate whether the error occurred during an IACK cycle). The current transaction in the FIFO is purged. The V_AMERR register also records if multiple errors have occurred (with the M_ERR bit), although the actual number of errors is not given. The error log is qualified by the value of the V_STAT bit. The address of the errored transaction is latched in the V_AERR register (Table A.108). When the Universe II receives a VMEbus error during a posted write, it generates an interrupt on the VMEbus and/or PCI bus depending upon whether the VERR and LERR interrupts are enabled (see “Interrupt Handling” on page 2-68, Table A.60 and Table A.61).

DTACK* signals the successful completion of the transaction.

2.2.3 Universe as VMEbus Slave

This section describes the VMEbus Slave Channel and other aspects of the Universe II as VMEbus slave. The following topics are discussed:

- “Coupled Transfers” on page 2-14,
- “Posted Writes” on page 2-15,
- “Prefetched Block Reads” on page 2-16,
- “VMEbus Lock Commands (ADOH Cycles)” on page 2-18,
- “VMEbus Read-Modify-Write Cycles (RMW Cycles)” on page 2-19,
- “Location Monitors” on page 2-19 and
- “Generating PCI Configuration Cycles” on page 2-20.

The Universe II becomes VMEbus slave when one of its eight programmed slave images or register images are accessed by a VMEbus master (note that the Universe II cannot reflect a cycle on the VMEbus and access itself). Depending upon the programming of the slave image, different possible transaction types result (see “VME Slave Images” on page 2-50 for a description of the types of accesses to which the Universe II responds).

For reads, the transaction can be coupled or prefetched. Similarly, write transactions can be coupled or posted. The type of read or write transaction allowed by the slave image depends on the programming of that particular VMEbus slave image (see Figure 2.3 below and “VME Slave Images” on page 2-50). To ensure sequential consistency, prefetched reads, coupled reads, and coupled write operations are only processed once all previously posted write operations have completed (i.e. the RXFIFO is empty).

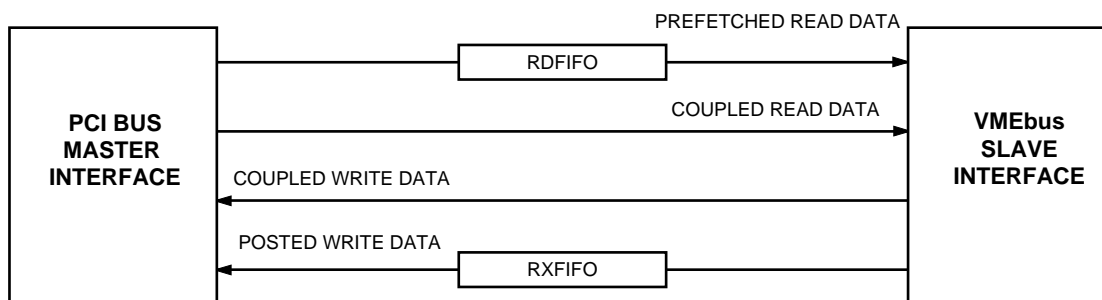


Figure 2.3 : VMEbus Slave Channel Dataflow

Incoming cycles from the VMEbus can have data widths of 8-bit, 16-bit, 32-bit, and 64-bit. Although the PCI bus supports only two port sizes (32-bit and 64-bit), the byte lanes on the PCI bus can be individually enabled, which allows each type of VMEbus transaction to be directly mapped to the PCI data bus.



In order for a VMEbus slave image to respond to an incoming cycle, the PCI Master Interface must be enabled (bit BM in the PCI_CSR register; Table A.3). If data is enqueued in the VMEbus Slave Channel FIFO and the PCI BM bit is cleared, the FIFO will empty but no additional transfers will be received.

2.2.3.1 Coupled Transfers

A coupled transfer means that no FIFO is involved in the transaction and handshakes are relayed directly through the Universe II. Coupled mode is the default setting for the VMEbus slave images. Coupled transfers only proceed once all posted write entries in the RXFIFO have completed (see “Posted Writes” below).

A coupled cycle with multiple data beats (i.e. block transfers) on the VMEbus side is always mapped to single data beat transactions on the PCI bus, where each data beat on the VMEbus is mapped to a single data beat transaction on the PCI bus regardless of data beat size. No packing or unpacking is performed. The only exception to this is when a D64 VMEbus transaction is mapped to D32 on the PCI bus. The data width of the PCI bus depends on the

programming of the VMEbus slave image (32-bit or 64-bit, see “VME Slave Images” on page 2-50). The Universe II enables the appropriate byte lanes on the PCI bus as required by the VMEbus transaction. For example, a VMEbus slave image programmed to generate 32-bit transactions on the PCI bus is accessed by a VMEbus D08 BLT read transaction (prefetching is not enabled in this slave image). The transaction is mapped to single data beat 32-bit transfers on the PCI bus with only one byte lane enabled.

Target-Retry from a PCI target is not communicated to the VMEbus master. PCI transactions terminated with Target-Abort or Master-Abort are terminated on the VMEbus with BERR*. Note that the Universe II sets the R_TA or R_MA bits in the PCI_CS register (Table A.3) when it receives a Target-Abort or Master-Abort.

2.2.3.2 Posted Writes

A posted write involves the VMEbus master writing data into the Universe II's RXFIFO, rather than directly to the PCI address. Write transactions from the VMEbus are processed as posted if the PWEN bit is set in the VMEbus slave image control register (see “VME Slave Images” on page 2-50). If the bit is cleared (the default setting) the transaction bypasses the FIFO and is performed as a coupled transfer (see above). Incoming posted writes from the VMEbus are queued in the 32-entry deep RXFIFO. (The RXFIFO is the same structure as the RDFIFO. The different names are used for the FIFO's two roles, only one of which it can implement at once.) Each entry in the RXFIFO can contain 64 address bits, or 64 data bits. Each incoming VMEbus address phase, whether it is 16-bit, 24-bit, or 32-bit, constitutes a single entry in the RXFIFO and is followed by subsequent data entries. The address entry contains the translated PCI address space and command information mapping relevant to the particular VMEbus slave image that has been accessed (see “VME Slave Images” on page 2-50). For this reason, any re-programming of VMEbus slave image attributes will only be reflected in RXFIFO entries queued after the re-programming. Transactions queued before the re-programming are delivered to the PCI bus with the VMEbus slave image attributes that were in use before the re-programming.

Incoming non-block write transactions from the VMEbus require two entries in the RXFIFO: one address entry (with accompanying command information) and one data entry. The size of the data entry corresponds to the data width of the VMEbus transfer. Block transfers require at least two entries: one entry for address and command information, and one or more data entries. The VMEbus Slave Channel packs data received during block transfers to the full 64-bit width of the RXFIFO. For example, a ten data phase D16 BLT transfer (20 bytes in total) does not require ten data entries in the RXFIFO. Instead, eight of the ten data phases (16 bits per data phase for a total of 128 bits) are packed into two 64-bit data entries in the RXFIFO. The final two data phases (32 bits combined) are queued in the next RXFIFO entry. When you add the address entry to the three data entries, this VMEbus block write has been stored in a total of four RXFIFO entries.

Unlike the PCI Target Channel (see page 2-38), the VMEbus Slave Channel does not retry the VMEbus if the RXFIFO does not have enough space to hold an incoming VMEbus write transaction. Instead, the DTACK* response from the VMEbus Slave Interface is delayed until space becomes available in the RXFIFO. Since single transfers require two entries in the RXFIFO, two entries must be freed up before the VMEbus Slave Interface asserts DTACK*. Similarly, the VMEbus Slave Channel requires two available RXFIFO entries before it can acknowledge the first data phase of a BLT or MBLT transfer (one entry for the address phase and one for the first data phase). If the RXFIFO has no available space for subsequent data phases in the block transfer, then the VMEbus Slave Interface delays assertion of DTACK* until a single entry is available for the next data phase in the block transfer.

The PCI Master Interface uses transactions queued in the RXFIFO to generate transactions on the PCI bus. No address phase deletion is performed, so the length of a transaction on the PCI bus corresponds to the length of the queued VMEbus transaction. Non-block transfers are generated on the PCI bus as single data beat transactions. Block transfers are generated as one or more burst transactions, where the length of the burst transaction is programmed by the (PABS field in the MAST_CTL register, Table A.81).

The Universe II always packs or unpacks data from the VMEbus transaction to the PCI bus data width programmed into the VMEbus slave image (with all PCI bus byte lanes enabled). For example, consider a VMEbus slave image programmed for posted writes and a D32 PCI bus that is accessed with a VMEbus D16 block write transaction. The VMEbus D16 write transaction is mapped to a D32 write transaction on the PCI bus with all byte lanes enabled. (However, note that a single D16 transaction from the VMEbus is mapped to the PCI bus as D32 with only two byte lanes enabled).

During block transfers, the Universe II will pack data to the full negotiated width of the PCI bus. This may imply that for block transfers that begin or end on addresses not aligned to the PCI bus width different byte lanes may be enabled during each data beat.

If an error occurs during a posted write to the PCI bus, the Universe II uses the L_CMDERR register (Table A.32) to log the command information for the transaction (CMDERR [3:0]). The L_CMDERR register also records if multiple errors have occurred (with the M_ERR bit) although the actual number is not given. The error log is qualified with the L_STAT bit. The address of the errored transaction is latched in the LAERR register (Table A.33). An interrupt is generated on the VMEbus and/or PCI bus depending upon whether the VERR and LERR interrupts are enabled (see “Bus Error Handling” on page 2-58 and “Interrupt Handling” on page 2-68).

2.2.3.3 Prefetched Block Reads

Prefetching of read data occurs for VMEbus block transfers (BLT, MBLT) in those slave images that have the prefetch enable (PREN) bit set (see “VME Slave Images” on page 2-50). In the VMEbus Slave Channel, prefetching is not supported for non BLT/MBLT transfers.

Without prefetching, block read transactions from a VMEbus master are handled by the VMEbus Slave Channel as coupled reads. This means that each data phase of the block transfer is translated to a single data beat transaction on the PCI bus. In addition, only the amount of data requested during the relevant data phase is fetched from the PCI bus. For example, a D16 block read transaction with 32 data phases on the VMEbus maps to 32 PCI bus transactions, where each PCI bus transaction has only two byte lanes enabled. Note the VMEbus lies idle during the arbitration time required for each PCI bus transaction, resulting in a considerable performance degradation.

With prefetching enabled, the VMEbus Slave Channel uses a 32-entry deep RDFIFO to provide read data to the VMEbus with minimum latency. (The RXFIFO is the same structure as the RDFIFO. The different names are used for the FIFO's two roles, only one of which it can implement at once.) The RDFIFO is 64 bits wide, with additional bits for control information. If a VMEbus slave image is programmed for prefetching (see “VME Slave Images” on page 2-50), then a block read access to that image causes the VMEbus Slave Channel to generate aligned burst read transactions on the PCI bus (the size of the burst read transactions is determined by the setting of the aligned burst size, PABS in the MAST_CTL register). These PCI burst read transaction are queued in the RDFIFO and the data is then delivered to the VMEbus. Note that the first data phase provided to the VMEbus master is essentially a coupled read, but subsequent data phases in the VMEbus block read are delivered from the RDFIFO and are essentially decoupled (see “Prefetched Reads” on page 2-60 for the impact on bus error handling).

The data width of the transaction on the PCI bus (32-bit or 64-bit) depends on the setting of the LD64EN bit in the VMEbus slave image control register (e.g. see Table A.85) and the capabilities of the accessed PCI target. Internally, the prefetched read data is packed to 64 bits, regardless of the width of the PCI bus or the data width of the original VMEbus block read (no address information is stored with the data). Once one entry is queued in the RDFIFO, the VMEbus Slave Interface delivers the data to the VMEbus, unpacking the data as necessary to fit with the data width of the original VMEbus block read (e.g. D16, or D32). The VMEbus Slave Interface continuously delivers data from the RDFIFO to the VMEbus master performing the block read transaction. Because PCI bus data transfer rates exceed those of the VMEbus, it is unlikely that the RDFIFO will ever be unable to deliver data to the VMEbus master. For this reason, block read performance on the VMEbus will be similar to that observed with block writes. However, should the RDFIFO be unable to deliver data to the VMEbus master (which may happen if there is considerable traffic on the PCI bus or the PCI bus target has a slow response) the VMEbus Slave Interface delays DTACK* assertion until an entry is queued and is available for the VMEbus block read.

On the PCI side, prefetching continues as long as there is room for another transaction in the RDFIFO and the initiating VMEbus block read is still active. The space required in the RDFIFO for another PCI burst read transaction is determined by the setting of the PCI aligned burst size (PABS in the MAST_CTL register, Table A.81). If PABS is set for 32 bytes, there must be four entries available in the RDFIFO; for aligned burst size set to 64 bytes, eight

entries must be available, for aligned burst size set to 128 bytes, there must be 16 entries available. When there is insufficient room in the RDFIFO to hold another PCI burst read, the read transactions on the PCI bus are terminated and only resume if room becomes available for another aligned burst AND the original VMEbus block read is still active. When the VMEbus block transfer terminates, any remaining data in the RDFIFO is purged.

Reading on the PCI side will not cross a 1024-byte boundary. The PCI Master Interface will release FRAME# and the VMEbus Slave Channel will relinquish internal ownership of the PCI Master Interface when it reaches this boundary. The VMEbus Slave Channel will re-request internal ownership of the PCI Master Interface as soon as possible, in order to continue reading from the external PCI target. (As described elsewhere, the PABS setting determines how much data must be available in the RDFIFO before the VMEbus Slave Channel continues reading.)

Regardless of the read request, the data width of prefetching on the PCI side is full width with all byte lanes enabled. If the request is unaligned, then the first PCI data beat will have only the relevant byte lanes enabled. Subsequent data beats will have full data width with all byte lanes enabled. If LD64EN is set in the VMEbus Slave image, the Universe II requests D64 on the PCI bus by asserting REQ64# during the address phase. If the PCI target does not respond with ACK64#, subsequent data beats are D32.

If an error occurs on the PCI bus, the Universe II does not translate the error condition into a BERR* on the VMEbus. Indeed, the Universe II does not directly map the error. By doing nothing, the Universe II forces the external VMEbus error timer to expire.

2.2.3.4 VMEbus Lock Commands (ADOH Cycles)

The Universe II supports VMEbus lock commands as described in the VME64 specification. Under the specification, ADOH cycles are used to execute the lock command (with a special AM code). Any resource locked on the VMEbus cannot be accessed by any other resource during the bus tenure of the VMEbus master.

When the Universe II receives a VMEbus lock command, it asserts LOCK# to the addressed resource on the PCI bus. The PCI Master Interface processes this as a read transfer (with no data). All subsequent slave VMEbus transactions are coupled while the Universe II owns PCI LOCK#. Note that the VMEbus Slave Channel has dedicated access to the PCI Master Interface during the locked transaction. The Universe II holds the PCI bus lock until the VMEbus lock command is terminated, i.e. when BBSY* is negated.

The Universe II accepts ADOH cycles in any of the slave images when the Universe II PCI Master Interface is enabled (BM bit in PCI_CSR register) and the images are programmed to map transactions into PCI Memory Space.

In the event that a Target-Abort or a Master-Abort occurs during a locked transaction on the PCI bus, the Universe II will relinquish its ownership of LOCK# in accord with the PCI bus Specification. It is the responsibility of the user to verify the R_MA and R_TA status bits of the PCI_CSR status register to determine whether or not ownership of LOCK# was lost.

Once an external VMEbus masters locks the PCI bus, the Universe II DMA will not perform transfers on the PCI bus until the bus is unlocked.

2.2.3.5 VMEbus Read-Modify-Write Cycles (RMW Cycles)

A read-modify-write (RMW) cycle allows a VMEbus master to read from a VMEbus slave and then write to the same resource without relinquishing bus tenure between the two operations. Each of the Universe II slave images can be programmed to map RMW transactions to PCI locked transactions. If the LLRMW enable bit is set in the selected VMEbus slave image control register (e.g. Table A.85), then every non-block slave read is mapped to a coupled PCI locked read. LOCK# will be held on the PCI bus until AS* is negated on the VMEbus. Every non-block slave read is assumed to be a RMW since there is no possible indication from the VMEbus master that the single cycle read is just a read or the beginning of a RMW.

If the LLRMW enable bit is not set and the Universe II receives a VMEbus RMW cycle, the read and write portions of the cycle will be treated as independent transactions on the PCI bus: i.e., a read followed by a write. The write may be coupled or decoupled depending on the state of the PWEN bit in the accessed slave image.



Note: There may be an adverse performance impact for reads that are processed through a RMW-capable slave image; this may be accentuated if LOCK# is currently owned by another PCI master.

RMW cycles are not supported with unaligned or D24 cycles.

When an external VMEbus Master begins a RMW cycle, at some point a read cycle will appear on the PCI bus. During the time between when the read cycle occurs on the PCI bus and when the associated write cycle occurs on the PCI bus, no DMA transfers will occur on the PCI bus.

2.2.3.6 Register Accesses

See “Registers” on page 2-100 for a full description of register mapping and register access.

2.2.3.7 Location Monitors

Universe II has four location monitors to support a VMEbus broadcast capability. The location monitors’ image is a 4-Kbyte image in A16, A24 or A32 space on the VMEbus. If enabled, an access to a location monitor causes the PCI Master Interface to generate an interrupt.

The Location Monitor Control Register (LM_CTL, Table A.101) controls the Universe II's location monitoring. The EN field of the LM_CTL register enables the capability. The PGM[1:0] field sets the Program/Data AM code. The SUPER[1:0] field of the LM_CTL register sets the Supervisor/User AM code to which the Universe II responds. The VAS[3:0] field of the LM_CTL register specifies the address space that is monitored. The BS[31:12] field of the location monitor Base Address Register (LM_BS, Table A.102) specifies the lowest address in the 4 Kbyte range that will be decoded as a location monitor access. While the Universe II is said to have four location monitors, they all share the same LM_CTL and LM_BS registers.

In address spaces A24 and A16, the respective upper address bits are ignored.

When an access to a location monitor is detected, an interrupt is generated on the PCI bus. VMEbus address bits [4:3] determine which Location Monitor will be used, and hence which of four PCI interrupts to generate. (See "Location Monitors" on page 2-75 for details on mapping the interrupts from the location monitor.)

The location monitors do not store write data. Read data from the location monitors is undefined. Location monitors do not support BLT or MBLT transfers.

Each Universe II on the VMEbus should be programmed to monitor the same 4 Kbytes of addresses on the VMEbus. Note that the Universe II may access its own location monitor. If the Universe II accesses its own (enabled) location monitor, the same Universe II generates DTACK* on the VMEbus and thereby terminates its own cycle. This removes the necessity of the system integrator ensuring that there is another card enabled to generate DTACK*. The generation of DTACK* happens after the Universe II has decoded and responded to the cycle. If the location monitor is accessed by a different master, the Universe II does not respond with DTACK*.

2.2.3.8 Generating PCI Configuration Cycles

PCI Configuration cycles can be generated by accessing a VMEbus slave image whose Local Address Space field (LAS) is set for Configuration Space.

Both Type 0 and Type 1 cycles are generated and handled through the same mechanism. Once a VMEbus cycle is received and mapped to a configuration cycle, the Universe II compares bits [23:16] of the incoming address with the value stored in the MAST_CTL Register's Bus Number field (BUS_NO[7:0] in Table A.81). If the bits are the same as the BUS_NO field, then a TYPE 0 access is generated. If they are not the same, a Type 1 configuration access is generated. The PCI bus-generated address then becomes an unsigned addition of the incoming VMEbus address and the VMEbus slave image translation offset.

Generating Configuration Type 0 Cycles

The Universe II asserts one of AD[31:11] on the PCI bus to select a device during a configuration Type 0 access. To perform a configuration Type 0 cycle on the PCI bus:

- Program the LAS field of VSIx_CTL for Configuration Space,
- Program the VSIx_BS, VSIx_BD registers to some suitable value,
- Program the VSIx_TO register to 0, and
- Program the BUS_NO field of the MAST_CTL register to some value.

Perform a VMEbus access where:

- VA[7:2] identifies the PCI Register Number and will be mapped directly to AD[7:2],
- VA[10:8] identifies the PCI Function Number and will be mapped directly to AD[10:8],
- VA[15:11] selects the device on the PCI bus and will be mapped to AD[31:12] according to Table 2.2,
- VA[23:16] matches the BUS_NO in MAST_CTL register, and
- Other address bits are irrelevant—they are not mapped to the PCI bus.

Table 2.2 : PCI Address Line Asserted as a Function of VA[15:11]

VA[15:11] ^a	PCI Address Line Asserted ^b
00000	11
00001	12
00010	13
00011	14
00100	15
00101	16
00110	17
00111	18
01000	19
01001	20
01010	21
01011	22
01100	23

Table 2.2 : PCI Address Line Asserted as a Function of VA[15:11]

VA[15:11] ^a	PCI Address Line Asserted ^b
01101	24
01110	25
01111	26
10000	27
10001	28
10010	29
10011	30
10100	31

a. The other values of VA[15:11] are not defined and must not be used.

b. Only one of AD[31:11] is asserted; the other address lines in AD[31:11] are negated.

Generating Configuration Type 1 Cycles

To generate a configuration Type 1 cycle on the VMEbus:

- Program LAS field of VSIx_CTL to Configuration Space,
- Program the VSIx_BS, VSIx_BD registers to some suitable value,
- Program the VSIx_TO register to 0 and
- Program the BUS_NO field of the MAST_CTL register to some value.

Perform a VMEbus access where:

- VMEbus Address[7:2] identifies the PCI Register Number,
- VMEbus Address[10:8] identifies the PCI Function Number,
- VMEbus Address[15:11] identifies the PCI Device Number,
- VMEbus Address[23:16] does not match the BUS_NO in MAST_CTL register, and
- VMEbus Address[31:24] are mapped directly through to the PCI bus.

2.2.4 VMEbus Configuration

The Universe II provides the following functions to assist in the initial configuration of the VMEbus system:

- First Slot Detector,
- Register Access at Power-up, and
- Auto Slot ID (two methods).

These are described separately below.

2.2.4.1 First Slot Detector

As specified by the VME64 specification the First Slot Detector module on the Universe II samples BG3IN* immediately after reset to determine whether the Universe II's host board resides in slot 1. The VMEbus specification requires that BG[3:0]* lines be driven high after reset. This means that if a card is preceded by another card in the VMEbus system, it will always sample BG3IN* high after reset. BG3IN* can only be sampled low after reset by the first card in the system (there is no preceding card to drive BG3IN* high). If BG3IN* is sampled at logic low immediately after reset (due to the Universe II's internal pull-down), then the Universe II's host board is in slot 1 and the Universe II becomes SYSCON; otherwise, the SYSCON module is disabled. This mechanism may be overridden by software through clearing or setting the SYSCON bit in the MISC_CTL register (Table A.82).

The Universe II monitors IACK* (rather than IACKIN*) when it is configured as SYSCON. This permits it to operate as SYSCON in a VMEbus chassis slot other than slot 1, provided there are only empty slots to its left. The slot with SYSCON in it becomes a virtual slot 1.

2.2.4.2 VMEbus Register Access at Power-up

The Universe II provides a VMEbus slave image that allows access to all Universe II Control and Status Registers (UCSR). The base address for this slave image is programmed through the VRAI_BS register (Table A.104). At power-up, the Universe II can program the VRAI_BS and VRAI_CTL (Table A.103) registers with information specifying the UCSR slave image (see "Power-Up Options" on page 2-115).

Register access at power-up would be used in systems where the Universe II's card has no CPU, or where register access for that card needs to be independent of the local CPU.

2.2.5 Automatic Slot Identification

The Universe II supports two types of Auto-ID functionality. One type uses the Auto Slot ID technique as described in the VME64 specification. The other type uses a proprietary method developed by DY4 Systems and implemented in the Tundra SCV64. Neither system identifies geographical addressing, only the relative position amongst the boards present in the system (i.e. fourth board versus fourth slot).

Auto-ID prevents the need for jumpers to uniquely identify cards in a system. This can:

- increase the speed of system level repairs in the field,
- reduce the possibility of incorrect configurations, and
- reduce the number of unique spare cards that must be stocked.

Both methods of Auto ID employed by the Universe II are described below.

2.2.5.1 Auto Slot ID: VME64 Specified

The VME64 auto ID cycle (described in the VME64 Specification) requires at power-up that the Auto ID slave

- generate IRQ2*, and
- negate SYSFAIL*.

When the Auto ID slave responds to the Monarch's IACK cycle, it will

- enable accesses to its CR/CSR space,
- provide a Status/ID to the Monarch indicating the interrupt is an Auto-ID request,
- assert DTACK*, and
- release IRQ2*.

The Universe II participates in the VME64 auto ID cycle in either an automatic or semi-automatic mode. In its fully automatic mode, it holds SYSFAIL* asserted until SYSRST* is negated. When SYSRST* is negated, the Universe II asserts IRQ2* and releases SYSFAIL*. In its semi-automatic mode, the Universe II still holds SYSFAIL* asserted until SYSRST* is negated. However, when SYSRST* is negated, the local CPU performs diagnostics and local logic sets the AUTOID bit in the MISC_CTL register (Table A.82). This asserts IRQ2* and releases SYSFAIL*.

After SYSFAIL* is released and the Universe II detects a level 2 IACK cycle, it responds with the STATUS/ID stored in its level 2 STATID register (which defaults to 0xFE).

The Universe II can be programmed so that it will not release SYSFAIL* until the SYSFAIL bit in the VCSR_CLR register (Table A.125) is cleared by local logic (SYSFAIL* is asserted if the SYSFAIL bit in the VCSR_SET register, Table A.126, is set at power-up). Since the system Monarch does not service the Auto-ID slave until after SYSFAIL* is negated, not clearing the SYSFAIL bit allows the Auto-ID process to be delayed until the CPU completes local diagnostics. Once local diagnostics are complete, the CPU clears the SYSFAIL bit and the Auto-ID cycle proceeds.

The Monarch can perform CR/CSR reads and writes at A[23:19]= 0x00 in CR/CSR space and re-locate the Universe II's CR/CSR base address.

Universe II and the Auto-ID Monarch

At power-up an Auto-ID Monarch waits to run a level 2 IACK cycle until after SYSFAIL* goes high. After the IACK cycle is performed and it has received a Status/ID indicating an Auto-ID request, the monarch software

- masks IRQ2* (so that it will not service other interrupters at that interrupt level until current Auto-ID cycle is completed),
- performs an access at 0x00 in CR/CSR space to get information about Auto-ID slave,
- moves the CR/CSR base address to a new location, and
- unmask IRQ2* (to allow it to service the next Auto-ID slave).

The Universe II supports monarch activity through its capability to be a level 2 interrupt handler. All other activity must be handled through software residing on the board.

2.2.5.2 Auto-ID: A Proprietary Tundra Method

The Universe II uses a proprietary Auto-ID scheme if this is selected as a power-up option (see "Auto-ID" on page 2-118). The Tundra proprietary Auto-ID function identifies the relative position of each board in the system, without using jumpers or on-board information. The ID number generated by Auto-ID can then be used to determine the board's base address.

After any system reset (assertion of SYSRST*), the Auto-ID logic responds to the first level one IACK cycle on the VMEbus.

After the level one IACK* signal has been asserted (either through IRQ1* or with a synthesized version), the Universe II in slot 1 counts five clocks from the start of the cycle and then asserts IACKOUT* to the second board in the system (see Figure 2.4). All other boards continue counting until they receive IACKIN*, then count four more clocks and assert IACKOUT* to the next board. Finally, the last board asserts IACKOUT* and the bus pauses until the data transfer time-out circuit ends the bus cycle by asserting BERR*.




Figure 2.4 : Timing for Auto-ID Cycle

Because all boards are four clocks “wide”, the value in the clock counter is divided by four to identify the slot in which the board is installed; any remainder is discarded. Note that since the start of the IACK cycle is not synchronized to SYSCLK, a one count variation from the theoretical value of the board can occur. However, in all cases the ID value of a board is greater than that of a board in a lower slot number. The result is placed in the DY4AUTOID [7:0] field and the DY4DONE bit is set (both are located in the MISC_STAT register, Table A.83).

2.2.6 System Controller Functions

When located in Slot 1 of the VMEbus system (see “First Slot Detector” on page 2-23), the Universe II assumes the role of SYSCON and sets the SYSCON status bit in the MISC_CTL register (Table A.82). In accordance with the VME64 specification, as SYSCON the Universe II provides:

- a system clock driver,
- an arbitration module,
- an IACK Daisy Chain Driver (DCD), and
- a bus timer.

2.2.6.1 System Clock Driver

The Universe II provides a 16 MHz SYSCLK signal derived from CLK64 when configured as SYSCON.

2.2.6.2 VMEbus Arbiter

When the Universe II is SYSCON, the Arbitration Module is enabled. The Arbitration Module supports the following arbitration modes:

- Fixed Priority Arbitration Mode (PRI),
- Single Level Arbitration (SGL) (a subset of PRI), or
- Round Robin Arbitration Mode (RRS) (default setting).

These are set with the VARB bit in the MISC_CTL register (Table A.82).

Fixed Priority Arbitration Mode (PRI)

In this mode, the order of priority is VRBR#[3], VRBR#[2], VRBR#[1], and VRBR#[0] as defined by the VME64 specification. The Arbitration Module issues a Bus Grant (VBGO [3:0]#) to the highest requesting level.

If a Bus Request of higher priority than the current bus owner becomes asserted, the Arbitration Module asserts VBCLR# until the owner releases the bus (VRBBSY# is negated).

Single Level Arbitration Mode (SGL)

In this mode, a subset of priority mode, all requests and grants are made exclusively on level 3. Set the Universe II in PRI mode to use this mode.

Round Robin Arbitration Mode (RRS)

This mode arbitrates all levels in a round robin mode, repeatedly scanning from levels 3 to 0. Only one grant is issued per level and one owner is never forced from the bus in favor of another requester (VBCLR# is never asserted).

Since only one grant is issued per level on each round robin cycle, several scans will be required to service a queue of requests at one level.

VMEbus Arbiter Time-out

The Universe II's VMEbus arbiter can be programmed to time-out if the requester does not assert BBSY* within a specified period. This allows BGOUT to be negated so that the arbiter may continue with other requesters. The timer is programmed using the VARBTO field in the MISC_CTL register (Table A.82), and can be set to 16 μ s, 256 μ s, or disabled. The default setting for the timer is 16 μ s. The arbitration time-out timer has a granularity of 8 μ s; setting the timer for 16 μ s means the timer may timeout in as little as 8 μ s.

2.2.6.3 IACK Daisy-Chain Driver Module

The IACK Daisy-Chain Driver module is enabled when the Universe II becomes system controller. This module guarantees that IACKIN* will stay high for at least 30 ns as specified in rule 40 of the VME64 specification.

2.2.6.4 VMEbus Time-out

A programmable bus timer allows users to select a VMEbus time-out period. The time-out period is programmed through the VBTO field in the MISC_CTL register (Table A.82) and can be set to 16 μ s, 32 μ s, 64 μ s, 128 μ s, 256 μ s, 512 μ s, 1024 μ s, or disabled. The default setting for the timer is 64 μ s. The VMEbus Timer module asserts VXBERR# if a VMEbus transaction times out (indicated by one of the VMEbus data strobes remaining asserted beyond the time-out period).

2.2.7 BI-Mode

BI-Mode[®] (Bus Isolation Mode) is a mechanism for logically isolating the Universe II from the VMEbus. This mechanism is useful for the following purposes:

- implementing hot-standby systems. A system may have two identically configured boards, one in BI-Mode. If the board that is not in BI-Mode fails, it can be put in BI-Mode while the spare board is removed from BI-Mode.
- system diagnostics for routine maintenance, or
- fault isolation in the event of a card failure, even if a spare board is not provided, at least the faulty board can be isolated.

While in BI-Mode, the Universe II data channels cannot be used to communicate between VMEbus and PCI (Universe II mailboxes do provide a means of communication). The only traffic permitted is to Universe II registers either through configuration cycles, the PCI register image, the VMEbus register image, or CR/CSR space. No IACK cycles will be generated or responded to. No DMA activity will occur. Any access to other PCI images will result in a Target-Retry. Access to other VMEbus images will be ignored.

Entering BI-Mode has the following effects.

- The VMEbus Master Interface becomes inactive. PCI Target Channel coupled accesses will thereafter be retried. The PCI Target Channel Posted Writes FIFO will continue to accept transactions but will eventually fill and no further posted writes will be accepted. The DMA FIFO will eventually empty or fill and no further DMA activity will take place on the PCI bus. The Universe II VMEbus Master will not service interrupts while in BI-Mode.
- The Universe II will not respond as a VMEbus slave, except for accesses to the register image and CR/CSR image.
- The Universe II will not respond to any interrupt it had outstanding. All VMEbus outputs from the Universe II will be tri-stated, so that the Universe II will not be driving any VMEbus signals. The only exception to this is the IACK and BG daisy chains which must remain in operation as before.

There are four ways to cause the Universe II to enter BI-Mode. The Universe II is put into BI-Mode:

1. if the BI-Mode power-up option is selected (See “Power-up Option Descriptions” on page 2-117 and Table 2.22 on page 2-116),
2. when SYSRST* or RST# is asserted any time after the Universe II has been powered-up in BI-Mode,
3. when VRIRQ# [1] is asserted, provided that the ENGBI bit in the MISC_CTL register has been set, or
4. when the BI bit in the MISC_CTL register is set.

Note that when the Universe II is put in BI-Mode, the BI bit in the MISC_CTL register (Table A.82) is set. Clearing this bit ends Bi-Mode.

There are two ways to remove the Universe II from BI-Mode:

1. power-up the Universe II with the BI-Mode option off (see “BI-Mode” on page 2-118), or
2. clear the BI bit in the MISC_CTL register, which will be effective only if the source of the BI-Mode is no longer active. That is, if VRIRQ# [1] is still being asserted while the ENGBI bit in the MISC_CTL register is set, then attempting to clear the BI bit in the MISC_CTL register will not be effective.

2.3 PCI Bus Interface

The PCI Bus Interface is organized as follows:

- “PCI Cycles—Overview” below,
- “Universe II as PCI Master” on page 2-35, and
- “Universe II as PCI Target” on page 2-38.

The Universe II PCI Bus Interface is electrically and logically directly connected to the PCI bus. For information concerning the different types of PCI accesses available, see “PCI Bus Target Images” on page 2-53.

2.3.1 PCI Cycles—Overview

The PCI bus port of the Universe II operates as a PCI compliant port with a 64-bit multiplexed address/data bus. The Universe II PCI bus Interface is configured as little-endian using address invariant translation when mapping between the VMEbus and the PCI bus. Address invariant translation preserves the byte ordering of a data structure in a little-endian memory map and a big-endian memory map (see Appendix-E).

The Universe II has all the PCI signals described in the PCI specification with the exception of SBO# and SDONE (since the Universe II does not provide cache support).

Universe II PCI cycles are synchronous, meaning that bus and control input signals are externally synchronized to the PCI clock (CLK). PCI cycles are divided into four phases:

1. request,
2. address phase,
3. data transfer, and
4. cycle termination.

2.3.1.1 32-Bit Versus 64-Bit PCI

The Universe II is configured with a 32-bit or 64-bit PCI data bus at power-up (see “PCI Bus Width” on page 2-119 for directions on how to configure the PCI bus width.)

Each of the Universe II's VMEbus slave images can be programmed so that VMEbus transactions are mapped to a 64-bit data bus on the PCI Interface (with the LD64EN bit, e.g. see Table A.85). If the VMEbus slave image is programmed with a 64-bit PCI bus data width and if the Universe II powered up in a 64-bit PCI environment, then the Universe II asserts REQ64# during the address phase of the PCI transaction. If the PCI target is 64-bit capable, then it will respond with ACK64# and the Universe II will pack data to the full width (64 bits) of the PCI bus. If the PCI target is not 64-bit capable, then it does not assert ACK64# and the Universe II will pack data to a 32-bit PCI bus.



Note that REQ64# will be asserted if LD64EN is set in a 64-bit PCI system independent of whether the Universe II has a full 64 bits to transfer. This may result in a performance degradation because of the extra clocks required to assert REQ64# and to sample ACK64#. Also, there can be some performance degradation when accessing 32-bit targets with LD64EN set. Do not set this bit unless there are 64-bit targets in the slave image window.

If the VMEbus slave images are not programmed for a 64-bit wide PCI data bus, then the Universe operates transparently in a 32-bit PCI environment.

Independent of the setting of the LD64EN bit, the Universe II will never attempt a 64-bit cycle on the PCI bus if it is powered up as 32-bit.

2.3.1.2 PCI Bus Request and Parking

The Universe II supports bus parking. If the Universe II requires the PCI bus it will assert REQ# only if its GNT# is not currently asserted. When the PCI Master Module is ready to begin a transaction and its GNT# is asserted, the transfer begins immediately. This eliminates a possible one clock cycle delay before beginning a transaction on the PCI bus which would exist if the Universe II did not implement bus parking. Bus parking is described in Section 3.4.3 of the PCI Specification (Rev. 2.1).

2.3.1.3 Address Phase

PCI transactions are initiated by asserting FRAME# and driving address and command information onto the bus. In the VMEbus Slave Channel, the Universe II calculates the address for the PCI transaction by adding a translation offset to the VMEbus address (see "Universe as VMEbus Slave" on page 2-13).

The command signals (on the C/BE# lines) contain information about Memory space, cycle type and whether the transaction is read or write. Table 2.3 below gives PCI the command type encoding implemented with the Universe II.

Table 2.3 : Command Type Encoding for Transfer Type

C/BE# [3:0] for PCI, C/BE# [7:4] for non-multiplexed	Command Type	Universe II Capability
0000	Interrupt Acknowledge	N/A
0001	Special Cycle	N/A
0010	I/O Read	Target/Master
0011	I/O Write	Target/Master
0100	Reserved	N/A
0101	Reserved	N/A
0110	Memory Read	Target/Master
0111	Memory Write	Target/Master
1000	Reserved	N/A
1001	Reserved	N/A
1010	Configuration Read	Target/Master
1011	Configuration Write	Target/Master
1100	Memory Read Multiple	(See Text)
1101	Dual Address Cycle	N/A
1110	Memory Read Line	(See Text)
1111	Memory Write and Invalidate	(See Text)

Memory Read Multiple and Memory Read Line transactions are aliased to Memory Read transactions when the Universe II is accessed as a PCI target with these commands. Likewise, Memory Write and Invalidate is aliased to Memory Write. As a PCI initiator, the Universe II may generate Memory Read Multiple but never Memory Read Line.

PCI targets are expected to assert DEVSEL# if they have decoded the access. During a Configuration cycle, the target is selected by its particular IDSEL. If a target does not respond with DEVSEL# within 6 clocks, a Master-Abort is generated. The role of configuration cycles is described in the PCI 2.1 Specification.

2.3.1.4 Data Transfer

Acknowledgment of a data phase occurs on the first rising clock edge after both IRDY# and TRDY# are asserted by the master and target, respectively. REQ64# may be driven during the address phase to indicate that the master wishes to initiate a 64-bit transaction. The PCI target asserts ACK64# if it is able to respond to the 64-bit transaction.

Wait cycles are introduced by either the master or the target by deasserting IRDY# or TRDY#. For write cycles, data is valid on the first rising edge after IRDY# is asserted. Data is acknowledged by the target on the first rising edge with TRDY# asserted. For read cycles, data is transferred and acknowledged on first rising edge with both IRDY# and TRDY# asserted.

A single data transfer cycle is repeated every time IRDY# and TRDY# are both asserted. The transaction only enters the termination phase when FRAME# is deasserted (master-initiated termination) or if STOP# is asserted (target-initiated). When both FRAME# and IRDY# are deasserted (final data phase is complete), the bus is defined as idle.

2.3.1.5 Termination Phase

The PCI Bus Interface permits all four types of PCI terminations:

1. Master-Abort: the PCI bus master negates FRAME# when no target responds (DEVSEL# not asserted) after 6 clock cycles.
2. Target-Disconnect: a termination is requested by the target (STOP# is asserted) because it is unable to respond within the latency requirements of the PCI specification or it requires a new address phase. Target-disconnect means that the transaction is terminated after data is transferred. The Universe II will deassert REQ# for at least two clock cycles if it receives STOP# from the PCI target.
3. Target-Retry: termination is requested (STOP# is asserted) by the target because it cannot currently process the transaction. Retry means that the transaction is terminated after the address phase without any data transfer.
4. Target-Abort: is a modified version of target-disconnect where the target requests a termination (asserts STOP#) of a transaction which it will never be able to respond to, or during which a fatal error occurred. Although there may be a fatal error for the initiating application, the transaction completes gracefully, ensuring normal PCI operation for other PCI resources.

2.3.1.6 Parity Checking

The Universe II both monitors and generates parity information using the PAR signal. The Universe II monitors PAR when it accepts data as a master during a read or a target during a write. The Universe II drives PAR when it provides data as a target during a read or a master during a write. The Universe II also drives PAR during the address phase of a transaction when it is a master and monitors PAR during an address phase when it is the PCI target. In both address and data phases, the PAR signal provides even parity for C/BE#[7:0] and AD[63:0]. The Universe II continues with a transaction independent of any parity error reported during the transaction.

The Universe II can also be programmed to report address parity errors. It does this by asserting the SERR# signal and setting a status bit in its registers. No interrupt is generated, and regardless of whether assertion of SERR# is enabled, the Universe II does not respond to the errored access. If powered up in a 64-bit PCI environment, the Universe II uses PAR64 in the same way as PAR, except for AD[63:32] and C/BE[7:4].

2.3.2 Universe II as PCI Master

The Universe II requests PCI bus mastership through its PCI Master Interface. The PCI Master Interface is available to either the VMEbus Slave Channel (access from a remote VMEbus master) or the DMA Channel.

The VMEbus Slave Channel makes an internal request for the PCI Master Interface when:

- the RXFIFO contains a complete transaction,
- sufficient data exists in the RXFIFO to generate a transaction of length defined by the programmable aligned burst size (PABS), or
- there is a coupled cycle request.

The DMA Channel makes an internal request for the PCI Master Interface when:

- the DMAFIFO has room for 128 bytes to be read from PCI,
- the DMAFIFO has queued 128 bytes to be written to PCI, or
- the DMA block is completely queued during a write to the PCI bus.

Arbitration between the two channels for the PCI Master Interface follows a round robin protocol. Each channel is given access to the PCI bus for a single transaction. Once that transaction completes, ownership of the PCI Master Interface is granted to the other channel if it requires the bus. The VMEbus Slave Channel and the DMA Channel each have a set of rules that determine when it is ‘done’ with the PCI Master Interface. The VMEbus Slave Channel is done under the following conditions:

- an entire transaction (no greater in length than the programmed aligned burst size) is emptied from the RXFIFO, or
- the coupled cycle is complete.

The DMA Channel is done when:

- the boundary programmed into the PCI aligned burst size is emptied from the DMAFIFO during writes to the PCI bus, or
- the boundary programmed into the PCI aligned burst size is queued to the DMAFIFO during reads from the PCI bus.

As discussed elsewhere (“Universe as VMEbus Slave” on page 2-13), access from the VMEbus may be either coupled or decoupled. For a full description of the operation of these data paths, see “Universe as VMEbus Slave” on page 2-13.

The PCI Master Interface can generate the following command types:

- I/O Read,
- I/O Write,
- Memory Read,
- Memory Read Multiple,
- Memory Write,
- Configuration Read (Type 0 and 1), and
- Configuration Write (Type 0 and 1).

The type of cycle the Universe II generates on the PCI bus depends on which VMEbus slave image is accessed and how it is programmed. For example, one slave image might be programmed as an I/O space, another as Memory space and another for Configuration space (see “VME Slave Images” on page 2-50). When generating a memory transaction, the implied addressing is either 32-bit or 64-bit aligned, depending upon the PCI target. When generating an I/O transaction, the implied addressing is 32-bit aligned and all incoming transactions are coupled.

2.3.2.1 PCI Burst Transfers

The Universe II generates aligned burst transfers of some maximum alignment, according to the programmed PCI aligned burst size (PABS field in the MAST_CTL register, Table A.81). The PCI aligned burst size can be programmed at 32, 64 or 128 bytes. Burst transfers will not cross the programmed boundaries. For example, when programmed for 32-byte boundaries, a new burst will begin at XXXX_XX20, XXXX_XX40, etc. If necessary, a new burst will begin at an address with the programmed alignment. To optimize PCI bus usage, the Universe II always attempts to transfer data in aligned bursts at the full negotiated width of the PCI bus.

The Universe II can perform a 64-bit data transfer over the AD [63:0] lines, if operated in a 64-bit PCI environment or against a 64-bit capable target or initiator. The LD64EN bit must be set if the access is being made through a VMEbus slave image; the LD64EN bit must be set if the access is being performed with the DMA.

The Universe II generates burst cycles on the PCI bus if it is:

- emptying the RXFIFO (the RXFE status bit in the MISC_STAT register is set when the RXFIFO empties),
- filling the RDFIFO (receives a block read request from a VMEbus master to an appropriately programmed VMEbus slave image), or
- performing DMA transfers

All other accesses are treated as single data beat transactions on the PCI bus.

During PCI burst transactions, the Universe II dynamically enables byte lanes on the PCI bus by changing the BE# signals during each data phase.

2.3.2.2 Termination

The Universe II performs a Master-Abort if the target does not respond within 6 clock cycles. Coupled PCI transactions terminated with Target-Abort or Master-Abort are terminated on the VMEbus with BERR*. The R_TA or R_MA bits in the PCI_CS register (Table A.3) are set when the Universe II receives a Target-Abort or generates a Master-Abort independent of whether the transaction was coupled, decoupled, prefetched, or initiated by the DMA.

If the Universe II receives a retry from the PCI target, then it relinquishes the PCI bus and re-requests within 2-3 PCI clock cycles. No other transactions are processed by the PCI Master Interface until the retry condition is cleared. The Universe II can be programmed to perform a maximum number of retries using the MAXRTRY field in the MAST_CTL register (Table A.81). When this number of retries has been reached, the Universe II responds in the same way as it does to a Target-Abort on the PCI bus. That is, the Universe II may issue a BERR* signal on the VMEbus. Target-Aborts are discussed in the next two paragraphs. All VMEbus slave coupled transactions and decoupled transactions will encounter a delayed DTACK once the FIFO fills until the condition clears either due to success or a retry time-out.

If the error occurs during a posted write to the PCI bus (see also “Bus Error Handling” on page 2-58), the Universe II uses the L_CMDERR register (Table A.32) to log the command information for the transaction (CMDERR [3:0]) and the address of the errored transaction is latched in the LAERR register (Table A.33). The L_CMDERR register also records if multiple errors occur (with the M_ERR bit) although the number of errors is not given. The error log is qualified with the L_STAT bit. The rest of the transaction will be purged from the RXFIFO if some portion of the write encounters an error. An interrupt is generated on the VMEbus and/or PCI bus depending upon whether the VERR and LERR interrupts are enabled (see “Interrupt Handling” on page 2-68).

If an error occurs on the PCI bus, the Universe II does not translate the error condition into a BERR* on the VMEbus. Indeed, the Universe II does not directly map the error. By doing nothing, the Universe II forces the external VMEbus error timer to expire.

2.3.2.3 Parity

The Universe II monitors PAR when it accepts data as a master during a read and drives PAR when it provides data as a master during a write. The Universe II also drives PAR during the address phase of a transaction when it is a master. In both address and data phases, the PAR signal provides even parity for C/BE#[3:0] and AD[31:0]. If the Universe II is powered up in a 64-bit PCI environment, then PAR64 provides even parity for C/BE#[7:4] and AD[63:32].

The PERESP bit in the PCI_CS register (Table A.3) determines whether or not the Universe II responds to parity errors as PCI master. Data parity errors are reported through the assertion of PERR# if the PERESP bit is set. Regardless of the setting of these two bits, the D_PE (Detected Parity Error) bit in the PCI_CS register is set if the Universe II encounters a parity error as a master. The DP_D (Data Parity Detected) bit in the same register is only set if parity checking is enabled through the PERESP bit and the Universe II detects a parity error while it is PCI master (i.e. it asserts PERR# during a read transaction or receives PERR# during a write).

No interrupts are generated by the Universe II in response to parity errors reported during a transaction. Parity errors are reported by the Universe II through assertion of PERR# and by setting the appropriate bits in the PCI_CS register. If PERR# is asserted to the Universe II while it is PCI master, the only action it takes is to set the DP_D. The Universe II continues with a transaction independent of any parity errors reported during the transaction.

As a master, the Universe II does not monitor SERR#. It is expected that a central resource on the PCI bus will monitor SERR# and take appropriate action.

2.3.3 Universe II as PCI Target

This section covers the following aspects of the Universe as PCI bus target:

- “Overview” on page 2-38,
- “Data Transfer” on page 2-39,
- “Coupled Transfers” on page 2-42,
- “Posted Writes” on page 2-44,
- “The Special Cycle Generator” on page 2-45,
- “Using the VOWN bit” on page 2-48,
- “Terminations” on page 2-49.

2.3.3.1 Overview

The Universe II becomes PCI bus target when one of its eight programmed PCI target images or one of its registers is accessed by a PCI bus master (the Universe II cannot be that PCI bus master). Register accesses are discussed elsewhere (see “Registers” on page 2-100); this section describes only those accesses destined for the VMEbus.

When one of its PCI target images is accessed, the Universe II responds with DEVSEL# within two clocks of FRAME# (making the Universe II a medium speed device, as reflected by the DEVSEL field in the PCI_CS register).

As PCI target, the Universe II responds to the following command types:

- I/O Read,
- I/O Write,
- Memory Read,
- Memory Write,
- Configuration Read (Type 0),
- Configuration Write (Type 0),
- Memory Read Multiple (aliased to Memory Read),
- Memory Line Read (aliased to Memory Read),
- Memory Write and Invalidate (aliased to Memory Write).

Type 0 Configuration accesses can only be made to the Universe II's PCI configuration registers. The PCI target images do not accept Type 0 accesses.

Address parity errors are reported if both PERESP and SERR_EN are set in the PCI_CS register (Table A.3). Address parity errors are reported by the Universe II by asserting the SERR# signal and setting the S_SERR (Signalled SERR#) bit in the PCI_CS register. Assertion of SERR# can be disabled by clearing the SERR_EN bit in the PCI_CS register. No interrupt is generated, and regardless of whether assertion of SERR# is enabled or not, the Universe II does not respond to the access with DEVSEL#. Typically the master of the transaction times out with a Master-Abort.

If the Universe II is accessed validly with REQ64# in Memory space as a 64-bit target, then it responds with ACK64# if it is powered up as a 64-bit device.

2.3.3.2 Data Transfer

Read transactions are always coupled (as opposed to VMEbus slave reads, which may be pre-fetched; see "Universe as VMEbus Slave" on page 2-13). Write transactions can be coupled or posted (see Figure 2.5 below and "PCI Bus Target Images" on page 2-53). To ensure sequential consistency, coupled operations (reads or writes) are only processed once all previously posted write operations have completed (i.e. the TXFIFO is empty).

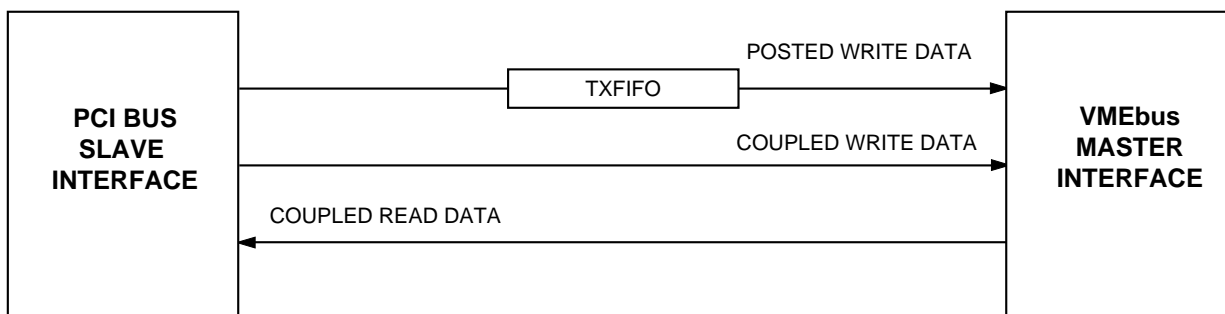


Figure 2.5 : PCI Bus Target Channel Dataflow

The data transfer between the PCI bus and VMEbus is perhaps best explained by Figure 2.6 below. The Universe II can be seen as a funnel where the mouth of the funnel is the data width of the PCI transaction. The end of the funnel is the maximum VMEbus data width programmed into the PCI target image (VDW bit in the PCI target image control register). For example, consider a 32-bit PCI transaction accessing a PCI target image with VDW set to 16 bits. A data beat with all byte lanes enabled will be broken into two 16-bit cycles on the VMEbus. If the PCI target image is also programmed with block transfers enabled, the 32-bit PCI data beat will result in a D16 block transfer on the VMEbus. Write data is unpacked to the VMEbus and read data is packed to the PCI bus data width.

If the data width of the PCI data beat is the same as the maximum data width of the PCI target image, then the Universe II maps the data beat to an equivalent VMEbus cycle. For example, consider a 32-bit PCI transaction accessing a PCI target image with VDW set to 32 bits. A data beat with all byte lanes enabled is translated to a single 32-bit cycle on the VMEbus.

As the general rule, if the PCI bus data width is less than the VMEbus data width then there is no packing or unpacking between the two buses. The only exception to this is during 32-bit PCI multi-data beat transactions to a PCI target image programmed with maximum VMEbus data width of 64 bits. In this case, packing/unpacking occurs to make maximum use of the full bandwidth on both buses.

Only aligned VMEbus transactions are generated, so if the requested PCI data beat has unaligned or non-contiguous byte enables, then it is broken into multiple aligned VMEbus transactions no wider than the programmed VMEbus data width. For example, consider a three-byte PCI data beat (on a 32-bit PCI bus) accessing a PCI target image with VDW set to 16 bits. The three-byte PCI data beat will be broken into two aligned VMEbus cycles: a single-byte cycle and a double-byte cycle (the ordering of the two cycles depends on the arrangement of the byte enables in the PCI data beat). If in the above example the PCI target image has a VDW set to 8 bits, then the three-byte PCI data beat will be broken into three single-byte VMEbus cycles.

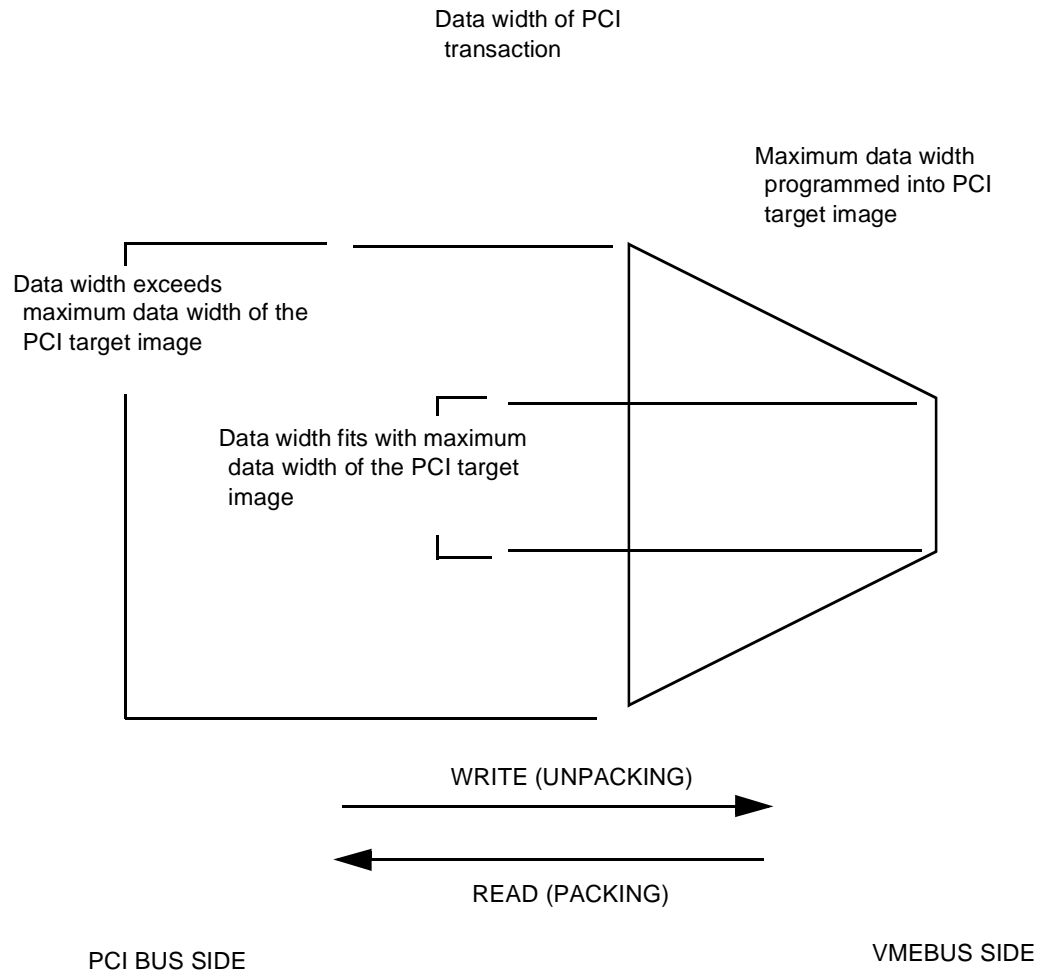


Figure 2.6 : Influence of Transaction Data Width and Target Image Data Width on Data Packing/Unpacking

2.3.3.3 Coupled Transfers

The PCI Target Channel supports “coupled transfers”. In a nutshell, a coupled transfer through the PCI Target Channel is a transfer between PCI and VME where the Universe II maintains ownership of the VMEbus from the beginning to the end of the transfer on the PCI bus (and possibly longer), and where the termination of the cycle on the VMEbus is relayed directly to the PCI initiator in the normal manner (i.e., Target-Abort, or Target Completion), rather than through error-logging and interrupts.

By default, all PCI target images are set for coupled transfers. Coupled transfers typically cause the Universe II to go through three phases: The Coupled Request Phase, the Coupled Data-Transfer Phase, and then the Coupled Wait Phase. When an external PCI Master attempts a data transfer through a slave image programmed for coupled cycles, then:

- If the Universe II currently owns the VMEbus, the PCI Target Channel moves directly to the Coupled Data-Transfer Phase; otherwise,
- the Universe II moves to the Coupled Request Phase. These three phases are described below.

Note that once the Coupled Request phase has begun, posted writes may traverse the PCI Target Channel without affecting coupled transfers.

Coupled Request Phase

During the Coupled Request Phase, the Universe II will attempt to acquire the VMEbus. But first it must empty any posted writes pending in the TXFIFO, and obtain ownership of the internal VMEbus Master Interface (see “VMEbus Release” on page 2-8 for more details on how the Universe II shares the VMEbus between channels.) The PCI Target Channel retries the PCI master until the PCI Target Channel obtains ownership of the VMEbus. Every time it issues such a retry, the Universe II restarts the Coupled Request Timer, which counts down a period of 2^{15} PCI clock cycles. The Coupled Request Timer co-determines how long the Universe II maintains the VMEbus since the last time the Universe II issued a Target-Retry during a Coupled Request Phase: the Universe II will release (or terminate its attempt to obtain) the VMEbus if a coupled transfer is not attempted before the Coupled Request Timer expires.

Usually, an external PCI Master will attempt a coupled cycle once the Universe II has acquired the VMEbus during its Coupled Request Phase. In this case the Universe will proceed to the “Coupled Data-Transfer Phase”. No address matching is performed to verify whether the current coupled cycle matches the initiating coupled cycle. If an external PCI Master requests a PCI I/O or RMW transfer with an illegal byte lane combination, the Universe II will exit the “Coupled Request Phase.”

Coupled Data-Transfer Phase

At the beginning of the Coupled Data-Transfer Phase, the Universe II latches the PCI command, byte enable, address and (in the case of a write) data. Regardless of the state of FRAME#, the Universe II retries¹ the master, and then performs the transaction on the VMEbus. The Universe II continues to signal Target-Retry to the external PCI master until the transfer completes (normally or abnormally) on the VMEbus.

If the transfer completes normally on the VMEbus, then in the case of a read, the data is transmitted to the PCI bus master. If a data phase of a coupled transfer requires packing or unpacking on the VMEbus, acknowledgment of the transfer is not given to the PCI bus master until all data has been packed or unpacked on the VMEbus. Successful termination is signalled on the PCI bus—the data beat is acknowledged with a Target-Disconnect, forcing all multi-beat transfers into single beat. At this point, the Universe II enters the Coupled Wait Phase.

If a bus error is signalled on the VMEbus or an error occurs during packing or unpacking, then the transaction is terminated on the PCI bus with Target-Abort.

See also “Data Transfer” on page 2-39.

Coupled Wait Phase

The Coupled Wait Phase is entered after the successful completion of a Coupled Data-Transfer phase. The Coupled Wait Phase allows consecutive coupled transactions to occur without releasing the VMEbus. If a new coupled transaction is attempted while the Universe II is in the Coupled Wait Phase, the Universe II will move directly to the Coupled Data-Transfer Phase without re-entering the Coupled Request Phase.

The Coupled Window Timer determines the maximum duration of the Coupled Wait Phase. When the Universe II enters the Coupled Wait Phase, the Coupled Window Timer starts. The period of this timer is specified in PCI clocks and is programmable through the CWT field of the LMISC register (Table A.30). If this field is programmed to 0000, the Universe II will do an early release of BBSY* during the coupled transfer on the VMEbus and will not enter the “Coupled Wait Phase.” In this case, VMEbus ownership is relinquished immediately by the PCI Target Channel after each coupled cycle.

Once the timer associated with the Coupled Wait Phase expires, the Universe II will release the VMEbus if release mode is set for RWD, or the release mode is set for ROR and there is a pending (external) request on the VMEbus.

1. PCI latency requirements (as described in revision 2.1 of the PCI Specification) require that only 16 clock cycles can elapse between the first and second data beat of a transaction. Since the Universe II cannot guarantee that data acknowledgment will be received from the VMEbus in time to meet these PCI latency requirements, the Universe II performs a target-disconnect after the first data beat of every coupled write transaction.

2.3.3.4 Posted Writes

Posted writes are enabled for a PCI target image by setting the PWEN bit in the control register of the PCI target image (see “PCI Bus Target Images” on page 2-53). Write transactions are relayed from the PCI bus to the VMEbus through a 32-entry deep TXFIFO. The TXFIFO allows each entry to contain 32 address bits (with extra bits provided for command information), or up to 64 data bits. For each posted write transaction received from the PCI bus, the PCI Target Interface queues an address entry in the FIFO. This entry contains the translated address space and mapped VMEbus attributes information relevant to the particular PCI target image that has been accessed (see “PCI Bus Target Images” on page 2-53). For this reason, any re-programming of PCI bus target image attributes will only be reflected in TXFIFO entries queued after the re-programming. Transactions queued before the re-programming are delivered to the VMEbus with the PCI bus target image attributes that were in use before the re-programming.



Caution: Care should be taken before reprogramming target images from one bus while that image is being accessed from the opposite bus. If there is a chance the image may be accessed while being reprogrammed, disable the image first before changing image attributes.

Once the address phase is queued in one TXFIFO entry, the PCI Target Interface may pack the subsequent data beats to a full 64-byte width before queuing the data into new entries in the TXFIFO.

For 32-bit PCI transfers in the Universe II, the TXFIFO will accept a single burst of one address phase and 59 data phases when it is empty. For 64-bit PCI, the TXFIFO will accept a single burst of one address phase and 31 data phases when it is empty. To improve PCI bus utilization, the TXFIFO does not accept a new address phase if it does not have room for a burst of one address phase and 128 bytes of data. If the TXFIFO does not have enough space for an aligned burst, then the posted write transaction is terminated with a Target-Retry immediately after the address phase.

When an external PCI Master posts writes to the PCI Target Channel of the Universe II, the Universe II will issue a disconnect if the implied address will cross a 256-byte boundary.

Before a transaction can be delivered to the VMEbus from the TXFIFO, the PCI Target Channel must obtain ownership of the VMEbus Master Interface. Ownership of the VMEbus Master Interface is granted to the different channels on a round robin basis (see “VMEbus Release” on page 2-8). Once the PCI Target Channel obtains the VMEbus through the VMEbus Master Interface, the manner in which the TXFIFO entries are delivered depends on the programming of the VMEbus attributes in the PCI target image (see “PCI Bus Target Images” on page 2-53). For example, if the VMEbus data width is programmed to 16 bits, and block transfers are disabled, then each data entry in the TXFIFO corresponds to four transactions on the VMEbus.

If block transfers are enabled in the PCI target image, then each transaction queued in the TXFIFO, independent of its length, is delivered to the VMEbus as a block transfer. This means that if a single data beat transaction is queued in the TXFIFO, it appears on the VMEbus as a single data phase block transfer.

Any PCI master attempting coupled transactions is retried while the TXFIFO contains data. If posted writes are continually written to the PCI Target Channel, and the FIFO does not empty, coupled transactions in the PCI Target Channel will not proceed and will be continually retried. This presents a potential starvation scenario.

2.3.3.5 The Special Cycle Generator

The Special Cycle Generator in the PCI Target Channel of the Universe II can be used in conjunction with one of the PCI Target Images to generate read-modify-write (RMW) and Address Only With Handshake (ADOH) cycles.

The address programmed into the SCYC_ADDR register (Table A.26), in the address space specified by the LAS field of the SCYC_CTL register (Memory or I/O), must appear on the PCI bus during the address phase of a transfer for the Special Cycle Generator to perform its function. Whenever this address on the PCI bus matches the address in the SCYC_ADDR register, the Universe II does not respond with ACK64# (since the Special Cycle Generator only processes up to 32-bit cycles).

The cycle that is produced on the VMEbus (if any) will use attributes programmed into the Image Control Register of the image that contains the address programmed in the SCYC_ADDR register.

The Special Cycle Generator is configured through the register fields shown in Table 2.4 and described below.

Table 2.4 : Register Fields for the Special Cycle Generator

Field	Register Bits	Description
32-bit address	ADDR in Table A.26	specifies PCI bus target image address
PCI Address Space	LAS in Table A.25	specifies whether the address specified in the ADDR field lies in PCI memory or I/O space
Special cycle	SCYC[1:0] in Table A.25	disabled, RMW or ADOH
32-bit enable	EN [31:0] in Table A.27	a bit mask to select the bits to be modified in the VMEbus read data during a RMW cycle
32-bit compare	CMP [31:0] in Table A.28	data which is compared to the VMEbus read data during a RMW cycle
32-bit swap	SWP [31:0] in Table A.29	data which is swapped with the VMEbus read data and written to the original address during a RMW cycle

The following sections describe the specific properties for each of the transfer types: RMW and ADOH.

Read-Modify-Write

When the SCYC field is set to RMW, any PCI bus read access to the specified PCI bus address (SCYC_ADDR register) will result in a RMW cycle on the VMEbus (provided the constraints listed below are satisfied). RMW cycles on the VMEbus consist of a single read followed by a single write operation. The data from the read portion of the RMW on the VMEbus is returned as the read data on the PCI bus.

RMW cycles make use of three 32-bit registers (see Table 2.5 above). The bit enable field is a bit mask which lets the user specify which bits in the read data are compared and modified in the RMW cycle. This bit enable setting is completely independent of the RMW cycle data width, which is determined by the data width of the initiating PCI transaction. During a RMW, the VMEbus read data is bitwise compared with the SCYC_CMP and SCYC_EN registers. The valid compared and enabled bits are then swapped using the SCYC_SWP register.

Each enabled bit that compares true is swapped with the corresponding bit in the 32-bit swap field. A false comparison results in the original bit being written back.

Once the RMW cycle completes, the VMEbus read data is returned to the waiting PCI bus master and the PCI cycle terminates.

Certain restrictions apply to the use of RMW cycles. If a write transaction is initiated to the VMEbus address when the special cycle field (SCYC in Table A.25) is set for RMW, then a standard write occurs with the attributes programmed in the PCI target image (in other words, the special cycle generator is not used). The Universe II performs no packing and unpacking of data on the VMEbus during a RMW operation. The following constraints must also be met.

1. The Special Cycle Generator will only generate a RMW if it is accessed with an 8-bit, aligned 16-bit, or aligned 32-bit read cycle.
2. The Special Cycle Generator will only generate a RMW if the size of the request is less than or equal to the programmed VMEbus Maximum Datawidth.
3. The destination VMEbus address space must be one of A16, A24 or A32.

In the event that the Special Cycle Generator is accessed with a read cycle that does not meet the three criteria described above, the Universe II generates a Target-Abort. Thus it is the user's responsibility to ensure that the Universe II is correctly programmed and accessed with correct byte-lane information.

VME Lock Cycles—Exclusive Access to VMEbus Resources

The VME Lock cycle is used in combination with the VOWN bit in the MAST_CTL register to lock resources on the VMEbus. The VME Lock cycle can be used by the Universe II to inform the resource that a locked cycle is intended (so that the VMEbus slave can prevent accesses from other masters on a different bus). The VOWN bit in the MAST_CTL register can be set to ensure that when the Universe II acquires the VMEbus, it is the only master given access to the bus (until the VOWN bit is cleared). It may also be necessary for the PCI master to have locked the Universe II using the PCI LOCK# signal.

When the SCYC field is set to VME Lock, any write access to the specified VMEbus address will result in a VME Lock cycle on the VMEbus. A VME Lock cycle is coupled: the cycle does not complete on the PCI bus until it completes on the VMEbus. Reads to the specified address translate to VMEbus reads in the standard fashion. The data during writes is ignored. The AM code generated on the VMEbus is determined by the PCI target image definition for the specified VMEbus address (see Table 2.12 on page 56).

However, after the VME Lock cycle is complete, there is no guarantee that the Universe II will remain VMEbus master unless it has set the VOWN bit. If the Universe II loses VMEbus ownership, then the VMEbus resource will no longer remain locked.

The following procedure is required to lock the VMEbus via an ADOH cycle:

- (If there is more than one master on the PCI bus, it may be necessary to use PCI LOCK# to ensure that the PCI master driving the ADOH cycle has sole PCI access to the Universe II registers and the VMEbus,)
- program the VOWN bit in the MAST_CTL register to a value of 1 (see “Using the VOWN bit” below),
- wait until the VOWN_ACK bit in the MAST_CTL register is a value of 1,
- generate an ADOH cycle with the Special Cycle Generator,
- perform transactions to be locked on the VMEbus,
- release the VMEbus by programming the VOWN bit in the MAST_CTL register to a value of 0, and
- wait until the VOWN_ACK bit in the MAST_CTL register is a value of 0.

In the event that BERR* is asserted on the VMEbus once the Universe II has locked and owns the VMEbus, it is the responsibility of the user to release ownership of the VMEbus by programming the VOWN bit in the MAST_CTL register to a value of 0.

The following restrictions apply to the use of VME Lock cycles:

1. All byte lane information is ignored for VME Lock cycles,
2. The Universe II will generate an VME Lock cycle on the VMEbus only if the PCI Target Image which subsumes the special cycle has posted writes disabled,
3. The Universe II Special Cycle Generator will not generate VME Lock cycles if the address space is not one of A16, A24 or A32. Instead it produces regular cycles.

2.3.3.6 Using the VOWN bit

The Universe II provides a VMEbus ownership bit (VOWN bit in the MAST_CTL register, Table A.81) to ensure that the Universe II has access to the locked VMEbus resource for an indeterminate period. The Universe II can be programmed to assert an interrupt on the PCI bus when it acquires the VMEbus and the VOWN bit is set (VOWN enable bit in the LINT_EN register, Table A.57). While the VMEbus is held using the VOWN bit, the Universe II sets the VOWN_ACK bit in the MAST_CTL register. The VMEbus Master Interface maintains bus tenure while the ownership bit is set, and only releases the VMEbus when the ownership bit is cleared. This function is important for the following two reasons.

If the VMEbus Master Interface is programmed for RWD (VREL bit in MAST_CTL register), it may release the VMEbus when the PCI Target Channel has completed a transaction (definition of ‘done’ for the PCI Target Channel, see “VMEbus Release” on page 2-8). Therefore, if exclusive access to the VMEbus resource is required for multiple transactions, then the VMEbus ownership bit will hold the bus until the exclusive access is no longer required.

Alternatively, if the VMEbus Master Interface is programmed for ROR, the VMEbus ownership bit will ensure VMEbus tenure even if other VMEbus requesters require the VMEbus.

2.3.3.7 Terminations

The Universe II performs the following terminations as PCI target:

1. Target-Disconnect
 - when registers are accessed with FRAME# asserted (no bursts allowed to registers),
 - after the first data beat of every coupled cycle, or
 - after the first data phase of a PCI Memory command (with FRAME# asserted) if AD[1:0] is not equal to 00, as recommended in Revision 2.1 of the PCI Specification (page 28).
2. Target-Retry
 - for 64-bit PCI, when a new posted write is attempted and the TXFIFO does not have room for a burst of one address phase and sixteen 64-bit data phases,
 - when a coupled transaction is attempted and the Universe II does not own the VMEbus,
 - when a coupled transaction is attempted while the TXFIFO has entries to process, or
 - when a master attempts to access the Universe II's registers while a VMEbus master owns the Register Channel (e.g., through a RMW access or another type of access).
3. Target-Abort
 - when the Universe II receives BERR* on the VMEbus during a coupled cycle (BERR* translated as Target-Abort on the PCI side and the S_TA bit is set in the PCI_CS register, Table A.3).

Whether to terminate a transaction or for retry purposes, the Universe II keeps STOP# asserted until FRAME# is deasserted, independent of the logic levels of IRDY# and TRDY#. If STOP# is asserted while TRDY# is deasserted, it means that the Universe II will not transfer any more data to the master.

If an error occurs during a posted write to the VMEbus, the Universe II uses the V_AMERR register (Table A.107) to log the AM code of the transaction (AMERR [5:0]), and the state of the IACK* signal (IACK bit, to indicate whether the error occurred during an IACK cycle). The FIFO entries for the offending cycle are purged. The V_AMERR register also records whether multiple errors have occurred (with the M_ERR bit) although the number is not given. The error log is qualified with the V_STAT bit (logs are valid if the V_STAT bit is set). The address of the errored transaction is latched in the VAERR register (Table A.108). When the Universe II receives a VMEbus error during a posted write, it generates an interrupt on the VMEbus and/or PCI bus depending upon whether the VERR and VERR interrupts are enabled (see "Interrupt Handling" on page 2-68).

2.4 Slave Image Programming

The Universe II recognizes two types of accesses on its bus interfaces: accesses destined for the other bus, and accesses decoded for its own register space. Address decoding for the Universe II's register space is described in "Registers" on page 2-100. This section describes the slave images used to map transactions between the PCI bus and VMEbus.

2.4.1 VME Slave Images

The Universe II accepts accesses from the VMEbus within specific programmed slave images. Each VMEbus slave image opens a window to the resources of the PCI bus and, through its specific attributes, allows the user to control the type of access to those resources. The tables below describe programming for the VMEbus slave images by dividing them into VMEbus, PCI bus and Control fields.

Table 2.5 : VMEbus Fields for VMEbus Slave Image

Field	Register Bits	Description
base	BS[31:12] or BS[31:16] in VSIx_BS	multiples of 4 or 64 Kbytes (base to bound: maximum of 4 GBytes)
bound	BD[31:12] or BD[31:16] in VSIx_BD	
address space	VAS in VSIx_CTL	A16, A24, A32, User 1, User 2
mode	SUPER in VSIx_CTL	supervisor and/or non-privileged
type	PGM in VSIx_CTL	program and/or data

Table 2.6 : PCI Bus Fields for VMEbus Slave Image

Field	Register Bits	Description
translation offset	TO[31:12] or TO[31:16] in VSIx_TO	offsets VMEbus slave address to a selected PCI address
address space	LAS in VSIx_CTL	Memory, I/O, Configuration
RMW	LLRMW in VSIx_CTL	RMW enable bit

Table 2.7 : Control Fields for VMEbus Slave Image

Field	Register Bits	Description
image enable	EN in VSIx_CTL	enable bit
posted write	PWEN in VSIx_CTL	posted write enable bit
prefetched read	PREN in VSIx_CTL	prefetched read enable bit
enable PCI D64	LD64EN in VSIx_CTL	enables 64-bit PCI bus transactions

Note that the Bus Master Enable (BM) bit of the PCI_CS register must be set in order for the image to accept posted writes from an external VMEbus master. If this bit is cleared while there is data in the VMEbus Slave Posted Write FIFO, the data will be written to the PCI bus but no further data will be accepted into this FIFO until the bit is set.

Tundra recommends that the attributes in a slave image not be changed while data is enqueued in the Posted Writes FIFO. To ensure data is dequeued from the FIFO, check the RXFE status bit in the MISC_STAT register (Table A.83) or perform a read from that image. If the programming for an image is changed after the transaction is queued in the FIFO, the transaction's attributes are not changed. Only subsequent transactions are affected by the change in attributes.

2.4.1.1 VMEbus Fields

Decoding for VMEbus accesses is based on the address, and address modifiers produced by the VMEbus master. Before responding to an external VMEbus master, the address must lie in the window defined by the base and bound addresses, and the Address Modifier must match one of those specified by the address space, mode, and type fields.

The Universe II's eight VMEbus slave images (images 0 to 7) are bounded by A32 space. The first and fifth of these images (VMEbus slave image 0 and 5) have a 4 Kbyte resolution while VMEbus slave images 1 to 3 and 6 to 8 have 64-Kbyte resolution (maximum image size of 4 GBytes). Typically, image 0 or 5 would be used as an A16 image since they provide the finest granularity of the eight images.



The address space of a VMEbus slave image must not overlap with the address space for the Universe II's control and status registers.

2.4.1.2 PCI Bus Fields

The PCI bus fields specify how the VMEbus transaction is mapped to the appropriate PCI bus transaction. The translation offset field allows the user to translate the VMEbus address to a different address on the PCI bus. The translation of VMEbus transactions beyond 4 Gbytes results in wrap-around to the low portion of the address range.

The PAS field controls generation of the PCI transaction command. The LLRMW bit allows indivisible mapping of incoming VMEbus RMW cycles to the PCI bus via the PCI LOCK# mechanism (see “VMEbus Read-Modify-Write Cycles (RMW Cycles)” on page 2-19). When the LLRMW bit is set, single cycle reads will always be mapped to single data beat locked PCI transactions. Setting this bit has no effect on non-block writes: they can be coupled or decoupled. However, note that only accesses to PCI Memory Space are decoupled, accesses to I/O or Configuration Space are always coupled.




Figure 2.7 : Address Translation Mechanism for VMEbus to PCI Bus Transfers

2.4.1.3 Control Fields

The control fields allow the user to enable a VMEbus slave image (using the EN bit), as well as specify how reads and writes will be processed. At power-up, all images are disabled and are configured for coupled reads and writes.

If the PREN bit is set, the Universe II will prefetch for incoming VMEbus block read cycles. It is the user's responsibility to ensure that prefetched reads are not destructive and that the entire image contains prefetchable resources.

If the PWEN bit is set, incoming write data from the VMEbus is loaded into the RXFIFO (see “Posted Writes” on page 2-15). Note that posted write transactions can only be mapped to Memory space on the PCI bus. Setting the PAS bit in the PCI fields to I/O or Configuration Space will force all incoming cycles to be coupled independent of this bit.

If the LD64EN bit is set, the Universe II will attempt to generate 64-bit transactions on the PCI bus by asserting REQ64#. The REQ64# line is asserted during the address phase in a 64-bit PCI system, and is the means of determining whether the PCI target is a 64-bit port. If the target asserts ACK64# with DEVSEL#, then the Universe II uses the 64-bit data bus. If the target does not assert ACK64# with DEVSEL#, then the Universe II uses a 32-bit data bus. However, note that use of REQ64# requires extra clocks internally. Therefore, if no 64-bit targets are expected on the PCI bus then performance can be improved by disabling LD64EN on the VMEbus slave images.



In order for a VMEbus slave image to respond to an incoming cycle, the PCI Master Interface must be enabled (bit BM in the PCI_CSR register, Table A.3).

2.4.2 PCI Bus Target Images

The Universe II accepts accesses from the PCI bus with specific programmed PCI target images. Each image opens a window to the resources of the VMEbus and allows the user to control the type of access to those resources. The tables below describe programming for the eight standard PCI bus target images (numbered 0 to 7) by dividing them into VMEbus, PCI bus and Control fields. One special PCI target image separate from the four discussed below is described in “Special PCI Target Image” on page 2-55.

Table 2.8 : PCI Bus Fields for the PCI Bus Target Image

Field	Register Bits	Description
base	BS[31:12] or BS[31:16] in LSIx_BS	multiples of 4 or 64 Kbytes (base to bound: maximum of 4 GBytes)
bound	BD[31:12] or BD[31:16] in LSIx_BD	
address space	LAS in LSIx_CTL	Memory or I/O

Table 2.9 : VMEbus Fields for the PCI Bus Target Image

Field	Register Bits	Description
translation offset	TO[31:12] or TO[31:16] in LSIx_TO	translates address supplied by PCI master to a specified VMEbus address
maximum data width	VDW in LSIx_CTL	8, 16, 32, or 64 bits
address space	VAS in LSIx_CTL	A16, A24, A32, CR/CSR, User1, User2
mode	SUPER in LSIx_CTL	supervisor or non-privileged
type	PGM in LSIx_CTL	program or data
cycle	VCT in LSIx_CTL	single or block

Table 2.10 : Control Fields for PCI Bus Target Image

Field	Register Bits	Description
image enable	EN in LSIx_CTL	enable bit
posted write	PWEN in LSIx_CTL	enable bit

Tundra recommends that the attributes in a target image not be changed while data is enqueued in the Posted Writes FIFO. To ensure data is dequeued from the FIFO, check the TXFE status bit in the MISC_STAT register (Table A.83) or perform a read from that image. If the programming for an image is changed after the transaction is queued in the FIFO, the transaction’s attributes are not changed. Only subsequent transactions are affected by the change in attributes.

2.4.2.1 PCI Bus Fields

All decoding for VMEbus accesses are based on the address and command information produced by a PCI bus master. The PCI Target Interface claims a cycle if there is an address match and if the command matches certain criteria.

All of the Universe II's eight PCI target images are A32-capable only. The first and fifth of them (i.e., PCI target images 0 and 4) have a 4 Kbyte resolution while PCI target images 1 to 3 and 5 to 8 have 64 Kbyte resolution. Typically, image 0 or image 4 would be used for an A16 image since they have the finest granularity.



The address space of a VMEbus slave image must not overlap with the address space for the Universe II's control and status registers.

2.4.2.2 VMEbus Fields

The VMEbus fields map PCI transactions to a VMEbus transaction, causing the Universe II to generate the appropriate VMEbus address, AM code, and cycle type. Some invalid combinations exist within the PCI target image definition fields. For example, A16 and CR/CSR spaces do not support block transfers, and A16 space does not support 64-bit transactions. Note that the Universe II does not attempt to detect or prevent these invalid programmed combinations, and that use of these combinations may cause illegal activity on the VMEbus.

The 21-bit translation offset allows the user to translate the PCI address to a different address on the VMEbus. The figure below illustrates the translation process:.




Figure 2.8 : Address Translation Mechanism for PCI Bus to VMEbus Transfers

Translations beyond the 4 Gbyte limit will wrap around to the low address range.

The AM code generated by the Universe II is a function of: the VMEbus fields, and the data width and alignment generated by the PCI bus master. For RMW and ADOH cycles, the AM code also depends on the settings for the Special Cycle Generator (see “The Special Cycle Generator” on page 2-45).

The address space, mode, type, and cycle fields control the VMEbus AM code for most transactions. Setting the cycle field to BLT enables the BLT cycle generation. MBLT cycles are generated when the maximum width is set to 64, the BLT bit is set, and the PCI master queues a transaction with at least 64 bits (aligned) of data.

The Universe II provides support for user defined AM codes. The USER_AM register (Table A.84) contains AM codes identified as User1 and User2. The USER_AM register can only be used to generate and accept AM codes 0x10 through 0x1F. These AM codes are designated as USERAM codes in the VMEbus specification. If the user selects one of these two, then the corresponding AM code from the global register is generated on the VMEbus. This approach results in standard single cycle transfers to A32 VMEbus address space independent of other settings in the VMEbus fields.

The VCT bits in the LSIx_CTL registers determine whether or not the VMEbus Master Interface will generate BLT transfers. The VCT bit will only be used if the VAS field is programmed for A24 or A32 space and the VDW bits are programmed for 8, 16, or 32 bits. If VAS bits of the control register are programmed to A24 or A32 and the VDW bits are programmed for 64-bit, the Universe II may perform MBLT transfers independent of the state of the VCT bit.

2.4.2.3 Control Fields

The control fields allow the user to enable a PCI target image (the EN bit), as well as specify how writes are processed. If the PWEN bit is set, then the Universe II will perform posted writes when that particular PCI target image is accessed. Posted write transactions are only decoded within PCI Memory space. Accesses from other spaces will result in coupled cycles independent of the setting of the PWEN bit.

2.4.3 Special PCI Target Image

The Universe II provides a special PCI target image located in Memory or I/O space. Its base address is aligned to 64-Mbyte boundaries and its size is fixed at 64 Mbytes (decoded using PCI address lines [31:26]). The Special PCI Target Image is divided into four 16Mbyte regions numbered 0 to 3 (see Figure 2.9 on page 2-57). These separate regions are selected with PCI address bits AD [25:24]. For example, if AD[25:24] = 01, then region 1 is decoded. Within each region, the upper 64Kbytes map to VMEbus A16 space, while the remaining portion of the 16 Mbytes maps to VMEbus A24 space. Note that no offsets are provided, so address information from the PCI transaction is mapped directly to the VMEbus.

The general attributes of each region are programmed according to the tables below.

Table 2.11 : PCI Bus Fields for the Special PCI Target Image

Field	Register Bits	Description
base	BS[5:0] in Table A.31	64 Mbyte aligned base address for the image
address space	LAS [1:0] in Table A.31	Places image in Memory or I/O

Table 2.12 : VMEbus Fields for the Special PCI Bus Target Image

Field	Register Bits	Description
maximum data width	VDW in Table A.31	separately sets each region for 16 or 32 bits
mode	SUPER in Table A.31	separately sets each region as supervisor or non-privileged
type	PGM in Table A.31	separately sets each region as program or data

Table 2.13 : Control Fields for the Special PCI Bus Target Image

Field	Register Bits	Description
image enable	EN in Table A.31	enable bit for the image
posted write	PWEN in Table A.31	enable bit for posted writes for the image

The special PCI target image provides access to all of A16 and most of A24 space (all except the upper 64 Kbytes). By using the special PCI target image for A16 and A24 transactions, it is possible to free the eight standard PCI target images (see “PCI Bus Target Images” on page 2-53), which are typically programmed to access A32 space.

Note that some address space redundancy is provided in A16 space. The VMEbus specification requires only two A16 spaces, while the special PCI target image allows for four A16 address spaces.

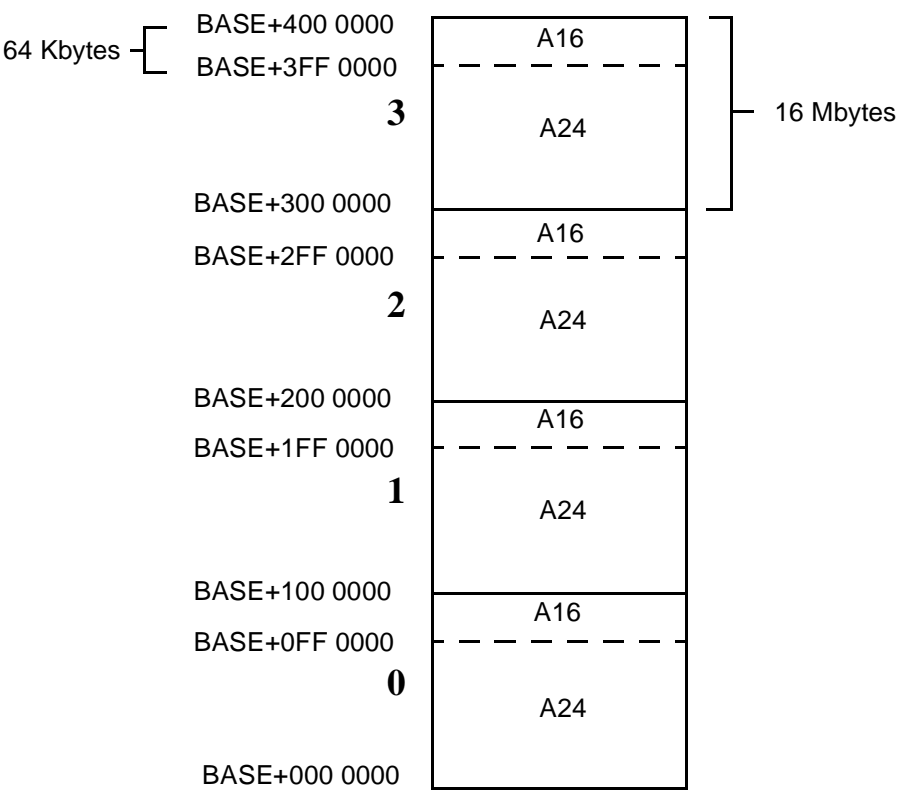


Figure 2.9 : Memory Mapping in the Special PCI Target Image

2.5 Bus Error Handling

There are two fundamentally different conditions under which bus errors may occur with the Universe II: during coupled cycles or during decoupled cycles. In a coupled transaction, the completion status is returned to the transaction master, which may then take some action. However, in a decoupled transaction, the master is not involved in the data acknowledgment at the destination bus and higher level protocols are required.

The error handling provided by the Universe II is described for both coupled and decoupled transactions below.

2.5.1 Coupled Cycles

During coupled cycles, the Universe II provides immediate indication of an errored cycle to the originating bus. VMEbus to PCI transactions terminated with Target-Abort or Master-Abort are terminated on the VMEbus with BERR*. The R_TA or R_MA bits in the PCI_CSR register (Table A.3 on page A-8) are set when the Universe II receives a Target-Abort or Master-Abort. For PCI to VMEbus transactions, a VMEbus BERR* received by the Universe II is communicated to the PCI master as a Target-Abort and the S_TA bit is set (Table A.3). No information is logged in either direction nor is an interrupt generated.

2.5.2 Decoupled Transactions

2.5.2.1 Posted Writes

The Universe II provides the option of performing posted writes in both the PCI Target Channel and the VMEbus Slave Channel. Once data is written into the RXFIFO or TXFIFO by the initiating master (VMEbus or PCI bus respectively), the Universe II provides immediate acknowledgment of the cycle's termination. When the data in the FIFO is written to the destination slave by the Universe II, the Universe II may subsequently receive a bus error instead of a normal termination. The Universe II handles this situation by logging the errored transactions in one of two error logs and generating an interrupt. Each error log (one for VMEbus errors and one for PCI bus errors) is comprised of two registers: one for address and one for command or address space logging.

If the error occurs during a posted write to the VMEbus, the Universe II uses the V_AMERR register (Table A.107) to log the AM code of the transaction (AMERR [5:0]). The state of the IACK* signal is logged in the IACK bit, to indicate whether the error occurred during an IACK cycle. The address of the errored transaction is latched in the V_AERR register (Table A.108). An interrupt is generated on the VMEbus and/or PCI bus depending upon whether the VERR and VERR interrupts are enabled (see “Interrupt Handling” on page 2-68). The remaining entries of the offending transaction are purged from the FIFO.

If the error occurs during a posted write to the PCI bus, the Universe II uses the L_CMDERR register (Table A.32) to log the command information for the transaction (CMDERR [3:0]). The address of the errored transaction is latched in the L_AERR register (Table A.33). An interrupt is generated on the VMEbus and/or PCI bus depending upon whether the VERR and LERR interrupts are enabled (see “Interrupt Handling” on page 2-68).

Under either of the above conditions (VMEbus to PCI, or PCI to VMEbus), the address that is stored in the log represents the most recent address the Universe II generated before the bus error was encountered. For single cycle transactions, the address represents the address for the actual errored transaction. However, for multi-data beat transactions (block transfers on the VMEbus or burst transactions on the PCI bus) the log only indicates that an error occurred somewhere after the latched address. For a VMEbus block transfer, the logged address will represent the start of the block transfer. In the PCI Target Channel, the Universe II generates block transfers that do not cross 256-byte boundaries, the error will have occurred from the logged address up to the next 256-byte boundary. In the VMEbus Slave Channel, the error will have occurred anywhere from the logged address up to the next burst aligned address.

In the case of PCI-initiated transactions, all data from the errored address up to the end of the initiating transaction is flushed from the TXFIFO. Since the Universe II breaks PCI transactions at 256-byte boundaries (or earlier if the TXFIFO is full), the data is not flushed past this point. If the PCI master is generating bursts that do not cross the 256-byte boundary, then (again) only data up to the end of that transaction is flushed.

In a posted write from the VMEbus, all data subsequent to the error in the transaction is flushed from the RXFIFO. However, the length of a VMEbus transaction differs from the length of the errored PCI bus transaction. For non-block transfers, the length always corresponds to one so only the errored data beat is flushed. However, if an error occurs on the PCI bus during a transaction initiated by a VMEbus block transfer, all data subsequent to the errored data beat in the block transfer is flushed from the RXFIFO. In the case of BLTs, this implies that potentially all data up to the next 256-byte boundary may be flushed. For MBLTs, all data up to the next 2-KByte boundary may be flushed.

Once an error is captured in a log, that set of registers is frozen against further errors until the error is acknowledged. The log is acknowledged and made available to latch another error by clearing the corresponding status bit in the VINT_STAT or LINT_STAT registers. Should a second error occur before the CPU has the opportunity to acknowledge the first error, another bit in the logs is set to indicate this situation (M_ERR bit).

2.5.2.2 Prefetched Reads

In response to a block read from the VMEbus, the Universe II initiates prefetching on the PCI bus (if the VMEbus slave image is programmed with this option, see “Slave Image Programming” on page 2-50). The transaction generated on the PCI bus is an aligned memory read transaction with multiple data beats extending to the aligned burst boundary (as programmed by PABS in the MAST_CTL register, Table A.82). Once an acknowledgment is given for the first data beat, an acknowledgment is sent to the VMEbus initiator by the assertion of DTACK*. Therefore, the first data beat of a prefetched read is coupled while all subsequent reads in the transaction are decoupled.

If an error occurs on the PCI bus, the Universe II does not translate the error condition into a BERR* on the VMEbus. Indeed, the Universe II does not directly map the error. By doing nothing, the Universe II forces the external VMEbus error timer to expire.

2.5.2.3 DMA Errors

How the Universe II responds to a bus error during a transfer controlled by the DMA Channel is described in “DMA Error Handling” on page 2-96.

2.5.2.4 Parity Errors

The Universe II both monitors and generates parity information using the PAR signal. The Universe II monitors PAR when it accepts data as a master during a read or as a target during a write. The Universe II drives PAR when it provides data as a target during a read or a master during a write. The Universe II also drives PAR during the address phase of a transaction when it is a master and monitors PAR during an address phase when it is the PCI target. In both address and data phases, the PAR signal provides even parity for C/BE#[3:0] and AD[31:0]. If the Universe II is powered up in a 64-bit PCI environment, then PAR64 provides even parity for C/BE#[7:4] and AD[63:32].

The PERESP and SERR_EN bits in the PCI_CS register (Table A.3) determine whether or not the Universe II responds to parity errors. Data parity errors are reported through the assertion of PERR# if the PERESP bit is set. Address parity errors, reported through the SERR# signal, are reported if both PERESP and SERR_EN are set. Regardless of the setting of these two bits, the D_PE (Detected Parity Error) bit in the PCI_CS register is set if the Universe II encounters a parity error as a master or as a target. The DP_D (Data Parity Detected) bit in the same register is only set if parity checking is enabled through the PERESP bit and the Universe II detects a parity error while it is PCI master (i.e. it asserts PERR# during a read transaction or receives PERR# during a write).

No interrupts are generated by the Universe II either as a master or as a target in response to parity errors reported during a transaction. Parity errors are reported by the Universe II through assertion of PERR# and by setting the appropriate bits in the PCI_CS register. If PERR# is asserted to the Universe II while it is PCI master, the only action it takes is to set the DP_D. Regardless of whether the Universe II is the master or target of the transaction, and regardless which agent asserted PERR#, the Universe II does not take any action other than to set bits in the PCI_CS register. The Universe II continues with a transaction independent of any parity errors reported during the transaction.

Similarly, address parity errors are reported by the Universe II (if the SERR_EN bit and the PERESP bit are set) by asserting the SERR# signal and setting the S_SERR (Signalled SERR#) bit in the PCI_CS register. Assertion of SERR# can be disabled by clearing the SERR_EN bit in the PCI_CS register. No interrupt is generated, and regardless of whether assertion of SERR# is enabled or not, the Universe II does not respond to the access with DEVSEL#. Typically the master of the transaction times out with a Master-Abort. As a master, the Universe II does not monitor SERR#. It is expected that a central resource on the PCI bus will monitor SERR# and take appropriate action.

2.6 Interrupt Generation

This is the first of two sections in this chapter which describe the Universe II's interrupt capabilities. This section describes the Universe II as a generator of interrupts. The following section, "Interrupt Handling" on page 2-68, describes the Universe II as an interrupt handler. Each of these sections has subsections which detail interrupt events on the PCI bus and VMEbus (i.e., how the Universe II can generate interrupts on the PCI bus and the VMEbus, and how the Universe II can respond to interrupt sources on the PCI bus and the VMEbus).

The Interrupt Channel handles the prioritization and routing of interrupt sources to interrupt outputs on the PCI bus and VMEbus. The interrupt sources are:

- the PCI LINT#[7:0] lines,
- the VMEbus IRQ*[7:1] lines,
- ACFAIL* and SYSFAIL*
- various internal events

These sources can be routed to either the PCI LINT# [7:0] lines or the VMEbus IRQ* [7:1] lines. Each interrupt source is individually maskable and can be mapped to various interrupt outputs. Most interrupt sources can be mapped to one particular destination bus. The PCI sources, LINT#[7:0], can only be mapped to the VMEbus interrupt outputs, while the VMEbus sources, VIRQ[7:1], can only be mapped to the PCI interrupt outputs. Some internal sources (for example, error conditions or DMA activity) can be mapped to either bus.

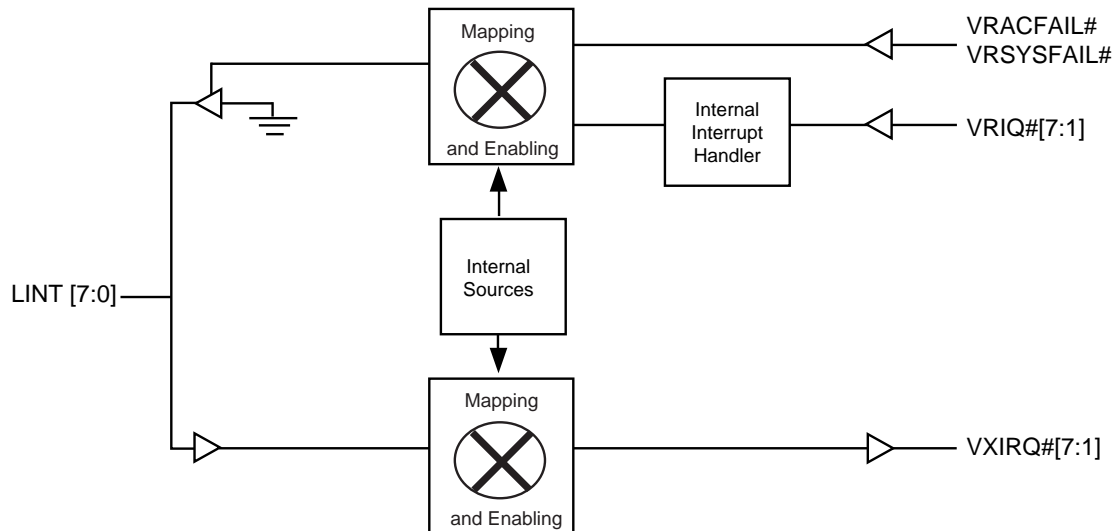


Figure 2.10 : Universe Interrupt Circuitry

Figure 2.10 above illustrates the circuitry inside the Universe II Interrupt Channel. The PCI hardware interrupts are listed on the left, and the VMEbus interrupt inputs and outputs are on the right. Internal interrupts are also illustrated. The figure shows that the interrupt sources may be mapped and enabled. The Internal Interrupt Handler is a block within the Universe II that detects assertion of the VRIRQ#[7:1] pins and generates the VME IACK through the VME Master. Upon completion of the IACK cycle, the Internal Interrupt Handler notifies the Mapping Block which in turn asserts the local LINT#, if enabled. (Whereas the Internal Interrupt Handler implies a delay between assertion of an interrupt condition to the Universe II and the Universe's mapping of the interrupt, all other interrupt sources get mapped immediately to their destination—assertion of LINT# immediately causes an IRQ, assertion of ACFAIL immediately causes an LINT#, etc.) This is described in further detail in the following sections.

2.6.1 PCI Interrupt Generation

The Universe II expands on the basic PCI specification which permits “single function” devices to assert only a single interrupt line. Eight PCI interrupt outputs provide maximum flexibility, although if full PCI compliancy is required, the user may route all interrupt sources to a single PCI interrupt output.



Only one of the PCI interrupt outputs, LINT#[0], has the drive strength to be fully compliant with the PCI specification. The other seven may require buffering if they are to be routed to PCI compliant interrupt lines. For most applications, however, the drive strength provided should be sufficient.

PCI interrupts may be generated from multiple sources:

- VMEbus sources of PCI interrupts
 - IRQ*[7:1]
 - SYSFAIL*
 - ACFAIL*

- internal sources of PCI interrupts
 - DMA
 - VMEbus bus error encountered
 - PCI Target-Abort or Master-Abort encountered
 - VMEbus ownership has been granted while the VOWN bit is set (see “VME Lock Cycles—Exclusive Access to VMEbus Resources” on page 2-47)
 - software interrupt
 - mailbox access
 - location monitor access
 - VMEbus IACK cycle performed in response to a software interrupt

Each of these sources may be individually enabled in the LINT_EN register (Table A.57) and mapped to a single LINT# signal through the LINT_MAP0, LINT_MAP1, and LINT_MAP2 registers (Table A.59, Table A.60, Table A.73). When an interrupt is received on any of the enabled sources, the Universe II asserts the appropriate LINT# pin and sets a matching bit in the LINT_STAT register (Table A.62). See Table 2.15 below for a list of the enable, mapping and status bits for PCI interrupt sources.

Table 2.14 : Source, Enabling, Mapping, and Status of PCI Interrupt Output

Interrupt Source	Enable Bit in LINT_EN (Table A.57)	Mapping Field in LINT_MAPx (Table A.59, Table A.60, Table A.73)	Status Bit in LINT_STAT (Table A.58)
ACFAIL*	ACFAIL	ACFAIL	ACFAIL
SYSFAIL*	SYSFAIL	SYSFAIL	SYSFAIL
PCI Software Interrupt	SW_INT	SW_INT	SW_INT
VMEbus Software IACK	SW_IACK	SW_IACK	SW_IACK
VMEbus Error	VERR	VERR	VERR
PCI Target-Abort or Master-Abort	LERR	LERR	LERR
DMA Event	DMA	DMA	DMA
VMEbus Interrupt Input	VIRQ7-1	VIRQ7-1	VIRQ7-1
Location Monitor	LM3-0	LM3-0	LM3-0
Mailbox Access	MBOX3-0	MBOX3-0	MBOX3-0
VMEbus Ownership	VOWN	VOWN	VOWN

The LINT_STAT register shows the status of all sources of PCI interrupts, independent of whether that source has been enabled. This implies that an interrupt handling routine must mask out those bits in the register that do not correspond to enabled sources on the active LINT# pin.

Except for SYSFAIL* and ACFAIL*, all sources of PCI interrupts are edge-sensitive. Enabling of the ACFAIL* or SYSFAIL* sources (ACFAIL and SYSFAIL bits in the LINT_EN register) causes the status bit and mapped PCI interrupt pin to assert synchronously with the assertion of the ACFAIL* or SYSFAIL* source. The PCI interrupt is negated once the ACFAIL or SYSFAIL status bit is cleared. The status bit cannot be cleared if the source is still active. Therefore, if SYSFAIL* or ACFAIL* is still asserted while the interrupt is enabled the interrupt will continue to be asserted. Both of these sources are synchronized and filtered with multiple edges of the 64 MHz clock at their inputs.

All other sources of PCI interrupts are edge-sensitive. Note that the VMEbus source for PCI interrupts actually comes out of the VMEbus Interrupt Handler block and reflects acquisition of a VMEbus STATUS/ID. Therefore, even though VMEbus interrupts externally are level-sensitive as required by the VMEbus specification, they are internally mapped to edge-sensitive interrupts (see “VMEbus Interrupt Handling” on page 2-68).

The interrupt source status bit (in the LINT_STAT register) and the mapped LINT# pin remain asserted with all interrupts. The status bit and the PCI interrupt output pin are only released when the interrupt is cleared by writing a “one” to the appropriate status bit.

2.6.2 VMEbus Interrupt Generation

This section details the conditions under which the Universe II generates interrupts to the VMEbus.

Interrupts may be generated on any combination of VMEbus interrupt lines (IRQ*[7:1]) from multiple sources:

- PCI sources of VMEbus interrupts
 - LINT#[7:0]
- Internal sources of VMEbus interrupts
 - DMA
 - VMEbus bus error encountered
 - PCI Target-Abort or Master-Abort encountered
 - Mailbox register access
 - software interrupt

Each of these sources may be individually enabled through the VINT_EN register (Table A.61) and mapped to a particular VMEbus Interrupt level using the VINT_MAPx registers (Table A.63, Table A.64, and Table A.74). Multiple sources may be mapped to any VMEbus level. Mapping interrupt sources to level 0 effectively disables the interrupt.

Once an interrupt has been received from any of the sources, the Universe II sets the corresponding status bit in the VINT_STAT register (Table A.62), and asserts the appropriate VMEbus interrupt output signal (if enabled). When a VMEbus interrupt handler receives the interrupt, it will perform an IACK cycle at that interrupt level. When the Universe II decodes that IACK cycle together with IACKIN* asserted, it provides the STATUS/ID previously stored in the STATID register (Table A.65), unless it is configured as SYSCON in which case it does not monitor IACKIN*. See Table 2.15 below for a list of the enable, mapping and status bits for VMEbus interrupt sources.

Table 2.15: Source, Enabling, Mapping, and Status of VMEbus Interrupt Outputs

Interrupt Source	Enable Bit in VINT_EN (Table A.61)	Mapping Field in VINT_MAPx (Table A.63, Table A.64, Table A.74)	Status Bit in VINT_STAT (Table A.62)
VMEbus Software Interrupt	SW_INT7-1	N/A ^a	SW_INT7-1
VMEbus Error	VERR	VERR (Table A.64)	VERR
PCI Target-Abort or Master-Abort	LERR	LERR (Table A.64)	LERR
DMA Event	DMA	DMA (Table A.64)	DMA
Mailbox Register	MBOX3-0	MBOX3-0 (Table A.74)	MBOX3-0
PCI bus Interrupt Input	LINT7-0	LINT7-0 (Table A.74)	LINT7-0
VMEbus Software Interrupt (mappable)	SW_INT	SW_INT (Table A.64)	SW_INT

a. This set of software interrupts cannot be mapped. That is, setting the SW_INT1 bit triggers VXIRQ1, setting the SW_INT2 bit triggers VXIRQ2 etc.

For all VMEbus interrupts, the Universe II interrupter supplies a pre-programmed 8-bit STATUS/ID: a common value for all interrupt levels. The upper seven bits are programmed in the STATID register. The lowest bit is cleared if the source of the interrupt was the software interrupt, and is set for all other interrupt sources. If a software interrupt source and another interrupt source are active and mapped to the same VMEbus interrupt level, the Universe II gives priority to the software source.




Figure 2.11 : STATUS/ID Provided by Universe II

Once the Universe II has provided the STATUS/ID to an interrupt handler during a software initiated VMEbus interrupt, it generates an internal interrupt, SW_IACK. If enabled, this interrupt feeds back to the PCI bus (through one of the LINT# pins) to signal a process that the interrupt started through software has been completed.

All VMEbus interrupts generated by the Universe II are RORA, except for the software interrupts which are ROAK. This means that if the interrupt source was a software interrupt, then the VMEbus interrupt output is automatically negated when the Universe II receives the IACK cycle. However, for any other interrupt, the VMEbus interrupt output remains asserted until cleared by a register access. Writing a “one” to the relevant bit in the VINT_STAT register clears that interrupt source. However, since PCI interrupts are level-sensitive, if an attempt is made to clear the VMEbus interrupt while the LINT# pin is still asserted, the VMEbus interrupt remains asserted. This causes a second interrupt to be generated to the VMEbus. For this reason, a VMEbus interrupt handler should clear the source of the PCI interrupt before clearing the VMEbus interrupt.

Since software interrupts are ROAK, the respective bits in the VINT_STAT register are cleared automatically on completion of the IACK cycle, simultaneously with the negation of the IRQ.

2.7 Interrupt Handling

This is the second of two sections in this chapter which describe the Universe II's interrupt capabilities. The previous section, "Interrupt Generation" on page 2-62, described the interrupt outputs of the Universe II on the PCI bus and the VMEbus. The current section describes how the Universe II responds to interrupt sources. In other words, this section describes the Universe II as an interrupt handler.

This section is broken down as follows:

- "PCI Interrupt Handling" on page 2-68 explains how the Universe II can respond to hardware interrupts on the PCI bus,
- "VMEbus Interrupt Handling" on page 2-68 explains how the Universe II can respond to hardware interrupts (or SYSFAIL* and ACFAIL*) on the VMEbus,
- "Internal Interrupt Handling" on page 2-71 explains how internal states of the Universe II can trigger interrupts.

2.7.1 PCI Interrupt Handling

This section explains how the Universe II can respond to hardware interrupts on the PCI bus.

All eight PCI interrupt lines, LINT#[7:0], can act as interrupt inputs to the Universe II. They are level-sensitive and, if enabled in the VINT_EN register (Table A.61), immediately generate an interrupt to the VMEbus. It is expected that when a VMEbus interrupt handler receives the Universe II's STATUS/ID from the Universe II, the interrupt handler will clear the VMEbus interrupt by first clearing the source of the interrupt on the PCI bus, and then clearing the VMEbus interrupt itself (by writing a "one" to the appropriate bit in the VINT_STAT register, Table A.62).

Note that since PCI interrupts are level-sensitive, if an attempt is made to clear the VMEbus interrupt while the LINT# pin is still asserted, the VMEbus interrupt remains asserted. This causes a second interrupt to be generated to the VMEbus. For this reason, a VMEbus interrupt handler should clear the source of the PCI interrupt before clearing the VMEbus interrupt.

2.7.2 VMEbus Interrupt Handling

This section explains how the Universe II can respond to hardware events on the VMEbus as an interrupt handler.

As a VMEbus interrupt handler, the Universe II can monitor any or all of the VMEbus interrupt levels. It can also monitor SYSFAIL* and ACFAIL*, although IACK cycles are not generated for these inputs. Each interrupt is enabled through the LINT_EN register (Table A.57).

Once enabled, assertion of any of the VMEbus interrupt levels, $IRQ[7:1]^*$, causes the internal interrupt handler circuitry to request ownership of the Universe II's VMEbus Master Interface on the level programmed in the `MAST_CTL` register (see “VMEbus Requester” on page 2-6). This interface is shared between several channels in the Universe II: the PCI Target Channel, the DMA Channel, and the Interrupt Channel. The Interrupt Channel has the highest priority over all other channels and, if an interrupt is pending, assumes ownership of the VMEbus Master Interface when the previous owner has relinquished ownership.

The Universe II latches the first interrupt that appears on the VMEbus and begins to process it immediately. Thus if an interrupt at a higher priority is asserted on the VMEbus before $BBSY^*$ is asserted the Universe II will perform an interrupt acknowledge for the first interrupt it detected. Upon completion of that IACK cycle, the Universe II will then perform IACK cycles for the higher of any remaining active interrupts.

There may be some latency between reception of a VMEbus interrupt and generation of the IACK cycle. This arises because of the latency involved in the Interrupt Channel gaining control of the VMEbus Master Interface, and because of possible latency in gaining ownership of the VMEbus if the VMEbus Master Interface is programmed for release-when-done. In addition, the Universe II only generates an interrupt on the PCI bus once the IACK cycle has completed on the VMEbus. Because of these combined latencies (time to acquire VMEbus and time to run the IACK cycle), systems should be designed to accommodate a certain worst case latency from VMEbus interrupt generation to its translation to the PCI bus.

When the Universe II receives a `STATUS/ID` in response to an IACK cycle, it stores that value in one of seven registers. These registers, `V1_STATID` through `V7_STATID` (Table A.66 to Table A.72), store the `STATUS/ID` corresponding to each IACK level (in the `STATID` field). Once an IACK cycle has been generated and the resulting `STATUS/ID` is latched, another IACK cycle will not be run on that level until the level has been re-armed by writing a “one” to the corresponding status bit in the `VINT_STAT` register (Table A.62). If other interrupts (at different levels) are pending while the interrupt is waiting to be re-armed, IACK cycles are run on those levels in order of priority and the `STATUS/IDs` stored in their respective registers.

Once the IACK cycle is complete and the `STATUS/ID` stored, an interrupt is generated to the PCI bus on one of `LINT#[7:0]` depending on the mapping for that VMEbus level in the `LINT_MAP0` register. The interrupt is cleared and the VMEbus interrupt level is re-armed by clearing the correct bit in the `LINT_STAT` register.

2.7.2.1 Bus Error During VMEbus IACK Cycle

A bus error encountered on the VMEbus while the Universe II is performing an IACK cycle is handled by the Universe II in two ways. The first is through the error logs in the VMEbus Master Interface. These logs store address and command information whenever the Universe II encounters a bus error on the VMEbus (see “Bus Error Handling” on page 2-58). If the error occurs during an IACK cycle, the IACK# bit is set in the V_AMERR register (Table A.107). The VMEbus Master Interface also generates an internal interrupt to the Interrupt Channel indicating a VMEbus error occurred. This internal interrupt can be enabled and mapped to either the VMEbus or PCI bus.

As well as generating an interrupt indicating an error during the IACK cycle, the Universe II also generates an interrupt as though the IACK cycle completed successfully. If an error occurs during the fetching of the STATUS/ID, the Universe II sets the ERR bit in the Vx_STATID register (Table A.66 to Table A.72), and generates an interrupt on the appropriate LINT# pin (as mapped in the LINT_MAP0 register, Table A.62). The PCI resource, upon receiving the PCI interrupt, is expected to read the STATUS/ID register, and take appropriate actions if the ERR bit is set. Note that the STATUS/ID cannot be considered valid if the ERR bit is set in the STATUS/ID register.

It is important to recognize that the IACK cycle error may generate two PCI interrupts: one through the VMEbus master bus error interrupt and another through the standard PCI interrupt translation. Should an error occur during acquisition of a STATUS/ID, the VINT_STAT register (Table A.62) will show that both VIRQx, and VERR are active.

2.7.3 Internal Interrupt Handling

The Universe II's internal interrupts are routed from several processes in the device. There is an interrupt from the VMEbus Master Interface to indicate a VMEbus error, another from the PCI Master Interface to indicate an error on that bus, another from the DMA to indicate various conditions in that channel, along with several others as indicated in Table 2.16 below. Table 2.16 shows to which bus each interrupt source may be routed (some sources may be mapped to both buses, but we recommend that you map interrupts to a single bus).

Table 2.16 : Internal Interrupt Routing

Interrupt Source	May be Routed to:	
	VMEbus	PCI Bus
PCI s/w interrupt		√
VMEbus s/w interrupt	√	
IACK cycle complete for s/w interrupt		√
DMA event	√	√
Mailbox access	√	√
Location monitor		√
PCI Target-Abort or Master-Abort	√	√
VMEbus bus error	√	√
VMEbus bus ownership granted		√

Figure 2.12 shows the sources of interrupts, and the interfaces from which they originate. Interrupt handling for each one of these sources is described in the following subsections.

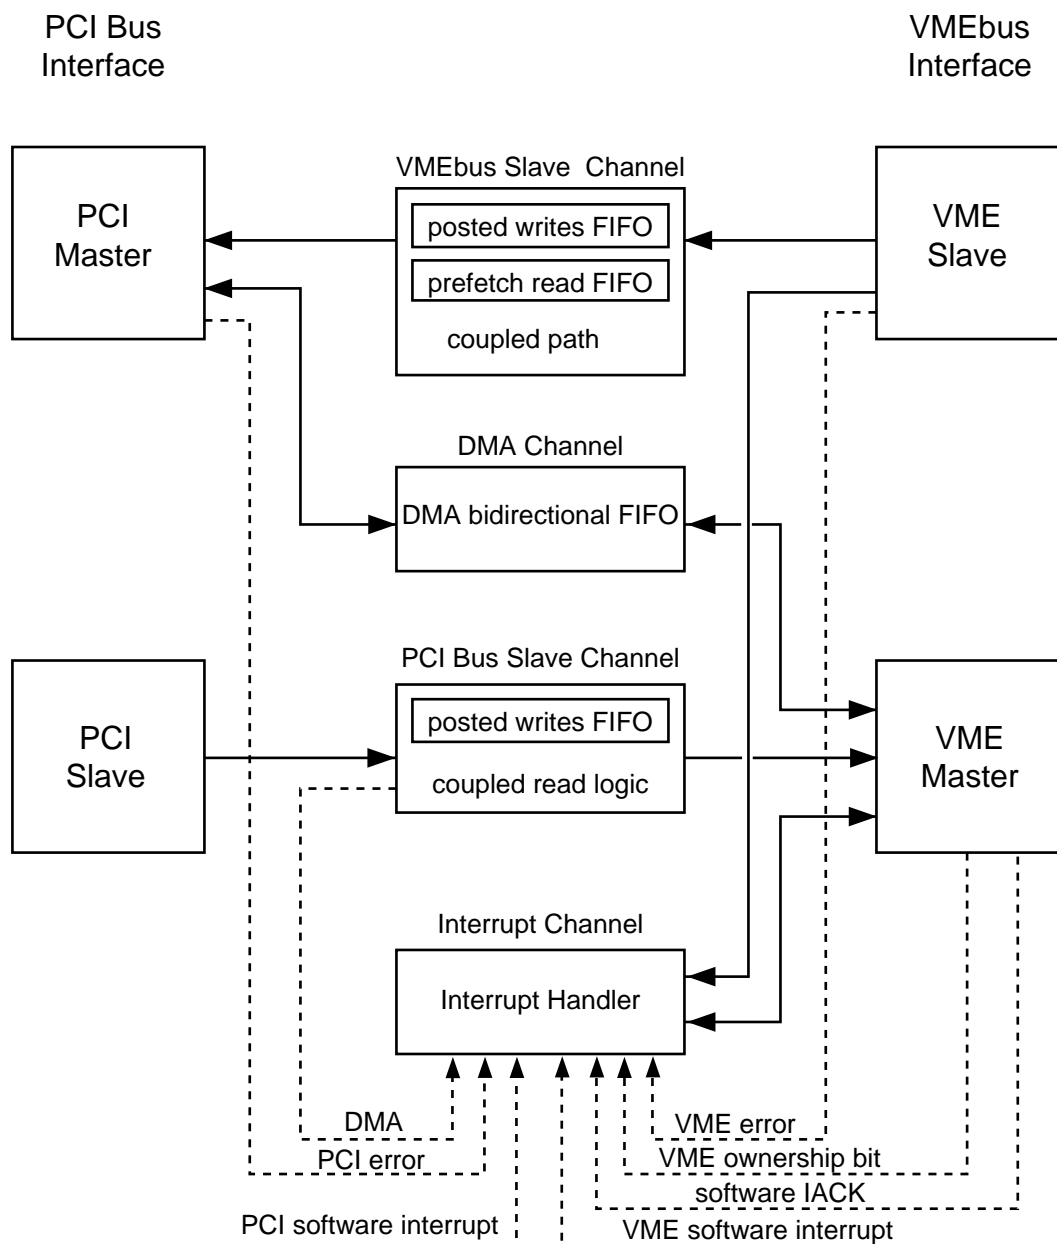


Figure 2.12 : Sources of Internal Interrupts

2.7.3.1 VMEbus and PCI Software Interrupts

It is possible to interrupt the VMEbus and the PCI bus through software. These interrupts may be triggered by writing a “one” to the respective enable bits.

Interrupting the VMEbus through software

There are two methods of triggering software interrupts on the VMEbus. The second method is provided for compatibility with the Universe I.

1. The first method for interrupting the VMEbus through software involves writing “one” to one of the SW_INT7-1 bits in the VINT_EN register (Table A.61) while the mask bit is zero.¹ This causes an interrupt to be generated on the corresponding IRQ7-1 line. That is, setting the SW_INT1 bit triggers VXIRQ1, setting the SW_INT2 bit triggers VXIRQ2, etc.
2. The second method for interrupting the VMEbus through software involves an extra step. Writing a “one” to the SW_INT bit in the VINT_EN register when this bit is “zero” (Table A.61) triggers one (and only one) interrupt on the VMEbus on the level programmed in the VINT_MAP1 register (Table A.64). Notice that this method requires that the user specify in the VINT_MAP1 register to which line the interrupt is to be generated. When the SW_INT interrupt (method 2) is active at the same level as one of SW_INT7-1 interrupts (method 1), the SW_INT interrupt (method 2) takes priority. While this interrupt source is active, the SW_INT status bit in the VINT_STAT register is set.

With both methods, the mask bit (SW_INTx or SW_INT) in the VINT_EN register must be zero in order for writing “one” to the bit to have any effect.

Regardless of the software interrupt method used, when an IACK cycle is serviced on the VMEbus, the Universe II can be programmed to generate an interrupt on the PCI bus by setting the SW_IACK enable bit in the LINT_EN register (see “Software IACK Interrupt” on page 2-74).

Interrupting the PCI bus through software

On the PCI bus, there is only one method of directly triggering a software interrupt. (This method is analogous to the second method described in the previous section.) Causing a “zero” to “one” transition in the SW_INT in the LINT_EN (Table A.57) register generates an interrupt to the PCI bus. While this interrupt source is active, the SW_INT status bit in LINT_STAT is set. The SW_INT field in the LINT_MAP1 register (Table A.60) determines which interrupt line is asserted on the PCI interface.

1. The term “enable” is more meaningful with respect to the other fields in this register, i.e., excluding the software interrupts. Writing to the software interrupt fields of this register does not enable an interrupt, it triggers an interrupt.

Termination of software interrupts

Any software interrupt may be cleared by clearing the respective bit in the VINT_EN or LINT_EN register. However, this method is not recommended for software VME bus interrupts because it may result in a spurious interrupt on that bus. That is, the Universe II will then not respond to the interrupt handler's IACK cycle, and the handler will be left without a STATUS/ID for the interrupt.

Since the software interrupt is edge-sensitive, the software interrupt bit in the VINT_EN or LINT_EN register should be cleared any time between the last interrupt finishing and the generation of another interrupt. It is recommended that the appropriate interrupt handler clear this bit once it has completed its operations. Alternatively, the process generating a software interrupt could clear this bit before re-asserting it.

Software interrupts on the VMEbus have priority over other interrupts mapped internally to the same level on the VMEbus. When a VMEbus interrupt handler generates an IACK cycle on a level mapped to both a software interrupt and another interrupt, the Universe II always provides the STATUS/ID for the software interrupt (bit zero of the Status/ID is cleared). If there are no other active interrupts on that level, the interrupt is automatically cleared upon completion of the IACK cycle (since software interrupts are ROAK).

While the software interrupt STATUS/ID has priority over other interrupt sources, the user can give other interrupt sources priority over the software interrupt. This is done by reading the LINT_STAT register (Table A.62) when handling a Universe II interrupt. This register indicates all active interrupt sources. Using this information, the interrupt handler can then handle the interrupt sources in any system-defined order.

2.7.3.2 Software IACK Interrupt

The Universe II generates an internal interrupt when it provides the software STATUS/ID to the VMEbus. This interrupt can only be routed to a PCI interrupt output. A PCI interrupt will be generated upon completion of an IACK cycle that had been initiated by the Universe II's software interrupt if:

- the SW_IACK bit in the LINT_EN register (Table A.57) is set, and
- the SW_IACK field in the LINT_MAP1 register (Table A.60) is mapped to a corresponding PCI interrupt line.

This interrupt could be used by a PCI process to indicate that the software interrupt generated to the VMEbus has been received by the slave device and acknowledged.

Like other interrupt sources, this interrupt source can be independently enabled through the LINT_EN register (Table A.57) and mapped to a particular LINT# pin using the LINT_MAP1 register (Table A.60). A status bit in the LINT_STAT register (Table A.62) indicates when the interrupt source is active, and is used to clear the interrupt once it has been serviced.

2.7.3.3 VMEbus Ownership Interrupt

The VMEbus ownership interrupt is generated when the Universe II acquires the VMEbus in response to programming of the VOWN bit in the MAST_CTL register (Table A.81). This interrupt source can be used to indicate that ownership of the VMEbus is ensured during an exclusive access (see “VME Lock Cycles—Exclusive Access to VMEbus Resources” on page 2-47). The interrupt is cleared by writing a one to the matching bit in the LINT_STAT register (Table A.62).

2.7.3.4 DMA Interrupt

The DMA module provides six possible interrupt sources:

- if the DMA is stopped (INT_STOP),
- if the DMA is halted (INT_HALT),
- if the DMA is done (INT_DONE),
- for PCI Target-Abort or Master-Abort (INT_LERR),
- for VMEbus errors (INT_VERR), or
- if there is a PCI protocol error or if the Universe II is not enabled as PCI master (INT_P_ERR).

All of these interrupt sources are ORed to a single DMA interrupt output line. When an interrupt comes from the DMA module, software must read the DMA status bits (Table A.55) to discover the originating interrupt source. The DMA interrupt can be mapped to either the VMEbus or one of the PCI interrupt output lines. See “DMA Interrupts” on page 2-95.

2.7.3.5 Mailbox Register Access Interrupts

The Universe II can be programmed to generate an interrupt on the PCI bus and/or the VMEbus when any one of its mailbox registers is accessed (see “Mailbox Registers” on page 2-108). The user may enable or disable an interrupt response to the access of any mailbox register (Table A.57). Each register access may be individually mapped to a specific interrupt on the PCI bus (LINT_MAP2, Table A.73) and/or the VMEbus (VINT_MAP2, Table A.74). The status of the PCI interrupt and the VMEbus are recorded in the LINT_STAT (Table A.58) and VINT_STAT registers (Table A.62), respectively.

2.7.3.6 Location Monitors

The Universe II can be programmed to generate an interrupt on the PCI bus when one of its four location monitors is accessed (see “Location Monitors” on page 2-19).

In order for an incoming VMEbus transaction to activate the location monitor of the Universe II, the location monitor must be enabled, the access must be within 4 kbytes of the location monitor base address (LM_BS, Table A.102), and it must be in the specified address space.

When an access to a location monitor is detected, an interrupt may be generated on the PCI bus (if the location monitor is enabled). There are four location monitors:

- VA[4:3] = 00 selects Location Monitor 1,
- VA[4:3] = 01 selects Location Monitor 2,
- VA[4:3] = 10 selects Location Monitor 3, and
- VA[4:3] = 11 selects Location Monitor 4.

The user may enable or disable an interrupt response to the access of any location monitor with bits in the LINT_EN register (Table A.57). Access to each location monitor may be individually mapped to a specific interrupt on the PCI bus (LINT_MAP2, Table A.73)—not to the VMEbus bus. The status of the PCI interrupt is logged in (LMn bit of the LINT_STAT, Table A.58).

2.7.3.7 PCI and VMEbus Error Interrupts

Interrupts from VMEbus errors, PCI Target-Aborts or Master-Aborts are generated only when bus errors arise during decoupled writes. The bus error interrupt (from either a PCI or VMEbus error) can be mapped to either a VMEbus or PCI interrupt output line.

2.7.4 VME64 Auto-ID

The Universe II includes a power-up option for participation in the VME64 Auto-ID process. When this option is enabled, the Universe II generates a level 2 interrupt on the VMEbus before release of SYSFAIL*. When the level 2 IACK cycle is run by the system Monarch, the Universe II responds with the Auto-ID Status/ID, 0xFE, and enables access to a CR/CSR image at base address 0x00_0000.

When the Monarch detects an Auto-ID STATUS/ID on level 2, it is expected to access the enabled CR/CSR space of the interrupter. From there it completes identification and configuration of the card. The Monarch functionality is typically implemented in software on one card in the VMEbus system. See “Automatic Slot Identification” on page 2-24.

2.8 DMA Controller

The Universe II has a DMA controller for high performance data transfer between the PCI bus and VMEbus. It is operated through a series of registers that control the source and destination for the data, length of the transfer and the transfer protocol to be used. There are two modes of operation for the DMA: Direct Mode, and Linked List Mode. In direct mode, the DMA registers are programmed directly by the external PCI master. In linked list mode, the registers are loaded from PCI memory by the Universe II, and the transfer described by these registers is executed. A block of DMA registers stored in PCI memory is called a command packet. A command packet may be linked to another command packet, such that when the DMA has completed the operations described by one command packet, it automatically moves on to the next command packet in the linked-list of command packets.

This section is broken into the following major sub-sections

1. “DMA Registers Outline” on page 2-77 describes in detail how the DMA is programmed from a register perspective.
2. “Direct Mode Operation” on page 2-83 describes how to operate the DMA when directly programming the DMA registers.
3. “Linked-List Operation” on page 2-86 describes how to operate the DMA when a linked-list of command packets describing DMA transfers is stored in PCI memory.
4. “FIFO Operation and Bus Ownership” on page 2-92 describes internally how the DMA makes use of its FIFO and how this affects ownership of the VMEbus and PCI bus.
5. “DMA Interrupts” on page 2-95 describes the interrupts generated by the DMA
6. “Interactions with Other Channels” on page 2-96 discusses the relations between the DMA Channel and the other data channels.
7. “DMA Error Handling” on page 2-96 describes how to handle errors encountered by the DMA

2.8.1 DMA Registers Outline

The DMA registers reside in a block starting at offset 0x200. They describe a single DMA transfer: where to transfer data from; where to transfer data to; how much data to transfer; and the transfer attributes to use on the PCI bus and VMEbus. A final register contains status and control information for the transfer. While the DMA is active, the registers are locked against any changes so that any writes to the registers will have no impact.

In direct-mode operation, these registers would be programmed directly by the user. In linked-list operation, they are repeatedly loaded by the Universe II from command packets residing in PCI memory until the end of the linked-list is reached (see “Linked-List Operation” on page 2-86).

2.8.1.1 Source and Destination Addresses

The source and destination addresses for the DMA reside in two registers: the DMA PCI bus Address Register (DLA register, Table A.52), and the DMA VMEbus Address Register (DVA register in Table A.53). The determination of which is the source address, and which is the destination is made by the L2V bit in the DCTL register (Table A.50). When set, the DMA transfers data from the PCI to the VMEbus. Hence DLA becomes the PCI source register and DVA becomes the VMEbus destination register. When cleared, the DMA transfers data from the VMEbus to PCI bus and DLA becomes the PCI destination register; DVA becomes the VMEbus source register.

The PCI address may be programmed to any byte address in PCI Memory space. It cannot transfer to or from PCI I/O or Configuration spaces.

The VMEbus address may also be programmed to any byte address, and can access any VMEbus address space from A16 to A32 in supervisory or non-privileged space, and data or program space. The setting of address space, A16, A24 or A32, is programmed in the VAS field of the DCTL register (Table A.50). The sub-spaces are programmed in the PGM and SUPER fields of the same register.



Although the PCI and VMEbus addresses may be programmed to any byte aligned address, they must be 8-byte aligned to each other (i.e. the low three bits of each must be identical). If not programmed with aligned source and destination addresses and an attempt to start the DMA is made, the DMA will not start, it will set the protocol error bit (P_ERR) in the DCSR register (Table A.55), and if enabled to, generate an interrupt. Linked-list operations will cease.

In direct mode the user must reprogram the source and destination address registers (DMA, DLA) before each transfer. These registers are not updated in direct mode. In linked-list mode, these registers are updated by the DMA when (and only when) the DMA is stopped, halted, or at the completion of processing a command packet. If read during DMA activity, the y will return the number of bytes remaining to transfer on the PCI side. All of the DMA registers are locked against any changes by the user while the DMA is active. When stopped due to an error situation, the DLA and DVA registers should not be used, but the DTBC is valid (see “DMA Error Handling” on page 2-96 for details). At the end of a successful linked-list transfer, the DVA and DLA registers will point to the next address at the end of the transfer block, and the DTBC register will be zero.

2.8.1.2 Transfer Size

The DMA may be programmed through the DMA Transfer Byte Count register (DTBC register in Table A.51) to transfer any number of bytes from 1 byte to 16 MBytes. There are no alignment requirements to the source or destination addresses. Should the width of the data turnovers (8- through 64-bit on VMEbus and 32- or 64-bit on PCI) not align to the length of the transfer or the source/destination addresses, the DMA will insert transfers of smaller width on the appropriate bus. For example, if a 15-byte transfer is programmed to start at address 0x1000 on the VMEbus, and the width is set for D32, the DMA will perform three D32 transfers, followed by a D16 transfer, followed by a D08 transfer. The Universe II does not generate unaligned transfers. On a 32-bit PCI bus, if the start address was 0x2000, the DMA would generate three data beats with all byte lanes enabled, and a fourth with three byte lanes enabled.

The DTBC register is not updated while the DMA is active (indicated by the ACT bit in the DGCS register). At the end of a transfer it will contain zero. However, if stopped by the user (via the STOP bit in the DGCS register) or the DMA encounters an error, the DTBC register contains the number of bytes remaining to transfer on the source side. See “DMA Error Handling” on page 2-96.

Starting the DMA while DTBC=0 will result in one of two situations. If the CHAIN bit in the DGCS register (Table A.55) is not set, the DMA will not start; it will perform no action. If the CHAIN bit is set, then the DMA loads the DMA registers with the contents of the command packet pointed to by the DCPD register (Table A.54), and starts the transfers described by that packet. Note that the DCPD[31:5] field of the DCPD register implies that the command packets be 32-byte aligned (bits 4:0 of this register must be 0).

2.8.1.3 Transfer Data Width

The VMEbus and PCI bus data widths are determined by three fields in the DCTL register (Table A.50). These fields affect the speed of the transfer. They should be set for the maximum allowable width that the destination device is capable of accepting.

On the VMEbus, the DMA supports the following data widths:

D08(E0), D16, D16BLT, D32, D64, D32BLT, and D64BLT (MBLT).

The width of the transfer is set with the VDW field in the DCTL register. The VCT bit determines whether or not the Universe II VMEbus Master will generate BLT transfers. The value of this bit only has meaning if the address space is A24 or A32 and the data width is not 64 bits. If the data width is 64 bits the Universe II may perform MBLT transfers independent of the state of the VCT bit.

The Universe II may perform data transfers smaller than that programmed in the VDW field in order to bring itself into alignment with the programmed width. For example if the width is set for D32 and the starting VMEbus address is 0x101, the DMA will perform a D08 cycle followed by a D16 cycle. Only once it has achieved the alignment set in the VDW field does it start D32 transfers. At the end of the transfer, the DMA will also have to perform more low-width transfers if the last address is not aligned to VDW. Similarly, if the VCT bit is set to enable block transfers, the DMA may perform non-block transfers to bring itself into alignment.

On the PCI bus, the DMA provides the option of performing 32- or 64-bit PCI transactions through the LD64EN bit in the DCTL register (Table A.50). If the Universe II has powered-up on a 32-bit bus (see “Power-up Option Descriptions” on page 2-117), this bit will have no effect. If powered-up on a 64-bit bus, this bit can provide some performance improvements when accessing 32-bit targets on that bus. Following the PCI specification, before a 64-bit PCI initiator starts a 64-bit transaction, it engages in a protocol with the intended target to determine if it is 64-bit capable. This protocol typically consumes one clock period. To save bandwidth, the LD64EN bit can be cleared to bypass this protocol when it is known that the target is only 32-bit capable

2.8.1.4 DMA Command Packet Pointer

The DMA Command Packet Pointer (DCPP in Table A.54) points to a 32-byte aligned address location in PCI Memory space that contains the next command packet to be loaded once the transfer currently programmed into the DMA registers has been successfully completed. When it has been completed (or the DTBC register is zero when the GO bit is set) the DMA reads the 32-byte command packet from PCI memory and executes the transfer it describes.

2.8.1.5 DMA Control and Status

The DMA General Control/Status Register (DGCS in Table A.55) contains a number of fields that control initiation and operation of the DMA as well as actions to be taken on completion.

DMA Initiation

Once all the parameters associated with the transfer have been programmed (source/destination addresses, transfer length and data widths, and if desired, linked lists enabled), the DMA transfer is started by setting the GO bit in the DGCS register. This causes the DMA first to examine the DTBC register. If it is non-zero, it latches the values programmed into the DCTL, DTBC, DLA, and DVA registers and initiates the transfer programmed into those registers. If DTBC=0, it checks the CHAIN bit in the DGCS register and if that bit is cleared it assumes the transfer to have completed and stops. Otherwise, if the CHAIN bit is set, it loads into the DMA registers the command packet pointed to by the DCPP register and initiates the transfer describe there.

If the GO bit is set, but the Universe II has not been enabled as a PCI master with the BM (bus master enable) bit in the PCI_CSR register, or if the DVA and DLA contents are not 64-bit aligned to each other, the transfer does not start, a protocol error is indicated by the P_ERR bit in the DGCS register and, if enabled, an interrupt is generated.

If the DMA has been terminated for some reason (stopped, halted, or error), all DMA registers contain values indicating where the DMA terminated. Once all status bits have been cleared, the DMA may be restarted from where it left off by simply setting the GO bit. The GO bit will only have an effect if all status bits have been cleared. These bits include STOP, HALT, DONE, LERR, VERR, and P_ERR; all in the DGCS register. These bits are all cleared by writing "one" to them, either before or while setting the GO bit.

The GO bit always returns a zero when read independent of the DMA's current state. Clearing the bit has no impact at any time. The ACT bit in the DGCS register indicates whether the DMA is currently active. It is set by the DMA once the GO bit is set, and cleared when the DMA is idle. Generally, when the ACT bit is cleared, one of the other status bits in the DGCS register is set (DONE, STOP, HALT, LERR, VERR, or P_ERR), indicating why the DMA is no longer active.

DMA VMEbus Ownership

Two fields in the DGCS register determine how the DMA will share the VMEbus with the other two potential masters in the Universe II (PCI Target Channel, and Interrupt Channel), and with other VMEbus masters on the bus. These fields are: VON and VOFF.

VON affects how much data the DMA will transfer before giving the opportunity to another master (either the Universe II or an external master) to assume ownership of the bus. The VON counter is used to temporarily stop the DMA from transferring data once a programmed number of bytes have been transferred (256 bytes, 512 bytes, 1K, 2K, 4K, 8K, or 16K). When performing MBLT transfers on the VMEbus, the DMA will stop performing transfers within 2048 bytes after the programmed VON limit has been reached. When not performing MBLT transfers, the DMA will stop performing transfers within 256 bytes once the programmed limit has been reached. When programmed for Release-When-Done operation, the Universe II will perform an early release of BBSY* when the VON counter reaches its programmed limit. VON may be disabled by setting the field to zero. When set as such, the DMA will continue transferring data as long as it is able.

There are other conditions under which the DMA may relinquish bus ownership. See "FIFO Operation and Bus Ownership" on page 2-92 for details on the VMEbus request and release conditions for the DMA.

VOFF affects how long the DMA will wait before re-requesting the bus after the VON limit has been reached. By setting VOFF to zero, the DMA will immediately re-request the bus once the VON boundary has been reached. Since the DMA operates in a round-robin fashion with the PCI Target Channel, and in a priority fashion with the Interrupt Channel, if either of these channels require ownership of the VMEbus, they will receive it at this time.

VOFF is only invoked when VMEbus tenure is relinquished due to encountering the VON boundary. When the VMEbus is released due to other conditions (e.g., the DMAFIFO has gone full while reading from the VMEbus), it will be re-requested as soon as that condition is cleared. The VOFF timer can be programmed to various time intervals from 0 μ s to 1024 μ s. See “FIFO Operation and Bus Ownership” on page 2-92 for details on the VMEbus request and release conditions for the DMA.

See “Interactions with Other Channels” on page 2-96 for information on other mechanisms which may delay the DMA Channel from acquiring the VMEbus or the PCI bus.

DMA Completion and Termination

Normally, the DMA will continue processing its transfers and command packets until either it completes everything has been requested to, or it encounters an error. There are also two methods for the user to interrupt this process and cause the DMA to terminate prematurely: Stop and Halt. Stop causes the DMA to terminate immediately, while halt causes the DMA to terminate when it has completed processing the current command packet in a linked list.

When the STOP_REQ bit in the DGCS register is set by the user, it tells the DMA to cease its operations on the source bus immediately. Remaining data in the FIFO continues to be written to the destination bus until the FIFO is empty. Once the FIFO is empty, the STOP bit in the same register is set and, if enabled, an interrupt generated. The DMA registers will contain the values that the DMA stopped at: the DTBC register contains the number of bytes remaining in the transfer, the source and destination address registers contain the next address to be read/written, the DCPD register contains the next command packet in the linked-list, and the DCTL register contains the transfer attributes.

If read transactions are occurring on the VMEbus, then setting a stop request can be affected by the VOFF timer. If the STOP_REQ bit is set while the DMA is lying idle waiting for VOFF to expire before recommencing reads, then the request remains pending until the VOFF timer has expired and the bus has been granted.

Halt provides a mechanism to interrupt the DMA at command packet boundaries during a linked-list transfer. In contrast, a stop requests the DMA to be interrupted immediately, while halt takes effect only when the current command packet is complete. A halt is requested of the DMA by setting the HALT_REQ bit in the DGCS register. This causes the DMA to complete the transfers defined by the current contents of the DMA registers and, if the CHAIN bit is set, load in the next command packet. The DMA then terminates, the HALT bit in the DGCS register is set, and, if enabled, an interrupt generated.

After a stop or halt, the DMA can be restarted from the point it left off by setting the GO bit; but before it can be re-started, the STOP and HALT bits must both be cleared.

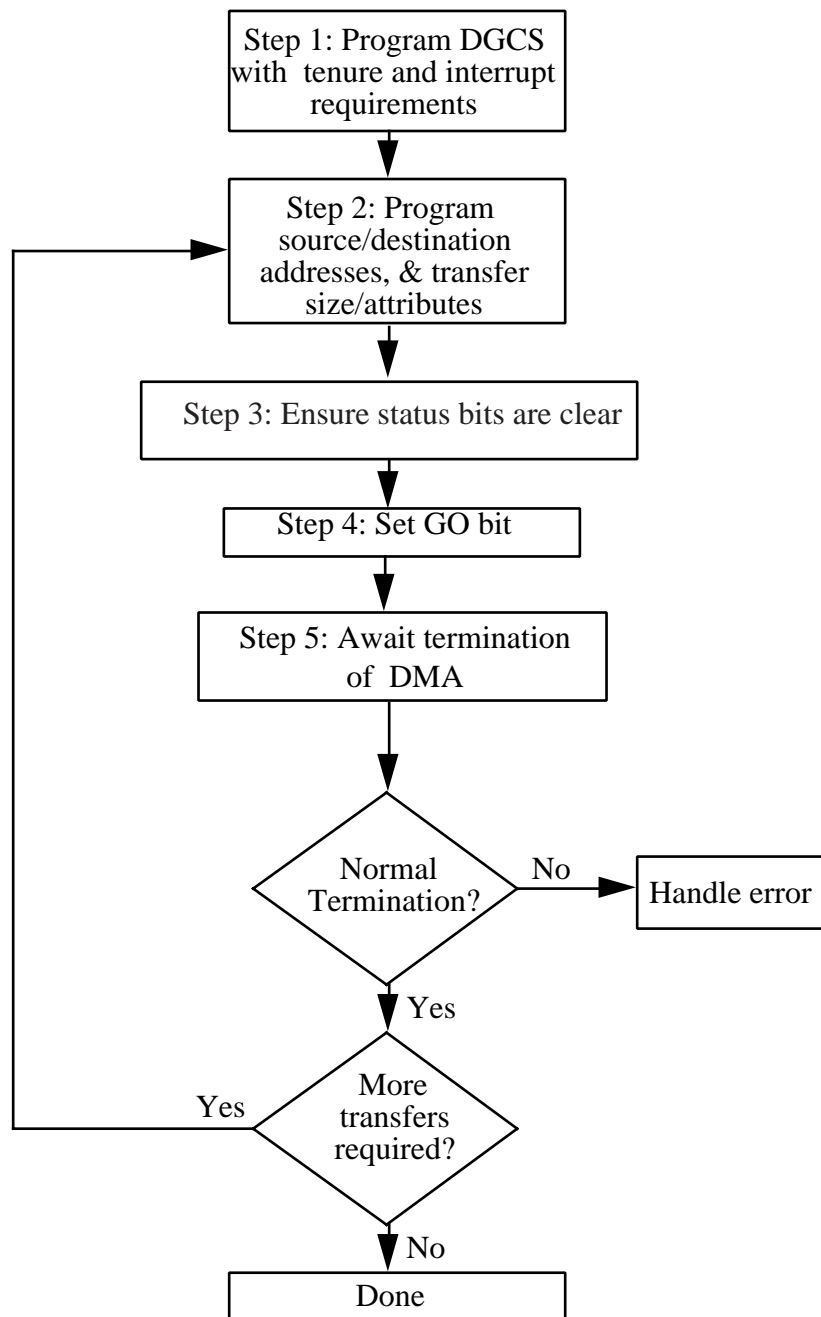
Regardless of how the DMA stops—whether normal, bus error or user interrupted—the DMA will indicate in the DGCS register why it stopped. The STOP and HALT bits get set in response to a stop or halt request. The DONE bit gets set when the DMA has successfully completed the DMA transfer, including all entries in the linked-list if operating in that mode. There are also three bits that are set in response to error conditions: LERR in the case of Target-Abort encountered on the PCI bus; VERR in the case of a bus error encountered on the VMEbus; and P_ERR in the case that the DMA has not been properly programmed (the DMA was started with the BM bit in the PCI_CSR register not enabled, or the DLA and DVA registers were not 64-bit aligned, (see “Source and Destination Addresses” on page 2-78). Before the DMA can be restarted, each of these status bits must be cleared.

When the DMA terminates, an interrupt may be generated to VMEbus or PCI bus. The user has control over which DMA termination conditions will cause the interrupt through the INT_STOP, INT_HALT, INT_DONE, INT_LERR, INT_VERR, and INT_P_ERR bits in the DGCS register.

2.8.2 Direct Mode Operation

When operated in direct mode, the Universe II DMA is set through manual register programming. Once the transfer described by the DVA, DLA, DTBC and DCTL registers has been completed, the DMA sits idle awaiting the next manual programming of the registers.

Figure 2.13 describes the steps involved in operating the DMA in direct mode.

**Figure 2.13 : Direct Mode DMA transfers**

In Step 1, the DGCS register is set up: the CHAIN bit is cleared, VON and VOFF are programmed with the appropriate values for controlling DMA VMEbus tenure, and the interrupt bits (INT_STOP, INT_HALT, INT_DONE, INT_LERR, INT_VERR, and INT_P_ERR) are programmed to enable generation of interrupts based on DMA termination events. DMA interrupt enable bits in the LINT_EN or VINT_EN bits should also be enabled as necessary (see “PCI Interrupt Generation” on page 2-63 and “VMEbus Interrupt Generation” on page 2-65 for details on generating interrupts).

In Step 2, the actual transfer is programmed into the DMA: source and destination start addresses into the DLA and DVA registers, transfer count into the DTBC register, and transfer width, direction and VMEbus address space into the DCTL register. These should be reprogrammed after each transfer.

In Step 3, ensure that if any status bits (DONE, STOP, HALT, LERR, VERR, or P_ERR) remain set from a previous transfer they are cleared. P_ERR must not be updated at the same time as Step 4, otherwise the P_ERR that may be generated by setting GO may be missed (see Step 4). These bits may be cleared as part of Step 1.

In Step 4, with the transfer programmed, the GO bit in DGCS must be set. If the DMA has been improperly programmed, either because the BM bit in the PCI_CSR has not been set to enable PCI bus mastership, or the source and destination start addresses are not aligned, then P_ERR will be asserted. Otherwise, the ACT bit will be set, and the DMA will then start transferring data, sharing ownership of the VMEbus with the PCI Target and Interrupt channels and the PCI bus with the VMEbus Slave Channel.

In Step 5, one waits for termination of the DMA transfers. The DMA will continue with the transfers until it:

- completes all transfers,
- is terminated early with the STOP_REQ bit, or
- encounters an error on the PCI bus or VMEbus.

Each of these conditions will cause the ACT bit to clear, and a corresponding status bit to be set in the DGCS register. If enabled in Step 1, an interrupt will also be generated. Once the software has set the GO bit, the software can monitor for DMA completion by either waiting for generation of an interrupt, or by polling the status bits. It is recommended that a background timer also be initiated to time-out the transfer. This will ensure that the DMA has not been hung up by a busy VMEbus, or other such system issues.

If an early termination is desired, perhaps because a higher priority operation is required, the STOP_REQ bit in the DGCS register can be set. This will stop all DMA operations on the source bus immediately, and set the STOP bit in the same register when the last piece of queued data in the DMA FIFO has been written to the destination bus. Attempting to terminate the transfer with the HALT_REQ bit will have no effect in direct mode operation since this bit only requests the DMA to stop between command packets in linked-list mode operation.

When the software has detected completion, it should verify the status bits in the DGCS register to see the reason for completion. If one of the error bits have been set it proceeds into an error handling routine (see “DMA Error Handling” on page 2-96). If the STOP bit was set, the software should take whatever actions were desired when it set the STOP_REQ bit. For example, if it was stopped for a higher priority transfer, it might record the DLA, DVA and DTBC registers, and then reprogram them with the higher priority transfer. When that has completed it can restore the DVA, DLA and DTBC registers to complete the remaining transfers.

If the DONE bit was set, it indicates that the DMA completed its requested transfer successfully, and if more transfers are required, the software can proceed to Step 2 to start a new transfer.

2.8.3 Linked-List Operation

Unlike direct mode, in which the DMA performs a single block of data at a time, linked-list mode allows the DMA to transfer a series of non-contiguous blocks of data without software intervention. Each entry in the linked-list is described by a command packet which parallels the DMA register layout. The data structure for each command packet is the same (see Figure 2.14 below), and contains all the necessary information to program the DMA address and control registers. It could be described in software as a record of eight 32-bit data elements. Four of the elements represent the four core registers required to define a DMA transfer: DCTL, DTBC, DVA, and DLA. A fifth element represents the DCPD register which points to the next command packet in the list. The least two significant bits of the DCPD element (the PROCESSED and NULL bits) provide status and control information for linked list processing.

The PROCESSED bit indicates whether a command packet has been PROCESSED or not. When the DMA processes the command packet and has successfully completed all transfers described by this packet, it sets the PROCESSED bit to “1” before reading in the next command packet in the list. This implies that the PROCESSED bit must be initially set for “0” by the user for it to be of use. This bit, when set to 1, indicates that this command packet has been disposed of by the DMA and its memory can be de-allocated or reused for another transfer description.

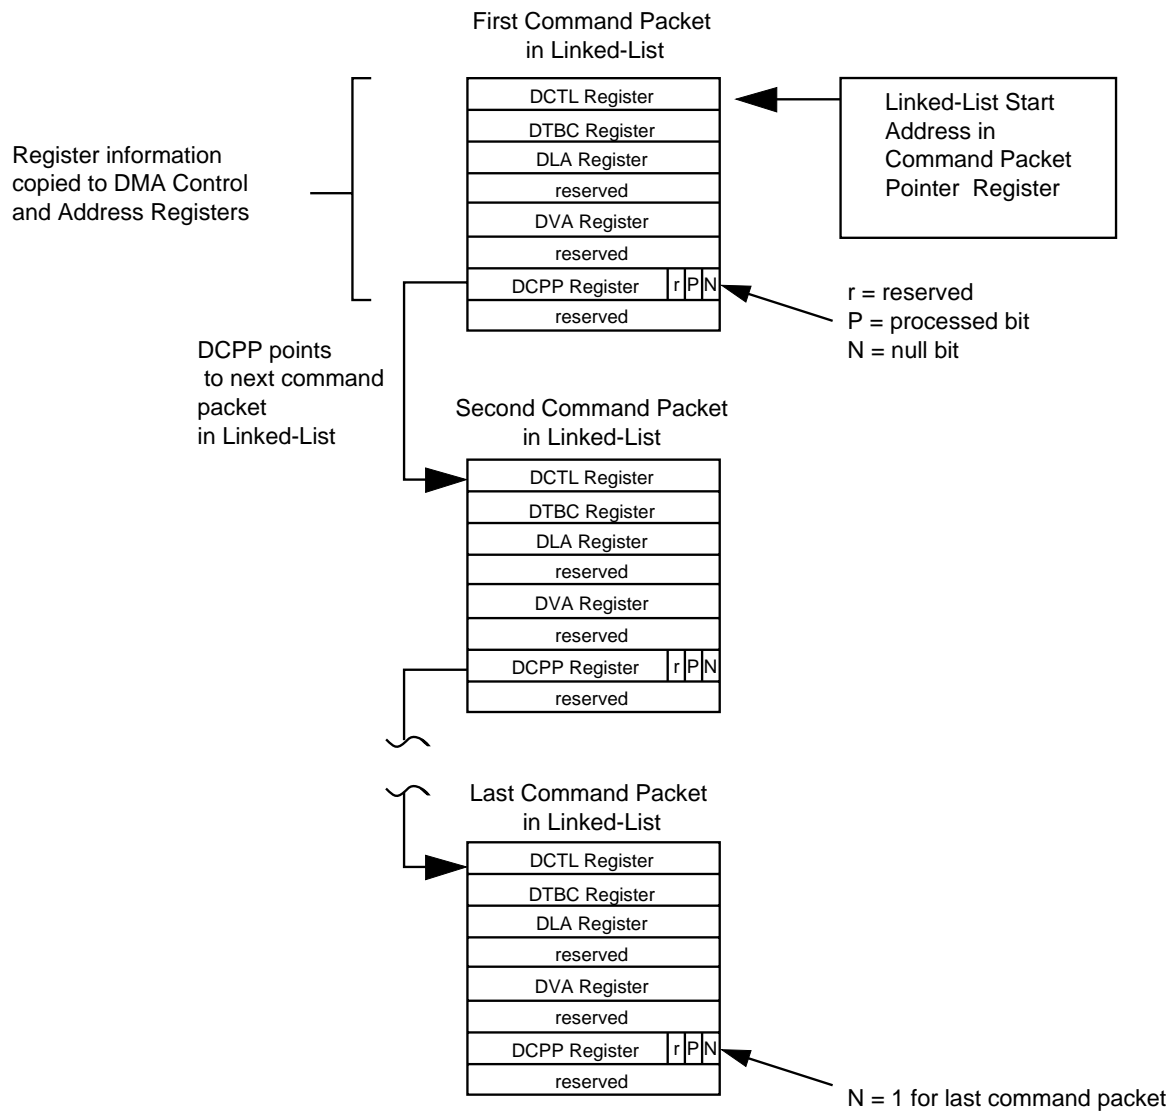


Figure 2.14 : Command Packet Structure and Linked List Operation

The NULL bit indicates the termination of the entire linked list. If the NULL bit is set to “0”, the DMA processes the next command packet pointed to by the command packet pointer. If the NULL bit is set to “1” then the address in the command packet pointer is considered invalid and the DMA stops at the completion of the transfer described by the current command packet.

Figure 2.15 outlines the steps in programming the DMA for linked-list operation.

In Step 1, the DGCS register is set up: the CHAIN bit is set, VON and VOFF are programmed with the appropriate values for controlling DMA VMEbus tenure, and the interrupt bits (INT_STOP, INT_HALT, INT_DONE, INT_LERR, INT_VERR, and INT_P_ERR) are programmed to enable generation of interrupts based on DMA termination events. DMA interrupt enable bits in the LINT_EN or VINT_EN bits should also be enabled as necessary (“PCI Interrupt Generation” on page 2-63 and “VMEbus Interrupt Generation” on page 2-65).

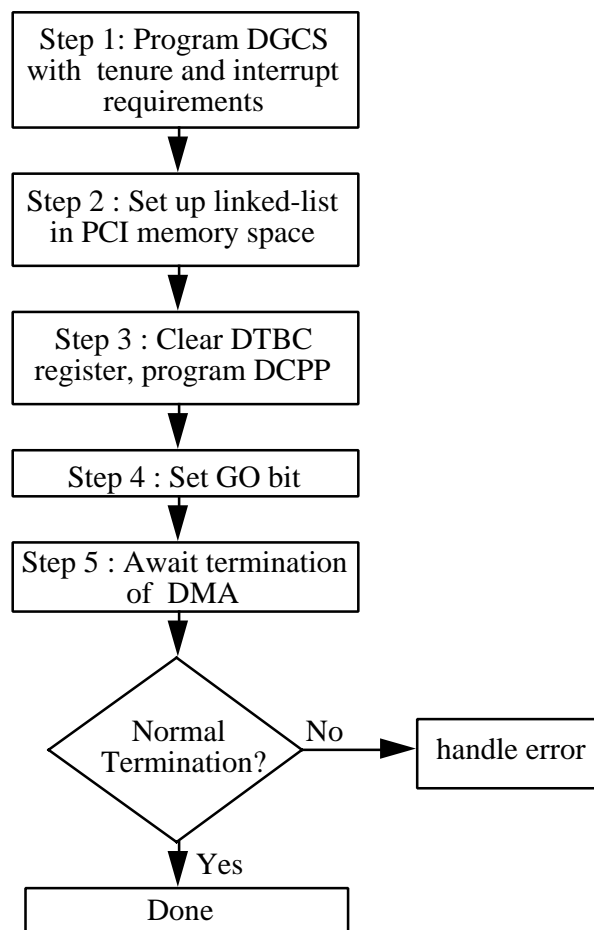


Figure 2.15 : DMA Linked List Operation

In Step 2, the linked-list structure is programmed with the required transfers. The actual structure may be set up at any time with command packet pointers pre-programmed and then only the remaining DMA transfer elements need be programmed later. One common way is to set up the command packets as a circular queue: each packet points to the next in the list, and the last points to the first. This allows continuous programming of the packets without having to set-up or tear down packets later.

Once the structure for the linked-list is established, the individual packets are programmed with the appropriate source and destination addresses, transfer sizes and attributes.

In Step 3, Clear the DTBC register and program the DCPD register to point to the first command packet in the list.



When using the DMA to perform linked-list transfers, it is important to ensure that the DTBC register contains a value of zero before setting the GO bit of the DGCS register. Otherwise, the DMA may not read the first command packet but instead perform a direct mode transfer based on the contents of the DCTL, DTBC, DLA, DVA and DGCS registers. After this direct mode transfer is completed, the PROCESSED bit of the first command packet is programmed with a value of 1 even though the packet was not actually processed. The DMA continues as expected with the next command packet.

In Step 4, to start the linked-list transfer, set the GO bit in the DGCS register. The DMA will first perform the transfers defined by the current contents of the DCTL, DTBC, DVA and DLA registers. Once that is complete it will then start the transfers defined by the linked-list pointed to in the DCPD register.

In Step 5, await and deal with termination of the DMA. Once the DMA channel is enabled, it processes the first command packet as specified by the DCPD register. The DMA transfer registers are programmed by information in the command packets and the DMA transfer steps along each command packet in sequence (see Figure 2.14 above). The DMA will terminate when it:

- processes a command packet with the NULL bit set indicating the last packet of the list,
- is stopped with the STOP_REQ bit in the DGCS register,
- is halted with the HALT_REQ bit in the DGCS register, or
- encounters an error on either the PCI bus or VMEbus.

Each of these conditions will cause the ACT bit to clear, and a corresponding status bit to be set in the DGCS register. If enabled in step 1, an interrupt will also be generated. Once the software has set the GO bit, the software can monitor for DMA completion by either waiting for generation of an interrupt, by polling the status bits in the DGCS register, or by polling the PROCESSED bits of the command packets. It is recommended that a background timer also be initiated to time-out the transfer. This will ensure that the DMA has not been hung up by a busy VMEbus, or other such system issues.

Linked-list operation can be halted by setting the HALT_REQ bit in the DGCS register (Table A.55). When the HALT_REQ bit is set, the DMA terminates when all transfers defined by the current command packet is complete. It then loads the next command packet into its registers. The HALT bit in the DGCS register is asserted, and the ACT bit in the DGCS register is cleared. The PROCESSED bit in the linked-list is set to “1” approximately 1 μ s after the HALT bit is set: therefore after a DMA halt the user should wait at least 1 μ s before checking the status of the PROCESSED bit.

The DMA can be restarted by clearing the HALT status bit and setting the GO bit if desired during the same register write. If the DMA is restarted, the ACT bit is set by the Universe II and execution continues as if no HALT had occurred: i.e., the Universe II processes the current command packet (see Figure 2.14 above).

In contrast to a halt, the DMA can also be immediately terminated through the STOP_REQ bit. This will stop all DMA operations on the source bus immediately, and set the STOP bit in the same register when the last piece of queued data in the DMA FIFO has been written to the destination bus.

Once stopped, the DVA, DLA and DTBC registers contain values indicating the next addresses to read/write and the number of bytes remaining in the transfer. Clearing the STOP bit and setting the GO bit will cause the DMA to start up again from where it left off, including continuing with subsequent command packets in the list.

If the DMA is being stopped to insert a high priority DMA transfer, the remaining portion of the DMA transfer may be stored as a new command packet inserted at the top of the linked list. A new command packet with the attributes of the high priority transfer is then placed before that one in the list. Now the linked list is set up with the high priority packet first, followed by the remainder of the interrupted packet, followed in turn by the rest of the linked list. Finally, the DTBC register is cleared and the DCPD programmed with a pointer to the top of the list where the high priority command packet has been placed. When the GO bit is set (after clearing the STOP status bit in the DGCS register), the DMA will perform the transfers in the order set in the linked list. For more details on updating the linked list see “Linked List Updating” on page 2-91.

DMA transfers continue until the DMA encounters a command packet with the NULL bit set to “1”, indicating that the last packet has been reached. At this point, the DMA stops, the DONE bit is set, and the ACT flag is cleared. As it completes the transfers indicated by each command packet, the DMA sets the PROCESSED bit in that command packet before reading in the next command packet and processing its contents.

2.8.3.1 Linked List Updating

The Universe II provides a mechanism which allows the linked list to be updated with additional linked list entries without halting or stopping the DMA. This takes place through the use of a semaphore in the device: the UPDATE bit in the D_LLUE register (Table A.56). This bit is meant to ensure that the DMA does not read a command packet into the DMA registers while the command packet (outside the Universe II) is being updated. This semaphore does not prevent external masters from updating the DMA registers.

Adding to a linked list begins by writing a “1” to the UPDATE bit. The DMA checks this bit before proceeding to the next command packet. If the UPDATE bit is “0”, then the DMA locks the UPDATE bit against writes and proceeds to the next command packet. If the UPDATE bit is “1”, then the DMA waits until the bit is cleared before proceeding to the next command packet. Therefore, setting the UPDATE bit is a means of stalling the DMA at command packet boundaries while local logic updates the linked list.

In order to ensure that the DMA is not currently reading a command packet during updates, the update logic must write a “1” to the UPDATE bit and read a value back. If a “0” is read back from the UPDATE bit, then the DMA is currently reading a command packet and has locked the UPDATE bit against writes. If a “1” is read back from the UPDATE bit, then the DMA is idle or processing a transaction and command packets can be updated. If the DMA attempts to proceed to the next command packet during the update, it will encounter the set UPDATE bit and wait until the bit is cleared.

If a set of linked command packets has already been created with empty packets at the end of new transfers, adding to the end of the current linked list takes the following procedure:

1. Get UPDATE valid (write “1”, read back “1”),
2. Program attributes for new transfer in next available packet in list.
3. Change “null” pointer (on previous tail of linked list),
4. Release update (clear the UPDATE bit).

After updating the linked list, the DMA controller will be in three possible conditions:

1. It may be active and working its way through the linked list. In this case, no further steps are required.
2. The DMA may be idle (done) because it reached the final command packet. If a full set of linked command packets had already been created ahead of time, then the DCPD register would point to the most recently programmed command packet, and the DTBC register would be zero. The DMA can be started on the new packet by simply clearing the DONE bit and setting the GO bit in the DGCS register. If a

set of command packets had not been created ahead of time, the DCPD register may not be programmed to any valid packet, and will need programming to the newly programmed packet.

3. The DMA has encountered an error. In this circumstance, see “DMA Error Handling” on page 2-96 for how to handle DMA errors.

Operation may be considerably simplified by ensuring that sufficient command packets have been created during system initialization, probably in a circular queue. In this fashion, when a new entry is added to the list, it is simply a matter of programming the next available entry in the list with the new transfer attributes and changing the previously last packet's NULL bit to zero. The DCPD register will be guaranteed to point to a valid command packet, so upon updating the list, both cases 1 and 2 above can be covered by clearing the DONE bit and setting the GO bit. This will have no effect for case 1 since the DMA is still active, and will restart the DMA for case 2.

If an error has been encountered by the DMA (case 3), setting the GO bit and clearing the DONE bit will not be sufficient to restart the DMA—the error bits in the DGCS register will also have to be cleared before operation can continue.

2.8.4 FIFO Operation and Bus Ownership

The DMA uses a 256-byte FIFO. (The DMA FIFO is 64 bits wide). This supports high performance DMA transfers. In general, the DMA reads data from the source, and stores it as transactions in the FIFO. On the destination side, the DMA requests ownership of the master and once granted begins transfers. Transfers stop on the source side when the FIFO fills, and on the destination side when the FIFO empties.

2.8.4.1 PCI to VMEbus Transfers

PCI to VMEbus transfers involve the Universe II reading from the PCI bus and writing to the VMEbus.

The PCI bus is requested for the current read once 128 bytes are available in the DMAFIFO. The DMA Channel fills the DMAFIFO using PCI read transactions with each transaction broken at address boundaries determined by the programmed PCI aligned burst size (PABS field in the MAST_CTL register, Table A.81). This ensures that the DMA makes optimal use of the PCI bus by always generating bursts of 32, 64 or 128 bytes with zero wait states.

The DMA packs read data into the DMAFIFO to the full 64-bit width of the FIFO, independent of the width of the PCI bus, or the data width of the ensuing VMEbus transaction. The PCI read transactions continue until either the DMA has completed the full programmed transfer, or there is insufficient room available in the DMAFIFO for a full transaction. The available space required for another burst read transaction is again 128 bytes. Since the VMEbus is typically much slower than the PCI bus, the DMAFIFO may fill frequently during

PCI to VMEbus transfers, though the depth of the FIFO helps to minimize this. When the DMAFIFO fills, the PCI bus is free for other transactions (for example, between other devices on the bus or possibly for use by the Universe II's VMEbus Slave Channel). The DMA only resumes read transactions on the PCI bus when the DMAFIFO has space for another aligned burst size transaction.



Caution: The DMA may prefetch extra read data from the external PCI target. This implies that the DMA should only be used with memory on the PCI bus which has no adverse side-effects when prefetched. The Universe II will prefetch up to the aligned address boundary defined in the PABS field of the MASC_CTL register. On the VMEbus, the actual programmed number of bytes in the DTBC register will be written. Prefetching can be avoided by programming the DMA for transfers that terminate at the PABS boundary. If further data is required beyond the boundary, but before the next boundary, the DTBC register may be programmed to eight byte transfers. The DMA will fetch the full eight bytes, and nothing more. Programming the DTBC to less than eight bytes will still result in eight bytes fetched from PCI.

The DMA requests ownership of the Universe II's VMEbus Master Interface once 64 bytes of data have been queued in the DMAFIFO (see “VMEbus Requester” on page 2-6 on how the VMEbus Master Interface is shared between the DMA, the PCI Target Channel, and the Interrupt Channel). The Universe II maintains ownership of the Master Interface until:

- the DMAFIFO is empty,
- the DMA block is complete,
- the DMA is stopped,
- a linked list is halted,
- the DMA encounters an error, or
- the DMA VMEbus tenure limit (VON in the DGCS register).

The DMA can be programmed to limit its VMEbus tenure to fixed block sizes using the VON field in the DGCS register (Table A.55). With VON enabled, the DMA will relinquish ownership of the Master Interface at defined address boundaries. See “DMA VMEbus Ownership” on page 2-81.

To further control the DMA's VMEbus ownership, the VOFF timer in the DGCS register can be used to program the DMA to remain off the VMEbus for a specified period when VMEbus tenure is relinquished. See “DMA VMEbus Ownership” on page 2-81.

The DMA Channel unpacks the 64-bit data queued in the DMAFIFO to whatever the programmed transfer width is on the VMEbus (e.g. D16, D32, or D64). The VMEbus Master Interface delivers the data in the DMAFIFO according to the VMEbus cycle type programmed into the DCTL register (Table A.50, see “DMA Controller” on page 2-77). The DMA provides data to the VMEbus until:

- the DMAFIFO empties, or
- the DMA VMEbus Tenure Byte Count (VON in the DMA_GCSR register, Table A.55) expires.

If the DMAFIFO empties, transfers on the VMEbus stop and, if the cycle being generated is a block transfer, then the block is terminated (AS* negated) and VMEbus ownership is relinquished by the DMA. The DMA does not re-request VMEbus ownership until another eight entries are queued in the DMAFIFO, or the DMA Channel has completed the current Transfer Block on the PCI bus (see “VMEbus Release” on page 2-8).

PCI bus transactions are the full width of the PCI data bus with appropriate byte lanes enabled. The maximum VMEbus data width is programmable to 8, 16, 32, or 64 bits. Byte transfers can be only of type DO8 (EO). Because the PCI bus has a more flexible byte lane enabling scheme than the VMEbus, the Universe II may be required to generate a variety of VMEbus transaction types to handle the byte resolution of the starting and ending addresses (see “Data Transfer” on page 2-39).

2.8.4.2 VMEbus to PCI Transfers

VMEbus to PCI transfers involve the Universe II reading from the VMEbus and writing to the PCI bus.

With DMA transfers in this direction, the DMA Channel begins to queue data in the DMAFIFO as soon as there is room for 64 bytes in the DMAFIFO. When this watermark is reached, the DMA will request the VMEbus (through the VMEbus Master Interface) and begin reading data from the VMEbus. The Universe II maintains VMEbus ownership until:

- the DMAFIFO is full,
- the DMA block is complete,
- the DMA is stopped,
- a linked list is halted,
- the DMA encounters an error, or
- the VMEbus tenure limit is reached (VON in the DGCS register).

The DMA can be programmed to limit its VMEbus tenure to fixed block sizes using the VON field in the DGCS register (Table A.55). With VON enabled, the DMA will relinquish ownership of the Master Interface at defined address boundaries. See “DMA VMEbus Ownership” on page 2-81.

To further control the DMA’s VMEbus ownership, the VOFF timer in the DGCS register can be used to program the DMA to remain off the VMEbus for a specified period when VMEbus tenure is relinquished. See “DMA VMEbus Ownership” on page 2-81.

Entries in the DMAFIFO are delivered to the PCI bus as PCI write transactions as soon as there are 128 bytes available in the DMAFIFO. If the PCI bus responds too slowly, the DMAFIFO runs the risk of filling before write transactions can begin at the PCI Master Interface. Once the DMAFIFO reaches a “nearly full” state (corresponding to three entries remaining) the DMA requests that the VMEbus Master Interface complete its pending operations and stop. The pending read operations typically fill the DMAFIFO. Once the pending VMEbus reads are completed (or the VON timer expires), the DMA relinquishes VMEbus ownership and only re-requests the VMEbus Master Interface once 64 bytes again become available in the DMAFIFO. If the bus was released due to encountering a VON boundary, the bus is not re-requested until the VOFF timer expires.

PCI bus transactions are the full width of the PCI data bus with appropriate byte lanes enabled. The maximum VMEbus data width is programmable to 8, 16, 32, or 64 bits. Byte transfers can be only of type DO8 (EO). Because the PCI bus has a more flexible byte lane enabling scheme than the VMEbus, the Universe II may be required to generate a variety of VMEbus transaction types to handle the byte resolution of the starting and ending addresses (see “Universe II as PCI Target” on page 2-38).

2.8.5 DMA Interrupts

The Interrupt Channel in the Universe II handles a single interrupt sourced from the DMA Channel which it routes to either the VMEbus or PCI bus via the DMA bits in the LINT_EN and VINT_EN registers. There are six internal DMA sources of interrupts and these are all routed to this single interrupt. Each of these six sources may be individually enabled, and are listed in Table 2.17 below. Setting the enable bit enables the corresponding interrupt source.

Table 2.17 : DMA Interrupt Sources and Enable Bits

Interrupt Source	Enable Bit
Stop Request	INT_STOP
Halt Request	INT_HALT
DMA Completion	INT_DONE
PCI Target-Abort or Master-Abort	INT_LERR
VMEbus Error	INT_VERR

Table 2.17 : DMA Interrupt Sources and Enable Bits

Interrupt Source	Enable Bit
Protocol Error	INT_M_ERR

Once an enabled DMA interrupt has occurred, regardless of whether the LINT_EN or VINT_EN enable bits have been set, the corresponding DMA bit in the LINT_STAT (Table A.58) and VINT_STAT (Table A.62) registers are set. Each one must be cleared independently. Clearing one does not clear the other. See “Interrupt Handling” on page 2-68.

2.8.6 Interactions with Other Channels

This section describes the impact that the PCI Bus Target Channel and the VMEbus Slave Channel may have on the DMA Channel.

The Universe II does not apply PCI 2.1 transaction ordering requirements to the DMA Controller. That is, reads and writes through the DMA Controller can occur independently of the other channels.

ADOH cycles and RMW cycles through the VMEbus Slave Channel do impact on the DMA Channel. Once an external VMEbus master locks the PCI bus, the DMA Controller will not perform transfers on the PCI bus until the Universe II is unlocked (see “VMEbus Lock Commands (ADOH Cycles)” on page 2-18). When an external VMEbus Master begins a RMW cycle, at some point a read cycle will appear on the PCI bus. During the time between when the read cycle occurs on the PCI bus and when the associated write cycle occurs on the PCI bus, no DMA transfers will occur on the PCI bus (see “VMEbus Read-Modify-Write Cycles (RMW Cycles)” on page 2-19).

If the PCI Target Channel locks the VMEbus using VOWN, no DMA transfers will take place on the VMEbus (see “Using the VOWN bit” on page 2-48).

2.8.7 DMA Error Handling

This section describes how the Universe II responds to errors involving the DMA, and how the user can recover from them. As described below, the software source of a DMA error is a protocol, and the hardware source of a DMA error is a VMEbus error, or PCI bus Target-Abort or Master-Abort.

2.8.7.1 DMA Software Response to Error

While the DMA is operating normally, the ACT bit in the DGCS register will be set (Table A.55). Once the DMA has terminated, it will clear this bit, and set one of six status bits in the same register. The DONE bit will be set if the DMA completed all its programmed operations normally. If the user interrupted the DMA, either the STOP or HALT bits will be set. If an error has occurred, one of the remaining three bits, LERR, VERR, or P_ERR, will be set. All six forms of DMA terminations can be optionally set to generate a DMA interrupt by setting the appropriate enable bit in the DGCS register (see “DMA Interrupts” on page 2-95).

- LERR is set if the DMA encounters an error on the PCI bus: either a Master-Abort or Target-Abort. Bits in the PCI_CSR register will indicate which of these conditions caused the error.
- VERR is set if the DMA encounters a bus error on the VMEbus. This will be exclusively through a detected assertion of BERR* during a DMA cycle.
- P_ERR is set if the GO bit in the DGCS register is set to start the DMA, and the DMA has been improperly programmed either because the BM bit in the PCI_CSR disables PCI bus mastership, or the source and destination start addresses are not aligned (see “Source and Destination Addresses” on page 2-78).

Whether the error occurs on the destination or source bus, the DMA_CTL register contains the attributes relevant to the particular DMA transaction. The DTBC register provides the number of bytes remaining to transfer on the PCI side. The DTBC register contains valid values after an error. The DLA and DVA registers should not be used for error recovery.

2.8.7.2 DMA Hardware Response to Error

This section describes how transfers proceed following a bus error, and how interrupts can be generated following DMA error conditions.

When the error condition (VMEbus Error, Target-Abort, or Master-Abort) occurs on the source bus while the DMA is reading from the source bus, the DMA stops reading from the source bus. Any data previously queued within the DMAFIFO is written to the destination bus. Once the DMAFIFO empties, the error status bit is set and the DMA generates an interrupt (if enabled by INT_LERR or INT_VERR in the DGCS register—see “DMA Interrupts” on page 2-95).

When the error condition (VMEbus Error, Target-Abort, or Master-Abort) occurs on the destination bus while the DMA is writing data to the destination bus, the DMA stops writing to the destination bus, and it also stops reading from the source bus. The error bit in the DGCS register is set and an interrupt asserted (if enabled).

Interrupt Generation During Bus Errors

To generate an interrupt from a DMA error, there are two bits in the DGCS register (and one bit each in the VINT_EN and LINT_EN registers). In the DGCS register the INT_LERR bit enables the DMA to generate an interrupt to the Interrupt Channel after encountering an error on the PCI bus. The INT_VERR enables the DMA to generate an interrupt to the Interrupt Channel upon encountering an error on the VMEbus. Upon reaching the Interrupt Channel, all DMA interrupts can be routed to either the PCI bus or VMEbus by setting the appropriate bit in the enable registers. All DMA sources of interrupts (Done, Stopped, Halted, VMEbus Error, and PCI Error) constitute a single interrupt into the Interrupt Channel.

2.8.7.3 Resuming DMA Transfers

When a DMA error occurs (on the source or destination bus), the user should read the status bits and determine the source of the error. If it is possible to resume the transfer, the transfer should be resumed at the address that was in place up to 256 bytes from the current byte count. The original addresses (i.e., DLA and DVA) are required in order to resume the transfer at the appropriate location. However, the values in the DLA and the DVA registers should not be used to reprogram the DMA, because they are not valid once the DMA begins. In direct mode, it is the user's responsibility to record the original state of the DVA and DLA registers for error recovery. In Linked-List mode, the user can refer to the current Command Packet stored on the PCI bus (whose location is specified by the DCPD register) for the location of the DVA and DLA information.

The DTBC register contains the number of bytes remaining to transfer on the source side. The Universe II does not store a count of bytes to transfer on the destination side. If the error occurred on the source side, then the location of the error is simply the latest source address plus the byte count. If the error occurred on the destination side, then one cannot infer specifically where the error occurred, because the byte count only refers to the number of data queued from the source, not what has been written to the destination. In this case, the error will have occurred up to 256 bytes before: the original address plus the byte count.

Given this background, the following procedure may be implemented to recover from errors.

1. Read the value contained in the DTBC register.
2. Read the record of the DVA and DLA that is stored on the PCI bus or elsewhere (not the value stored in the Universe II registers of the same name, see above).
3. If the difference between the value contained in the DTBC register and the original value is less than 256 bytes (the FIFO depth of the Universe II), reprogram all the DMA registers with their original values.
4. If the difference between the value contained in the DTBC register and the original value is greater than 256 bytes (the FIFO depth of the Universe II), add 256 (the FIFO depth of the Universe II) to the value contained in the DTBC register.
5. Add the difference between the original value in the DTBC and the new value in the DTBC register to the original value in the DLA register.
6. Add the difference between the original value in the DTBC and the new value in the DTBC register to the original value in the DVA register.
7. Clear the status flags.
8. Restart the DMA (see “DMA Initiation” on page 2-80).

2.9 Registers

This section is organized as follows:

- “Overview of Universe II Registers” on page 2-100,
- “Register Access from the PCI Bus” on page 2-101,
- “Register Access from the VMEbus” on page 2-104,
- “Mailbox Registers” on page 2-108, and
- “Semaphores” on page 2-109

2.9.1 Overview of Universe II Registers

The Universe II Control and Status Registers (UCSR) occupy 4 Kbytes of internal memory. This 4 Kbytes is logically divided into three groups (see Figure 2.16 below):

- PCI Configuration Space (PCICS),
- Universe II Device Specific Registers (UDSR), and
- VMEbus Control and Status Registers (VCSR).

The Universe II registers are little-endian.

The access mechanisms for the UCSR are different depending upon whether the register space is accessed from the PCI bus or VMEbus. Register access from the PCI bus and VMEbus is discussed below.

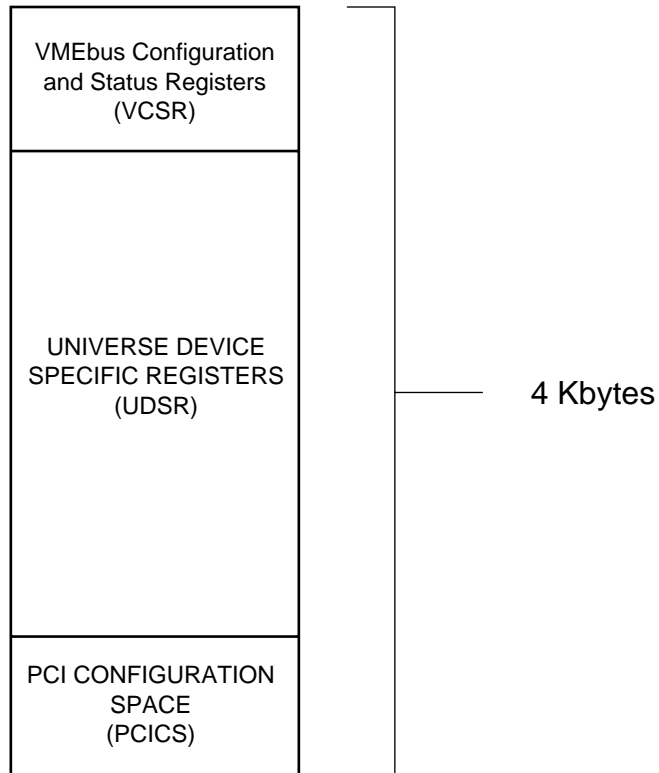


Figure 2.16 : Universe II Control and Status Register Space

2.9.2 Register Access from the PCI Bus

There are two mechanisms to access the UCSR space from the PCI bus: through Configuration space or through PCI Memory or I/O space (Table A.6).

2.9.2.1 PCI Configuration Access

When the UCSR space is accessed as Configuration space, it means that the access is externally decoded and the Universe II is notified via IDSEL (much like a standard chip select signal). Since the register location is encoded by a 6-bit register number (a value used to index a 32-bit chunk of Configuration space), only the lower 256 bytes of the UCSR can be accessed as Configuration space (this corresponds to the PCICS in the UCSR space, see Figure 2.17 on page 2-102). Thus, only the PCI configuration registers are accessible through PCI Configuration cycles.

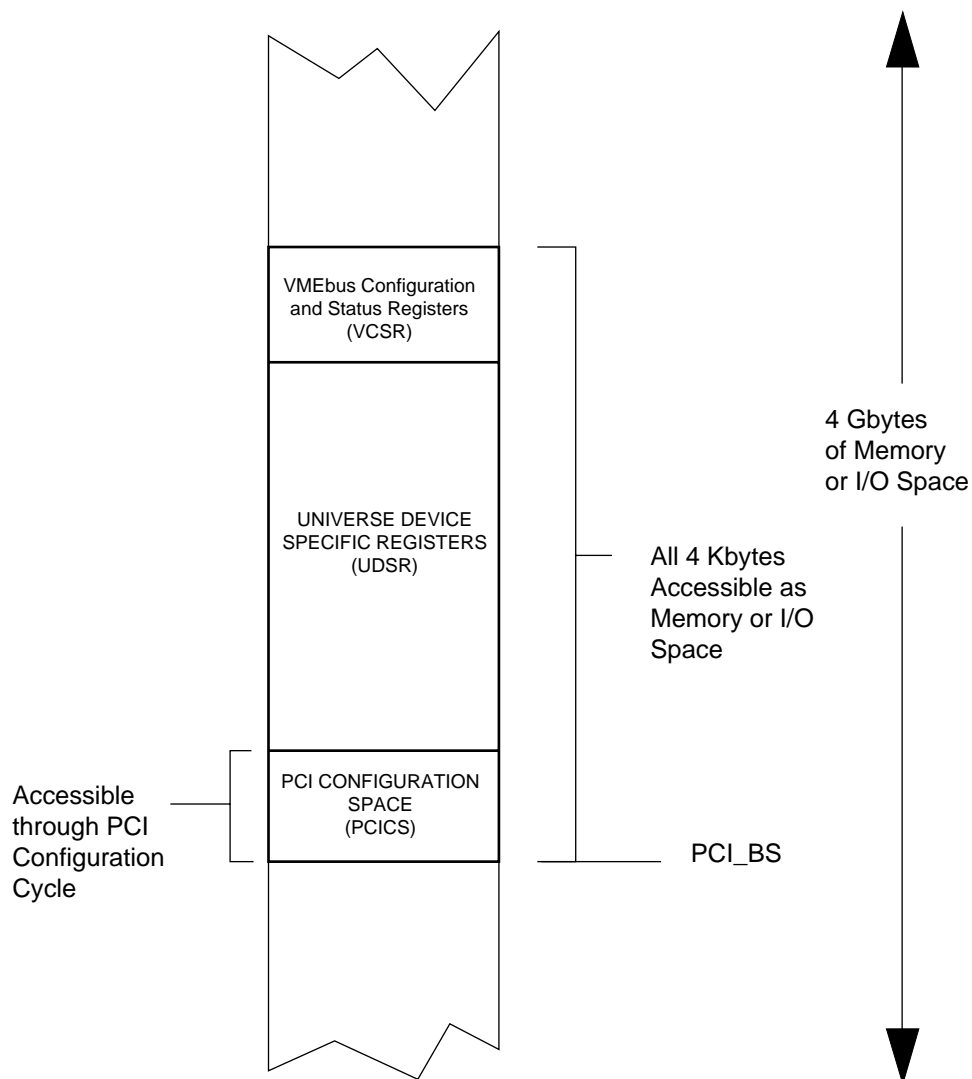


Figure 2.17 : PCI Bus Access to UCSR as Memory or I/O Space

2.9.2.2 Memory or I/O Access

Exactly two 4-Kbyte ranges of addresses in PCI Memory space and/or PCI I/O space can be dedicated to the Universe II registers. The Universe II has two programmable registers (PCI_BS0 and PCI_BS1) that each specify the base address and address space for PCI access to the Universe II's registers. The PCI_BSx registers can be programmed through PCI Configuration space or through a VMEbus access, to make the Universe II registers available anywhere in the 32-bit Memory space and in I/O space (as offsets of the BS[31:12] field in PCIBSx).

The SPACE bit of the PCI_BSx registers specifies whether the address lies in Memory space or I/O space. The SPACE bit of these two registers are read-only. There is a power-up option that determines the value of the SPACE bit of the PCI_BSx registers. At power-up the SPACE bit of the PCI_BS1 register is the negation of the SPACE bit of the PCI_BS0 register.

- When the VA[1] pin is sampled low at power-up, the PCI_BS0 register's SPACE bit is set to "1", which signifies I/O space, and the PCI_BS1 register's SPACE bit is set to "0", which signifies Memory space.
- When VA[1] is sampled high at power-up, the PCI_BS0 register's SPACE register's bit is set to "0", which signifies Memory space, and the PCI_BS1 register's SPACE bit is set to "1", which signifies I/O space.

Universe II registers are not prefetchable. The Universe II does not accept burst writes to its registers.

Eliciting Conditions of Target-Retry

Attempts to access UCSR space from the PCI bus will be retried by the Universe II under the following conditions:

- While UCSR space is being accessed by a VMEbus master, PCI masters will be retried.
- If a VMEbus master is performing a RMW access to the UCSRs then PCI attempts to access the UCSR space will result in a Target-Retry until AS* is negated.
- If the Universe II registers are accessed through an ADOH cycle from the VMEbus, any PCI attempt to access the UCSRs will be retried until BBSY* is negated.

2.9.2.3 Locking the Register Block from the PCI bus

The Universe II registers can be locked by a PCI master by using a PCI locked transaction. When an external PCI master locks the register block of the Universe II, an access to the register block from the VMEbus will not terminate with the assertion of DTACK* until the register block is unlocked. Hence a prolonged lock of the register block by a PCI resource may cause the VMEbus to timeout with a BERR*.

2.9.3 Register Access from the VMEbus

There are two mechanisms to access the UCSR space from the VMEbus. One method uses a VMEbus Register Access Image (VRAI) which allows the user to put the UCSR in an A16, A24 or A32 address space. The VRAI approach is useful in systems not implementing CR/CSR space as defined in the VME64 specification. The other way to access the UCSR is as CR/CSR space, where each slot in the VMEbus system is assigned 512 Kbytes of CR/CSR space.

Each method is discussed below.

2.9.3.1 VMEbus Register Access Image (VRAI)

The VMEbus register access image is defined by the following register fields:

Table 2.18 : Programming the VMEbus Register Access Image

Field	Register Bits	Description
address space	VAS in Table A.103	one of A16, A24, A32
base address	BS[31:12] in Table A.104	lowest address in the 4Kbyte slave image
slave image enable	EN in Table A.103	enables VMEbus register access image
mode	SUPER in Table A.103	Supervisor and/or Non-Privileged
type	PGM in Table A.103	Program and/or Data

The VMEbus Register Access Image occupies 4 Kbytes in A16, A24 or A32 space (depending upon the programming of the address space described in Table 2.18 above, see Figure 2.18 below). All registers are accessed as address offsets from the VRAI base address programmed in the VRAI_BS register (Table A.104). The image can be enabled or disabled using the EN bit in the VRAI_CTL register (Table A.103).

Note that the VRAI base address can be configured as a power-up option (see “Power-up Option Descriptions” on page 2-117).

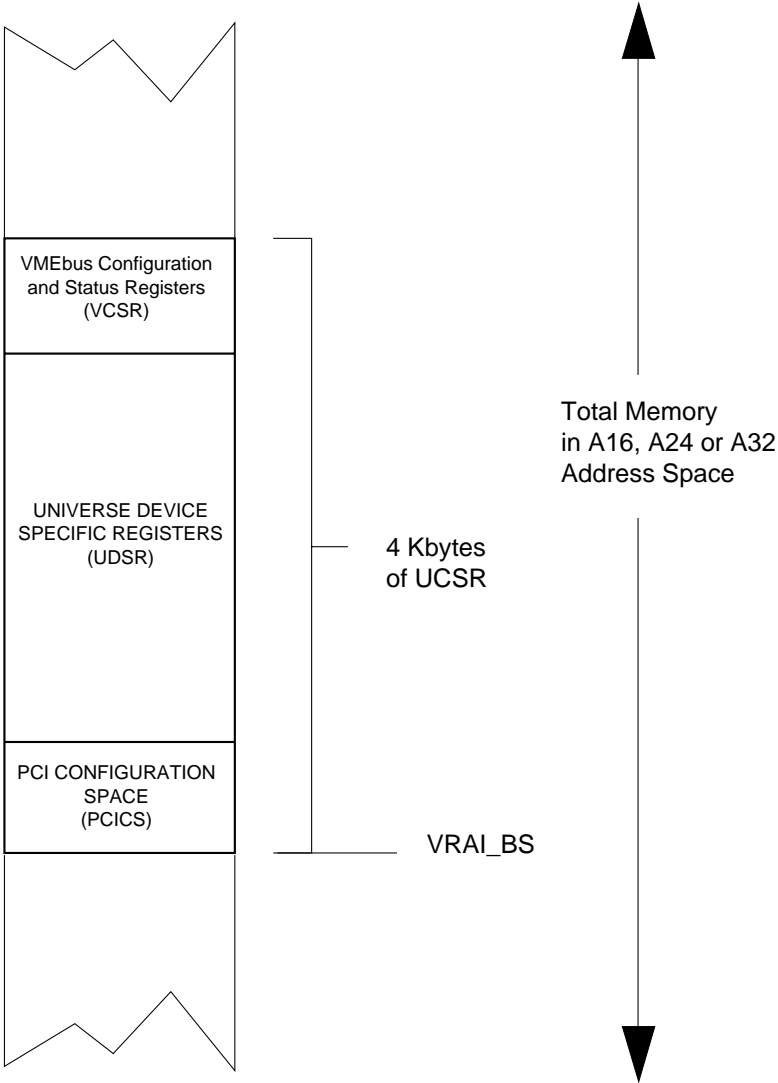


Figure 2.18 : UCSR Access from the VMEbus Register Access Image

2.9.3.2 CR/CSR Accesses

The VME64 specification assigns a total of 16 Mbytes of CR/CSR space for the entire VMEbus system. The CR/CSR image is enabled with the EN bit in the VCSR_CTL register (Table A.105). This 16 Mbytes is broken up into 512 Kbytes per slot for a total of 32 slots. The first 512 Kbyte block is reserved for use by the Auto-ID mechanism. The UCSR space occupies the upper 4 Kbytes of the 512 Kbytes available for its slot position (see Figure 2.19 below). The base address of the CR/CSR space allocated to the Universe II's slot is programmed in the VCSR_BS register (Table A.127). For CSRs not supported in the Universe II and for CR accesses, the LAS field in the VCSR_CTL register specifies the PCI bus command that is generated when the cycle is mapped to the PCI bus. There is also a translation offset added to the 24-bit VMEbus address to produce a 32-bit PCI bus address (programmed in the VCSR_TO register, Table A.106).

Note that the registers in the UCSR space are located as address offsets from VCSR_BS. These offsets are different from those used in the VRAI mechanisms, where the first register in the UCSR has address offset of zero (see Table A.1 in Appendix A). When accessing the UCSR in CR/CSR space, the first register will have an address offset of 508 Kbytes (512 Kbytes minus 4 Kbytes). A simple approach for determining the register offset when accessing the UCSR in CR/CSR space is to add 508 Kbytes (0x7F000) to the address offsets given in Table A.1.

2.9.3.3 RMW and ADOH Register Access Cycles

The Universe II supports RMW and ADOH accesses to its registers.

A read-modify-write (RMW) cycle allows a VMEbus master to read from a VMEbus slave and then write to the same resource without relinquishing VMEbus tenure between the two operations. The Universe II accepts RMW cycles to any of its registers. This prevents an external PCI Master from accessing the registers of the Universe II until VMEbus AS* is asserted. This is useful if a single RMW access to the ADHO is required.

If a sequence of accesses to the Universe registers must be performed without intervening PCI access to UCSR is required, then the VMEbus master should lock the Universe II through the use of ADOH. This prevents an external PCI Master from accessing the registers of the Universe II until VMEbus BBSY* is negated. It also prevents other VMEbus masters from accessing the Universe II registers.




Figure 2.19 : UCSR Access in VMEbus CR/CSR Space

2.9.4 Mailbox Registers

The Universe II has four 32-bit mailbox registers which provide an additional communication path between the VMEbus and the PCI bus (see Table A.75 to Table A.78). The mailboxes support read and write accesses from either bus, and may be enabled to generate interrupts on either bus when they are written to. The mailboxes are accessible from the same address spaces and in the same manner as the other Universe II registers, as described above.

Mailbox registers are useful for the communication of concise command, status, and parameter data. The specific uses of mailboxes depend on the application. For example, they can be used when a master on one bus needs to pass information (a message) on the other bus, without knowing where the information should be stored in the other bus's address space. Or they can be used to store the address of a longer message written by the processor on one bus to the address space on the other bus, through the Universe II. They can also be used to initiate larger transfers through the FIFO, in a user-defined manner.

Often users will enable and map mailbox interrupts, so that when the processor writes to a mailbox from one bus, the Universe II will interrupt the opposite bus. The interrupt service routine on the opposite bus would then cause a read from this same mailbox.

Reading a mailbox cannot automatically trigger an interrupt. However, a similar effect can be achieved by reading the mailbox and then triggering an interrupt through hardware or software. Or one may use a "polling" approach, where one designates a bit in a mailbox register to indicate whether one has read from the mailbox.

For details on how the mailbox interrupts are enabled and mapped, see "Interrupt Handling" on page 2-68 and "Mailbox Register Access Interrupts" on page 2-75.

Applications will sometimes designate two mailboxes on one interface as being read/write from the PCI bus, and read-only from the VMEbus, and the two other mailboxes as read/write from the VMEbus and read-only from the PCI bus. This eliminates the need to implement locking. The Universe II provides semaphores which can be also be used to synchronise access to the mailboxes. Semaphores are described in the next section.

2.9.5 Semaphores

The Universe II has two general-purpose semaphore registers each containing four semaphores. The registers are SEMA0 (Table A.79) and SEMA1 (Table A.80). To gain ownership of a semaphore, a process writes a logic one to the semaphore bit and a unique pattern to the associated tag field. If a subsequent read of the tag field returns the same pattern, the process can consider itself the owner of the semaphore. A process writes a value of 0 to the semaphore to release it.

When a semaphore bit is a value of 1, the associated tag field cannot be updated. Only when a semaphore is a value of 0 can the associated tag field be updated.

These semaphores allow the user to share resources in the system. While the Universe II provides the semaphores, it is up to the user to determine access to which part of the system will be controlled by semaphores, and to design the system to enforce these rules.

An example of a use of the semaphore involves gating access to the Special Cycle Generator (page 2-45). It may be necessary to ensure that while one process uses the Special Cycle Generator on an address, no other process accesses this address. Before performing a Special Cycle, a process would be required to obtain the semaphore. This process would hold the semaphore until the Special Cycle completes. A separate process that intends to modify the same address would need to obtain the semaphore before proceeding (it need not verify the state of the SCYC[1:0] bit). This mechanism requires that processes know which addresses might be accessed through the Special Cycle Generator.

2.10 Utility Functions

This section discusses miscellaneous utility functions that are provided by the Universe II, including:

- “Resets” on page 2-110,
- “Power-Up Options” on page 2-115,
- “Hardware Initialization (Normal Operating Mode)” on page 2-121,
- “Test Modes” on page 2-122, and
- “Clocks” on page 2-123.

2.10.1 Resets

This section is divided in three sections. The first section lists the pins and registers that are involved in the reset circuitry. The second section presents and explains the reset circuitry diagram. The third section provides important suggestions and warnings about configuring the Universe II reset circuitry.

2.10.1.1 Overview of Reset Support

The Universe II provides a number of pins and registers for reset support. Pin support is summarized in Table 2.19.

Table 2.19 : Hardware Reset Mechanism

Interface and Direction	Pin Name	Long Name	Effects ^a
VMEbus Input	VRSYSRST#	VMEbus Reset Input	Asserts LRST# on the local bus, resets the Universe II, and reconfigures power-up options.
VMEbus Output	VXSYSRST	VMEbus System Reset	Universe II output for SYSRST* (resets the VMEbus)
PCI Input	PWRRST#	Power-up Reset	Resets the Universe II and reconfigures power-up options.
	RST#	PCI Reset Input	Resets the Universe II from the PCI bus.
	VME_RESET#	VMEbus Reset Initiator	Causes Universe II to assert VXSYSRST
PCI Output	LRST#	PCI Bus Reset Output	Resets PCI resources
JTAG Input	TRST#	JTAG Test Reset	Provides asynchronous initialization of the TAP controller in the Universe II.

a. A more detailed account of the effects of reset signals is provided in “Reset Implementation Cautions” on page 2-114

The Universe II is only reset through hardware. Software can only make the Universe II assert its reset outputs. In order to reset the Universe II through software, the Universe II reset outputs must be connected to the Universe II reset inputs. For example, the SW_LRST bit in the MISC_CTL register, which asserts the LRST# output, will not reset the Universe II itself unless LRST# is looped back to RST#. As described in “Reset Implementation Cautions” on page 2-114, there are potential loopback configurations resulting in permanent reset.

Table 2.20 : Software Reset Mechanism

Register and Table	Name	Type	Function
MISC_CTL Table A.82	SW_LRST	W	Software PCI Reset 0=No effect, 1=Initiate LRST# A read always returns 0.
	SW_SYSRST	W	Software VMEbus SYSRESET 0=No effect, 1=Initiate SYSRST* A read always returns 0.
VCSR_SET Table A.126	RESET	R/W	Board Reset Reads: 0=LRST# not asserted, 1=LRST# asserted Writes: 0=no effect, 1=assert LRST#
	SYSFAIL	R/W	VMEbus SYSFAIL Reads: 0=VXSYSFAIL not asserted, 1=VXSYSFAIL asserted Writes: 0=no effect, 1=assert VXSYSFAIL
VCSR_CLR Table A.125	RESET	R/W	Board Reset Reads: 0=LRST# not asserted, 1=LRST# asserted Writes: 0=no effect, 1=negate LRST#
	SYSFAIL	R/W	VMEbus SYSFAIL Reads: 0=VXSYSFAIL not asserted, 1=VXSYSFAIL asserted Writes: 0=no effect, 1=negate VXSYSFAIL

More detailed information about the effects of various reset events is provided in the next section.

2.10.1.2 Universe II Reset Circuitry

Table 2.21 below shows how to reset various aspects of the Universe II. For example, it shows that in order to reset the clock services (SYSCLK, CLK64 enables, and PLL divider), PWRRST# should be asserted. If the table is read from left to right, it indicates the effects of various reset sources. Notice that PWRRST# resets all aspects of Universe II listed in column 1 of Table 2.21. Table 2.21 also indicates the reset effects that are extended in time. For example, VXSYSRST# remains asserted for 256 ms after all initiators are removed—this satisfies VMEbus rule 5.2 (minimum of 200 ms SYSRST*). The external 64 MHz clock controls this assertion time. LRST# is asserted for 5 ms or more from all sources except VRSYSRST#. The same information is presented in pictorial format in Figure 2.20.

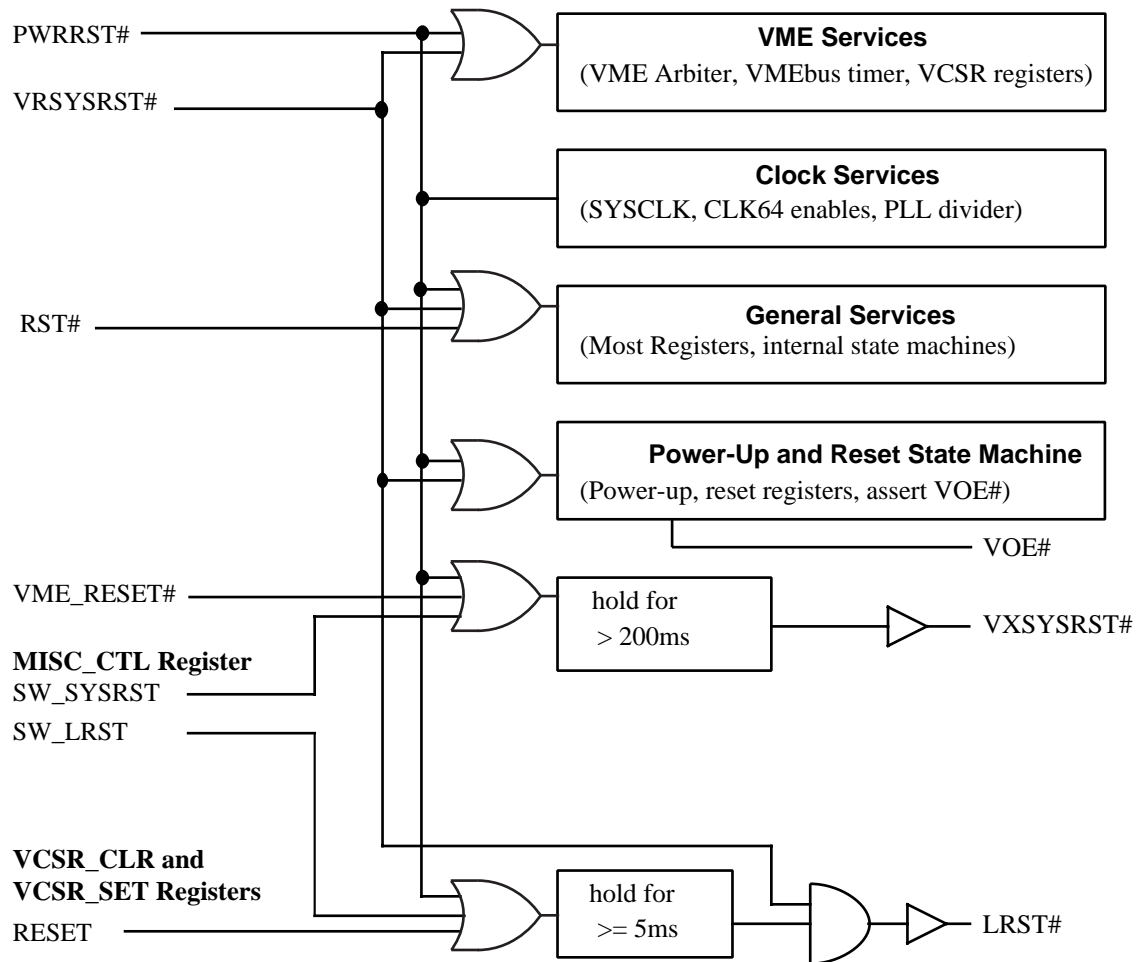
Table 2.21 : Functions Affected by Reset Initiators

Effect of Reset ^{a,b}	Reset Source
Clock Services SYSCLK CLK64 enables PLL Divider	PWRRST#
VMEbus Services VMEbus Arbiter VMEbus Timer VCSR Registers	PWRRST#, or VRSYSRST#
General Services Most registers Internal state machines	PWRRST#, RST# or VRSYSRST#
Power-Up and Reset State Machine Power-up the device Reset Registers	PWRRST#, or VRSYSRST#
VMEbus Reset Output VXSYSRST# (asserted for more than 200 ms)	PWRRST#, or VME_RESET#, or SW_SYSRST bit in MISC_CTL register
PCI Bus Reset Output LRST# (asserted for at least 5 ms)	PWRRST#, or SW_LRST bit in MISC_CTL register, or RESET bit in VCSR_SET register ^c

a. On PWRRST#, options are loaded from pins. On SYSRST and RST#, options are loaded from values that were latched at the previous PWRRST#.

b. Refer to Appendix-A to find the effects of various reset events

c. LRST# may be cleared by writing 1 to the RESET bit in the CSR_CLR register.



Note 1: On PWRRST#, options are loaded from pins. On SYSRST and RST#, options are loaded from values that were latched at the previous PWRRST#.

Note 2: Refer to Appendix-A to find the effects of various reset events.

Figure 2.20 : Reset Circuitry

2.10.1.3 Reset Implementation Cautions

To prevent the Universe II from resetting the PCI bus, the LRST# output may be left unconnected. Otherwise, LRST# should be grouped with other PCI reset generators to assert the RST# signal such that:

$RST\# = LRST\# \& \text{reset_source1} \& \text{reset_source2} \& \dots$

If the Universe II is the only initiator of PCI reset, LRST# may be directly connected to RST#.

Assertion of VME_RESET# causes the Universe II to assert VXSYSRST#.



Caution: Since VME_RESET# causes assertion of SYSRST, and since SYSRST* causes assertion of LRST#, tying both VME_RESET# and LRST# to RST# will put the Universe II into permanent reset. If VME_RESET# is to be driven by PCI reset logic, ensure that the logic is designed to break this feedback path.*

The PWRRST# input keeps the Universe II in reset until the power supply has reached a stable level (see Table 2.21). It should be held asserted for over 100 milliseconds after power is stable. Typically this can be achieved through a resistor/capacitor combination (see Figure 2.21); however, a more reliable solution using under voltage sensing circuits (e.g. MC34064) is common.

The Universe II supports the VMEbus CSR Bit Clear and Bit Set registers (Table A.125 and Table A.126). The VCSR_SET registers allows the user to assert LRST# or SYSFAIL by writing to the RESET or SYSFAIL bits, respectively. LRST# or SYSFAIL remains asserted until the corresponding bit is cleared in the VCSR_CLR register. The FAIL bit in each of these registers is a status bit and is set by the software to indicate board failure.




Figure 2.21 : Resistor-Capacitor Circuit Ensuring Power-Up Reset Duration

2.10.2 Power-Up Options

The Universe II may be automatically configured at power-up to operate in different functional modes. These power-up options allow the Universe II to be set in a particular mode independent of any local intelligence. The Universe II power-up options are listed in Table 2.22 and described below.

The majority of the Universe II power-up options (listed below) are loaded from the VMEbus address and data lines after any PWRRST#. There are two power-up options that are not initiated by PWRRST#. The first of these is PCI bus width (a power-up option required by the PCI bus specification), and this is loaded on any RST# event from the REQ64# pin. The second special power-up option is VMEbus SYSCON enabling, required by the VMEbus specification. The SYSCON option is loaded during a SYSRST* event from the BG3IN* signal.

All power-up options are latched from the state of a particular pin or group of pins on the rising edge of PWRST#. Each of these pins *except* REQ64# has a weak internal pull-down to put the Universe II into a default configuration. (REQ64# has an internal *pull-up*). If a non-default configuration is required, a pull-up of approximately 10k Ω (or active drive) is required on the signal. See “PCI Bus Width” on page 2-119 and the VMEbus Specification concerning Auto-Syscon Detect for the exceptions to the rule described in this paragraph.

The Universe II may be restored to the state it was in immediately following the previous power-up without re-asserting PWRRST#: after SYSRST* or RST# (with PWRRST# negated), the values that were originally latched at the rising edge of PWRRST# will be reloaded into the Universe II (except for PCI bus width and VMEbus SYSCON enabling, which are loaded from their pins).

Table 2.22 lists the power-up options of the Universe II, the pins which determine the options, and the register settings that are set by this option. Each option is described in more detail in “Power-up Option Descriptions” on page 2-117.

Table 2.22 : Power-Up Options^a

Option	Register	Field	Default	Pins
VMEbus Register Access Slave Image	VRAI_CTL	EN	disabled	VA[31]
		VAS	A16	VA[30:29]
	VRAI_BS	BS	0x00	VA[28:21]
VMEbus CR/CSR slave image	VCSR_CTL	LAS[0] ^b	memory	VA[20]
	VCSR_TO	TO	0x00	VA[19:15]
Auto-ID	MISC_STAT	DY4AUTO	disabled	VD[30]
	MISC_CTL	V64AUTO	disabled	VD[29]
	VINT_EN	SW_INT	0	
	VINT_STAT	SW_INT	0	
	VINT_MAP1	SW_INT	000	
BI-Mode®	MISC_CTL	BI	disabled	VD[28]
Auto-Syscon Detect	MISC_CTL	SYSCON	enabled	VBGIN[3]*
SYSFAIL* Assertion	VCSR_SET	SYSFAIL	asserted	VD[27]
	VCSR_CLR	SYSFAIL		
PCI Target Image	LSIO_CTL	EN	disabled	VA[13]
		LAS[0]	memory	VA[12]
		VAS	A16	VA[11:10]
	LSIO_BS	BS	0x0	VA[9:6]
	LSIO_BD	BD	0x0	VA[5:2]
PCI Register Access	PCI_BS0, PCI_BS1	SPACE	See Table A.6 and Table A.7	VA[1]
PCI Bus Size ^c	MISC_STAT	LCLSIZE	32-bit	REQ64#
PCI CSR Image Space	PCI_CSR	BM	disabled	VA[14]

a. All power-up options are latched only at the rising-edge of PWRRST#. They are loaded when PWRRST#, SYSRST* and RST# are negated.

b. The LAS field will enable the PCI_CSR register's MS or IOS field if the EN FIELD of the LSIO_CTL register is set.

c. The PCI Bus Size is loaded on any RST# event, as per the PCI 2.1 Specification.

2.10.2.1 Power-up Option Descriptions

This section describes each of the groups of power-up options that were listed in Table 2.22.

VMEbus Register Access Image

The Universe II has several VMEbus slave images, each of which may provide a different mapping of VMEbus cycles to PCI cycles. All VMEbus slave images are configurable through a set of VMEbus slave registers: VSIxCTL, VSIx_BS, VSIx_BA, and VSIx_IO. No VMEbus to PCI transaction is possible until these registers are programmed.

The VMEbus Register Access Image (VRAI) power-up option permits access from the VMEbus to the Universe II internal registers at power-up. The power-up option allows programming of the VMEbus register slave image address space and the upper five bits of its base address; all other bits will be zero (see Table 2.23 below). Once access is provided to the registers, then all other Universe II features (such as further VMEbus slave images) can be configured from the VMEbus.

Table 2.23 : VRAI Base Address Power-up Options

VRAI_CTL: VAS	BS [31:24]	BS [23:16]	BS [15:12]
A16	0	0	Power-up Option VA [28:25]
A24	0	Power-up Option VA [28:21]	0
A32	Power-up Option VA [28:21]	0	0

Table 2.23 above shows how the upper bits in the VRAI base address are programmed for A16, A24, and A32 VMEbus register access images.

VMEbus CR/CSR Slave Image

CR/CSR space is an address space introduced in the VME64 specification. The CR/CSR space on any VMEbus device is 512 Kbytes in size: the upper region of the 512 Kbytes dedicated to register space, and the lower region is dedicated to configuration ROM. The Universe II maps its internal registers to the upper region of the CR/CSR space, and passes all other accesses through to the PCI bus (see “Registers” on page 2-100).

The VMEbus CR/CSR Slave Image power-up option maps CR/CSR accesses to the PCI bus. CR/CSR space can be mapped to memory or I/O space with a 5-bit offset. This allows mapping to any 128Mbyte page on the PCI bus. As part of this implementation, ensure that the PCI Master Interface is enabled through the MAST_EN bit power-up option (see below) or configured through a register access before accessing configuration ROM.

Auto-ID

There are two Auto-ID mechanisms provided by the Universe II. One is the VME64 specified version which relies upon use of the CR/CSR space for configuration of the VMEbus system, and a Tundra proprietary system which uses the IACK daisy chain for identifying cards in a system. Either of these mechanisms can be enabled at power-up (see “Automatic Slot Identification” on page 2-24).

Because VME64 Auto-ID relies upon SYSFAIL to operate correctly, this power-up option overrides the SYSFAIL power-up option described below.

BI-Mode

BI-Mode (Bus Isolation Mode) is a mechanism for logically isolating the Universe II from the VMEbus for diagnostic, maintenance and failure recovery purposes. BI-Mode may be enabled as a power-up option (see “BI-Mode” on page 2-28). When the Universe II has been powered-up in BI-Mode, then any subsequent SYSRST* or RST# restores the Universe II to to BI-Mode,

Auto-Syscon Detect

The VMEbus SYSCON enabling, required by the VMEbus specification, is a special power-up option in that it does not return to its after-power-up state following RST# or SYSRST#. The SYSCON option is loaded during a SYSRST* event from the VBG3IN* signal.

SYSFAIL* Assertion

This power-up option causes the Universe II to assert SYSFAIL* immediately upon entry into reset. The SYSFAIL* pin is released through a register access. Note that this power-up option is over-ridden if VME64 Auto-ID has been enabled. This option would be used when extensive on-board diagnostics need to be performed before release of SYSFAIL*. After completion of diagnostics, SYSFAIL* may be released through software or through initiation of the VME64 Auto-ID sequence if that mechanism is to be used (see “Auto Slot ID: VME64 Specified” on page 2-24).

PCI Target Image

The PCI Target Image power-up option provides for default enabling of a PCI target image (automatically mapping PCI cycles to the VMEbus). The default target image can be mapped with base and bounds at 256MB resolution in Memory or I/O space, and map PCI transactions to different VMEbus address spaces. Beyond the settings provided for in this power-up option, the target image will possess its other default conditions: the translation offset will be zero, posted writes will be disabled, and only 32-bit (maximum) non-block VMEbus cycles in the non-privileged data space will be generated. This option would typically be used to access permits the use of Boot ROM on another card in the VMEbus system.

PCI Register Access

A power-up option determines if the registers are mapped into Memory or I/O space.

PCI Bus Width

The PCI Interface can be used as a 32-bit bus or 64-bit bus. The PCI bus width is determined during a PCI reset (see Section 4.3.2 of the PCI Specification, Rev. 2.1). The Universe II is configured as 32-bit PCI if REQ64# is high on RST#; it is configured as 64-bit if REQ64# is low. The Universe II has an internal pull-up on REQ64#, so the Universe II defaults to 32-bit PCI. On a 32-bit PCI bus, the Universe II drives all its 64-bit extension bi-direct signals at all times; these signals include: C/BE[7:4]#, AD[63:32], REQ64#, PAR64 and ACK64# to unknown values. If used as a 32-bit interface, the 64-bit pins, AD[63:32], C/BE[7:4], PAR64 and ACK64# may be left unterminated.

PCI CSR Image Space

There is a power-up option (using the VA[1] pin) that determines the value of the SPACE bit of the PCI_BSx registers. At power-up the SPACE bit of the PCI_BS1 register is the negation of the SPACE bit of the PCI_BS0 register.

- When the VA[1] pin is sampled low at power-up, the PCI_BS0 register's SPACE bit is set to "1", which signifies I/O space, and the PCI_BS1 register's SPACE bit is set to "0", which signifies Memory space.
- When VA[1] is sampled high at power-up, the PCI_BS0 register's SPACE register's bit is set to "0", which signifies Memory space, and the PCI_BS1 register's SPACE bit is set to "1", which signifies I/O space.

Once set, this mapping persists until the next power-up sequence.

See "Memory or I/O Access" on page 2-103, Table A.6 and Table A.7.

2.10.2.2 Power-Up Option Implementation

This section describes pull-up requirements and timing relevant to the power-up options.

The pull-ups for the general power-up options (if other than default values are required) must be placed on the VA[31:1] and VD[31:27] lines. During reset, the Universe II will negate VOE#, putting these signals into a high-impedance state. While VOE# is negated the pull-ups (or internal pull-downs) will bring the option pins (on A[31:1] and D[31:27]) to their appropriate state.



Caution: The internal pull-downs are very weak. The leakage current on many transceivers may be sufficient to override these pull-downs. To ensure proper operation designers should ensure power-up option pins will go to the correct state.

Within two CLK64 periods after PWRRST# is negated, the Universe II latches the levels on the option pins, and then negates VOE# one clock later. This enables the VMEbus transceivers inwards.




Figure 2.22 : Power-up Options Timing

The power-up options are subsequently loaded into their respective registers several PCI clock periods after PWRRST#, SYSRST* and RST# have all been negated.



Because of the power-up configuration, the VMEbus buffers are not enabled until several CLK64 periods after release of SYSRST (approximately 45 ns). Allowing for worst case backplane skew of 25 ns, the Universe II will not be prepared to receive a slave access until 70 ns after release of SYSRST*.*

2.10.3 Hardware Initialization (Normal Operating Mode)

The Universe II has I/O capabilities that are specific to manufacturing test functions. These pins are not required in a non-manufacturing test setting. Table 2.24 below shows how these pins should be terminated.

Table 2.24 : Manufacturing Pin Requirements for Normal Operating Mode

Pin Name	Pin Value
TMODE[2]	V _{SS} (or pulled-down if board tests will occasionally be performed, see “Auxiliary Test Modes” on page 2-122)
TMODE[1]	
TMODE[0]	
PLL_TESTSEL	V _{SS}
ENID	V _{SS}
PLL_TESTOUT	N/C
VCCTL	V _{SS}

2.10.4 Test Modes

The Universe II provides two types of test modes: auxiliary modes (NAND tree simulation and High Impedance) and JTAG (IEEE 1149.1).

2.10.4.1 Auxiliary Test Modes

Two auxiliary test modes are supported: NAND tree and high impedance. The Universe II has three test mode input pins (TMODE[2:0]). For normal operations these inputs should be tied to ground (or pulled to ground through resistors). Table 2.25 below indicates the 3 operating modes of the Universe II. At reset the TMODE[2:0] inputs are latched by the Universe II to determine the mode of operation. The Universe II remains in this mode until the TMODE[2:0] inputs have changed and a reset event has occurred. PLL_TESTSEL must be high for any test mode.

Table 2.25 : Test Mode Operation

Operation Mode	TMODE[2:0]	PLL_TESTSEL
Normal Mode	000	0
Accelerate	001	0
Reserved	010	1
Reserved	011	1
NAND Tree Simulation	100	1
Reserved	101	1
High Impedance	110	0/1
Reserved	111	1

For NAND Tree Simulation, the values of the TMODE pins are latched during the active part of PWRRST#. These pins can change state during the NAND Tree tests. The timers are always accelerated in this mode. All outputs are tristated in this mode, except for the VXSYSFAIL output pin.

For High Impedance mode, the values of the TMODE pins are also latched during the active part of PWRRST#. All outputs are tristated in this mode, except for the VXSYSFAIL output pin.

2.10.4.2 JTAG support

The Universe II includes dedicated user-accessible test logic that is fully compatible with the IEEE 1149.1 Standard Test Access Port (TAP) and Boundary Scan Architecture. This standard was developed by the Test Technology Technical Committee of IEEE Computer Society and the Joint Test Action Group (JTAG). The Universe II's JTAG support includes:

- A five-pin JTAG interface (TCK, TDI, TDO, TMS, and TRST#),
- a JTAG TAP controller,
- a three-bit instruction register,
- a boundary scan register,
- a bypass register,
- and an IDCODE register.

The following required public instructions are supported: BYPASS (3'b111), SAMPLE(3'b100), and EXTEST(3'b000). The optional public instruction IDCODE(3'b011) selects the IDCODE register which returns 32'b01e201d. The following external pins are not part of the boundary scan register: LCLK, PLL_TESTOUT, PLL_TESTSEL, TMODE[3:0], and VCOCTL.

A BSDL file is available upon request from Tundra Semiconductor Corporation.

2.10.5 Clocks

CLK64 is a 64 MHz clock that is required by the Universe II in order to synchronize internal Universe II state machines and to produce the VMEbus system clock (VSYSClk) when the Universe II is system controller (SYSCON). This clock is specified to have a minimum 60-40 duty cycle with a maximum rise time of 5 ns. Using a different frequency is not recommended as it will alter various internal timers and change some VME timing.

3 Description of Signals

The following detailed description of the Universe II signals is organized according to these functional groups:

- VMEbus Signals
- PCI Signals

3.1 VMEbus Signals

CLK64	Input
Reference Clock – this 64MHz clock is used to generate fixed timing parameters. It requires a 50-50 duty cycle ($\pm 20\%$) with a 5ns maximum rise time. CLK64 is required to synchronize the internal state machines of the VME side of the Universe II.	
VA [31:1]	Bidirectional
VMEbus Address Lines 31 to 01 – during MBLT transfers, VA 31-01 serve as data bits D63-D33. VA03-01 are used to indicate interrupt level on the VMEbus.	
VA_DIR	Output
VMEbus Address Transceiver Direction Control – the Universe II controls the direction of the address (VA31-01, VLWORD#) transceivers as required for master, slave and bus isolation modes. When the Universe II is driving lines on the VMEbus, this signal is driven high; when the VMEbus is driving the Universe II, this signal is driven low.	
VAM [5:0]	Bidirectional
VMEbus Address Modifier Codes – these codes indicate the address space being accessed (A16, A24, A32), the privilege level (user, supervisor), the cycle type (standard, BLT, MBLT) and the data type (program, data).	
VAM_DIR	Output
VMEbus AM Code Direction Control – controls the direction of the AM code transceivers as required for master, slave and bus isolation modes. When the Universe II is driving lines on the VMEbus, this signal is driven high; when the VMEbus is driving the Universe II, this signal is driven low.	
VAS#	Bidirectional
VMEbus Address Strobe – the falling edge of VAS# indicates a valid address on the bus. By continuing to assert VAS#, ownership of the bus is maintained during a RMW cycle.	

3.1 VMEbus Signals (Continued)

VAS_DIR	Output
VMEbus Address Strobe Direction Control – controls the direction of the address strobe transceiver as required for master, slave and bus isolation modes. When the Universe II is driving lines on the VMEbus, this signal is driven high; when the VMEbus is driving the Universe II, this signal is driven low.	
VBCLR#	Output
VMEbus Bus Clear – requests that the current owner release the bus. Asserted by the Universe II when configured as SYSCON and the arbiter detects a higher level pending request.	
VBGI# [3:0]	Input
VMEbus Bus Grant Inputs – The VME arbiter awards use of the data transfer bus by driving these bus grant lines low. The signal propagates down the bus grant daisy chain and is either: <ul style="list-style-type: none"> • accepted by a requester if it requesting at the appropriate level, or • passed on as a VBGO [3:0]# to the next board in the bus grant daisy chain. 	
VBGO# [3:0]	Output
VMEbus Bus Grant Outputs – Only one output is asserted at any time, according to the level at which the VMEbus is being granted.	
VD [31:0]	Bidirectional
VMEbus Data Lines 31 through 00	
VD_DIR	Output
VMEbus Data Transceiver Direction Control – the Universe II controls the direction of the data (VD [31:0]) transceivers as required for master, slave and bus isolation modes. When the Universe II is driving lines on the VMEbus, this signal is driven high; when the VMEbus is driving the Universe II, this signal is driven low.	
VDS# [1:0]	Bidirectional
VMEbus Data Strokes – the level of these signals are used to indicate active byte lanes: <ul style="list-style-type: none"> • During write cycles, the falling edge indicates valid data on the bus. • During read cycles, assertion indicates a request to a slave to provide data. 	
VDS_DIR	Output
VMEbus Data Strobe Direction Control – controls the direction of the data strobe transceivers as required for master, slave and bus isolation modes. When the Universe II is driving lines on the VMEbus, this signal is driven high; when the VMEbus is driving the Universe II, this signal is driven low.	
VDTACK#	Bidirectional
VMEbus Data Transfer Acknowledge – VDTACK# driven low indicates that the addressed slave has responded to the transfer. The Universe II always rescinds DTACK*. It is tristated once the initiating master negates AS*.	

3.1 VMEbus Signals (Continued)

VIACK#	Bidirectional
VMEbus Interrupt Acknowledge – indicates that the cycle just beginning is an interrupt acknowledge cycle.	
VIACKI#	Input
VMEbus Interrupt Acknowledge In – Input for IACK daisy chain driver. If interrupt acknowledge is at same level as interrupt currently generated by the Universe II, then the cycle is accepted. If interrupt acknowledge is not at same level as current interrupt or Universe II is not generating an interrupt, then the Universe II propagates VIACKO#.	
VIACKO#	Output
VMEbus Interrupt Acknowledge Out– generated by the Universe II if it receives VIACKI# and is not currently generating an interrupt at the level being acknowledged.	
VLWORD#	Bidirectional
VMEbus Longword Data Transfer Size Indicator – this signal is used in conjunction with the two data strobes VDS [1:0]# and VA 01 to indicate the number of bytes (1 – 4) in the current transfer. During MBLT transfers VLWORD# serves as data bit D32.	
VOE#	Output
VMEbus Transceiver Output Enable – used to control transceivers to isolate the Universe II from the VMEbus during a reset or BI-mode. On power-up, VOE# is high (to disable the buffers).	
VRACFAIL#	Input
VMEbus ACFAIL Input signal – warns the VMEbus system of imminent power failure. This gives the modules in the system time to shut down in an orderly fashion before powerdown. ACFAIL is mapped to a PCI interrupt.	
VRBBSY#	Input
VMEbus Receive Bus Busy – allows the Universe II to monitor whether the VMEbus is owned by another VMEbus master	
VRBERR#	Input
VMEbus Receive Bus Error – a low level signal indicates that the addressed slave has not responded, or is signalling an error.	
VRBR# [3:0]	Input
VMEbus Receive Bus Request Lines – if the Universe II is the Syscon, the Arbiter logic monitors these signals and generates the appropriate Bus Grant signals. Also monitored by requester in ROR mode.	
VRIRQ# [7:1]	Input
VMEbus Receive Interrupts 7 through 1 – these interrupts can be mapped to any of the Universe II's PCI interrupt outputs. VRIRQ7-1# are individually maskable, but cannot be read.	

3.1 VMEbus Signals (Continued)

VRSYSFAIL#	Input
VMEbus Receive SYSFAIL – asserted by a VMEbus system to indicate some system failure. VRSYSFAIL# is mapped to a PCI interrupt.	
VRSYSRST#	Input
VMEbus Receive System Reset – causes assertion of LRST# on the local bus and resets the Universe II.	
VSLAVE_DIR	Output
VMEbus Slave Direction Control – transceiver control that allow the Universe II to drive DTACK* on the VMEbus. When the Universe II is driving lines on the VMEbus, this signal is driven high; when the VMEbus is driving the Universe II, this signal is driven low.	
VSYSCLK	Bidirectional
VMEbus System Clock – generated by the Universe II when it is the Syscon and monitored during DY4 Auto ID sequence	
VSCON_DIR	Output
Syscon Direction Control – transceiver control that allows the Universe II to drive VBCLR# and SYSCLK. When the Universe II is driving lines on the VMEbus, this signal is driven high; when the VMEbus is driving the Universe II, this signal is driven low.	
VWRITE#	Bidirectional
VMEbus Write signal – indicates the direction of data transfer.	
VXBBSY	Output
VMEbus Transmit Bus Busy Signal – generated by the Universe II when it is VMEbus master	
VXBERR	Output
VMEbus Transmit Bus Error Signal – generated by the Universe II when PCI target generates target abort on coupled PCI access from VMEbus.	
VXBR [3:0]	Output
VMEbus Transmit Bus Request – the Universe II requests the VMEbus when it needs to become VMEbus master.	
VXIRQ [7:1]	Output
VMEbus Transmit Interrupts – the VMEbus interrupt outputs are individually maskable.	
VXSYSFAIL	Output
VMEbus System Failure – asserted by the Universe II during reset and plays a role in VME64 Auto ID.	
VXSYSRST	Output
VMEbus System Reset – the Universe II output for SYSRST*.	

3.2 PCI Bus Signals

ACK64#	Bidirectional
Acknowledge 64-bit Transfer – when driven by the PCI slave (target), it indicates slave can perform a 64-bit transfer.	
AD [31:0]	Bidirectional
PCI Address/Data Bus – address and data are multiplexed over these pins providing a 32-bit address/data bus.	
AD [63:32]	Bidirectional
PCI Address/Data Bus – address and data are multiplexed over these pins providing 64-bit address and data capability.	
C/BE# [7:0]	Bidirectional
PCI Bus Command and Byte Enable Lines – command and byte enable information is multiplexed over all eight C/BE lines. C/BE [7:4]# are only used in a 64-bit PCI bus	
DEVSEL#	Bidirectional
PCI Device Select – is driven by the Universe II when it is accessed as PCI slave.	
ENID	Input
Enable IDD Tests – required for ASIC manufacturing test, tie to ground for normal operation.	
FRAME#	Bidirectional
Cycle Frame – is driven by the Universe II when it is PCI initiator, and is monitored by the Universe II when it is PCI target	
GNT#	Input
PCI Grant – indicates to the Universe II that it has been granted ownership of the PCI bus.	
IDSEL	Input
PCI Initialization Device Select – is used as a chip select during configuration read and write transactions	
LINT# [7:0]	Bidirectional (Open Drain)
PCI Interrupt Inputs – these PCI interrupt inputs can be mapped to any PCI bus or VMEbus interrupt output.	
IRDY#	Bidirectional
Initiator Ready – is used by the Universe II as PCI master to indicate that is ready to complete a current data phase.	
LCLK	Input
PCI Clock – provides timing for all transactions on the PCI bus. PCI signals are sampled on the rising edge of CLK, and all timing parameters are defined relative to this signal. The PCI clock frequency of the Universe II must be between 25 and 33MHz. Lower frequencies will result in invalid VME timing.	

3.2 PCI Bus Signals (Continued)

LOCK#	Bidirectional
Lock – used by the Universe II to indicate an exclusive operation with a PCI device. While the Universe II drives LOCK#, other PCI masters are excluded from accessing that particular PCI device. Likewise, when the Universe II samples LOCK#, it may be excluded from a particular PCI device.	
LRST#	Output
PCI Reset Output – used to reset PCI resources.	
PAR	Bidirectional
Parity – parity is even across AD [31:0] and C/BE [3:0] (the number of 1s summed across these lines and PAR equal an even number).	
PAR64	Bidirectional
Parity Upper DWORD – parity is even across AD [63:32] and C/BE [7:4] (the number of 1s summed across these lines and PAR equal an even number).	
PERR#	Bidirectional
Parity Error – reports parity errors during all transactions. The Universe II drives PERR# high within two clocks of receiving a parity error on incoming data, and holds PERR# for at least one clock for each errored data phase.	
PLL_TESTOUT	Output
Manufacturing Test Output—No connect	
PLL_TESTSEL	Input
Manufacturing Test Select—tie to ground for normal operation	
PWRRST#	Input
Power-up Reset – all Universe II circuitry is reset by this input.	
REQ#	Output
Bus Request – used by the Universe II to indicate that it requires the use of the PCI bus.	
REQ64#	Bidirectional
64-Bit Bus Request– used to request a 64-bit PCI transaction. If the target does not respond with ACK64#, 32-bit operation is assumed.	
RST#	Input
PCI Reset Input—resets the Universe II from the PCI bus.	
SERR#	Bidirectional
System Error – reports address parity errors or any other system error.	

3.2 PCI Bus Signals (Continued)

STOP#	Bidirectional
Stop – used by the Universe II as PCI slave when it wishes to signal the PCI master to stop the current transaction. As PCI master, the Universe II will terminate the transaction if it receives STOP# from the PCI slave.	
TCK	Input
JTAG Test Clock Input – used to clock the Universe II TAP controller. Tie to any logic level if JTAG is not used in the system.	
TDI	Input
JTAG Test Data Input – used to serially shift test data and test instructions into the Universe II. Tie to any logic level if JTAG is not used in the system.	
TDO	Output
JTAG Test Data Output – used to serially shift test data and test instructions out of the Universe II	
TMODE [2:0]	Input
Test Mode Enable – used for chip testing, tie to ground for normal operation.	
TMS	Input
JTAG Test Mode Select – controls the state of the Test Access Port (TAP) controller in the Universe II. Tie to any logic level if JTAG is not used in the system.	
TRDY#	Bidirectional
Target Ready – used by the Universe II as PCI slave to indicate that it is ready to complete the current data phase. During a read with Universe II as PCI master, the slave asserts TRDY# to indicate to the Universe II that valid data is present on the data bus.	
TRST#	Input
JTAG Test Reset – provides asynchronous initialization of the TAP controller in the Universe II. Tie to ground if JTAG is not used in the system.	
VCOCTL	Input
Manufacturing testing, tie to ground for normal operation	
VME_RESET#	Input
VMEbus Reset Input — generates a VME bus system reset.	

4 Signals and DC Characteristics

4.1 Terminology

The I/O type abbreviations used in the pin list in Table 4.2 on page 4-3 are defined below. A numbered suffix indicates the current rating of the output (in mA).

Analog	Analog input signal
I	Input only
I/O	Input and output
O	Output only
OD	Open drain output
PD	Pulled-down internally
PU	Pulled-up internally
TP	Totem pole output
TTL	Input with TTL thresholds
TTL SCH	Schmitt trigger input with TTL thresholds
3S	Tri-state output

4.2 DC Characteristics and Pin Assignments

Table 4.1 : DC Electrical Characteristics ($V_{DD} = 5\text{ V} \pm 10\%$)

Symbol	Parameter	Signal Type	Test Conditions	Tested at 0°C to 70°C	
				Min	Max
V_{IH}	Min. high-level input	CTTL	$V_{OUT} = 0.1\text{ V or } V_{DD} - 0.1\text{ V}; [I_{OUT}] = 20\text{ }\mu\text{A}$	2.2 V	$V_{DD} + 0.3\text{ V}$
		CMOS	$V_{OUT} = 0.1\text{ V or } V_{DD} - 0.1\text{ V}; [I_{OUT}] = 20\text{ }\mu\text{A}$	$0.7V_{DD}$	$V_{DD} + 0.3\text{ V}$
V_{IL}	Max. low-level input	CTTL	$V_{OUT} = 0.1\text{ V or } V_{DD} - 0.1\text{ V}; [I_{OUT}] = 20\text{ }\mu\text{A}$	-0.3 V	0.8V
		CMOS	$V_{OUT} = 0.1\text{ V or } V_{DD} - 0.1\text{ V}; [I_{OUT}] = 20\text{ }\mu\text{A}$	-0.3 V	$0.3V_{DD}$
V_{T+}	Positive going Schmitt trigger voltage	CTTL/SCH	$V_{OUT} = 0.1\text{ V or } V_{DD} - 0.1\text{ V}; [I_{OUT}] = 20\text{ }\mu\text{A}$	–	2.4 V
		CMOS/SCH	$V_{OUT} = 0.1\text{ V or } V_{DD} - 0.1\text{ V}; [I_{OUT}] = 20\text{ }\mu\text{A}$	–	$0.7V_{DD}$
V_{T-}	Negative going Schmitt trigger voltage	CTTL/SCH	$V_{OUT} = 0.1\text{ V or } V_{DD} - 0.1\text{ V}; [I_{OUT}] = 20\text{ }\mu\text{A}$	0.8 V	–
		CMOS/SCH	$V_{OUT} = 0.1\text{ V or } V_{DD} - 0.1\text{ V}; [I_{OUT}] = 20\text{ }\mu\text{A}$	$0.25V_{DD}$	–
$V_{Hysteresis}$	Schmitt trigger hysteresis voltage	CTTL/SCH	V_{T+} to V_{T-}	$0.05V_{DD}$	–
		CMOS/SCH	V_{T+} to V_{T-}	$0.12V_{DD}$	–
I_{IN}	Maximum input leakage current	CMOS and CTTL	With no pull-up resistor ($V_{IN} = V_{SS}$ or V_{DD})	-5.0 μA	5.0 μA
I_{OZ}	Maximum output leakage current	3S	$(V_{OUT} = V_{SS} \text{ or } V_{DD})$	-10.0 μA	10.0 μA
		OD	$(V_{OUT} = V_{DD})$	-10.0 μA	10.0 μA

Table 4.2 : Pin List and DC Characteristics for Universe II Signals

Pin Name	PBGA Pin Number	CBGA Pin Number	Type	Input Type	Output Type	I _{OL} (mA)	I _{OH} (mA)	Signal Description ¹
ACK64#	W11	W8	I/O	TTL	3S	6	−2	PCI Acknowledge 64 Bit Transfer
AD [63:0]	TABLE 4.3		I/O	TTL	3S	6	−2	PCI Address/Data Pins
C/BE# [0]	Y14	T11	I/O	TTL	3S	6	−2	PCI Command and Byte Enables
C/BE# [1]	V14	Y11						
C/BE# [2]	T14	U11						
C/BE# [3]	W13	W10						
C/BE# [4]	AE15	Y12						
C/BE# [5]	AD14	V11						
C/BE# [6]	T12	V10						
C/BE# [7]	AD12	Y9						
CLK64	C23	D16	I	TTL	−	−	−	VME Clock 64 MHz—60-40 duty, 5 ns rise time
DEVSEL#	AC7	V6	I/O	−	3S	6	−2	PCI Device Select
ENID	AE21	Y16	I	CMOS	−	−	−	Enable IDD Tests
FRAME#	W17	T12	I/O	TTL	3S	6	−2	PCI Cycle Frame
GNT#	AE17	Y14	I	TTL	−	−	−	PCI Grant
IDSEL	AB16	Y15	I	TTL	−	−	−	PCI Initialization Device Select
LINT# [0]	K20	H15	I/O	TTL	OD	12	−12	PCI Interrupt
LINT# [1]	AA5	U3	I/O	TTL	OD	4	−4	
LINT# [2]	L9	J3						
LINT# [3]	V6	R4						
LINT# [4]	M4	J4						
LINT# [5]	L3	J6						
LINT# [6]	M8	J5						
LINT# [7]	L1	K4						
IRDY#	AC15	V12	I/O	TTL	3S	6	−2	PCI Initiator Ready
LCLK	AA3	W3	I	TTL	−	−	−	PCI Clock Signal
LOCK#	AA23	T18	I/O	TTL	3S	6	−2	PCI Lock
LRST#	R1	M1	O	−	3S	6	−2	PCI Reset Output
PAR	P8	L1	I/O	TTL	3S	6	−2	PCI parity
PAR64	AE5	W4	I/O	TTL	3S	6	−2	PCI Parity Upper DWORD
PERR#	AB4	W5	I/O	TTL	3S	6	−2	PCI Parity Error
PLL_TESTOUT	AB2	V1	FOR FACTORY TESTING					

Table 4.2 : Pin List and DC Characteristics for Universe II Signals (Continued)

Pin Name	PBGA Pin Number	CBGA Pin Number	Type	Input Type	Output Type	I _{OL} (mA)	I _{OH} (mA)	Signal Description ¹
PLL_TESTSEL	AC1	T2	FOR FACTORY TESTING					
PWRRST#	T4	R1	I	TTL/SCHM	–	–	–	Power-up Reset
REQ#	K22	F20	O	–	3S	6	–2	PCI Request
REQ64#	AD18	T13	I/O	TTL	3S	6	–2	PCI Request 64 Bit Transfer
RST#	AA19	W17	I	TTL	–	–	–	PCI Reset
SERR#	AA7	R7	O	TTL	OD	12	–12	PCI System Error
STOP#	AB18	W15	I/O	TTL	3S	6	–2	PCI Stop
TCK	H12	A10	I	TTL	–	–	–	JTAG Test Clock Input
TDI	A13	F10	I	TTL (PU)	–	–	–	JTAG Test Data Input
TDO	C13	E11	O	–	3S	–	–	JTAG Test Data OUTput
TMODE [0]	AA13	W11	I	TTL	–	–	–	Test Mode Enable
TMODE [1]	AA21	V17						
TMODE [2]	W23	R18						
TMS	C11	C9	I	TTL (PU)	–	–	–	JTAG Test Mode Select
TRDY#	AD8	W7	I/O	TTL	3S	6	–2	PCI Target Ready
TRST#	E13	B10	I	TTL (PU)	–	–	–	JTAG Test Reset
VA [31:1]	TABLE 4.4		I/O	TTL (PD)	3S	3	–3	VMEbus Address Pins
VAM [0]	E11	D8	I/O	TTL	3S	3	–3	VMEbus Address Modifier Signals
VAM [1]	D10	A6						
VAM [2]	G9	E9						
VAM [3]	B10	B8						
VAM [4]	H10	D9						
VAM [5]	A9	A7						
VAM_DIR	B8	E8	O	–	3S	6	–6	VMEbus AM Signal Direction Control
VAS#	B14	A12	I/O	TTL/SCHM (PU)	3S	3	–3	VMEbus Address Strobe
VAS_DIR	K12	D10	O	–	3S	6	–6	VMEbus AS Direction Control
VA_DIR	G13	B11	O	–	3S	12	–12	VMEbus Address Direction Control

Table 4.2 : Pin List and DC Characteristics for Universe II Signals (Continued)

Pin Name	PBGA Pin Number	CBGA Pin Number	Type	Input Type	Output Type	I _{OL} (mA)	I _{OH} (mA)	Signal Description ¹
VBCLR#	N3	K5	O	–	3S	3	–3	VMEbus BCLR* Signal
VBGI# [0]	N21	K19	I	TT	–	–	–	VMEbus Bus Grant In
VBGI# [1]	M16	K17						
VBGI# [2]	N25	K15						
VBGI# [3]	N23	L16		TT (PD)				
VBGO# [0]	M20	K16	O	–	3S	12	–12	VMEbus Bus Grant Out
VBGO# [1]	L25	J20						
VBGO# [2]	M18	K20						
VBGO# [3]	M24	K18						
VCOCTL	AE3	T5	I	–	–	–	–	Factory testing
VD [31:0]	TABLE 4.4		I/O	TTL	3S	3	–3	VMEbus Data Pins
VD_DIR	F10	F8	O	–	3S	12	–12	VMEbus Data Direction Control
VDS# [0]	F12	E10	I/O	TTL (PU)	3S	3	–3	VMEbus Data Strobes
VDS# [1]	A11	A9						
VDS_DIR	J11	F9	O	–	3S	6	–6	VMEbus Data Strobe Direction Control
VDTACK#	G15	B13	I/O	TTL/ SCHM (PU)	3S	3	–3	VMEbus DTACK* Signal
VIACK#	E7	E6	I/O	TTL	3S	3	–3	VMEbus IACK* Signal
VIACKI#	AE23	W16	I	TTL	–	–	–	VMEbus IACKIN* Signal
VIACKO#	L21	H17	O	–	3S	12	–12	VMEbus IACKOUT* Signal
VLWORD#	K14	C11	I/O	TTL (PD)	3S	3	–3	VMEbus LWORD* Signal
VME_RESET#	V22	T19	I	TTL				VMEbus Reset Input
VOE#	B12	C10	O	–	3S	24	–24	VMEbus Transceiver Output Enable
VRACFAIL#	P18	M16	I	TTL/ SCHM	–	–	–	VMEbus ACFAIL* Signal
VRBBSY#	M6	J2	I	TTL/ SCHM	–	–	–	VMEbus Received BBSY* Signal
VRBERR#	A7	B7	I	TTL/ SCHM	–	–	–	VMEbus Receive Bus Error

Table 4.2 : Pin List and DC Characteristics for Universe II Signals (Continued)

Pin Name	PBGA Pin Number	CBGA Pin Number	Type	Input Type	Output Type	I _{OL} (mA)	I _{OH} (mA)	Signal Description ¹
VRBR# [0]	W5	U2	I	TTL/ SCHM	–	–	–	VMEbus Receive Bus Request
VRBR# [1]	U1	P1						
VRBR# [2]	R3	M3						
VRBR# [3]	L7	H2						
VRIRQ# [1]	H22	F19	I	TTL/ SCHM	–	–	–	VMEbus Receive Interrupts
VRIRQ# [2]	H20	F17						
VRIRQ# [3]	E25	E20						
VRIRQ# [4]	J21	G18						
VRIRQ# [5]	V16	U12						
VRIRQ# [6]	P20	M19						
VRIRQ# [7]	R17	M18						
VRSYSFAIL#	AC13	T10	I	TTL	–	–	–	VMEbus Receive SYSFAIL Signal
VRSYSRST#	C21	C16	I	TTL/ SCHM	–	–	–	VMEbus Receive SYSRESET* Signal
VSCON_DIR	M2	J1	O	–	3S	6	–6	SYSCON signals direction control
VSLAVE_DIR	C15	F12	O	–	3S	6	–6	DTACK/BERR direction control
VSYSCLK	N7	K2	I/O	TTL	3S	3	–3	VMEbus SYSCLK Signal
VWRITE#	D8	B6	I/O	TTL	3S	3	–3	VMEbus Write
VXBBSY	P2	L3	O	–	3S	3	–3	VMEbus Transmit BBSY* Signal
VXBERR	D12	B9	O	–	3S	3	–3	VMEbus Transmit Bus Error (BERR*)
VXBR [0]	G25	G19	O	–	3S	3	–3	VMEbus Transmit Bus Request
VXBR [1]	H24	H16						
VXBR [2]	P24	M20						
VXBR [3]	G23	C19						

Table 4.2 : Pin List and DC Characteristics for Universe II Signals (Continued)

Pin Name	PBGA Pin Number	CBGA Pin Number	Type	Input Type	Output Type	I _{OL} (mA)	I _{OH} (mA)	Signal Description ¹
VXIRQ [1]	J19	J16	O	–	3S	3	–3	VMEbus Transmit Interrupts
VXIRQ [2]	K24	H19						
VXIRQ [3]	K18	J17						
VXIRQ [4]	J25	G20						
VXIRQ [5]	L23	J18						
VXIRQ [6]	M22	J19						
VXIRQ [7]	R25	L17						
VXSYSFAIL	M10	K3	O	–	3S	3	–3	VMEbus Transmit SYSFAIL Signal
VXSYSRST	A23	E16	O	–	3S	3	–3	VMEbus Transmit SYSRESET* Signal

Note 1: All PCI pins meet PCI's AC current specifications.

Table 4.3 : PCI Bus Address/Data Pins

Signal	PBGA	CBGA	Signal	PBGA	CBGA
AD [0]	P16	L18	AD [32]	L17	J15
AD [1]	P22	M17	AD [33]	N19	L19
AD [2]	R19	N19	AD [34]	R23	M15
AD [3]	T18	U20	AD [35]	U19	N18
AD [4]	T22	N15	AD [36]	U23	R20
AD [5]	T20	T20	AD [37]	W25	P16
AD [6]	AA25	U19	AD [38]	U21	P17
AD [7]	AB24	R17	AD [39]	V20	R19
AD [8]	AB22	R16	AD [40]	Y22	U18
AD [9]	AE25	T17	AD [41]	W21	P15
AD [10]	AC21	V19	AD [42]	AD22	Y18
AD [11]	AB20	V15	AD [43]	Y20	T15
AD [12]	AC19	W18	AD [44]	AD20	T14
AD [13]	AA17	V14	AD [45]	Y18	U15
AD [14]	AA15	U13	AD [46]	AE19	W14
AD [15]	U15	R12	AD [47]	AD16	W13
AD [16]	AE11	U10	AD [48]	V12	T9

Table 4.3 : PCI Bus Address/Data Pins

Signal	PBGA	CBGA		Signal	PBGA	CBGA
AD [17]	AB12	U9		AD [49]	Y12	W9
AD [18]	W9	V8		AD [50]	AC11	R9
AD [19]	AD10	U8		AD [51]	V10	Y4
AD [20]	AE7	T7		AD [52]	AB10	R8
AD [21]	Y8	W6		AD [53]	AA9	U7
AD [22]	AD4	U6		AD [54]	AB8	T6
AD [23]	Y6	R5		AD [55]	AB6	V4
AD [24]	Y2	P5		AD [56]	Y4	R3
AD [25]	V4	R2		AD [57]	W3	V2
AD [26]	U5	P3		AD [58]	AA1	T1
AD [27]	W1	P2		AD [59]	V2	N5
AD [28]	U7	M5		AD [60]	R5	N4
AD [29]	T8	M4		AD [61]	T2	N2
AD [30]	P6	L5		AD [62]	R9	M6
AD [31]	P10	L4		AD [63]	N5	L2

Table 4.4 : VMEbus Address Pins^a

Signal	PBGA	CBGA		Signal	PBGA	CBGA
VA [1]	H14	E12		VA [17]	D18	B16
VA [2]	A15	D11		VA [18]	C19	C15
VA [3]	F14	B12		VA [19]	B20	D17
VA [4]	J15	C12		VA [20]	B22	D15
VA [5]	D14	D12		VA [21]	D20	C17
VA [6]	G17	C13		VA [22]	F20	E15
VA [7]	H16	A17		VA [23]	E19	F14
VA [8]	B16	D13		VA [24]	A25	C20
VA [9]	C17	A15		VA [25]	E23	B18
VA [10]	D16	F13		VA [26]	C25	E19
VA [11]	A19	E14		VA [27]	G21	F16
VA [12]	B18	B14		VA [28]	E21	D18
VA [13]	F16	A16		VA [29]	F22	F18
VA [14]	E17	D14		VA [30]	D24	D19
VA [15]	A21	B17		VA [31]	F24	G16
VA [16]	F18	B15				

a. All VA pins have an internal pull-down.

Table 4.5 : VMEbus Data Pins^a

Signal	PBGA	CBGA		Signal	PBGA	CBGA
VD [0]	J7	H3		VD [16]	F6	E2
VD [1]	K8	D1		VD [17]	G5	G6
VD [2]	K2	H4		VD [18]	B2	E5
VD [3]	J3	F1		VD [19]	C1	E3
VD [4]	K4	H6		VD [20]	E3	E4
VD [5]	G1	G5		VD [21]	C3	A3
VD [6]	H2	G2		VD [22]	A1	C2
VD [7]	K6	E1		VD [23]	A3	B5
VD [8]	J5	G4		VD [24]	C5	C4
VD [9]	E1	D2		VD [25]	D6	C6
VD [10]	H6	F2		VD [26]	B4	B4
VD [11]	H4	F5		VD [27]	B6	E7
VD [12]	G3	F3		VD [28]	C7	B3
VD [13]	F2	D4		VD [29]	F8	D6
VD [14]	D2	F4		VD [30]	A5	A5
VD [15]	F4	D3		VD [31]	E9	C7

a. VD[30:27] have internal pull-downs.

Table 4.6 : Pin Assignments for Power and Ground

V _{SS} Pins				V _{DD} Pins					
PBGA		CBGA		PBGA			CBGA		
AB14	N15	A14	R10	A17	H8	W15	A4	G17	U1
AC9	N17	A18	R11	AA11	H18	W19	A8	H1	U14
AC25	P4	B2	R13	AC3	J1	Y24	A11	H5	U16
AD6	P12	B19	R14	*AC5	J9		A13	H18	U17
*AE1	P14	C1	T16	AC17	J17		C3	H20	V3
AE13	R11	F7	U4	AC23	J23		C5	K1	*V5
J13	R13	F11	*U5	AD2	L5		C8	L20	V7
K10	R15	G1	V9	AD24	L19		C14	N1	V13
K16	T6	G15	V20	AE9	R7		C18	N3	V16
L11	T16	K6	W2	B24	R21		D5	N16	V18
L13	T24	L6	W12	C9	T10		D7	N20	Y8
L15	U11	L15	W19	D4	U3		D20	P4	Y10
M12	U13	M2	Y3	D22	U9		E13	P18	Y13
M14	V24	N6	Y5	E5	U17		E17	R6	Y17
N1	Y10	N17	Y6	E15	U25		E18	R15	
N9	Y16	P6	Y7	G7	V8		F6	T3	
N11		P19		G11	V18		F15	T4	
N13		P20		G19	W7		G3	T8	

*AVDD and AVSS are power pins specifically used for powering the analog circuitry in the Universe II. Extra care should be taken to avoid noise and ground shifting on these pins through the use of decoupling capacitors or isolated ground and power planes.



Table 4.7 and Table 4.8 below are tables that map pin numbers to signal names. These tables should not be read as figures. For layout purposes, please see Appendix-G.

Table 4.7 : Pinout for 313-pin Plastic BGA Package

	A	B	C	D	E	F	G	H	J	K	L	M	N	P	R	T	U	V	W	Y	AA	AB	AC	AD	AE	
1	VD[22]		VD[19]		VD[9]		VD[5]		VDD		INT#[7]		VSS		LRST#		VRBR#[1]		AD[27]		AD[58]		PLL_TESTSEL		AVSS	1
2		VD[18]		VD[14]		VD[13]		VD[6]		VD[2]		VSCON_DIR		VXBBSY		AD[61]		AD[59]		AD[24]		PLL_TESTOUT		VDD		2
3	VD[23]		VD[21]		VD[20]		VD[12]		VD[3]		INT#[5]		VBCLR#		VRBR#[2]		VDD		AD[57]		LCLK		VDD		VCOCTL	3
4		VD[26]		VDD		VD[15]		VD[11]		VD[4]		INT#[4]		VSS		PWRRST#		AD[25]		AD[56]		PERR#		AD[22]		4
5	VD[30]		VD[24]		VDD		VD[17]		VD[8]		VDD		AD[63]		AD[60]		AD[26]		VRBR#[0]		INT#[11]		AVDD		PAR64	5
6		VD[27]		VD[25]		VD[16]		VD[10]		VD[7]		VRBBSY#		AD[30]		VSS		INT#[3]		AD[23]		AD[55]		VSS		6
7	VRBERR#		VD[28]		VIACK#		VDD		VD[0]		VRBR#[3]		VSYSCLK		VDD		AD[28]		VDD		SERR#		DEVSEL#		AD[20]	7
8		VAM_DIR		VWRITE#		VD[29]		VDD		VD[1]		INT#[6]		PAR		AD[29]		VDD		AD[21]		AD[54]		TRDY#		8
9	VAM[5]		VDD		VD[31]		VAM[2]		VDD		INT#[2]		VSS		AD[62]		VDD		AD[18]		AD[53]		VSS		VDD	9
10		VAM[3]		VAM[1]		VD_DIR		VAM[4]		VSS		VXSYSFAIL		AD[31]		VDD		AD[51]		VSS		AD[52]		AD[19]		10
11	VDS#[1]		TMS		VAM[0]		VDD		VDS_DIR		VSS		VSS		VSS		VSS		ACK64#		VDD		AD[50]		AD[16]	11
12		VOE#		VXBERR		VDS#[0]		TCK		VAS_DIR		VSS		VSS		CBE[6]		AD[48]		AD[49]		AD[17]		CBE[7]		12
13	TDI		TDO		TRST#		VA_DIR		VSS		VSS		VSS		VSS		VSS		CBE[3]		TMODE[0]		VRSYS-FAIL#		VSS	13
14		VAS#		VA[5]		VA[3]		VA[1]		VLWORD#		VSS		VSS		CBE[2]		CBE[1]		CBE[0]		VSS		CBE[5]		14
15	VA[2]		VSLAVE_DIR		VDD		VDTACK#		VA[4]		VSS		VSS		VSS		AD[15]		VDD		AD[14]		IRDY#		CBE[4]	15
16		VA[8]		VA[10]		VA[13]		VA[7]		VSS		VBGI#[1]		AD[0]		VSS		VRIRQ#[5]		VSS		IDSEL		AD[47]		16
17	VDD		VA[9]		VA[14]		VA[6]		VDD		AD[32]		VSS		VRIRQ#[7]		VDD		FRAME#		AD[13]		VDD		GNT#	17
18		VA[12]		VA[17]		VA[16]		VDD		VXIRQ[3]		VBGO#[2]		VRACFAIL#		AD[3]		VDD		AD[45]		STOP#		REQ64#		18
19	VA[11]		VA[18]		VA[23]		VDD		VXIRQ[1]		VDD		AD[33]		AD[2]		AD[35]		VDD		RST#		AD[12]		AD[46]	19
20		VA[19]		VA[21]		VA[22]		VRIRQ#[2]		INT#[0]		VBGO#[0]		VRIRQ#[6]		AD[5]		AD[39]		AD[43]		AD[11]		AD[44]		20
21	VA[15]		VRSYSRST#		VA[28]		VA[27]		VRIRQ#[4]		VIACKO#		VBGI#[0]		VDD		AD[38]		AD[41]		TMODE[1]		AD[10]		ENID	21
22		VA[20]		VDD		VA[29]		VRIRQ#[1]		REQ#		VXIRQ[6]		AD[1]		AD[4]		VME_RST#		AD[40]		AD[8]		AD[42]		22
23	VXSYSRST		CLK64		VA[25]		VXBR[3]		VDD		VXIRQ[5]		VBGI#[3]		AD[34]		AD[36]		TMODE[2]		LOCK#		VDD		VIACKI#	23
24		VDD		VA[30]		VA[31]		VXBR[1]		VXIRQ[2]		VBGO#[3]		VXBR[2]		VSS		VSS		VDD		AD[7]		VDD		24
25	VA[24]		VA[26]		VRIRQ#[3]		VXBR[0]		VXIRQ[4]		VBGO#[1]		VBGI#[2]		VXIRQ[7]		VDD		AD[37]		AD[6]		VSS		AD[9]	25

Table 4.8 : Pinout for 324-pin Ceramic BGA Package

	A	B	C	D	E	F	G	H	J	K	L	M	N	P	R	T	U	V	W	Y	
1			VSS	VD[1]	VD[7]	VD[3]	VSS	VDD	VSCON_DIR	VDD	PAR	LRST#	VDD	VRBR#[1]	PWRRST#	AD[58]	VDD	PLL_TESTOUT			1
2		VSS	VD[22]	VD[9]	VD[16]	VD[10]	VD[6]	VRBR#[3]	VRBBSY#	VSYCLK	AD[63]	VSS	AD[61]	AD[27]	AD[25]	PLL_TESTSEL	VRBR#[0]	AD[57]	VSS		2
3	VD[21]	VD[28]	VDD	VD[15]	VD[19]	VD[12]	VDD	VD[0]	INT#[2]	VXSYSFAIL	VXBBSY	VRBR#[2]	VDD	AD[26]	AD[56]	VDD	INT#[1]	VDD	LCLK	VSS	3
4	VDD	VD[26]	VD[24]	VD[13]	VD[20]	VD[14]	VD[8]	VD[2]	INT#[4]	INT#[7]	AD[31]	AD[29]	AD[60]	VDD	INT#[3]	VDD	VSS	AD[55]	PAR64	AD[51]	4
5	VD[30]	VD[23]	VDD	VDD	VD[18]	VD[11]	VD[5]	VDD	INT#[6]	VBCLR#	AD[30]	AD[28]	AD[59]	AD[24]	AD[23]	VCOCTL	AVSS	AVDD	PERR#	VSS	5
6	VAM[1]	VWRITE#	VD[25]	VD[29]	VIACK#	VDD	VD[17]	VD[4]	INT#[5]	VSS	VSS	AD[62]	VSS	VSS	VDD	AD[54]	AD[22]	DEVSEL#	AD[21]	VSS	6
7	VAM[5]	VRBERR#	VD[31]	VDD	VD[27]	VSS									SERR#	AD[20]	AD[53]	VDD	TRDY#	VSS	7
8	VDD	VAM[3]	VDD	VAM[0]	VAM_DIR	VD_DIR									AD[52]	VDD	AD[19]	AD[18]	ACK64#	VDD	8
9	VDS#[1]	VXBERR	TMS	VAM[4]	VAM[2]	VDS_DIR									AD[50]	AD[48]	AD[17]	VSS	AD[49]	CBE[7]	9
10	TCK	TRST#	VOE#	VAS_DIR	VDS#[0]	TDI									VSS	VRSYSFAIL#	AD[16]	CBE[6]	CBE[3]	VDD	10
11	VDD	VA_DIR	VLWORD#	VA[2]	TDO	VSS									VSS	CBE[0]	CBE[2]	CBE[5]	TMODE[0]	CBE[1]	11
12	VAS#	VA[3]	VA[4]	VA[5]	VA[1]	VSLAVE_DIR									AD[15]	FRAME#	VRIRQ#[5]	IRDY#	VSS	CBE[4]	12
13	VDD	VDTACK#	VA[6]	VA[8]	VDD	VA[10]									VSS	REQ64#	AD[14]	VDD	AD[47]	VDD	13
14	VSS	VA[12]	VDD	VA[14]	VA[11]	VA[23]									VSS	AD[44]	VDD	AD[13]	AD[46]	GNT#	14
15	VA[9]	VA[16]	VA[18]	VA[20]	VA[22]	VDD	VSS	INT#[0]	AD[32]	VBGI#[2]	VSS	AD[34]	AD[4]	AD[41]	VDD	AD[43]	AD[45]	AD[11]	STOP#	IDSEL	15
16	VA[13]	VA[17]	VRSYSRST#	CLK64	VXSYSRST	VA[27]	VA[31]	VXBR[1]	VXIRQ[1]	VBGO#[0]	VBGI#[3]	VRACFAIL#	VDD	AD[37]	AD[8]	VSS	VDD	VDD	VIACK#	ENID	16
17	VA[7]	VA[15]	VA[21]	VA[19]	VDD	VRIRQ#[2]	VDD	VIACKO#	VXIRQ[3]	VBGI#[1]	VXIRQ[7]	AD[1]	VSS	AD[38]	AD[7]	AD[9]	VDD	TMODE[1]	RST#	VDD	17
18	VSS	VA[25]	VDD	VA[28]	VDD	VA[29]	VRIRQ#[4]	VDD	VXIRQ[5]	VBGO#[3]	AD[0]	VRIRQ#[7]	AD[35]	VDD	TMODE[2]	LOCK#	AD[40]	VDD	AD[12]	AD[42]	18
19		VSS	VXBR[3]	VA[30]	VA[26]	VRIRQ#[1]	VXBR[0]	VXIRQ[2]	VXIRQ[6]	VBGI#[0]	AD[33]	VRIRQ#[6]	AD[2]	VSS	AD[39]	VME_RESET#	AD[6]	AD[10]	VSS		19
20			VA[24]	VDD	VRIRQ#[3]	REQ#	VXIRQ[4]	VDD	VBGO#[1]	VBGO#[2]	VDD	VXBR[2]	VDD	VSS	AD[36]	AD[5]	AD[3]	VSS			20



Table 4.7 and Table 4.8 map pin numbers to signal names. These tables should not be read as figures. For layout purposes, please see Appendix-G.

Appendix A Registers

The Universe II Control and Status Registers facilitate host system configuration and allow the user to control Universe II operational characteristics. The registers are divided into three groups:

- PCI Configuration Space,
- VMEbus Configuration and Status Registers, and
- Universe II Device Specific Status Registers.

The Universe II registers have little-endian byte-ordering.

Figure A.1 below summarizes the supported register access mechanisms.

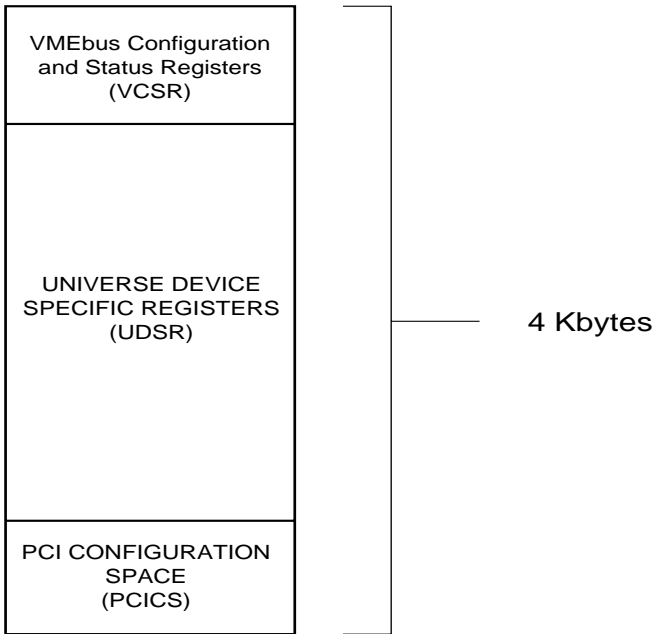


Figure A.1 : UCSR Access Mechanisms

The bit combinations listed as "Reserved" must not be programmed. All bits listed as "Reserved" must read back a value of zero.

Table A.1 below lists the Universe II registers by address offset. The tables following the register map (Table A.2 to Table A.127) provide detailed descriptions of each register.

Table A.1 : Universe II Register Map

Offset	Register	Name
000	PCI Configuration Space ID Register	PCI_ID
004	PCI Configuration Space Control and Status Register	PCI_CSR
008	PCI Configuration Class Register	PCI_CLASS
00C	PCI Configuration Miscellaneous 0 Register	PCI_MISC0
010	PCI Configuration Base Address Register	PCI_BS0
014	PCI Configuration Base Address 1 Register	PCI_BS1
018-024	PCI Unimplemented	
028	PCI Reserved	
02C	PCI Reserved	
030	PCI Unimplemented	
034	PCI Reserved	
038	PCI Reserved	
03C	PCI Configuration Miscellaneous 1 Register	PCI_MISC1
040-0FF	PCI Unimplemented	
100	PCI Target Image 0 Control Register	LSI0_CTL
104	PCI Target Image 0 Base Address Register	LSI0_BS
108	PCI Target Image 0 Bound Address Register	LSI0_BD
10C	PCI Target Image 0 Translation Offset Register	LSI0_TO
110	Reserved	
114	PCI Target Image 1 Control Register	LSI1_CTL
118	PCI Target Image 1 Base Address Register	LSI1_BS
11C	PCI Target Image 1 Bound Address Register	LSI1_BD
120	PCI Target Image 1 Translation Offset Register	LSI1_TO
124	Reserved	
128	PCI Target Image 2 Control Register	LSI2_CTL
12C	PCI Target Image 2 Base Address Register	LSI2_BS
130	PCI Target Image 2 Bound Address Register	LSI2_BD
134	PCI Target Image 2 Translation Offset Register	LSI2_TO
138	Reserved	
13C	PCI Target Image 3 Control Register	LSI3_CTL
140	PCI Target Image 3 Base Address Register	LSI3_BS
144	PCI Target Image 3 Bound Address Register	LSI3_BD
148	PCI Target Image 3 Translation Offset Register	LSI3_TO

Table A.1 : Universe II Register Map (Continued)

Offset	Register	Name
14C-16C	Reserved	
170	Special Cycle Control Register	SCYC_CTL
174	Special Cycle PCI Bus Address Register	SCYC_ADDR
178	Special Cycle Swap/Compare Enable Register	SCYC_EN
17C	Special Cycle Compare Data Register	SCYC_CMP
180	Special Cycle Swap Data Register	SCYC_SWP
184	PCI Miscellaneous Register	LMISC
188	Special PCI Target Image Register	SLSI
18C	PCI Command Error Log Register	L_CMDERR
190	PCI Address Error Log Register	LAERR
194-19C	Reserved	
1A0	PCI Target Image 4 Control Register	LSI4_CTL
1A4	PCI Target Image 4 Base Address Register	LSI4_BS
1A8	PCI Target Image 4 Bound Address Register	LSI4_BD
1AC	PCI Target Image 4 Translation Offset Register	LSI4_TO
1B0	Reserved	
1B4	PCI Target Image 5 Control Register	LSI5_CTL
1B8	PCI Target Image 5 Base Address Register	LSI5_BS
1BC	PCI Target Image 5 Bound Address Register	LSI5_BD
1C0	PCI Slave Image 5 Translation Offset Register	LSI5_TO
1C4	Reserved	
1C8	PCI Target Image 6 Control Register	LSI6_CTL
1CC	PCI Target Image 6 Base Address Register	LSI6_BS
1D0	PCI Target Image 6 Bound Address Register	LSI6_BD
1D4	PCI Target Image 6 Translation Offset Register	LSI6_TO
1D8	Reserved	
1DC	PCI Target Image 7 Control Register	LSI7_CTL
1E0	PCI Target Image 7 Base Address Register	LSI7_BS
1E4	PCI Target Image 7 Bound Address Register	LSI7_BD
1E8	PCI Target Image 7 Translation Offset Register	LSI7_TO
1EC-1FC	Reserved	
200	DMA Transfer Control Register	DCTL
204	DMA Transfer Byte Count Register	DTBC
208	DMA PCI Bus Address Register	DLA
20C	Reserved	
210	DMA VMEbus Address Register	DVA
214	Reserved	

Table A.1 : Universe II Register Map (Continued)

Offset	Register	Name
218	DMA Command Packet Pointer Register	DCPP
21C	Reserved	
220	DMA General Control and Status Register	DGCS
224	DMA Linked List Update Enable Register	D_LLUE
228-2FC	Reserved	
300	PCI Interrupt Enable Register	LINT_EN
304	PCI Interrupt Status Register	LINT_STAT
308	PCI Interrupt Map 0 Register	LINT_MAP0
30C	PCI Interrupt Map 1 Register	LINT_MAP1
310	VMEbus Interrupt Enable Register	VINT_EN
314	VMEbus Interrupt Status Register	VINT_STAT
318	VMEbus Interrupt Map 0 Register	VINT_MAP0
31C	VMEbus Interrupt Map 1 Register	VINT_MAP1
320	Interrupt Status/ID Out Register	STATID
324	VIRQ1 STATUS/ID Register	V1_STATID
328	VIRQ2 STATUS/ID Register	V2_STATID
32C	VIRQ3 STATUS/ID Register	V3_STATID
330	VIRQ4 STATUS/ID Register	V4_STATID
334	VIRQ5 STATUS/ID Register	V5_STATID
338	VIRQ6 STATUS/ID Register	V6_STATID
33C	VIRQ7 STATUS/ID Register	V7_STATID
340	PCI Interrupt Map 2 Register	LINT_MAP2
344	VME Interrupt Map 1 Register	VINT_MAP2
348	Mailbox 0 Register	MBOX0
34C	Mailbox 1 Register	MBOX1
350	Mailbox 2 Register	MBOX2
354	Mailbox 3 Register	MBOX3
358	Semaphore 0 Register	SEMA0
35C	Semaphore 1 Register	SEMA1
360-3FC	Reserved	
400	Master Control Register	MAST_CTL
404	Miscellaneous Control Register	MISC_CTL
408	Miscellaneous Status Register	MISC_STAT
40C	User AM Codes Register	USER_AM
410-EFC	Reserved	
F00	VMEbus Slave Image 0 Control Register	VSI0_CTL
F04	VMEbus Slave Image 0 Base Address Register	VSI0_BS

Table A.1 : Universe II Register Map (Continued)

Offset	Register	Name
F08	VMEbus Slave Image 0 Bound Address Register	VSI0_BD
F0C	VMEbus Slave Image 0 Translation Offset Register	VSI0_TO
F10	Reserved	
F14	VMEbus Slave Image 1 Control Register	VSI1_CTL
F18	VMEbus Slave Image 1 Base Address Register	VSI1_BS
F1C	VMEbus Slave Image 1 Bound Address Register	VSI1_BD
F20	VMEbus Slave Image 1 Translation Offset Register	VSI1_TO
F24	Reserved	
F28	VMEbus Slave Image 2 Control Register	VSI2_CTL
F2C	VMEbus Slave Image 2 Base Address Register	VSI2_BS
F30	VMEbus Slave Image 2 Bound Address Register	VSI2_BD
F34	VMEbus Slave Image 2 Translation Offset Register	VSI2_TO
F38	Reserved	
F3C	VMEbus Slave Image 3 Control Register	VSI3_CTL
F40	VMEbus Slave Image 3 Base Address Register	VSI3_BS
F44	VMEbus Slave Image 3 Bound Address Register	VSI3_BD
F48	VMEbus Slave Image 3 Translation Offset Register	VSI3_TO
F4C-F60	Reserved	
F64	Location Monitor Control Register	LM_CTL
F68	Location Monitor Base Address Register	LM_BS
F6C	Reserved	
F70	VMEbus Register Access Image Control Register	VRAI_CTL
F74	VMEbus Register Access Image Base Address Register	VRAI_BS
F78	Reserved	
F7C	Reserved	
F80	VMEbus CSR Control Register	VCSR_CTL
F84	VMEbus CSR Translation Offset Register	VCSR_TO
F88	VMEbus AM Code Error Log Register	V_AMERR
F8C	VMEbus Address Error Log Register	VAERR
F90	VMEbus Slave Image 4 Control Register	VSI4_CTL
F94	VMEbus Slave Image 4 Base Address Register	VSI4_BS
F98	VMEbus Slave Image 4 Bound Address Register	VSI4_BD
F9C	VMEbus Slave Image 4 Translation Offset Register	VSI4_TO
FA0	Reserved	
FA4	VMEbus Slave Image 5 Control Register	VSI5_CTL
FA8	VMEbus Slave Image 5 Base Address Register	VSI5_BS
FAC	VMEbus Slave Image 5 Bound Address Register	VSI5_BD

Table A.1 : Universe II Register Map (Continued)

Offset	Register	Name
FB0	VMEbus Slave Image 5 Translation Offset Register	VSI5_TO
FB4	Reserved	
FB8	VMEbus Slave Image 6 Control Register	VSI6_CTL
FBC	VMEbus Slave Image 6 Base Address Register	VSI6_BS
FC0	VMEbus Slave Image 6 Bound Address Register	VSI6_BD
FC4	VMEbus Slave Image 6 Translation Offset Register	VSI6_TO
FC8	Reserved	
FCC	VMEbus Slave Image 7 Control Register	VSI7_CTL
FD0	VMEbus Slave Image 7 Base Address Register	VSI7_BS
FD4	VMEbus Slave Image 7 Bound Address Register	VSI7_BD
FD8	VMEbus Slave Image 7 Translation Offset Register	VSI7_TO
FDC-FEC	Reserved	
FF0	VME CR/CSR Reserved	
FF4	VMEbus CSR Bit Clear Register	VCSR_CLR
FF8	VMEbus CSR Bit Set Register	VCSR_SET
FFC	VMEbus CSR Base Address Register	VCSR_BS

Table A.2 : PCI Configuration Space ID Register (PCI_ID)

Register Name: PCI_ID		Offset:000
Bits	Function	
31-24	DID	
23-16	DID	
15-08	VID	
07-00	VID	

PCI_ID Description

Name	Type	Reset By	Reset State	Function
DID[15:0]	R	all	0	Device ID - Tundra allocated device identifier
VID[15:0]	R	all	10E3	Vendor ID - PCI SIG allocated vendor identifier

Table A.3 : PCI Configuration Space Control and Status Register (PCI_CSR)

Register Name: PCI_CSR							Offset:004	
Bits	Function							
31-24	D_PE	S_SERR	R_MA	R_TA	S_TA	DEVSEL		DP_D
23-16	TFBBC	PCI Reserved						
15-08	PCI Reserved						MFBBC	SERR_EN
07-00	WAIT	PERESP	VGAPS	MWI_EN	SC	BM	MS	IOS

PCI_CSR Description

Name	Type	Reset By	Reset State	Function
D_PE	R/Write 1 to Clear	all	0	Detected Parity Error 0=No parity error, 1=Parity error This bit is always set by the Universe II when the PCI master interface detects a data parity error or the PCI target interface detects address or data parity errors.
S_SERR	R/Write 1 to Clear	all	0	Signalled SERR# 0=SERR# not asserted, 1=SERR# asserted. The Universe II PCI target interface sets this bit when it asserts SERR# to signal an address parity error. SERR_EN must be set before SERR# can be asserted.
R_MA	R/Write 1 to Clear	all	0	Received Master-Abort 0=Master did not generate Master-Abort, 1=Master generated Master-Abort The Universe II PCI master interface sets this bit when a transaction it initiated had to be terminated with a Master-Abort.
R_TA	R/Write 1 to Clear	all	0	Received Target-Abort 0=Master did not detect Target-Abort, 1=Master detected Target-Abort. The Universe II PCI master interface sets this bit when a transaction it initiated was terminated with a Target-Abort.
S_TA	R/Write 1 to Clear	all	0	Signalled Target-Abort 0=Target did not terminate transaction with Target-Abort, 1=Target terminated transaction with Target-Abort.
DEVSEL	R	all	01	Device Select Timing The Universe II is a medium speed device
DP_D	R/Write 1 to Clear	all		Data Parity Detected 0=Master did not detect/generate data parity error, 1=Master detected/generated data parity error. The Universe II PCI master interface sets this bit if the Parity Error Response bit is set, it is the master of transaction in which it asserts PERR#, or the addressed target asserts PERR#.
TFBBC	R	all	0	Target Fast Back to Back Capable Universe II cannot accept Back to Back cycles from a different agent.

PCI_CSR Description

Name	Type	Reset By	Reset State	Function
MFBBC	R	all	0	Master Fast Back to Back Enable 0=no fast back-to-back transactions The Universe II master never generates fast back to back transactions.
SERR_EN	R/W	all	0	SERR# Enable 0=Disable SERR# driver, 1=Enable SERR# driver. Setting this and PERESP allows the Universe II PCI target interface to report address parity errors with SERR#.
WAIT	R	all	0	Wait Cycle Control 0=No address/data stepping
PERESP	R/W	all	0	Parity Error Response 0=Disable, 1=Enable Controls the Universe II response to data and address parity errors. When enabled, it allows the assertion of PERR# to report data parity errors. When this bit and SERR_EN are asserted, the Universe II can report address parity errors on SERR#. Universe II parity generation is unaffected by this bit.
VGAPS	R	all	0	VGA Palette Snoop 0=Disable The Universe II treats palette accesses like all other accesses.
MWI_EN	R	all	0	Memory Write and Invalidate Enable 0=Disable The Universe II PCI master interface never generates a Memory Write and Invalidate command.
SC	R	all	0	Special Cycles 0=Disable The Universe II PCI target interface never responds to special cycles.
BM	R/W	PWR VME	see note 1	Master Enable 0=Disable, 1=Enable For a VMEbus slave image to respond to an incoming cycle, this bit must be set. If this bit is cleared while there is data in the VMEbus Slave Posted Write FIFO, the data will be written to the PCI bus but no further data will be accepted into this FIFO until the bit is set.
MS	R/W	PWR VME	see note 1	Target Memory Enable 0=Disable, 1=Enable
IOS	R/W	PWR VME	see note 1	Target IO Enable 0=Disable, 1=Enable

If the VCSR or LSI0 power-up options are enabled, these bits are not disabled after reset.

The Universe II only rejects PCI addresses with parity errors in the event that both the PERESP and SERR_EN bits are programmed to a value of 1.

Table A.4 : PCI Configuration Class Register (PCI_CLASS)

Register Name: PCI_CLASS		Offset:008
Bits	Function	
31-24	BASE	
23-16	SUB	
15-08	PROG	
07-00	RID	

PCI_CLASS Description

Name	Type	Reset By	Reset State	Function
BASE [7:0]	R	all	06	Base Class Code The Universe II is defined as a PCI bridge device
SUB [7:0]	R	all	80	Sub Class Code The Universe II sub-class is "other bridge device"
PROG [7:0]	R	all	00	Programming Interface The Universe II does not have a standardized register-level programming interface
RID [7:0]	R	all	01	Revision ID

Table A.5 : PCI Configuration Miscellaneous 0 Register (PCI_MISC0)

Register Name: PCI_MISC0					Offset:00C		
Bits	Function						
31-24	BISTC	SBIST	PCI Reserved	CCODE			
23-16	MFUNCT	LAYOUT					
15-08	LTIMER				0	0	0
07-00	PCI Unimplemented						

PCI_MISC0 Description

Name	Type	Reset By	Reset State	Function
BISTC	R	all	0	The Universe II is not BIST Capable
SBIST	R	all	0	Start BIST The Universe II is not BIST capable
CCODE	R	all	0	Completion Code The Universe II is not BIST capable
MFUNCT	R	all	0	Multifunction Device 0=No, 1=Yes The Universe II is not a multi-function device.
LAYOUT	R	all	0	Configuration Space Layout
LTIMER [7:3]	R/W	all	0	Latency Timer: The latency timer has a resolution of 8 clocks

The Universe II is not a multi-function device.

Table A.6 : PCI Configuration Base Address Register (PCI_BS0)

Register Name: PCI_BS0								Offset:010
Bits	Function							
31-24	BS							
23-16	BS							
15-08	BS				0	0	0	0
07-00	0	0	0	0	0	0	0	SPACE

PCI_BS0 Description

Name	Type	Reset By	Reset State	Function
BS[31:12]	R/W	all	0	Base Address
SPACE	R	all	Power-up Option	PCI Bus Address Space 0=Memory, 1=I/O

This register specifies the 4 Kbyte aligned base address of the 4 Kbyte Universe II register space on PCI.

A power-up option determines if the registers are mapped into Memory or I/O space in relation to this base address. (See “Power-Up Options” on page 2-115). If mapped into Memory space, the user is free to locate the registers anywhere in the 32-bit address space.

- When the VA[1] pin is sampled low at power-up, the PCI_BS0 register’s SPACE bit is set to “1”, which signifies I/O space, and the PCI_BS1 register’s SPACE bit is set to “0”, which signifies memory space.
- When VA[1] is sampled high at power-up, the PCI_BS0 register’s SPACE register’s bit is set to “0”, which signifies Memory space, and the PCI_BS1 register’s SPACE bit is set to “1”, which signifies I/O space.

A write must occur to this register before the Universe II Device Specific Registers can be accessed. This write can be performed with a PCI configuration transaction or a VMEbus register access.

Table A.7 : PCI Configuration Base Address 1 Register (PCI_BS1)

Register Name: PCI_BS1								Offset:014
Bits	Function							
31-24	BS							
23-16	BS							
15-08	BS				0	0	0	0
07-00	0	0	0	0	0	0	0	SPACE

PCI_BS1 Description

Name	Type	Reset By	Reset State	Function
BS[31:12]	R/W	all	0	Base Address
SPACE	R	all	Power-up Option	PCI Bus Address Space 0=Memory, 1=I/O

This register specifies the 4 KByte aligned base address of the 4 KByte Universe II register space in PCI.

A power-up option determines the value of the SPACE bit. This determines whether the registers are mapped into Memory or I/O space in relation to this base address. (See “Power-Up Options” on page 2-115). If mapped into Memory space, the user is free to locate the Universe registers anywhere in the 32-bit address space. If PCI_BS0 is mapped to Memory space, PCI_BS1 is mapped to I/O space; if PCI_BS0 is mapped to I/O space, then PCI_BS1 is mapped to Memory space.

- When the VA[1] pin is sampled low at power-up, the PCI_BS0 register’s SPACE bit is set to “1”, which signifies I/O space, and the PCI_BS1 register’s SPACE bit is set to “0”, which signifies memory space.
- When VA[1] is sampled high at power-up, the PCI_BS0 register’s SPACE register’s bit is set to “0”, which signifies Memory space, and the PCI_BS1 register’s SPACE bit is set to “1”, which signifies I/O space.

A write must occur to this register before the Universe II Device Specific Registers can be accessed. This write can be performed with a PCI configuration transaction or a VMEbus register access.

The SPACE bit in this register is an inversion of the SPACE field in PCI_BS0.

Table A.8 : PCI Configuration Miscellaneous 1 Register (PCI_MISC1)

Register Name: PCI_MISC1		Offset:03C
Bits	Function	
31-24	MAX_LAT [7:0]	
23-16	MIN_GNT [7:0]	
15-08	INT_PIN [7:0]	
07-00	INT_LINE [7:0]	

PCI_MISC1 Description

Name	Type	Reset By	Reset State	Function
MAX_LAT[7:0]	R	all	0	Maximum Latency: This device has no special latency requirements
MIN_GNT[7:0]	R	all	00000011	Minimum Grant:.250 ns units
INT_PIN[7:0]	R	all	00000001	Interrupt Pin: Universe II pin INT# [0] has a PCI compliant I/O buffer
INT_LINE[7:0]	R/W	all	0	Interrupt Line: used by some PCI systems to record interrupt routing information

The MIN_GNT parameter assumes the Universe II master is transferring an aligned burst size of 64 bytes to a 32-bit target with no wait states. This would require roughly 20 clocks (at a clock frequency of 33 MHz, this is about 600 ns). MIN_GNT is set to three, or 750 ns.

Table A.9 : PCI Target Image 0 Control (LSI0_CTL)

Register Name: LSI0_CTL					Offset:100
Bits	Function				
31-24	EN	PWEN	Reserved		
23-16	VDW		Reserved		VAS
15-08	Reserved	PGM	Reserved	SUPER	Reserved VCT
07-00	Reserved				LAS

LSI0_CTL Description

Name	Type	Reset By	Reset State	Function
EN	R/W	all	Power-up Option	Image Enable 0=Disable, 1=Enable
PWEN	R/W	all	0	Posted Write Enable 0=Disable, 1=Enable
VDW	R/W	all	10	VMEbus Maximum Datawidth 00=8-bit data width, 01=16 bit data width, 10=32-bit data width, 11=64-bit data width
VAS	R/W	all	Power-up Option	VMEbus Address Space 000=A16, 001=A24, 010=A32, 011= Reserved, 100=Reserved 101=CR/CSR, 110=User1, 111=User2
PGM	R/W	all	0	Program/Data AM Code 0=Data, 1=Program
SUPER	R/W	all	0	Supervisor/User AM Code 0=Non-Privileged, 1=Supervisor
VCT	R/W	all	0	VMEbus Cycle Type 0=No BLTs on VMEbus, 1=Single BLTs on VMEbus
LAS	R/W	all	Power-up Option	PCI Bus Memory Space 0=PCI Bus Memory Space, 1=PCI Bus I/O Space

In the PCI Target Image Control register, setting the VCT bit will only have effect if the VAS bits are programmed for A24 or A32 space and the VDW bits are programmed for 8-bit, 16-bit, or 32-bit.

If VAS bits are programmed to A24 or A32 and the VDW bits are programmed for 64-bit, the Universe II may perform MBLT transfers independent of the state of the VCT bit.

The setting of the PWEN bit is ignored if the LAS bit is programmed for PCI Bus I/O Space, forcing all transactions through this image to be coupled.

Table A.10 : PCI Target Image 0 Base Address Register (LSI0_BS)

Register Name: LSI0_BS		Offset:104
Bits	Function	
31-24	BS	
23-16	BS	
15-08	BS	Reserved
07-00	Reserved	

LSI0_BS Description

Name	Type	Reset By	Reset State	Function
BS[31:28]	R/W	all	Power-up Option	Base Address
BS[27:12]	R/W	all	0	Base Address

The base address specifies the lowest address in the address range that will be decoded.

The base address for PCI Target Image 0 and PCI Target Image 4 have a 4Kbyte resolution. PCI Target Images 1, 2, 3, 5, 6, and 7 have a 64Kbyte resolution.

Table A.11 : PCI Target Image 0 Bound Address Register (LSI0_BD)

Register Name: LSI0_BD		Offset:108
Bits	Function	
31-24	BD	
23-16	BD	
15-08	BD	Reserved
07-00	Reserved	

LSI0_BD Description

Name	Type	Reset By	Reset State	Function
BD[31:28]	R/W	all	Power-up Option	Bound Address
BD[27:12]	R/W	all	0	Bound Address

The addresses decoded in a slave image are those which are greater than or equal to the base address and less than the bound register. If the bound address is 0, then the addresses decoded are those greater than or equal to the base address.

The bound address for PCI Target Image 0 and PCI Target Image 4 have a 4Kbyte resolution. PCI Target Images 1, 2, 3, 5, 6, and 7 have a 64Kbyte resolution.

Table A.12 : PCI Target Image 0 Translation Offset (LSI0_TO)

Register Name: LSI0_TO		Offset:10C
Bits	Function	
31-24	TO	
23-16	TO	
15-08	TO	Reserved
07-00	Reserved	

LSI0_TO Description

Name	Type	Reset By	Reset State	Function
TO[31:12]	R/W	all	0	Translation Offset

The translation offset for PCI Target Image 0 and PCI Target Image 4 have a 4Kbyte resolution. PCI Target Images 1, 2, 3, 5, 6, and 7 have a 64Kbyte resolution.

Address bits [31:12] generated on the VMEbus in response to an image decode are a two's complement addition of address bits [31:12] on the PCI Bus and bits [31:12] of the image's translation offset.

Table A.13 : PCI Target Image 1 Control (LSI1_CTL)

Register Name: LSI1_CTL					Offset:114
Bits	Function				
31-24	EN	PWEN	Reserved		
23-16	VDW		Reserved		VAS
15-08	Reserved	PGM	Reserved	SUPER	Reserved VCT
07-00	Reserved				LAS

LSI1_CTL Description

Name	Type	Reset By	Reset State	Function
EN	R/W	all	0	Image Enable 0=Disable, 1=Enable
PWEN	R/W	all	0	Posted Write Enable 0=Disable, 1=Enable
VDW	R/W	all	10	VMEbus Maximum Datawidth 00=8-bit data width, 01=16 bit data width, 10=32-bit data width, 11=64-bit data width
VAS	R/W	all	0	VMEbus Address Space 000=A16, 001=A24, 010=A32, 011= Reserved, 100=Reserved, 101=CR/CSR, 110=User1, 111=User2
PGM	R/W	all	0	Program/Data AM Code 0=Data, 1=Program
SUPER	R/W	all	0	Supervisor/User AM Code 0=Non-Privileged, 1=Supervisor
VCT	R/W	all	0	VMEbus Cycle Type 0=no BLTs on VMEbus, 1=BLTs on VMEbus
LAS	R/W	all	0	PCI Bus Memory Space 0=PCI Bus Memory Space, 1=PCI Bus I/O Space

In the PCI Target Image Control register, setting the VCT bit will only have effect if the VAS bits are programmed for A24 or A32 space and the VDW bits are programmed for 8-bit, 16-bit, or 32-bit.

If VAS bits are programmed to A24 or A32 and the VDW bits are programmed for 64-bit, the Universe II may perform MBLT transfers independent of the state of the VCT bit.

The setting of the PWEN bit is ignored if the LAS bit is programmed for PCI Bus I/O Space, forcing all transactions through this image to be coupled.

Table A.14 : PCI Target Image 1 Base Address Register (LSI1_BS)

Register Name: LSI1_BS		Offset:118
Bits	Function	
31-24	BS	
23-16	BS	
15-08	Reserved	
07-00	Reserved	

LSI1_BS Description

Name	Type	Reset By	Reset State	Function
BS[31:16]	R/W	all	0	Base Address

The base address specifies the lowest address in the address range that will be decoded.

Table A.15 : PCI Target Image 1 Bound Address Register (LSI1_BD)

Register Name: LSI1_BD		Offset:11C
Bits	Function	
31-24	BD	
23-16	BD	
15-08	Reserved	
07-00	Reserved	

LSI1_BD Description

Name	Type	Reset By	Reset State	Function
BD[31:16]	R/W	all	0	Bound Address

The addresses decoded in a slave image are those which are greater than or equal to the base address and less than the bound register. If the bound address is 0, then the addresses decoded are those greater than or equal to the base address.

The bound address for PCI Target Image 0 and PCI Target Image 4 have a 4Kbyte resolution. PCI Target Images 1, 2, 3, 5, 6, and 7 have a 64Kbyte resolution.

Table A.16 : PCI Target Image 1 Translation Offset (LSI1_TO)

Register Name: LSI1_TO		Offset:120
Bits	Function	
31-24	TO	
23-16	TO	
15-08	Reserved	
07-00	Reserved	

LSI1_TO Description

Name	Type	Reset By	Reset State	Function
TO[31:16]	R/W	all	0	Translation offset

Address bits [31:16] generated on the VMEbus in response to an image decode are a two's complement addition of address bits [31:16] on the PCI Bus and bits [31:16] of the image's translation offset.

Table A.17 : PCI Target Image 2 Control (LSI2_CTL)

Register Name: LSI2_CTL					Offset:128
Bits	Function				
31-24	EN	PWEN	Reserved		
23-16	VDW		Reserved		VAS
15-08	Reserved	PGM	Reserved	SUPER	Reserved VCT
07-00	Reserved				LAS

LSI2_CTL Description

Name	Type	Reset By	Reset State	Function
EN	R/W	all	0	Image Enable 0=Disable, 1=Enable
PWEN	R/W	all	0	Posted Write Enable 0=Disable, 1=Enable
VDW	R/W	all	10	VMEbus Maximum Datawidth 00=8-bit data width, 01=16 bit data width, 10=32-bit data width, 11=64-bit data width
VAS	R/W	all	0	VMEbus Address Space 000=A16, 001=A24, 010=A32, 011= Reserved, 100=Reserved, 101=CR/CSR, 110=User1, 111=User2
PGM	R/W	all	0	Program/Data AM Code 0=Data, 1=Program
SUPER	R/W	all	0	Supervisor/User AM Code 0=Non-Privileged, 1=Supervisor
VCT	R/W	all	0	VMEbus Cycle Type 0=no BLTson VMEbus, 1=BLTs on VMEbus
LAS	R/W	all	0	PCI Bus Memory Space 0=PCI Bus Memory Space, 1=PCI Bus I/O Space

In the PCI Target Image Control register, setting the VCT bit will only have effect if the VAS bits are programmed for A24 or A32 space and the VDW bits are programmed for 8-bit, 16-bit, or 32-bit.

If VAS bits are programmed to A24 or A32 and the VDW bits are programmed for 64-bit, the Universe II may perform MBLT transfers independant of the state of the VCT bit.

The setting of the PWEN bit is ignored if the LAS bit is programmed for PCI Bus I/O Space, forcing all transactions through this image to be coupled.

Table A.18 : PCI Target Image 2 Base Address Register (LSI2_BS)

Register Name: LSI2_BS		Offset:12C
Bits	Function	
31-24	BS	
23-16	BS	
15-08	Reserved	
07-00	Reserved	

LSI2_BS Description

Name	Type	Reset By	Reset State	Function
BS[31:16]	R/W	all	0	Base Address

The base address specifies the lowest address in the address range that will be decoded.

Table A.19 : PCI Target Image 2 Bound Address Register (LSI2_BD)

Register Name: LSI2_BD		Offset:130
Bits	Function	
31-24	BD	
23-16	BD	
15-08	Reserved	
07-00	Reserved	

LSI2_BD Description

Name	Type	Reset By	Reset State	Function
BD[31:16]	R/W	all	0	Bound Address

The addresses decoded in a slave image are those which are greater than or equal to the base address and less than the bound register. If the bound address is 0, then the addresses decoded are those greater than or equal to the base address.

Table A.20 : PCI Target Image 2 Translation Offset (LSI2_TO)

Register Name: LSI2_TO		Offset:134
Bits	Function	
31-24	TO	
23-16	TO	
15-08	Reserved	
07-00	Reserved	

LSI2_TO Description

Name	Type	Reset By	Reset State	Function
TO[31:16]	R/W	all	0	Translation offset

Address bits [31:16] generated on the VMEbus in response to an image decode are a two's complement addition of address bits [31:16] on the PCI Bus and bits [31:16] of the image's translation offset.

Table A.21 : PCI Target Image 3 Control (LSI3_CTL)

Register Name: LSI3_CTL					Offset:13C
Bits	Function				
31-24	EN	PWEN	Reserved		
23-16	VDW		Reserved		VAS
15-08	Reserved	PGM	Reserved	SUPER	Reserved VCT
07-00	Reserved				LAS

LSI3_CTL Description

Name	Type	Reset By	Reset State	Function
EN	R/W	all	0	Image Enable 0=Disable, 1=Enable
PWEN	R/W	all	0	Posted Write Enable 0=Disable, 1=Enable
VDW	R/W	all	10	VMEbus Maximum Datawidth 00=8-bit data width, 01=16 bit data width, 10=32-bit data width, 11=64-bit data width
VAS	R/W	all	0	VMEbus Address Space 000=A16, 001=A24, 010=A32, 011= Reserved, 100=Reserved, 101=CR/CSR, 110=User1, 111=User2
PGM	R/W	all	0	Program/Data AM Code 0=Data, 1=Program
SUPER	R/W	all	0	Supervisor/User AM Code 0=Non-Privileged, 1=Supervisor
VCT	R/W	all	0	VMEbus Cycle Type 0=no BLTs on VMEbus, 1=BLTs on VMEbus
LAS	R/W	all	0	PCI Bus Memory Space 0=PCI Bus Memory Space, 1=PCI Bus I/O Space

In the PCI Target Image Control register, setting the VCT bit will only have effect if the VAS bits are programmed for A24 or A32 space and the VDW bits are programmed for 8-bit, 16-bit, or 32-bit.

If VAS bits are programmed to A24 or A32 and the VDW bits are programmed for 64-bit, the Universe II may perform MBLT transfers independent of the state of the VCT bit.

The setting of the PWEN bit is ignored if the LAS bit is programmed for PCI Bus I/O Space, forcing all transactions through this image to be coupled.

Table A.22 : PCI Target Image 3 Base Address Register (LSI3_BS)

Register Name: LSI3_BS		Offset:140
Bits	Function	
31-24	BS	
23-16	BS	
15-08	Reserved	
07-00	Reserved	

LSI3_BS Description

Name	Type	Reset By	Reset State	Function
BS[31:16]	R/W	all	0	Base Address

The base address specifies the lowest address in the address range that will be decoded.

Table A.23 : PCI Target Image 3 Bound Address Register (LSI3_BD)

Register Name: LSI3_BD		Offset:144
Bits	Function	
31-24	BD	
23-16	BD	
15-08	Reserved	
07-00	Reserved	

LSI3_BD Description

Name	Type	Reset By	Reset State	Function
BD[31:16]	R/W	all	0	Bound Address

The addresses decoded in a slave image are those which are greater than or equal to the base address and less than the bound register. If the bound address is 0, then the addresses decoded are those greater than or equal to the base address.

Table A.24 : PCI Target Image 3 Translation Offset (LSI3_TO)

Register Name: LSI3_TO		Offset:148
Bits	Function	
31-24	TO	
23-16	TO	
15-08	Reserved	
07-00	Reserved	

LSI3_TO Description

Name	Type	Reset By	Reset State	Function
TO[31:16]	R/W	all	0	Translation offset

Address bits [31:16] generated on the VMEbus in response to an image decode are a two's complement addition of address bits [31:16] on the PCI Bus and bits [31:16] of the image's translation offset.

Table A.25 : Special Cycle Control Register (SCYC_CTL)

Register Name: SCYC_CTL	Offset: 170
--------------------------------	--------------------

Bits	Function		
31-24	Reserved		
23-16	Reserved		
15-08	Reserved		
07-00	Reserved	LAS	SCYC

SCYC_CTL Description

Name	Type	Reset By	Reset State	Function
LAS	R/W	all	0	PCI Bus Address Space 0=PCI Bus Memory Space, 1=PCI Bus I/O Space
SCYC	R/W	all	0	Special Cycle 00=Disable, 01=RMW, 10=ADOH, 11=Reserved

The special cycle generator will generate an ADOH or RMW cycle for the 32-bit PCI Bus address which matches the programmed address in SCYC_ADDR, in the address space specified in the LAS field of the SCYC_CTL register. A Read-Modify-Write command is initiated by a read to the specified address. Address-Only cycles are initiated by either read or write cycles.

Table A.26 : Special Cycle PCI Bus Address Register (SCYC_ADDR)

Register Name: SCYC_ADDR		Offset: 174
Bits	Function	
31-24	ADDR	
23-16	ADDR	
15-08	ADDR	
07-00	ADDR	Reserved

SCYC_ADDR Description

Name	Type	Reset By	Reset State	Function
ADDR[31:2]	R/W	all	0	Address

This register designates the special cycle address. This address must appear on the PCI Bus during the address phase of a transfer for the Special Cycle Generator to perform its function. Whenever the addresses match, the Universe II does not respond with ACK64#

Table A.27 : Special Cycle Swap/Compare Enable Register (SCYC_EN)

Register Name: SCYC_EN		Offset: 178
Bits	Function	
31-24	EN	
23-16	EN	
15-08	EN	
07-00	EN	

SCYC_EN Description

Name	Type	Reset By	Reset State	Function
EN[31:0]	R/W	all	0	Bit Enable 0=Disable, 1=Enable

The bits enabled in this register determine the bits that will be involved in the compare and swap operations for VME RMW cycles.

Table A.28 : Special Cycle Compare Data Register (SCYC_CMP)

Register Name: SCYC_CMP		Offset: 17C
Bits	Function	
31-24	CMP	
23-16	CMP	
15-08	CMP	
07-00	CMP	

SCYC_CMP Description

Name	Type	Reset By	Reset State	Function
CMP[31:0]	R/W	all	0	The data returned from the VMEbus is compared with the contents of this register.

The data returned from the read portion of a VMEbus RMW is compared with the contents of this register. SCYC_EN is used to control which bits are compared.

Table A.29 : Special Cycle Swap Data Register (SCYC_SWP)

Register Name: SCYC_SWP		Offset: 180
Bits	Function	
31-24	SWP	
23-16	SWP	
15-08	SWP	
07-00	SWP	

SCYC_SWP Description

Name	Type	Reset By	Reset State	Function
SWP[31:0]	R/W	all	0	Swap data

If enabled bits matched with the value in the compare register, then the contents of the swap data register is written back to VME. SCYC_EN is used to control which bits are written back to VME.

Table A.30 : PCI Miscellaneous Register (LMISC)

Register Name: LMISC			Offset:184
Bits	Function		
31-24	CRT[3:0]	Reserved	CWT
23-16	Reserved		
15-08	Reserved		
07-00	Reserved		

SLSI Description

Name	Type	Reset By	Reset State	Function
CRT[3:0]	R/W	all	0000	CRT This field is provided for backward compatibility with the Universe I. It has no effect on the operation of the Universe II.
CWT [2:0]	R/W	all	000	Coupled Window Timer 000=Disable - release after first coupled transaction, 001=16 PCI Clocks, 010=32 PCI Clocks, 011=64 PCI Clocks, 100=128 PCI Clocks, 101=246 PCI Clocks, 110=512 PCI Clocks, others=Reserved

The Universe II uses CWT to determine how long to hold ownership of the VMEbus after processing a coupled transaction. The timer is restarted each time the Universe II processes a coupled transaction. If this timer expires, the PCI Slave Channel releases the VMEbus.

Device behaviour is unpredictable if CWT is changed during coupled cycle activity.

This register can only be set at configuration or after disabling all PCI Slave Images.

Table A.31 : Special PCI Target Image (SLSI)

Register Name: SLSI				Offset:188
Bits	Function			
31-24	EN	PWEN	Reserved	
23-16	VDW			Reserved
15-08	PGM			SUPER
07-00	BS			Reserved LAS

SLSI Description

Name	Type	Reset By	Reset State	Function
EN	R/W	all	0	Image Enable 0=Disable, 1=Enable
PWEN	R/W	all	0	Posted Write Enable 0=Disable, 1=Enable
VDW [3:0]	R/W	all	0	VMEbus Maximum Datawidth Each of the four bits specifies a data width for the corresponding 16 MByte region. Low order bits correspond to the lower address regions. 0=16-bit, 1=32-bit
PGM [3:0]	R/W	all	0	Program/Data AM Code Each of the four bits specifies Program/Data AM code for the corresponding 16 MByte region. Low order bits correspond to the lower address regions. 0=Data, 1=Program
SUPER [3:0]	R/W	all	0	Supervisor/User AM Code Each of the four bits specifies Supervisor/User AM code for the corresponding 16 MByte region. Low order bits correspond to the lower address regions. 0=Non-Privileged, 1=Supervisor
BS [5:0]	R/W	all	0	Base Address Specifies a 64 MByte aligned base address for this 64 MByte image.
LAS	R/W	all	0	PCI Bus Address Space 0=PCI Bus Memory Space, 1=PCI Bus I/O Space

This register fully specifies an A32 capable special PCI Target Image. The base is programmable to a 64 Mbyte alignment, and the size is fixed at 64 Mbytes. Incoming address lines [31:26] (in Memory or I/O) must match this field for the Universe II to decode the access. This special PCI Target Image has lower priority than any other PCI Target Image.

The 64 Mbytes of the SLSI is partitioned into four 16 Mbyte regions, numbered 0 to 3 (0 is at the lowest address). PCI address bits [25:24] are used to select regions. The top 64 Kbyte of each region is mapped to VMEbus A16 space, and the rest of each 16 Mbyte region is mapped to A24 space.

The user can use the PGM, SUPER and VDW fields to specify the AM code and the maximum port size for each region. The PGM field is ignored for the portion of each region mapped to A16 space.

No block transfer AM codes are generated.

Table A.32 : PCI Command Error Log Register (L_CMDERR)

Register Name: L_CMDERR				Offset: 18C
Bits	Function			
31-24	CMDERR		M_ERR	Reserved
23-16	L_STAT	Reserved		
15-08	Reserved			
07-00	Reserved			

L_CMDERR Description

Name	Type	Reset By	Reset State	Function
CMDERR [3:0]	R	all	0111	PCI Command Error Log
M_ERR	R	all	0	Multiple Error Occurred 0=Single error, 1=At least one error has occurred since the logs were frozen.
L_STAT	R/W	all	0	PCI Error Log Status Reads: 0=logs invalid, 1=logs are valid and error logging halted Writes: 0=no effect, 1=clears L_STAT and enables error logging

The Universe II PCI Master Interface is responsible for logging errors under the following conditions:

- a posted write transaction results in a target abort,
- a posted write transaction results in a master abort, or
- a maximum retry counter expires during retry of posted write transaction.

This register logs the command information.

Table A.33 : PCI Address Error Log (LAERR)

Register Name: LAERR		Offset: 190
Bits	Function	
31-24	LAERR	
23-16	LAERR	
15-08	LAERR	
07-00	LAERR	

LAERR Description

Name	Type	Reset By	Reset State	Function
LAERR [31:0]	R	all	0	PCI address error log

The starting address of an errored PCI transaction is logged in this register under the following conditions:

- a posted write transaction results in a target abort,
- a posted write transaction results in a master abort, or
- a maximum retry counter expires during retry of posted write transaction.

Contents are qualified by bit L_STAT of the L_CMDERR register.

Table A.34 : PCI Target Image 4 Control Register (LSI4_CTL)

Register Name: LSI4_CTL							Offset:1A0
Bits	Function						
31-24	EN	PWEN	Reserved				
23-16	VDW		Reserved			VAS	
15-08	Reserved	PGM	Reserved	SUPER	Reserved	VCT	
07-00	Reserved					LAS	

LSI4_CTL Description

Name	Type	Reset By	Reset State	Function
EN	R/W	all	0	Image Enable 0=Disable, 1=Enable
PWEN	R/W	all	0	Posted Write Enable 0=Disable, 1=Enable
VDW	R/W	all	10	VMEbus Maximum Datawidth 00=8-bit data width, 01=16 bit data width, 10=32-bit data width, 11=64-bit data width
VAS	R/W	all	0	VMEbus Address Space 000=A16, 001=A24, 010=A32, 011= Reserved, 100=Reserved, 101=CR/CSR, 110=User1, 111=User2
PGM	R/W	all	0	Program/Data AM Code 0=Data, 1=Program
SUPER	R/W	all	0	Supervisor/User AM Code 0=Non-Privileged, 1=Supervisor
VCT	R/W	all	0	VMEbus Cycle Type 0=no BLTs on VMEbus, 1=BLTs on VMEbus
LAS	R/W	all	0	PCI Bus Memory Space 0=PCI Bus Memory Space, 1=PCI Bus I/O Space

In the PCI Target Image Control register, setting the VCT bit will only have effect if the VAS bits are programmed for A24 or A32 space and the VDW bits are programmed for 8-bit, 16-bit, or 32-bit.

If VAS bits are programmed to A24 or A32 and the VDW bits are programmed for 64-bit, the Universe II may perform MBLT transfers independent of the state of the VCT bit.

The setting of the PWEN bit is ignored if the LAS bit is programmed for PCI Bus I/O Space, forcing all transactions through this image to be coupled.

Table A.35 : PCI Target Image 4 Base Address Register (LSI4_BS)

Register Name: LSI4_BS		Offset:1A4
Bits	Function	
31-24	BS	
23-16	BS	
15-08	BS	Reserved
07-00	Reserved	

LSI4_BS Description

Name	Type	Reset By	Reset State	Function
BS[31:12]	R/W	all	0	Base Address

The base address specifies the lowest address in the address range that will be decoded.

Table A.36 : PCI Target Image 4 Bound Address Register (LSI4_BD)

Register Name: LSI4_BD		Offset:1A8
Bits	Function	
31-24	BD	
23-16	BD	
15-08	BD	Reserved
07-00	Reserved	

LSI4_BD Description

Name	Type	Reset By	Reset State	Function
BD[31:12]	R/W	all	0	Bound Address

The addresses decoded in a slave image are those which are greater than or equal to the base address and less than the bound register. If the bound address is 0, then the addresses decoded are those greater than or equal to the base address.

The bound address for PCI Target Image 0 and PCI Target Image 4 have a 4Kbyte resolution. PCI Target Images 1, 2, 3, 5, 6, and 7 have a 64Kbyte resolution.

Table A.37 : PCI Target Image 4 Translation Offset (LSI4_TO)

Register Name: LSI4_TO		Offset:1B0
Bits	Function	
31-24	TO	
23-16	TO	
15-08	TO	Reserved
07-00	Reserved	

LSI4_TO Description

Name	Type	Reset By	Reset State	Function
TO[31:12]	R/W	all	0	Translation offset

Address bits [31:12] generated on the VMEbus in response to an image decode are a two's complement addition of address bits [31:12] on the PCI Bus and bits [31:12] of the image's translation offset.

Table A.38 : PCI Target Image 5 Control Register (LSI5_CTL)

Register Name: LSI5_CTL					Offset:1B4
Bits	Function				
31-24	EN	PWEN	Reserved		
23-16	VDW		Reserved		VAS
15-08	Reserved	PGM	Reserved	SUPER	Reserved VCT
07-00	Reserved				LAS

LSI5_CTL Description

Name	Type	Reset By	Reset State	Function
EN	R/W	all	0	Image Enable 0=Disable, 1=Enable
PWEN	R/W	all	0	Posted Write Enable 0=Disable, 1=Enable
VDW	R/W	all	10	VMEbus Maximum Datawidth 00=8-bit data width, 01=16 bit data width, 10=32-bit data width, 11=64-bit data width
VAS	R/W	all	0	VMEbus Address Space 000=A16, 001=A24, 010=A32, 011= Reserved, 100=Reserved, 101=CR/CSR, 110=User1, 111=User2
PGM	R/W	all	0	Program/Data AM Code 0=Data, 1=Program
SUPER	R/W	all	0	Supervisor/User AM Code 0=Non-Privileged, 1=Supervisor
VCT	R/W	all	0	VMEbus Cycle Type 0=no BLTs on VMEbus, 1=BLTs on VMEbus
LAS	R/W	all	0	PCI Bus Memory Space 0=PCI Bus Memory Space, 1=PCI Bus I/O Space

In the PCI Target Image Control register, setting the VCT bit will only have effect if the VAS bits are programmed for A24 or A32 space and the VDW bits are programmed for 8-bit, 16-bit, or 32-bit.

If VAS bits are programmed to A24 or A32 and the VDW bits are programmed for 64-bit, the Universe II may perform MBLT transfers independent of the state of the VCT bit.

The setting of the PWEN bit is ignored if the LAS bit is programmed for PCI Bus I/O Space, forcing all transactions through this image to be coupled.

Table A.39 : PCI Target Image 5 Base Address Register (LSI5_BS)

Register Name: LSI5_BS		Offset:1B8
Bits	Function	
31-24	BS	
23-16	BS	
15-08	Reserved	
07-00	Reserved	

LSI5_BS Description

Name	Type	Reset By	Reset State	Function
BS[31:16]	R/W	all	0	Base Address

The base address specifies the lowest address in the address range that will be decoded.

Table A.40 : PCI Target Image 5 Bound Address Register (LSI5_BD)

Register Name: LSI5_BD		Offset:1BC
Bits	Function	
31-24	BD	
23-16	BD	
15-08	Reserved	
07-00	Reserved	

LSI5_BD Description

Name	Type	Reset By	Reset State	Function
BD[31:16]	R/W	all	0	Bound Address

The addresses decoded in a slave image are those which are greater than or equal to the base address and less than the bound register. If the bound address is 0, then the addresses decoded are those greater than or equal to the base address.

The bound address for PCI Target Image 0 and PCI Target Image 4 have a 4Kbyte resolution. PCI Target Images 1, 2, 3, 5, 6, and 7 have a 64Kbyte resolution.

Table A.41 : PCI Target Image 5 Translation Offset (LSI5_TO)

Register Name: LSI5_TO		Offset:1C0
Bits	Function	
31-24	TO	
23-16	TO	
15-08	Reserved	
07-00	Reserved	

LSI5_TO Description

Name	Type	Reset By	Reset State	Function
TO[31:16]	R/W	all	0	Translation offset

Address bits [31:16] generated on the VMEbus in response to an image decode are a two's complement addition of address bits [31:16] on the PCI Bus and bits [31:16] of the image's translation offset.

Table A.42 : PCI Target Image 6 Control Register (LSI6_CTL)

Register Name: LSI6_CTL					Offset:1C8
Bits	Function				
31-24	EN	PWEN	Reserved		
23-16	VDW		Reserved		VAS
15-08	Reserved	PGM	Reserved	SUPER	Reserved VCT
07-00	Reserved				LAS

LSI6_CTL Description

Name	Type	Reset By	Reset State	Function
EN	R/W	all	0	Image Enable 0=Disable, 1=Enable
PWEN	R/W	all	0	Posted Write Enable 0=Disable, 1=Enable
VDW	R/W	all	10	VMEbus Maximum Datawidth 00=8-bit data width, 01=16 bit data width, 10=32-bit data width, 11=64-bit data width
VAS	R/W	all	0	VMEbus Address Space 000=A16, 001=A24, 010=A32, 011= Reserved, 100=Reserved, 101=CR/CSR, 110=User1, 111=User2
PGM	R/W	all	0	Program/Data AM Code 0=Data, 1=Program
SUPER	R/W	all	0	Supervisor/User AM Code 0=Non-Privileged, 1=Supervisor
VCT	R/W	all	0	VMEbus Cycle Type 0=no BLTs on VMEbus, 1=BLTs on VMEbus
LAS	R/W	all	0	PCI Bus Memory Space 0=PCI Bus Memory Space, 1=PCI Bus I/O Space

In the PCI Target Image Control register, setting the VCT bit will only have effect if the VAS bits are programmed for A24 or A32 space and the VDW bits are programmed for 8-bit, 16-bit, or 32-bit.

If VAS bits are programmed to A24 or A32 and the VDW bits are programmed for 64-bit, the Universe II may perform MBLT transfers independent of the state of the VCT bit.

The setting of the PWEN bit is ignored if the LAS bit is programmed for PCI Bus I/O Space, forcing all transactions through this image to be coupled.

Table A.43 : PCI Target Image 6 Base Address Register (LSI6_BS)

Register Name: LSI6_BS		Offset:1CC
Bits	Function	
31-24	BS	
23-16	BS	
15-08	Reserved	
07-00	Reserved	

LSI1_BS Description

Name	Type	Reset By	Reset State	Function
BS[31:16]	R/W	all	0	Base Address

The base address specifies the lowest address in the address range that will be decoded.

Table A.44 : PCI Target Image 6 Bound Address Register (LSI6_BD)

Register Name: LSI6_BD		Offset:1D0
Bits	Function	
31-24	BD	
23-16	BD	
15-08	Reserved	
07-00	Reserved	

LSI6_BD Description

Name	Type	Reset By	Reset State	Function
BD[31:16]	R/W	all	0	Bound Address

The addresses decoded in a slave image are those which are greater than or equal to the base address and less than the bound register. If the bound address is 0, then the addresses decoded are those greater than or equal to the base address.

The bound address for PCI Target Image 0 and PCI Target Image 4 have a 4Kbyte resolution. PCI Target Images 1, 2, 3, 5, 6, and 7 have a 64Kbyte resolution.

Table A.45 : PCI Target Image 6 Translation Offset (LSI6_TO)

Register Name: LSI6_TO		Offset:1D4
Bits	Function	
31-24	TO	
23-16	TO	
15-08	Reserved	
07-00	Reserved	

LSI6_TO Description

Name	Type	Reset By	Reset State	Function
TO[31:16]	R/W	all	0	Translation offset

Address bits [31:16] generated on the VMEbus in response to an image decode are a two's complement addition of address bits [31:16] on the PCI Bus and bits [31:16] of the image's translation offset.

Table A.46 : PCI Target Image 7 Control Register (LSI7_CTL)

Register Name: LSI7_CTL							Offset:1DC
Bits	Function						
31-24	EN	PWEN	Reserved				
23-16	VDW		Reserved			VAS	
15-08	Reserved	PGM	Reserved	SUPER	Reserved	VCT	
07-00	Reserved					LAS	

LSI7_CTL Description

Name	Type	Reset By	Reset State	Function
EN	R/W	all	0	Image Enable 0=Disable, 1=Enable
PWEN	R/W	all	0	Posted Write Enable 0=Disable, 1=Enable
VDW	R/W	all	10	VMEbus Maximum Datawidth 00=8-bit data width, 01=16 bit data width, 10=32-bit data width, 11=64-bit data width
VAS	R/W	all	0	VMEbus Address Space 000=A16, 001=A24, 010=A32, 011= Reserved, 100=Reserved, 101=CR/CSR, 110=User1, 111=User2
PGM	R/W	all	0	Program/Data AM Code 0=Data, 1=Program
SUPER	R/W	all	0	Supervisor/User AM Code 0=Non-Privileged, 1=Supervisor
VCT	R/W	all	0	VMEbus Cycle Type 0=no BLTs on VMEbus, 1=BLTs on VMEbus
LAS	R/W	all	0	PCI Bus Memory Space 0=PCI Bus Memory Space, 1=PCI Bus I/O Space

In the PCI Target Image Control register, setting the VCT bit will only have effect if the VAS bits are programmed for A24 or A32 space and the VDW bits are programmed for 8-bit, 16-bit, or 32-bit.

If VAS bits are programmed to A24 or A32 and the VDW bits are programmed for 64-bit, the Universe II may perform MBLT transfers independent of the state of the VCT bit.

The setting of the PWEN bit is ignored if the LAS bit is programmed for PCI Bus I/O Space, forcing all transactions through this image to be coupled.

Table A.47 : PCI Target Image 7 Base Address Register (LSI7_BS)

Register Name: LSI7_BS		Offset:1E0
Bits	Function	
31-24	BS	
23-16	BS	
15-08	Reserved	
07-00	Reserved	

LSI7_BS Description

Name	Type	Reset By	Reset State	Function
BS[31:16]	R/W	all	0	Base Address

The base address specifies the lowest address in the address range that will be decoded.

Table A.48 : PCI Target Image 7 Bound Address Register (LSI7_BD)

Register Name: LSI7_BD		Offset:1E4
Bits	Function	
31-24	BD	
23-16	BD	
15-08	Reserved	
07-00	Reserved	

LSI7_BD Description

Name	Type	Reset By	Reset State	Function
BD[31:16]	R/W	all	0	Bound Address

The addresses decoded in a slave image are those which are greater than or equal to the base address and less than the bound register. If the bound address is 0, then the addresses decoded are those greater than or equal to the base address.

The bound address for PCI Target Image 0 and PCI Target Image 4 have a 4Kbyte resolution. PCI Target Images 1, 2, 3, 5, 6, and 7 have a 64Kbyte resolution.

Table A.49 : PCI Target Image 7 Translation Offset (LSI7_TO)

Register Name: LSI7_TO		Offset:1E8
Bits	Function	
31-24	TO	
23-16	TO	
15-08	Reserved	
07-00	Reserved	

LSI7_TO Description

Name	Type	Reset By	Reset State	Function
TO[31:16]	R/W	all	0	Translation offset

Address bits [31:16] generated on the VMEbus in response to an image decode are a two's complement addition of address bits [31:16] on the PCI Bus and bits [31:16] of the image's translation offset.

Table A.50 : DMA Transfer Control Register (DCTL)

Register Name: DCTL				Offset:200
Bits	Function			
31-24	L2V	Reserved		
23-16	VDW	Reserved	VAS	
15-08	PGM	SUPER	Reserved	VCT
07-00	LD64EN	Reserved		

DCTL Description

Name	Type	Reset By	Reset State	Function
L2V	R/W	all	0	Direction 0=Transfer from VMEbus to PCI Bus, 1=Transfer from PCI Bus to VMEbus
VDW	R/W	all	0	VMEbus Maximum Datawidth 00=8-bit data width, 01=16 bit data width, 10=32-bit data width, 11=64-bit data width
VAS	R/W	all	0	VMEbus Address Space 000=A16, 001=A24, 010=A32, 011= Reserved, 100=Reserved, 101=Reserved, 110=User1, 111=User2
PGM	R/W	all	0	Program/Data AM Code 00=Data, 01=Program, others=Reserved
SUPER	R/W	all	0	Supervisor/User AM Code 00=Non-Privileged, 01=Supervisor, others=Reserved
VCT	R/W	all	0	VMEbus Cycle Type 0=no BLTs on VMEbus, 1=BLTs on VMEbus
LD64EN	R/W	all	1	Enable 64-bit PCI Bus Transactions 0=Disable, 1=Enable

This register is programmed from either bus or is programmed by the DMAC when it loads the command packet. The DMA only accesses PCI Bus Memory space.

The VCT bit determines whether or not the Universe II VME Master will generate BLT transfers. The value of this bit only has meaning if the address space is A24 or A32 and the data width is not 64 bits. If the data width is 64 bits the Universe II may perform MBLT transfers independent of the state of the VCT bit.

Table A.51 : DMA Transfer Byte Count Register (DTBC)

Register Name: DTBC		Offset:204
Bits	Function	
31-24	Reserved	
23-16	DTBC	
15-08	DTBC	
07-00	DTBC	

DTBC Description

Name	Type	Reset By	Reset State	Function
DTBC[23:0]	R/W	all	0	DMA Transfer Byte Count

This register specifies the number of bytes to be moved by the DMA before the start of the DMA transfer, or the number of remaining bytes in the transfer while the DMA is active. This register is programmed from either bus or is programmed by the DMA Controller when it loads a command packet from a linked-list.

In direct mode the user must reprogram the DTBC register before each transfer.

When using the DMA to perform linked-list transfers, it is essential that the DTBC register contains a value of zero before setting the GO bit of the DGCS register or undefined behaviors may occur.

Table A.52 : DMA PCI Bus Address Register (DLA)

Register Name: DLA		Offset:208
Bits	Function	
31-24	LA	
23-16	LA	
15-08	LA	
07-00	LA	

DLA Description

Name	Type	Reset By	Reset State	Function
LA[31:3]	R/W	all	0	PCI Bus Address
LA[2:0]	R/W	all	0	PCI Bus Address

This register is programmed from either bus or by the DMA Controller when it loads a command packet. In direct mode the user must reprogram the DLA register before each transfer. In linked-list mode, this register is only updated when the DMA is stopped, halted, or at the completion of processing a command packet.

After a Bus Error, a Target-Abort, or a Master-Abort, the value in the DLA register must not be used to reprogram the DMA because it has no useable information. Some offset from its original value must be used.

Address bits [2:0] must be programmed the same as those in the DVA.

Table A.53 : DMA VMEbus Address Register (DVA)

Register Name: DVA		Offset:210
Bits	Function	
31-24	VA	
23-16	VA	
15-08	VA	
07-00	VA	

DVA Description

Name	Type	Reset By	Reset State	Function
VA[31:0]	R/W	all	0	VMEbus Address

This register is programmed from either bus or is programmed by the DMA Controller when it loads a command packet. In direct mode the user must reprogram the DVA register before each transfer. In linked-list operation, this register is only updated when the DMA is stopped, halted, or at the completion of processing a command packet.

After a Bus Error, a Target-Abort, or a Master-Abort, the value in the DLA register must not be used to reprogram the DMA because it has no useable information. Some offset from its original value must be used.

Address bits [2:0] must be programmed the same as those in the DLA.

Table A.54 : DMA Command Packet Pointer (DCPP)

Register Name: DCPP		Offset:218
Bits	Function	
31-24	DCPP	
23-16	DCPP	
15-08	DCPP	
07-00	DCPP	Reserved

DCPP Description

Name	Type	Reset By	Reset State	Function
DCPP[31:5]	R/W	all	0	DMA Command Packet Pointer

This register contains the pointer into the current command packet. Initially it is programmed to the starting packet of the linked-list, and is updated with the address to a new command packet at the completion of a packet. The packets must be aligned to a 32-byte address.

Table A.55 : DMA General Control/Status Register (DGCS)

Register Name: DGCS								Offset: 220
Bits	Function							
31-24	GO	STOP_REQ	HALT_REQ	0	CHAIN	0	0	0
23-16	Reserved	VON			VOFF			
15-08	ACT	STOP	HALT	0	DONE	LERR	VERR	P_ERR
07-00	0	INT_STOP	INT_HALT	0	INT_DONE	INT_LERR	INT_VERR	INT_P_ERR

DGCS Description

Name	Type	Reset By	Reset State	Function
GO	W/Read 0 always	all	0	DMA Go Bit 0=No effect, 1=Enable DMA Transfers
STOP_REQ	W/Read 0 always	all	0	DMA Stop Request 0=No effect, 1=Stop DMA transfer when all buffered data has been written
HALT_REQ	W/Read 0 always	all	0	DMA Halt Request 0=No effect, 1=Halt the DMA transfer at the completion of the current command packet
CHAIN	R/W	all	0	DMA Chaining 0=DMA Direct Mode, 1=DMA Linked List mode
VON [2:0]	R/W	all	0	VMEbus "On" counter 000=Until done, 001=256 bytes, 010=512 bytes, 011=1024 bytes, 100=2048 bytes, 101=4096 bytes, 110=8192 bytes, 111=16384 bytes, others=Reserved
VOFF [3:0]	R/W	all	0	VMEbus "Off" Counter 0000=0μs, 0001=16μs, 0010=32μs, 0011=64μs, 0100=128μs, 0101=256μs, 0110=512μs, 0111=1024μs, 1000=2μs, 1001=4μs, 1010=8μs, others=Reserved The DMA will not re-request the VME Master until this timer expires.
ACT	R	all	0	DMA Active Status Bit 0=Not Active, 1=Active
STOP	R/Write 1 to Clear	all	0	DMA Stopped Status Bit 0=Not Stopped, 1=Stopped
HALT	R/Write 1 to Clear	all	0	DMA Halted Status Bit 0=Not Halted, 1=Halted
DONE	R/Write 1 to Clear	all	0	DMA Done Status Bit 0=Not Complete, 1=Complete
LERR	R/Write 1 to Clear	all	0	DMA PCI Bus Error Status Bit 0=No Error, 1=Error
VERR	R/Write 1 to Clear	all	0	DMA VMEbus Error Status Bit 0=No Error, 1=Error

DGCS Description

Name	Type	Reset By	Reset State	Function
P_ERR	R/Write 1 to Clear	all	0	DMA Programming Protocol Error Status Bit Asserted if PCI master interface disabled or lower three bits of PCI and VME addresses differ 0=No Error, 1=Error
INT_STOP	R/W	all	0	Interrupt when Stopped 0=Disable, 1=Enable
INT_HALT	R/W	all	0	Interrupt when Halted 0=Disable, 1=Enable
INT_DONE	R/W	all	0	Interrupt when Done 0=Disable, 1=Enable
INT_LERR	R/W	all	0	Interrupt on LERR 0=Disable, 1=Enable
INT_VERR	R/W	all	0	Interrupt on VERR 0=Disable, 1=Enable
INT_P_ERR	R/W	all	0	Interrupt on Master Enable Error 0=Disable, 1=Enable

STOP, HALT, DONE, LERR, VERR, and P_ERR must be cleared before the GO bit is enabled.

Table A.56 : DMA Linked List Update Enable Register (D_LLUE)

Register Name: D_LLUE		Offset:224
Bits	Function	
31-24	UPDATE	Reserved
23-16	Reserved	
15-08	Reserved	
07-00	Reserved	

D_LLUE Description

Name	Type	Reset By	Reset State	Function
UPDATE	R/W	all	0	DMA Linked List Update Enable 0=PCI Resource not Updating Linked List 1=PCI Resource Updating Linked List

The PCI Resource must read back a logic 1 in the UPDATE field before proceeding to modify the linked list. After the Linked List has been modified the PCI Resource must clear the UPDATE field by writing a logic 0. the Universe II does not prevent an external master, from the PCI bus or the VMEbus, from writing to the other DMA registers. See “Linked List Updating” on page 2-91.

Table A.57 : PCI Interrupt Enable Register (LINT_EN)

Register Name: LINT_EN								Offset:300
Bits	Function							
31-24	Reserved							
23-16	LM3	LM2	LM1	LM0	MBOX3	MBOX2	MBOX1	MBOX0
15-08	ACFAIL	SYSFAIL	SW_INT	SW_IACK	Reserved	VERR	LERR	DMA
07-00	VIRQ7	VIRQ6	VIRQ5	VIRQ4	VIRQ3	VIRQ2	VIRQ1	VOWN

LINT_EN Description

Name	Type	Reset By	Reset State	Function
LM3	R/W	all	0	Location Monitor 3 Mask 0=LM3 Interrupt Masked, 1=LM3 Interrupt Enabled
LM2	R/W	all	0	Location Monitor 2 Mask 0=LM2 Interrupt Masked, 1=LM2 Interrupt Enabled
LM1	R/W	all	0	Location Monitor 1 Mask 0=LM1 Interrupt Masked, 1=LM1 Interrupt Enabled
LM0	R/W	all	0	Location Monitor 0 Mask 0=LM0 Interrupt Masked, 1=LM0 Interrupt Enabled
MBOX3	R/W	all	0	Mailbox 3 Mask 0=MBOX3 Interrupt Masked, 1=MBOX3 Interrupt Enabled
MBOX2	R/W	all	0	Mailbox 2 Mask 0=MBOX2 Interrupt Masked, 1=MBOX2 Interrupt Enabled
MBOX1	R/W	all	0	Mailbox 1 Mask 0=MBOX1 Interrupt Masked, 1=MBOX1 Interrupt Enabled
MBOX0	R/W	all	0	Mailbox 0 Mask 0=MBOX0 Interrupt Masked, 1=MBOX0 Interrupt Enabled
ACFAIL	R/W	all	0	ACFAIL Interrupt Mask 0=ACFAIL Interrupt masked 1=ACFAIL Interrupt enabled
SYSFAIL	R/W	all	0	SYSFAIL Interrupt Mask 0=SYSFAIL Interrupt masked 1=SYSFAIL Interrupt enabled
SW_INT	R/W	all	0	Local Software Interrupt Mask 0=PCI Software Interrupt masked 1=PCI Software Interrupt enabled A zero-to-one transition will cause the PCI software interrupt to be asserted. Subsequent zeroing of this bit will cause the interrupt to be masked, but will not clear the PCI Software Interrupt Status bit.
SW_IACK	R/W	all	0	“VME Software IACK” Mask 0 = “VME Software IACK” Interrupt masked 1 = “VME Software IACK” Interrupt enabled
VERR	R/W	all	0	PCI VERR Interrupt Mask 0 =PCI VERR Interrupt masked 1=PCI VERR Interrupt enabled

LINT_EN Description

LERR	R/W	all	0	PCI LERR Interrupt Mask 0 =PCI LERR Interrupt masked 1 =PCI LERR Interrupt enabled
DMA	R/W	all	0	PCI DMA Interrupt Mask 0=PCI DMA Interrupt masked 1=PCI DMA Interrupt enabled
VIRQ7-VIRQ1	R/W	all	0	VIRQ _x Interrupt Mask 0=VIRQ _x Interrupt masked 1 =VIRQ _x Interrupt enabled
VOWN	R/W	all	0	VOWN Interrupt Mask 0=VOWN Interrupt masked 1=VOWN Interrupt Enabled

Bits VIRQ7-VIRQ1 enable the Universe II to respond as a VME Interrupt Handler to interrupts on the VIRQ[x] lines. When a VIRQ_x interrupt is enabled, and the corresponding VIRQ[x] pin is asserted, the Universe II requests the VMEbus and performs a VME IACK cycle for that interrupt level. When the interrupt acknowledge cycle completes, the STATUS/ID is stored in the corresponding VINT_ID register, the VIRQ_x bit of the LINT_STAT register is set, and a PCI interrupt is generated. The Universe II does not acquire further interrupt STATUS/ID vectors at the same interrupt level until the VIRQ_x bit in the LINT_STAT register is cleared.

The other bits enable the respective internal or external sources to interrupt the PCI side.

Table A.58 : PCI Interrupt Status Register (LINT_STAT)

Register Name: LINT_STAT								Offset: 304
Bits	Function							
31-24	Reserved							
23-16	LM3	LM2	LM1	LM0	MBOX3	MBOX2	MBOX1	MBOX0
15-08	ACFAIL	SYSFAIL	SW_INT	SW_IACK	Reserved	VERR	LERR	DMA
07-00	VIRQ7	VIRQ6	VIRQ5	VIRQ4	VIRQ3	VIRQ2	VIRQ1	VOWN

LINT_STAT Description

Name	Type	Reset By	Reset State	Function
LM3	R/Write 1 to Clear	all	0	Location Monitor 3 Status/Clear 0=no Location Monitor 3 Interrupt, 1=Location Monitor 3 Interrupt active
LM2	R/Write 1 to Clear	all	0	Location Monitor 2 Status/Clear 0=no Location Monitor 2 Interrupt, 1=Location Monitor 2 Interrupt active
LM1	R/Write 1 to Clear	all	0	Location Monitor 1 Status/Clear 0=no Location Monitor 1 Interrupt, 1=Location Monitor 1 Interrupt active
LM0	R/Write 1 to Clear	all	0	Location Monitor 0 Status/Clear 0=no Location Monitor 0 Interrupt, 1=Location Monitor 0 Interrupt active
MBOX3	R/Write 1 to Clear	all	0	Mailbox 3 Status/Clear 0=no Mailbox 3 Interrupt, 1=Mailbox 3 Interrupt active
MBOX2	R/Write 1 to Clear	all	0	Mailbox 2 Status/Clear 0=no Mailbox 2 Interrupt, 1=Mailbox 2 Interrupt active
MBOX1	R/Write 1 to Clear	all	0	Mailbox 1 Status/Clear 0=no Mailbox 1 Interrupt, 1=Mailbox 1 Interrupt active
MBOX0	R/Write 1 to Clear	all	0	Mailbox 0 Status/Clear 0=no Mailbox 0 Interrupt, 1=Mailbox 0 Interrupt active
ACFAIL	R/Write 1 to Clear	all	0	ACFAIL Interrupt Status/Clear 0=no ACFAIL Interrupt, 1=ACFAIL Interrupt active
SYSFAIL	R/Write 1 to Clear	all	0	SYSFAIL Interrupt Status/Clear 0=no SYSFAIL Interrupt, 1=SYSFAIL Interrupt active
SW_INT	R/Write 1 to Clear	all	0	Local Software Interrupt Status/Clear 0=no PCI Software Interrupt, 1=PCI Software Interrupt active

LINT_STAT Description

SW_IACK	R/Write 1 to Clear	all	0	“VME Software IACK” Status/Clear 0=no “VME Software IACK” Interrupt, 1=“VME Software IACK” Interrupt active
VERR	R/Write 1 to Clear	all	0	Local VERR Interrupt Status/Clear 0=Local VERR Interrupt masked, 1=Local VERR Interrupt enabled
LERR	R/Write 1 to Clear	all	0	Local LERR Interrupt Status/Clear 0=Local LERR Interrupt masked, 1=Local LERR Interrupt enabled
DMA	R/Write 1 to Clear	all	0	Local DMA Interrupt Status/Clear 0=Local DMA Interrupt masked, 1=Local DMA Interrupt enabled
VIRQ7-VIRQ1	R/Write 1 to Clear	all	0	VIRQ _x Interrupt Status/Clear 0=VIRQ _x Interrupt masked, 1=VIRQ _x Interrupt enabled
VOWN	R/Write 1 to Clear	all	0	VOWN Interrupt Status/Clear 0=no VOWN Interrupt masked, 1=VOWN Interrupt enabled

Status bits indicated as “R/Write 1 to Clear” are edge sensitive: the status is latched when the interrupt event occurs. These status bits can be cleared independently of the state of the interrupt source by writing a “1” to the status register. Clearing the status bit does not imply the source of the interrupt is cleared.

However, ACFAIL and SYSFAIL are level-sensitive. Clearing ACFAIL or SYSFAIL while their respective pins are still asserted will have no effect.

Table A.59 : PCI Interrupt Map 0 Register (LINT_MAP0)

Register Name: LINT_MAP0				Offset: 308
Bits	Function			
31-24	Reserved	VIRQ7	Reserved	VIRQ6
23-16	Reserved	VIRQ5	Reserved	VIRQ4
15-08	Reserved	VIRQ3	Reserved	VIRQ2
07-00	Reserved	VIRQ1	Reserved	VOWN

LINT_MAP0 Description

Name	Type	Reset By	Reset State	Function
VIRQ7-VIRQ1	R/W	all	0	PCI interrupt destination (LINT[7:0]) for VIRQx
VOWN	R/W	all	0	VMEbus ownership bit interrupt map to PCI interrupt

This register maps various interrupt sources to one of the eight PCI interrupt pins. A value of 000 maps the corresponding interrupt source to LINT# [0], a value of 001 maps to LINT# [1], etc.

Table A.60 : PCI Interrupt Map 1 Register (LINT_MAP1)

Register Name: LINT_MAP1				Offset: 30C
Bits	Function			
31-24	Reserved	ACFAIL	Reserved	SYSFAIL
23-16	Reserved	SW_INT	Reserved	SW_IACK
15-08	Reserved			VERR
07-00	Reserved	LERR	Reserved	DMA

LINT_MAP1 Description

Name	Type	Reset By	Reset State	Function
ACFAIL	R/W	all	0	ACFAIL interrupt destination
SYSFAIL	R/W	all	0	SYSFAIL interrupt destination
SW_INT	R/W	all	0	PCI software interrupt destination
SW_IACK	R/W	all	0	VMEbus Software IACK interrupt destination
VERR	R/W	all	0	VMEbus Error interrupt destination
LERR	R/W	all	0	PCI Bus Error interrupt destination
DMA	R/W	all	0	DMA interrupt destination

This register maps various interrupt sources to one of the eight PCI interrupt pins. A value of 000 maps the corresponding interrupt source to LINT# [0], a value of 001 maps to LINT# [1], etc.

Table A.61 : VMEbus Interrupt Enable Register (VINT_EN)

Register Name: VINT_EN								Offset:310
Bits	Function							
31-24	SW_INT7	SW_INT6	SW_INT5	SW_INT4	SW_INT3	SW_INT2	SW_INT1	Reserved
23-16	Reserved				MBOX3	MBOX2	MBOX1	MBOX0
15-08	Reserved			SW_INT	Reserved	VERR	LERR	DMA
07-00	LINT7	LINT6	LINT5	LINT4	LINT3	LINT2	LINT1	LINT0

VINT_EN Description

Name	Type	Reset By	Reset State	Function
SW_INT7	R/W	all	0	VME Software 7 Interrupt Mask 0=VME Software 7 Interrupt masked, 1=VME Software 7 Interrupt enabled A zero-to-one transition will cause a VME level 7 interrupt to be generated. Subsequent zeroing of this bit will cause the interrupt to be masked, but will not clear the VME Software 7 Interrupt Status bit.
SW_INT6	R/W	all	0	VME Software 6 Interrupt Mask 0=VME Software 6 Interrupt masked, 1=VME Software 6 Interrupt enabled A zero-to-one transition will cause a VME level 6 interrupt to be generated. Subsequent zeroing of this bit will cause the interrupt to be masked, but will not clear the VME Software 6 Interrupt Status bit.
SW_INT5	R/W	all	0	VME Software 5 Interrupt Mask 0=VME Software 5 Interrupt masked, 1=VME Software 5 Interrupt enabled A zero-to-one transition will cause a VME level 5 interrupt to be generated. Subsequent zeroing of this bit will cause the interrupt to be masked, but will not clear the VME Software 5 Interrupt Status bit.
SW_INT4	R/W	all	0	VME Software 4 Interrupt Mask 0=VME Software 4 Interrupt masked, 1=VME Software 4 Interrupt enabled A zero-to-one transition will cause a VME level 4 interrupt to be generated. Subsequent zeroing of this bit will cause the interrupt to be masked, but will not clear the VME Software 4 Interrupt Status bit.
SW_INT3	R/W	all	0	VME Software 3 Interrupt Mask 0=VME Software 3 Interrupt masked, 1=VME Software 3 Interrupt enabled A zero-to-one transition will cause a VME level 3 interrupt to be generated. Subsequent zeroing of this bit will cause the interrupt to be masked, but will not clear the VME Software 3 Interrupt Status bit.

VINT_EN Description

SW_INT2	R/W	all	0	VME Software 2 Interrupt Mask 0=VME Software 2 Interrupt masked, 1=VME Software 2 Interrupt enabled A zero-to-one transition will cause a VME level 2 interrupt to be generated. Subsequent zeroing of this bit will cause the interrupt to be masked, but will not clear the VME Software 2 Interrupt Status bit.
SW_INT1	R/W	all	0	VME Software 1 Interrupt Mask 0=VME Software 1 Interrupt masked, 1=VME Software 1 Interrupt enabled A zero-to-one transition will cause a VME level 1 interrupt to be generated. Subsequent zeroing of this bit will cause the interrupt to be masked, but will not clear the VME Software 1 Interrupt Status bit.
MBOX3	R/W	all	0	Mailbox 3 Mask 0=MBOX3 Interrupt masked, 1=MBOX3 Interrupt enabled
MBOX2	R/W	all	0	Mailbox 2 Mask 0=MBOX2 Interrupt masked, 1=MBOX2 Interrupt enabled
MBOX1	R/W	all	0	Mailbox 1 Mask 0=MBOX1 Interrupt masked, 1=MBOX1 Interrupt enabled
MBOX0	R/W	all	0	Mailbox 0 Mask 0=MBOX0 Interrupt masked, 1=MBOX0 Interrupt enabled
SW_INT	R/W	all	Power-up Option	“VME Software Interrupt” Mask 0 = VME Software Interrupt masked 1 = VME Software Interrupt enabled A zero-to-one transition causes the VME software interrupt to be asserted. Subsequent zeroing of this bit causes the interrupt to be masked and the VMEbus interrupt negated, but does not clear the VME software interrupt status bit.
VERR	R/W	all	0	VERR Interrupt Mask 0 =PCI VERR Interrupt masked 1=PCI VERR Interrupt enabled
LERR	R/W	all	0	LERR Interrupt Mask 0 =PCI LERR Interrupt masked 1 =PCI LERR Interrupt enabled
DMA	R/W	all	0	DMA Interrupt Mask 0=PCI DMA Interrupt masked 1=PCI DMA Interrupt enabled
LINT7-LINT0	R/W	all	0	PCI Interrupt Mask 0=LINT _x Interrupt masked 1 =LINT _x Interrupt enabled

This register enables the various sources of VMEbus interrupts. SW_INT can be enabled with the VME64AUTO power-up option.

Table A.62 : VMEbus Interrupt Status Register (VINT_STAT)

Register Name: VINT_STAT								Offset:314
Bits	Function							
31-24	SW_INT7	SW_INT6	SW_INT5	SW_INT4	SW_INT3	SW_INT2	SW_INT1	Reserved
23-16	Reserved				MBOX3	MBOX2	MBOX1	MBOX0
15-08	Reserved			SW_INT	Reserved	VERR	LERR	DMA
07-00	LINT7	LINT6	LINT5	LINT4	LINT3	LINT2	LINT1	LINT0

VINT_STAT Description

Name	Type	Reset By	Reset State	Function
SW_INT7	R/Write 1 to clear	all	0	VME Software 7 Interrupt Status/Clear 0=no VME Software 7 Interrupt, 1=VME Software 7 Interrupt active
SW_INT6	R/Write 1 to clear	all	0	VME Software 6 Interrupt Status/Clear 0=no VME Software 6 Interrupt, 1=VME Software 6 Interrupt active
SW_INT5	R/Write 1 to clear	all	0	VME Software 5 Interrupt Status/Clear 0=no VME Software 5 Interrupt, 1=VME Software 5 Interrupt active
SW_INT4	R/Write 1 to clear	all	0	VME Software 4 Interrupt Status/Clear 0=no VME Software 4 Interrupt, 1=VME Software 4 Interrupt active
SW_INT3	R/Write 1 to clear	all	0	VME Software 3 Interrupt Status/Clear 0=no VME Software 3 Interrupt, 1=VME Software 3 Interrupt active
SW_INT2	R/Write 1 to clear	all	0	VME Software 2 Interrupt Status/Clear 0=no VME Software 2 Interrupt, 1=VME Software 2 Interrupt active
SW_INT1	R/Write 1 to clear	all	0	VME Software 1 Interrupt Status/Clear 0=no VME Software 1 Interrupt, 1=VME Software 1 Interrupt active
MBOX3	R/Write 1 to clear	all	0	Mailbox 3 Status/Clear 0=no Mailbox 3 Interrupt, 1=Mailbox 3 Interrupt active
MBOX2	R/Write 1 to clear	all	0	Mailbox 2 Status/Clear 0=no Mailbox 2 Interrupt, 1=Mailbox 2 Interrupt active
MBOX1	R/Write 1 to clear	all	0	Mailbox 1 Status/Clear 0=no Mailbox 1 Interrupt, 1=Mailbox 1 Interrupt active
MBOX0	R/Write 1 to clear	all	0	Mailbox 0 Status/Clear 0=no Mailbox 0 Interrupt, 1=Mailbox 0 Interrupt active

VINT_STAT Description

SW_INT	R/Write 1 to Clear	all	Power-up Option	VME Software Interrupt Status/Clear 0=VME Software Interrupt inactive, 1=VME Software Interrupt active
VERR	R/Write 1 to Clear	all	0	VERR Interrupt Status/Clear 0=VME VERR Interrupt masked, 1=VME VERR Interrupt enabled
LERR	R/Write 1 to Clear	all	0	LERR Interrupt Status/Clear 0=VME LERR Interrupt masked, 1=VME LERR Interrupt enabled
DMA	R/Write 1 to Clear	all	0	DMA Interrupt Status/Clear 0=VME DMA Interrupt masked, 1=VME DMA Interrupt enabled
LINT7-LINT0	R/Write 1 to Clear	all	0	LINT _x Interrupt Status/Clear 0=LINT _x Interrupt masked, 1=LINT _x Interrupt enabled

SW_INT can be set with the VME64AUTO power-up option.

Table A.63 : VME Interrupt Map 0 Register (VINT_MAP0)

Register Name: VINT_MAP0				Offset: 318
Bits	Function			
31-24	Reserved	LINT7	Reserved	LINT6
23-16	Reserved	LINT5	Reserved	LINT4
15-08	Reserved	LINT3	Reserved	LINT2
07-00	Reserved	LINT1	Reserved	LINT0

VINT_MAP0 Description

Name	Type	Reset By	Reset State	Function
LINT7-LINT0	R/W	all	0	VMEbus destination of PCI Bus interrupt source

This register maps various interrupt sources to one of the seven VMEbus interrupt pins. A value of 001 maps the corresponding interrupt source to VIRQ*[1], a value of 002 maps to VIRQ*[2], etc. A value of 000 effectively masks the interrupt since there is no corresponding VIRQ*[0].

Table A.64 : VME Interrupt Map 1 Register (VINT_MAP1)

Register Name: VINT_MAP1				Offset: 31C
Bits	Function			
31-24	Reserved			
23-16	Reserved			SW_INT
15-08	Reserved			VERR
07-00	Reserved	LERR	Reserved	DMA

VINT_MAP1 Description

Name	Type	Reset By	Reset State	Function
SW_INT	R/W	all	Power-up Option	VMEbus Software interrupt destination
VERR	R/W	all	0	VMEbus Error interrupt destination
LERR	R/W	all	0	PCI Bus Error interrupt destination
DMA	R/W	all	0	DMA interrupt destination

This register maps various interrupt sources to one of the seven VMEbus interrupt pins. A value of 001 maps the corresponding interrupt source to VIRQ*[1], a value of 002 maps to VIRQ*[2], etc. A value of 000 effectively masks the interrupt since there is no corresponding VIRQ*[0].

SW_INT is set to 010 with the VME64AUTO power-up option.

Table A.65 : Interrupt STATUS/ID Out Register (STATID)

Register Name: STATID		Offset: 320
Bits	Function	
31-24	STATID [7:0]	
23-16	Reserved	
15-08	Reserved	
07-00	Reserved	

STATID Description

Name	Type	Reset By	Reset State	Function
STATID [7:1]	R/W	all	1111111	Bits [7:1] of the STATUS/ID byte are returned when the Universe II responds to a VMEbus IACK cycle.
STATID [0]	R	all	See below	0 = the Universe II is generating a SW_IACK at the same level as the interrupt acknowledge cycle. 1 = the Universe II is not generating a SW_IACK at the same level as the interrupt acknowledge cycle.

When the Universe II responds to an interrupt acknowledge cycle on VMEbus it returns an 8-bit STATUS/ID. STATID [7:1] can be written by software to uniquely identify the VMEbus module within the system. STATID [0] is a value of 0 if the Universe II is generating a software interrupt (SW_IACK) at the same level as the interrupt acknowledge cycle, otherwise it is a value of 1.

The reset state is designed to support the VME64 Auto ID STATUS/ID value.

Table A.66 : VIRQ1 STATUS/ID Register (V1_STATID)

Register Name: V1_STATID			Offset: 324
Bits	Function		
31-24	Reserved		
23-16	Reserved		
15-08	Reserved		ERR
07-00	STATID [7:0]		

V1_STATID Description

Name	Type	Reset By	Reset State	Function
ERR	R	all	0	Error Status Bit 0=STATUS/ID was acquired without bus error 1=bus error occurred during acquisition of the STATUS/ID
STATID [7:0]	R	all	0	STATUS/ID acquired during IACK cycle for level 1 VMEbus interrupt

The V_x_STATID registers are read-only registers that hold the 8-bit VMEbus STATUS/ID that is acquired when the Universe II performs a IACK cycle for a given interrupt level. The Universe II is enabled as the interrupt handler for a given interrupt level via the VIRQ_x bits of the LINT_EN register. Once a vector for a given level is acquired, the Universe II does not perform a subsequent interrupt acknowledge cycle at that level until the corresponding VIRQ_x bit in the LINT_STAT register is cleared.

The acquisition of a level *x* STATUS/ID by the Universe II updates the STATUS/ID field of the corresponding V_x_STATID register and generation of a PCI interrupt. A VMEbus error during the acquisition of the STATUS/ID vector sets the ERR bit, which means the STATUS/ID field may not contain a valid vector.

Table A.67 : VIRQ2 STATUS/ID Register (V2_STATID)

Register Name: V2_STATID		Offset: 328
Bits	Function	
31-24	Reserved	
23-16	Reserved	
15-08	Reserved	ERR
07-00	STATID [7:0]	

V2_STATID Description

Name	Type	Reset By	Reset State	Function
ERR	R	all	0	Error Status Bit 0=STATUS/ID was acquired without bus error 1=bus error occurred during acquisition of the STATUS/ID
STATID [7:0]	R	all	0	STATUS/ID acquired during IACK cycle for level 1 VMEbus interrupt

The Vx_STATID registers are read-only registers that hold the 8-bit VMEbus STATUS/ID that is acquired when the Universe II performs a IACK cycle for a given interrupt level. The Universe II is enabled as the interrupt handler for a given interrupt level via the VIRQx bits of the LINT_EN register. Once a vector for a given level is acquired, the Universe II does not perform a subsequent interrupt acknowledge cycle at that level until the corresponding VIRQx bit in the LINT_STAT register is cleared.

The acquisition of a level x STATUS/ID by the Universe II updates the STATUS/ID field of the corresponding Vx_STATID register and generation of a PCI interrupt. A VMEbus error during the acquisition of the STATUS/ID vector sets the ERR bit, which means the STATUS/ID field may not contain a valid vector.

Table A.68 : VIRQ3 STATUS/ID Register (V3_STATID)

Register Name: V3_STATID			Offset: 32C
Bits	Function		
31-24	Reserved		
23-16	Reserved		
15-08	Reserved		ERR
07-00	STATID [7:0]		

V3_STATID Description

Name	Type	Reset By	Reset State	Function
ERR	R	all	0	Error Status Bit 0=STATUS/ID was acquired without bus error 1=bus error occurred during acquisition of the STATUS/ID
STATID [7:0]	R	all	0	STATUS/ID acquired during IACK cycle for level 3VMEbus interrupt

The Vx_STATID registers are read-only registers that hold the 8-bit VMEbus STATUS/ID that is acquired when the Universe II performs a IACK cycle for a given interrupt level. The Universe II is enabled as the interrupt handler for a given interrupt level via the VIRQx bits of the LINT_EN register. Once a vector for a given level is acquired, the Universe II does not perform a subsequent interrupt acknowledge cycle at that level until the corresponding VIRQx bit in the LINT_STAT register is cleared.

The acquisition of a level x STATUS/ID by the Universe II updates the STATUS/ID field of the corresponding Vx_STATID register and generation of a PCI interrupt. A VMEbus error during the acquisition of the STATUS/ID vector sets the ERR bit, which means the STATUS/ID field may not contain a valid vector.

Table A.69 : VIRQ4 STATUS/ID Register (V4_STATID)

Register Name: V4_STATID		Offset: 330
Bits	Function	
31-24	Reserved	
23-16	Reserved	
15-08	Reserved	ERR
07-00	STATID [7:0]	

V4_STATID Description

Name	Type	Reset By	Reset State	Function
ERR	R	all	0	Error Status Bit 0=STATUS/ID was acquired without bus error 1=bus error occurred during acquisition of the STATUS/ID
STATID [7:0]	R	all	0	STATUS/ID acquired during IACK cycle for level 4 VMEbus interrupt

The Vx_STATID registers are read-only registers that hold the 8-bit VMEbus STATUS/ID that is acquired when the Universe II performs a IACK cycle for a given interrupt level. The Universe II is enabled as the interrupt handler for a given interrupt level via the VIRQx bits of the LINT_EN register. Once a vector for a given level is acquired, the Universe II does not perform a subsequent interrupt acknowledge cycle at that level until the corresponding VIRQx bit in the LINT_STAT register is cleared.

The acquisition of a level x STATUS/ID by the Universe II updates the STATUS/ID field of the corresponding Vx_STATID register and generation of a PCI interrupt. A VMEbus error during the acquisition of the STATUS/ID vector sets the ERR bit, which means the STATUS/ID field may not contain a valid vector.

Table A.70 : VIRQ5 STATUS/ID Register (V5_STATID)

Register Name: V5_STATID			Offset: 334
Bits	Function		
31-24	Reserved		
23-16	Reserved		
15-08	Reserved		ERR
07-00	STATID [7:0]		

V5_STATID Description

Name	Type	Reset By	Reset State	Function
ERR	R	all	0	Error Status Bit 0=STATUS/ID was acquired without bus error 1=bus error occurred during acquisition of the STATUS/ID
STATID [7:0]	R	all	0	STATUS/ID acquired during IACK cycle for level 5 VMEbus interrupt

The Vx_STATID registers are read-only registers that hold the 8-bit VMEbus STATUS/ID that is acquired when the Universe II performs a IACK cycle for a given interrupt level. The Universe II is enabled as the interrupt handler for a given interrupt level via the VIRQx bits of the LINT_EN register. Once a vector for a given level is acquired, the Universe II does not perform a subsequent interrupt acknowledge cycle at that level until the corresponding VIRQx bit in the LINT_STAT register is cleared.

The acquisition of a level x STATUS/ID by the Universe II updates the STATUS/ID field of the corresponding Vx_STATID register and generation of a PCI interrupt. A VMEbus error during the acquisition of the STATUS/ID vector sets the ERR bit, which means the STATUS/ID field may not contain a valid vector.

Table A.71 : VIRQ6 STATUS/ID Register (V6_STATID)

Register Name: V6_STATID		Offset: 338
Bits	Function	
31-24	Reserved	
23-16	Reserved	
15-08	Reserved	ERR
07-00	STATID [7:0]	

V6_STATID Description

Name	Type	Reset By	Reset State	Function
ERR	R	all	0	Error Status Bit 0=STATUS/ID was acquired without bus error 1=bus error occurred during acquisition of the STATUS/ID
STATID [7:0]	R	all	0	STATUS/ID acquired during IACK cycle for level 6 VMEbus interrupt

The Vx_STATID registers are read-only registers that hold the 8-bit VMEbus STATUS/ID that is acquired when the Universe II performs a IACK cycle for a given interrupt level. The Universe II is enabled as the interrupt handler for a given interrupt level via the VIRQx bits of the LINT_EN register. Once a vector for a given level is acquired, the Universe II does not perform a subsequent interrupt acknowledge cycle at that level until the corresponding VIRQx bit in the LINT_STAT register is cleared.

The acquisition of a level x STATUS/ID by the Universe II updates the STATUS/ID field of the corresponding Vx_STATID register and generation of a PCI interrupt. A VMEbus error during the acquisition of the STATUS/ID vector sets the ERR bit, which means the STATUS/ID field may not contain a valid vector.

Table A.72 : VIRQ7 STATUS/ID Register (V7_STATID)

Register Name: V7_STATID			Offset: 33C
Bits	Function		
31-24	Reserved		
23-16	Reserved		
15-08	Reserved		ERR
07-00	STATID [7:0]		

V7_STATID Description

Name	Type	Reset By	Reset State	Function
ERR	R	all	0	Error Status Bit 0=STATUS/ID was acquired without bus error 1=bus error occurred during acquisition of the STATUS/ID
STATID [7:0]	R	all	0	STATUS/ID acquired during IACK cycle for level 7 VMEbus interrupt

The Vx_STATID registers are read-only registers that hold the 8-bit VMEbus STATUS/ID that is acquired when the Universe II performs a IACK cycle for a given interrupt level. The Universe II is enabled as the interrupt handler for a given interrupt level via the VIRQx bits of the LINT_EN register. Once a vector for a given level is acquired, the Universe II does not perform a subsequent interrupt acknowledge cycle at that level until the corresponding VIRQx bit in the LINT_STAT register is cleared.

The acquisition of a level x STATUS/ID by the Universe II updates the STATUS/ID field of the corresponding Vx_STATID register and generation of a PCI interrupt. A VMEbus error during the acquisition of the STATUS/ID vector sets the ERR bit, which means the STATUS/ID field may not contain a valid vector.

Table A.73 : PCI Interrupt Map 2 Register (LINT_MAP2)

Register Name: LINT_MAP2				Offset: 340
Bits	Function			
31-24	Reserved	LM3	Reserved	LM2
23-16	Reserved	LM1	Reserved	LM0
15-08	Reserved	MBOX3	Reserved	MBOX2
07-00	Reserved	MBOX1	Reserved	MBOX0

LINT_MAP2 Description

Name	Type	Reset By	Reset State	Function
LM3 [2:0]	R/W	all	0	Location Monitor 3 Interrupt destination
LM2 [2:0]	R/W	all	0	Location Monitor 2 Interrupt destination
LM1 [2:0]	R/W	all	0	Location Monitor 1 Interrupt destination
LM0 [2:0]	R/W	all	0	Location Monitor 0 Interrupt destination
MBOX3 [2:0]	R/W	all	0	Mailbox 3 Interrupt destination
MBOX2 [2:0]	R/W	all	0	Mailbox 2 Interrupt destination
MBOX1 [2:0]	R/W	all	0	Mailbox 1 Interrupt destination
MBOX0 [2:0]	R/W	all	0	Mailbox 0 Interrupt destination

This register maps various interrupt sources to one of the eight PCI interrupt pins. A value of 000 maps the corresponding interrupt source to LINT# [0], a value of 001 maps to LINT# [1], etc.

Table A.74 : VME Interrupt Map 2 Register (VINT_MAP2)

Register Name: VINT_MAP2					Offset: 344
Bits		Function			
31-24		Reserved			
23-16		Reserved			
15-08	Reserved	MBOX3		Reserved	MBOX2
07-00	Reserved	MBOX1		Reserved	MBOX0

VINT_MAP2 Description

Name	Type	Reset By	Reset State	Function
MBOX3 [2:0]	R/W	all	0	Mailbox 3 Interrupt destination
MBOX2 [2:0]	R/W	all	0	Mailbox 2 Interrupt destination
MBOX1 [2:0]	R/W	all	0	Mailbox 1 Interrupt destination
MBOX0 [2:0]	R/W	all	0	Mailbox 0 Interrupt destination

This register maps various interrupt sources to one of the seven VMEbus interrupt pins. A value of 001 maps the corresponding interrupt source to VIRQ*[1], a value of 002 maps to VIRQ*[2], etc. A value of 000 effectively masks the interrupt since there is no corresponding VIRQ*[0].

Table A.75 : Mailbox 0 Register (MBOX0)

Register Name: MBOX0		Offset:348
Bits	Function	
31-24	MBOX0	
23-16	MBOX0	
15-08	MBOX0	
07-00	MBOX0	

DVA Description

Name	Type	Reset By	Reset State	Function
MBOX0 [31:0]	R/W	all	0	Mailbox

General purpose mailbox register. Writes to this register will cause interrupt generation on PCI or VMEbus if enabled in LINT_EN register.

Table A.76 : Mailbox 1 Register (MBOX1)

Register Name: MBOX1		Offset:34C
Bits	Function	
31-24	MBOX1	
23-16	MBOX1	
15-08	MBOX1	
07-00	MBOX1	

DVA Description

Name	Type	Reset By	Reset State	Function
MBOX1 [31:0]	R/W	all	0	Mailbox

General purpose mailbox register. Writes to this register will cause interrupt generation on PCI or VMEbus if enabled in LINT_EN register.

Table A.77 : Mailbox 2 Register (MBOX2)

Register Name: MBOX2		Offset:350
Bits	Function	
31-24	MBOX2	
23-16	MBOX2	
15-08	MBOX2	
07-00	MBOX2	

DVA Description

Name	Type	Reset By	Reset State	Function
MBOX2 [31:0]	R/W	all	0	Mailbox

General purpose mailbox register. Writes to this register will cause interrupt generation on PCI or VMEbus if enabled in LINT_EN register.

Table A.78 : Mailbox 3 Register (MBOX3)

Register Name: MBOX3		Offset:354
Bits	Function	
31-24	MBOX3	
23-16	MBOX3	
15-08	MBOX3	
07-00	MBOX3	

DVA Description

Name	Type	Reset By	Reset State	Function
MBOX3 [31:0]	R/W	all	0	Mailbox

General purpose mailbox register. Writes to this register will cause interrupt generation on PCI or VMEbus if enabled in LINT_EN register.

Table A.79 : Semaphore 0 Register (SEMA0)

Register Name: SEMA0		Offset: 358
Bits	Function	
31-24	SEM3	TAG3
23-16	SEM2	TAG2
15-08	SEM1	TAG1
07-00	SEM0	TAG0

SEMA0 Description

Name	Type	Reset By	Reset State	Function
SEM3	R/W	all	0	Semaphore 3
TAG3 [6:0]	R/W	all	0	Tag 3
SEM2	R/W	all	0	Semaphore 2
TAG2 [6:0]	R/W	all	0	Tag2
SEM1	R/W	all	0	Semaphore 1
TAG1 [6:0]	R/W	all	0	Tag 1
SEM0	R/W	all	0	Semaphore 0
TAG0 [6:0]	R/W	all	0	Tag 0

If a semaphore bit is a value of 0, the associated tag field can be written to. If a semaphore bit is a value of 1, the associated tag field cannot be written to.

This register can only be accessed via byte-wide access.

See “Semaphores” on page 2-109.

Table A.80 : Semaphore 1 Register (SEMA1)

Register Name: SEMA1		Offset: 35C
Bits	Function	
31-24	SEM7	TAG6
23-16	SEM6	TAG6
15-08	SEM5	TAG5
07-00	SEM4	TAG4

SEMA1 Description

Name	Type	Reset By	Reset State	Function
SEM3	R/W	all	0	Semaphore 7
TAG3 [6:0]	R/W	all	0	Tag 7
SEM2	R/W	all	0	Semaphore 6
TAG2 [6:0]	R/W	all	0	Tag 6
SEM1	R/W	all	0	Semaphore 5
TAG1 [6:0]	R/W	all	0	Tag 5
SEM0	R/W	all	0	Semaphore 4
TAG0 [6:0]	R/W	all	0	Tag 4

If a semaphore bit is a value of 0, the associated tag field can be written to. If a semaphore bit is a value of 1, the associated tag field cannot be written to.

This register can only be accessed via byte-wide access.

See “Semaphores” on page 2-109.

Table A.81 : Master Control Register (MAST_CTL)

Register Name: MAST_CTL							Offset: 400
Bits	Function						
31-24	MAXRTRY			PWON			
23-16	VRL	VRM	VREL	VOWN	VOWN_ACK	Reserved	
15-08	Reserved	PABS		Reserved			
07-00	BUS_NO						

MAST_CTL Description

Name	Type	Reset By	Reset State	Function
MAXRTRY [3:0]	R/W	all	1000	Maximum Number of Retries 0000=Retry Forever, Multiples of 64 (0001 through 1111). Maximum Number of retries before the PCI master interface signals error condition
PWON [3:0]	R/W	all	0000	Posted Write Transfer Count 0000=128 bytes, 0001=256 bytes, 0010=512 bytes, 0011=1024 bytes, 0100=2048 bytes, 0101=4096 bytes, 0110 - 1110 = Reserved, 1111=Early release of BBSY*. Transfer count at which the PCI Slave Channel Posted Writes FIFO gives up the VME Master Interface.
VRL [1:0]	R/W	all	11	VMEbus Request Level 00=Level 0, 01=Level 1, 10=Level 2, 11=Level 3
VRM	R/W	all	0	VMEbus Request Mode 0=Demand, 1=Fair
VREL	R/W	all	0	VMEbus Release Mode 0=Release When Done (RWD), 1=Release on Request (ROR)
VOWN	W	all	0	VME Ownership Bit 0=Release VMEbus, 1=Acquire and Hold VMEbus
VOWN_ACK	R	all	0	VME Ownership Bit Acknowledge 0=VMEbus not owned, 1=VMEbus acquired and held due to assertion of VOWN
PABS [1:0]	R/W	all	00	PCI Aligned Burst Size 00=32-byte, 01=64-byte, 10=128-byte, 11=Reserved Controls the PCI address boundary at which the Universe II breaks up a PCI transaction in the VME Slave channel (see “VME Slave Images” on page 2-50) and the DMA Channel (see “FIFO Operation and Bus Ownership” on page 2-92). This field also determines when the PCI Master Module as part of the VME Slave Channel will request the PCI bus (i.e., when 32, 64, or 128 bytes are available). It does not have this effect on the DMA Channel, which has a fixed watermark of 128 bytes (see “FIFO Operation and Bus Ownership” on page 2-92).
BUS_NO [7:0]	R/W	all	0000 0000	PCI Bus Number

Writing a 1 to the VOWN bit in the MAST_CTL register has the effect of asserting BBSY* until a 0 is written to the VOWN bit. It does not affect the transactions in the PCI Target Channel. The Universe II will not do an early release of BBSY* if the VMEbus was owned during a transaction by means of VOWN, regardless of the value of PWON.

It is important to wait until VOWN_ACK is a value of 0 before writing a value of 1 to the VOWN bit.

In the event that BERR* is asserted on the VMEbus once the Universe II owns the VMEbus, the user must release ownership by programming the VOWN bit to a value of 0, if the VMEbus was gained by setting the VOWN bit. VMEbus masters must not write a value of 1 to the VOWN bit since this will lock up the VMEbus.

Once the value programmed in the PWON field is reached during dequeuing of posted writes, the Universe II will do an early release of BBSY*. If the PWON field is programmed to a value of 1111, the Universe II will do an early release of BBSY* at the completion of each transaction. Note that the VOWN setting described above overrides the PWON setting.

BUS_NO is used by the VMEbus Slave Channel when mapping VME transactions into PCI Configuration space. If the bus number of the VMEbus address (bits [23:16]) is equal to the BUS_NO field, then the Universe II generates a Type 0 configuration cycle, otherwise Type 1 is generated.

The PABS[1:0] field also determines when the PCI Master Module as part of the VME Slave Channel will request the PCI bus (i.e., when 32, 64, or 128 bytes are available). It does not have this effect on the DMA Channel, which has a fixed watermark of 128 bytes (see “FIFO Operation and Bus Ownership” on page 2-92).

Table A.82 : Miscellaneous Control Register (MISC_CTL)

Register Name: MISC_CTL								Offset: 404
Bits	Function							
31-24	VBTO				Reserved	VARB	VARBTO	
23-16	SW_LRST	SW_SRST	Reserved	BI	ENGBI	RESCIND	SYSCON	V64AUTO
15-08	Reserved							
07-00	Reserved							

MISC_CTL Description

Name	Type	Reset By	Reset State	Function
VBTO	R/W	all	0011	VME Bus Time-out 0000=Disable, 0001=16 μ sec, 0010=32 μ sec, 0011=64 μ sec, 0100=128 μ sec, 0101=256 μ sec, 0110=512 μ sec, 0111=1024 μ sec, others=RESERVED
VARB	R/W	all	0	VMEbus Arbitration Mode 0=Round Robin, 1=Priority
VARBTO	R/W	all	01	VMEbus Arbitration Time-out 00=Disable Timer, 01=16 μ s (minimum 8 μ s), 10=256 μ s, others=Reserved
SW_LRST	W	all	0	Software PCI Reset 0=No effect, 1=Initiate LRST# A read always returns 0.
SW_SYSRST	W	all	0	Software VMEbus SYSRESET 0=No effect, 1=Initiate SYSRST* A read always returns 0.
BI	R/W	all	Power-up Option	BI-Mode 0=Universe II is not in BI-Mode, 1=Universe II is in BI-Mode Write to this bit to change the Universe II BI-Mode status. This bit is also affected by the global BI-Mode initiator VRIRQ1*, if this feature is enabled.
ENGBI	R/W	all	0	Enable Global BI-Mode Initiator 0=Assertion of VIRQ1 ignored, 1=Assertion of VIRQ1 puts device in BI-Mode
RESCIND	R/W	all	1	RESCIND is unused in the Universe II.
SYSCON	R/W	all	Power-up Option	SYSCON 0=Universe II is not VMEbus System Controller, 1=Universe II is VMEbus System Controller
V64AUTO	R/W	all	Power-up Option	VME64 Auto ID Write: 0=No effect, 1=Initiate sequence This bit initiates Universe II VME64 Auto ID Slave participation.

VMEbus masters must not write to SW_SYSRST, and PCI masters must not write to SW_LRST.

The bits VBTO, VARB and VARBTO support SYSCON functionality.

Universe II participation in the VME64 Auto ID mechanism is controlled by the VME64AUTO bit. When this bit is detected high, the Universe II uses the SW_IACK mechanism to generate VXIRQ2 on the VMEbus, then releases VXSYSFAIL. Access to the CR/CSR image is enabled when the level 2 interrupt acknowledge cycle completes. This sequence can be initiated with a power-up option or by software writing a 1 to this bit.

Table A.83 : Miscellaneous Status Register (MISC_STAT)

Register Name: MISC_STAT							Offset: 408	
Bits	Function							
31-24	Reserved	LCLSIZE	Reserved		DY4AUTO	Reserved		
23-16	Reserved		MYBBSY	Reserved	DY4_DONE	TXFE	RXFE	Reserved
15-08	DY4AUTOID							
07-00	Reserved							

MISC_STAT Description

Name	Type	Reset By	Reset State	Function
LCLSIZE	R	all	Power-up Option	PCI Bus Size At the trailing edge of RST#, the Universe II samples REQ64# to determine the PCI Bus size. This bit reflects the result. 0=32-bit, 1=64-bit
DY4AUTO	R	all	Power-up Option	DY4 Auto ID Enable 0=Disable, 1=Enable
MYBBSY	R	all	1	Universe II BBSY 0=Asserted, 1=Negated
DY4DONE	R	all	0	DY4 Auto ID Done 0=Not done, 1=Done
TXFE	R	all	1	PCI Target Channel Posted Writes FIFO 0=no data in the FIFO, 1=data in the FIFO
RXFE	R	all	1	VME Slave Channel Posted Writes FIFO 0=no data in the FIFO, 1=data in the FIFO
DY4AUTOID	R	none	0	DY4 Auto ID

Table A.84 : User AM Codes Register (USER_AM)

Register Name: USER_AM				Offset: 40C
Bits	Function			
31-24	0	1	USER1AM	Reserved
23-16	0	1	USER2AM	Reserved
15-08	Reserved			
07-00	Reserved			

USER_AM Description

Name	Type	Reset By	Reset State	Function
USER1AM [3:0]	R/W	all	0000	User AM Code 1
USER2AM [3:0]	R/W	all	0000	User AM Code 2

The reset state is one of the VME64 user-defined AM codes.

The USER_AM register can only be used to generate and accept AM codes 0x10 through 0x1F. These AM codes are designated as USERAM codes in the VMEbus specification.

Table A.85 : VMEbus Slave Image 0 Control (VSI0_CTL)

Register Name: VSI0_CTL						Offset: F00
Bits	Function					
31-24	EN	PWEN	PREN	Reserved		
23-16	PGM		SUPER		Reserved VAS	
15-08	Reserved					
07-00	LD64EN	LLRMW	Reserved		LAS	

VSI0_CTL Description

Name	Type	Reset By	Reset State	Function
EN	R/W	PWR VME	0	Image Enable 0=Disable, 1=Enable
PWEN	R/W	PWR VME	0	Posted Write Enable 0=Disable, 1=Enable
PREN	R/W	PWR VME	0	Prefetch Read Enable 0=Disable, 1=Enable
PGM	R/W	PWR VME	11	Program/Data AM Code 00=Reserved, 01=Data, 10=Program, 11=Both
SUPER	R/W	PWR VME	11	Supervisor/User AM Code 00=Reserved, 01=Non-Privileged, 10=Supervisor, 11=Both
VAS	R/W	PWR VME	0	VMEbus Address Space 000=A16, 001=A24, 010=A32, 011= Reserved, 100=Reserved, 101=Reserved, 110=User1, 111=User2
LD64EN	R/W	PWR VME	0	Enable 64-bit PCI Bus Transactions 0=Disable, 1=Enable
LLRMW	R/W	PWR VME	0	Enable PCI Bus Lock of VMEbus RMW 0=Disable, 1=Enable
LAS	R/W	PWR VME	0	PCI Bus Address Space 00=PCI Bus Memory Space, 01=PCI Bus I/O Space, 10=PCI Bus Configuration Space, 11=Reserved

This register provides the general, VMEbus and PCI controls for this slave image. Note that only transactions destined for PCI Memory space are decoupled (the posted write RXFIFO generates on Memory space transactions on the PCI Bus). This image has 4 Kbyte resolution.

In order for a VMEbus slave image to respond to an incoming cycle, the BM bit in the PCI_CSR register must be enabled.

The state of PWEN and PREN are ignored if LAS is not programmed memory space.

Table A.86 : VMEbus Slave Image 0 Base Address Register (VSI0_BS)

Register Name: VSI0_BS		Offset: F04
Bits	Function	
31-24	BS	
23-16	BS	
15-08	BS	Reserved
07-00	Reserved	

VSI0_BS Description

Name	Type	Reset By	Reset State	Function
BS[31:12]	R/W	PWR VME	0	Base Address

The base address specifies the lowest address in the address range that will be decoded.

This image has 4 Kbyte resolution.

Table A.87 : VMEbus Slave Image 0 Bound Address Register (VSI0_BD)

Register Name: VSI0_BD		Offset: F08
Bits	Function	
31-24	BD	
23-16	BD	
15-08	BD	Reserved
07-00	Reserved	

VSI0_BD Description

Name	Type	Reset By	Reset State	Function
BD[31:12]	R/W	PWR VME	0	Bound Address

The addresses decoded in a slave image are those which are greater than or equal to the base address and less than the bound register.

This image has 4 Kbyte resolution.

Table A.88 : VMEbus Slave Image 0 Translation Offset (VSI0_TO)

Register Name: VSI0_TO		Offset: F0C
Bits	Function	
31-24	TO	
23-16	TO	
15-08	TO	Reserved
07-00	Reserved	

VSI0_TO Description

Name	Type	Reset By	Reset State	Function
TO[31:12]	R/W	PWR VME	0	Translation Offset

This image has 4 Kbyte resolution.

Table A.89 : VMEbus Slave Image 1 Control (VSI1_CTL)

Register Name: VSI1_CTL						Offset: F14
Bits	Function					
31-24	EN	PWEN	PREN	Reserved		
23-16	PGM		SUPER		Reserved VAS	
15-08	Reserved					
07-00	LD64EN	LLRMW	Reserved		LAS	

VSI1_CTL Description

Name	Type	Reset By	Reset State	Function
EN	R/W	PWR VME	0	Image Enable 0=Disable, 1=Enable
PWEN	R/W	PWR VME	0	Posted Write Enable 0=Disable, 1=Enable
PREN	R/W	PWR VME	0	Prefetch Read Enable 0=Disable, 1=Enable
PGM	R/W	PWR VME	11	Program/Data AM Code 00=Reserved, 01=Data, 10=Program, 11=Both
SUPER	R/W	PWR VME	11	Supervisor/User AM Code 00=Reserved, 01=Non-Privileged, 10=Supervisor, 11=Both
VAS	R/W	PWR VME	0	VMEbus Address Space 000=Reserved, 001=A24, 010=A32, 011= Reserved, 100=Reserved, 101=Reserved, 110=User1, 111=User2
LD64EN	R/W	PWR VME	1	Enable 64-bit PCI Bus Transactions 0=Disable, 1=Enable
LLRMW	R/W	PWR VME	1	Enable PCI Bus Lock of VMEbus RMW 0=Disable, 1=Enable
LAS	R/W	PWR VME	0	PCI Bus Address Space 00=PCI Bus Memory Space, 01=PCI Bus I/O Space, 10=PCI Bus Configuration Space, 11=Reserved

This register provides the general, VMEbus and PCI controls for this slave image. Note that only transactions destined for PCI Memory space are decoupled (the posted write RXFIFO generates on Memory space transactions on the PCI Bus).

In order for a VMEbus slave image to respond to an incoming cycle, the BM bit in the PCI_CSR register must be enabled.

The state of PWEN and PREN are ignored if LAS is not programmed memory space.

Table A.90 : VMEbus Slave Image 1 Base Address Register (VSI1_BS)

Register Name: VSI1_BS		Offset: F18
Bits	Function	
31-24	BS	
23-16	BS	
15-08	Reserved	
07-00	Reserved	

VSI1_BS Description

Name	Type	Reset By	Reset State	Function
BS[31:16]	R/W	PWR VME	0	Base Address

The base address specifies the lowest address in the address range that will be decoded.

Table A.91 : VMEbus Slave Image 1 Bound Address Register (VSI1_BD)

Register Name: VSI1_BD		Offset: F1C
Bits	Function	
31-24	BD	
23-16	BD	
15-08	Reserved	
07-00	Reserved	

VSI1_BD Description

Name	Type	Reset By	Reset State	Function
BD[31:16]	R/W	PWR VME	0	Bound Address

The addresses decoded in a slave image are those which are greater than or equal to the base address and less than the bound register.

Table A.92 : VMEbus Slave Image 1 Translation Offset (VSI1_TO)

Register Name: VSI1_TO		Offset: F20
Bits	Function	
31-24	TO	
23-16	TO	
15-08	Reserved	
07-00	Reserved	

VSI1_TO Description

Name	Type	Reset By	Reset State	Function
TO[31:16]	R/W	PWR VME	0	Translation Offset

Table A.93 : VMEbus Slave Image 2 Control (VSI2_CTL)

Register Name: VSI2_CTL						Offset: F28	
Bits	Function						
31-24	EN	PWEN	PREN	Reserved			
23-16	PGM		SUPER		Reserved	VAS	
15-08	Reserved						
07-00	LD64EN	LLRMW	Reserved				LAS

VSI2_CTL Description

Name	Type	Reset By	Reset State	Function
EN	R/W	PWR VME	0	Image Enable 0=Disable, 1=Enable
PWEN	R/W	PWR VME	0	Posted Write Enable 0=Disable, 1=Enable
PREN	R/W	PWR VME	0	Prefetch Read Enable 0=Disable, 1=Enable
PGM	R/W	PWR VME	11	Program/Data AM Code 00=Reserved, 01=Data, 10=Program, 11=Both
SUPER	R/W	PWR VME	11	Supervisor/User AM Code 00=Reserved, 01=Non-Privileged, 10=Supervisor, 11=Both
VAS	R/W	PWR VME	0	VMEbus Address Space 000=Reserved, 001=A24, 010=A32, 011= Reserved, 100=Reserved, 101=Reserved, 110=User1, 111=User2
LD64EN	R/W	PWR VME	0	Enable 64-bit PCI Bus Transactions 0=Disable, 1=Enable
LLRMW	R/W	PWR VME	0	Enable PCI Bus Lock of VMEbus RMW 0=Disable, 1=Enable
LAS	R/W	PWR VME	0	PCI Bus Address Space 00=PCI Bus Memory Space, 01=PCI Bus I/O Space, 10=PCI Bus Configuration Space, 11=Reserved

This register provides the general, VMEbus and PCI controls for this slave image. Note that only transactions destined for PCI Memory space are decoupled (the posted write RXFIFO generates on Memory space transactions on the PCI Bus).

In order for a VMEbus slave image to respond to an incoming cycle, the BM bit in the PCI_CSR register must be enabled.

The state of PWEN and PREN are ignored if LAS is not programmed memory space.

Table A.94 : VMEbus Slave Image 2 Base Address Register (VSI2_BS)

Register Name: VSI2_BS		Offset: F2C
Bits	Function	
31-24	BS	
23-16	BS	
15-08	Reserved	
07-00	Reserved	

VSI2_BS Description

Name	Type	Reset By	Reset State	Function
BS[31:16]	R/W	PWR VME	0	Base Address

Table A.95 : VMEbus Slave Image 2 Bound Address Register (VSI2_BD)

Register Name: VSI2_BD		Offset: F30
Bits	Function	
31-24	BD	
23-16	BD	
15-08	Reserved	
07-00	Reserved	

VSI2_BD Description

Name	Type	Reset By	Reset State	Function
BD[31:16]	R/W	PWR VME	0	Bound Address

Table A.96 : VMEbus Slave Image 2 Translation Offset (VSI2_TO)

Register Name: VSI2_TO		Offset: F34
Bits	Function	
31-24	TO	
23-16	TO	
15-08	Reserved	
07-00	Reserved	

VSI2_TO Description

Name	Type	Reset By	Reset State	Function
TO[31:16]	R/W	PWR VME	0	Translation Offset

Table A.97 : VMEbus Slave Image 3 Control (VSI3_CTL)

Register Name: VSI3_CTL						Offset: F3C	
Bits	Function						
31-24	EN	PWEN	PREN	Reserved			
23-16	PGM		SUPER		Reserved	VAS	
15-08	Reserved						
07-00	LD64EN	LLRMW	Reserved				LAS

VSI3_CTL Description

Name	Type	Reset By	Reset State	Function
EN	R/W	PWR VME	0	Image Enable 0=Disable, 1=Enable
PWEN	R/W	PWR VME	0	Posted Write Enable 0=Disable, 1=Enable
PREN	R/W	PWR VME	0	Prefetch Read Enable 0=Disable, 1=Enable
PGM	R/W	PWR VME	11	Program/Data AM Code 00=Reserved, 01=Data, 10=Program, 11=Both
SUPER	R/W	PWR VME	11	Supervisor/User AM Code 00=Reserved, 01=Non-Privileged, 10=Supervisor, 11=Both
VAS	R/W	PWR VME	0	VMEbus Address Space 000=Reserved, 001=A24, 010=A32, 011= Reserved, 100=Reserved, 101=Reserved, 110=User1, 111=User2
LD64EN	R/W	PWR VME	0	Enable 64-bit PCI Bus Transactions 0=Disable, 1=Enable
LLRMW	R/W	PWR VME	0	Enable PCI Bus Lock of VMEbus RMW 0=Disable, 1=Enable
LAS	R/W	PWR VME	0	PCI Bus Address Space 00=PCI Bus Memory Space, 01=PCI Bus I/O Space, 10=PCI Bus Configuration Space, 11=Reserved

This register provides the general, VMEbus and PCI controls for this slave image. Note that only transactions destined for PCI Memory space are decoupled (the posted write RXFIFO generates on Memory space transactions on the PCI Bus).

In order for a VMEbus slave image to respond to an incoming cycle, the BM bit in the PCI_CSR register must be enabled.

The state of PWEN and PREN are ignored if LAS is not programmed memory space.

Table A.98 : VMEbus Slave Image 3 Base Address Register (VSI3_BS)

Register Name: VSI3_BS		Offset: F40
Bits	Function	
31-24	BS	
23-16	BS	
15-08	Reserved	
07-00	Reserved	

VSI3_BS Description

Name	Type	Reset By	Reset State	Function
BS[31:16]	R/W	PWR VME	0	Base Address

Table A.99 : VMEbus Slave Image 3 Bound Address Register (VSI3_BD)

Register Name: VSI3_BD		Offset: F44
Bits	Function	
31-24	BD	
23-16	BD	
15-08	Reserved	
07-00	Reserved	

VSI3_BD Description

Name	Type	Reset By	Reset State	Function
BD[31:16]	R/W	PWR VME	0	Bound Address

Table A.100 : VMEbus Slave Image 3 Translation Offset (VSI3_TO)

Register Name: VSI3_TO		Offset: F48
Bits	Function	
31-24	TO	
23-16	TO	
15-08	Reserved	
07-00	Reserved	

VSI3_TO Description

Name	Type	Reset By	Reset State	Function
TO[31:16]	R/W	PWR VME	0	Translation Offset

Table A.101 : Location Monitor Control Register (LM_CTL)

Register Name: LM_CTL				Offset: F64
Bits	Function			
31-24	EN	Reserved		
23-16	PGM	SUPER	Reserved	VAS
15-08	Reserved			
07-00	Reserved			

LM_CTL Description

Name	Type	Reset By	Reset State	Function
EN	R/W	PWR VME	0	Image Enable 0=Disable, 1=Enable
PGM	R/W	PWR VME	11	Program/Data AM Code 00=Reserved, 01=Data, 10=Program, 11=Both
SUPER	R/W	PWR VME	11	Supervisor/User AM Code 00=Reserved, 01=Non-Privileged, 10=Supervisor, 11=Both
VAS	R/W	PWR VME	0	VMEbus Address Space 000=A16, 001=A24, 010=A32, 011=Reserved, 100=Reserved, 101=Reserved, others=Reserved

This register specifies the VMEbus controls for the location monitor image. This image has a 4 Kbyte resolution and a 4 Kbyte size. The image responds to a VME read or write within the 4 Kbyte space and matching one of the address modifier codes specified. BLTs and MBLTs are not supported.

VMEbus address bits [4:3] are used to set the status bit in LINT_STAT for one of the four location monitor interrupts. If the Universe II VMEbus master is the owner of the VMEbus, the Universe II VMEbus slave will generate DTACK* to terminate the transaction.

The Location Monitor does not store write data and read data is undefined.

Table A.102 : Location Monitor Base Address Register (LM_BS)

Register Name: LM_BS		Offset: F68
Bits	Function	
31-24	BS	
23-16	BS	
15-08	Reserved	
07-00	Reserved	

LM_BS Description

Name	Type	Reset By	Reset State	Function
BS [31:16]	R/W	PWR VME	0	Base Address

The base address specifies the lowest address in the 4 Kbyte range that will be decoded as a location monitor access.

Table A.103 : VMEbus Register Access Image Control Register (VRAI_CTL)

Register Name: VRAI_CTL				Offset: F70
Bits	Function			
31-24	EN	Reserved		
23-16	PGM	SUPER	Reserved	VAS
15-08	Reserved			
07-00	Reserved			

VRAI_CTL Description

Name	Type	Reset By	Reset State	Function
EN	R/W	PWR VME	Power-up Option	Image Enable 0=Disable, 1=Enable
PGM	R/W	PWR VME	11	Program/Data AM Code 00=Reserved, 01=Data, 10=Program, 11=Both
SUPER	R/W	PWR VME	11	Supervisor/User AM Code 00=Reserved, 01=Non-Privileged, 10=Supervisor, 11=Both
VAS	R/W	PWR VME	Power-up Option	VMEbus Address Space 00=A16, 01=A24, 10=A32, all others are reserved

The VME Register Access Image allows access to the Universe II registers with standard VMEbus cycles. Only single cycle and lock AM codes are accepted. When a register is accessed with a RMW, it is locked for the duration of the transaction.

Table A.104 : VMEbus Register Access Image Base Address Register (VRAI_BS)

Register Name: VRAI_BS		Offset: F74
Bits	Function	
31-24	BS	
23-16	BS	
15-08	BS	Reserved
07-00	Reserved	

VRAI_BS Description

Name	Type	Reset By	Reset State	Function
BS[31:12]	R/W	PWR VME	Power-up Option	The base address specifies the lowest address in the 4 Kbyte VMEbus Register Access Image.

The base address specifies the lowest address in the 4 Kbyte VMEbus Register Access Image. The reset state is a function of the Power-up Option behaviour of the VAS field in VRAI_CTL:

VRAI_CTL: VAS	BS [31:24]	BS [23:16]	BS [15:12]
A16	0	0	Power-up Option VA [28:25]
A24	0	Power-up Option VA [28:21]	0
A32	Power-up Option VA [28:21]	0	0

Table A.105 : VMEbus CSR Control Register (VCSR_CTL)

Register Name: VCSR_CTL		Offset: F80
Bits	Function	
31-24	EN	Reserved
23-16	Reserved	
15-08	Reserved	
07-00	Reserved	LAS

VCSR_CTL Description

Name	Type	Reset By	Reset State	Function
EN	R/W	PWR,VME	0	Image Enable 0=Disable, 1=Enable
LAS	R/W	PWR VME	Power-up Option	PCI Bus Address Space 00=PCI Bus Memory Space, 01=PCI Bus I/O Space, 10=PCI Bus Configuration Space, 11=Reserved

The EN bit allows software to enable or disable the Universe II CR/CSR image. This image can also be enabled by the VME64 Auto ID process.

For CSR's not supported in the Universe II and for CR accesses, the LAS field determines the PCI Bus command that will be generated when the cycle is mapped to the PCI Bus.

Table A.106 : VMEbus CSR Translation Offset (VCSR_TO)

Register Name: VCSR_TO		Offset: F84
Bits	Function	
31-24	TO	
23-16	TO	Reserved
15-08	Reserved	
07-00	Reserved	

VCSR_TO Description

Name	Type	Reset By	Reset State	Function
TO [31:24]	R/W	PWR VME	0	Translation Offset
TO [23:19]	R/W	PWR VME	Power-up Option	Translation Offset

For CSR's not supported in the Universe II and for CR accesses, the translation offset is added to the 24-bit VMEbus address to produce a 32-bit PCI Bus address. Negative offsets are not supported.

Table A.107 : VMEbus AM Code Error Log (V_AMERR)

Register Name: V_AMERR			Offset: F88	
Bits	Function			
31-24	AMERR		IACK	M_ERR
23-16	V_STAT	Reserved		
15-08	Reserved			
07-00	Reserved			

V_AMERR Description

Name	Type	Reset By	Reset State	Function
AMERR [5:0]	R	PWR, VME	0	VMEbus AM Code Error Log
IACK	R	PWR, VME	0	VMEbus IACK Signal
M_ERR	R	PWR, VME	0	Multiple Error Occurred 0=Single error, 1=At least one error has occurred since the logs were frozen
V_STAT	R/W	PWR, VME	0	VME Error Log Status Reads: 0=logs invalid, 1=logs are valid and error logging halted Writes: 0=no effect, 1=clears V_STAT and enables error logging

The Universe II VMEbus Master Interface is responsible for logging the parameters of a posted write transaction that results in a bus error. This register holds the address modifier code and the state of the IACK* signal. The register contents are qualified by the V_STAT bit.

Table A.108 : VMEbus Address Error Log (VAERR)

Register Name: VAERR		Offset: F8C
Bits	Function	
31-24	VAERR	
23-16	VAERR	
15-08	VAERR	
07-00	VAERR	

VAERR Description

Name	Type	Reset By	Reset State	Function
VAERR [31:0]	R	PWR, VME	0	VMEbus address error log

The Universe II PCI Master Interface is responsible for logging the parameters of a posted write transaction that results in a bus error. This register holds the address. The register contents are qualified by the V_STAT bit of the V_AMERR register.

Table A.109 : VMEbus Slave Image 4 Control (VSI4_CTL)

Register Name: VSI4_CTL					Offset: F90
Bits	Function				
31-24	EN	PWEN	PREN	Reserved	
23-16	PGM		SUPER	Reserved	VAS
15-08	Reserved				
07-00	LD64EN	LLRMW	Reserved		LAS

VSI4_CTL Description

Name	Type	Reset By	Reset State	Function
EN	R/W	PWR, VME	0	Image Enable 0=Disable, 1=Enable
PWEN	R/W	PWR, VME	0	Posted Write Enable 0=Disable, 1=Enable
PREN	R/W	PWR, VME	0	Prefetch Read Enable 0=Disable, 1=Enable
PGM	R/W	PWR, VME	11	Program/Data AM Code 00=Reserved, 01=Data, 10=Program, 11=Both
SUPER	R/W	PWR, VME	11	Supervisor/User AM Code 00=Reserved, 01=Non-Privileged, 10=Supervisor, 11=Both
VAS	R/W	PWR, VME	0	VMEbus Address Space 000=A16, 001=A24, 010=A32, 011= Reserved, 100=Reserved, 101=Reserved, 110=User1, 111=User2
LD64EN	R/W	PWR, VME	0	Enable 64-bit PCI Bus Transactions 0=Disable, 1=Enable
LLRMW	R/W	PWR, VME	0	Enable PCI Bus Lock of VMEbus RMW 0=Disable, 1=Enable
LAS	R/W	PWR, VME	0	PCI Bus Address Space 00=PCI Bus Memory Space, 01=PCI Bus I/O Space, 10=PCI Bus Configuration Space, 11=Reserved

This register provides the general, VMEbus and PCI controls for this slave image. Note that only transactions destined for PCI Memory space are decoupled (the posted write RXFIFO generates on Memory space transactions on the PCI Bus). This image has 4 Kbyte resolution.

In order for a VMEbus slave image to respond to an incoming cycle, the BM bit in the PCI_CSR register must be enabled.

The state of PWEN and PREN are ignored if LAS is not programmed memory space.

Table A.110 : VMEbus Slave Image 4 Base Address Register (VSI4_BS)

Register Name: VSI4_BS		Offset: F94
Bits	Function	
31-24	BS	
23-16	BS	
15-08	BS	Reserved
07-00	Reserved	

VSI4_BS Description

Name	Type	Reset By	Reset State	Function
BS[31:12]	R/W	PWR VME	0	Base Address

The base address specifies the lowest address in the address range that will be decoded.
This image has 4 Kbyte resolution.

Table A.111 : VMEbus Slave Image 4 Bound Address Register (VSI4_BD)

Register Name: VSI4_BD		Offset: F98
Bits	Function	
31-24	BD	
23-16	BD	
15-08	BD	Reserved
07-00	Reserved	

VSI4_BD Description

Name	Type	Reset By	Reset State	Function
BD[31:12]	R/W	PWR VME	0	Bound Address

The addresses decoded in a slave image are those which are greater than or equal to the base address and less than the bound register. If the bound register is 0, then the addresses decoded are those greater than or equal to the base address.

This image has 4 Kbyte resolution.

Table A.112 : VMEbus Slave Image 4 Translation Offset (VSI4_TO)

Register Name: VSI4_TO		Offset: F9C
Bits	Function	
31-24	TO	
23-16	TO	
15-08	TO	Reserved
07-00	Reserved	

VSI4_TO Description

Name	Type	Reset By	Reset State	Function
TO[31:12]	R/W	PWR VME	0	Translation Offset

The translation offset is added to the source address that is decoded and this new address becomes the destination address. If a negative offset is desired, the offset must be expressed as a two's complement.

This image has 4 Kbyte resolution.

Table A.113 : VMEbus Slave Image 5 Control (VSI5_CTL)

Register Name: VSI5_CTL						Offset: FA4	
Bits	Function						
31-24	EN	PWEN	PREN	Reserved			
23-16	PGM		SUPER		Reserved	VAS	
15-08	Reserved						
07-00	LD64EN	LLRMW	Reserved				LAS

VSI5_CTL Description

Name	Type	Reset By	Reset State	Function
EN	R/W	PWR VME	0	Image Enable 0=Disable, 1=Enable
PWEN	R/W	PWR VME	0	Posted Write Enable 0=Disable, 1=Enable
PREN	R/W	PWR VME	0	Prefetch Read Enable 0=Disable, 1=Enable
PGM	R/W	PWR VME	11	Program/Data AM Code 00=Reserved, 01=Data, 10=Program, 11=Both
SUPER	R/W	PWR VME	11	Supervisor/User AM Code 00=Reserved, 01=Non-Privileged, 10=Supervisor, 11=Both
VAS	R/W	PWR VME	0	VMEbus Address Space 000=Reserved, 001=A24, 010=A32, 011= Reserved, 100=Reserved, 101=Reserved, 110=User1, 111=User2
LD64EN	R/W	PWR VME	0	Enable 64-bit PCI Bus Transactions 0=Disable, 1=Enable
LLRMW	R/W	PWR VME	0	Enable PCI Bus Lock of VMEbus RMW 0=Disable, 1=Enable
LAS	R/W	PWR VME	0	PCI Bus Address Space 00=PCI Bus Memory Space, 01=PCI Bus I/O Space, 10=PCI Bus Configuration Space, 11=Reserved

This register provides the general, VMEbus and PCI controls for this slave image. Note that only transactions destined for PCI Memory space are decoupled (the posted write RXFIFO generates on Memory space transactions on the PCI Bus).

In order for a VMEbus slave image to respond to an incoming cycle, the BM bit in the PCI_CSR register must be enabled.

The state of PWEN and PREN are ignored if LAS is not programmed memory space.

Table A.114 : VMEbus Slave Image 5 Base Address Register (VSI5_BS)

Register Name: VSI5_BS		Offset: FA8
Bits	Function	
31-24	BS	
23-16	BS	
15-08	Reserved	
07-00	Reserved	

VSI5_BS Description

Name	Type	Reset By	Reset State	Function
BS[31:16]	R/W	PWR VME	0	Base Address

The base address specifies the lowest address in the address range that will be decoded.

Table A.115 : VMEbus Slave Image 5 Bound Address Register (VSI5_BD)

Register Name: VSI5_BD		Offset: FAC
Bits	Function	
31-24	BD	
23-16	BD	
15-08	Reserved	
07-00	Reserved	

VSI5_BD Description

Name	Type	Reset By	Reset State	Function
BD[31:16]	R/W	PWR VME	0	Bound Address

The addresses decoded in a slave image are those which are greater than or equal to the base address and less than the bound register. If the bound register is 0, then the addresses decoded are those greater than or equal to the base address,

Table A.116 : VMEbus Slave Image 5 Translation Offset (VSI5_TO)

Register Name: VSI5_TO		Offset: FB0
Bits	Function	
31-24	TO	
23-16	TO	
15-08	Reserved	
07-00	Reserved	

VSI5_TO Description

Name	Type	Reset By	Reset State	Function
TO[31:16]	R/W	PWR VME	0	Translation Offset

The translation offset is added to the source address that is decoded and this new address becomes the destination address. If a negative offset is desired, the offset must be expressed as a two's complement.

Table A.117 : VMEbus Slave Image 6 Control (VSI6_CTL)

Register Name: VSI6_CTL						Offset: FB8	
Bits	Function						
31-24	EN	PWEN	PREN	Reserved			
23-16	PGM		SUPER		Reserved	VAS	
15-08	Reserved						
07-00	LD64EN	LLRMW	Reserved				LAS

VSI6_CTL Description

Name	Type	Reset By	Reset State	Function
EN	R/W	PWR VME	0	Image Enable 0=Disable, 1=Enable
PWEN	R/W	PWR VME	0	Posted Write Enable 0=Disable, 1=Enable
PREN	R/W	PWR VME	0	Prefetch Read Enable 0=Disable, 1=Enable
PGM	R/W	PWR VME	11	Program/Data AM Code 00=Reserved, 01=Data, 10=Program, 11=Both
SUPER	R/W	PWR VME	11	Supervisor/User AM Code 00=Reserved, 01=Non-Privileged, 10=Supervisor, 11=Both
VAS	R/W	PWR VME	0	VMEbus Address Space 000=Reserved, 001=A24, 010=A32, 011= Reserved, 100=Reserved, 101=Reserved, 110=User1, 111=User2
LD64EN	R/W	PWR VME	0	Enable 64-bit PCI Bus Transactions 0=Disable, 1=Enable
LLRMW	R/W	PWR VME	0	Enable PCI Bus Lock of VMEbus RMW 0=Disable, 1=Enable
LAS	R/W	PWR VME	0	PCI Bus Address Space 00=PCI Bus Memory Space, 01=PCI Bus I/O Space, 10=PCI Bus Configuration Space, 11=Reserved

This register provides the general, VMEbus and PCI controls for this slave image. Note that only transactions destined for PCI Memory space are decoupled (the posted write RXFIFO generates on Memory space transactions on the PCI Bus).

In order for a VMEbus slave image to respond to an incoming cycle, the BM bit in the PCI_CSR register must be enabled.

The state of PWEN and PREN are ignored if LAS is not programmed memory space.

Table A.118 : VMEbus Slave Image 6 Base Address Register (VSI6_BS)

Register Name: VSI6_BS		Offset: FBC
Bits	Function	
31-24	BS	
23-16	BS	
15-08	Reserved	
07-00	Reserved	

VSI6_BS Description

Name	Type	Reset By	Reset State	Function
BS[31:16]	R/W	PWR VME	0	Base Address

The base address specifies the lowest address in the address range that will be decoded.

Table A.119 : VMEbus Slave Image 6 Bound Address Register (VSI6_BD)

Register Name: VSI6_BD		Offset: FC0
Bits	Function	
31-24	BD	
23-16	BD	
15-08	Reserved	
07-00	Reserved	

VSI6_BD Description

Name	Type	Reset By	Reset State	Function
BD[31:16]	R/W	PWR VME	0	Bound Address

The addresses decoded in a slave image are those which are greater than or equal to the base address and less than the bound register. If the bound register is 0, then the addresses decoded are those greater than or equal to the base address,

Table A.120 : VMEbus Slave Image 6 Translation Offset (VSI6_TO)

Register Name: VSI6_TO		Offset: FC4
Bits	Function	
31-24	TO	
23-16	TO	
15-08	Reserved	
07-00	Reserved	

VSI6_TO Description

Name	Type	Reset By	Reset State	Function
TO[31:16]	R/W	PWR VME	0	Translation Offset

The translation offset is added to the source address that is decoded and this new address becomes the destination address. If a negative offset is desired, the offset must be expressed as a two's complement.

Table A.121 : VMEbus Slave Image 7 Control (VSI7_CTL)

Register Name: VSI7_CTL						Offset: FCC	
Bits	Function						
31-24	EN	PWEN	PREN	Reserved			
23-16	PGM		SUPER		Reserved	VAS	
15-08	Reserved						
07-00	LD64EN	LLRMW	Reserved				LAS

VSI7_CTL Description

Name	Type	Reset By	Reset State	Function
EN	R/W	PWR VME	0	Image Enable 0=Disable, 1=Enable
PWEN	R/W	PWR VME	0	Posted Write Enable 0=Disable, 1=Enable
PREN	R/W	PWR VME	0	Prefetch Read Enable 0=Disable, 1=Enable
PGM	R/W	PWR VME	11	Program/Data AM Code 00=Reserved, 01=Data, 10=Program, 11=Both
SUPER	R/W	PWR VME	11	Supervisor/User AM Code 00=Reserved, 01=Non-Privileged, 10=Supervisor, 11=Both
VAS	R/W	PWR VME	0	VMEbus Address Space 000=Reserved, 001=A24, 010=A32, 011= Reserved, 100=Reserved, 101=Reserved, 110=User1, 111=User2
LD64EN	R/W	PWR VME	0	Enable 64-bit PCI Bus Transactions 0=Disable, 1=Enable
LLRMW	R/W	PWR VME	0	Enable PCI Bus Lock of VMEbus RMW 0=Disable, 1=Enable
LAS	R/W	PWR VME	0	PCI Bus Address Space 00=PCI Bus Memory Space, 01=PCI Bus I/O Space, 10=PCI Bus Configuration Space, 11=Reserved

This register provides the general, VMEbus and PCI controls for this slave image. Note that only transactions destined for PCI Memory space are decoupled (the posted write RXFIFO generates on Memory space transactions on the PCI Bus).

In order for a VMEbus slave image to respond to an incoming cycle, the BM bit in the PCI_CSR register must be enabled.

The state of PWEN and PREN are ignored if LAS is not programmed memory space.

Table A.122 : VMEbus Slave Image 7 Base Address Register (VSI7_BS)

Register Name: VSI7_BS		Offset: FD0
Bits	Function	
31-24	BS	
23-16	BS	
15-08	Reserved	
07-00	Reserved	

VSI7_BS Description

Name	Type	Reset By	Reset State	Function
BS[31:16]	R/W	PWR VME	0	Base Address

The base address specifies the lowest address in the address range that will be decoded.

Table A.123 : VMEbus Slave Image 7 Bound Address Register (VSI7_BD)

Register Name: VSI7_BD		Offset: FD4
Bits	Function	
31-24	BD	
23-16	BD	
15-08	Reserved	
07-00	Reserved	

VSI7_BD Description

Name	Type	Reset By	Reset State	Function
BD[31:16]	R/W	PWR VME	0	Bound Address

The addresses decoded in a slave image are those which are greater than or equal to the base address and less than the bound register. If the bound register is 0, then the addresses decoded are those greater than or equal to the base address,

Table A.124 : VMEbus Slave Image 7 Translation Offset (VSI7_TO)

Register Name: VSI7_TO		Offset: FD8
Bits	Function	
31-24	TO	
23-16	TO	
15-08	Reserved	
07-00	Reserved	

VSI7_TO Description

Name	Type	Reset By	Reset State	Function
TO[31:16]	R/W	PWR VME	0	Translation Offset

Table A.125 : VMEbus CSR Bit Clear Register (VCSR_CLR)

Register Name: VCSR_CLR				Offset: FF4
Bits	Function			
31-24	RESET	SYSFAIL	FAIL	Reserved
23-16	Reserved			
15-08	Reserved			
07-00	Reserved			

VCSR_CLR Description

Name	Type	Reset By	Reset State	Function
RESET	R/W	PWR VME	0	Board Reset Reads: 0=LRST# not asserted, 1=LRST# asserted Writes: 0=no effect, 1=negate LRST#
SYSFAIL	R/W	all	Power-up Option	VMEbus SYSFAIL Reads: 0=VXSYSFAIL not asserted, 1=VXSYSFAIL asserted Writes: 0=no effect, 1=negate VXSYSFAIL
FAIL	R	PWR VME	0	Board Fail 0=Board has not failed

This register implements the Bit Clear Register as defined in the VME64 specification. Note that the RESET bit can be written to only from the VMEbus.

Table A.126 : VMEbus CSR Bit Set Register (VCSR_SET)

Register Name: VCSR_SET				Offset: FF8
Bits	Function			
31-24	RESET	SYSFAIL	FAIL	Reserved
23-16	Reserved			
15-08	Reserved			
07-00	Reserved			

VCSR_SET Description

Name	Type	Reset By	Reset State	Function
RESET	R/W	PWR VME	0	Board Reset Reads: 0=LRST# not asserted, 1=LRST# asserted Writes: 0=no effect, 1=assert LRST#
SYSFAIL	R/W	all	Power-up Option	VMEbus SYSFAIL Reads: 0=VXSYSFAIL not asserted, 1=VXSYSFAIL asserted Writes: 0=no effect, 1=assert VXSYSFAIL
FAIL	R	PWR VME	0	Board Fail 0=Board has not failed

This register implements the Bit Set Register as defined in the VME64 specification. Note that the RESET bit can be written to only from the VMEbus. Writing 1 to the RESET bit asserts LRST#. The PCI reset remains asserted until a 1 is written to the RESET bit of the VCSR_CLR register.

Table A.127 : VMEbus CSR Base Address Register (VCSR_BS)

Register Name: VCSR_BS		Offset: FFC
Bits	Function	
31-24	BS	Reserved
23-16	Reserved	
15-08	Reserved	
07-00	Reserved	

VCSR_BS Description

Name	Type	Reset By	Reset State	Function
BS [23:19]	R/W	PWR VME	0	Base Address

The base address specifies one of thirty-one available CR/CSR windows as defined in the VME64 specification. Each window consumes 512 Kbytes of CR/CSR space.

Appendix B Performance

As a VMEbus bridge, the Universe II's most important function is data transfer. This function is performed by its three channels: the PCI Slave Channel, the VME Slave Channel, and the DMA Channel. Since each channel operates independently of the others and because each has its own unique characteristics, the following analysis reviews the data transfer performance for each channel:

- “PCI Slave Channel” on page B-2
- “VME Slave Channel” on page B-6
- “DMA Channel” on page B-14
- “Summary” on page B-17

Where relevant, descriptions of factors affecting performance and how they might be controlled in different environments are discussed.

The decoupled nature of the Universe II can cause some confusion in discussing performance parameters. This is because, in a fully decoupled bus bridge each of the two opposing buses operates at its peak performance independently of the other. The Universe II, however, because of the finite size of its FIFOs does not represent a 100% decoupled bridge. As the FIFOs fill or empty (depending on the direction of data movement) the two buses tend to migrate to matched performance where the higher performing bus is forced to slow down to match the other bus. This limits the sustained performance of the device. Some factors such as the PCI Aligned Burst Size and VME request/release modes can limit the effect of FIFO size and enhance performance.

Another aspect in considering the performance of a device is bandwidth consumption. The greater bandwidth consumed to transfer a given amount of data, the less is available for other bus masters. Decoupling significantly improves the Universe II's bandwidth consumption, and on the PCI bus allows it to use the minimum permitted by the PCI specification.

To simplify the analysis and allow comparison with other devices, Universe II performance has been calculated using the following assumptions:

As a PCI master:

- one clock bus grant latency
- zero wait state PCI target

As a VME master:

- ideal VME slave response (DS^* to $DTACK^* = 30\text{ns}$)

Assumed as part of any calculation on VME performance is the inclusion of VME transceivers with propagation delay of 4 ns.

This appendix presents sustained performance values. In contrast, the Universe User manual (9000000.MD303.01) provided *peak* performance numbers. This explains why some of the performance numbers in this document appear to be lower than for the original Universe.

B.1 PCI Slave Channel

B.1.1 Coupled Cycles

B.1.1.1 Request of VMEbus

The Universe II has a "Coupled Window Timer" (CWT in the LMISC register) which permits the coupled channel to maintain ownership of the VMEbus for an extended period beyond the completion of a cycle. This permits subsequent coupled accesses to the VMEbus to occur back-to-back without requirement for re-arbitration.

The CWT should be set for the expected latency between sequential coupled accesses attempted by the CPU. In calculating the latency expected here, the designer needs to account for latency across their host PCI bridge as well as latency encountered in re-arbitration for the PCI bus between each coupled access. Care must be taken not to set the CWT greater than necessary as the Universe II blocks all decoupled write transactions with target-retry, while the coupled channel owns the VMEbus. It is only when the CWT has expired that the PCI bus is permitted to enqueue transactions in the TXFIFO.

When a coupled access to the VMEbus is attempted, the Universe II generates a target-retry to the PCI initiator if the coupled path does not currently own the VMEbus. This occurs if the Universe II is not currently VMEbus master, or if the DMA is currently VMEbus master or if entries exist in the TXFIFO.

If the Universe II does not have ownership of the VMEbus when a coupled access is attempted, the Universe II generates a target-retry with a single wait state (See Figure B.1). The request for the VMEbus occurs shortly after the cycle is retried.

B.1.1.2 Read Cycles

Once the coupled channel owns the VMEbus, the Universe II propagates the cycle out to the VMEbus. Figure B.1 shows such a coupled read cycle against an ideal VME slave. There are 10 wait states inserted by the Universe II on the PCI bus before it responds with TRDY#. Further wait states are inserted for each extra 30ns in slave response.

Performing 32-bit PCI reads from VME gives a sustained performance of approximately 8.5 MB/s. Figure B.2 shows several of these accesses occurring consecutively.

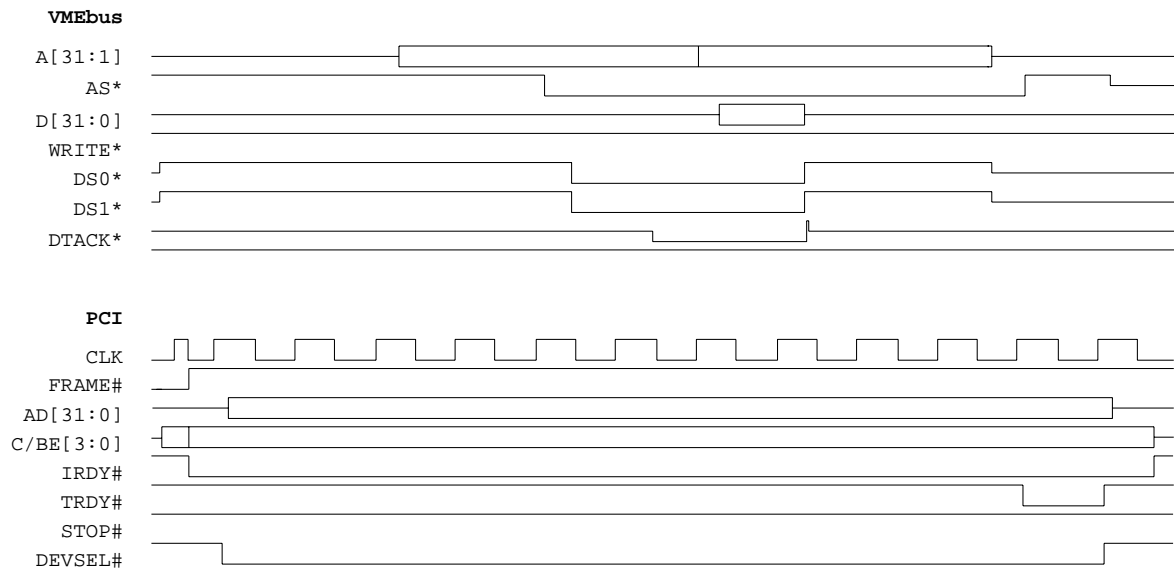


Figure B.1 : Coupled Read Cycle - Universe II as VME Master

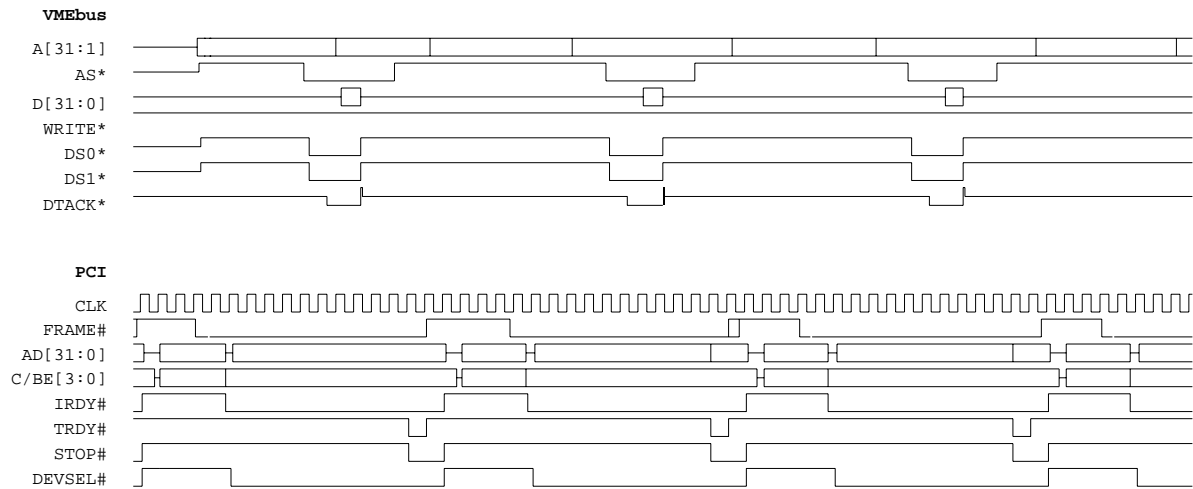


Figure B.2 : Several Coupled Read Cycles - Universe II as VME Master

B.1.1.3 Write Cycles

The performance of coupled write cycles is similar to that of coupled read cycles except that an extra wait state is inserted. Figure B.3 shows a coupled write cycle against an ideal VME slave. Ten wait states are inserted on the PCI bus by the Universe II before it responds with TRDY#.

A slower VME slave response translates directly to more wait states on the PCI bus.

The sustained performance, when generating write cycles from a 32-bit PCI bus against an ideal VME slave is approximately 9.3 MB/s.

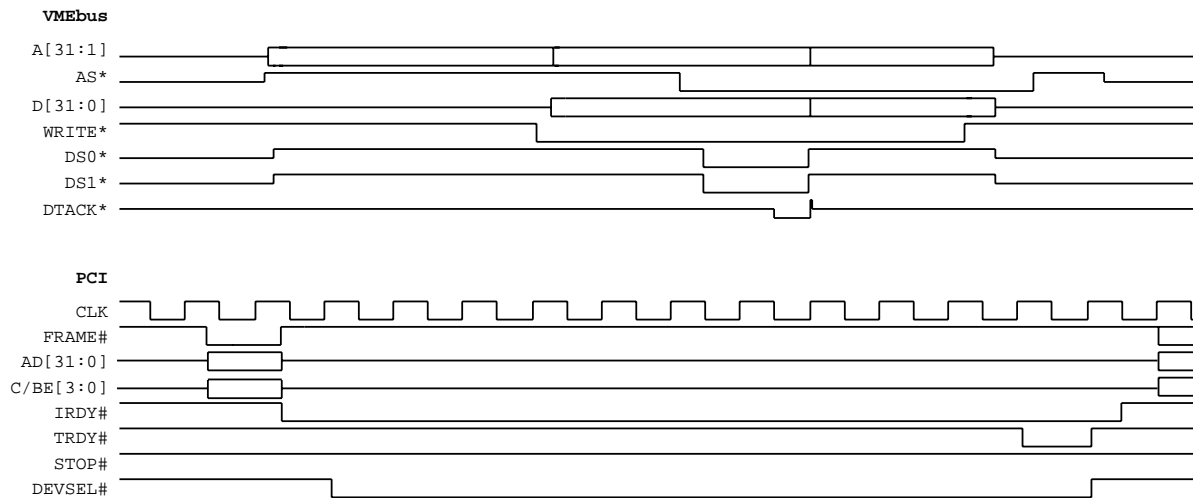


Figure B.3 : Coupled Write Cycle - Universe II as VME Master

B.1.2 Decoupled Cycles

Only write transactions can be decoupled in the PCI Target Channel.

Effect of the PWON Counter

The Posted Write On Counter (PWON in the MAST_CTL register) controls the maximum tenure that the PCI Slave Channel will have on the VMEbus. Once this channel has gained ownership of the VMEbus for use by the TXFIFO, it only relinquishes it if the FIFO becomes empty or if the number of bytes programmed in the counter expires. In most situations, the FIFO empties before the counter expires. However, if a great deal of data is being transferred by a PCI initiator to the VMEbus, then this counter ensures that only a fixed amount of VME bandwidth is consumed.

Limiting the size of the PWON counter imposes greater arbitration overhead on data being transferred out from the FIFO. This is true even when programmed for ROR mode since an internal arbitration cycle will still occur. The value for the PWON counter must be weighed from the system perspective with the impact of imposing greater latency on other channels (the DMA and Interrupt Channels) and other VME masters in gaining ownership of the VMEbus. On a Universe II equipped card which is only performing system control functions, the counter would be set to minimum. On a card which is responsible for transferring considerable amounts of performance-critical data the counter will be set much higher at the expense of system latency.

PCI Target Response

As the PCI target during decoupled write operations to the VMEbus, the Universe II responds in one of two manners:

1. It immediately issues a target retry because the FIFO does not have sufficient room for a burst of one address phase and 128 bytes of data. (There are no programmable watermarks in the PCI Target Channel. The PCI Aligned Burst Size (PABS) does not affect the PCI Target Channel.)
2. It responds as a zero-wait state target receiving up to 256 bytes in a transaction. When the FIFO is full or a 256-byte boundary has been reached, the Universe II issues a Target-Disconnect.

In either case, the Universe II will consume the minimum possible PCI bandwidth, never inserting wait states.

VME Master Performance

As a VME master, the Universe II waits until a full transaction has been enqueued in the Tx-FIFO before requesting the VMEbus and generating a VME cycle. If the VMEbus is already owned by the decoupled path (see “Effect of the PWON Counter” on page B-4), the Universe II still waits until a full transaction is enqueued in the FIFO before processing it.

If configured to generate non-block transfers, the Universe II can generate back-to-back VME transfers with cycle times of approximately 180ns (AS* to AS*) against an ideal VME slave (30-45 ns). A greater cycle time is required between the termination of one full enqueued transaction and the start of the next. This inter-transaction time is approximately 210ns. As such, the longer the PCI transaction, the greater the sustained performance on the VMEbus. With 64-byte PCI transactions, the sustained rate is 43 MB/s. With 32-byte transactions, this drops to 23 MB/s. Each of these numbers is calculated with no initial arbitration or re-arbitration for the bus. Figure B.4 shows the Universe II dequeuing a transaction with multiple non-block VME transfers.

Block transfers significantly increase performance. The inter-transaction period remains at approximately 210 ns for BLTs and MBLTs, but the data beat cycle time (DS* to DS*) drops to

about 120ns against the same ideal slave. Again the length of the burst size affects the sustained performance because of the inter-transaction time. For BLTs operating with a burst size of 64 bytes, the sustained performance is 37 MB/s, dropping to 33 MB/s for a burst size of 32 bytes. MBLTs operating with 64-byte bursts perform at a sustained rate of 66 MB/s, dropping to 50 MB/s for 32 bytes.




Figure B.4 : Several Non-Block Decoupled Writes - Universe II as VME Master




Figure B.5 : BLT Decoupled Write - Universe II as VME Master

B.2 VME Slave Channel

B.2.1 Coupled Cycles

B.2.1.1 Block vs. non-Block Transfers

The Universe II VME Slave Channel handles both block and non-block coupled accesses in similar manners. Each data beat is translated to a single PCI transaction. Once the transaction has been acknowledged on the PCI bus, the Universe II asserts DTACK* to terminate the VME data beat.

A non-block transfer and the first beat of a BLT transfer have identical timing. In each, the Universe II decodes the access and then provides a response to the data beat. Subsequent data

beats in the BLT transfer are shorter than the first due to the fact that no address decoding need be performed in these beats.

MBLT transfers behave somewhat differently. The first beat of an MBLT transfer is address only, and so the response is relatively fast. Subsequent data beats require acknowledgment from the PCI bus. With a 32-bit PCI bus, the MBLT data beat (64 bits of data) requires a two data beat PCI transaction. Because of this extra data beat required on the PCI bus, the slave response of the Universe II during coupled MBLT cycles is at least one PCI clock greater (depending upon the response from the PCI target) than that during BLT cycles.

B.2.1.2 Read Cycles

During coupled cycles, the Universe II does not acknowledge a VME transaction until it has been acknowledged on the PCI bus. Because of this the VME slave response during coupled reads is directly linked to the response time for the PCI target. Each clock of latency in the PCI target response translates directly to an extra clock of latency in the Universe II's VME coupled slave response.

The address of an incoming VME transaction is decoded and translated to an equivalent PCI transaction. Typically, four PCI clock periods elapse between the initial assertion of AS* on the VMEbus and the assertion of REQ# on the PCI bus. During the data only portion of subsequent beats in block transfers, the time from DS* assertion to REQ# is about 4 clocks. If the PCI bus is parked at the Universe II, no REQ# is asserted and FRAME# is asserted 4 clocks after AS*.

From assertion of REQ#, the Universe II does not insert any extra wait states in its operations as an initiator on the PCI bus. Upon receiving GNT# asserted, the Universe II asserts FRAME# in the next clock and after the required turn-around phase, asserts IRDY# to begin data transfer.

Once TRDY# is sampled asserted, the Universe II responds back to the VMEbus by asserting DTACK*. If the initiating VME transaction is 64-bit and the PCI bus or PCI bus target are 32 bit, then two data transfers are required on PCI before the Universe II can respond with DTACK*. No wait states are inserted by the Universe II between these two data beats on PCI. The assertion of DTACK* from the assertion of TRDY# has a latency of 1 clock. Figure B.6 shows a typical non-block coupled read cycle.

When accessing a PCI target with a zero wait state response, the Universe II VME response becomes approximately 10 PCI clock periods (about 301ns in a 33MHz system) during single cycles, and the first beat of a BLT. During pure data beats in both BLT and MBLTs, the slave response becomes 8 clocks.

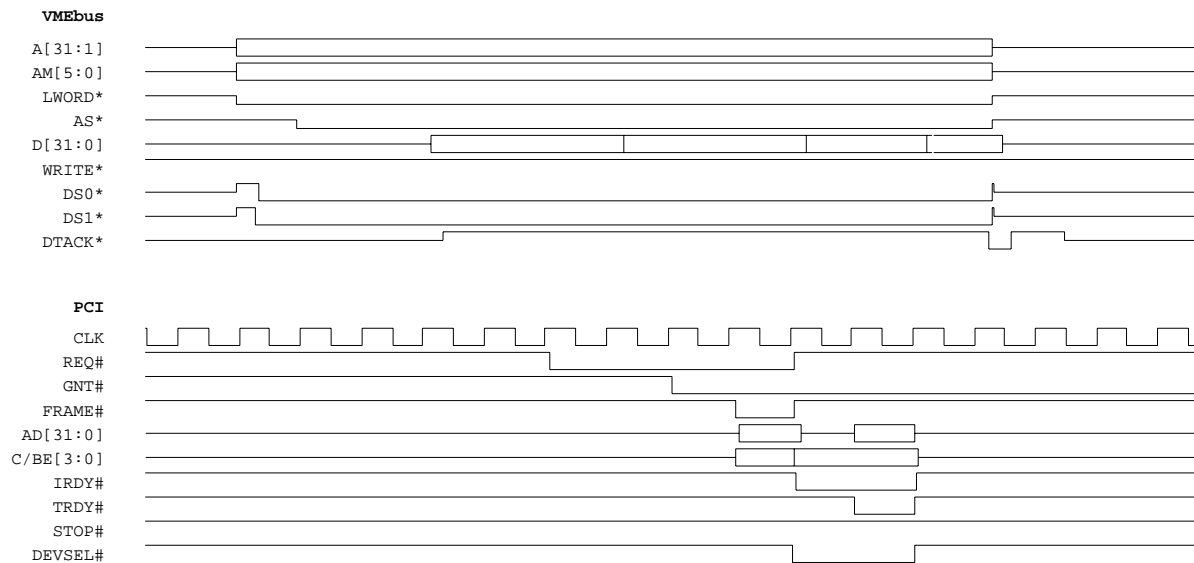


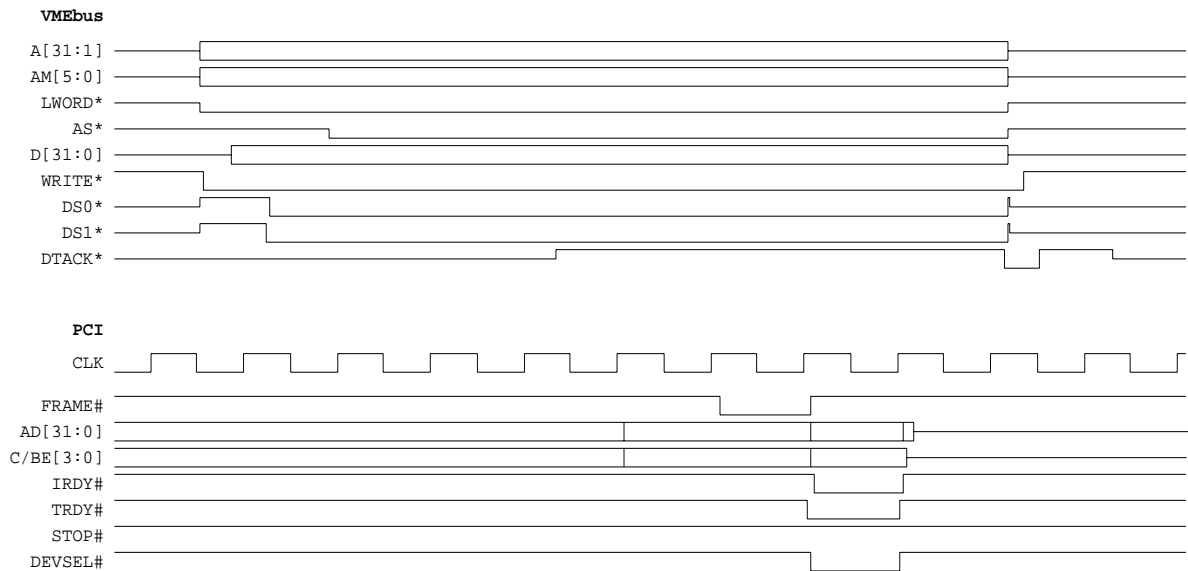
Figure B.6 : Coupled Read Cycle - Universe II as VME Slave

B.2.1.3 Write Cycles

Coupled writes in the VME Slave Channel operate in a similar fashion to the coupled reads. The VME slave response is directly linked to the response of the PCI target. In generating the request to the PCI bus, coupled write cycles require one further clock over reads. Hence, during single cycles, or the first beat of a BLT, the time from AS* to REQ# asserted is 3-4 PCI clocks, while DS* to REQ# is 3 clocks for the data beat portion of a block transfer. If the PCI bus is parked at the Universe II, REQ# is not asserted and the transaction begins immediately with assertion of FRAME#.

As with reads, the response from the PCI target's assertion of TRDY# to DTACK* assertion by the Universe II adds one clock to the transfer. Figure B.7 shows a typical non-block coupled write cycle.

Because write cycles on the PCI bus require one less clock than reads, due to the absence of the turn-around phase between address and data phases, the overall slave response during coupled writes works out to the same as coupled reads against an identical target. In accessing a zero-wait state PCI target, the Universe II's coupled write slave response then is approximately 10 PCI clocks. During subsequent data beats of a block transfer (either BLT or MBLT), the slave response (DS* to DTACK*) is 8 clocks.



**Figure B.7 : Coupled Write Cycle - Universe II as VME Slave
(bus parked at Universe II)**

B.2.2 Decoupled Cycles

B.2.2.1 Write Cycles

Effect of the PCI Aligned Burst Size

The PCI Aligned Burst Size (PABS in the MAST_CTL register) affects the maximum burst size that the Universe II generates onto the PCI bus; either 32, 64, or 128 bytes. Note that the VME Slave Channel only generates PCI bursts in response to incoming block transfers.

The greater burst size means less arbitration and addressing overhead. However, incumbent in this is the greater average latency for other devices in the PCI system. Hence, in the VME Slave Channel, the burst size is a trade-off between performance and latency.

VME Slave Response

As a VME slave, the Universe II accepts data into its RXFIFO with minimum delay provided there is room in the FIFO for a further data beat. Assertion of DTACK* is delayed if there is insufficient room in the FIFO for the next data beat.

During non-block transfers, the Universe II must both decode the address and enqueue the data before asserting DTACK* to acknowledge the transfer. Because of this, the slave response

during non-block transfers is considerably slower than block transfers. This slave response time is 127ns.

During BLT transfers, the slave response in the first data beat being both address decode and data transfer is the same as a non-block transfer, i.e., 127ns. Subsequent data beats, however, are much faster. Response time for these is 50 to 56ns.

During MBLT transfers, the first phase is address only and the slave response is 127ns. Subsequent phases are data only and so the slave response is the same as with BLTs i.e., 50 to 56ns.

Note that the slave response is independent of the data size. D16 non-block transfers have a slave response identical to D32. BLT data beats have slave responses identical to MBLT data beats.

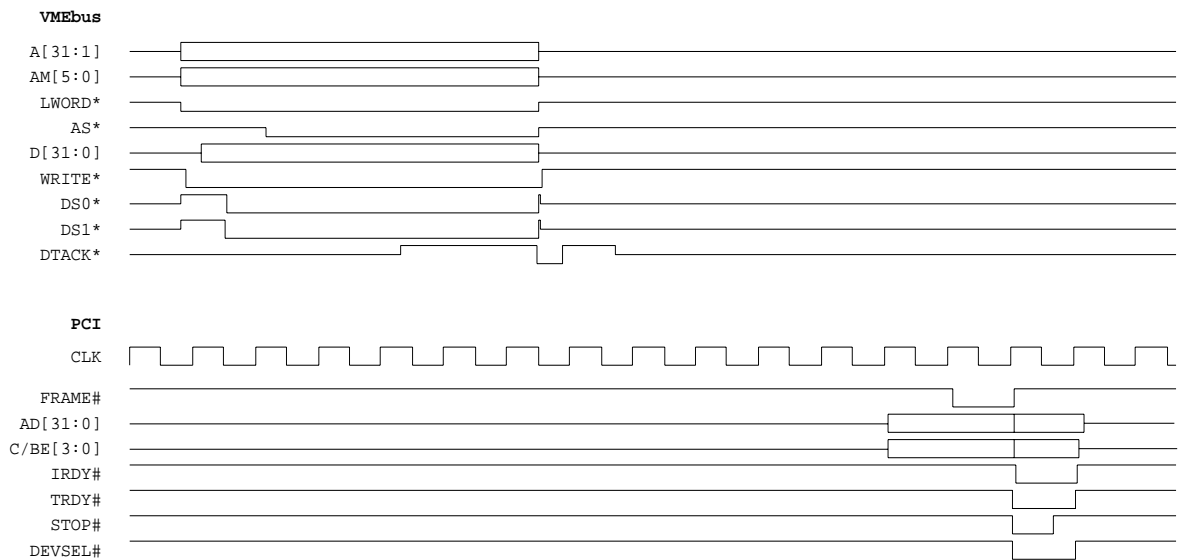


Figure B.8 : Non-Block Decoupled Write Cycle - Universe II as VME Slave

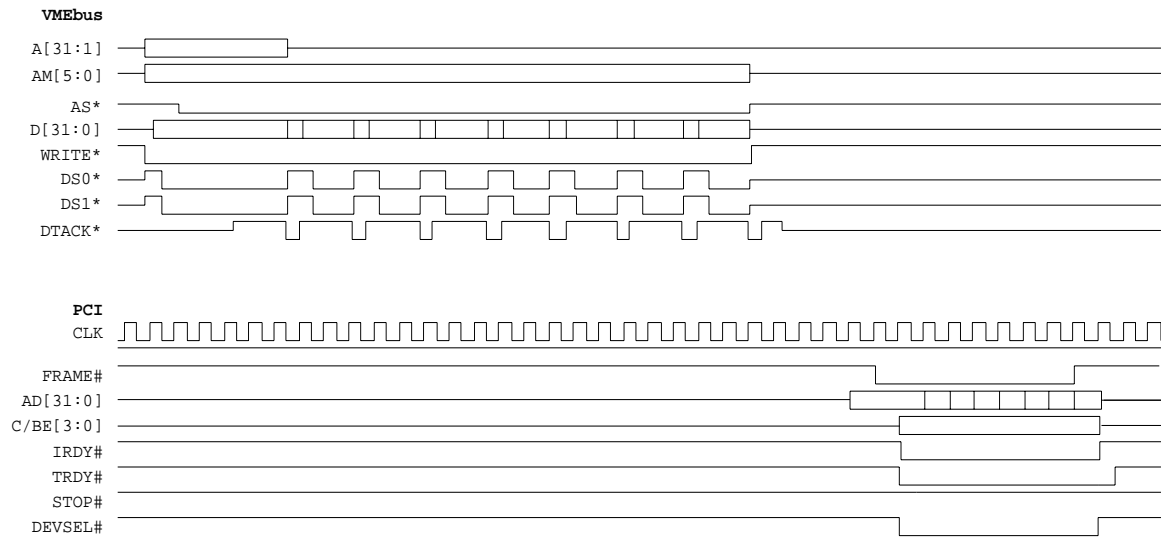


Figure B.9 : BLT Decoupled Write Cycle - Universe II as VME Slave

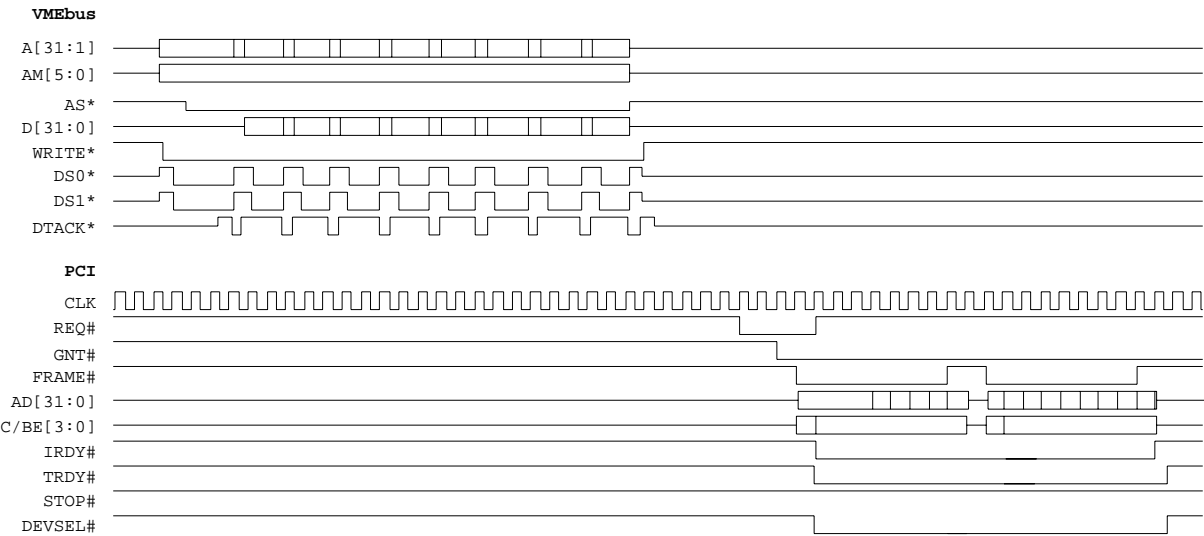


Figure B.10 : MBLT Decoupled Write Cycle - Universe II as VME Slave

PCI Master Performance

The Universe II supports bus parking. If the Universe II requires the PCI bus it will assert REQ# only if its GNT# is not currently asserted. When the PCI Master Module is ready to begin a transaction and its GNT# is asserted, the transfer begins immediately. This eliminates a possible one clock cycle delay before beginning a transaction on the PCI bus which would exist if the Universe II did not implement bus parking. Bus parking is described in Section 3.4.3 of the PCI Specification (Rev. 2.1).

On the PCI bus, the Universe II dequeues data from the RXFIFO once a complete VME transaction has been enqueued or once sufficient data has been enqueued to form a PCI transaction of length defined by the PABS field.

Since the Universe II does not perform any address phase deletion, non-block transfers are dequeued from the RXFIFO as single data beat transactions. Only block transfers result in multi-data beat PCI transactions; typically 8, 16 or 32 data beats. In either case, the Universe II does not insert any wait states as a PCI master. The clock, after the bus has been granted to the Universe II, drives out FRAME# to generate the address phase. The data phases begin immediately on the next clock. If there is more than one data phase, each phase will immediately follow the acknowledgment of the previous phase.

In each case, because of the lack of any wait states as a PCI master, the Universe II is consuming the minimum possible bandwidth on the PCI bus, and data will be written to the PCI bus at an average sustained rate equal to the rate at which the VME master is capable of writing it.

The sustained performance on the PCI bus performing single data beat write transactions to a 32-bit PCI bus is 15 MB/s; double this for a 64-bit bus. When performing 32-byte transactions the sustained performance increases to 106 MB/s; 120 MB/s with 64-byte transactions. Again, these can be doubled for a 64-bit PCI bus. Bear in mind that the PCI bus can only dequeue data as fast as it is being enqueued on the VMEbus. Hence, as the RXFIFO empties, the sustained performance on the PCI will drop down to match the lower performance on the VME side. However, even with the decreased sustained performance, the consumed bandwidth will remain constant (no extra wait states are inserted while the Universe II is master of the PCI bus.)

These numbers assume the PCI bus is granted to the Universe II immediately and that the writes are to a zero-wait state PCI target capable of accepting the full burst length. Figure B.1 through Figure B.10 show the Universe II responding to non-block, BLT and MBLT write transactions to a 32-bit PCI bus. Even better performance is obtained with PCI bus parking.

B.2.2.2 Prefetched Read Cycles

To minimize its slave response, the Universe II generates prefetched reads to the PCI bus in response to BLT and MBLT reads coming in from the VMEbus. This option must first be enabled on a per image basis.

When enabled, the Universe II will respond to a block read by performing burst reads on the

PCI bus of length defined by the PCI Aligned Burst Size (PABS in the MAST_CTL register). These burst reads continue while the block transfer is still active on the VMEbus (AS* not negated) and there is room in the RDFIFO. If there is insufficient room in the RDFIFO to continue (a common occurrence since the Universe II is capable of fetching data from the PCI bus at a much faster rate than a VME master is capable of receiving it), then pre-fetching stops and only continues once enough room exists in the RDFIFO for another full burst size.

The first data beat of a block transfer must wait for the first data beat to be retrieved from the PCI bus—this is essentially a coupled transfer. See the section on coupled transfers for details on coupled performance. However, once the pre-fetching begins, data is provided by the Universe II in subsequent data beats with a slave response of 57ns. This continues while there is data in the RDFIFO. If the RDFIFO empties because data is being fetched from the PCI bus too slowly, wait states are inserted on the VMEbus awaiting the enqueueing of more data.

On the PCI bus, the Universe II fetches data at 89 MB/s with PABS set to 32-byte transactions; 106 MB/s when set to 64-byte transactions. Even better performance is obtained if PABS is set for 128-byte transactions. Once the RDFIFO fills, pre-fetching slows to match the rate at which it is being read by the external VMEbus master. Bandwidth consumption, however, remains constant, only the idle time between transactions increases.

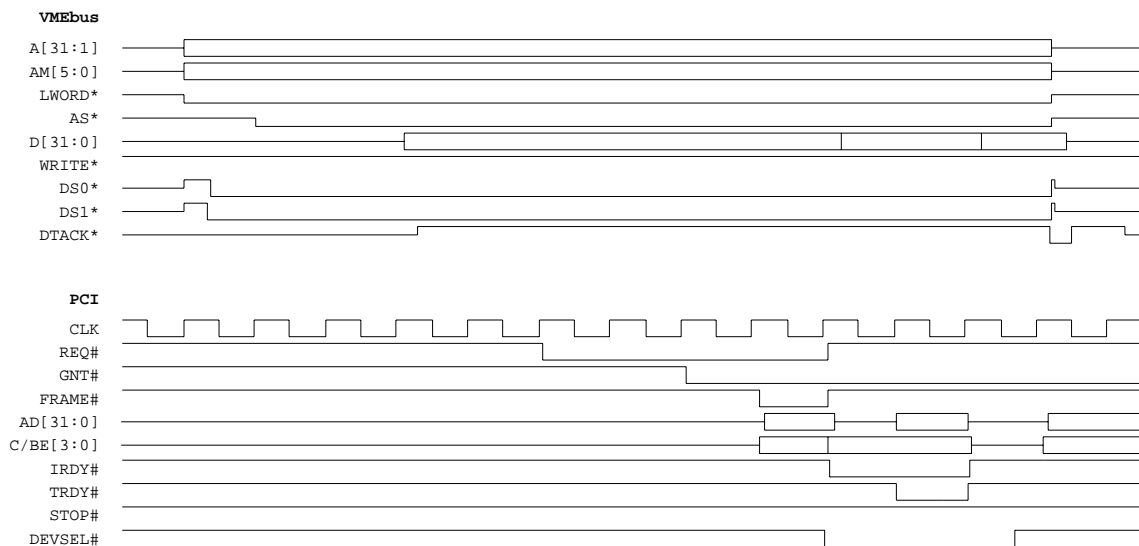


Figure B.11 : BLT Pre-fetched Read Cycle - Universe II as VME Slave

B.3 DMA Channel

B.3.1 Relative FIFO sizes

Two fixed “watermarks” in the DMA Channel control the Universe’s II requisition of the PCI bus and VMEbus. The DMAFIFO PCI Watermark is 128 bytes. This means that during reads from the PCI bus, the Universe II will wait for 128 bytes to be free in the DMAFIFO before requesting the PCI bus. For PCI writes, the Universe II waits for 128 bytes of data to be in the FIFO before requesting the PCI bus. The DMAFIFO VMEbus watermark is 64 bytes. This means that during reads from the VMEbus, the Universe II will wait for 64 bytes to be free in the DMAFIFO before requesting the Vmebus. For VMEbus writes, the Universe II waits for 64 bytes of data to be in the FIFO before requesting the VMEbus.

These watermarks have been tailored for the relative speeds of each bus, and provide near optimal use of the DMA channel.

B.3.2 VMEbus Ownership Modes

The DMA has two counters that control its access to the VMEbus: the VON (VMEbus On) counter and the VOFF (VMEbus Off) timer. The VON counter controls the number of bytes that are transferred by the DMA during any VMEbus tenure, while the VOFF timer controls the period before the next request after a VON time-out.

While the bus is more optimally shared between various masters in the system, and average latency drops as the value programmed for the VON counter drops, the sustained performance of the DMA also drops. The DMA is typically limited by its performance on the VMEbus. As this drops off with greater re-arbitration cycles, the average VMEbus throughput will drop. Even if the Universe II is programmed for ROR mode, and no other channels or masters are requesting the bus, there will be a period of time during which the DMA will pause its transfers on the bus, due to the VON counter expiring.

An important point to consider when programming these timers is the more often the DMA relinquishes its ownership of the bus, the more frequently the PCI Slave Channel will have access to the VMEbus. If DMA tenure is too long, the TXFIFO may fill up causing any further accesses to the bus to be retried. In the same fashion, all coupled accesses will be retried while the DMA has tenure on the bus. This can significantly affect transfer latency and should be considered when calculating the overall system latency.

B.3.3 VME Transfers

On the VMEbus, the Universe II can perform D08 through D64 transactions in either block or non-block mode. The time to perform a single beat, however, is independent of the bus width

being used. Hence, a D08 transaction will transfer data at 25% the rate of a D32, which in turn is half that for D64.

There is a significant difference between the performance for block vs. non-block operations. Because of the extra addressing required for each data transfer in non-block operations, the DMA performance is about half that compared to operating in block mode. Moreover, considering that most VME slaves respond less quickly in non-block mode, the overall performance may drop to one-quarter of that achievable in block mode.

When programmed for Release-When-Done operation, the Universe II will perform an early release of BBSY* when the VON counter reaches its programmed limit. This gives other masters a chance to use the VMEbus (and possibly access the VME Slave Channel), but may decrease performance of the DMA Channel; this factor may also play in favor of the DMA Channel, by pausing the PCI Target Channel's use of the VMEbus.

B.3.3.1 Read Transfers

When performing non-block reads on the VMEbus, the Universe II cycle time (AS* to next AS*) is approximately 209ns, which translates to about 20 MB/s when performing D32 transfers. For block transfers the cycle time (DS* to next DS*) falls to about 156ns, or 25 MB/s for D32 transfers. For multiplexed block transfers (MBLTs) the cycle time remains the same, but because the data width doubles, the transfer rate increases to about 50MB/s.

B.3.3.2 Write Transfers

Non-block writes to the VMEbus occur at 180ns cycle time (AS* to next AS*), or 23MB/s during D32 transfers. Block writes, however, are significantly faster with a 116ns cycle time (DS* to next DS*), or 36 MB/s. Multiplexed block transfers have slightly longer cycle times at about 112ns (DS* to next DS*), or 62 MB/s with D64 MBLTs.

B.3.4 PCI Transfers

As a master on the PCI bus, the Universe II DMA follows the same general set of rules as the VME Slave channel does: it never inserts any wait states into the transfer (i.e., it never negates IRDY# until the transaction is complete) and will whenever possible, generate full aligned bursts as set in the PABS field of the MAST_CTL register.

Between transactions on the PCI bus, the Universe II DMA typically sits idle for 6 clocks. Hence, minimizing the number of idle periods and re-arbitration times by setting PABS to its maximum value of 128 bytes may increase the performance of the DMA on this bus. Higher PABS values imply that the Universe II will hold on to both the PCI bus and the VMEbus for longer periods of time. The reason that PABS also may impact on VMEbus tenure is that (in the case of PCI writes), the DMA FIFO is less likely to fill, and (in the case of PCI reads) the

DMA is less likely to go empty. However, given the relative speeds of the buses, and the relative watermarks, the effect of PABS on VMEbus utilization is not as significant as its effects on the PCI bus.

While higher values of PABS increase DMA throughput, they may increase system latency. That is, there will be a longer latency for other PCI transactions, including possible transactions coming through the VME Slave Channel (since the DMA channel will own the PCI bus for longer periods of time). Also, accesses between other PCI peripherals will, on average, have a longer wait before being allowed to perform their transactions. PCI latency must be traded off against possible DMA performance.

Although both read and write transactions occur on the PCI bus with zero wait states, there is a period of six PCI clocks during which the Universe II remains idle before re-requesting the bus for the next transaction. PCI bus parking may be used to eliminate the need for re-arbitration.

With PABS set for 32-byte transactions on a 32-bit PCI bus, this translates to a peak transfer rate of 97 MB/s for reads (including pre-fetching), 98 MB/s for writes, doubling to 194 and 196 for a 64-bit PCI bus. With PABS set for 64-byte transactions, the peak transfer rate increases to 118 MB/s for reads, 125 MB/s for writes on a 32-bit PCI bus—236 MB/s and 250 MB/s respectively for 64-bit PCI buses. The numbers for writes to PCI assume that data are read from VME using BLTs.




Figure B.12 : PCI Read Transactions During DMA Operation




Figure B.13 : Multiple PCI Read Transactions During DMA Operation

B.4 Summary

Table B.1 : PCI Slave Channel Performance

Cycle Type	Performance
Coupled Read - PCI target response	8 PCI clocks
Coupled Write - PCI target response	9 PCI clocks
Decoupled Write - non-block D32 - VME cycle time - sustained perf (32-byte PABS) - sustained perf (64-byte PABS) - D32 BLT - VME cycle time - sustained perf (32-byte PABS) - sustained perf (64-byte PABS) - D64 MBLT - VME cycle time - sustained perf (32-byte PABS) - sustained perf (64-byte PABS)	180 ns 23 MB/s 43 MB/s 119 ns 32 MB/s 35 MB/s 119 ns 53 MB/s 59 MB/s

Table B.2 : VME Slave Channel Performance

Cycle Type	Performance
	VME Slave Response (ns)
Coupled Read	
- non-block	301
- D32 BLT	293
- D64 BLT	322
Coupled Write	
- non-block	278
- D32 BLT	264
- D64 BLT	292
Pre-fetched Read	
- VME slave response (1st data beat)	293
- VME slave response (other data beats)	57
Decoupled Write	
- non-block slave response	127
- block slave response (1st data beat)	127
- block slave response (other data beats)	50

Table B.3 : DMA Channel Performance

Cycle Type	Performance MB/s
PCI Reads - 32-byte PABS - 64byte PABS	97 (194) ^a 118 (236)
PCI Writes - 32-byte PABS - 64byte PABS	98 (196) 125 (250)
VME Reads - non-block D32 - D32 BLT - D64 MBLT	18 22 45
VME Writes - non-block D32 - D32 BLT - D64 MBLT	22 32 65

a. 64-bit PCI performance in brackets.

Appendix C Typical Applications

Being a bridge between standard interfaces, the Universe II requires minimal external logic to interface to either the VMEbus or to the PCI bus. In most applications, only transceivers to buffer the Universe II from the VMEbus, plus some reset logic are all that is required. The following information should be used only as a guide in designing the Universe II into a PCI/VME application. Each application will have its own set-up requirements.

C.1 VME Interface

C.1.1 Transceivers

The Universe II has been designed such that it requires full buffering from VMEbus signals. Necessary drive current to the VMEbus is provided by the transceivers while at the same time isolating the Universe II from potentially noisy VMEbus backplanes. In particular, complete isolation of the Universe II from the VMEbus backplane allows use of ETL transceivers which provide high noise immunity as well as use in live insertion environments. The VME community has recently standardized "VME64 Extensions" (ANSI VITA 1.1) which among other new VME features, facilitates live insertion environments.

If neither live insertion nor noise immunity are a concern, those buffers that provide input only (U15 and U17 in Figure C.1, below) may be omitted. The daisy chain input signals, BGIN[3:0] and IACKIN, have Schmitt trigger inputs, which should rectify any minor noise on these signals. If considerable noise is expected, the designer may wish to put external filters on these signals. Bear in mind that any filtering done on these signals will detrimentally affect the propagation of bus grants down the daisy chain. Only extremely noisy systems or poorly designed backplanes should require these filters.

Figure C.1 shows one example of how to connect the Universe II to the VMEbus. The transceivers in this example were chosen to meet the following criteria:

- provide sufficient drive strength as required by the VME specification (see Table C.1 on page C-4),
- meet Universe II skew requirements, and
- minimize part counts.

U15 and U17 in Figure C.1 are optional devices. They will provide better noise immunity.

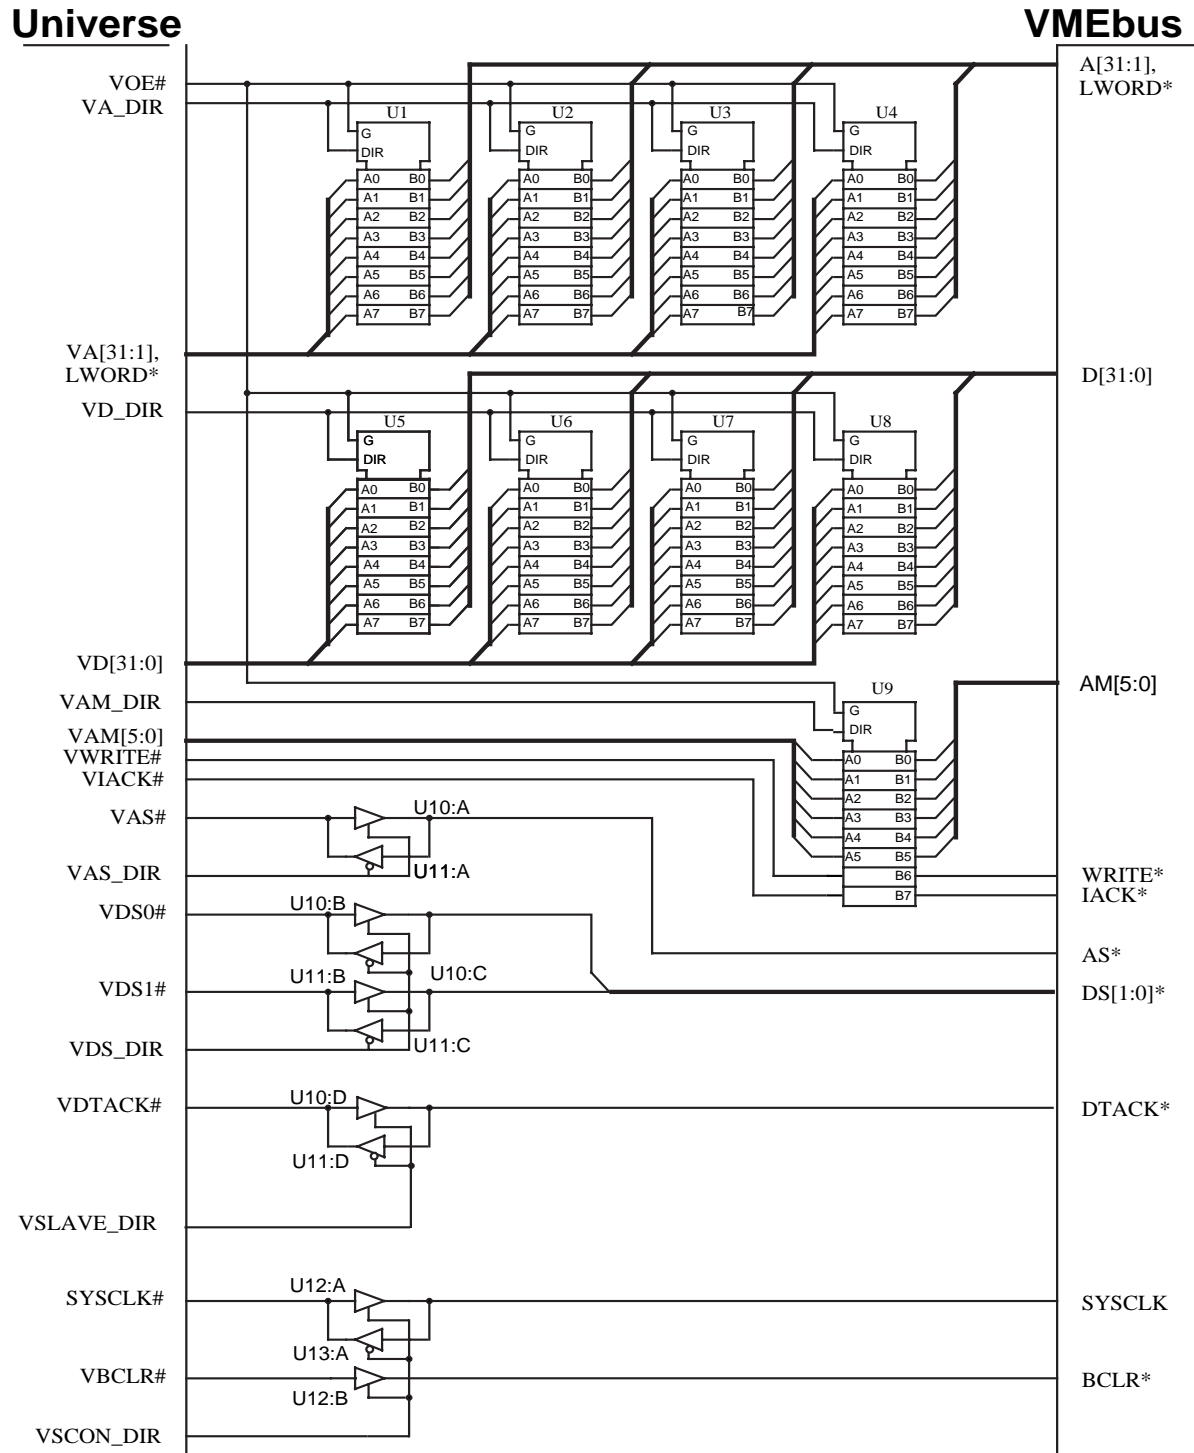


Figure C.1 : Universe II Connections to the VMEbus Through TTL Buffers

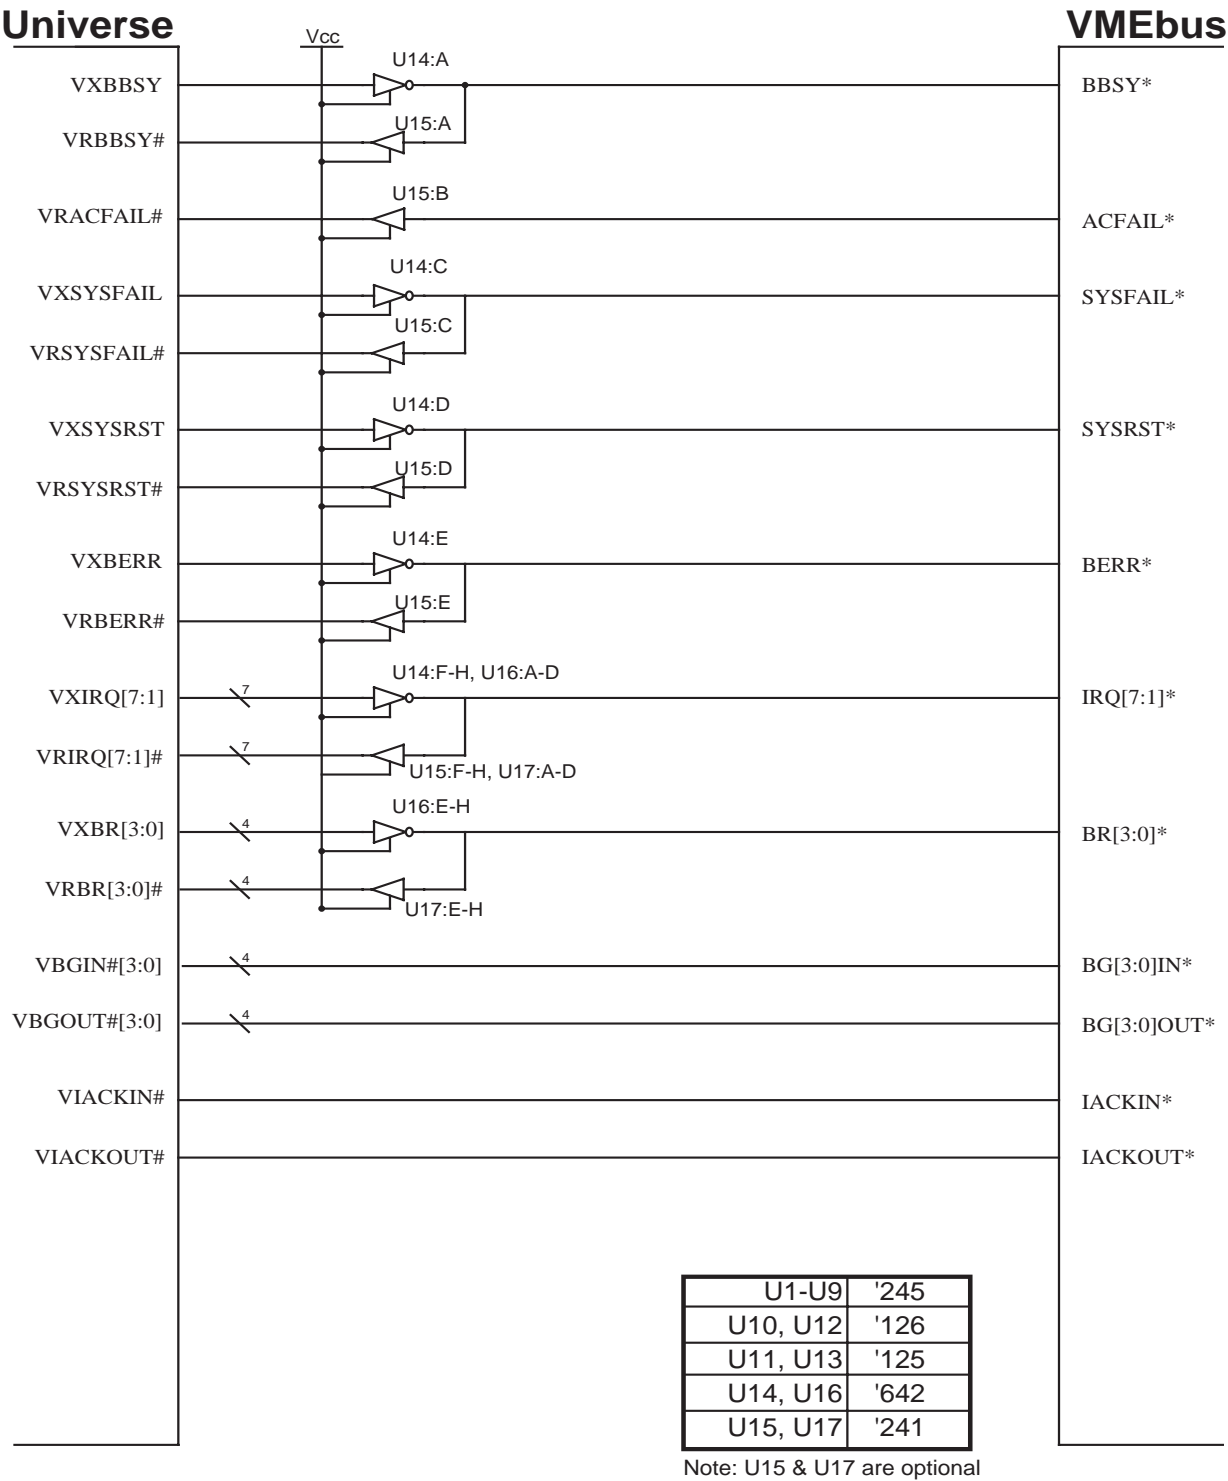


Figure C.1 (continued): Universe II Connections to the VMEbus Through TTL Buffers

The Universe II, with the addition of external transceivers, is designed to meet the timing requirements of the VME specification. Refer to the VME64 specification (ANSI VITA 1.0) for details on the VME timing. In order to meet the requirements outlined in this specification, the external transceivers must meet certain characteristics as outlined in Table C.2.

Table C.1 : VMEbus Signal Drive Strength Requirements

VME bus Signal	Required Drive Strength
A[31:1], D[31:0], AM[5:0], IACK*, LWORD*, WRITE*, DTACK*	IOL ≥ 48mA IOH ≥ 3mA
AS*, DS[1:0]*,	IOL ≥ 64mA IOH ≥ 3mA
SYSCLK*	IOL ≥ 64mA IOH ≥ 3mA
BR[3:0]*, BSY*, IRQ[7:0]*, BERR*, SYSFAIL*, SYSRESET*	IOL ≥ 48mA

Table C.2 : VMEbus Transceiver Requirements

Parameter	From (Input)	To (Output)	Timing	
			Min	Max
VA, VD, VAM, VIACK, VLWORD, VWRITE, VAS, VDSx, VDTACK ^a				
skew (pkg to pkg)	A	B		8 ns
skew (pkg to pkg)	B	A		4 ns
t _{Prop}	DIR	A	1 ns	5ns
t _{Prop}	DIR	B	2 ns	10ns
Cin	A			25 pf

a. There are no limits on propagation delay or skew on the remaining buffered VME signals: VSYSCLK, VBCLR, VXBBSY, VRBBSY, VRACFAIL, VXSYSFAIL, VRSYSFAIL, VXSYSRST, VRSYSRST, VXBERR, VRBERR, VXIRQ, VRIRQ, VXBR, VRBR.

F Series transceivers meet the requirements specified in Table C.1 and Table C.2. A faster family such as ABT, may also be used. Care should be taken in the choice of transceivers to avoid ground bounces and also to minimize crosstalk incurred during switching. To limit the effects of crosstalk, the amount of routing under these transceivers must be kept to a minimum. Daisy chain signals can be especially susceptible to crosstalk.

Should the designer wish to put any further circuitry between the Universe II and the VMEbus, that circuitry must meet the same timing requirements as the transceivers in order for the combined circuit to remain compliant with the VME64 specification.

C.1.1.1 Pull-down resistors

The Universe II has internal pull-down resistors which are used for its default power-up option state. The pins requiring pull-ups or pull-downs are indicated in Table 4.2 on page 4-3. (Note that REQ64# has an internal pull-up.) These internal pull-down resistors, ranging from 25k Ω -500k Ω , are designed to sink between 10 μ A-200 μ A. F-series buffers, however, can source up to 650 μ A of current (worst case). This sourced current has the ability to override the internal power up resistors on the Universe II. This may cause the Universe II to incorrectly sample a logic "1" on the pins. To counteract this potential problem, assuming a worst case scenario of a 650 μ A current, Tundra recommends connecting a 1K resistor to ground, in parallel, with the internal pull-down resistor.

Tundra recommends that any pins controlling the power-up options which are critical to the application at powerup be connected to ground with a pull-down resistor as described above. If these options are not critical and if it is possible to reprogram these options after reset, additional resistors need need not be added.

C.1.2 Direction control

When the Universe II is driving VMEbus lines, it drives the direction control signals high (i.e., VA_DIR, VAM_DIR, VAS_DIR, VD_DIR, VDS_DIR, VSLAVE_DIR, and VSCON_DIR). When the VMEbus is driving the Universe II, these signals are driven low. The control signals in the Universe II do not all have the same functionality. Since the Universe II implements early bus release, VAS_DIR must be a separate control signal.

Contention between the Universe II and the VME buffers is handled since the Universe II tristates its outputs one 64MHz clock period before the buffer direction control is faced inwards.

C.1.3 Power-up Options

Power-up options for the automatic configuration of slave images and other Universe II features are provided through the state of the VME address and data pins, VA[31:1] and VD[31:27]. All of these signals are provided with internal pull-downs to bias these signals to their default conditions. Should values other than the defaults be required here, either pull-ups or active circuitry may be applied to these signals to provide alternate configurations.

Power-up options are described in "Power-Up Options" on page 2-115.

Since the power-up configurations lie on pins that may be driven by the Universe II or by the VME transceivers, care must be taken to ensure that there is no conflict. During any reset event, the Universe II does not drive the VA or VD signals. As well, during any VMEbus reset (SYSRST*) and for several CLK64 periods after, the Universe II negates VOE# to tri-state the transceivers. During the period that these signals are tri-stated, the power-up options are loaded with their values latched on the rising edge of PWRRST#.

Configuration of power-up options is most easily accomplished through passive 10k pull-up resistors on the appropriate VA and VD pins. The configurations may be made user-configurable through jumpers or switches as shown in Figure C.2




Figure C.2 : Power-up Configuration Using Passive Pull-ups

Alternatively, an active circuit may be designed which drives the VA and VD pins with pre-set (or pre-programmed) values. This sort of circuit would be of value when power-up configurations such as the register access slave image are stored in an external programmable register. To implement this circuit, the VOE# output from the Universe II must be monitored. When the Universe II negates this signal, the appropriate VA and VD signals may be driven

and upon re-assertion the drive must be removed. To avoid conflict with the transceivers, logic must be designed such that the enabling of the transceivers does not occur until some point after the configuration options have been removed from the VD and VA signals. Figure C.3 shows one such implementation. The delay for enabling of the VMEbus transceivers could be implemented though clocked latches.

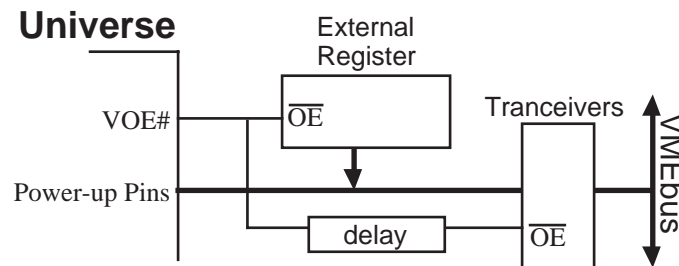


Figure C.3 : Power-up Configuration Using Active Circuitry

Auto-Syscon and PCI Bus Width Power-up Options

The VME64 specification provides for automatic enabling of the system controller in a VME system through monitoring of the BGIN3* signal. If at the end of SYSRST* this pin is low, then the system controller is enabled; otherwise it is disabled. The Universe II provides an internal pull-down resistor for this function. If it is in slot one, this pin will be sampled low. If not in slot one, then it will be driven high by the previous board in the system and system controller functions will be disabled. No external logic is required to implement this feature.

C.2 PCI Bus Interface

The Universe II provides a fully standard PCI bus interface compliant for both 32-bit and 64-bit designs. No external transceivers or glue logic is required in interfacing the Universe II to any other PCI compliant devices. All signals may be routed directly to those devices.

The Universe II's PCI interface can be used as a 32-bit bus or 64-bit bus. If used as a 32-bit interface, the 64-bit pins, AD[32:63] and ACK64# are left unterminated. On a 32-bit PCI bus, the Universe II drives all its 64-bit extension bi-direct signals (C/BE[7:4]#, AD[63:32], REQ64#, PAR64 and ACK64#) at all times to unknown values. Independent of the setting of the LD64EN bit, the Universe II will never attempt a 64-bit cycle on the PCI bus if it is powered up as 32-bit.

REQ64# must be pulled-down (with a 4.7k Ω resistor) at reset for 64-bit PCI (see “PCI Bus Width” on page 2-119). There is an internal pull-up on this pin which causes the Universe II to default to 32-bit PCI. This power-up option provides the necessary information to the Universe II so that these unused pins may be left unterminated.

C.2.1 Resets

The Universe II provides several reset input and outputs which are asserted under various conditions. These can be grouped into three types as shown in Table C.3.

Table C.3 : Reset Signals

Group	Signal Name	Direction
VMEbus	VXSYSRST	output
	VRSYSRST#	input
PCI bus	LRST#	output
	RST#	input
	VME_RESET#	input
Power-up	PWRRST#	input

VMEbus Resets

The VMEbus resets are connected to the VMEbus as indicated in Figure C.1 on page C-2 through external buffers.

PCI bus Resets

Use of the PCI bus resets will be application dependent. The RST# input to the Universe II should typically be tied in some fashion to the PCI bus reset signal of the same name. This will ensure that all Universe II PCI related functions are reset together with the PCI bus.

The LRST# pin is a totem-pole output which is asserted due to any of the following initiators:

- PWRRST#,
- VRSYSRST#,
- local software reset (in the MISC_CTL register), or
- VME CSR reset (in the VCSR_SET register).

The designer may wish to disallow the Universe II from resetting the PCI bus in which case this output may be left unconnected. Otherwise LRST# should be grouped with other PCI reset generators to assert the RST# signal such that:

$RST\# = LRST\# \& \text{reset_source1} \& \text{reset_source2} \& \dots$

If the Universe II is the only initiator of PCI reset, LRST# may be directly connected to RST#. Assertion of VME_RESET causes the Universe II to assert VXSYSRST.



This signal must not be tied to the PCI RST# signal unless the Universe II LRST# output will not generate a PCI bus reset. Connecting both LRST# and VME_RESET# to RST# will cause a feedback loop on the reset circuitry forcing the entire system into a endless reset.

To reset the VMEbus through this signal it is recommended that it be asserted for several clock cycles, until the Universe II asserts RST#, and then released. This ensures a break is made in the feedback path.

Power-Up Reset

The PWRRST# input is used to provide reset to the Universe II until the power supply has reached a stable level. It should be held asserted for 100 milliseconds after power is stable. Typically this can be achieved through a resistor/capacitor combination although more accurate solutions using under voltage sensing circuits (e.g. MC34064) are often implemented. The power-up options are latched on the rising edge of PWRRST#.

JTAG Reset

The JTAG reset, TRST#, should be tied into the master system JTAG controller. It resets the Universe II internal JTAG controller. If JTAG is not being used, this pin should be tied to ground.

C.2.2 Local Interrupts

The Universe II provides eight local bus interrupts, only one of which has drive strength that is fully PCI compliant. If any of the other seven interrupts are to be used as interrupt outputs to the local bus (all eight may be defined as either input or output), an analysis must be done on the design to determine whether the 4 mA of drive that the Universe II provides on these lines is sufficient for the design. If more drive is required, the lines may simply be buffered.

All Universe II interrupts are initially defined as inputs. To prevent excess power dissipation, any interrupts defined as inputs should always be driven to either high or low. Pull-ups should be used for this purpose rather than direct drive since a mis-programming of the interrupt registers may cause the local interrupts to be configured as outputs and potentially damage the device.

C.3 Manufacturing Test Pins

The Universe II has several signals used for manufacturing test purposes. They are listed in Table 2.24 on page 2-121, along with the source to which they should be tied.

C.4 Decoupling V_{DD} and V_{SS} on the Universe II

This section is intended to be a guide for decoupling the power and ground pins on the Universe II. A separate analog power and ground plane is not required to provide power to the analog portion of the Universe II. However, to ensure a jitter free PLL operation, the analog AV_{DD} and AV_{SS} pins must be noise free. The following are recommended solutions for noise free PLL operation. The design could implement one of these solutions, but not both.

The Analog Isolation Scheme consists of the following:

- a $0.1\mu\text{F}$ capacitor between the AV_{DD} and AV_{SS} pins, and
- corresponding inductors between the pins and the board power and ground planes (See Figure C.4). These inductors are not necessary, but they are recommended.




Figure C.4 : Analog Isolation Scheme

The Noise Filter Scheme filters out the noise using two capacitors to filter high and low frequencies (See Figure C.5).




Figure C.5 : Noise Filter Scheme

For both schemes, it is recommended that the components involved be tied as close as possible to the associated analog pins.

In addition to the decoupling schemes shown above, it is recommended that $0.1\mu\text{F}$ bypass capacitors should be tied between every three pairs of V_{DD} pins and the board ground plane. These bypass capacitors should also be tied as close as possible to the package.

Appendix D

Reliability Prediction

This section is designed to help the user to estimate the inherent reliability of the Universe II, as based on the requirements of MIL HDBK217F. This information is recommended for personnel who are familiar with the methods and limitations contained in MIL HDBK217F. The information serves as a guide only; meaningful results will be obtained only through careful consideration of the device, its operating environment, and its application.

D.1 Physical characteristics

- CMOS gate array
- 120,000 two-input nand gate equivalence
- 0.65 μm feature size
- 476 mils x 476 mils scribed die size

D.2 Thermal characteristics

Idle power consumption:	1.50 Watts
Typical power consumption* (32-bit PCI):	2.00 Watts
Maximum power consumption (32-bit PCI):	2.70 Watts
Typical power consumption (64-bit PCI):	2.20 Watts
Maximum power consumption* (64-bit PCI):	3.20 Watts

*NOTE: Maximum power consumption is worst case consumption when the Universe II is performing DMA reads from the VME bus with alternating worst case data patterns (\$FFFF_FFFF, \$0000_0000 on consecutive cycles), and 100pF loading on the PCI bus.

In the majority of system applications, the Universe II will consume typical values or less. Typical power consumption numbers are based on the Universe II remaining idle 30%-50% of the time, which is significantly less than what is considered likely in most systems. For this reason, it is recommended that typical power consumption numbers be used for power estimation and ambient temperature calculations, as described below.

Reliability calculations of the Universe II design in Motorola's H4EPlus Gate Array family show that the Failure In Time (FIT) rate is 68 at a junction temperature of 125°C (maximum junction temperature). (Failure in time is the basic reliability rate expressed as failures per billion (1e-9) device hours. Mean Time Between Failures (MTBF) is the reciprocal of FIT. MTBF is the predicted number of device hours before a failure will occur.)

D.3 Universe II Ambient Operating Calculations

The maximum ambient temperature of the Universe II can be calculated as follows:

$$T_a \leq T_j - \theta_{ja} * P$$

Where,

T_a = Ambient temperature (°C)

T_j = Maximum Universe II Junction Temperature (°C)

θ_{ja} = Ambient to Junction Thermal Impedance (°C / Watt)

P = Universe II power consumption (Watts)

The ambient to junction thermal impedance (θ_{ja}) is dependent on the air flow in linear feet per minute over the Universe II. The values for θ_{ja} over different values of air flow are as follows:

Table D.1 : Ambient to Junction Thermal Impedance

Air Flow (LFPM)	0	100	300
313 PBGA	20.10	17.0	15.1
324 CBGA	17.80	15.4	13.5

For example, the maximum ambient temperature of the 313 PBGA, 32-bit PCI environment with 100 LFPM blowing past the Universe II is:

$$T_a \leq T_j - \theta_{ja} * P$$

$$T_a \leq 125 - 17.0 * 2.70$$

$$T_a \leq 79.1 \text{ } ^\circ\text{C}$$

Hence the maximum rated ambient temperature for the Universe II in this environment is 79.1°C. The thermal impedance can be improved by approximately 10% by adding thermal conductive tape to the top of the packages and through accounting for heat dissipation into the ground planes. This would improve the maximum ambient temperature to 87°C in the above example. Further improvements can be made by adding heat sinks to the PBGA package.

T_j values of Universe II are calculated as follows ($T_j = \theta_{ja} * P + T_a$)

Table D.2 : Maximum Universe II Junction Temperature

Extended (125 C Ambient)	Industrial (85 C Ambient)	Commercial: (70 C Ambient)
$T_j = 17.0 * 2.70 + 125 \text{ C}$ $= 170.9 \text{ C}$	$T_j = 17.0 * 2.70 + 85 \text{ C}$ $= 130.9 \text{ C}$	$T_j = 17.0 * 2.70 + 70 \text{ C}$ $= 115.9 \text{ C}$

D.4 Thermal vias

The 313-pin plastic BGA package contains thermal vias which directly pipe heat from the die to the solder balls on the underside of the package. The solder balls use the capabilities of the power and ground planes of the printed circuit board to draw heat out of the package.

Appendix E Cycle Mapping

The Universe II always performs Address Invariant translation between the PCI and VMEbus ports. Address Invariant mapping preserves the byte ordering of a data structure in a little-endian memory map and a big-endian memory map.

E.1 Little-endian Mode

Table E.1 below shows the byte lane swapping and address translation between a 32-bit little-endian PCI bus and the VMEbus for the address invariant translation scheme.

Table E.1 : Mapping of 32-bit Little-Endian PCI Bus to 32-bit VMEbus

PCI Bus						Byte Lane Mapping	VMEbus			
Byte Enables				Address			DS1	DS0	A1	LW
3	2	1	0	1	0					
1	1	1	0	0	0	D0-D7 <-> D8-D15	0	1	0	1
1	1	0	1	0	1	D8-D15 <-> D0-D7	1	0	0	1
1	0	1	1	1	0	D16-D23 <-> D8-D15	0	1	1	1
0	1	1	1	1	1	D24-D31 <-> D0-D7	1	0	1	1
1	1	0	0	0	0	D0-D7 <-> D8-D15	0	0	0	1
						D8-D15 <-> D0-D7				
1	0	0	1	0	1	D8-D15 <-> D16-D23	0	0	1	0
						D16-D23 <-> D8-D15				
0	0	1	1	1	0	D16-D23 <-> D8-D15	0	0	1	1
						D24-D31 <-> D0-D7				
1	0	0	0	0	0	D0-D7 <-> D24-D31	1	0	0	0
						D8-D15 <-> D16-D23				
						D16-D23 <-> D8-D15				
0	0	0	1	0	1	D8-D15 <-> D16-D23	0	1	0	0
						D16-D23 <-> D8-D15				
						D24-D31 <-> D0-D7				

Table E.1 : Mapping of 32-bit Little-Endian PCI Bus to 32-bit VMEbus (Continued)

PCI Bus						Byte Lane Mapping	VMEbus			
Byte Enables				Address			DS1	DS0	A1	LW
3	2	1	0	1	0					
0	0	0	0	0	0	D0-D7 <-> D24-D31	0	0	0	0
						D8-D15 <-> D16-D23				
						D16-D23 <-> D8-D15				
						D24-D31 <-> D0-D7				

The unpacking of multiplexed 64-bit data from the VMEbus into two 32-bit quantities on a little-endian PCI bus is outlined in Table E.2 below.

Table E.2 : Mapping of 32-bit Little-Endian PCI Bus to 64-bit VMEbus

Byte Enables				Address			PCI to VME Byte Lane Mapping
3	2	1	0	2	1	0	
First Transfer (D32-D63)							
0	0	0	0	0	0	0	D0-D7 <-> A24-A31 (D56-D63)
							D8-D15 <-> A16-A23 (D48-D55)
							D16-D23 <-> A8-A15 (D40-D47)
							D24-D31 <-> LWORD, A1-A7 (D32-D39)
Second Transfer (D0-D31)							
0	0	0	0	1	0	0	D0-D7 <-> D24-D31
							D8-D15 <-> D16-D23
							D16-D23 <-> D8-D15
							D24-D31 <-> D0-D7

Appendix F

Operating and Storage Conditions

Table F.1 : Recommended Operating Conditions

DC Supply Voltage (V_{DD})	5 V
Ambient Operating Temperature (T_A Commercial)	0°C to +70°C
Ambient Operating Temperature (T_A Industrial)	-40°C to +85°C
Ambient Operating Temperature (T_A Extended)	-55°C to +125°C

Table F.2 : Absolute Maximum Ratings

DC Supply Voltage (V_{SS} to V_{DD})	-0.5 to 6.0 V
Input Voltage (V_{IN})	-0.5 to $V_{DD}+0.5$ V
DC Current Drain per Pin, Any Single Input or Output	± 50 mA
DC Current Drain per Pin, Any Paralleled Outputs	± 100 mA
DC Current Drain V_{DD} and V_{SS} Pins	± 75 mA
Storage Temperature, (T_{STG})	0°C to 70°C



WARNING: Stresses beyond those listed above may cause permanent damage to the devices. These are stress ratings only, and functional operation of the devices at these or any other conditions beyond those indicated in the operational sections of this document is not implied. Exposure to maximum rating conditions for extended periods may affect device reliability.

Table F.3 : Power Dissipation

IDLE	
Power Dissipation (32-bit PCI)	1.97W
Power Dissipation (64-bit PCI)	2.12W
Typical	
Power Dissipation (32-bit PCI)	2.65W
Power Dissipation (64-bit PCI)	3.15W

Appendix G

Mechanical and
Ordering Information

G.1 Mechanical Information






Figure G.1 : Mechanical Dimensions for the 324-Pin Ceramic BGA Package

Note that dimensions and tolerancing for all mechanical drawings follow ANSI Y14.5M, 1982.

The 313-PBGA ships in two varieties: OMPAC and GTPAC. Tundra cannot predict or guarantee which variety of the 313-PBGA we will ship to a given customer. The differences between these packages are minor and should not impact on your board layout. This manual includes mechanicals for both varieties of 313-PBGA. The bottom and side views of these packages are identical. However, the top views of the OMPAC and GTPAC differ slightly. Figure G.2 presents a single mechanical diagram for the 313-PBGA, which covers both types of package. It gives the top, side, and bottom view of these packages. Figure G.2 should be used for board layout.

Figure G.3 presents the specific top views of the 313 OMPAC and the 313 GTPAC. This information is provided for customer convenience only. Figure G.2 is the figure that should be used for board layout.

**Figure G.2 : 313-PBGA (Generic)**



NOTES:

1. Zone T delineates area of flatness for the GTPAC package.

Figure G.3: 313-PBGA Top View (OMPAC and GTPAC Drawings)

G.2 Ordering Information

Tundra Semiconductor Corporation products are designated by a Product Code. When ordering, refer to products by their full code. For detailed mechanical drawings or alternative packaging requirements, please contact our factory directly

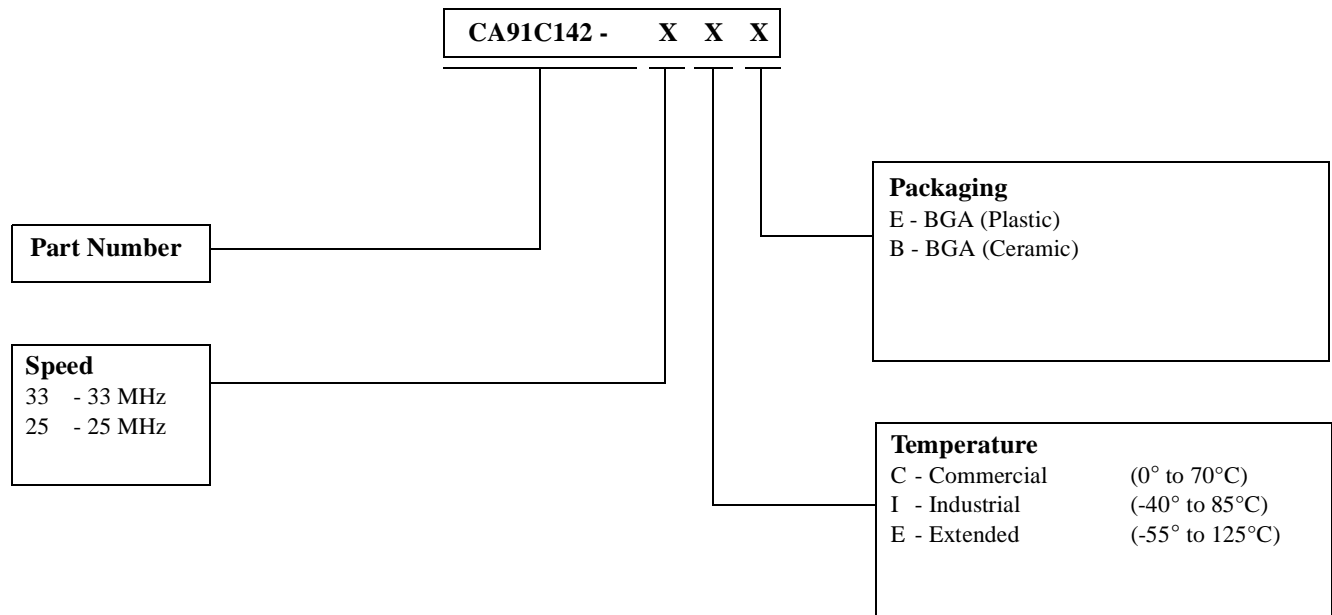


Table G.1 below details the available part numbers.

Table G.1 : Standard Ordering Information

Part Number	PCI Frequency	Temperature	Package
CA91C142-33CE	33	0° to 70°C	BGA (Plastic)
CA91C142-33IE	33	-40° to 85°C	BGA (Plastic)
CA91C142-25EE	25	-55° to 125°C	BGA (Plastic)
CA91C142-33CB	33	0° to 70°C	BGA (Ceramic)
CA91C142-33IB	33	-40° to 85°C	BGA (Ceramic)
CA91C142-25EB	25	-55° to 125°C	BGA (Ceramic)

Index

A

Absolute Maximum Ratings App F-1
ACFAIL*
 interrupt source 2-63
 interrupts 2-62
ACK64# 2-18, 2-31, 2-33, 2-39, 2-52, 2-119,
 3-5, 3-6, 4-3, App C-7
 Special Cycle Generator 2-45
AD 3-5, 4-3, 4-7
 and Configuration Cycles 2-21, 2-22
 parity checking 2-34
 Special PCI target image 2-55
 Target-Disconnects 2-49
Address Translation App E-1
 PCI to VME 2-54
 VME to PCI 2-52
Addressing Capabilities
 PCI Master Interface 2-36
 VMEbus Master Interface 2-9
ADOH Cycles 2-18, 2-45, 2-47
 DMA Channel 2-96
 generating 2-47
 Target-Retry 2-103
Ambient Operating Temperature App F-1
Architectural Diagram 2-3
AS* 2-19
Auto Slot ID
 proprietary method 2-25
 VME64 specified 2-24
A40 2-2

B

BBSY* 2-8
 ADOH cycles 2-18
 coupled-cycles 2-43
 DMA VMEbus ownership 2-81
BCLR* 2-9
BERR 2-13
BERR* 2-13, 2-15, 2-37
 coupled cycles 2-58
BGIN App C-1
BGIN3* App C-7
BG3IN* 2-115
 and First Slot Detector 2-23
BI-Mode 2-28
Block Diagram 2-3
BR3-0* 2-7
BSDL 2-123
Bus Errors
 DMA controller 2-83, 2-97
 IACK cycle 2-70
 parity 2-34
 RXFIFO posted writes 2-16
 TXFIFO posted writes 2-49
Bus Ownership 2-92
Bus Parking 2-31
Byte Enables App E-1
Byte Lane Mapping App E-1
Byte Ordering App E-1

C

CBGA App G-1
Ceramic BGA Package App G-1

CLK64 2-26, 2-112, 2-123, 3-1, 4-3
Configuration Cycles 2-20, 2-32, 2-36, 2-78,
2-100, 2-101, 2-102, 2-103
Coupled Request Phase 2-42
Coupled Request Timer 2-8
Coupled Transactions
 error handling 2-58
 PCI Target Channel 2-42
 VMEbus Slave Channel 2-14
Coupled Wait Phase 2-43
Coupled Window Timer 2-8, 2-43
Cycle Mapping App E-1
Cycle Terminations
 PCI Master Interface 2-37
 PCI Target Channel 2-49
 VMEbus Master Interface 2-13
 VMEbus Slave Channel 2-15, 2-16
C/BE# 2-31, 2-34, 2-37, 2-60, 2-119, 3-5, 4-3

D

Data Packing/Unpacking
 PCI Target Channel 2-41
 VMEbus Master Interface 2-12
Data Transfer
 PCI Master Interface 2-36
 PCI Target Channel 2-39
 VMEbus Master Interface 2-10
 VMEbus Slave Channel 2-14, 2-15, 2-16,
 2-19
Data Width 2-79
DC Current Drain App F-1
DC Supply Voltage App F-1
DCPP Register App A-61
 DCPP field 2-80

DCTL Register App A-57
 LD64EN bit 2-80
 L2V bit 2-78
 PGM field 2-78
 SUPER field 2-78
 VAS field 2-78
 VCT bit 2-80
 VDW bit 2-80
Designer's Resource Center 1-5
DEVSEL# 2-32, 2-33, 2-38, 2-52, 2-61, 3-5, 4-3
DGCS Register App A-62
 ACT bit 2-79, 2-81, 2-85, 2-89, 2-97
 CHAIN bit 2-79
 DONE bit 2-81, 2-83, 2-85, 2-86, 2-90, 2-91,
 2-92, 2-97
 GO bit 2-80, 2-81, 2-82, 2-85, 2-89, 2-90,
 2-91, 2-92, 2-97
 HALT bit 2-82, 2-83, 2-90, 2-97
 INT_DONE bit 2-75
 INT_HALT bit 2-75
 INT_LERR bit 2-75, 2-83, 2-85, 2-88, 2-95,
 2-97, 2-98
 INT_M_ERR bit 2-75, 2-96
 INT_STOP bit 2-75, 2-83, 2-85, 2-88, 2-95
 INT_VERR bit 2-75, 2-83, 2-85, 2-88, 2-95,
 2-97, 2-98
 LERR bit 2-81, 2-83, 2-85, 2-97
 P_ERR bit 2-81, 2-83, 2-85, 2-97
 STOP bit 2-79, 2-81, 2-82, 2-83, 2-85, 2-86,
 2-90, 2-97
 STOP_REQ bit 2-82, 2-85, 2-86, 2-90
 VERR bit 2-85, 2-97
 VOFF field 2-81, 2-85, 2-88, 2-95
 VON field 2-81, 2-82, 2-85, 2-88, 2-93, 2-94
Direction Control App C-5
DLA Register 2-78, App A-59

DMA Channel
 PCI requests 2-35
 PCI to VME transfers 2-92
 VME tot PCI transfers 2-94
 VMEbus release 2-9
 VMEbus requests 2-7

DMA Completion 2-82

DMA Controller
 defined 2-5
 direct mode operation 2-83
 error handling 2-60, 2-96
 FIFO operation and bus ownership 2-92
 interrupts 2-95
 linked-list operation 2-86

DMA Interrupts 2-95

DMAFIFO 2-5
 packing 2-92, 2-94
 PCI bus watermark 2-95
 VMEbus watermark 2-94

DTACK 2-13

DTACK* 2-13, 2-17, 2-103, 4-5
 Location Monitors 2-20
 RXFIFO 2-16

DTBC Register App A-58
 DTBC field 2-78

DVA Register 2-78, App A-60, App A-87

DY4 Systems 2-24

D_LLUE Register App A-64
 UPDATE bit 2-91

D32 2-16, 2-79

E

Endian Modes App E-1

ENID 3-5, 4-3

Error Handling
 coupled transactions 2-58
 DMA controller 2-60
 parity 2-60
 posted writes 2-58
 prefetched reads 2-60

F

FIFOs
 DMAFIFO 2-92
 RDFIFO 2-17
 RXFIFO 2-15
 TXFIFO 2-44

First Slot Detector 2-23

FIT rate App D-1

FRAME# 2-18, 2-31, 2-33, 2-38, 2-43, 2-49,
 3-5, 4-3

G

GNT# 2-31, 3-5, 4-3, App B-12

H

High Impedance Mode 2-122

I

IACK 2-58

IACKIN App C-1

IACKIN* 2-25
 SYSCON 2-23

IACKOUT* 2-25

IACK* 2-13, 2-23
 error logging 2-58

IDSEL 2-32, 2-102, 3-5, 4-3

Input Voltage App F-1
Interrupt Acknowledge Cycles
 auto-ID 2-25
 bus errors 2-70
 STATUS/ID 2-70
Interrupt Channel
 VMEbus release 2-8
 VMEbus requests 2-6
Interrupt Generation
 PCI bus 2-63
 VMEbus 2-65
Interrupt Handling
 internal sources 2-71
 PCI bus 2-68
 VMEbus 2-68
Interrupter
 defined 2-4
IRDY# 2-33, 2-49, 3-5, 4-3
IRQ* 2-62, 2-65
 interrupt source 2-63
IRQ2*
 Auto Slot ID 2-24

J

JTAG 2-123, App C-9

L

LAERR Register App A-40
 LAERR field 2-16
LCLK 2-33, 2-123, 3-5, 4-3
LERR 2-37, 2-59
Linked-List Operation of DMA 2-86
LINT# 2-64, 2-69, 2-70, 2-74, 3-5, 4-3

LINT_EN Register App A-65
 ACFAIL bit 2-64, 2-65
 DMA bit 2-64
 interrupt sources 2-64
 LERR bit 2-37, 2-64
 LMn bit 2-64
 MBOXn bit 2-64
 SW_IACK bit 2-64, 2-73, 2-74
 SW_INT bit 2-64, 2-73
 SYSFAIL bit 2-64, 2-65
 VERR bit 2-13, 2-49, 2-64
 VIR1x bits 2-64
 VMEbus interrupt inputs 2-68
 VOWN bit 2-64
LINT_MAP0 Register App A-69
 interrupt sources 2-64
 VERR field 2-64
 VIRQ7-1 fields 2-64, 2-69, 2-70
 VMEbus interrupt handling 2-64, 2-69
 VMEbus interrupt inputs 2-64
 VMEbus ownership bit 2-64
 VOWN field 2-64
LINT_MAP1 Register App A-70
 ACFAIL field 2-64
 DMA field 2-64
 interrupt sources 2-64
 LERR field 2-64
 SW_IACK field 2-64, 2-74
 SW_INT field 2-64
 SYSFAIL field 2-64
LINT_MAP2 Register 2-64, App A-85
 LM3-0 fields 2-76
 MBOX3-0 fields 2-75

- LINT_STAT Register 2-59, App A-67
 - ACFAIL bit 2-64
 - DMA bit 2-64
 - interrupt sources 2-64
 - LERR bit 2-64
 - LMx bit 2-76
 - MBOXn bit 2-64, 2-75
 - Software interrupts 2-73
 - SW_INT bit 2-73
 - SYSFAIL bit 2-64
 - VERR bit 2-64
 - VMEbus interrupt handling 2-69
 - VMEbus interrupt inputs 2-64
 - VOWN bit 2-64
- Little-Endian Mode App E-1
- LMISC Register App A-36
 - CWT field 2-43
- LM_BS Register 2-20, 2-75, App A-116
- LM_CTL Register App A-115
 - EN bit 2-20
 - SUPER field 2-20
 - VAS field 2-20
- Location Monitors 2-19-2-20
 - and interrupts 2-20
 - interrupts 2-19, 2-75, 2-76
- Locks
 - VMEbus Slave Channel 2-19
- LOCK# 2-18, 2-19, 2-47, 2-51, 3-6, 4-3
- LRST# 2-110, 2-111, 2-112, 2-114, 3-6, 4-3, App C-8
- LSIn_BD Registers
 - BD field 2-53
 - Power-up options 2-116
- LSIn_BS Registers
 - BS field 2-53
 - Power-up options 2-116
- LSIn_CTL Registers
 - EN bit 2-53
 - LAS field 2-53
 - PGM field 2-53
 - Power-up options 2-116
 - PWEN bit 2-44, 2-53
 - SUPER field 2-53
 - VAS field 2-53
 - VCT field 2-53
 - VDW field 2-53
- LSIn_TO Registers
 - TO field 2-53
- LSI0_BD Register App A-17
- LSI0_BS Register App A-16
- LSI0_CTL Register App A-15
- LSI0_TO Register App A-18
- LSI1_BD Register App A-21
- LSI1_BS Register App A-20
- LSI1_CTL Register App A-19
- LSI1_TO Register App A-22
- LSI2_BD Register App A-25
- LSI2_BS Register App A-24
- LSI2_CTL Register App A-23
- LSI2_TO Register App A-26
- LSI3_BD Register App A-29
- LSI3_BS Register App A-28
- LSI3_CTL Register App A-27
- LSI3_TO Register App A-30
- LSI4_BD Register App A-43
- LSI4_BS Register App A-42
- LSI4_CTL Register App A-41
- LSI4_TO Register App A-44
- LSI5_BD Register App A-47
- LSI5_BS Register App A-46
- LSI5_CTL Register App A-45
- LSI5_TO Register App A-48
- LSI6_BD Register App A-51
- LSI6_BS Register App A-50

LSI6_CTL Register App A-49
LSI6_TO Register App A-52
LSI7_BD Register App A-55
LSI7_BS Register App A-54
LSI7_CTL Register App A-53
LSI7_TO Register App A-56
L_CMDERR Register App A-39
 CMDERR field 2-16, 2-37
 L_STAT bit 2-16, 2-37
 M_ERR bit 2-16, 2-37

M

Mailbox Registers 2-108
 Interrupts 2-65
Master-Abort
 defined 2-33
MAST_CTL Register App A-93
 BUS_NO field 2-20, 2-21, 2-22
 MAXRTRY field 2-37
 PABS field 2-16, 2-17, 2-35, 2-36, 2-60
 PWON field 2-8, 2-11
 VOWN bit 2-7, 2-8, 2-47, 2-48, 2-75
 VOWN_ACK 2-47
 VOWN_ACK bit 2-7, 2-48
 VREL bit 2-8, 2-48
 VRL field 2-7
 VRM field 2-7
MBOXn Registers 2-75, 2-108
MBOX0 Register App A-87
MBOX1 Register App A-88
MBOX2 Register App A-89
MBOX3 Register App A-90
MD32 2-2
Mechanical App G-1

MISC_CTL Register App A-95
 AUTOID bit 2-24
 BI bit 2-29, 2-116
 ENGBI bit 2-29
 SW_LRST bit 2-111, 2-112
 SW_SYSRST bit 2-112
 SYSCON bit 2-23, 2-26, 2-116
 VARB bit 2-27
 VARBTO field 2-27
 VBTO field 2-28
 V64AUTO bit 2-116
MISC_STAT Register App A-97
 DY4AUTO bit 2-116
 DY4AUTOID field 2-26
 DY4DONE bit 2-26
 LCLSIZE bit 2-116
 MYBBSY bit 2-8
 RXFE bit 2-36
 TXFE bit 2-8

Monarch 2-24
MTBF App D-1

N

NAND Tree Simulation 2-122
Noise Filter App C-10
Normal Mode 2-122

O

Operating Conditions App F-1
Ordering Information App G-5

P

Packaging App G-5
PAR 2-34, 2-37, 2-60, 3-6, 4-3

- Parity
 - error handling 2-60
 - PCI Master Interface 2-37
- Parity Checking
 - Universe capability 2-34
- Part Number App G-5
- PAR64 2-34, 2-37, 2-60, 3-6, 4-3
- PBGA App D-2
- PCI Aligned Burst Size (PABS) 2-17, 2-35, 2-36, 2-60, 2-92
- PCI Cycle Types
 - Universe II capability 2-36
- PCI Interface
 - cycle types 2-32
 - defined 2-4
 - Universe II as master 2-35
 - Universe II as slave 2-38
 - 32-bit versus 64-bit 2-30
- PCI Master Interface
 - cycle terminations 2-37
 - data transfer 2-36
 - parity 2-37
- PCI Requests
 - DMA Channel 2-35
 - VME Slave Channel 2-35
- PCI Slave Images
 - defined 2-53
- PCI Target Channel
 - ADOH cycles 2-47
 - coupled transactions 2-42
 - cycle terminations 2-49
 - data packing/unpacking 2-41
 - data transfer 2-39
 - posted writes 2-44
 - read-modify-writes 2-46
 - TXFIFO 2-44
 - VMEbus release 2-8
 - VMEbus requests 2-6
- PCI Target Image
 - Power-up Option 2-116
- PCI Terminations
 - defined 2-33
- PCI_BS_n Register
 - BS field 2-103
 - SPACE bit 2-103
- PCI_BS_n Registers
 - BS field 2-103
 - SPACE bit 2-103
- PCI_BS0 Register App A-12
 - SPACE bit 2-116
- PCI_BS1 Register App A-13
 - SPACE bit 2-116
- PCI_CLASS Register App A-10
- PCI_CSR Register App A-8
 - BM bit 2-14, 2-18, 2-51, 2-52, 2-81, 2-83
 - power-up option 2-116
 - DEVSEL field 2-38
 - DP_D bit 2-38
 - D_PE bit 2-38, 2-60
 - PERESP bit 2-38, 2-60
 - R_MA bit 2-15, 2-37
 - R_TA bit 2-15, 2-37
 - SERR_EN bit 2-39, 2-60
 - S_SERR bit 2-39
 - S_TA bit 2-49
- PCI_ID Register App A-7
- PCI_MISC0 Register App A-11
- PCI_MISC1 Register App A-14
- PERR# 2-38, 2-60, 2-61, 3-6, 4-3
- PLL App C-10
- PLL_TESTOUT 2-123, 3-6, 4-3
- PLL_TESTSEL 2-122, 3-6, 4-4
- Posted Writes
 - and coupled transfers 2-42
 - error handling 2-58
 - errors 2-16
 - PCI Target Channel 2-44
 - VMEbus Slave Channel 2-15
- power consumption App D-1
- Power Dissipation App F-2

- Power-up
 - register access 2-23
- Power-up Options 2-117, App C-5
 - auto-ID 2-118
 - BI-mode 2-118
 - PCI bus width 2-119
 - PCI CSR image space 2-119
 - PCI slave image 2-118
 - SYSFAIL assertion 2-118
 - VME CR/CSR slave image 2-117
 - VME register access slave image 2-117
- Prefetched Reads
 - error handling 2-60
 - VMEbus Slave Channel 2-16
- Product Code App G-5
- Pull-down resistors App C-5
- PWON App B-4
- PWRRST# 2-110, 2-112, 2-113, 2-114, 2-115, 2-116, 2-120, 2-122, 3-6, 4-4, App C-6, App C-8, App C-9

R

- RDFIFO 2-3, 2-17
 - size 2-17
- Read-Modify-Writes
 - PCI Target Channel 2-46
 - VMEbus Slave Channel 2-19
- Register Access
 - at power-up 2-23
 - configuration space 2-102
 - CR/CSR access 2-106
 - from VMEbus 2-104
 - I/O space 2-103
 - memory space 2-103
 - VMEbus register access image 2-104
- Register Map App A-2
- Registers 2-100-2-109
- Reliability App D-1

- Request Modes 2-7
- REQ# 2-31, 2-33, 3-6, 4-4, App B-12
- REQ64# 2-18, 2-31, 2-33, 2-39, 2-52, 2-115, 2-116, 2-119, 3-6, 4-4, App C-7
- Resets 2-110
- RETRY* 2-4, 2-13
- RST# 2-110, 2-112, 3-6, 4-4, App C-8, App C-9
 - BI-MODE 2-29
- RXFIFO 2-15

S

- SCV64 1-3, 2-24
- SCYC_ADDR Register App A-32
 - ADDR field 2-45
- SCYC_CMP Register App A-34
 - CMP field 2-45
- SCYC_CTL Register App A-31
 - SCYC field 2-45, 2-46
- SCYC_EN Register App A-33
 - EN field 2-45
- SCYC_SWP Register App A-35
 - SWP field 2-45
- SDONE 2-30
- SEMA_n Register 2-109
- Semaphores 2-91, 2-109
- SEMA0 Register 2-109, App A-91
- SEMA1 Register 2-109, App A-92
- SERR# 2-34, 2-38, 2-39, 2-60, 2-61, 3-6, 4-4
- Signal Conventions 1-6

SLSI Register App A-37
 BS field 2-56
 EN bit 2-56
 LAS field 2-56
 PGM field 2-56
 PWEN bit 2-56
 SUPER field 2-56
 VDW field 2-56
Special Cycle Generator 2-4, 2-45-2-48
 semaphores 2-109
Special PCI Slave Image
 defined 2-55
 memory mapping 2-57
STATID Register 2-24, 2-66, App A-77
 STATID field 2-67
STATUS/ID
 provided by Universe II 2-67
STOP# 2-33, 2-49, 3-7, 4-4
Storage Conditions App F-1
Storage Temperature App F-1
SYSCON 2-123
SYSFAIL* 2-68, 2-76, 2-118
 and auto ID cycle 2-24
 Auto ID cycle 2-24
 Auto Slot ID 2-24
 interrupt source 2-63
 interrupts 2-62
SYSRST* 2-24, 2-112, 2-115, 2-118, App C-6, App C-7
 Auto ID cycle 2-24
 BI-Mode 2-29

T

Target-Abort
 defined 2-33
Target-Disconnect
 defined 2-33

Target-Retry
 defined 2-33
TCK 2-123, 3-7, 4-4
TDI 2-123, 3-7, 4-4
TDO 2-123, 3-7, 4-4
Terminology 1-7
Thermal vias App D-3
Time-Outs
 VMEbus 2-28
 VMEbus arbiter 2-27
TMODE 2-122, 2-123, 3-7, 4-4
TMS 2-123, 3-7, 4-4
TRDY# 2-33, 2-49, 3-7, 4-4
TRST# 2-110, 2-123, 3-7, 4-4
TTL Buffers App C-2
Tundra web site 1-5
TXFIFO 2-44

U

USER_AM Register 2-55, App A-98
 USER1AM, USER2AM fields 2-10

V

VA 3-1, 4-4, App C-5
 and Configuration Cycles 2-21
 Configuration Cycles 2-21
 Location Monitor 2-76
 power-up options 2-116
VAERR Register App A-122
 VAERR field 2-13, 2-49
VAM 3-1, 4-4
VAM_DIR 3-1, 4-4, App C-5
VAS# 3-1, 4-4
VAS_DIR 3-2, 4-4, App C-5
VA_DIR 3-1, 4-4, App C-5

- VBCLR# 2-27, 3-2, 4-5
- VBGIN 2-116
- VBGI# 3-2, 4-5
- VBGO# 2-27, 3-2
- VBG0# 4-5
- VCOCTL 2-121, 2-123, 3-7, 4-5
- VCSR_BS Register App A-141
 - BS field 2-106
- VCSR_CLR Register 2-114, App A-139
 - FAIL bit 2-114
 - RESET bit 2-111
 - SYSFAIL bit 2-25, 2-111, 2-114, 2-116
- VCSR_CTL Register App A-119
 - EN bit 2-106
 - LAS field 2-106, 2-116
- VCSR_SET Register App A-140
 - FAIL bit 2-114
 - SYSFAIL bit 2-25, 2-114, 2-116
- VCSR_TO Register App A-120
 - TO field 2-106, 2-116
- VD 3-2, 4-5
 - power-up options 2-116
 - pull-ups 2-119
- VDS# 3-2, 4-5
- VDS_DIR 3-2, 4-5, App C-5
- VDTACK# 3-2, 4-5
- VD_DIR 3-2, 4-5, App C-5
- VERR 2-37, 2-59
- VIACKI# 4-5
- VIACKO# 3-3, 4-5
- VIACK# 3-3, 4-5
- VINT_EN Register App A-71
 - DMA bit 2-66
 - interrupt sources 2-66
 - LERR bit 2-16, 2-37, 2-66
 - MBOX3-0 bits 2-66
 - PCI interrupt inputs 2-68
 - SW_INT bit 2-66, 2-73
 - SW_INT7-1 bits 2-66, 2-73, 2-116
 - VERR bit 2-13, 2-49, 2-66
- VINT_MAP0 Register 2-66, App A-75
- VINT_MAP1 Register 2-66, App A-76
 - DMA field 2-66
 - interrupt sources 2-66
 - LERR field 2-66
 - SW_INT bit 2-116
 - SW_INT field 2-66, 2-73
 - VERR field 2-66
- VINT_MAP2 Register 2-66, App A-86
 - MBOX3-0 fields 2-75
- VINT_STAT Register 2-59, 2-66, 2-67, 2-68, App A-73
 - DMA bit 2-66, 2-96
 - IACK cycle error 2-70
 - interrupt sources 2-66
 - LERR bit 2-66, 2-70
 - LINT7-0 bitS 2-66
 - LINT7-0 bits 2-75
 - MBOX3-0 bitS 2-66
 - MBOX3-0 bits 2-75
 - PCI interrupt inputs 2-68
 - SW_INT bit 2-66, 2-73, 2-116
 - SW_INT7-1 bits 2-66
 - VERR bit 2-66, 2-70
 - VMEbus interrupt handling 2-69
- VLWORD# 3-3, 4-5
- VME Slave Images
 - defined 2-50

- VMEbus Arbitration 2-27
 - arbiter time-out 2-27
 - priority mode 2-27
 - round robin 2-27
 - single level 2-27
- VMEbus Interface
 - BI-mode 2-28
 - configuration 2-23
 - CR/CSR access 2-106
 - defined 2-2
 - first slot detector 2-23
 - requester 2-6
 - system clock 2-26
 - system controller 2-26
 - Universe as master 2-9
 - Universe as slave 2-13
 - VMEbus release 2-8
 - VMEbus time-out 2-28
- VMEbus Master Interface
 - addressing capabilities 2-9
 - cycle terminations 2-13
 - data packing/unpacking 2-12
 - data transfer 2-10
- VMEbus Register Access Image 2-104
- VMEbus Requester
 - demand mode 2-7
 - DMA Channel 2-7
 - fair mode 2-7
 - Interrupt Channel 2-6
 - PCI Target Channel 2-6
 - request levels 2-7
- VMEbus Slave Channel
 - coupled transactions 2-14
 - errors 2-16
 - locks 2-18
 - PCI requests 2-35
 - posted writes 2-15
 - prefetched reads 2-16
 - RDFIFO 2-17
 - read-modify-writes 2-19
 - RXFIFO 2-15
- VMEbus System Controller
 - IACK daisy chain 2-27
 - VMEbus arbitration 2-27
- VME_RESET App C-9
- VME_RESET# 2-110, 2-112, 2-114, 3-7, 4-5, App C-8, App C-9
- Vn_STATID Registers 2-69
 - ERR bit 2-70
 - STATID field 2-69
- VOE# 2-119, 2-120, 3-3, 4-5, App C-6
- VOFF timer 2-82, 2-93, 2-95
- Voltage App F-1
- VON 2-81
- VRACFAIL# 3-3, 4-5
- VRAI_BS Register App A-118
 - BS field 2-23, 2-104, 2-116
- VRAI_CTL Register App A-117
 - EN bit 2-104, 2-116
 - PGM field 2-104
 - SUPER field 2-104
 - VAS field 2-104
- VRBBSY# 2-27, 3-3, 4-5
- VRBERR# 3-3, 4-5
- VRBR# 2-27, 3-3, 4-6
- VRIRQ 2-29, 2-63
- VRIRQ# 3-3, 4-6
- VRSYSFAIL# 3-4, 4-6
- VRSYSRST# 2-110, 2-112, 3-4, 4-6, App C-8
- VSCON_DIR 3-4, 4-6, App C-5
- VSI_n_BD Registers
 - BD field 2-50
- VSI_n_BS Registers
 - BS field 2-50

- VSI_n_CTL Registers
 - and Type 0 configuration cycles 2-22
 - EN bit 2-50
 - LAS field 2-50
 - LD64EN bit 2-17, 2-50
 - LLRMW bit 2-19, 2-50
 - PGM field 2-50
 - PREN bit 2-16, 2-50
 - PWEN bit 2-15, 2-50
 - SUPER field 2-50
 - Type 0 configuration cycles 2-21
 - VAS field 2-50
- VSI_n_TO Registers
 - TO field 2-50
- VSI0_BD Register App A-101
- VSI0_BS Register App A-100
- VSI0_CTL Register App A-99
- VSI0_TO Register App A-102
- VSI1_BD Register App A-105
- VSI1_BS Register App A-104
- VSI1_CTL Register App A-103
- VSI1_TO Register App A-106
- VSI2_BD Register App A-109
- VSI2_BS Register App A-108
- VSI2_CTL Register App A-107
- VSI2_TO Register App A-110
- VSI3_BD Register App A-113
- VSI3_BS Register App A-112
- VSI3_CTL Register App A-111
- VSI3_TO Register App A-114
- VSI4_BD Register App A-125
- VSI4_BS Register App A-124
- VSI4_CTL Register App A-123
- VSI4_TO Register App A-126
- VSI5_BD Register App A-129
- VSI5_BS Register App A-128
- VSI5_CTL Register App A-127
- VSI5_TO Register App A-130
- VSI6_BD Register App A-133
- VSI6_BS Register App A-132
- VSI6_CTL Register App A-131
- VSI6_TO Register App A-134
- VSI7_BD Register App A-137
- VSI7_BS Register App A-136
- VSI7_CTL Register App A-135
- VSI7_TO Register App A-138
- VSLAVE_DIR 3-4, 4-6, App C-5
- VSYSCLK 2-123, 3-4, 4-6
- VWRITE# 3-4, 4-6
- VXBBSY 3-4, 4-6
- VXBERR 3-4, 4-6
- VXBERR# 2-28
- VXBR 3-4, 4-6
- VXIRQ 3-4, 4-7
- VXSYSFAIL 3-4, 4-7
- VXSYSRST 3-4, 4-7, App C-8, App C-9
- V_AMERR Register 2-58, App A-121
 - AMERR field 2-13, 2-49
 - IACK bit 2-13, 2-70
 - M_ERR bit 2-13, 2-49
 - V_STAT bit 2-13, 2-49
- V1_STATID Register App A-78
- V2_STATID Register App A-79
- V3_STATID Register App A-80
- V4_STATID Register App A-81
- V5_STATID Register App A-82
- V6_STATID Register App A-83
- V7_STATID Register App A-84

Worldwide Sales Network

USA Sales Representatives

Eastern

D.C., Maryland, Virginia

Chesapeake Technologies

3905 National Dr., Suite 425
Burtonsville, MD 20866
Tel: 301-236-0530
Fax: 301-384-9596

Chesapeake Technologies

12-616E Hampton Dr.
Midlothian, VA 23113
Tel: 804-379-1816
Fax: 804-379-3474

*Connecticut, Maine, Massachusetts,
New Hampshire, New Jersey, New
York, Rhode Island, Vermont, Eastern
Pennsylvania*

Daner Hayes Inc.

300 Eliot St., Suite 342
Ashland, MA 01721
Tel: 508-881-0400
Fax: 508-881-8227

Upstate New York, Delaware, Ohio

Supertech

103-A Kingsberry Dr.
Rochester, NY 14626
Tel: 716-720-0915
Fax: 716-720-0916

EMA, Inc.

Aurora Office Plaza, Suite 291
251 West Garfield Rd.
Aurora, OH 44202
Tel: 216-562-6104
Fax: 216-562-7498

EMA, Inc.

35 Compark Road,
Centerville, OH 45459
Tel: 937-433-2800
Fax: 937-433-3147

Florida

Apollo Technical Sales Co.

1275S Patrick Dr., #M2
Satellite Beach, FL 32937-3963
Tel: 407-777-7511
Fax: 407-777-5251

Central

Illinois, Michigan, Wisconsin

Horizon Technical Sales

4340 Cross Street,
Downer's Grove, IL 60515
Tel: 630-852-2500
Fax: 630-852-2520

Horizon Technical Sales

6011 Hilldale Dr.
Hartford, WI 53027-9541
Tel: 414-670-6776
Fax: 414-670-6778

*Alabama, Georgia, Mississippi,
North Carolina, South Carolina,
Tennessee*

BITS, Inc.

2707 Artie St., Suite 10
Huntsville, AL 35805
Tel: 205-534-4020
Fax: 205-534-0410

BITS, Inc.

1 Meca Way
Norcross, GA 30093
Tel: 770-564-5599
Fax: 770-564-5588

BITS, Inc.

7706 Six Forks Rd., Suite 101
Raleigh, NC 27615
Tel: 919-676-1880
Fax: 919-676-1881

BITS, Inc.

2305 Keaton Ave., Suite E
Charlotte, NC 28269
Tel: 704-510-1064
Fax: 704-510-1065

*Arkansas, Louisiana, Oklahoma,
Texas*

Quad State Sales and Marketing

12160 Abrams Rd., Suite 406
Dallas, TX 75243
Tel: 972-669-8567
Fax: 972-669-8834

Quad State Sales and Marketing

8310 Capital of Texas Highway
North, Suite 365
Austin, TX 78731
Tel: 512-346-7002
Fax: 512-346-3601

Quad State Sales and Marketing

10565 Katy Freeway, Suite 212
Houston, TX 77024
Tel: 713-467-7749
Fax: 713-467-5942

Quad State Sales and Marketing

110 West Commercial, Suite 210
Broken Arrow, OK 74013
Tel: 918-258-7723
Fax: 918-258-7653

Minnesota, North Dakota, South
Dakota, Western Wisconsin

Electromec Sales Inc.

1601 East Highway 13, Suite 102
Burnsville, MN 55337
Tel: 612-894-8200
Fax: 612-894-9352

Missouri, Southern Illinois, Iowa,
Kansas, Nebraska

Please contact Tundra Semiconductor
directly at: 1-800-267-7231

Mountain

Arizona, New Mexico

Please contact Tundra Semiconductor
directly at: 1-800-267-7231

USA Sales Representatives Cont'd

*Colorado, Idaho, Montana, Utah,
Wyoming*

Intermountain Component Sales

12665 South Minuteman Dr.
Drasper, UT 84020
Tel: 801-572-4010
Fax: 801-572-8180

Intermountain Component Sales

2801 Youngfield St., Suite 300
Golden, CO 80401
Tel: 303-275-7150
Fax: 303-275-7142

Western

*Northern California,
Northern Nevada*

Quorum Technical Sales

4701 Patrick Henry Dr.
Santa Clara, CA 95054
Tel: 408-980-0812
Fax: 408-748-1163

*Southern California,
Southern Nevada*

Please contact Tundra Semiconductor
directly at: 1-800-267-7231

Oregon, Wahington State

Electronic Sources, Inc.

6865 S.W. 105th, Suite B
Beaverton, OR
97008

Tel: 503-627-0838
Fax: 503-627-0238

Electronic Sources, Inc.

1603 116th Ave. NE, Suite 115,
Bellevue, WA 98004
Tel: 425-451-3500
Fax: 425-451-1038

International Sales Representatives

Australia

ACD

Unit 2, 17-19 Melrich Rd.
Bayswater, Victoria
Australia 3153

Tel: 61-3-9762-7644
Fax: 61-3-9762-5446

ACD

106 Belmore Rd.
North Riverwood, NSW
Australia 2210

Tel: 61-2-534-6200
Fax: 61-2-534-4910

ACD

20D William Street,
Norwood, South Australia
Australia

Tel: 61-8-364-2844
Fax: 61-8-364-2811

ACD

Unit 1, Technology Park
Bently, West Australia
Australia

Tel: 61-9-472-3232
Fax: 61-9-470-2303

Belgium, The Netherlands

Microtron

Generaal Dewittelaan 7
2800 Mechelen
Belgium

Tel: 32-15-292929
Fax: 32-15-292900

Canada

Hewetson Repping, Inc.

161 Herzberg Rd
Kanata, ON K2K 2Y3
Tel: 613-271-9044
Fax: 613-271-9040

Hewetson Repping, Inc.

215 Traders Blvd, Unit 2
Mississauga, ON L4Z 3K5
Tel: 905-501-8800
Fax: 905-501-8818

Hewetson Repping, Inc.

100 boul. Alexis Nihon, suite 945
St-Laurent, PQ H4M 2P5
Tel: 514-744-6511
Fax: 514-744-0904

Hewetson Repping, Inc.

103 Douglas Woods Grove SE
Calgary AL T2Z 2H6

Tel: 403-236-4622
Fax: 403-236-4366

Denmark, Norway, Sweden

IE Komponenter AB

Ulv sundavagen 106C
Box 11113
161 11 Bromma, Sweden

Tel: 46 8 804685
Fax: 46 8 262286

Finland

Integrated Electronics OyAb

Laurinmaenkuja 3
FIN-00440 Helsinki
Finland

Tel: 358-9-586-1770
Fax: 358-9-586-1771

France

Compress

47 Rue de L'Estérel
539 94 633 Rungis CEDEX
FRANCE

Tel: 33 1 41 802902
Fax: 33 1 46 866763

Germany

P I Electronic Engineering, GmbH

Keitlanderstrasse 48
74354 Besenheim-Ottmarsheim,
Germany

Tel: 49 71 4381 1153
Fax: 49 71 4358 853

Hong Kong

Please contact Tundra Semiconductor
directly at: 613-592-0714

India

Hynetic Electronics

No. 50, 2nd Cross
Gavipuram Extension,
Bangalore, 560019, India

Tel: 91-80-620852
Fax: 91-80-624073

Israel

Migvan Technologies and Engineering Ltd.

13 Hashiloach St.
PO Box 7022
Petach-Tikva 49170, Israel

Tel: 972-3-9240784
Fax: 972-3-9240787

Italy

EL. CO. MI. S.r.l.

Viale G Matteotti, 26
20095 Cusano Milanino (MI),
Italy

Tel: 39 2 619 6452
Fax: 39 2 613 4836

Japan

K.K. Rocky

4-44-11, Nogata
Nakano-ku,
Tokyo, 165 Japan

Tel: 813 3228 4511
Fax: 813 3388 1391

Tokyo Electron Ltd.

TBS Broadcast Center
3-6 Akasaka, 5-Chome
Minato-Ku
Tokyo, 107 Japan

Tel: 81-35561-7216
Fax: 81-35561-7389

Korea

VME Tech Inc.

627-19 Shinsa-dong Kangnam-gu,
5th Floor, Shinsa Building,
Seoul, Korea 135-120

Tel: 82-2-3443-1491
Fax: 82-2-3443-1495

International Sales Representatives Cont'd

I & C Microsystems Co. Ltd.

8th Floor, Bethel Building,
324-1 Yangjae-Dong Seocho-Ku,
Seoul, Korea

Tel: 822-577-9131
Fax: 822-577-9130

Welvar Electronics Ltd.

MBM Box 66
Manor Lane, Rochester
Kent, England ME5 1HS

Tel: 44 1 634 815033
Fax: 44 1 634 832133

New Zealand

ACD

University Paton NZ Ltd.
Unit 7, 110 Mays Road,
Penrose
Auckland, NZ

Tel: 649-636-5984
Fax: 649-636-5985

Singapore

BBS Electronics

1 Genting Link
#05-04 Perfect Industrial Building
Singapore, 349 518
Singapore

Tel: 65 748 9851
Fax: 65 748 5965

Switzerland

Egli, Fischer & Co, Ltd. Zurich

Gotthardstrasse 6
Ch-8022
Zurich, Switzerland

Tel: 41 1 209 81 11
Fax: 41 1 201 22 75

Taiwan

Data Fast Electronics

6 Fl, No 27-1, Alley 169
Kang Ning Street
Hsi Chih
Taipei, Taiwan

Tel: 886 2 6921051
Fax: 886 2 6958240

United Kingdom

Thame Components

Thame Park Road
Thame, Oxfordshire OX9 3UQ

Tel: 44 1 844 261188
Fax: 44 1 844 261681