

ANSI/VITA 1-1994

Approved as an American National Standard by 

American National Standard for VME64

Secretariat
VMEbus International Trade Association

Approved April 10, 1995
American National Standards Institute, Inc.



VMEbus INTERNATIONAL TRADE ASSOCIATION

7825 E. Gelding Drive, Suite 104, Scottsdale, AZ 85260-3415

PH: 602-951-8866, FAX: 602-951-0720

E-mail: info@vita.com, URL: <http://www.vita.com>

American National Standard
for VME64

Secretariat
VMEbus International Trade Association

Approved April 10, 1995
American National Standards Institute, Inc.

Abstract

The VME64 specification establishes a framework for 8-, 16-, 32-, and 64-bit parallel-bus computer architectures that can implement single and multiprocessor systems. It is based on the VMEbus specification released by the VMEbus Manufacturers Group (now VITA) in August of 1982. This bus includes the initial four basic subbuses: (1) data transfer bus, (2) priority interrupt bus, (3) arbitration bus, and (4) utility bus. Other architectures with other subbuses are possible within this VME framework.

The data transfer bus will support 8-, 16-, 32-, and 64-bit data transfers in multiplexed and non multiplexed form. The transfer protocols are asynchronous with varying degrees of handshaking dependent on the speeds required. The priority interrupt subsystem provides real-time interrupt services to the system. The allocation of bus mastership is performed by the arbitration subsystem which allows the implementation of several prioritization algorithms. The utility bus provides the system with power plus power-up and power-down synchronization. The mechanical specifications of boards, backplanes, subracks, and enclosures are based on IEC 297 and IEEE 1101.1 specifications, also known as the Eurocard form factor. Additional standards exist that can be used as sub-busses to this architecture for data transfers transactions, peripheral interfaces and intra-crate communications among compatible modules.

American National Standard

Approval of an American National Standard requires verification by ANSI that the requirements for due process, consensus, and other criteria for approval have been met by the standards developer.

Consensus is established when, in the judgment of the ANSI Board of Standards Review, substantial agreement has been reached by directly and materially affected interests. Substantial agreement means much more than a simple majority, but not necessarily unanimity. Consensus requires that all views and objections be considered, and that a concerted effort be made toward their resolution.

The use of American National Standards is completely voluntary; their existence does not in any respect preclude anyone, whether they have approved the standards or not, from manufacturing, marketing, purchasing, or using products, processes, or procedures not conforming to the standards.

The American National Standards Institute does not develop standards and will in no circumstances give an interpretation of any American National Standard. Moreover, no person shall have the right or authority to issue an interpretation of an American National Standard in the name of the American National Standards Institute. Requests for interpretations should be addressed to the secretariat or sponsor whose name appears on the title page of this standard.

CAUTION NOTICE: This American National Standard may be revised or withdrawn at any time. The procedures of the American National Standards Institute require that action be taken periodically to reaffirm, revise, or withdraw this standard. Purchasers of American National Standards may receive current information on all standards by calling or writing the American National Standards Institute.

Published by

VMEbus International Trade Association
7825 E. Gelding Drive, Suite 104, Scottsdale, AZ 85260

Certain portions of this document are derived from portions of IEEE Std 1014-1987 IEEE Standard for a Versatile Backplane Bus: VMEbus, Copyright © 1987 by the Institute of Electrical and Electronics Engineers, Inc. and are adapted and reprinted with the permission of the IEEE. The IEEE disclaims any responsibility or liability resulting from the placement and use in this publication.

Copyright © 1995 by VMEbus International Trade Association
All rights reserved.

No part of this publication may be reproduced in any form, in an electronic retrieval system or otherwise, without prior written permission of the publisher.

NOTE: The user's attention is called to the possibility that compliance with this standard may require use of one or more inventions covered by patent rights.

By publication of this standard, no position is taken with respect to the validity of such claims or of any patent rights in connection therewith. The patent holders have, however, filed a statement of willingness to grant a license under these rights on reasonable and non-discriminatory terms and conditions to applicants desiring to obtain such a license for use of this standard. Details may be obtained from the publisher.

Printed in the United States of America - R1.13
ISBN 1-885731-02-7

TABLE OF CONTENTS

Abstract	x
Foreword	xi
VMEbus Specification Genealogy.....	xiii
Chapter 1	1
Introduction to the VMEbus Specification.....	1
1.1 VMEbus Specification Objectives.....	1
1.2 VMEbus Interface System Elements.....	1
1.2.1 Basic Definitions.....	1
1.2.1.1 Terms Used to Describe the VMEbus Mechanical Structure.....	1
1.2.1.2 Terms Used to Describe the VMEbus Functional Structure.....	2
1.2.1.3 Types of Cycles on the VMEbus.....	5
1.2.2 Basic VMEbus Structure.....	6
1.3 VMEbus Specification Diagrams.....	7
1.4 Specification Terminology.....	9
1.4.1 Signal Line States.....	10
1.4.2 Use of the Asterisk (*).....	10
1.4.3 Keyword Numbering.....	11
1.5 Protocol Specification.....	11
1.5.1 Interlocked Bus Signals.....	11
1.5.2 Broadcast Bus Signal.....	12
1.6 System Examples and Explanations.....	12
Chapter 2	13
Data Transfer Bus.....	13
2.1 Introduction.....	13
2.1.1 Enhancements.....	13
2.2 Data-Transfer-Bus Lines.....	14
2.2.1 Addressing Lines.....	16
2.2.2 Address Modifier Lines.....	20
2.2.3 Data Lines.....	23
2.2.4 Data Transfer Bus Control Lines.....	26
2.2.4.1 AS*.....	26
2.2.4.2 DS0* and DS1*.....	26
2.2.4.3 DTACK*.....	26
2.2.4.3.1 Rescinding DTACK*.....	27
2.2.4.4 BERR*.....	27
2.2.4.5 WRITE*.....	27
2.2.4.6 RETRY*.....	27
2.3 DTB Modules -- Basic Description.....	28

Table of Contents

2.3.1 Master.....	28
2.3.2 Slave.....	32
2.3.3 Bus Timer.....	34
2.3.4 Location Monitor.....	36
2.3.5 Addressing Phases and Modes.....	37
2.3.6 Basic Data Transfer Capabilities.....	39
2.3.7 Block Transfer Capabilities.....	42
2.3.8 Read-Modify-Write Capability.....	45
2.3.9 Unaligned Transfer Capability.....	47
2.3.10 Address-Only Capability.....	50
2.3.11 Lock Capability.....	51
2.3.12 Configuration ROM / Control And Status Register Capability.....	52
2.3.13 Retry Capability.....	57
2.3.14 Interaction Between DTB Functional Modules.....	58
2.4 Typical Operation.....	60
2.4.1 Typical Data-Transfer Cycles.....	60
2.4.2 Address Pipelining.....	67
2.5 Data-Transfer-Bus Acquisition.....	68
2.6 DTB Timing Rules and Observations.....	70
Chapter 3	121
Data Transfer Bus Arbitration.....	121
3.1 Bus Arbitration Philosophy.....	121
3.1.1 Types Of Arbitration.....	121
3.2 Arbitration Bus Lines.....	123
3.2.1 Bus Request And Bus Grant Lines.....	125
3.2.2 Bus Busy Line (BBSY*).....	125
3.2.3 Bus Clear Line (BCLR*).....	125
3.3 Functional Modules.....	126
3.3.1 Arbiter.....	126
3.3.3 Data Transfer Bus Master.....	133
3.3.3.1 Release Of The DTB.....	133
3.3.3.2 Acquisition Of The DTB.....	134
3.3.3.3 Other Information.....	134
3.4 Typical Operation.....	135
3.4.1 Arbitration Of Two Different Levels Of Bus Request.....	135
3.4.2 Arbitration Of Two Bus Requests On The Same Bus Request Line.....	139
3.5 Race Conditions Between Master Requests and Arbiter Grants.....	144
Chapter 4	145
Priority Interrupt Bus.....	145
4.1 Introduction.....	145
4.1.1 Single Handler Systems.....	145
4.1.2 Distributed Systems.....	145
4.2 Priority Interrupt Bus Lines.....	148
4.2.1 Interrupt Request Lines.....	149
4.2.2 Interrupt Acknowledge Line.....	149
4.2.3 Interrupt Acknowledge Daisy-Chain - IACKIN*/IACKOUT*.....	149
4.3 Priority Interrupt Bus Modules - Basic Description.....	150
4.3.1 Interrupt Handler.....	151

Table of Contents

4.3.2	Interrupter.....	153
4.3.3	IACK Daisy-Chain Driver.....	157
4.3.4	Interrupt Request Handling Capabilities.....	158
4.3.5	Interrupt Request Generation Capabilities.....	158
4.3.6	Status/Id Transfer Capabilities.....	159
4.3.7	Interrupt Request Release Capabilities.....	160
4.3.8	Interaction Between Priority Interrupt Bus Modules.....	162
4.4	Typical Operation.....	165
4.4.1	Single Handler Interrupt Operation.....	166
4.4.2	Distributed Interrupt Operation.....	166
4.4.2.1	Distributed Interrupt Systems With Seven Interrupt Handlers.....	166
4.4.2.2	Distributed Interrupt Systems With Two To Six Interrupt Handlers.....	168
4.4.3	Example: Typical Single Handler Interrupt System Operation.....	169
4.4.4	Example: Prioritization of Two Interrupts In A Distributed Interrupt System.....	172
4.5	Race Conditions.....	174
4.6	Priority Interrupt Bus Timing Rules and Observations.....	175
Chapter 5		199
Utility Bus		199
5.1	Introduction.....	199
5.2	Utility Bus Signal Lines.....	199
5.3	Utility Bus Modules.....	199
5.3.1	System Clock Driver.....	199
5.3.2	The Serial Bus Lines.....	199
5.3.3	Power Monitor.....	201
5.4	System Initialization and Diagnostics.....	202
5.5	Power and Ground pins.....	207
5.6	Reserved Line.....	207
5.7	Auto Slot Id.....	208
5.8	Auto System Controller.....	215
Chapter 6		217
Electrical Specifications.....		217
6.1	Introduction.....	217
6.2	Power Distribution.....	217
6.2.1	DC Voltage Specifications.....	218
6.2.2	Pin and Socket Connector Electrical Ratings.....	219
6.3	Electrical Signal Characteristics.....	219
6.4	Bus Driving and Receiving Requirements.....	220
6.4.1	Bus Driver Definitions.....	220
6.4.2	Driving And Loading RULEs For All VMEbus Lines.....	221
6.4.2.1	Driving And Loading RULEs For High Current Three-State Lines.....	222
6.4.2.2	Driving And Loading RULEs For Standard Three-State Lines.....	222
6.4.2.3	Driving And Loading RULEs For High Current Totem-Pole Lines.....	223
6.4.2.4	Driving And Loading RULEs For Standard Totem-Pole Lines.....	224
6.4.2.5	Driving And Loading RULEs For Open-Collector Lines.....	225
6.5	Backplane Signal Line Interconnections.....	226
6.5.1	Termination Networks.....	226
6.5.2	Characteristic Impedance.....	228

Table of Contents

6.5.3 Additional Information.....	229
6.6 User Defined Signals.....	230
6.7 Signal Line Drivers and Terminations.....	230
Chapter 7	233
Mechanical Specifications.....	233
7.1 Introduction.....	233
7.2 VMEbus Boards.....	234
7.2.1 Single Height Boards.....	234
7.2.2 Double Height Boards.....	235
7.2.3 Board Connectors.....	236
7.2.4 Board Assemblies.....	237
7.2.5 Board Widths.....	237
7.2.6 VMEbus Board Warpage, Lead Length and Component Height.....	238
7.3 Front Panels.....	238
7.3.1 Handles.....	239
7.3.2 Front Panel Mounting.....	240
7.3.3 Front Panel Dimensions.....	240
7.3.4 Filler Panels.....	241
7.3.5 Board Ejectors/Injectors.....	242
7.4 Backplanes.....	242
7.4.1 Backplane Dimensional Requirements.....	243
7.4.2 Signal Line Termination Networks.....	244
7.5 Assembly of VMEbus Subracks.....	244
7.5.1 Subracks And Slot Widths.....	244
7.5.2 Subrack Dimensions.....	245
7.6 Conduction Cooled VMEbus Systems.....	245
7.7 VMEbus Backplane Connectors and VMEbus Board Connectors.....	245
7.7.1 Pin Assignments For The J1/P1 Connector.....	246
7.7.2 Pin Assignments For The J2/P2 Connector.....	247
Appendix A	267
Glossary of VMEbus Terms.....	267
Appendix B	275
VMEbus Connector/Pin Description.....	275
Appendix C	279
Manufacturer's Board Identification.....	279
Appendix D	281
Rule Index	281

List of Figures

Figure 1 - 1	System Elements Defined By This Document.....	3
Figure 1 - 2	Functional Modules and Buses Defined by the Document.....	8
Figure 1 - 3	Signal Timing Notation.....	12
Figure 2 - 1	Data Transfer Bus Functional Block Diagram.....	15
Figure 2 - 2	Block Diagram.....	29
Figure 2 - 3	Block Diagram.....	32
Figure 2 - 4	Block Diagram.....	35
Figure 2 - 5	Block Diagram.....	36
Figure 2 - 6	Four Ways That 32 Bits Of Data Might Be Stored In Memory.....	47
Figure 2 - 7	Four Ways That 16 Bits of Data Might Be Stored in Memory.....	48
Figure 2 - 34	Block Diagram.....	53
Figure 2 - 8	An Example of a Non-multiplexed Address, Single-Byte Read Cycle.....	61
Figure 2 - 9	An Example Of A Multiplexed Address Double-Byte Write Cycle.....	62
Figure 2 - 10	An Example Of A Non-multiplexed Address Quad-Byte Write Cycle.....	64
Figure 2 - 35	An Example Of An Eight-Byte Block Read Cycle.....	65
Figure 2 - 11	Data Transfer Bus Master Exchange Sequence.....	69
Figure 2 - 12	Address Broadcast Timing - All Cycles.....	95
Figure 2 - 13	A16, A24, A32 Master, Responding Slave, And Location Monitor.....	96
Figure 2 - 14	Master, Slave, and Location Monitor - A16, A24 and A32 Address Broadcast Timing.....	97
Figure 2 - 15	Master, Slave, And Location Monitor A16, A24, And A32 Address Broadcast Timing.....	98
Figure 2 - 27	Master, Slave And Location Monitor - A64 , A40, And ADOH Address Broadcast Timing.....	99
Figure 2 - 16	Master, Slave, And Location Monitor Data Transfer Timing.....	100
Figure 2 - 17	Master, Slave, And Location Monitor Data Transfer Timing.....	102
Figure 2 - 28	Master, Slave And Location Monitor Data Transfer Timing.....	104
Figure 2 - 18	Master, Slave, And Location Monitor Data Transfer Timing.....	106
Figure 2 - 19	Master, Slave, And Location Monitor Data Transfer Timing.....	108
Figure 2 - 29	Master, Slave And Location Monitor Data Transfer Timing.....	110
Figure 2 - 20	Master, Slave, And Location Monitor Data Transfer Timing Single-Byte RMW Cycles.....	112
Figure 2 - 21	Master, Slave, and Location Monitor Data Transfer Timing.....	113
Figure 2 - 22	Address Strobe Inter-Cycle Timing.....	114
Figure 2 - 23	Data Strobe Inter-Cycle Timing.....	114
Figure 2 - 24	Data Strobe Inter-Cycle Timing.....	115
Figure 2 - 25	Master, Slave, And Bus Timer Data Transfer Timing.....	115
Figure 2 - 26	Master DTB Control Transfer Timing.....	116
Figure 2 - 30	Master And Slave Data Transfer Timing.....	117
Figure 2 - 31	Master And Slave Data Transfer Timing.....	118
Figure 2 - 32	A40, MD32 Read-Modify-Write.....	119
Figure 2 - 33	Rescinding DTACK Timing.....	120
Figure 3 - 1	Arbitration Functional Block Diagram.....	122
Figure 3 - 2	Illustration Of The Daisy Chain Bus Grant Lines.....	124
Figure 3 - 3	Block Diagram.....	128
Figure 3 - 4	Block Diagram.....	132
Figure 3 - 5	Arbitration Flow Diagram Two Requesters.....	136
Figure 3 - 6	Arbitration Sequence Diagram Two Requesters, Two Request Levels.....	138
Figure 3 - 7	Arbitration Flow Diagram.....	140
Figure 3 - 8	Arbitration Sequence Diagram.....	143
Figure 4 - 1	Priority Interrupt Bus Functional Diagram.....	146
Figure 4 - 2	Interrupt Subsystem Structure.....	147
Figure 4 - 3	Interrupt Subsystem Structure.....	148

Table of Contents

Figure 4 - 4	IACKIN*/IACKOUT* DAISY-CHAIN.....	150
Figure 4 - 5	Block Diagram.....	152
Figure 4 - 6	Block Diagram.....	156
Figure 4 - 7	Block Diagram.....	157
Figure 4 - 8	Release Of Interrupt Request Lines By ROAK And RORA Interrupters.....	162
Figure 4 - 9	An IACK Daisy-Chain Driver And Interrupter On The Same Board.....	164
Figure 4 - 10	Two Interrupters On The Same Board.....	165
Figure 4 - 11	The Three Phases Of An Interrupt Sequence.....	166
Figure 4 - 12	Two Interrupt Handlers, Each Monitoring One Interrupt.....	167
Figure 4 - 13	Two Interrupt Handlers, Each Monitoring Several Interrupt.....	168
Figure 4 - 14	Typical Single Handler Interrupt System Operation.....	170
Figure 4 - 15	Typical Distributed Interrupt System With Two Interrupt Handlers, Flow Diagram.....	173
Figure 4 - 16	Interrupt Handler And Interrupter--Interrupter Selection Timing.....	189
Figure 4 - 17	IACK Daisy-Chain Driver--Interrupter Selection Timing.....	190
Figure 4 - 18	Participating Interrupter--Interrupter Selection Timing.....	191
Figure 4 - 19	Responding Interrupter--Interrupter Selection Timing.....	192
Figure 4 - 20	Interrupt Handler--Status/Id Transfer Timing.....	193
Figure 4 - 21	Interrupt Handler--Status/Id Transfer Timing.....	194
Figure 4 - 22	Responding Interrupter--Status/Id Transfer Timing.....	195
Figure 4 - 23	Responding Interrupter--Status/Id Transfer Timing.....	196
Figure 4 - 24	IACK Daisy-Chain Driver, Responding Interrupter, And Participating Interrupter IACK Daisy-Chain Inter-Cycle Timing.....	197
Figure 5 - 1	Utility Bus Block Diagram.....	200
Figure 5 - 2	System Clock Driver Timing.....	201
Figure 5 - 3	Block Diagram Of Power Monitor Module.....	202
Figure 5 - 4	Power Monitor Power Failure Timing.....	203
Figure 5 - 5	Power Monitor System Restart Timing.....	203
Figure 5 - 6	SYSRESET* and SYSFAIL* Timing Diagram.....	206
Figure 5 - 7	Current Rating For Power Pins.....	208
Figure 5 - 9	First Slot Detector (FSD).....	215
Figure 6 - 1	VMEbus Signal Levels.....	220
Figure 6 - 2	Standard Bus Termination.....	227
Figure 7 - 1	Subrack with Mixed Board Sizes.....	248
Figure 7 - 2	Single Height Board.....	249
Figure 7 - 3	Double Height Board.....	250
Figure 7 - 4	Connector Position On Single And Double Height Boards.....	251
Figure 7 - 5	Cross Sectional View Of Board, Connector, Backplane, And Front Panel.....	252
Figure 7 - 6	Optional Enhanced DIN Connector.....	253
Figure 7 - 7	Component Height, Lead Length And Board Warpage.....	254
Figure 7 - 8	Single Height, Single Width Front Panel.....	255
Figure 7 - 9	Double Height, Single Width Front Panel.....	256
Figure 7 - 10	Front Panel Mounting Brackets And Dimensions Of Single Height Boards.....	257
Figure 7 - 11	Front Panel Mounting Brackets And Dimensions Of Double Height Boards.....	258
Figure 7 - 12	Single Height Filler Panel.....	259
Figure 7 - 13	Double Height Filler Panel.....	260
Figure 7 - 14	Detailed Dimensions of a J1 and a J2 Backplane.....	261
Figure 7 - 15	Detailed Dimensions of a J1/J2 Backplane.....	262
Figure 7 - 16	"Off-Board Type" Backplane Termination.....	263
Figure 7 - 17	"On-Board Type" Backplane Termination.....	264
Figure 7 - 18	21 Slot Sub Rack.....	265
Figure 7 - 19	Board Guide Detail.....	266

List of Tables

Table 2 - 1	The Eight Categories Of Byte Locations.....	16
Table 2 - 28	Address Alignment on Bus.....	16
Table 2 - 2	Signal Levels During Data Transfers.....	18
Table 2 - 3	Address Modifier Codes.....	21
Table 2 - 4	Use Of Data Lines To Move Data During Nonmultiplexed Data Transfers.....	24
Table 2 - 29	Use Of The Address And Data Lines For Multiplexed Data Cycles.....	25
Table 2 - 5	RULES And PERMISSIONS That Specify The Use Of The Dotted Lines By The Various Types Of Masters.....	30
Table 2 - 6	Slaves - RULES And PERMISSIONS That Specify The Use of The.....	33
Table 2 - 7	Use Of The BTO() Mnemonic.....	34
Table 2 - 8	Location Monitors - RULES And PERMISSIONS That Specify The Use Of The Dotted Lines By The Various Types Of Location Monitors.....	37
Table 2 - 9	Mnemonics That Specify Addressing Capabilities.....	38
Table 2 - 10	Mnemonics That Specify Basic Data Transfer Capabilities.....	41
Table 2 - 11	Mnemonics That Specify Block Transfer Capabilities.....	44
Table 2 - 12	The Mnemonic That Specifies Read-Modify-Write Capabilities.....	46
Table 2 - 13	Transferring 32 Bits Of Data Using Multiple-Byte Transfer Cycles.....	48
Table 2 - 14	Transferring 16 Bits Of Data Using Multiple-Byte Transfer Cycles.....	49
Table 2 - 15	Mnemonic That Specifies Unaligned Transfer Capability.....	50
Table 2 - 16	Mnemonics That Specify Address Only Capability.....	51
Table 2 - 30	Configuration ROM / Control & Status Registers.....	54
Table 2 - 31	Control And Status Register Base Definition.....	54
Table 2 - 32	Configuration ROM Definition.....	55
Table 2 - 17	Timing Diagrams That Define Master, Slave, And Location Monitor Operation (See Table 2-22 For Timing Values).....	72
Table 2 - 18	Definitions Of Mnemonics Used In Tables 2 -19, 2 - 20, And 2 - 21.....	74
Table 2 - 19	Use Of The Address And Data Lines To Select A Byte Group.....	75
Table 2 - 20	Use of DS1*, DS0*, A1, A2, and LWORD* During the Address Phase of the.....	76
Table 2 - 21	Use Of The Data Lines To Transfer Data.....	79
Table 2 - 22	Master, Slave, And Location Monitor Timing Parameters.....	82
Table 2 - 23	Bus-Timer Timing Parameters.....	83
Table 2 - 24	Master, Timing RULEs and OBSERVATIONS.....	84
Table 2 - 25	Slave, Timing RULEs and OBSERVATIONS.....	89
Table 2 - 26	Location Monitor, Timing OBSERVATIONS.....	93
Table 2 - 27	Bus Timer, Timing RULEs.....	94
Table 4 - 1	RULES And PERMISSIONS That Specify The Use Of The Dotted Lines.....	153
Table 4 - 2	RULES And PERMISSIONS That Specify The Use Of The Dotted.....	155
Table 4 - 3	Use Of The IH() Mnemonic To Specify Interrupt.....	158
Table 4 - 4	Use Of The I() Mnemonic To Specify Interrupt Request.....	158
Table 4 - 5	Mnemonics That Specify Status/Id Transfer Capabilities.....	159
Table 4 - 6	Mnemonics That Specify Interrupt Request Release Capabilities.....	161
Table 4 - 7	3-Bit Interrupt Acknowledge Code.....	172
Table 4 - 8	Timing Diagrams That Define Interrupt Handler And.....	176
Table 4 - 9	Timing Diagrams That Define IACK Daisy-Chain Driver Operation.....	177
Table 4 - 10	Timing Diagrams That Define Participating Interrupter Operation.....	177
Table 4 - 11	Timing Diagrams That Define Responding Interrupter Operation.....	177

Table of Contents

Table 4 - 12	Definitions Of Mnemonics Used In Tables 4-13, 4-14, And 4-15.....	178
Table 4 - 13	Use Of Addressing Lines During Interrupt Acknowledge Cycles.....	179
Table 4 - 14	Use Of The DS1*, DS0*, LWORD*, And WRITE* Lines During.....	179
Table 4 - 15	Use of the data bus lines to transfer the Status/Id.....	179
Table 4 - 16	Interrupt Handler, Interrupter, And IACK DAISY-CHAIN.....	180
Table 4 - 17	Interrupt Handler, Timing RULES And OBSERVATIONS.....	181
Table 4 - 18	Interrupter, Timing RULES And OBSERVATIONS.....	184
Table 4 - 19	IACK Daisy-Chain Driver, Timing RULES And OBSERVATIONS.....	188
Table 5 - 1	Module Drive During Power-Up and Power-Down Sequences.....	205
Table 6 - 1	Bus Voltage Specification.....	218
Table 6 - 2	Bus Driving And Receiving Requirements.....	221
Table 6 - 3	Bus Driver Summary.....	230
Table 7 - 1	J1/P1 Pin Assignments.....	246
Table 7 - 2	J2/P2 Pin Assignments.....	247

Abstract

The VME64 specification establishes a framework for 8-, 16-, 32-, and 64-bit parallel-bus computer architectures that can implement single and multiprocessor systems. It is based on the VMEbus specification released by the VMEbus Manufacturers Group (now VITA) in August of 1982. This bus includes the initial four basic subbuses: (1) data transfer bus, (2) priority interrupt bus, (3) arbitration bus, and (4) utility bus. Other architectures with other subbuses are possible within this VME framework.

The data transfer bus will support 8-, 16-, 32-, and 64-bit data transfers in multiplexed and non multiplexed form. The transfer protocols are asynchronous with varying degrees of handshaking dependent on the speeds required. The priority interrupt subsystem provides real-time interrupt services to the system. The allocation of bus mastership is performed by the arbitration subsystem which allows the implementation of several prioritization algorithms. The utility bus provides the system with power plus power-up and power-down synchronization. The mechanical specifications of boards, backplanes, subracks, and enclosures are based on IEC 297 and IEEE 1101.1 specifications, also known as the Eurocard form factor. Additional standards exist that can be used as sub-busses to this architecture for data transfer transactions, peripheral interfaces and intra-crate communications among compatible modules.

FOREWORD

(This foreword is not part of ANSI/VITA 1-1994.)

The architectural concepts of VMEbus are based on the VERSAbus developed by Motorola in the late 1970s. Motorola's European Microsystems group in Munich, West Germany proposed the development of a VERSAbus-like product line of computers and controllers based on the Eurocard mechanical standard. To demonstrate the concept, Max Loesel and Sven Rau developed three prototype boards: (1) a 68000 CPU card, (2) a dynamic memory card, and (3) a static memory card. They named the new bus VERSAbus-E, which was later renamed "VME" by Lyman Hevle, then VP of the Motorola Microsystems Operation (and later the founder of VITA). VME is the acronym for VERSA-module Europe. Motorola, Mostek, and Signetics agreed to jointly develop and support the new bus architecture in early 1981.

John Black of Motorola, Craig McKenna of Mostek, and Cecil Kaplinsky of Signetics developed the first draft of the VMEbus specification. In October of 1981, at the Systems 81 trade show in Munich, West Germany, Motorola, Mostek, and Signetics announced their joint support for VMEbus, and placed Revision A of the specification in the public domain.

In August of 1982, Revision B of the VMEbus specification was published by the newly formed VMEbus Manufacturers Group (now VITA). This new revision refined the electrical specifications for the signal line drivers and receivers, and also brought the mechanical specifications more in line with the developing IEC 297 standard, the formal specifications for Eurocard mechanical formats.

In the latter part of 1982, the French delegation of the International Electrotechnical Commission (IEC) proposed Revision B of the VMEbus specification as an international standard. The IEC SC47B subcommittee nominated Mira Pauker of Philips, France, as the chairperson of an editorial committee, formally starting international standardization of the VMEbus.

In March of 1983, the IEEE Microprocessor Standards Committee (MSC) requested authorization to establish a working group to standardize the VMEbus in the US. This request was approved by the IEEE Standards Board, and the P1014 Working Group was established. Wayne Fischer was appointed first chairman of the working group. John Black served as chairman of the P1014 Technical Subcommittee.

The IEC, IEEE, and VMEbus Manufacturers Group (now VITA) distributed copies of Revision B for comment, and received requests for changes to the document as a result. These comments made it clear that it was time to go forward past revision B. In December of 1983, a meeting was held that included John Black, Mira Pauker, Wayne Fischer, and Craig McKenna. It was agreed that a revision C should be created, and that it should take into consideration all comments received by the three organizations. John Black and Shlomo Pri-Tal of Motorola incorporated the changes from all sources into a common document. The VMEbus Manufacturers Group (now VITA) labeled the document Revision C.1 and placed it in the public domain. The IEEE labeled it P1014 Draft 1.2, and the IEC labeled it IEC 821 Bus. Subsequent ballots in the IEEE P1014 group and in the MSC resulted in more comments, and required that the IEEE P1014 draft be updated. This work resulted in the ANSI/IEEE 1014-1987 specification.

Table of Contents

In 1989, John Peters of Performance Technologies, Inc. (Rochester, NY) developed the initial concept of VME64: multiplexing address and data lines (A64/D64) on the VMEbus. This concept was shown for the first time in 1989 and placed in the VITA Technical Committee in 1990 as a performance enhancement to the VMEbus specification. In 1991, the PAR (Project Authorization Request) for P1014R (revisions to the VMEbus specification) was granted by the IEEE. Ray Alderman, Technical Director of VITA, co-chaired the activity with Kim Clohessy of DY 4 Systems (Nepean, Ontario, Canada).

At the end of 1992, the additional enhancements to VMEbus (A40/D32, Locked Cycles, Rescinding DTACK*, Autoslot-ID, Auto System Controller, and enhanced DIN connector mechanicals) required more work to complete this document. In 1992, the VITA Technical Committee suspended work with the IEEE and sought accreditation as a standards developer organization (SDO) with the American National Standards Institute. The original IEEE Par P1014R was subsequently withdrawn by the IEEE. The VITA Technical Committee returned to using the public domain VMEbus C.1 specification as its base level document to which it added new enhancements. This enhancement work was undertaken entirely by the VITA Technical Committee resulting in this document. The tremendous undertaking of the document editing was accomplished by Kim Clohessy of DY 4 Systems, the technical co-chair of the activity with great help from Frank Hom who created the mechanical drawings, and with exceptional contributions by each chapter editor.

Additional enhancements proposed to the VME64 Subcommittee have been placed in another VITA subcommittee: the VME64 Extensions Document. Two other activities began in late 1992: (1) BLLI (VMEbus Board-level Live Insertion Specifications), and (2) VSLI (VMEbus System-level Live Insertion with Fault Tolerance).

New activities begun in 1993 using the base-VME architecture involve the implementation of high-speed serial and parallel sub-buses for use as I/O interconnections and data mover subsystems. These architectures can be used as message switches, routers, and small multiprocessor parallel architectures.

VITA's application for recognition as an accredited standards developer organization of ANSI (American National Standards Institute) was granted in June 1993. Numerous other documents, including mezzanine, P2, and serial bus standards, have been placed with VITA as the Public Domain Administrator of these technologies.

VMEBUS Specification Genealogy

VMEbus	Revision B and C.1 (Public Domain)
IEEE 1014-1987	Versatile Backplane Bus VMEbus
VITA 1-1994	VME64 Specification
IEEE 1096-1988	VSBbus Specification (IEEE)
IEC 821	A Bus Specification
IEEE 1101.1	IEEE Standard for Mechanical Core Specifications for Microcomputers Using IEC 603-2 Connectors
IEEE 1101.2	IEEE Standard for Mechanical Core Specification for Conduction-Cooled Eurocards

This standard was constructed through the many hours of hard work by the members of the VME64 Subcommittee (of the VITA Technical Committee) and the commitment of their companies to this standard.

NAME	COMPANY
Ray Alderman	PEP Modular Computers
Michael Humphrey	VITA
John Rynearson	VITA
Martin Blake	BICC-VERO UK
Kim Clohessy	DY 4 Systems, Inc.
Jing Kwok	DY 4 Systems, Inc.
Clarence Peckham	Heurikon Corporation
Dennis Terry	Heurikon Corporation
Wayne Fischer	Force Computers Inc.
Jack Regula	Consultant
Tad Kubic	Dawn VME Products, Inc.
Will Hamsher	AMP, Inc.
Doug Reubendall	AMP, Inc.
William Mahusen	Performance Technologies, Inc.
Thanos Mentzelopoulos	Ironics, Inc.
Joel Silverman	Radstone Technology Corp.
Colin Davies	Radstone Technology UK
Frank Hom	Electronic Solutions
Keith Burgess	Mizar
John Black	Micrology PBT
Greg Hill	Hewlett-Packard Co.
Sam Babb	Hewlett Packard Co.
Ben LaPointe	Motorola GEG
Chau Pham	Motorola MCG
Mac Rush	Motorola MCG
Mike Hasenfratz	Micro Memory, Inc.
Richard DeBock	Matrix Corp.
Richard O'Connor	Newbridge Microsystems
Michael Bryant	PEP Modular Computers

Table of Contents

Dr. Chris Eck	CERN-Geneva
Mike Thompson	Schroff, Inc.
Eike Waltz	Schroff, Inc.
Tom Baillio	Mercury Computer Systems, Inc.
Steve Corbesero	Mupac Corp.
Dave Horton	Cypress Semiconductor
Mike Maas	Cypress Semiconductor
Mike Munroe	Hybricon Corp.

The following are acknowledged for their extra contributions as chapter editors.

Kim Clohessy	DY-4 Systems Inc.
Wayne Fischer	Force Computers, Inc.
Chau Pham	Motorola MCG
Frank Hom	Electronic Solutions
Richard DeBock	Matrix Corp.
Jing Kwok	DY 4 Systems, Inc.
Mike Hasenfratz	Micro Memory, Inc.

Canvass Ballot

Consensus for this standard was achieved by use of the Canvass Method.

The following organizations, recognized as having an interest in the standardization of VME64, participated in the Canvass Ballot process. Inclusion in this list does not necessarily imply that the organization concurred with the submittal of the proposed standard to ANSI.

767 AWACS	Technology Consulting
Adept Technology	Texas Instruments
AMP	VERO Electronics
AT&T Bell Laboratories	VITA
Berg Electronics	VME MEMBER
Bit 3 Computer	Winchester Electronics
CERN	
CSPI	
Cypress Semiconductor	
Dawn VME Products	
Dialogic Corporation	
Digital Equipment Corp.	
DY 4 Systems	
Electronic Solutions	
Force Computers	
Harting Elektronik	
Heurikon Corporation	
Hewlett-Packard	
Hughes Aircraft Company	
Hybricon Corporation	
IBM	
IXTHOS	
Loral Western Devel. Lab	
Los Alamos Nat'l Lab	
Matrix Corporation	
Micro Memory	
Micrology	
MITRE Corporation	
Motorola Computer Group	
Mupac Corporation	
Newbridge Microsystems	
Object Technology Inc.	
PEP Modular Computers	
Philips Ind. Automation	
Picosoft	
Radstone Technology	
Schroff	

Chapter 1

Introduction to the VMEBUS Specification

1.1 VMEBUS Specification Objectives

The VMEbus specification defines an interfacing system used to interconnect microprocessors, data storage, and peripheral control devices in a closely coupled hardware configuration. The system has been conceived with the following objectives:

- a. To allow communication between devices on the VMEbus without disturbing the internal activities of other devices interfaced to the VMEbus.
- b. To specify the electrical and mechanical system characteristics required to design devices that will reliably and unambiguously communicate with other devices interfaced to the VMEbus.
- c. To specify protocols that precisely define the interaction between the VMEbus and devices interfaced to it.
- d. To provide terminology and definitions that describe system protocol.
- e. To allow a broad range of design latitude so that the designer can optimize cost and/or performance without affecting system compatibility.
- f. To provide a system where performance is primarily device limited, rather than system interface limited.

1.2 VMEBUS Interface System Elements

1.2.1 Basic Definitions

The VMEbus structure can be described from two points of view: its mechanical structure and its functional structure. The mechanical specification describes the physical dimensions of subracks, backplanes, front panels, plug-in boards, etc. The VMEbus functional specification describes how the bus works, what functional modules are involved in each transaction, and the rules which govern their behavior. This section provides informal definitions for some basic terms used to describe both the mechanical and the functional structure of the VMEbus.

1.2.1.1 Terms Used to Describe the VMEbus Mechanical Structure

VMEbus Backplane — A printed circuit (PC) board with 96 or 160 pin connectors and signal paths that bus the connector pins. Some VMEbus systems have a single PC board, called the J1 backplane. It provides the signal paths needed for basic operation. Other VMEbus systems also have an optional second PC board, called a J2 backplane. It provides the additional 96 or 160 pin connectors and signal paths needed for wider data and address transfers. Still others have a single PC board that provides the signal conductors and connectors of both the J1 and J2 backplanes.

Board — A printed circuit (PC) board, its collection of electronic components, with either one or two 96 or 160 pin connectors that can be plugged into VMEbus backplane connectors.

Slot — A position where a board can be inserted into a VMEbus backplane. If the VMEbus system has both a J1 and a J2 backplane (or a combination J1/J2 backplane) each slot provides a pair of 96 or 160 pin connectors. If the system has only a J1 backplane, then each slot provides a single 96 or 160 pin connector.

Subrack — A rigid framework that provides mechanical support for boards inserted into the backplane, ensuring that the connectors mate properly and that adjacent boards do not contact each other. It also guides the cooling airflow through the system, and ensures that inserted boards do not disengage themselves from the backplane due to vibration or shock.

1.2.1.2 Terms Used to Describe the VMEbus Functional Structure

Figure 1-1 shows a simplified block diagram of the functional structure, including the VMEbus signal lines, backplane interface logic, and functional modules.

Backplane Interface Logic — Special interface logic that takes into account the characteristics of the backplane: its signal line impedance, propagation time, termination values, etc. The VMEbus specification prescribes certain rules for the design of this logic based on the maximum length of the backplane and its maximum number of board slots.

Functional Module — A collection of electronic circuitry that resides on one VMEbus board and works together to accomplish a task.

Data Transfer Bus — One of the four buses provided by the VMEbus backplane. The Data Transfer Bus allows Masters to direct the transfer of binary data between themselves and Slaves. (Data Transfer Bus is often abbreviated DTB.)

Data Transfer Bus Cycle — A sequence of level transitions on the signal lines of the DTB that result in the transfer of an address or an address and data between a Master and a Slave. The Data Transfer Bus cycle is divided into two portions, the address broadcast and then zero or more data transfers. There are 34 types of Data Transfer Bus cycles. They are defined later in this chapter.

Master — A functional module that initiates DTB cycles in order to transfer data between itself and a Slave module.

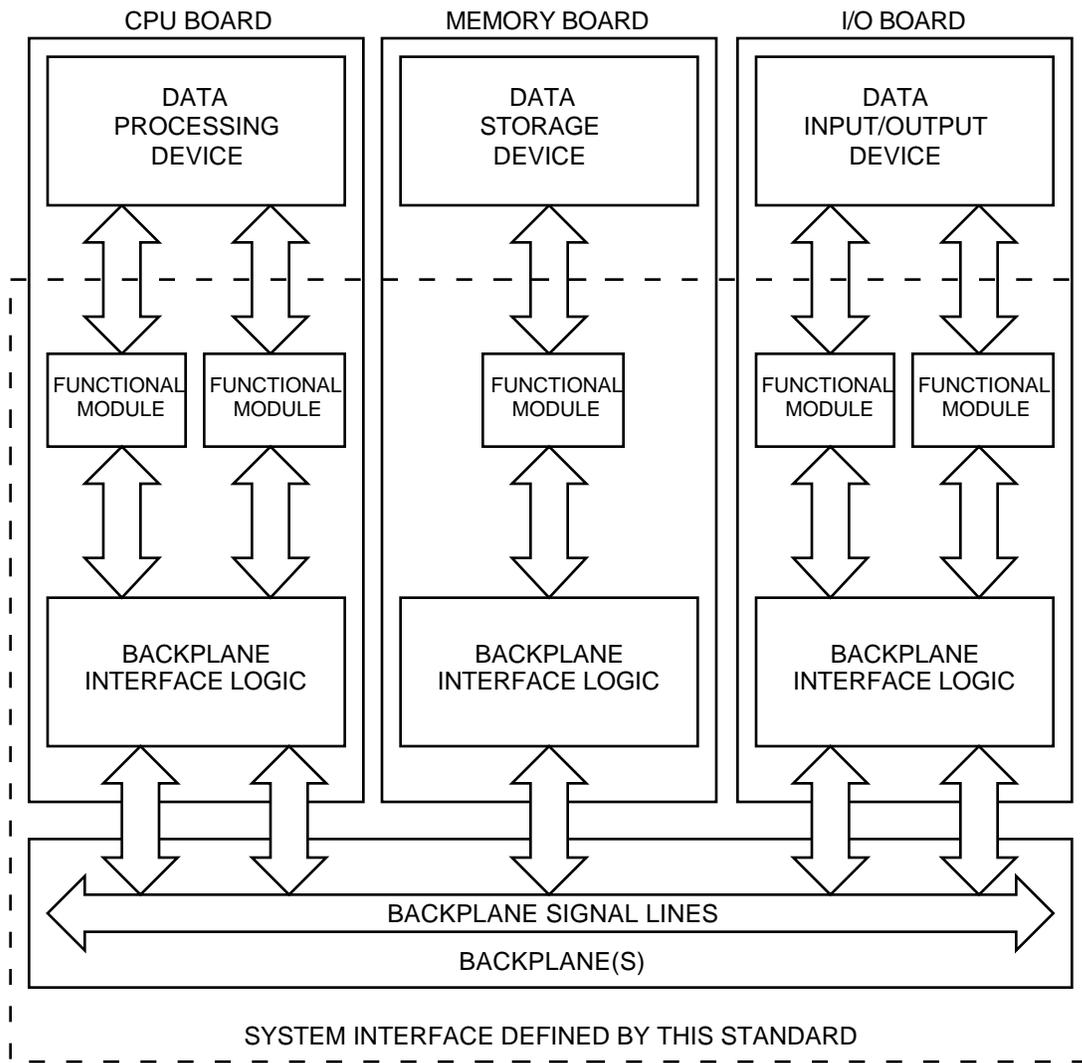


Figure 1 - 1 System Elements Defined By This Document

Slave — A functional module that detects DTB cycles initiated by a Master and, when those cycles specify its participation, transfers data between itself and the Master.

Location Monitor — A functional module that monitors data transfers over the DTB in order to detect accesses to the locations it has been assigned to watch. When an access to one of these assigned locations occurs, the Location Monitor generates an on-board signal.

Bus Timer — A functional module with a preset time-out period which terminates the DTB cycle if a transfer exceeds the time-out period. Without this module, if the Master tries to transfer data to or from a nonexistent Slave location it might wait forever. The Bus Timer prevents this by terminating the cycle.

Priority Interrupt Bus — One of the four buses provided by the VMEbus backplane. The Priority Interrupt Bus allows Interrupter modules to send interrupt requests to Interrupt Handlers.

Interrupter — A functional module that generates an interrupt request on the Priority Interrupt Bus and then provides Status/ID information when the Interrupt Handler requests it.

Interrupt Handler — A functional module that detects interrupt requests generated by Interrupters and responds to those requests by asking for Status/ID information.

Status/Id -- An eight, sixteen, or thirty two bit value returned by an interrupter to an interrupt handler during an interrupt acknowledge cycle.

Daisy-Chain — A special type of VMEbus signal line that is used to propagate a signal level from board to board, starting with the first slot and ending with the last slot. There are four bus grant daisy-chains and one interrupt acknowledge daisy-chain on the VMEbus.

IACK Daisy-Chain Driver — A functional module which activates the interrupt acknowledge daisy-chain whenever an Interrupt Handler acknowledges an interrupt request. This daisy-chain ensures that only one Interrupter will respond with its STATUS/ID when more than one has generated an interrupt request.

Arbitration Bus — One of the four buses provided by the VMEbus backplane. This bus allows an Arbiter module and several Requester modules to coordinate use of the DTB.

Requester — A functional module that resides on the same board as a Master or Interrupt Handler and requests use of the DTB whenever its Master or Interrupt Handler needs it.

Arbiter — A functional module that accepts bus requests from Requester modules and grants control of the DTB to one Requester at a time.

Utility Bus — One of the four buses provided by the VMEbus backplane. This bus includes signals that provide periodic timing and coordinate the power-up and power-down of VMEbus systems.

CR/CSR -- A functional module that provides Configuration ROM information and Control and Status Registers. The module provides manufacturing and board ID and other important board information. The CSRs are used for software configuration of a VMEbus system.

System Clock Driver — A functional module that provides a 16 MHz timing signal on the Utility Bus.

Serial Bus — A functional module that provides a bused 2 wire interface between cards in the backplane, independent of the other VMEbus modules.

Power Monitor Module — A functional module that monitors the status of the primary power source to the VMEbus system, and signals when that power has strayed outside the limits required for reliable system operation. Since most systems are powered by an AC source, the Power Monitor is typically designed to detect drop-out or brown-out conditions on AC lines.

System Controller Board — A board which resides in slot 1 of a VMEbus backplane and has a System Clock Driver, an Arbiter, an IACK Daisy-Chain Driver, and a Bus Timer. Some also have a Power Monitor.

1.2.1.3 Types of Cycles on the VMEbus

Read Cycle — A DTB cycle used to transfer 1, 2, 3, 4 or 8 bytes from a Slave to a Master. The cycle begins when the Master broadcasts an address and an address modifier. Each Slave captures the address modifier and address and checks to see if it is to respond to the cycle. If so, it retrieves the data from its internal storage, places it on the data bus and acknowledges the transfer. The Master then terminates the cycle.

Write Cycle — A DTB cycle used to transfer 1, 2, 3, 4 or 8 bytes from a Master to a Slave. The cycle begins when the Master broadcasts an address and address modifier and places data on the DTB. Each Slave captures the address and address modifier and checks to see if it is to respond to the cycle. If so, it stores the data and then acknowledges the transfer. The Master then terminates the cycle.

Block Read Cycle — A DTB cycle used to transfer a block of 1 to 256 bytes from a Slave to a Master. This transfer is done using a string of 1, 2, or 4 byte data transfers. Once the block transfer is started, the Master does not release the DTB until all of the bytes have been transferred. It differs from a string of read cycles in that the Master broadcasts only one address and address modifier (at the beginning of the cycle). Then the Slave increments this address on each transfer so that the data for the next transfer is retrieved from the next higher location.

Block Write Cycle — A DTB cycle used to transfer a block of 1 to 256 bytes from a Master to a Slave. The block write cycle is very similar to the block read cycle. It uses a string of 1, 2, or 4 byte data transfers. The Master does not release the DTB until all of the bytes have been transferred. It differs from a string of write cycles in that the Master broadcasts only one address and address modifier (at the beginning of the cycle). Then the Slave increments this address on each transfer so that the data from the next transfer is stored in the next higher location.

Multiplexed Cycle — A DTB cycle that transfers address information and/or data information using both the address and data buses. Multiplexed cycles are used in four cases.

- a: A64 - the full address bus and the full data bus are combined to create a 64 bit address.
- b: MBLT - the full address bus and the full data bus are combined to create a 64 bit data word.
- c: A40 - The full 24 bit address bus and the full 16 bit data bus on the P1/J1 connector are combined to create a 40 bit address. This mode is especially useful for 3U boards which have a J1 connector only.
- d: MD32 - The lower 16 address lines and the lower 16 data lines are combined to create a 32 bit data word. This mode is especially useful for 3U modules.

Multiplexed cycles are used in both basic transfers and block transfers. A64 and A40 basic transfers support 1 byte, 2 byte and 4 byte transfers. In addition MBLT cycles support 8 byte transfers.

A Multiplexed Cycle will have an Address phase that is separate from the Data phase. The Address phase may include (i.e. A64 and A40 cycles) or may not include (i.e. A32, A24 cycles) the use of the Data bus. The Data phase may include (MBLT, MD32 cycles) or may not include (i.e. D32, D16, D08(OE) cycles) the use of the Address bus.

Multiplexed block transfers can have 1 to 256 transfers between the Master and a Slave. With 8 byte transfers, up to 2,048 bytes can be transferred in one block.

Read-Modify-Write Cycle — A DTB cycle that is used to both read from, and write to, a Slave location without permitting any other Master to access that location. This cycle is most useful in multiprocessing systems where certain memory locations are used to provide semaphore functions.

Address-Only Cycle — A DTB cycle that consists of an address broadcast, but no data transfer. Slaves do not acknowledge ADDRESS-ONLY cycles and Masters terminate the cycle without waiting for an acknowledgment. No data strobes or acknowledge strobes are asserted in an ADDRESS-ONLY cycle.

Address-Only-With-Handshake Cycle -- A DTB cycle that consists of an address broadcast, but no data transfer. The addressed Slave responds in the same manner as a standard access cycle.

Interrupt Acknowledge Cycle — A DTB cycle, initiated by an Interrupt Handler, which reads a STATUS/ID from an Interrupter. An Interrupt Handler generates this cycle whenever it detects an interrupt request from an Interrupter and it has control of the DTB.

1.2.2 Basic VMEbus Structure

The VMEbus interface system consists of backplane interface logic, four groups of signal lines called buses, and a collection of functional modules which can be configured as required. The functional modules communicate with each other using the backplane signal lines.

The functional modules defined in this document are used as vehicles for discussion of the bus protocol and need not be considered a constraint to logic design. For example, the designer might choose to design logic which interacts with the VMEbus in the manner described, but uses different on-board signals, or monitors other VMEbus signals. VMEbus boards might be designed to include any combination of the functional modules defined by this document.

The VMEbus functional structure can be divided into four categories. Each consists of a bus and its associated functional modules which work together to perform specific duties. Figure 1-2 shows the VMEbus functional modules and buses. Each category is briefly summarized below.

Data Transfer — Devices transfer data over the Data Transfer Bus (DTB), which contains data and address pathways and associated control signals. Functional modules called Masters, Slaves, Interrupters, and Interrupt Handlers use the DTB to transfer data between each other. Two other modules, called Bus Timer and IACK Daisy-Chain Driver also assist them in this process.

DTB Arbitration — Since a VMEbus system can be configured with more than one Master or Interrupt Handler, a means is provided to transfer control of the DTB between them in an orderly manner and to guarantee that only one Master controls the DTB at a given time. The Arbitration Bus modules (Requesters and Arbiter) coordinate the control transfer.

Priority Interrupt — The priority interrupt capability of the VMEbus provides a means by which devices can request services from an Interrupt Handler. These interrupt requests can be prioritized into a maximum of seven levels. Interrupters and Interrupt Handlers use the Priority Interrupt Bus signal lines.

Utilities — Periodic clocks, initialization, and failure detection are provided by the Utility Bus. It includes a general purpose system clock line, a system reset line, a system fail line, an AC fail line, and two serial lines. Utilities also include power and ground pins.

1.3 VMEbus Specification Diagrams

As aids to defining or describing VMEbus operation, several types of diagrams are used, including:

- a. Timing diagrams that show the timing relationships between signal transitions. The times involved will have minimum and/or maximum limits associated with them. Some of the times specified on these diagrams specify the behavior of the backplane interface logic, while other times specify the interlocked behavior of the functional modules.
- b. Sequence diagrams that are similar to timing diagrams but show only the interlocked timing relationships of the functional modules. These diagrams are intended to show a sequence of events, rather than to specify the times involved. For example, a sequence diagram might indicate that module A cannot generate signal transition B until it detects module C's generation of signal transition D.
- c. Flow diagrams that show a stream of events as they would occur during a VMEbus operation. The events are stated in words and result from interaction of two or more functional modules. The flow diagram describes VMEbus operations in a sequential manner and, at the same time, shows interaction of the functional modules.

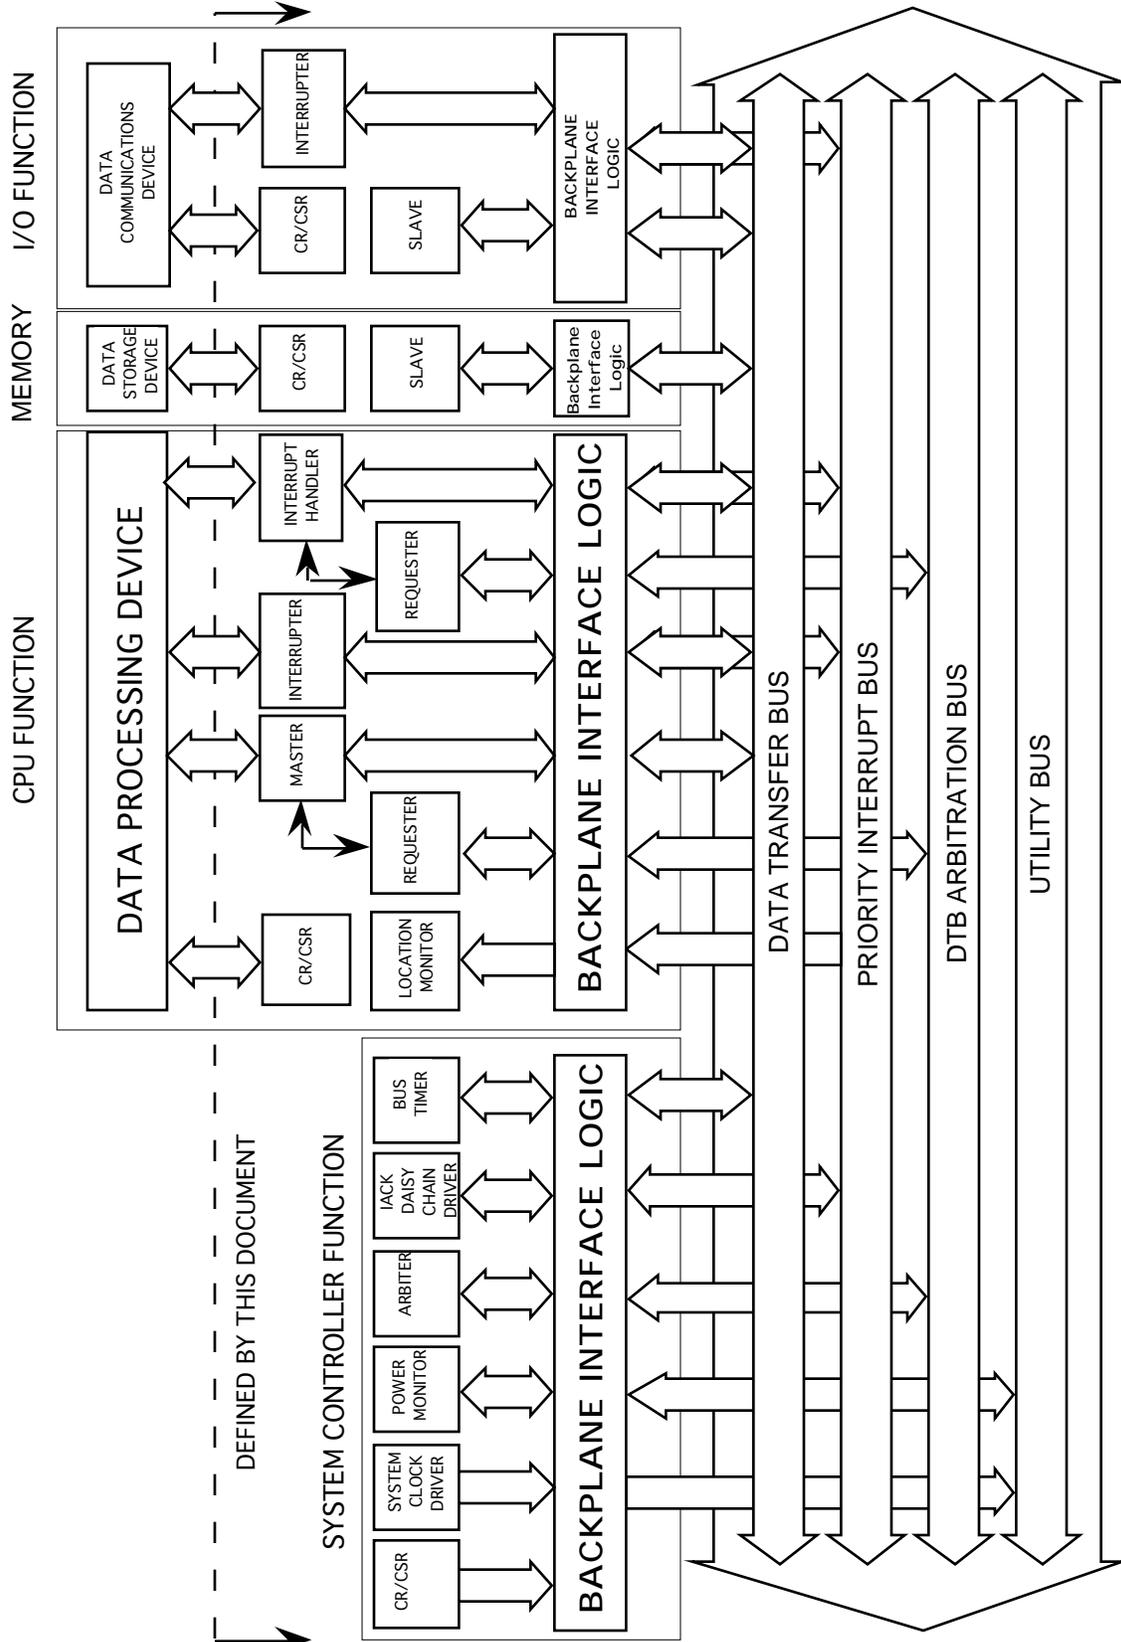


Figure 1 - 2 Functional Modules and Buses Defined by the Document

F

1.4 Specification Terminology

To avoid confusion, and to make very clear what the requirements for compliance are, many of the paragraphs in this document are labeled with keywords that indicate the type of information they contain. The keywords are listed below

RULE
RECOMMENDATION
SUGGESTION
PERMISSION
OBSERVATION

Any text not labeled with one of these keywords describes the VMEbus structure or operation. It is written in either a descriptive or a narrative style. These keywords are used as follows:

RULE chapter.number:

Rules form the basic framework of the VMEbus specification. They are sometimes expressed in text form and sometimes in the form of figures, tables, or drawings. All VMEbus rules **MUST** be followed to ensure compatibility between VMEbus designs. Rules are characterized by an imperative style. The upper-case words **MUST** and **MUST NOT** are reserved exclusively for stating rules in this document and are not used for any other purpose.

RECOMMENDATION chapter.number:

Wherever a recommendation appears, designers would be wise to take the advice given. Doing otherwise might result in some awkward problems or poor performance. While the VMEbus has been designed to support high performance systems, it is possible to design a VMEbus system that complies with all the rules, but has abysmal performance. In many cases, a designer needs a certain level of experience with the VMEbus in order to design boards that deliver top performance. Recommendations found in this document are based on this kind of experience and are provided to designers to speed their traversal of the learning curve.

SUGGESTION chapter.number:

In the VMEbus specification, a suggestion contains advice which is helpful but not vital. The reader is encouraged to consider the advice before discarding it. Some design decisions that need to be made in designing VMEbus boards are difficult until experience has been gained with the VMEbus. Suggestions are included to help a designer who has not yet gained this experience. Some suggestions have to do with designing boards that can be easily reconfigured for compatibility with other boards, or with designing the board to make the job of system debugging easier.

PERMISSION chapter.number:

In some cases a VMEbus rule does not specifically prohibit a certain design approach, but the reader might be left wondering whether that approach might violate the spirit of the rule, or whether it might lead to some subtle problem. Permissions reassure the reader that a certain approach is acceptable and will cause no problems. The upper-case word **MAY** is reserved exclusively for stating permissions in this document and is not used for any other purpose.

OBSERVATION chapter.number:

Observations do not offer any specific advice. They usually follow naturally from what has just been discussed. They spell out the implications of certain VMEbus rules and bring attention to things that might otherwise be overlooked. They also give the rationale behind certain rules, so that the reader understands why the rule must be followed.

1.4.1 Signal Line States

The VMEbus specification describes its protocol in terms of levels and transitions on bus lines.

A signal line is always assumed to be in one of two levels or in transition between these levels. Whenever the term “high” is used, it refers to a high TTL voltage level. The term “low” refers to a low TTL voltage level. A signal line is “in transition” when its voltage is moving between these levels. (See Chapter 6 for voltage thresholds used on the VMEbus.)

There are two possible transitions which can appear on a signal line, and these are called “edges”. A rising edge is the time during which a signal level makes its transition from a low level to a high level. The falling edge is the time during which a signal level makes its transition from a high level to a low level.

Some bus specifications prescribe maximum or minimum rise and fall times for these edges. The problem with doing this is that board designers have very little control over these times. If the backplane is heavily loaded, the rise and fall times will be long. If it is lightly loaded, these times might be short. Even if designers know what the maximum and minimum loading will be, they still need to spend time in the lab, experimenting to find out which drivers will provide the needed rise and fall times.

In fact, rise and fall times are the result of a complex set of interactions involving the signal line impedances of the backplane, its terminations, the source impedance of the drivers, and the capacitive loading of the signal line. In order to trade off all of these factors the board designer would have to study transmission line theory, as well as certain specific parameters of drivers and receivers which are not normally found in most manufacturers' data sheets.

Recognizing all of this, the VMEbus standard doesn't specify rise and fall times. Instead, it specifies the electrical characteristics for drivers and receivers and specifies the backplane design. It also tells designers how the worst case bus loading will affect the propagation delay of these drivers so that they can ensure that the VMEbus timing is met before building a board. If VMEbus designers follow these propagation delay guidelines, their boards will operate reliably with other VMEbus compatible boards under worst case conditions.

System performance is influenced by the rise time of one of the open collector control signals. To improve the performance of the system, the concept of a rescinding signal has been introduced. A rescinding signal is an open collector type output that is initially driven high and then released back to open collector mode within a short period of time.

1.4.2 Use of the Asterisk (*)

To help define usage, signal mnemonics have an asterisk suffix where required:

- a. An asterisk (*) following the signal name of signals which are level significant denotes that the signal is true or valid when the signal is low.
- b. An asterisk (*) following the signal name of signals which are edge significant denotes that the actions initiated by that signal occur on a high to low transition.

OBSERVATION 1.1:

The asterisk is inappropriate for the asynchronously running clock line SYSCLK. There is no fixed phase relationship between this clock line and other VMEbus timing.

1.4.3 Keyword Numbering

This standard is based on previous standards. To preserve keyword numbering, new keywords (rules, recommendation, suggestions, permissions, observations) have been numbered starting from the last available number. Hence, keywords as presented throughout this standard are not necessary in sequential order. See keyword cross reference index to determine where a specific keyword is referenced.

1.5 Protocol Specification

There are two layers of VMEbus protocol. The lowest VMEbus layer, called the backplane access layer, is composed of the backplane interface logic, the Utility Bus modules, and the Arbitration Bus modules. The VMEbus data transfer layer, is composed of the Data Transfer Bus and Priority Interrupt Bus modules. Figure 1-2 shows this layering.

OBSERVATION 1.2:

The signal lines used by the data transfer layer modules form a special class because they are driven by different modules at different times. They are driven with line drivers that can be turned on and off at each board based upon signals generated in the backplane access layer. It is very important that their turn-on and turn-off times be carefully controlled to prevent two drivers from attempting to drive the same signal line to different levels. Special timing diagram notation is used in this document to specify their turn-on and turn-off times. It is shown in Figure 1-3.

There are two basic kinds of protocol used on the VMEbus: closed loop protocols and open loop protocols. Closed loop protocols use interlocked bus signals while open loop protocols use broadcast bus signals.

1.5.1 Interlocked Bus Signals

An interlocked bus signal is sent from a specific module to another specific module. The signal is acknowledged by the receiving module. An interlocked relationship exists between the two modules until the signal is acknowledged.

For example, an Interrupter can send an interrupt request which is answered later with an interrupt acknowledge signal (no time limit is prescribed by the VMEbus specification). The Interrupter doesn't remove the interrupt request until the Interrupt Handler acknowledges it.

Interlocked bus signals coordinate internal functions of the VMEbus system, as opposed to interacting with external stimuli. Each interlocked signal has a source module and a destination module within the VMEbus system.

The address strobe and data strobes are especially important interlocking signals. They are interlocked with the data transfer acknowledge and bus error signals and coordinate the transfer of addresses and data which are the basis for all information flow between modules in the data transfer layer.

1.5.2 Broadcast Bus Signal

A module generates a broadcast signal in response to an event. There is no protocol for acknowledging a broadcast signal. Instead, the broadcast is maintained for a minimum specified time, long enough to assure that all appropriate modules detect the signal. Broadcast signals might be activated at any time, irrespective of any other activity taking place on the bus. They are each sent over a dedicated signal line. Some examples are the system reset and AC failure lines. These signal lines are not sent to any specific module, but announce special conditions to all modules.

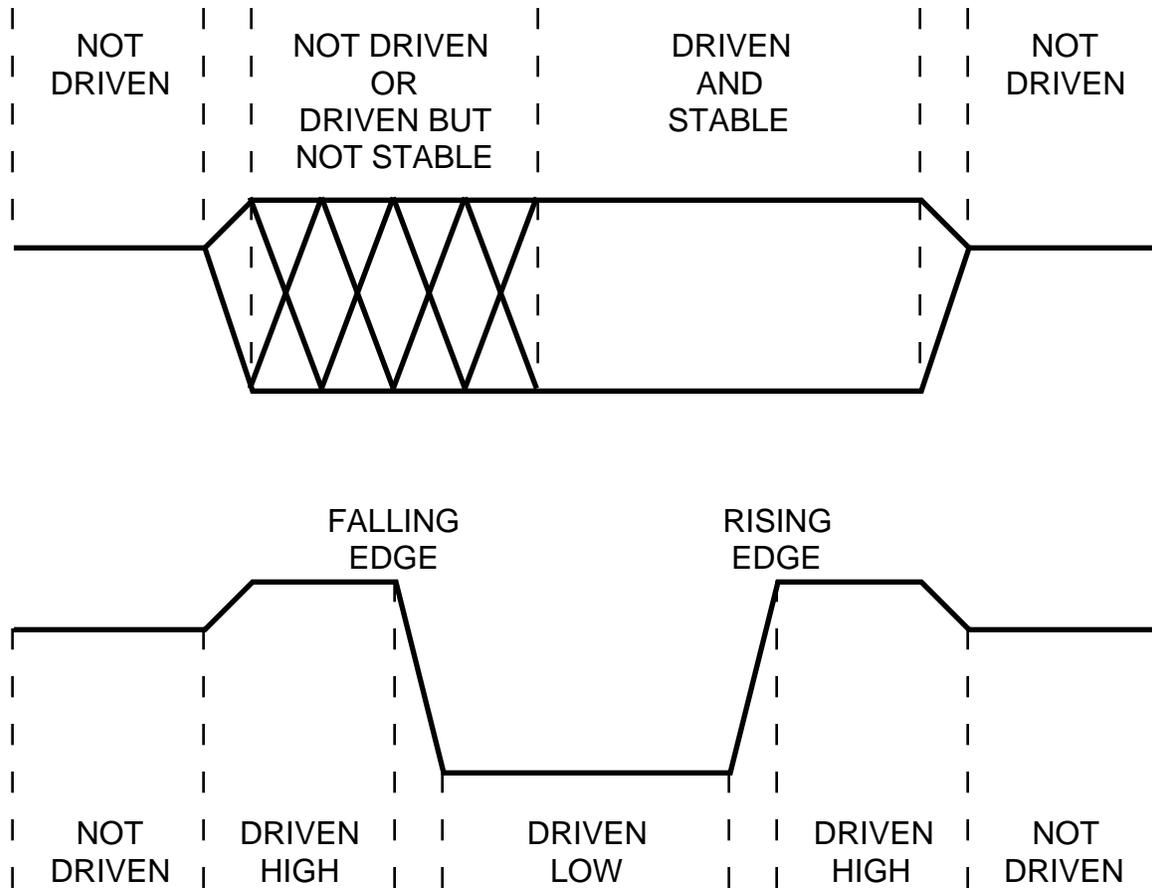


Figure 1 - 3 Signal Timing Notation

1.6 System Examples and Explanations

A protocol specification describes, in detail, the behavior of the various functional modules. It discusses how a module responds to a signal without saying where the signal came from. Because of this, a protocol specification does not give the reader a complete picture of what is going on over the bus. To help the reader, the VMEbus specification provides examples of typical VMEbus operations. Each example shows one possible sequence of events; other sequences are also possible. In providing these examples, there is the danger that readers will assume that the sequence shown in the example is the only legal one. To help readers avoid this trap, all examples are given in a narrative style, using the present tense. This is in contrast to the imperative style used when giving rules for compliance with the VMEbus specification.

Chapter 2

Data Transfer Bus

2.1 Introduction

This standard describes an asynchronous parallel Data Transfer Bus (DTB). Figure 2-1 shows a typical VMEbus system, including all of the DTB functional module types. Masters use the DTB to select storage locations provided by Slaves and to transfer data to or from those locations. Some Masters and Slaves use all of the DTB lines while others use only a subset.

Location Monitors monitor data transfers between Masters and Slaves. When an access is done to one of the byte location(s) that it monitors, a Location Monitor generates an on-board signal. For example, it might signal its on-board processor by means of an interrupt request. In such a configuration, if processor board A writes into a location of the global VMEbus memory that is monitored by processor B's Location Monitor, processor B will be interrupted.

After a Master initiates a data transfer cycle it waits for the responding Slave to respond before finishing the cycle. The asynchronous definition of the VMEbus allows a Slave to take as long as it needs to respond. If a Slave fails to respond because of some malfunction, or if the Master accidentally addresses a location where there is no Slave, the Bus Timer intervenes, allowing the cycle to be terminated.

2.1.1 Enhancements

This standard is based on VMEbus Rev C.1 and includes all the capabilities described in that standard. In addition the following enhancements have been added. For faster and wider transfers, the multiplexed block transfer (MBLT) feature is included in this standard. MBLT provides 64-bit data transfers by using both the data and address lines. MBLT protocols use the same basic asynchronous protocol as in D08(E0), D16 and D32 transfers.

Additionally, 64-bit addressing modes has been added to the existing A16, A24 and A32 modes. As memory densities increase, a single VME board could use the complete A32 address range. A64 also provides support for RISC and CISC processors with full A64 addressing.

For enhanced support of P1/J1 only systems (e.g. 3U VMEbus cards), a 40-bit address mode has been included. This mode provides access to a one terabyte address space using only the P1/J1 connector. A multiplexed data cycle (MD32) has been added to allow 32 bit transfer on P1 only. MD32 uses the data bus and the address bus for 32 bit transfers. A block transfer mode (A40BLT) allows 8, 16 and 32 bit transfers. These modes are grouped naturally with the A24 and A16 address modes. Although intended for P1/J1 only systems nothing precludes its use in full P1/P2 systems (6U VMEbus cards).

Five lock commands were added to provide multiport resource locking. This feature better supports the RISC and CISC microprocessors with a lock line output capability. A series of operations can now be performed on a resource without interference from another processor.

The Configuration ROM / Control & Status Registers (CR/CSR) provides processor, software and operating system independence for initialization, test and configuration of a board plugged into a VME backplane. The board's manufacturer and identification are also contained in the ROM area. A board's Control and Status Registers (CSR) can be implemented in the most expedient method deemed necessary by the board designer. Use of jumpers and other manual configuration are no longer needed.

VME cycles are defined to have either one phase or two phases. Single phase cycles are typically writes where both address and data are presented on the bus at the same time. Read cycles are divided into two phases, where the address is first broadcast to all Slaves followed by a data phase where the selected Slave presents data to the bus for the Master to read. For D08(E), D08(O), D16, D32 and MD32 reads there is an address broadcast phase followed by a data phase with a single data transfer. In BLT (except the first write cycle), MBLT and A40BLT cycles there is an address broadcast phase followed by a data phase with multiple data transfers. All A40 and A64 transactions also have two phases. Except for address only cycles (without handshake), all VME transfers end with one or more combinations of DTACK*, RETRY*, and BERR* being asserted low. Note the Rules 2.61 through 2.65 and 2.67 are not used in this document to avoid confusion with other VMEbus specifications.

2.2 Data-Transfer-Bus Lines.

The Data Transfer Bus lines are grouped into three categories:

Addressing Lines	Data Lines	Control Lines
A[31..1]	D[31..0]	AS*
D[31..0]	A[31..1]	WRITE*
LWORD*	LWORD*	BERR*
AM[5..0]		DTACK*
DS0*		DS0*
DS1*		DS1*
		RETRY*

OBSERVATION 2.1:

The two data strobes (DS0* and DS1*) serve a dual function:

- The levels of these two data strobe lines are used to select which byte(s) are accessed.
- The edges of the data strobes are also used as timing signals which coordinate the transfer of data between the Master and Slave.

OBSERVATION 2.86:

A[31..1] and LWORD* serve a dual function:

- During the address broadcast phase of all cycles, all or some of these signals are used to carry addressing information.
- During the data transfer phase of 64-bit block transfers, all of these signals are used to carry data.
- During the data transfer phase of A40 and A40BLT cycles, some of these signals may be used to carry data.

OBSERVATION 2.87:

The data lines D[31..0] serve a dual function:

- During the data transfer of all cycles, all or some of these signals are used to carry data.
- During the address broadcast phase of cycles where A40 and A64 address modes are used, all or some of these signals are used to carry addressing information.

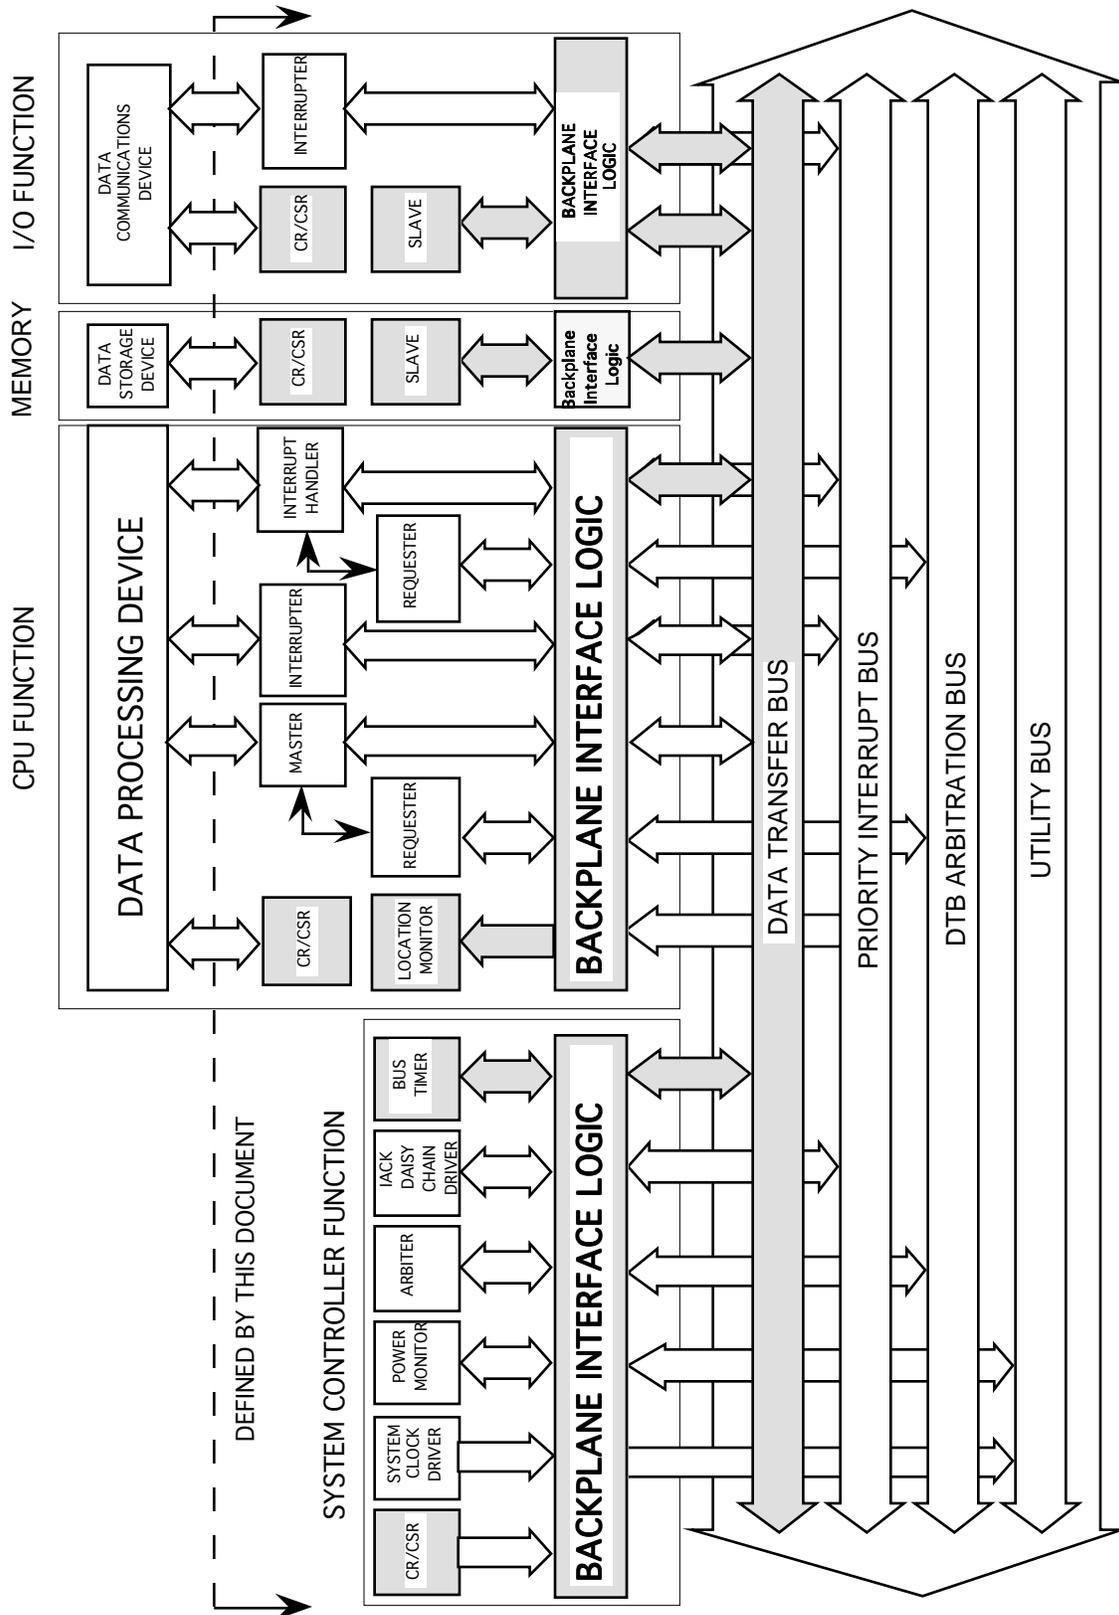


Figure 2 - 1 Data Transfer Bus Functional Block Diagram

2.2.1 Addressing Lines

The smallest addressable unit of storage is a byte location. Each byte location is assigned a unique binary address. Each byte location can be assigned to one of eight categories, according to the least significant bits of its address. (See Table 2-1)

Table 2 - 1 The Eight Categories Of Byte Locations

Category in a 4-Byte Group	Category in an 8-Byte Group	Byte Address
Byte(0)	Byte(0)	XXXXXX.....XXXXXX000
Byte(1)	Byte(1)	XXXXXX.....XXXXXX001
Byte(2)	Byte(2)	XXXXXX.....XXXXXX010
Byte(3)	Byte(3)	XXXXXX.....XXXXXX011
Byte(0)	Byte(4)	XXXXXX.....XXXXXX100
Byte(1)	Byte(5)	XXXXXX.....XXXXXX101
Byte(2)	Byte(6)	XXXXXX.....XXXXXX110
Byte(3)	Byte(7)	XXXXXX.....XXXXXX111

Table 2 - 28 Address Alignment on Bus

Address Mode	D[31..24]	D[23..16]	D[15..8]	D[7..0]	A[31..24]	A[23..16]	A[15..8]	A[7..1]
A16							A[15..8]	A[7..1]
A24						A[23..16]	A[15..8]	A[7..1]
A32					A[31..24]	A[23..16]	A[15..8]	A[7..1]
A40			A[31..24]	A[39..32]		A[23..16]	A[15..8]	A[7..1]
A64	A[63..56]	A[55..48]	A[47..40]	A[39..32]	A[31..24]	A[23..16]	A[15..8]	A[7..1]

Notes:

1. Byte 0 is selected by DS1* and Byte 1 is selected by DS0*
2. Note that in A40 mode, Byte(1) of the 40-bit address uses D[15..08]. This reduces the cost of implementation of mixed A64 and A40 bus capabilities.

A set of byte locations whose address differs only in the two least significant bits, is referred to as a 4-byte group or a Byte(0-3) group. Some or all of the bytes in a 4-byte group can be accessed simultaneously during DTB cycles.

A set of byte locations whose address differs only in the three least significant bits, is referred to as an 8-byte group or a Byte(0-7) group. All of the bytes in an 8-byte group are always accessed simultaneously. The 8-byte group is signaled by unique AM codes.

Assignment of the address bytes and bits for A16, A24, A32, A40 and A64 is shown in Table 2 - 28.

When accessing a 4-byte group, Masters use address lines A[63..2] or a subset of these lines to select which 4-byte group is accessed. Four additional lines, DS1*, DS0*, A1 and LWORD*, are used during the address broadcast phase to select which byte locations within the 4-byte group will be accessed during the data transfer. Using these four lines, a Master can access 1, 2, 3 or 4 byte locations simultaneously as shown in Table 2-2.

OBSERVATION 2.2:

In cycles where both data strobes are driven low, one data strobe might go low slightly after the other. In this case the signal levels indicated in Table 2-2 are the final levels.

OBSERVATION 2.3:

Given the 4 signal line levels shown in Table 2-2, there are 16 possible combinations of levels. Of these 16, there are two illegal combinations that are not used (see Table in RULE 2.1a).

RULE 2.1a:

Masters **MUST** ensure that none of the following combinations are generated during the address phase:

DS1*	DS0*	A1	LWORD*
high	low	high	low
low	high	high	low

PERMISSION 2.1:

Masters that generate Byte(1-2) Read or Byte(1-2) Write cycles **MAY** generate either of the two combinations described in RULE 2.1a briefly as transition states (i.e. while one data strobe has fallen but the other has not).

OBSERVATION 2.4:

Whenever a Master drives LWORD* low and A1 high it drives both data strobes low. (Any other combination is illegal.) VMEbus board designers can take advantage of this to simplify the logic on Slaves.

PERMISSION 2.2:

To simplify the required logic, Slaves which respond to Byte(1-2) Read and Byte(1-2) Write cycles **MAY** be designed without logic to distinguish between these cycles and the two illegal cycles described in RULE 2.1a.

As will be explained in paragraph 2.2.2, Masters use special address modifier codes to indicate that they wish to access an 8-byte group. Address lines A[63..3] or a subset of these lines are used during the address broadcast phase to select which 8-byte group will be accessed during the data transfer. DS0*, DS1*, A1 and LWORD* lines are not used to select byte-locations when selecting an 8-byte group. During the MBLT data transfer phases, LWORD* and A[31..1] are used to carry data.

Also special address modifier codes are used to indicate A40 transfer modes. During A40 modes, A[15..1] and LWORD* are used along with D[15..0] to transfer a 4-byte group.

RULE 2.69:

Whenever a Master identifies a cycle as an access to an 8-byte group by using one of the 64-bit block transfer AM codes (see 2.2.2), it **MUST** drive DS0*, DS1*, LWORD*, A1 and A2 low during the address phase.

Table 2 - 2 Signal Levels During Data Transfers
Used To Select Which Byte Location(s) Are Accessed During A Data Transfer

Type of cycle	DS1*	DS0*	A1	LWORD*	A2
Address-Only with No Handshake	high	high	Note 1	Note 1	high or low
Single even byte transfers					
Byte(0) Read or Write	low	high	low	high	high or low
Byte(2) Read or Write	low	high	high	high	high or low
Single odd byte transfers					
Byte(1) Read or Write	high	low	low	high	high or low
Byte(3) Read or Write	high	low	high	high	high or low
Double byte transfers					
Byte(0-1) Read or Write	low	low	low	high	high or low
Byte(2-3) Read or Write	low	low	high	high	high or low
Quad byte transfers					
Byte(0-3) Read or Write	low	low	low	low	high or low
Single byte block transfers					
Single Byte Block Read or Write	<-----Note 2 ----->			high	high or low
Double byte block transfers					
Double Byte Block Read or Write	low	low	Note 3	high	high or low
Quad byte block transfers					
Quad Byte Block Read or Write	low	low	low	low	high or low
Single byte RMW transfers					
Byte(0) Read-Modify-Write	low	high	low	high	high or low
Byte(1) Read-Modify-Write	high	low	low	high	high or low
Byte(2) Read-Modify-Write	low	high	high	high	high or low
Byte(3) Read-Modify-Write	high	low	high	high	high or low
Double byte RMW transfers					
Byte(0-1) Read-Modify-Write	low	low	low	high	high or low
Byte(2-3) Read-Modify-Write	low	low	high	high	high or low
Quad byte RMW transfers					
Byte(0-3) Read-Modify-Write	low	low	low	low	high or low

Table 2 - 2 Signal Levels During Data Transfers Used To Select Which Byte Location(s) Are Accessed During A Data Transfer (continued)

Type of cycle	DS1*	DS0*	A1	LWORD*	A2
Unaligned transfers					
Byte(0-2) Read or Write	low	high	low	low	high or low
Byte(1-3) Read or Write	high	low	low	low	high or low
Byte(1-2) Read or Write	low	low	high	low	high or low
Eight Byte Access					
Byte(0-7)	low	low	low	low	low

Notes:

1. During Address-Only cycles, both data strobes are maintained high, but the A1 and LWORD* lines might be either high or low.
2. During single byte block transfers, the two data strobes are alternately driven low. Either data strobe might be driven low on the first transfer. If the first accessed byte location is Byte(0) or Byte(2), then DS1* is driven low first. If the first accessed byte location is Byte(1) or Byte(3), then DS0* is driven low first. A1 is valid only on the first data transfer (i.e. until the Slave drives DTACK* or BERR* low the first time) and might be either high or low depending upon which byte the single byte block transfer begins with. If the first byte location is Byte(0) or Byte(1), then A1 is low. If the first byte location is Byte(2) or Byte(3), then A1 is high.

An example of a single byte block transfer cycle which starts with Byte(2) is given below:

		DS1*	DS0*	A1	LWORD*
First data transfer	Byte (2)	low	high	high	high
	Byte (3)	high	low	X	X
	Byte (0)	low	high	X	X
	Byte (1)	high	low	X	X
Last data transfer	Byte (2)	low	high	X	X

X = high or low.

3. During a double byte block transfer, the data strobes are driven low on each data transfer. A1 is valid only on the first data transfer (i.e. until the Slave drives DTACK* or BERR* low the first time) and might be either high or low depending upon what double byte group the double byte block transfer begins with. If the first double byte group is Byte(0-1), then A1 is low. If the first double byte group is Byte(2-3), then A1 is high.

An example of a double byte block transfer cycle which starts with Byte(2-3) is given below:

		DS1*	DS0*	A1	LWORD*
First data transfer	Byte (2-3)	low	low	high	high
	Byte (0-1)	low	low	X	X
	Byte (2-3)	low	low	X	X
Last data transfer	Byte (0-1)	low	low	X	X

X = high or low.

2.2.2 Address Modifier Lines.

There are 6 address-modifier lines. The Master uses these lines to pass additional binary information to the Slave during data transfers. Table 2-3 lists all of the 64 possible address modifier (AM) codes and classifies each into one of three categories:

- Defined
- Reserved
- User defined

The defined address modifier codes can be further classified into five categories:

- a) A16 addressing AM codes, which indicate that address lines A[15..2] are being used to select a Byte(0-3) group.
- b) A24 addressing AM codes, which indicate that address lines A[23..2] are being used to select a Byte(0-3) group or address lines A[23..3] for selecting a Byte(0-7) group.
- c) A32 addressing AM codes, which indicate that address lines A[31..2] are being used to select a Byte(0-3) group or address lines A[31..3] for selecting a Byte(0-7) group.
- d) A40 addressing AM codes, which indicate that address lines A[23..2] and data lines D[15..0] are being used to select a Byte(0-3) group.
- e) A64 addressing AM codes, which indicate that address lines A[31..2] and data lines D[31..0] are being used to select a Byte(0-3) group or address lines A[31..3] and data lines D[31..0] for selecting a Byte(0-7) group.

RULE 2.2 :

Except for the user-defined codes, the codes defined in Table 2-3 MUST NOT be used for purposes other than those specified.

RULE 2.3

VMEbus Slaves boards MUST NOT respond to reserved address-modifier codes.

OBSERVATION 2.5:

Reserved address-modifier codes are for future enhancements. If Slave boards respond to these codes, incompatibilities might result at some future date, when the use of these codes is defined.

PERMISSION 2.3:

User-defined codes MAY be used for any purpose that board vendors or board users deem appropriate, (page switching, memory protection, Master or task identification, privileged access to resources, etc.).

Table 2 - 3 Address Modifier Codes

ADDRESS MODIFIER							FUNCTION
HEX CODE	5	4	3	2	1	0	
3F	H	H	H	H	H	H	A24 supervisory block transfer (BLT)
3E	H	H	H	H	H	L	A24 supervisory program access
3D	H	H	H	H	L	H	A24 supervisory data access
3C	H	H	H	H	L	L	A24 supervisory 64-bit block transfer (MBLT)
3B	H	H	H	L	H	H	A24 non privileged block transfer (BLT)
3A	H	H	H	L	H	L	A24 non privileged program access
39	H	H	H	L	L	H	A24 non privileged data access
38	H	H	H	L	L	L	A24 non privileged 64-bit block transfer (MBLT)
37	H	H	L	H	H	H	A40BLT
36	H	H	L	H	H	L	Reserved
35	H	H	L	H	L	H	A40 lock command (LCK)
34	H	H	L	H	L	L	A40 access
33	H	H	L	L	H	H	Reserved
32	H	H	L	L	H	L	A24 lock command (LCK)
31	H	H	L	L	L	H	Reserved
30	H	H	L	L	L	L	Reserved
2F	H	L	H	H	H	H	Configuration ROM/Control & Status Register (CR/CSR)
2E	H	L	H	H	H	L	Reserved
2D	H	L	H	H	L	H	A16 supervisory access
2C	H	L	H	H	L	L	A16 lock command (LCK)
2B	H	L	H	L	H	H	Reserved
2A	H	L	H	L	H	L	Reserved
29	H	L	H	L	L	H	A16 non-privileged access
28	H	L	H	L	L	L	Reserved
27	H	L	L	H	H	H	Reserved
26	H	L	L	H	H	L	Reserved
25	H	L	L	H	L	H	Reserved
24	H	L	L	H	L	L	Reserved
23	H	L	L	L	H	H	Reserved
22	H	L	L	L	H	L	Reserved
21	H	L	L	L	L	H	Reserved
20	H	L	L	L	L	L	Reserved

L = low-signal level

H = high-signal level

Table 2 - 3 Address Modifier Codes (continued)

ADDRESS MODIFIER							FUNCTION
HEX CODE	5	4	3	2	1	0	
1F	L	H	H	H	H	H	User-defined
1E	L	H	H	H	H	L	User-defined
1D	L	H	H	H	L	H	User-defined
1C	L	H	H	H	L	L	User-defined
1B	L	H	H	L	H	H	User-defined
1A	L	H	H	L	H	L	User-defined
19	L	H	H	L	L	H	User-defined
18	L	H	H	L	L	L	User-defined
17	L	H	L	H	H	H	User-defined
16	L	H	L	H	H	L	User-defined
15	L	H	L	H	L	H	User-defined
14	L	H	L	H	L	L	User-defined
13	L	H	L	L	H	H	User-defined
12	L	H	L	L	H	L	User-defined
11	L	H	L	L	L	H	User-defined
10	L	H	L	L	L	L	User-defined
0F	L	L	H	H	H	H	A32 supervisory block transfer (BLT)
0E	L	L	H	H	H	L	A32 supervisory program access
0D	L	L	H	H	L	H	A32 supervisory data access
0C	L	L	H	H	L	L	A32 supervisory 64-bit block transfer (MBLT)
0B	L	L	H	L	H	H	A32 non privileged block transfer (BLT)
0A	L	L	H	L	H	L	A32 non privileged program access
09	L	L	H	L	L	H	A32 non privileged data access
08	L	L	H	L	L	L	A32 non privileged 64-bit block transfer (MBLT)
07	L	L	L	H	H	H	Reserved
06	L	L	L	H	H	L	Reserved
05	L	L	L	H	L	H	A32 lock command (LCK)
04	L	L	L	H	L	L	A64 lock command (LCK)
03	L	L	L	L	H	H	A64 block transfer (BLT)
02	L	L	L	L	H	L	Reserved
01	L	L	L	L	L	H	A64 single transfer access
00	L	L	L	L	L	L	A64 64-bit block transfer (MBLT)

L = low-signal level

H = high-signal level

RECOMMENDATION 2.1:

VMEbus users MAY tailor the use of the user-defined address-modifier codes to their own needs and decode them in a flexible way on Slave boards. Users can then configure the board to give any decoding required for their system.

2.2.3 Data Lines.

VMEbus systems can be built with a backplane configuration that provides either 16 data lines, D[15..0], or 32 data lines, D[31..0]. Backplane configurations that provide 16 data lines allow a Master to access up to 2 or 4 byte locations simultaneously, while those with 32 data lines allow it to access up to 8 byte locations simultaneously. When the Master has selected 1, 2, 3, 4 or 8 byte locations, using the method described in Section 2.2.1 and 2.2.2, it can transfer binary data between itself and those locations. Table 2-4 shows how the data lines and the address lines are used to access byte locations.

RULE 2.70

When accessing Byte(0-7), LWORD* MUST carry the least significant bit of Byte(3) and A7 MUST carry the most significant bit of Byte(3). When accessing Byte(0-3) in MD32 Mode, LWORD* MUST carry the least significant bit of Byte(1) and A7 MUST carry the most significant bit of Byte(1). (See Table 2-29.)

PERMISSION 2.4:

The data sender (Master for a write cycle; Slave for a read cycle) MAY drive data lines which are not used to transfer data.

Table 2 - 4 Use Of Data Lines To Move Data During Nonmultiplexed Data Transfers

During the following types of cycles...	the data lines are used to transfer data as shown below:			
	D[31..24]	D[23..16]	D[15..8]	D[7..0]
Address-Only	<----- no bytes transferred ----->			
Single even byte transfers				
Byte(0) Read, Write or RMW	Byte(0)			
Byte(2) Read, Write or RMW	Byte(2)			
Single odd byte transfers				
Byte(1) Read, Write or RMW	Byte(1)			
Byte(3) Read, Write or RMW	Byte(3)			
Double byte transfers				
Byte(0-1) Read, Write or RMW	Byte(0) Byte(1)			
Byte(2-3) Read, Write or RMW	Byte(2) Byte(3)			
Quad byte transfers				
Byte(0-3) Read, Write or RMW	Byte(0)	Byte(1)	Byte(2)	Byte(3)
Single byte block transfers				
Single Byte Block Read or Write	<----- Note 1 ----->			
Double byte block transfers				
Double Byte Block Read or Write	<----- Note 2 ----->			
Quad byte block transfers				
Quad Byte Block Read or Write	Byte(0)	Byte(1)	Byte(2)	Byte(3)
Unaligned transfers				
Byte(0-2) Read or Write	Byte(0)	Byte(1)	Byte(2)	
Byte(1-3) Read or Write		Byte(1)	Byte(2)	Byte(3)
Byte(1-2) Read or Write		Byte(1)	Byte(2)	

*Note: RMW is Read-Modify-Write

Table 2 - 4 Use Of Data Lines To Move Data During Nonmultiplexed Data Transfers (Continued)

Notes:

1. During single byte block transfers, data is transferred 8 bits at a time over D[15..8] or D[7..0]. One example of this is given below:

	D[31..24]	D[23..16]	D[15..8]	D[7..0]
First data transfer			Byte(2)	Byte(1)
			Byte(0)	Byte(3)
			Byte(2)	Byte(1)
Last data transfer			Byte(2)	Byte(3)

2. During double byte block transfers, data is transferred 16 bits at a time over D[15..0]. One example of this is given below:

	D[31..24]	D[23..16]	D[15..8]	D[7..0]
First data transfer			Byte(2)	Byte(3)
			Byte(0)	Byte(1)
			Byte(2)	Byte(3)
			Byte(0)	Byte(1)
			Byte(2)	Byte(3)
Last data transfer			Byte(0)	Byte(1)

Table 2 - 29 Use Of The Address And Data Lines For Multiplexed Data Cycles

A[31..24]	A[23..16]	A[15..8]	A[7..1]	D[31..24]	D[23..16]	D[15..8]	D[7..0]
LWORD*							
Multiplexed Quad Byte transfers (MD32)							
Byte (0-3) Read, Write or RMW							
		Byte(0)	Byte(1)			Byte(2)	Byte(3)
Multiplexed Quad Byte Block transfers (MD32)							
Byte (0-3) Read or Write							
		Byte(0)	Byte(1)			Byte(2)	Byte(3)
Multiplexed Eight Byte Block Transfer (MBLT)							
Byte (0-7) Read or Write							
	Byte(0)	Byte(1)	Byte(2)	Byte(3)	Byte(4)	Byte(5)	Byte(6) Byte(7)

2.2.4 Data Transfer Bus Control Lines.

The following signal lines are used to control the transfer of data over the data transfer lines:

AS*	Address Strobe
DS0*	Data Strobe Zero
DS1*	Data Strobe One
BERR*	Bus Error
DTACK*	Data Transfer Acknowledge
RETRY*	Retry
WRITE*	Read/Write

2.2.4.1 AS*

A falling edge on AS* informs all Slaves that the addressing information on all or some of the A[31..1] lines, LWORD* and AM(5..0) are stable and can be captured.

2.2.4.2 DS0* and DS1*

In addition to their function in selecting byte locations for data transfer, as described in 2.2.1, the data strobes also serve additional functions. On write cycles, the first falling edge of a data strobe indicates that the Master has placed valid data on the data bus. On read cycles, the first rising edge tells the Slave that it can remove valid data from the data bus.

When using the 64-bit or 40-bit addressing mode, the first falling edge of a data strobe indicates that the Master has placed valid addressing information on the data lines.

OBSERVATION 2.7:

VMEbus Masters are not permitted to drive either of the data strobes low before driving AS* low. However, due to the fact that AS* might be more heavily loaded on the backplane than the data strobes, Slaves and Location Monitors might detect a falling edge on a data strobe, before they detect the falling edge on AS*.

PERMISSION 2.5a:

Slaves and Location Monitors that do not support block transfer capability as described in 2.3.7 MAY be designed to capture AM[5..0], A[31..1] and LWORD* when they detect a falling edge on a data strobe instead of on the falling edge of AS*.

OBSERVATION 2.8:

VMEbus Slaves and Location Monitors that capture the address on the falling edge of the data strobe(s) need not monitor AS*.

OBSERVATION 2.9a:

To perform block read and write cycles as described in 2.3.7 or to take full advantage of address pipelining as described in 2.4.2, a Slave should capture the address on the falling edge of the data strobe (DS0* or DS1*) that falls first while AS* is low.

2.2.4.3 DTACK*.

On a Write cycle, the Slave drives DTACK* low to indicate that it has successfully received all of the data called for by the transfer type. On a read cycle, the Slave drives DTACK* low to indicate that it has placed data on the data bus. During A64, A40, MBLT, and A40BLT address phases, DTACK* is driven low by the Slave to indicate it can receive or send data in the data phase. For lock commands the Slave drives DTACK* low indicating it has accepted the lock command and has locked out the other ports.

2.2.4.3.1 Rescinding DTACK*.

Slaves are allowed to drive DTACK* high before releasing it at the end of a cycle, or during BLT, MBLT and A40BLT operations. Driving DTACK* high, versus letting DTACK* be pulled high by the backplane termination resistors provides less delay at the end of the cycle or data transfer and improves system performance. Figure 2 - 33 shows the timing parameters for Rescinding DTACK*. The timing parameter values are listed in Table 2 - 22.

PERMISSION 2.14:

A selected Slave MAY drive DTACK* high for a short time, as specified by Timing Parameter 30A at the end of a cycle and during BLT, MBLT and MD32 data transfer phases after the data strobes go high.

2.2.4.4 BERR*.

BERR* is driven low by the Slave or by the Bus Timer to indicate to the Master that the data transfer or address broadcast was unsuccessful or an error was detected. For example, when a Master tries to write to a location that contains Read-Only Memory, the responding Slave should drive BERR* low. When the Master tries to access a location that is not provided by any Slave, the Bus Timer drives BERR* low after a time-out period has expired.

RECOMMENDATION 2.2

Design Slaves to drive BERR* low when they detect an uncorrectable error in the data they retrieve from their internal storage during a read cycle. Also, design Slaves to drive BERR* low when a bus operation tries to write to devices that are read only.

PERMISSION 2.6:

VMEbus Slaves MAY be designed without a BERR* driver.

2.2.4.5 WRITE*.

WRITE* is a level-significant signal line. It is used by the Master to indicate the direction of data transfer operations. When WRITE* is driven low, the data transfer direction is from the Master to the Slave. When WRITE* is driven high, the data transfer direction is from the Slave to the Master.

2.2.4.6 RETRY*

RETRY* can be driven low by the responding Slave to indicate to the Master that the requested data transfer cannot be executed at this time. The Master that supports RETRY* should retry the data transfer at a later time. This operation may be used to break deadlock situations.

The Master waits for a short period of time before trying the operation again. In the interim, the bus can be used for other transactions.

See 2.3.13 for rules, recommendations, permissions and observations governing the use of the RETRY* line.

2.3 DTB Modules -- Basic Description.

The DTB protocols define the methods used to transfer data. Five addressing modes are provided:

- a. A16 (16-bit)
- b. A24 (24-bit)
- c. A32 (32-bit)
- d. A40 (40-bit)
- e. A64 (64-bit).

The capabilities of Masters, Slaves, Location Monitors and CR/CSRs are described by a list of mnemonics that show what cycle types they can generate, accept or monitor respectively. (This will be described in more detail later in this chapter.)

Sections 2.3.1 through 2.3.5 provide block diagrams for the five types of DTB functional modules: Master, Slave, Location Monitor, Bus Timer and CR/CSR.

RULE 2.7:

Output signal lines shown with solid lines in Figures 2 - 2 through 2 - 5 MUST be driven by the respective module, unless it would always drive them high.

OBSERVATION 2.11:

IF an output signal line is not driven,
THEN terminators on the backplane ensure that it is high.

RULE 2.8:

Input signal lines shown with solid lines in Figures 2-2 through 2-5 MUST be monitored and responded to in the appropriate fashion.

OBSERVATION 2.12:

RULES and PERMISSIONS for driving and monitoring signal lines shown with dotted lines in Figures 2-2 through 2-5 are given in Tables 2-5, 2-6, and 2-8.

2.3.1 Master.

The block diagram of the Master is shown in Figure 2 - 2. The dotted lines in the diagram show signals whose use varies among the various types of Masters. Table 2 - 5 specifies how the various types of Masters use these lines. Further information about how the various types of Masters drive the address line, the data lines, LWORD*, DS0*, DS1* and A1 is given in Tables 2 - 19, 2 - 20, 2 - 21.

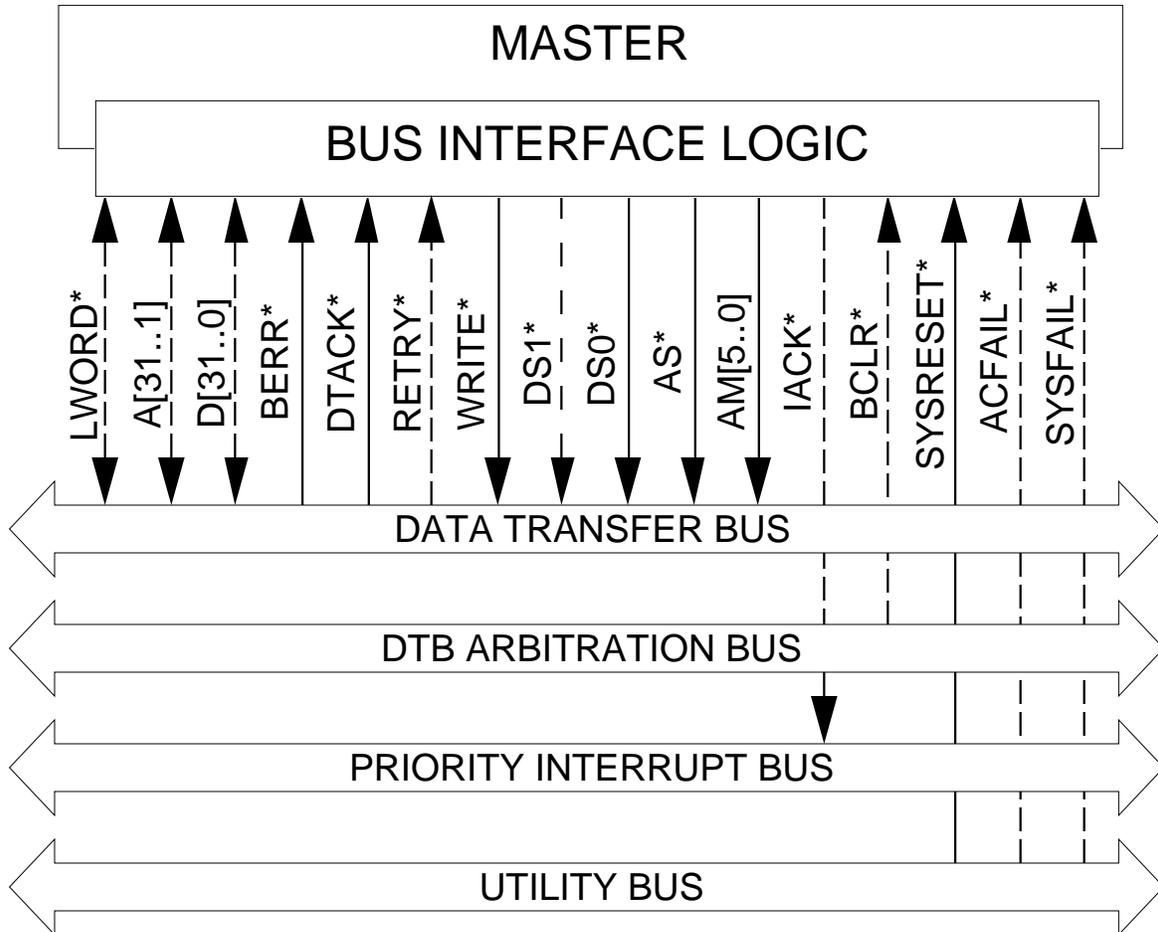


Figure 2 - 2 Block Diagram: Master

Table 2 - 5 RULES And PERMISSIONS That Specify The Use Of The Dotted Lines By The Various Types Of Masters

Type of Master	Use of Dotted Lines
D08(O)	MUST monitor and drive D[7..0] MAY monitor and drive LWORD* MAY monitor and drive D[31..8]
D08(E0)	MUST drive DS0* or DS1, but not both low on the same data transfer MUST monitor and drive D[15..0] MAY monitor and drive LWORD* MAY monitor and drive D[31..16]
D16 Single & BLT	MUST drive DS1* MUST monitor and drive D[15..0] MAY monitor and drive LWORD* MAY monitor and drive D[31..16]
D32 Single & BLT	MUST drive DS1* MUST drive LWORD* MUST monitor and drive D[31..0] MAY monitor LWORD*
MD32	MUST drive DS1* MUST monitor and drive LWORD* MUST monitor and drive D[15..0] MUST monitor and drive A[15..1]
MBLT	MUST drive DS1* MUST monitor and drive LWORD* MUST monitor and drive D[31..0] MUST monitor and drive A[31..1]
A16	MUST drive A[15..1] MAY drive A[31..16]
A24	MUST drive A[23..1] MAY drive A[31..24]
A32	MUST drive A[31..1]

Table 2 - 5 RULES And PERMISSIONS That Specify The Use Of The Dotted Lines By The Various Types Of Masters (Continued)

Type of Master	Use of Dotted Lines
A40 & A40BLT	MUST drive DS1* MUST drive A[23..1] MUST drive D[15..0] MUST drive LWORD*
A64	MUST drive DS1* MUST drive A[31..1] MUST drive D[31..0] MUST drive LWORD*
ALL	MAY monitor RETRY* MAY monitor BCLR* (see Section 3) MAY monitor ACFAIL* (see Section 5) MAY monitor SYSFAIL* (see Section 5) MUST NOT drive IACK* low

Notes:

1. The mnemonics A16, A24, A32, A40 and A64 are defined in Table 2 - 9.
2. The mnemonics D08(E0), D16, D32 and MD32 are defined in Table 2 - 10.
3. The mnemonics BLT, MBLT and A40BLT are defined in Table 2 - 11.

2.3.2 Slave.

The block diagram of the Slave is shown in Figure 2-3. The dotted lines in the diagram show signals whose use varies among the various types of Slaves. Table 2-6 shows how the various types of Slaves use these lines. Further information about how the various types of Slaves drive the data lines is given in Table 2-22.

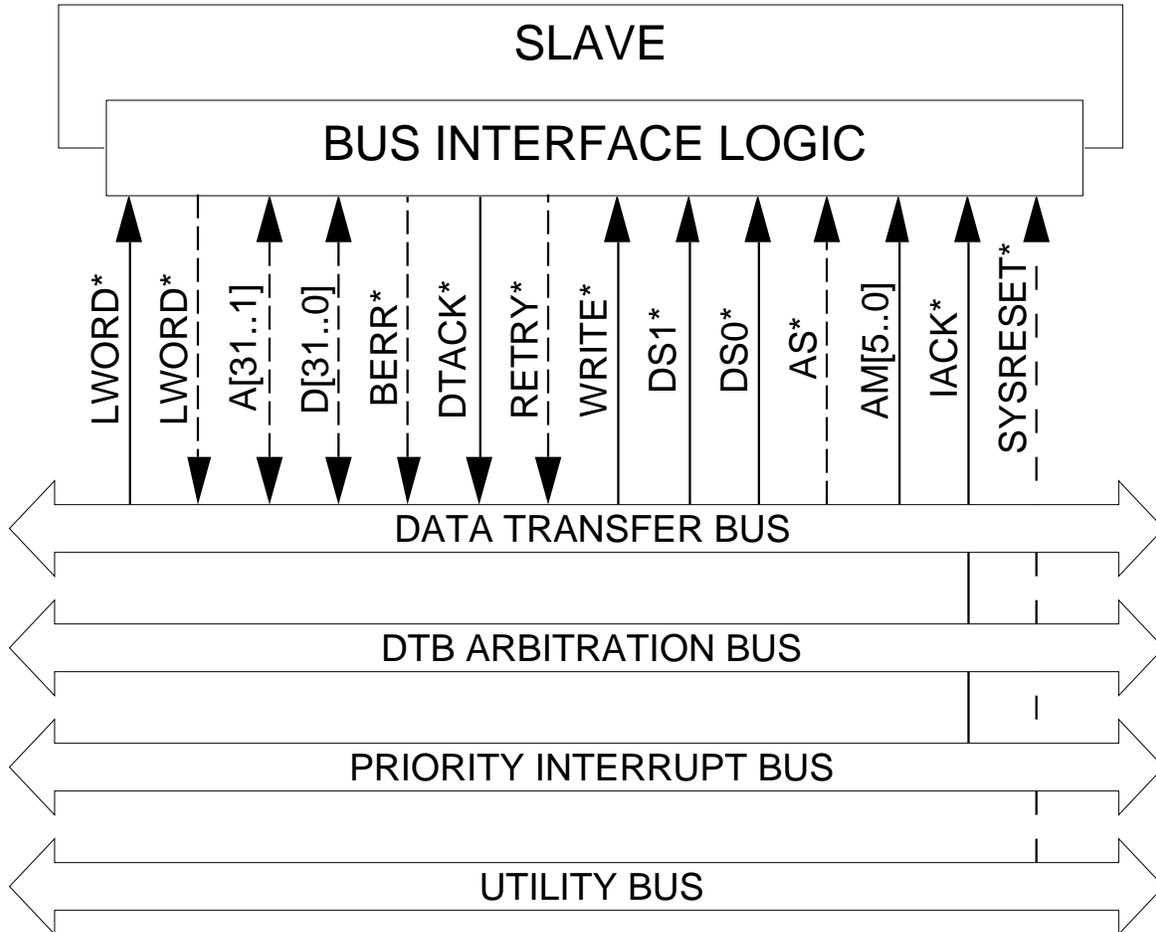


Figure 2 - 3 Block Diagram: Slave

Table 2 - 6 Slaves - RULES And PERMISSIONS That Specify The Use of The Dotted Lines By The Various Type Of Slaves

Type of Slave	Use of Dotted Lines
D08(O)	MUST monitor and/or drive D[7..0] MAY monitor AS* MAY monitor or drive D[31..8]
D08(EO) and D16	MUST monitor and/or drive D[15..0] MAY monitor AS* MAY monitor or drive D[31..16]
D32	MUST monitor and/or drive D[31..0] MAY monitor AS*
MD32	MUST monitor and/or drive A[15..1] MUST monitor and/or drive LWORD* MUST monitor and/or drive D[15..0] MUST monitor AS*
MBLT	MUST monitor and/or drive D[31..0] MUST monitor and/or drive A[31..1] MUST monitor AS* MUST monitor and/or drive LWORD*
A16	MUST monitor A[15..1] MAY monitor A[31..16]
A24	MUST monitor A[23..1] MAY monitor A[31..24]
A32	MUST monitor A[31..1]
A40 & A40BLT	MUST monitor D[15..0] MUST monitor A[23..1] MUST monitor AS*
A64	MUST monitor A[31..1] MUST monitor D[31..0] MUST monitor AS*
ALL	MAY drive RETRY* MAY drive BERR* MAY monitor SYSRESET*

Notes:

- (1) The mnemonics D08(O), D08(EO), D16 and D32 are defined in Table 2 - 10.
- (2) The mnemonics BLT, MBLT and A40BLT are defined in Table 2 - 11.
- (3) The mnemonics A16, A24, A32, A40 and A64 are defined in Table 2 - 9.

2.3.3 Bus Timer

The block diagram of the Bus Timer is shown in Figure 2 - 4. Bus Timers can be designed to drive BERR* low after various periods of time. Table 2-7 shows how the BTO() mnemonic is used to describe the various types of Bus Timers.

OBSERVATION 2.13:

The dashed DTACK* and BERR* lines shown in Figure 2 - 4 allow designers to implement a bus timer in one of two ways:

- a. To drive BERR* low when the first data strobe stays low for longer than the bus time-out period, regardless of the levels on the DTACK* and BERR* lines
- b. To drive BERR* low when the first data strobe stays low for longer than the bus time-out period, but only when both DTACK* and BERR* are high at the point of time-out.

Table 2 - 7 Use Of The BTO() Mnemonic
To Specify The Time-Out Period Of Bus Timers

The Following Mnemonic	When Applied to a	Means that it
BTO(x)	Bus Timer	drives BERR* low when the first data strobe stays low for longer than x microseconds, without DTACK* being asserted low

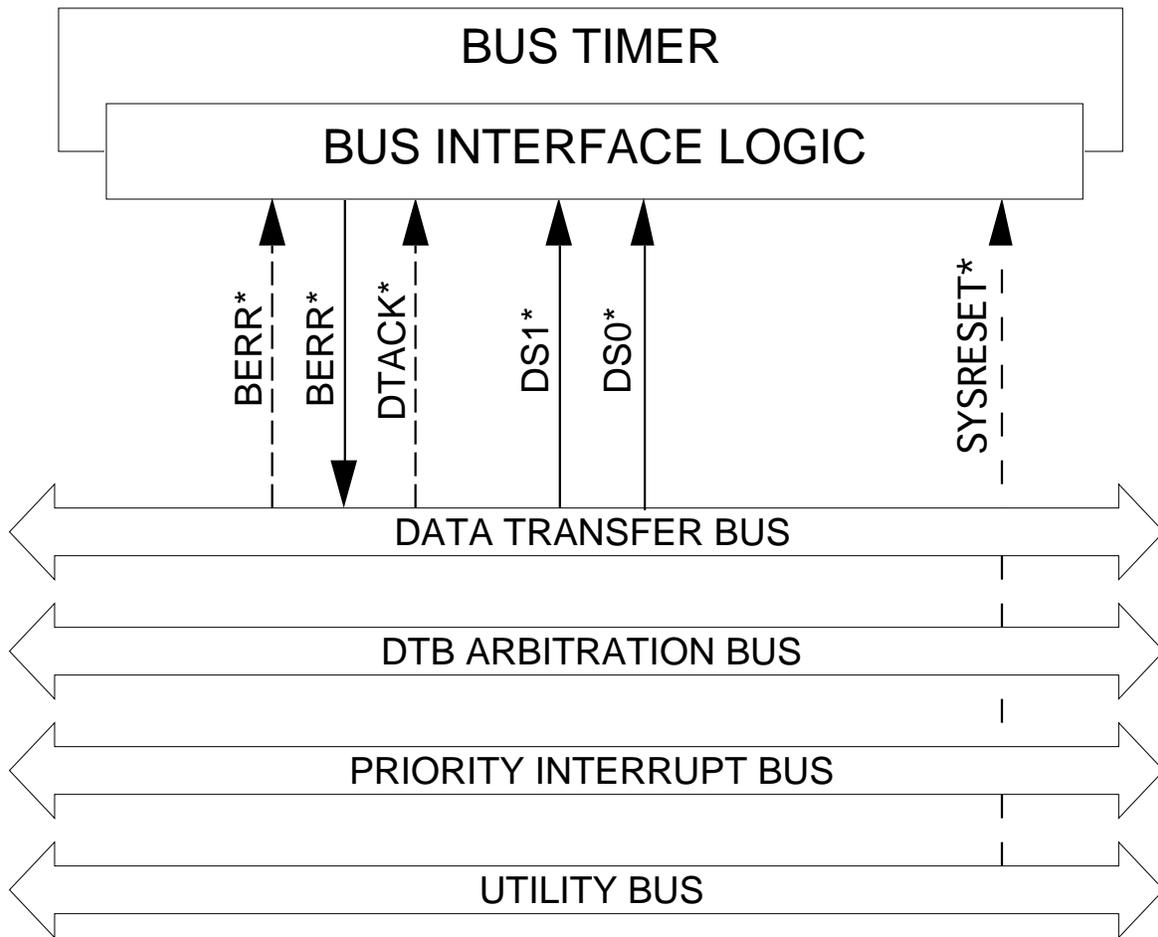


Figure 2 - 4 Block Diagram: Bus Timer

2.3.4 Location Monitor

The block diagram of the Location Monitor is shown in Figure 2-5. The dotted lines in the diagram show signals whose use varies among the various types of Location Monitors. Table 2-8 shows how the various types of Location Monitors use these lines.

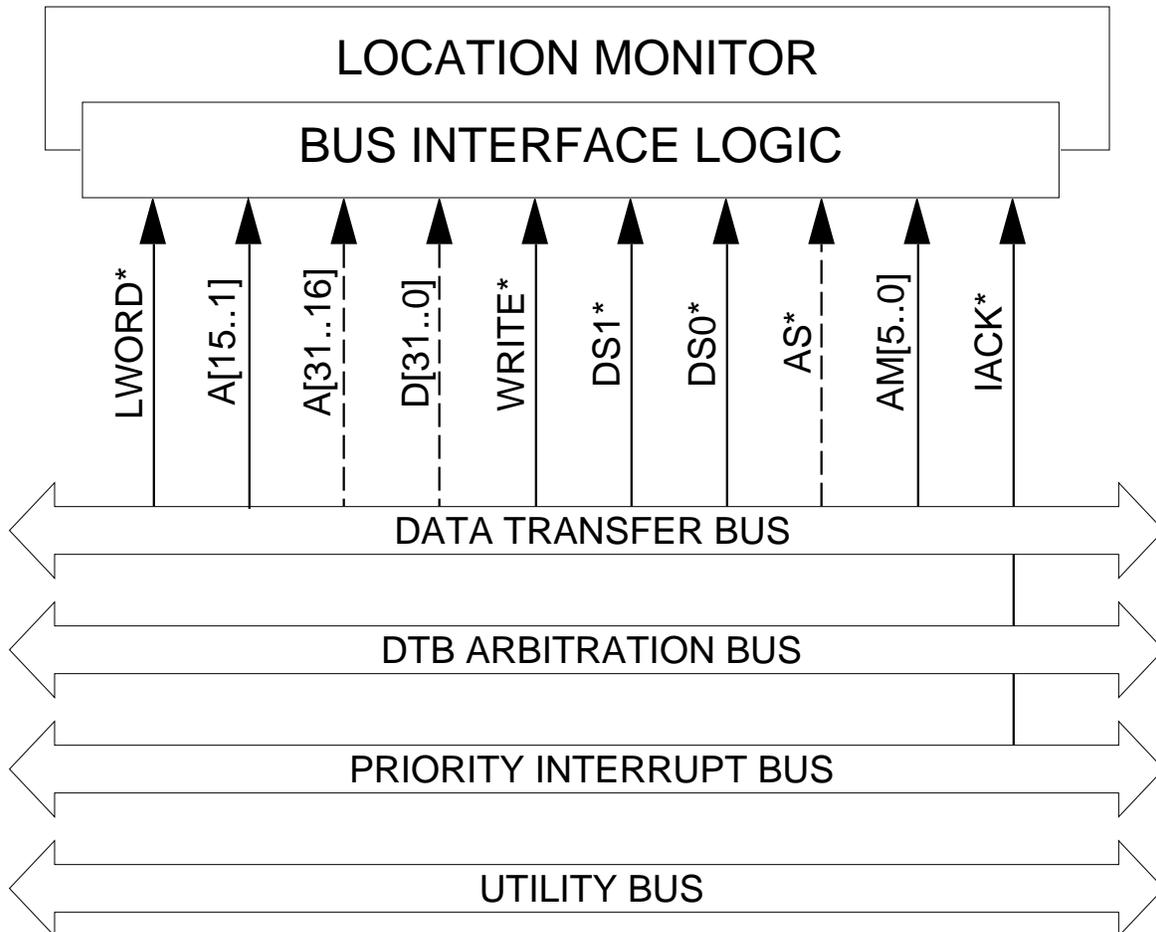


Figure 2 - 5 Block Diagram: Location Monitor

Table 2 - 8 Location Monitors - RULES And PERMISSIONS That Specify The Use Of The Dotted Lines By The Various Types Of Location Monitors

Type of Location Monitor	Use of dotted lines
A16	MUST monitor A[15..1] MAY monitor A[31..1]
A24	MUST monitor A[23..1] MAY monitor A[31..24]
A32	MUST monitor A[31..1]
A40	MUST monitor A[23..1] MUST monitor D[15..0] MUST monitor AS*
A64	MUST monitor A[31..1] MUST monitor D[31..0] MUST monitor AS*
ALL	MAY monitor AS*

Note: The mnemonics A16, A24, A32, A40 and A64 are defined in Table 2 - 9.

2.3.5 Addressing Phases and Modes.

Masters broadcast an address over the DTB at the beginning of each cycle. This broadcast address might be a 16-bit, a 24-bit, a 32-bit, a 40-bit or a 64-bit address, depending on the capabilities of the Master broadcasting it. 16-bit addresses are referred to as A16 addresses, 24-bit addresses are referred to as A24 addresses, 32-bit addresses are referred to as A32 addresses, 40-bit addresses are referred to as A40 addresses and 64-bit addresses are referred to as A64 addresses.

Table 2 - 9 shows the various mnemonics used to describe the addressing capabilities and how each is used to describe Masters, Slaves, Location Monitors and CR/CSR's.

The Master broadcasts an Address Modifier (AM) code along with each address to tell Slaves whether the address is A16, A24, A32, A40 or A64.

RULE 2.9:

Slave boards MUST decode all of the address modifier lines.

RULE 2.71:

When broadcasting a 64-bit address, A64 Masters MUST drive the least significant 32-bits of the address on LWORD* and A[31..1] and the most significant 32-bits of address on D[31..0]. Refer to Table 2-28.

OBSERVATION 2.14a :

Decoding all the address modifier lines allows a Slave to differentiate A16, A24, A32, A40 or A64 address mode as well as the specific operation within each address mode.

OBSERVATION 2.15a:

In addition to the five modes of addressing described here, there is a sixth mode that is used on interrupt acknowledge cycles (see Chapter 4). These interrupt acknowledge cycles can be distinguished from data transfer cycles by the fact that IACK* is low instead of high.

RULE 2.10:

Whenever a Master broadcasts an address, it MUST ensure that IACK* is high.

PERMISSION 2.7:

The Master MAY either drive IACK* high during the address broadcast or it MAY leave IACK* undriven (the bus terminators will then hold it high).

RULE 2.11:

Slaves MUST NOT respond to DTB cycles when IACK* is low.

OBSERVATION 2.88:

Systems may include a mixture of A16, A24, A32, A40 and A64 Masters and Slaves.

Table 2 - 9 Mnemonics That Specify Addressing Capabilities

The Following Mnemonic	When Applied to a	Means that it
A16	Master	MUST generate cycles with A16 addresses
	Slave	MUST accept cycles with A16 addresses
	Location Monitor	MUST monitor cycles with A16 addresses
A24	Master	MUST generate cycles with A24 & A16 addresses
	Slave	MUST accept cycles with A24 addresses
	Location Monitor	MUST monitor cycles with A24 addresses
	CR/CSR	MUST monitor cycles with A24 addresses
A32	Master	MUST generate cycles with A32, A24 & A16 addresses
	Slave	MUST accept cycles with A32 addresses
	Location Monitor	MUST monitor cycles with A32 addresses
A40	Master	MUST generate cycles with A40, A24 & A16 addresses
	Slave	MUST accept cycles with A40 addresses
	Location Monitor	MUST monitor cycles with A40 addresses
A64	Master	MUST generate cycles with A64, A32, A24 & A16 addresses
	Slave	MUST accept cycles with A64 addresses
	Location Monitor	MUST monitor cycles with A64 addresses

RULE 2.72

A64 Masters MUST include the A32, A24 and A16 capabilities.

RULE 2.73

A40 Masters MUST include the A24 and A16 capabilities.

RULE 2.74

A32 Masters MUST include the A24 and A16 capabilities.

RULE 2.75

A24 Masters MUST include the A16 capability.

SUGGESTION 2.6a

Do not assume that the above RULEs are known to the readers of product specifications. Rather, specify an A64 Master product as a "A64, A32, A24 and A16 Master". Specify an A32 Master product as a "A32, A24 and A16 Master". Similarly, specify an A24 Master product as a "A24 and A16 Master".

OBSERVATION 2.89

A64 and A32 Masters MAY include A40 capability for interoperability with 3U products.

2.3.6 Basic Data Transfer Capabilities.

There are five basic data transfer capabilities associated with the DTB: D08(O) (Odd byte only), D08(EO) (Even and Odd byte), D16, D32 and MD32. These capabilities allow flexibility when interfacing different types of processors and peripherals to the bus.

Eight-bit processors can be interfaced to the bus as D08(EO) Masters. Sixteen-bit processors can be interfaced to the bus as D16 Masters. The D16 Slave is useful for interfacing 16-bit memory devices or 16-bit I/O devices to the DTB.

Many existing peripheral chips have registers that are only 8 bits wide. While these chips often have several of these registers, they cannot provide the contents of two registers simultaneously when a D16 Master attempts to access two adjacent locations with a double-byte read cycle. These 8-bit peripheral ICs can be interfaced to the DTB as D08(O) Slaves, which provide only Byte(1) and Byte(3) locations and respond only to single-odd-byte accesses. This simplifies the D08(O) Slave's interface logic, since single-odd-byte accesses always take place over D[7..0].

RULE 2.76

D16 Slaves MUST include D08(EO) capability.

RECOMMENDATION 2.3

D16 Masters should include D08(EO) capability.

RULE 2.77

D32 and MD32 Slaves MUST include D16 and D08(EO) capabilities.

RECOMMENDATION 2.4

D32 and MD32 Masters should include D16 and D08(EO) capabilities.

RECOMMENDATION 2.5

MBLT Masters should include D32, D16 and D08(EO) capabilities.

SUGGESTION 2.9

Do not assume that the above RULEs are known to the readers of product specifications. Specify products that have the D32 capability as "D32, D16 and D08(EO)". Similarly, specify D16 products as "D16 and D08(EO)".

RULE 2.4:

D08(O), D08(EO), and D16 Slaves **MUST NOT** respond with a falling edge on DTACK* during a quad byte cycle if they do not have quad-byte capability.

RULE 2.5:

D08(O) and D08(EO) Slaves **MUST NOT** respond with a falling edge on DTACK* during a double-byte cycle if it does not have double-byte capability.

SUGGESTION 2.8a:

Design Slaves to respond with a falling edge on BERR* in the following situations:

- a. When a D08(O), D08(EO) or D16 only Slave is requested to do a quad-byte cycle.
- b. When a D08(O) or D08(EO) Slave is requested to do a double-byte cycle.
- c. When a D08(O), D08(EO) or D16 Slave is requested to do an unaligned transfer (that is, a triple-byte transfer or a double-byte Byte(1-2) transfer).
- d. When a D08(O) only Slave is requested to do an even single-byte transfer (that is, a single-byte Byte(0) or Byte(2) transfer).

Table 2 - 10 shows the various mnemonics used to describe the basic data transfer capabilities and how each is used to describe Masters, Slaves and Location Monitors.

Table 2 - 10 Mnemonics That Specify Basic Data Transfer Capabilities

The Following Mnemonic	When Applied to a	Means that it
D08(E0)	Master Slave Location Monitor	MUST generate the following cycles: MUST accept the following cycles: MUST monitor the following cycles: Single byte read cycles: Byte(0) Read Byte(1) Read Byte(2) Read Byte(3) Read Single byte write cycles: Byte(0) Write Byte(1) Write Byte(2) Write Byte(3) Write
D08(O)	Slave	MUST accept the following cycles: Single byte read cycles: Byte(1) Read Byte(3) Read Single byte write cycles: Byte(1) Write Byte(3) Write
D16	Master Slave Location Monitor	MUST generate the following cycles: MUST accept the following cycles: MUST monitor the following cycles: Double byte read cycles: Byte(0-1) Read Byte(2-3) Read Double byte write cycles: Byte(0-1) Write Byte(2-3) Write
D32	Master Slave Location Monitor	MUST generate the following cycles: MUST accept the following cycles: MUST monitor the following cycles: Quad byte read cycle: Byte(0-3) Read Quad byte write cycle: Byte(0-3) Write
MD32	Master Slave Location Monitor	MUST generate the following cycles: MUST accept the following cycles: MUST monitor the following cycles: Multiplexed Quad byte read cycle: Byte(0-3) Read Multiplexed Quad byte write cycle: Byte(0-3) Write

Note : (EO) is Even and Odd, (O) is Odd only.

2.3.7 Block Transfer Capabilities.

Masters often access several memory locations in ascending order. When this is the case, block transfer cycles are very useful. They allow the Master to provide a single address and then access data in that location and those at higher addresses without providing additional addresses.

When a Master initiates a block transfer cycle, the responding Slave latches the address into an on-board address counter. The Master, upon completing the first data transfer, (that is, driving the data strobes high) does not allow the address strobe to go high. Instead, it repeatedly drives the data strobe(s) low in response to data transfer acknowledgments from the Slave and transfers data to or from sequential memory locations in ascending order.

To access the next location(s), the Slave increments an on-board counter that generates the address for each transition of the data strobe(s).

Two types of block data transfer cycles are defined: basic block transfer cycles (BLT and A40BLT) and a multiplexed block transfer (MBLT) cycle. During standard block transfer cycles the Master can simultaneously access either 1, 2 or 4 byte locations in the course of each data transfer. During multiplexed block transfers, the Master simultaneously accesses 8 byte locations (MBLT) in each data transfer.

Both the data lines and the address lines are used to transfer data during MBLT and MD32 data transfers.

The block read cycle is very similar to a string of read cycles. Likewise, the block write cycle is very similar to a string of write cycles. The difference is that only the initial address is broadcast by the Master and the address strobe is held low during all of the data transfers.

The number of transfers in a block is limited to 256. The intent is to limit the length of time a block transfer can occupy the bus and prevent an arbitration cycle.

OBSERVATION 2.90:

As described in paragraph 2.2.2, the Master differentiates BLT, MBLT, and A40BLT cycles by driving special codes on the address modifier lines.

OBSERVATION 2.18a:

Block transfer cycles of indefinite length complicate the design of memory boards. Specifically, all block transfer Slaves (the one that responds and those that do not) need to latch the initial address and then increment the address counter on each bus transfer. All Slaves then have to decode the incremented address to determine whether the transfer has crossed a board boundary into their address range. While this is certainly possible, such address decoding typically limits access times of the Slave. The following RULEs were formulated to simplify the design of these Slaves and to permit faster access time.

RULE 2.12a:

D08(E0), D16, D32 and MD32 block transfer cycles (BLT) MUST NOT cross any 256 byte boundary.

RULE 2.78 :

MBLT cycles MUST NOT cross any 2048 byte boundary.

OBSERVATION 2.19a:

The above RULEs, which establish block transfer boundaries, limit the maximum length of block transfers. However, knowing that only A[7..1] or A[10..3] can change during the block transfer simplifies the design of block transfer Slaves. The upper address lines only have to be decoded once, at the beginning of the block transfer cycle, allowing faster access times on all subsequent data transfers.

OBSERVATION 2.20a:

In some cases it might be necessary to transfer a large block of data that crosses one or more block transfer boundaries. In such a case, when the Master that does the block transfer is designed to recognize the arrival at a block transfer boundary, it can drive AS* high for the required minimum period and then initiate another block transfer without the intervention of system software.

OBSERVATION 2.21 :

Control of the DTB cannot be transferred during block transfer cycles because AS* is held low through all of the data transfers and control of the DTB can only be transferred while AS* is high.

RULE 2.66:

Slaves that include block transfer capability MUST monitor AS* and capture the addressing information when they detect a falling edge on AS*.

OBSERVATION 2.91:

Address is stable from AS* low to DTACK* low. AS* will stay low when transitioning from an address phase to a data transfer phase.

OBSERVATION 2.92:

D32 BLT Masters and Slaves and MD32 BLT Masters and Slaves are not required to support D16 BLT and D08(E0) BLT. Likewise, D16 BLT Masters and Slaves are not required to support D08(E0) BLT.

RULE 2.79:

During the course of a single BLT cycle, a Master MUST NOT mix single-byte, double-byte and quad-byte data transfers.

Table 2 - 11 lists the mnemonics used to describe block transfer capabilities and how they are used to describe the various types of Masters, Slaves and Location Monitors.

OBSERVATION 2.93 :

A block transfer could span a range of addresses which may include one that is monitored by a Location Monitor. Therefore, when the mnemonics BLT and MBLT are applied to a Location Monitor, it means that the Location Monitor includes address counters that allow it to recognize when one of the locations it was configured to monitor was accessed during a block transfer cycle.

RULE 2.80 :

As shown in Table 2 - 11, a Master MUST NOT execute a block transfer cycle that accesses byte locations using unaligned transfers (see paragraph 2.3.9).

Table 2 - 11 Mnemonics That Specify Block Transfer Capabilities

The Following Mnemonic	When Applied to a	Means that it
BLT and A40BLT	D08(E0) Master D08(E0) Slave D08(E0) Loc Mon	MUST generate the following cycles: MUST accept the following cycles: MUST monitor the following cycles: Block read cycles: Single-Byte Block Read Block write cycles: Single-Byte Block Write
BLT and A40BLT	D16 Master D16 Slave D16 Loc Mon	MUST generate the following cycles: MUST accept the following cycles: MUST monitor the following cycles: Block read cycles: Double-Byte Block Read Block write cycles: Double-Byte Block Write
BLT	D32 Master D32 Slave D32 Loc Mon	MUST generate the following cycles: MUST accept the following cycles: MUST monitor the following cycles: Block read cycles: Quad-Byte Block Read Block write cycles: Quad-Byte Block Write
A40BLT	MD32 Master MD32 Slave MD32 Loc Mon	MUST generate the following cycles: MUST accept the following cycles: MUST monitor the following cycles: Multiplexed Quad Block read cycles: Quad-Byte Block Read Multiplexed Quad Block write cycles: Quad-Byte Block Write
MBLT	MBLT Master MBLT Slave MBLT Loc Mon	MUST generate the following cycles: MUST accept the following cycles: MUST monitor the following cycles: Multiplexed Block read cycles: Eight-Byte Block Read Multiplexed Block write cycles: Eight-Byte Block Write

2.3.8 Read-Modify-Write Capability

In multiprocessor systems which share resources such as memory and I/O, a method is needed to allocate these resources. One very important goal of this allocation algorithm is to ensure that a resource being used by one task cannot be used by another at the same time. The problem is best described by an example:

Two processors in a distributed processing system share a common resource (e.g., a printer). Only one processor can use the resource at a time. The resource is allocated by a bit in memory — i.e., if the bit is set, the resource is busy; if it is cleared, the resource is available. To gain use of the resource, processor A reads the bit and tests it to determine whether it is cleared. If the bit is cleared, processor A sets the bit to lock out processor B. This operation takes two data transfers: a read to test the bit, and a write to set the bit. However, a difficulty might arise if the bus is given to processor B between these two transfers. Processor B might then also find the bit cleared and assume the resource is available. Both processors will then set the bit in the next available cycle and attempt to use the resource.

This conflict is avoided by defining a Read-Modify-Write Cycle or Lock command which prevents transferring control of the DTB between the read portion and the write portion of the cycle. This cycle is very similar to a read cycle immediately followed by a write cycle. The difference is that the address strobe is held low during both transfers. This ensures that, unlike a read cycle followed by a write cycle, control of the DTB cannot be transferred during a Read-Modify-Write Cycle, as this is only possible while the address strobe is high.

The same applies for a resource that can be accessed locally and globally over the bus. Whenever a resource is being accessed by the bus, with AS* low, local access is to be locked out. With this type of operation, a global processor can be assured that when it runs a Read-Modify-Write cycle on a dual-ported resource it will be run atomically.

Table 2-12 lists the various mnemonics used to describe read-modify-write capabilities and how each is used to describe Masters, Slaves, and Location Monitors.

Table 2 - 12 The Mnemonic That Specifies Read-Modify-Write Capabilities

The Following Mnemonic	When Applied to a	Means that it
RMW	D08(E0) Master D08(E0) Slave D08(E0) Location Monitor	MUST generate the following cycles: MUST accept the following cycles: MUST monitor the following cycles: Single byte Read-Modify-Write cycles: Byte(0) Read-Modify-Write Byte(1) Read-Modify-Write Byte(2) Read-Modify-Write Byte(3) Read-Modify-Write
	D08(O) Slave	MUST accept the following cycles: Single byte Read-Modify-Write cycles: Byte(1) Read-Modify-Write Byte(3) Read-Modify-Write
	D16 Master	MUST generate the following cycles:
	D16 Slave	MUST accept the following cycles:
	D16 Location Monitor	MUST monitor the following cycles: Double byte Read-Modify-Write cycles: Byte(0-1) Read-Modify-Write Byte(2-3) Read-Modify-Write
	D32 Master D32 Slave D32 Location Monitor	MUST generate the following cycles: MUST accept the following cycles: MUST monitor the following cycles: Quad byte Read-Modify-Write cycles: Byte(0-3) Read-Modify-Write
	MD32 Master MD32 Slave MD32 Location Monitor	MUST generate the following cycles: MUST accept the following cycles: MUST monitor the following cycles: Multiplexed Quad byte Read-Modify-Write cycles: Byte(0-3) Read-Modify-Write

2.3.9 Unaligned Transfer Capability

Some 32-bit microprocessors store and retrieve data in an unaligned fashion. For example, a 32-bit value might be stored in four different ways, as shown in Figure 2 - 6.

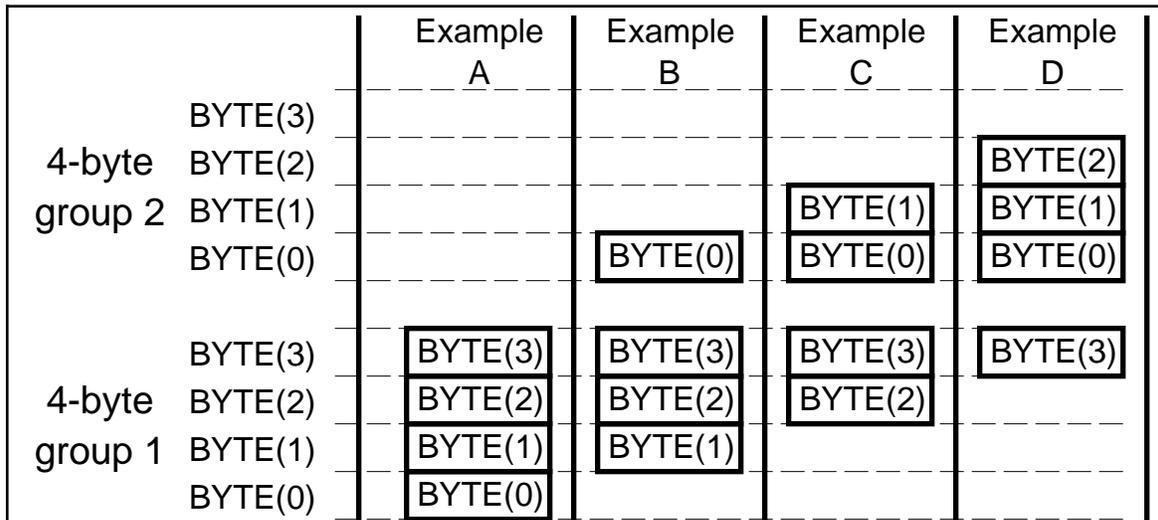


Figure 2 - 6 Four Ways That 32 Bits Of Data Might Be Stored In Memory

The Master can transfer the 32 bits of data using several different sequences of DTB cycles. For example, it can transfer the data one byte at a time, using four single-byte data transfers. However, a Master can accomplish the transfer much quicker by using one of the cycle sequences shown in Table 2 -13.

OBSERVATION 2.22:

The sequences shown in Table 2 -13 are typical of a Master that accesses the byte locations in ascending order. The DTB protocols do not require this.

As shown in Table 2 -13, each of these 32-bit transfers can be accomplished with a combination of single-byte and double-byte transfers. However, examples B and D require three bus cycles when this procedure is followed. Because of this, the DTB protocol also includes two triple-byte transfer cycles. When used in combination with a single-byte cycle, these triple-byte cycles allow data to be stored as shown in examples B and D using only two bus cycles.

Some 32-bit microprocessors also store and retrieve data 16 bits at a time, in an unaligned fashion, as shown in Figure 2 - 7.

Table 2 - 13 Transferring 32 Bits Of Data Using Multiple-Byte Transfer Cycles

Example	Cycle Sequences Used to Accomplish the Transfer	Data Bus Lines Used	Byte Locations Accessed (See Figure 2-6)
A	Quad-byte transfer	D[31..0]	Grp 1, Byte(0-3)
B	Single-byte transfer	D[7..0]	Grp 1, Byte(1)
	Double-byte transfer	D[15..0]	Grp 1, Byte(2-3)
	Single-byte transfer	D[15..8]	Grp 2, Byte(0)
	or		
	Triple-byte transfer	D[23..0]	Grp 1, Byte(1-3)
	Single-byte transfer	D[15..8]	Grp 2, Byte(0)
C	Double-byte transfer	D[15..0]	Grp 1, Byte(2-3)
	Double-byte transfer	D[15..0]	Grp 2, Byte(0-1)
D	Single-byte transfer	D[7..0]	Grp 1, Byte(3)
	Double-byte transfer	D[15..0]	Grp 2, Byte(0-1)
	Single-byte transfer	D[15..8]	Grp 2, Byte(2)
	or		
	Single-byte transfer	D[7..0]	Grp 1, Byte(3)
	Triple-byte transfer	D[31..8]	Grp 2, Byte(0-2)

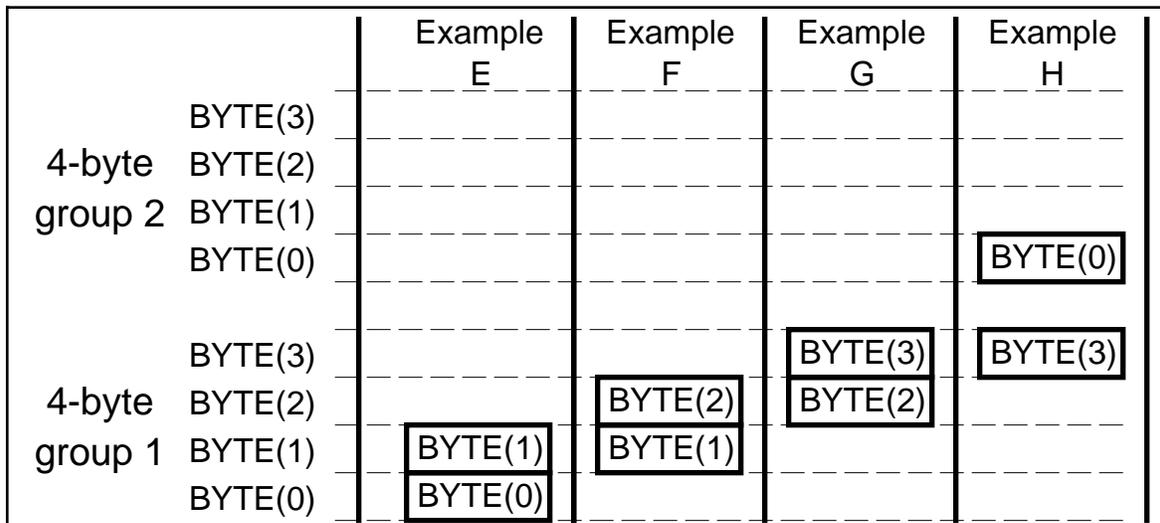


Figure 2 - 7 Four Ways That 16 Bits of Data Might Be Stored in Memory

The Master can transfer the 16 bits of data using several different sequences of DTB cycles as listed in Table 2 - 14.

OBSERVATION 2.23:

The sequences listed in Table 2 - 14 is typical of a Master that accesses the byte locations in ascending order. The VMEbus protocol do not require this.

As shown in Table 2 -13, the 16-bit transfer in example F can be accomplished with two single-byte transfers. However, this requires two bus cycles. Because of this, the DTB protocol also includes a double-byte transfer cycle that allows data to be stored as shown in example F using only one bus cycle.

OBSERVATION 2.24:

Since unaligned transfers make use of all 32 data lines, only D32 Masters and D32 Slaves can do unaligned transfers.

Table 2 -15 lists how the Unaligned Transfer mnemonic (UAT) is used to describe Masters, Slaves, and Location Monitors.

Table 2 - 14 Transferring 16 Bits Of Data Using Multiple-Byte Transfer Cycles

Example	Cycle Sequences Used to Accomplish the Transfer	Data Bus Lines Used	Byte Locations Accessed (See Figure 10)
E	Double-byte transfer	D[15..0]	Grp 1, Byte(0-1)
F	Single-byte transfer Single-byte transfer or Double-byte transfer	D[7..0] D[15..8] D[23..8]	Grp 1, Byte(1) Grp 1, Byte(2) Grp 1, Byte(1-2)
G	Double-byte transfer	D[15..8]	Grp 1, Byte(2-3)
H	Single-byte transfer Single-byte transfer	D[7..0] D[15..8]	Grp 1, Byte(3) Grp 2, Byte(0)

RECOMMENDATION 2.6: (used to be RULE 2.67)

D32 Masters, Slaves and Location Monitors should support UAT capability when appropriate.

RULE 2.6:

A Slave **MUST NOT** respond with a falling edge on DTACK* during an unaligned transfer cycle, if it does not have UAT capability.

Table 2 - 15 Mnemonic That Specifies Unaligned Transfer Capability

The Following Mnemonic	When Applied to a	Mean that it
UAT	D32 Master D32 Slave D32 Location Monitor	can generate the following cycles: can accept the following cycles can monitor the following cycles Triple-byte read cycles: Byte(0-2) Read Byte(1-3) Read Triple-byte write cycles: Byte(0-2) Write Byte(1-3) Write Double-byte read cycle: Byte(1-2) Read Double-byte write cycle: Byte(1-2) Write

2.3.10 Address-Only Capability.

The Address-Only cycle (ADO) on the DTB is not used to transfer data, but can be used to transfer special coded information. It begins as a typical DTB cycle, with the address, address modifier code, IACK* and LWORD* lines becoming valid and AS* falling after a setup time. However, the data strobes are never driven low. After holding the various lines strobed by AS* stable for a prescribed minimum period, the Master terminates the cycle without waiting for DTACK* or BERR* to go low. The ADO cycle is also the only type of DTB cycle that does not require a response to complete the cycle.

A second type of address-only cycle is ADDRESS-ONLY-WITH-HANDSHAKE (ADOH). It begins as a typical DTB cycle, with the address, address modifier code, IACK* and LWORD* lines becoming valid and AS* falling after a setup time. The data strobes are also driven low. After a Slave recognizes the address, it asserts DTACK*, BERR* or RETRY* low in response to the specific request. The WRITE* line may or may not be asserted during an ADOH cycle. The ADOH cycle is typically used for the lock commands.

Table 2 - 16 shows how the mnemonics ADO and ADOH are used to describe Masters and Slaves.

OBSERVATION 2.25:

Address-only cycles (ADO) can be used to enhance board performance by allowing a CPU board that is already in control of the bus to broadcast an address before it has determined whether or not that address selects a Slave on the bus. Broadcasting the address in this fashion allows Slaves to decode the address concurrently with the CPU board.

(Note this capability was included for historical reasons related to the use of certain CPUs in use at the time. This practice is not recommended.)

Table 2 - 16 Mnemonics That Specify Address Only Capability

The Following Mnemonic	When Applied to a	Means that it
ADO	Master	can generate Address-Only cycles.
ADOH	Master	can generate Address-Only -With-Handshake cycles
	Slaves	can accept Address-Only-With-Handshake cycles

RULE 2.68:

All Slaves **MUST** be designed to accommodate ADO cycles without loss of data or erroneous operation.

PERMISSION 2.15:

All Masters **MAY** generate address-only cycles (ADO) with no handshake.

OBSERVATION 2.94:

In all ADO cycles, neither of the two data strobes (DS0* and DS1*) are asserted low.

2.3.11 Lock Capability

The lock commands are Address-Only-Cycles-With-Handshake (ADOH) cycles that are used to lock the other ports(s) of multiported resources, where one port is on VMEbus. Each Slave's resource that is addressed with a lock command is to lock out all other accesses to that resource. The resource may be dual port memory, registers, special flags or control bits. Note that most Slaves have multiple resources. Only the resource that is addressed is to be locked and the other resources should remain unlocked. The extent of the locked resource is implementation dependent.

Each of the lock commands has an address phase with the same protocol as MBLT's address phase. The address (A16, A24, A32, A40 or A64) is presented to the bus and fully handshaked with the targeted Slave.

All locked resources are to remain locked until the end of the Master's bus mastership. End of bus mastership is signaled by the release of both BBSY* and AS* going high. This allows a Master to lock multiple ports in one lock sequence.

Use of a Lock Command signifies the start of a locked sequence. Each Slave that was addressed with a lock command records the start of the locked sequence. All subsequent transfers within the Master's mastership are guaranteed to be indivisible with assurance that no other operation can be performed on that resource. The locked sequence ends at the end of the current Master's bus mastership.

The effect of BERR* or RETRY* assertion is to end bus mastership and therefore unlock any locked resources. Error recovery and retry operations are not specified and are left up to higher level protocols.

A Master with its bus requester configured for ROR (release-on-request) requires a special operation. After performing a bus operation the Master effectively parks on the bus by not releasing the bus since no other Master wants to use the bus. If the Master ran any locked cycles it should release the bus even though no other Master wants the bus. This will ensure that all locks are released after any locked sequence is performed. The Master can optionally request to use the bus again.

RULE 2.81:

Slaves that accept a lock command **MUST** lock out all other accesses to that resource until the bus is released.

RULE 2.82:

Slaves **MUST** unlock any locks when bus mastership is released, which is signaled by both AS* and BBSY* high.

RECOMMENDATION 2.7:

Slaves should drive RETRY* low in response to a lock command if any of its other port(s) have locked out bus accesses to the selected resource.

RULE 2.83:

A Master **MUST** release BBSY* high after the last rising edge of AS* when it finishes its locked sequence.

OBSERVATION 2.95:

Rule 2.83 precludes the early release of BBSY* during the last data transfer of a locked transfer.

RULE 2.84:

All lock commands **MUST** be terminated with DTACK*, BERR*, or a RETRY* being driven low.

OBSERVATION 2.96:

A locked Slave can only be unlocked when the bus is released. Parking on the bus **MAY** maintain the lock beyond the point at which it was intended to be released.

RECOMMENDATION 2.8:

Masters that can generate lock commands should be labeled as A16:LCK, A24:LCK, A32:LCK, A40:LCK or A64:LCK.

RECOMMENDATION 2.9:

Slaves that can accept and support lock commands should be labeled as A16:LCK, A24:LCK, A32:LCK, A40:LCK or A64:LCK.

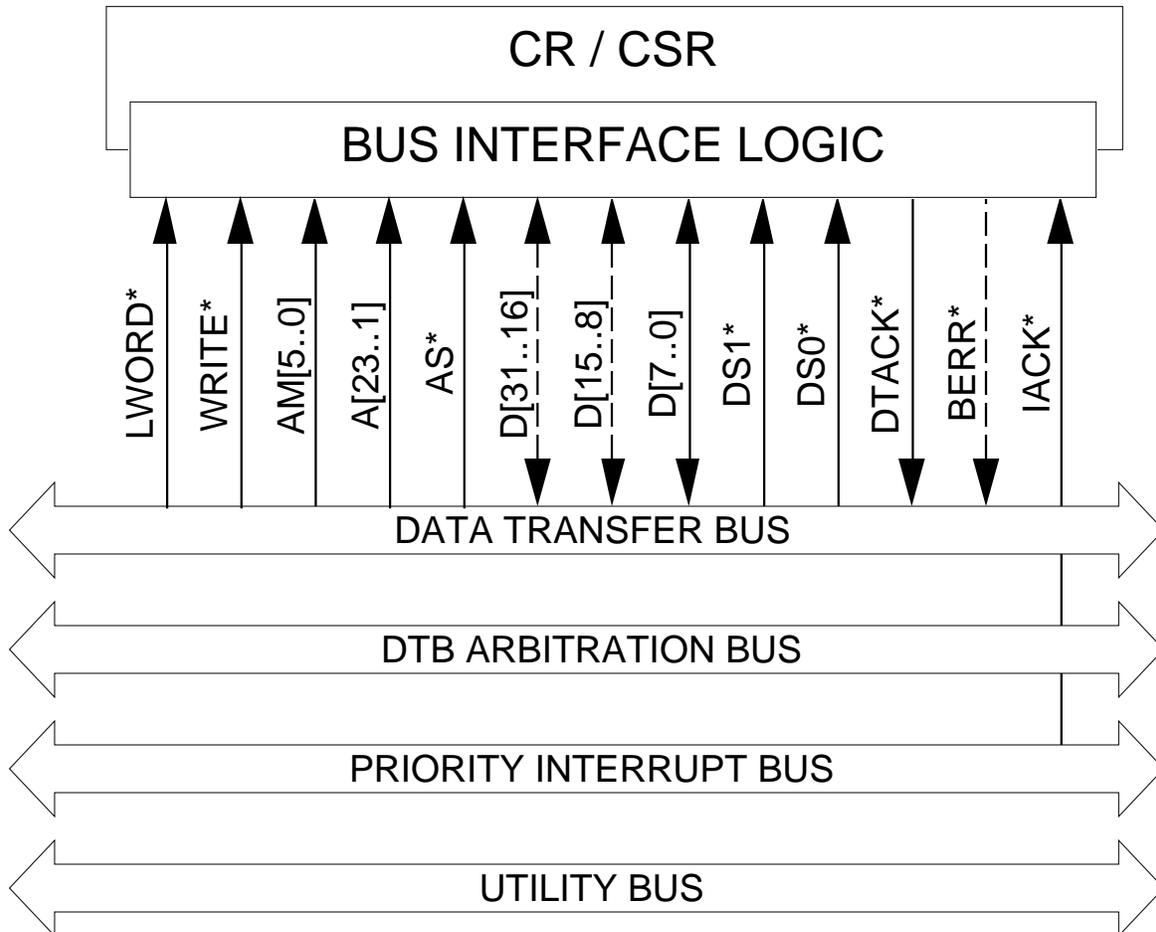
2.3.12 Configuration ROM / Control And Status Register Capability

The Configuration ROM / Control and Status Registers (CR/CSR) provide a mechanism for manufacturer identification, board identification, automatic board initialization, board test and board configuration.

The CR/CSR space is accessed using a specific address modifier code (refer to table 2-3). Accesses to CR/CSR space use the A24 address mode to provide compatibility between 6U and 3U boards. The CR/CSR space is divided into 32 independent 512KBytes regions. Each board in a system may occupy one of these CR/CSR regions. The method used for establishing unique addresses for each board is not defined. The method can be jumpers, Autoslot ID, etc. The 5 most significant bits of the A24 address range must be unique, A(23..19). Address lines, A(18..1), are used to address within each CR/CSR region. Data width modes D08(O), D08(E0), D16 and D32 are used to access a board's CR/CSRs. 3U boards are limited to D08(O), D08(E0) and D16 data width modes.

The lower portion (lowest address) of each CR/CSR region is to be used for configuration ROM and the upper portion (highest address) is to be used for a board's control and status registers (CSRs).

The CR/CSR block diagram is shown in Figure 2 - 34. The dashed lines show signals whose use varies among the various implementations of CR/CSRs. Table 2 - 30 specifies the requirements to monitor these lines



NOTE: Accessing CR/CSR is via a slave interface. All appropriate Slave rules, recommendations, observations, and timing parameters are to be used for CR/CSR accesses.

Figure 2 - 34 Block Diagram: Configuration ROM / Control & Status Registers

RULE 2.85:

Each board with CR/CSR capability MUST NOT occupy more than one 512 KB CR/CSR region. The block MUST be divided into two portions, with the first portion being a configuration ROM area and the second portion a CSR area. The Configuration ROM MUST start at the bottom (lowest) address. The CSR MUST start at the highest address and works its way towards the configuration ROM. Table 2 - 31 defines the CSR Base Register.

RULE 2.86:

Bytes 0x00 through 0x7F of the Configuration ROM MUST be programmed as defined in Table 2- 32.

Table 2 - 30 Configuration ROM / Control & Status Registers: RULEs And PERMISSIONs Or Monitoring The Dashed Lines

Type of CR/CSR	RULEs and PERMISSIONs
D08(O)	MAY monitor D[31..8]
D08(EO)	MUST monitor D[15..8] MAY monitor D[31..16]
D16	MUST monitor D[15..8] MAY monitor D[31..16]
D32	MUST monitor D[31..8]
All	MAY drive BERR*

Note: The mnemonics D08(E), D08(EO), D16 and D32 are defined in Table 2 - 10.

RULE 2.87:

The minimum Configuration ROM access data width MUST be D08(O), with access to every fourth byte, starting at byte 0x03.

PERMISSION 2.16:

Boards may be designed without a CR/CSR capability.

Table 2 - 31 Control And Status Register Base Definition

Address Offset	Definition												
0x7FFFF	CR/CSR Base Address Register (BAR). The BAR is used to select one of the 31 available CR/CSR regions (region 0x00 is reserved for use in Auto Slot Id). Thus the BAR values are in the range 1 to 31 (decimal). The BAR bits 7 to 3 define the CR/CSR region and correspond to A[23..19] in CR/CSR space.												
0x7FFFB	<p style="text-align: center;">Bit Set Register</p> <table style="width: 100%; border: none;"> <thead> <tr> <th></th> <th style="text-align: center;">On Writes</th> <th style="text-align: center;">On Reads</th> </tr> </thead> <tbody> <tr> <td>Bit 7</td> <td>1: put board in reset mode 0: no effect</td> <td>1: board is in reset mode 0: board is not in reset mode</td> </tr> <tr> <td>Bit 6</td> <td>1: enable SYSFAIL driver 0: no effect</td> <td>1: SYSFAIL driver is enabled 0: SYSFAIL driver not enabled</td> </tr> <tr> <td>Bit 5</td> <td>1: no effect 0: no effect</td> <td>1: board has failed 0: board has not failed</td> </tr> </tbody> </table>		On Writes	On Reads	Bit 7	1: put board in reset mode 0: no effect	1: board is in reset mode 0: board is not in reset mode	Bit 6	1: enable SYSFAIL driver 0: no effect	1: SYSFAIL driver is enabled 0: SYSFAIL driver not enabled	Bit 5	1: no effect 0: no effect	1: board has failed 0: board has not failed
	On Writes	On Reads											
Bit 7	1: put board in reset mode 0: no effect	1: board is in reset mode 0: board is not in reset mode											
Bit 6	1: enable SYSFAIL driver 0: no effect	1: SYSFAIL driver is enabled 0: SYSFAIL driver not enabled											
Bit 5	1: no effect 0: no effect	1: board has failed 0: board has not failed											
0x7FFF7	<p style="text-align: center;">Bit Clear Register</p> <table style="width: 100%; border: none;"> <thead> <tr> <th></th> <th style="text-align: center;">On Writes</th> <th style="text-align: center;">On Reads</th> </tr> </thead> <tbody> <tr> <td>Bit 7</td> <td>1: remove board from reset mode 0: no effect</td> <td>1: board is in reset mode 0: board is not in reset mode</td> </tr> <tr> <td>Bit 6</td> <td>1: disable SYSFAIL driver 0: no effect</td> <td>1: SYSFAIL driver is enabled 0: SYSFAIL driver not enabled</td> </tr> <tr> <td>Bit 5</td> <td>1: no effect 0: no effect</td> <td>1: board has failed 0: board has not failed</td> </tr> </tbody> </table>		On Writes	On Reads	Bit 7	1: remove board from reset mode 0: no effect	1: board is in reset mode 0: board is not in reset mode	Bit 6	1: disable SYSFAIL driver 0: no effect	1: SYSFAIL driver is enabled 0: SYSFAIL driver not enabled	Bit 5	1: no effect 0: no effect	1: board has failed 0: board has not failed
	On Writes	On Reads											
Bit 7	1: remove board from reset mode 0: no effect	1: board is in reset mode 0: board is not in reset mode											
Bit 6	1: disable SYSFAIL driver 0: no effect	1: SYSFAIL driver is enabled 0: SYSFAIL driver not enabled											
Bit 5	1: no effect 0: no effect	1: board has failed 0: board has not failed											

* Setting Bit 7 (RESET) must not affect normal VMEbus operation.(i.e. daisy chains and all the bus protocols are unaffected. System Controller functions are also unaffected by the RESET bit). Further definition of reset mode is implementation dependent. Except as previously defined, locations 0x7FFF0 through 0x7FFFF are reserved.

Table 2 - 32 Configuration ROM Definition

Address	Offset	Definition
0x03		Checksum. An eight bit 2s complement binary checksum. The sum of bytes is from offset 0x07 for the number of bytes specified in the length field (inclusive) allowing for the CSR data access width. Area to address 0x7F is assumed to be every fourth byte regardless of the ROM data access width specified in 0x13
0x07, 0x0B 0x0F		Length of ROM that is to be checksummed. The length is a three byte binary value with the MSB in byte 0x07. Note that bits 7 through 3 of the MSB should be all zeros. The range of length is 0x1F to 0x7FFF0.
0x13		Configuration ROM data access width: 0x00 Not to be used 0x01-0x80 Reserved for future use 0x81 Only use D08(E0), every fourth byte 0x82 Only use D08(E0), every other byte 0x83 Use D16 or D08(E0), every byte used 0x84 Use D32, D16 or D08(E0), every byte used 0x85-0xFE Reserved for future use 0xFF Not to be used
0x17		CSR data access width: 0x00 Not to be used 0x01-0x80 Reserved for future use 0x81 Only use D08(E0), every fourth byte 0x82 Only use D08(E0), every other byte 0x83 Use D16 or D08(E0), every byte used 0x84 Use D32, D16 or D08(E0), every byte used 0x85-0xFE Reserved for future use 0xFF Not to be used
0x1B		CSR /CSR Space Specification ID 0x00 Not to be used 0x01 VME64 -1994 version 0x02-0xFE Reserved for future use 0xFF Not to be used
0x1F	0x43	'C'. Used to identify a valid CR

Table 2 - 32 Configuration ROM Definition (continued)

Address	Definition
0x23	0x52 'R'. Used to identify a valid CR
0x27, 0x2B, 0x2F	Manufacturer's ID The 24 bit ID assigned by IEEE. 0x27 is the MSB, 0x2F is LSB. See Appendix C.
0x33, 0x37, 0x3B & 0x3F	Board ID (4 bytes) (Binary number, supplied & controlled by manufacturer)
0x43, 0x47, 0x4B & 0x4F	Revision ID (4 bytes) (Binary number, format & information controlled by manufacturer)
0x53, 0x57 & 0x5B	Pointer to a null terminated ASCII type printable string or a null value (0x000000) if no is string available. The pointer is relative to a base address of 0x000000 and is assumed to be in the CR address region.
0x5F to 0x7B	Reserved for future use
0x7F	Program ID Code
	0x00 Not used
	0x01 No program, ID ROM only
	0x02-0x4F Manufacturer defined
	0x50-7F User defined
	0x80-EF Reserved for future use
	0xF0-FE Reserved by Boot Firmware (P1275)
	0xFF Not to be used
0x80	Start of user defined area, including special programs

OBSERVATION 2.97

The method used to configure a module and the system via the CR/CSR is beyond the scope of this standard.

RULE 2.88:

Accessing CR/CSRs is the same as accessing a slave and **MUST** follow the same rules as specified for slaves.

OBSERVATION 2.98:

All Configuration ROM byte locations between 0x00 and 0x7F not defined in Table 2 - 32 are reserved.

RECOMMENDATION 2.10:

All reserved locations in the Configuration ROM should read as either 0x00 or 0xFF.

Note that "Boot Firmware" accesses only one byte at a time. Therefore, the Configuration ROM interface implementation may be simplified by using only a byte wide ROM/PROM.

The definition and usage of the control and status registers (CSR) is user defined. Interface access to CSRs can be any data width: D08(O), D08(E0), D16 and/or D32. CSRs are to be placed at the top end of the board's CR/CSR region. This allows for up to 512 KB of Configuration ROM, minus the area used for CSRs.

2.3.13 Retry Capability

RETRY* can be driven low by a responding Slave to indicate to the Master that the requested data transfer cannot be executed at this time. The Master that supports Retry should retry the data transfer at a later time. In the interim, the bus can be used for other transactions.

If a Slave supports Retry capability, it can assert RETRY* low to signal the Master that either the requested resource is busy or that a deadlock condition has occurred. If the Master supports retry capability, it terminates the bus cycle when it detects RETRY* low without waiting for either DTACK* or BERR*.

In the case of deadlock condition, if the bus Master does not support Retry and does not terminate the cycle, the Slave drives BERR* to end the bus cycle.

In the case of a busy condition, if the Master does not support Retry and does not terminate the cycle the slave waits and asserts DTACK* low once the busy resource becomes available.

Figure 2-30 shows the timing for the case when the Master terminates the bus cycle when it detects RETRY* low. Figure 2-31 shows the timing for the case when the Master ignores RETRY* low. Additional RULES and OBSERVATIONS regarding the use of RETRY* signal are covered in Tables 2-24, 2-25, and 2-26.

PERMISSION 2.17:

VMEbus Masters and Slaves MAY be designed without support for RETRY* signal.

OBSERVATION 2.99:

The bus cycle is always terminated with either DTACK* low or BERR* low if either the Master or the selected Slave does not support retry capability.

RULE 2.89:

The RETRY* driver MUST be a rescinding driver.

RULE 2.90:

If the Master responds to a RETRY* low, then it MUST release the bus and allow another Master to use the bus if there is another request pending, regardless of the request level.

RULE 2.91:

If a Slave wishes to RETRY* a block transfer cycle it MUST do so during the address phase of the cycle - not the data phase.

PERMISSION 2.18:

After releasing control of the DTB as a result of detecting RETRY* low, the Master MAY signal its requester to request the bus again.

RECOMMENDATION 2.11:

After releasing the control of the DTB as a result of detecting **RETRY*** low, the Master should wait for a user-defined period of time before requesting the bus again, if there are other requests pending.

2.3.14 Interaction Between DTB Functional Modules

Data transfers take place between Masters and Slaves. The Master is the module controlling the transfer. The Slave that recognizes the address as its own is the responding Slave and all other Slaves are nonresponding Slaves.

After initiating a data transfer cycle, the Master waits for a response from the responding Slave. When the Master detects that response, it drives its data strobes and address strobe high, terminating the cycle. The Slave responds by releasing its response line.

OBSERVATION 2.26a:

Although the address and data timing are largely independent, there are two exceptions. First, the Master waits until it has driven **AS*** low before driving either of the data strobes low. Second, the Slave acknowledges both the address strobe and the data strobes with either **DTACK***, **RETRY*** or **BERR***.

RULE 2.13a:

IF a Slave responds to a data transfer cycle with **DTACK*** or **BERR***,
THEN it **MUST** either drive **DTACK*** low or it **MUST** drive **BERR*** low, but not both.

OBSERVATION 2.27:

Because of possible bus skew, due to different loading of the address strobe and the data strobes, the falling edge of the data strobes might be detected by the Slave slightly before the address strobe falling edge.

RULE 2.14a:

Before driving the data bus, the Master **MUST** ensure that the previous responding Slave has stopped driving the data bus. It does so by verifying that **DTACK***, and **BERR*** are high before it drives the data strobe(s) to low on any cycle, and before it drives any of the data lines during a write cycle.

RULE 2.15a:

At the end of a read cycle and at the end of the last read of a block read cycle, the responding Slave **MUST** release the lines that carry data before allowing **DTACK*** or **RETRY*** to go high or driving **DTACK*** or **RETRY*** to high.

RULE 2.16:

When the Master reads data from the Slave, the Slave **MUST** maintain valid data on the data bus until the Master returns the first data strobe to high.

SUGGESTION 2.3a:

For optimum performance, design Masters so that they drive the data strobes high as soon as possible after **DTACK***, **BERR***, or **RETRY*** goes low. Also, design Slaves so that they release the data bus, **BERR***, **RETRY***, and **DTACK*** high as soon as possible after detecting the data strobes high. This allows the maximum data transfer rate on the bus. (Also refer to **OBSERVATION 2.102**)

RECOMMENDATION 2.12:

For higher performance, it is desirable to drive DTACK* high for 30 nanoseconds or less before releasing DTACK*.

RULE 2.92:

Masters that drive DS0* and/or DS1* and AS* high at the end of a cycle, MUST do so for no longer than 30 nanoseconds, if the Master also releases bus mastership.

OBSERVATION 2.29 :

Addressing information on the bus might change soon after a module drives DTACK* or BERR* low, and before the Master drives the data strobes high.

A third type of module, the Location Monitor, monitors the data transfer and generates on-board signal(s) whenever a location that it monitors is accessed.

If the cycle takes too long, a fourth module, called a Bus Timer intervenes by driving BERR* low, completing the data transfer handshake, and allowing the bus to resume operation.

RULE 2.17a:

There is a strict interlock between the rising and falling edges of the data strobes and DTACK*/BERR*/RETRY*. Once a Master has driven its data strobe(s) low it MUST NOT drive its data strobe(s) high and finish a transfer without first receiving a data transfer acknowledge, a bus error response, or a retry request.

OBSERVATION 2.30a:

A board containing a processor, which needs to direct data transfers between itself and other VMEbus boards, would contain a Master module. If the same board also contained memory accessible from the VMEbus, it would also contain a Slave module. A floating point processor or intelligent peripheral controller might receive commands through a Slave interface from a general purpose processor board. It then might act as a Master to access global VMEbus memory to execute the command it has been given.

2.4 Typical Operation

Masters initiate data transfers over the DTB. The addressed Slave then acknowledges the transfer. After receiving the data-transfer acknowledge, the Master terminates the data transfer cycle. The asynchronous nature of the DTB allows the Slave to control the time taken for the transfer.

Before doing any data transfers, a Master has to be granted exclusive control of the DTB. This ensures that multiple Masters will not try to use the DTB at the same time. The Master gains control of the DTB using the modules and signal lines of the Arbitration Bus (Section 2.5 explains this further). The following discussion presumes that the Master has already been granted and has assumed control of the DTB.

2.4.1 Typical Data-Transfer Cycles

Figure 2 - 8 shows a typical single-byte read cycle. To start the transfer, the Master drives the addressing lines with the desired address and address modifier code. Since this example is a Byte(1) read cycle, the Master drives LWORD* high and A1 low. Since it is not an interrupt acknowledge cycle, it does not drive IACK* low. The Master then waits for a specified setup time before driving AS* low, to allow the address lines and the address modifier lines to stabilize before the Slaves sample them.

Each Slave determines whether it should respond by examining the levels on the address lines, the address modifier lines and IACK*. While this is happening, the Master drives WRITE* high to indicate a read operation. The Master then verifies that DTACK* and BERR* are high to ensure that the Slave from the previous cycle is no longer driving the data bus. If this is the case, the Master then drives DS0* low, while keeping DS1* high.

The responding Slave then determines which 4-byte group and which byte of that group is to be accessed and starts the transfer. After it has retrieved the data from its internal storage and placed it on data lines D[7..0], the Slave signals the Master by driving DTACK* low. The Slave then holds DTACK* low and maintains the data valid for as long as the Master holds DS0* low.

When the Master receives DTACK* driven low, it waits the prescribed amount of time and then latches the data on D[7..0], releases the address lines and drives DS0* and AS* high. The Slave responds by releasing D[7..0] and releasing DTACK* high.

OBSERVATION 2.31:

The Master in Figure 2-8 releases all of the DTB lines at the end of the data transfer. This is not required unless the Master's requester released BBSY* during the data transfer as described in Section 2.5.

The cycle flow for double byte and quad byte data transfer cycles are very similar to the single byte cycle. Flow diagrams for these cycles are shown in Figures 2-9 and 2-10.

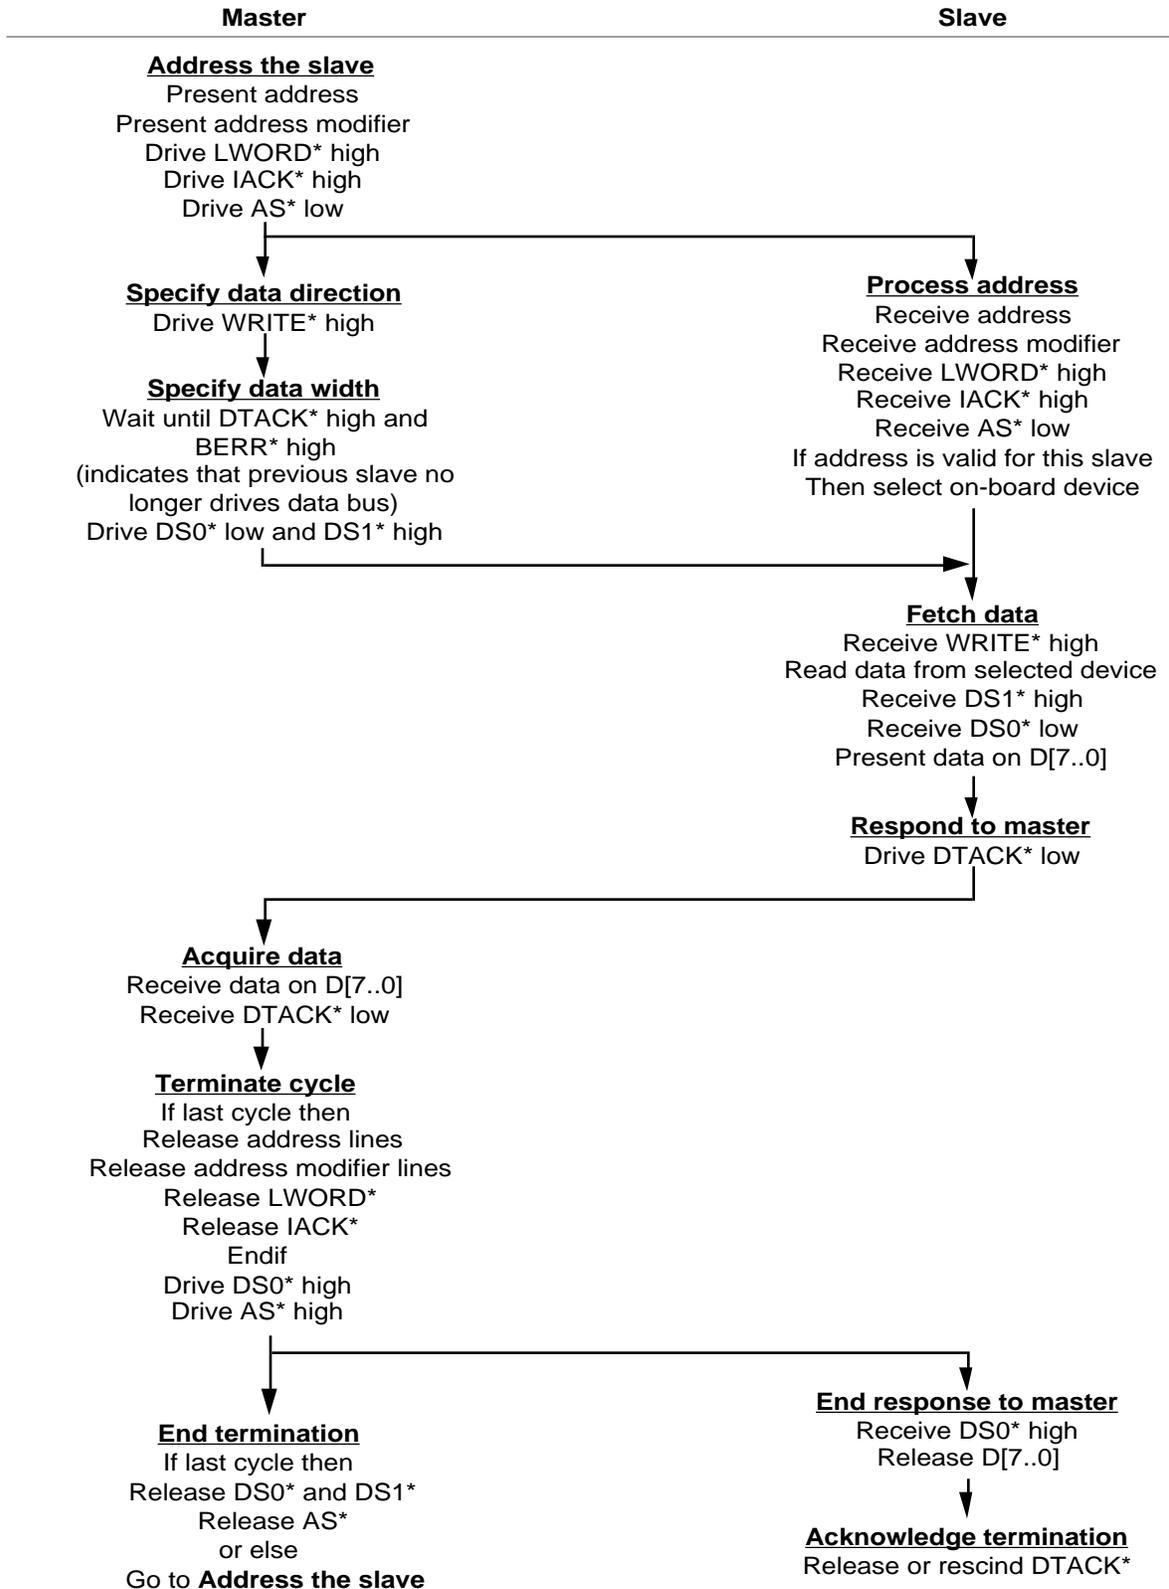


Figure 2 - 8

Non-multiplexed Address, Single-Byte Read Cycle

An Example of a

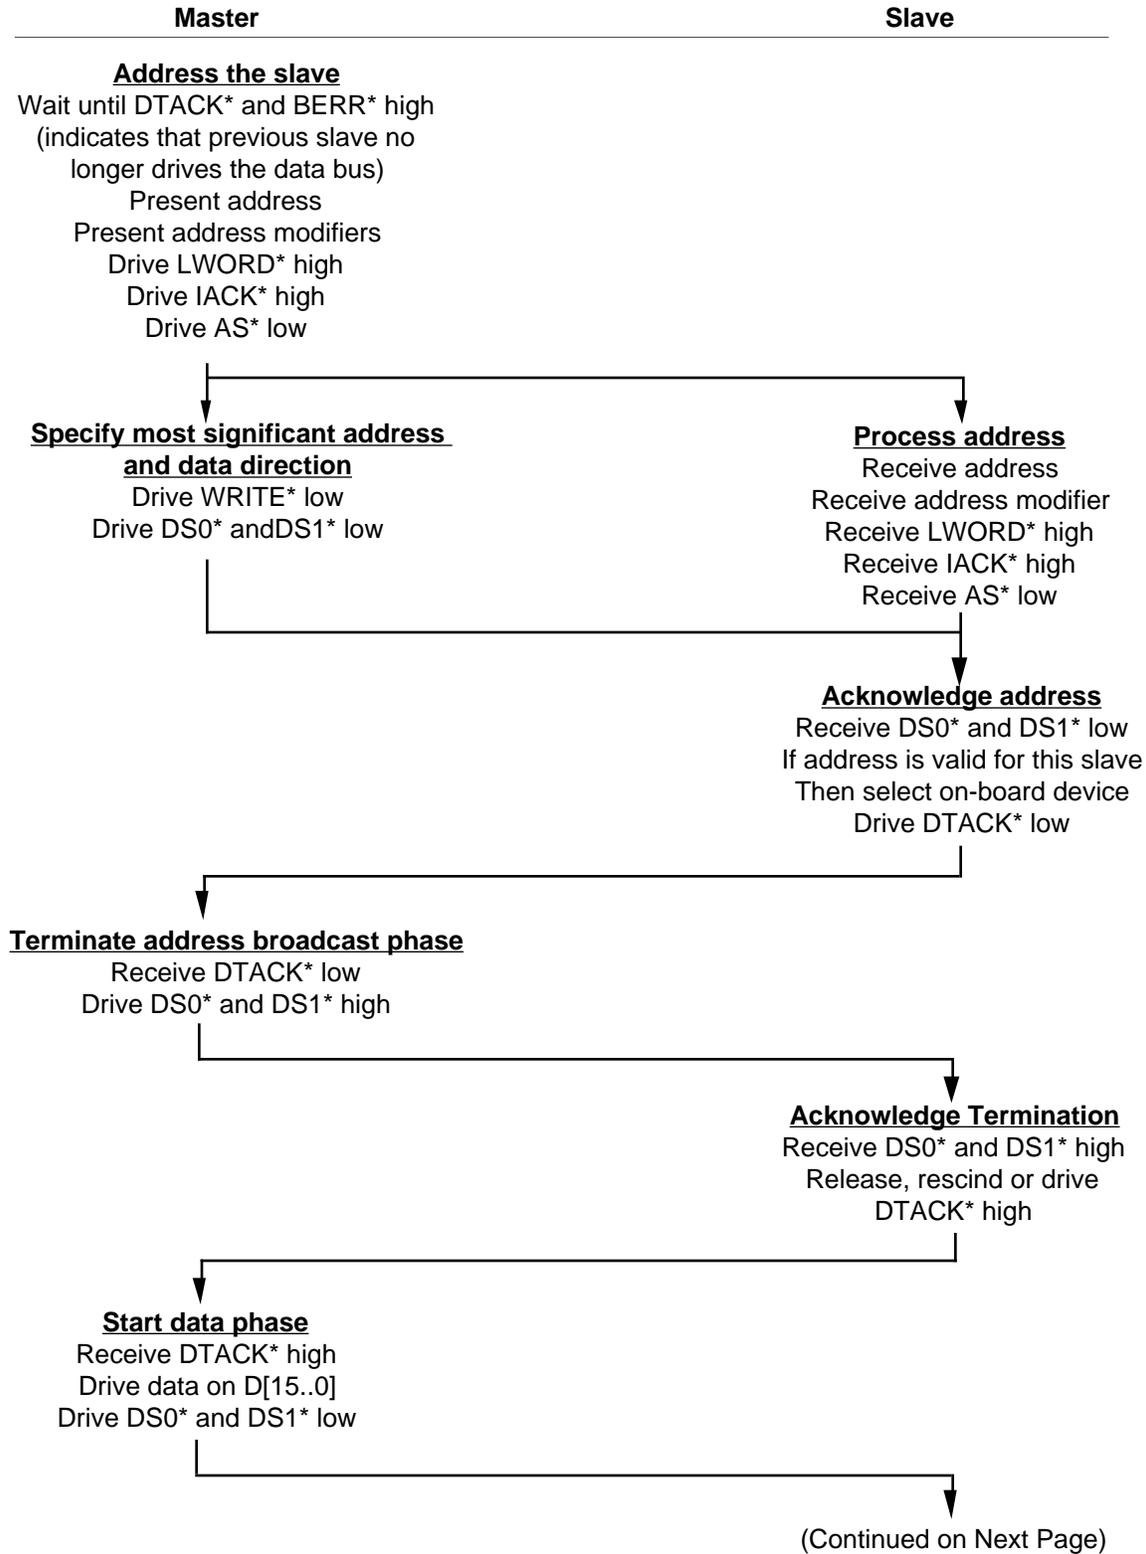


Figure 2 - 9 An Example Of A Multiplexed Address Double-Byte Write Cycle

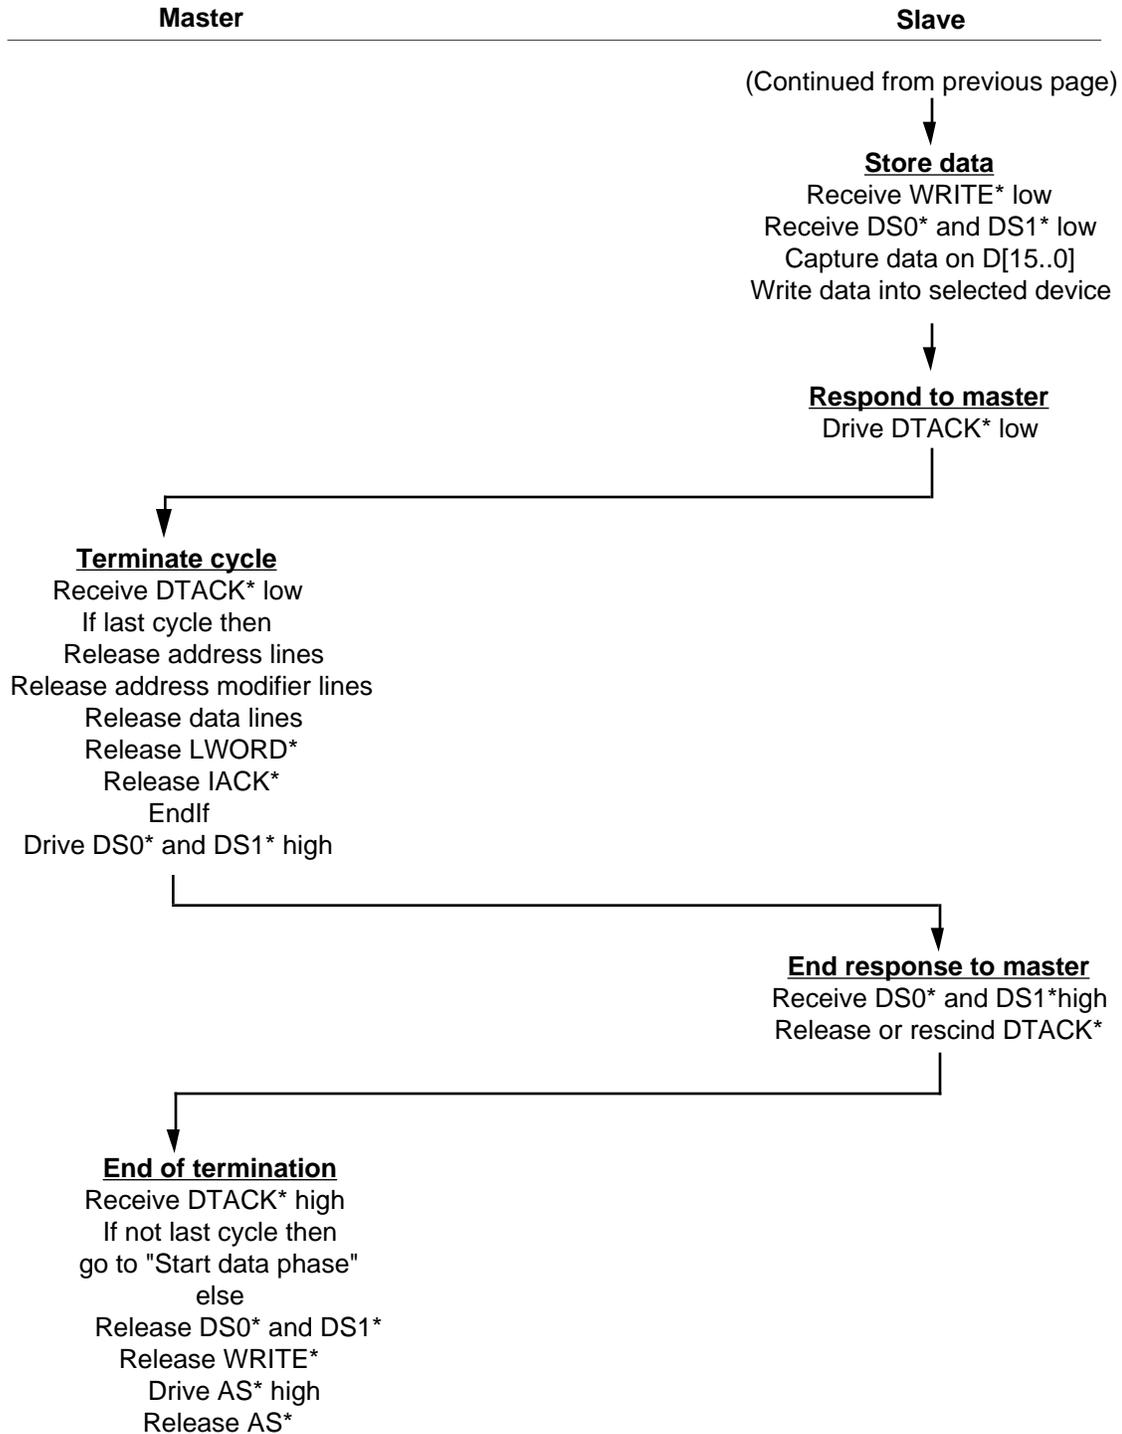


Figure 2 - 9 An Example Of A Multiplexed Address Double-Byte Write Cycle (Cont'd)

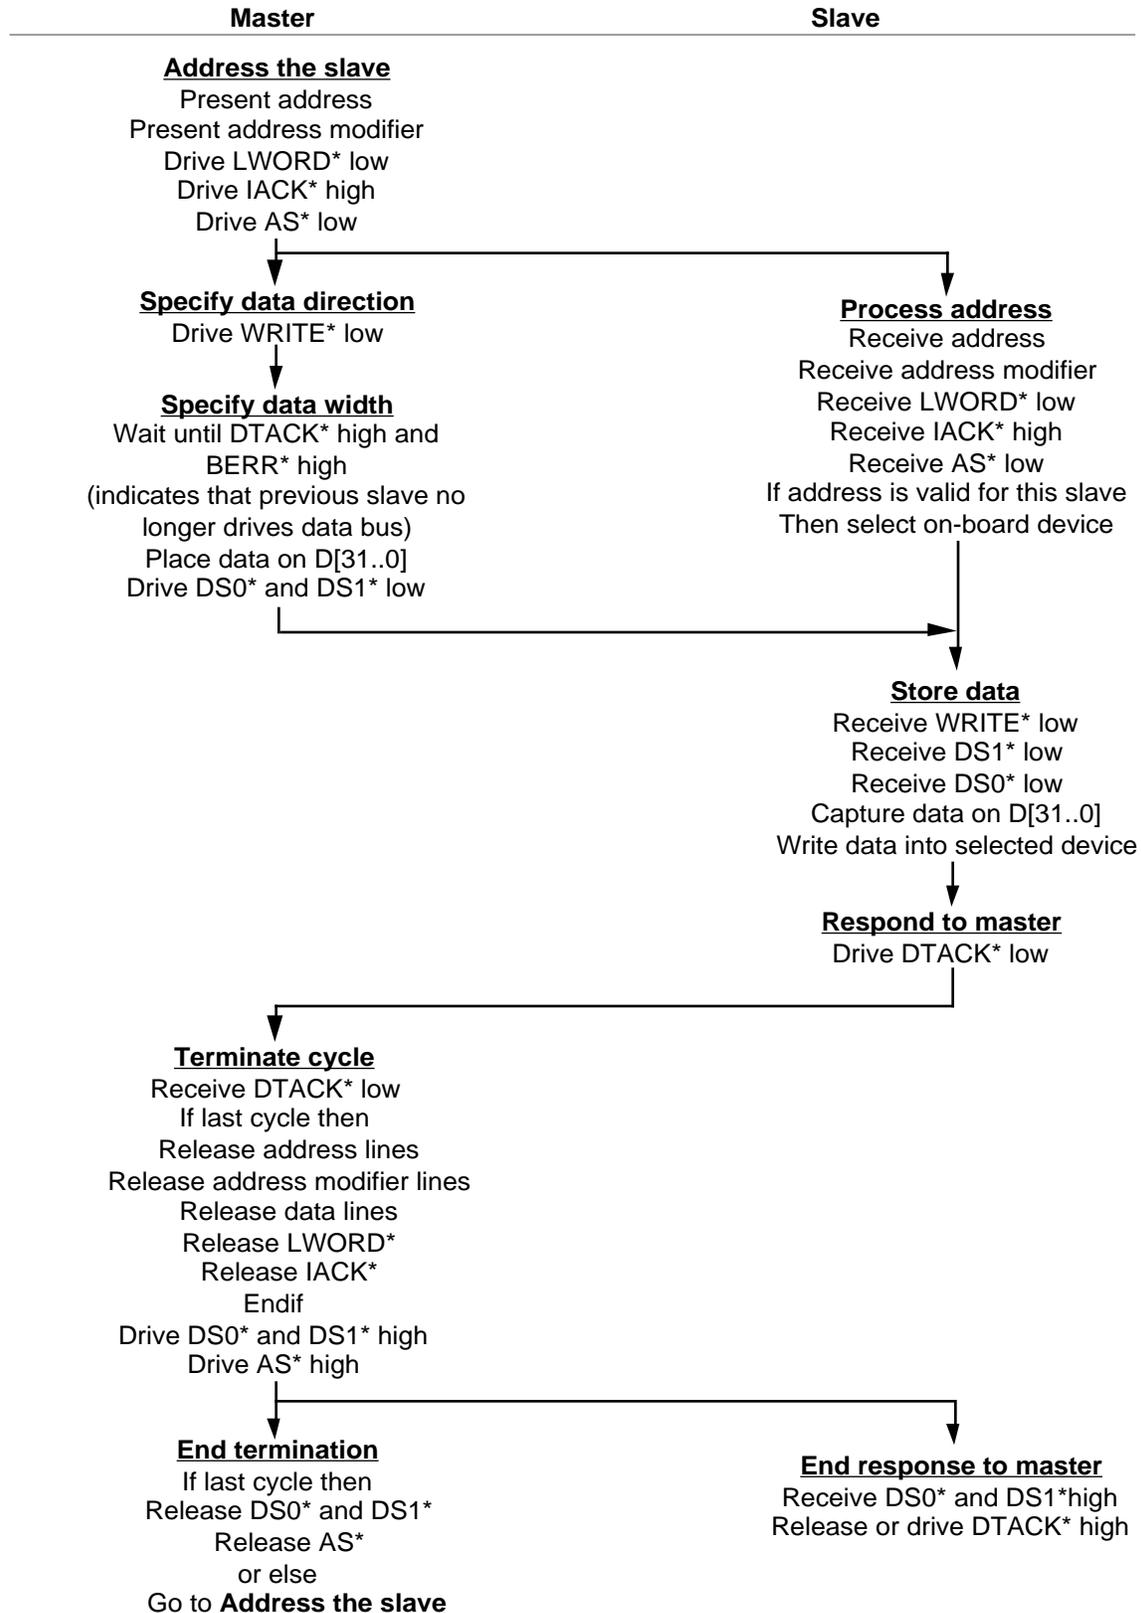


Figure 2 - 10 An Example Of A Non-multiplexed Address Quad-Byte Write Cycle

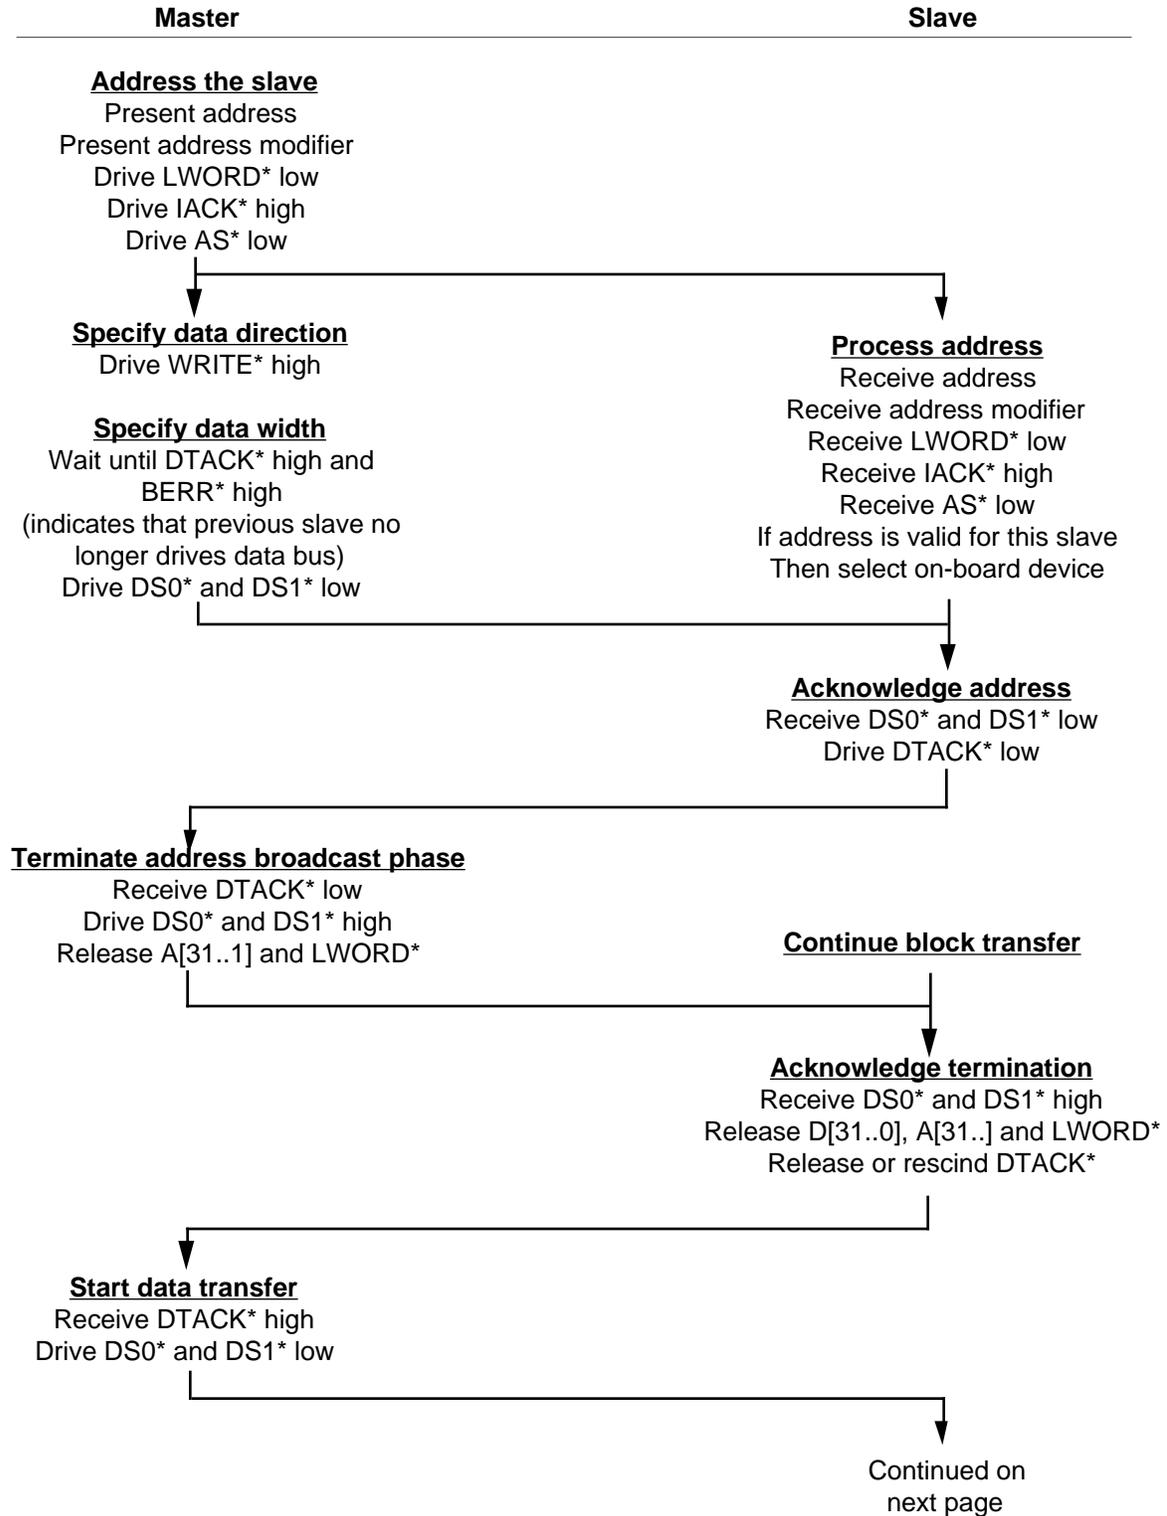


Figure 2 - 35 An Example Of An Eight-Byte Block Read Cycle

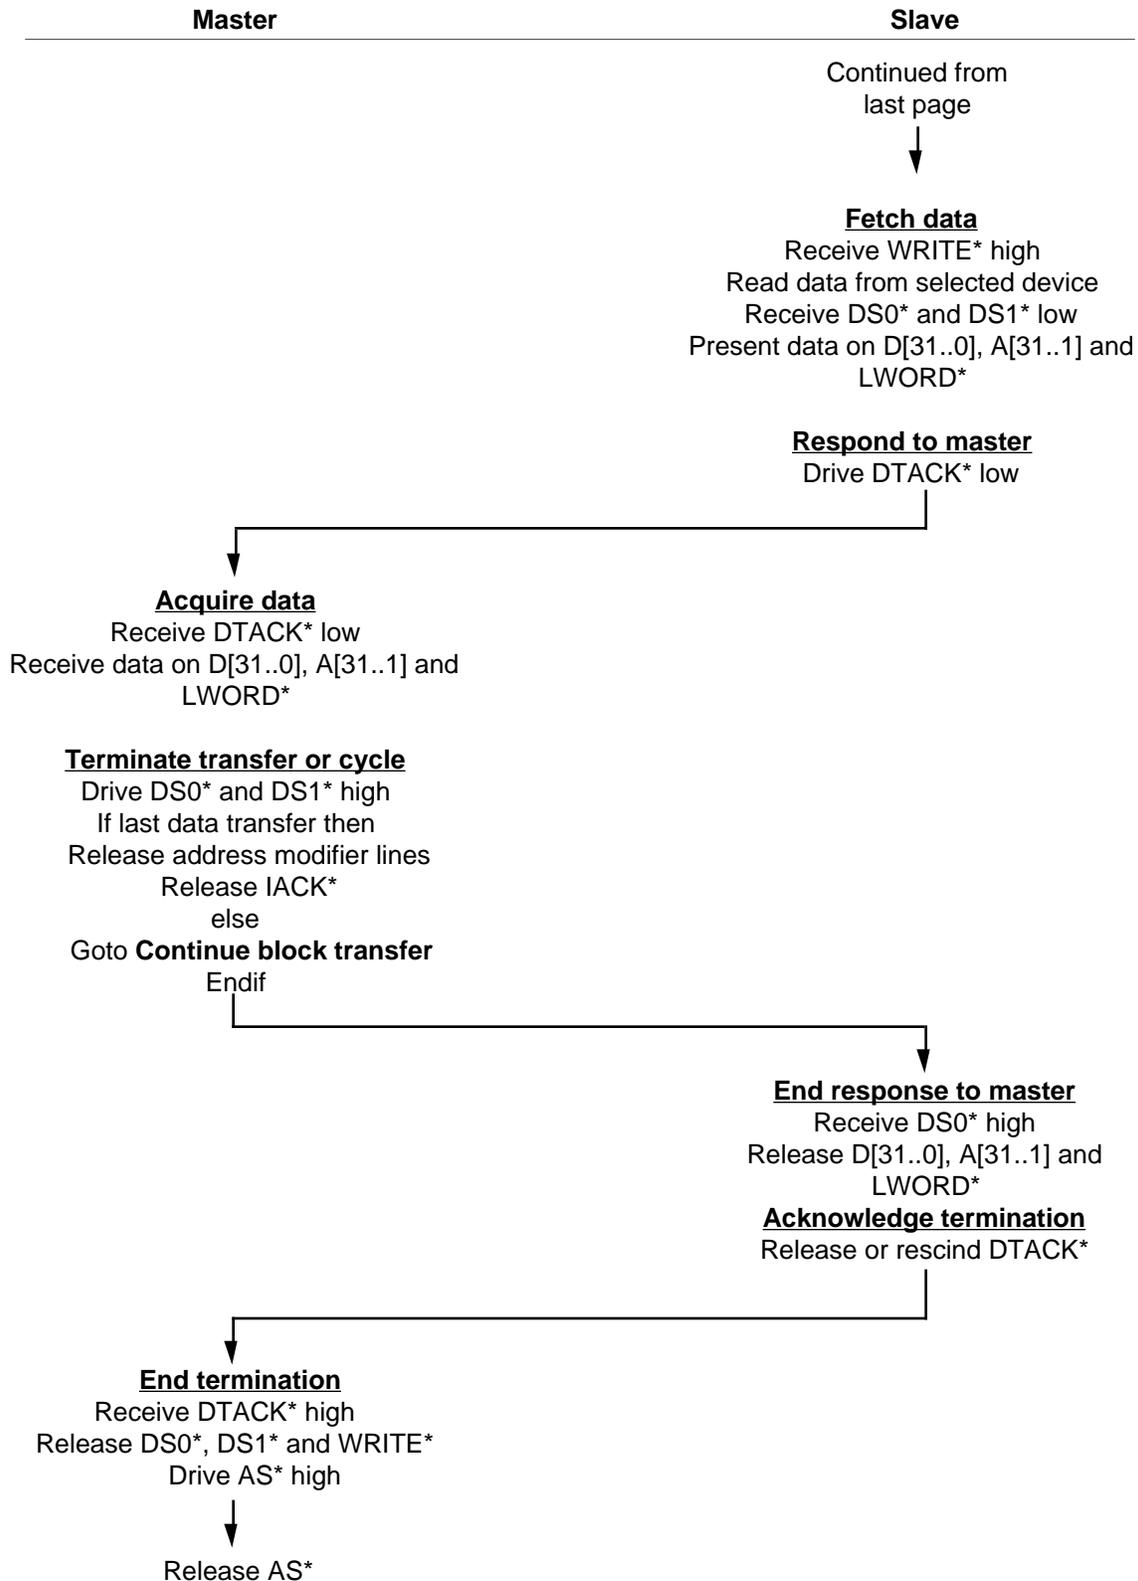


Figure 2 - 35 An Example Of An Eight-Byte Block Read Cycle (continued)

2.4.2 Address Pipelining

The bus strobbs the address and data with separate strobe signals. This allows a Master to broadcast the address for the next cycle while the data transfer for the previous cycle is still in progress. This is called "address pipelining".

PERMISSION 2.8a:

During all cycles except A40, A64, MD32 and MBLT cycles, when the Master detects that the responding Slave has driven DTACK* or BERR* low, it MAY change the address and, after driving AS* high for a minimum time, drive AS* low again.

For example, when a Slave drives DTACK* low on a read cycle, the Master can place a new address on the address bus while it is reading the data from the bus. This amounts to an overlapping of one cycle with the next one, and permits greater speeds on the bus.

RULE 2.93:

Masters MUST NOT use address pipelining during A40, A64 or MBLT and MD32 transfers.

OBSERVATION 2.100:

During MBLT and MD32 read cycles, Slaves use address lines to transfer data. Therefore, the address for the next cycle cannot be broadcast until after the responding Slave releases DTACK* and BERR*, signifying that it no longer drives the address lines.

OBSERVATION 2.101:

During A64 and A40 address phases, Masters use both the address and data lines to broadcast the address. Therefore, the address for the next cycle cannot be broadcast until after the responding Slave releases DTACK* and BERR*, signifying that it no longer is driving the data lines.

RULE 2.18:

Since address pipelining can occur, Slaves MUST NOT be designed on the assumption that they will never encounter pipelined cycles.

OBSERVATION 2.32:

The responding Slave may recognize its address and respond very quickly on the DTACK* or BERR* lines. Since the Master is permitted to remove the address after the responding Slave drives DTACK* or BERR* low, nonresponding Slaves might not be able to decode the addressing information before the Master removes it from the bus.

SUGGESTION 2.4:

Design all Slaves to monitor AS*, and to capture (latch) the addressing information from A[31..1] and LWORD* on the falling edge of AS*.

OBSERVATION 2.33:

Because a Master might broadcast a new address while a previous cycle is finishing, the designer of Slave boards needs to ensure that the second assertion of AS* does not invalidate the first address if that address is needed by on-board logic to maintain the data on the bus.

PERMISSION 2.9:

Masters MAY be designed without the ability for address pipelining. (e.g. They MAY wait until the responding Slave releases DTACK* or BERR* before driving AS* low for the next cycle.)

OBSERVATION 2.34:

A Master might drive AS* low for a new cycle before it drives data strobes high from the previous cycle. Because of this, there might be a period when the address strobe for the new cycle, as well as at least one data strobe from the previous cycle coincide during the cycle overlap.

SUGGESTION 2.5a:

Design Slaves to respond to data transfers when one or both of DS0* and DS1* are low and when DTACK* and BERR* are both high, instead of a simultaneous low level on AS* and one or both of DS0* and DS1*.

2.5 Data-Transfer-Bus Acquisition

RULE 2.19:

Before transferring any data on the DTB, a Master MUST get permission to use it.

Several Masters might want to use the DTB at the same time. The process that determines which Master can use the DTB is called arbitration and is discussed in Chapter 3. Because arbitration is closely tied to the operation of the DTB, it is briefly described here.

Figure 2-11 provides two examples that show possible sequences when a Master (called Master A) finishes using the DTB and allows arbitration to take place.

In Example 1, Master A, part way into its last transfer, indicates that it no longer needs the DTB. It does this by having its Requester release the bus busy (BBSY*) signal line. Since Master A gives this early notice that the DTB will soon be available, the arbitration takes place during the last data transfer. The arbitration is completed and Master B is granted permission to use the DTB before Master A has terminated its cycle, but it waits until Master A releases AS* (This ensures that Master B will not start driving the DTB before Master A has finished with its last data transfer).

In Example 2, Master A waits until after its last transfer (i.e., after AS* is released) before releasing BBSY*. In this case the DTB is idle while the arbitration is done. Master B is then granted the bus and, since AS* is already high, it begins using the DTB immediately.

RULE 2.20:

Once a Master's Requester releases BBSY* to high, the Master MUST NOT drive AS* from high to low (i.e., it MUST NOT begin a new cycle) until its Requester receives a new bus grant.

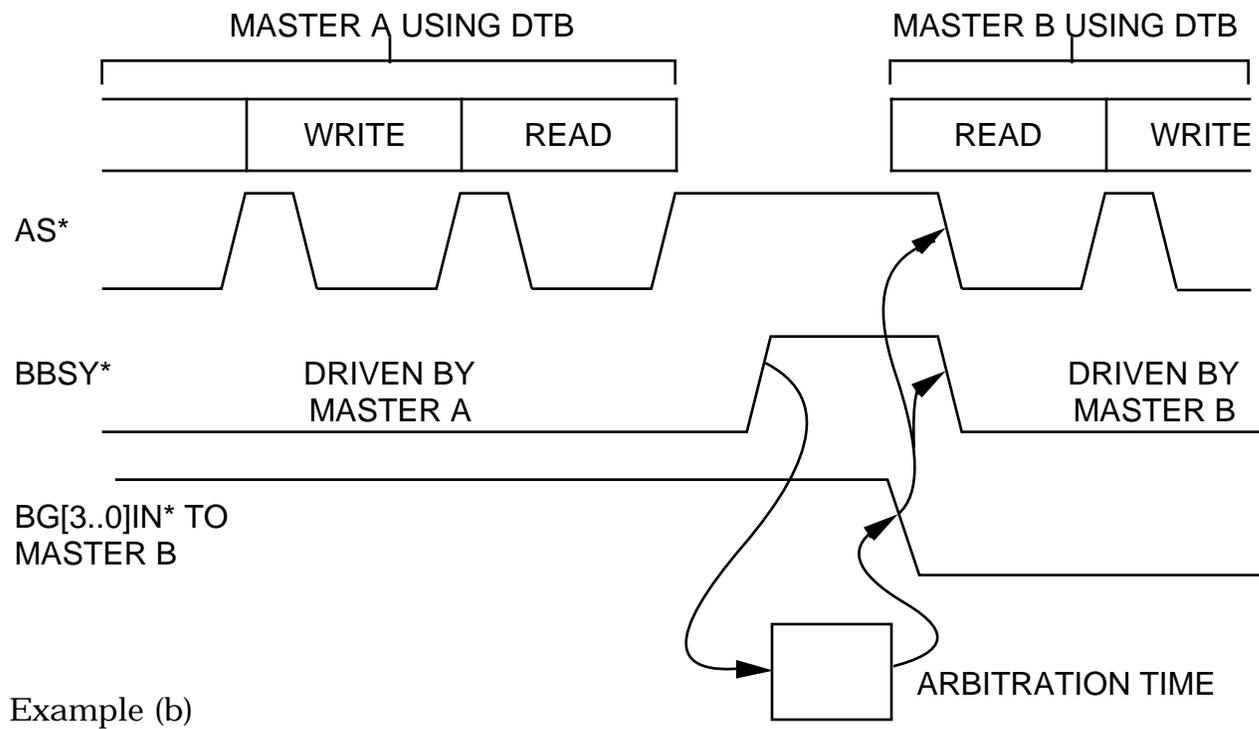
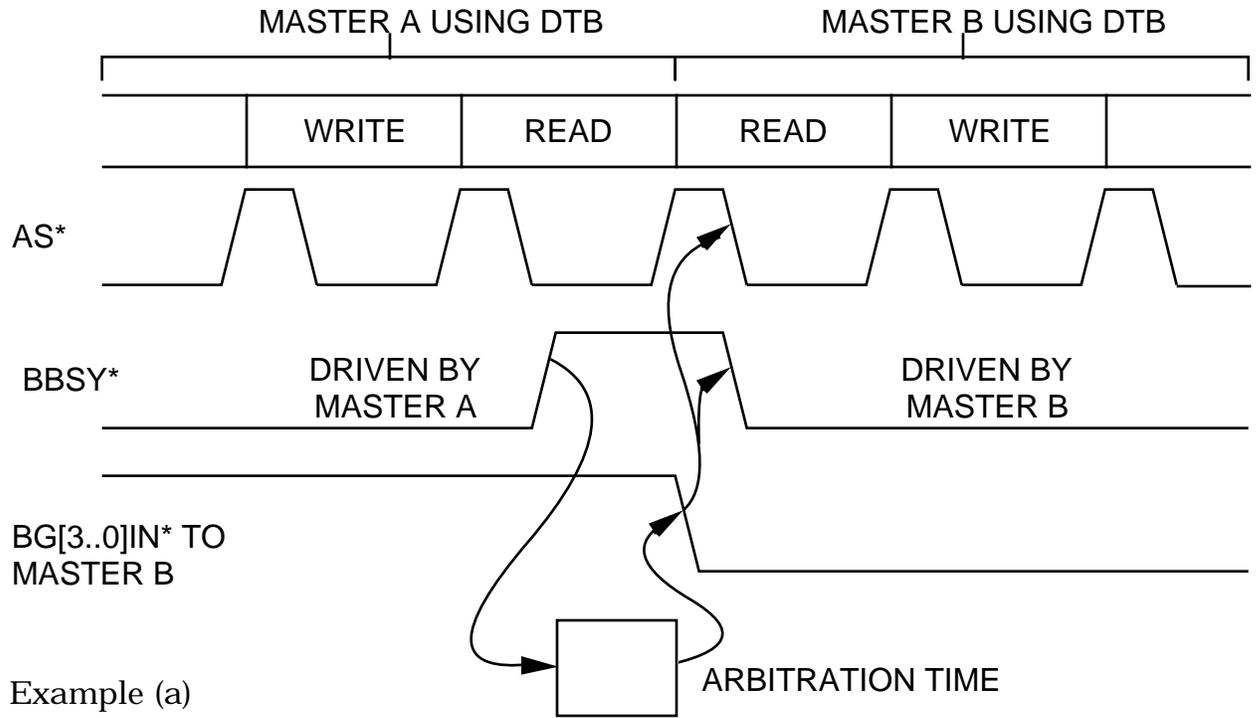


Figure 2 - 11 Data Transfer Bus Master Exchange Sequence

2.6 DTB Timing Rules and Observations

This section describes the timing rules and observations that govern the behavior of Masters and Slaves. This timing information is in the form of figures and tables.

Table 2-17 lists a timing table and timing diagrams that specify Master, Slave, and Location Monitor operation.

Table 2-18 defines the various mnemonics that are used in Tables 2-19 through 2-21.

Tables 2-19 through 2-21 specify the use of the DTB signals.

Tables 2-22 through 2-27 specify the timing parameters of the DTB. (The reference numbers used in Tables 2-24 through 2-27 correspond to the timing parameter numbers in Tables 2-22 and 2-23.)

Figures 2-27 and 2-12 through 2-15 specify the timing RULEs and OBSERVATIONs during address broadcasting time.

Figures 2-28, 2-29, 2-16 through 2-21 and 2-32 specify the timing RULEs and OBSERVATIONs for Masters, Slaves, and Location Monitors during data transfer time.

Figures 2-22 through 2-24 specify the timing RULEs and OBSERVATIONs for Masters and Slaves between DTB cycles.

Figure 2-25 is the timing diagram for Master, Slave, and Bus Timer during a timed-out cycle.

Figure 2-26 shows the timing during the mastership transfer of the DTB.

Figures 2-30 and 2-31 specify the timing RULEs and OBSERVATIONs for Masters and Slaves when RETRY * is used.

Figure 2 - 33 shows the timing for Rescinding DTACK*.

In order to meet the timing RULEs, board designers need to take into account the worst and best case propagation delays of bus drivers and receivers used on their VMEbus boards. The propagation delay of the drivers depends on their output loads, but manufacturer's specifications do not always give enough information to calculate the propagation delays under various loads. To help the board designer, some suggestions are offered in Chapter 6.

The OBSERVATIONs specify the timing of incoming lines signal transitions. These times can be relied upon as long as the backplane loading RULEs in Chapter 6 are not violated. The RULEs for the bus terminators in Chapter 6 guarantee that the timing parameters for signal lines that are released after they have been driven, are met.

Typically, for each timing RULE there is a corresponding OBSERVATION. However, the time that is guaranteed in the OBSERVATION might differ from the time specified by the RULE. For example, a careful inspection of the timing diagrams shows that the Master is required to provide 35 nanoseconds of address and data set-up time, but the Slave is only guaranteed 10 nanoseconds. This is because the address and data bus drivers are not always able to drive the backplane's signal lines completely through the threshold region from low to high until the transition propagates to the end of the backplane and is reflected back. The falling edge of the address and data strobes, however, typically cross the 0.8-volt threshold without waiting for a reflection. The resulting set-up time at the Slave is the Master's set-up time less two bus propagation times.

A special notation has been used to describe the data strobe timing. The two data strobes (DS0* and DS1*) will not always make their transitions simultaneously. For purposes of these timing diagrams, DSA* represents the first data strobe to make its transition (whether that is DS0* or DS1*). The broken line shown while the data strobes are stable is to indicate that the first data strobe to make a falling transition might not be the first to make its rising transition - i.e., DSA* can represent DS0* on its falling edge and DS1* on its rising edge.

Tables 2 -19, 2 - 20, and 2 - 21 show how the various signal lines of the DTB are used to broadcast addresses and to transfer data. These tables are referenced by the various timing diagrams that follow. To keep these Tables compact, mnemonics are used to describe when and how the various lines are driven. These mnemonics are defined in Table 18.

Note that the CR/CSR is a functional module defined in this chapter, but is not specifically called out in this section. Access to CR/CSRs is treated as a slave access. All the appropriate slave RULES therefore apply.

Table 2 - 17 Timing Diagrams That Define Master, Slave, And LOCATION MONITOR Operation (See Table 2-22 For Timing Values)

Mnemonic	Type of Cycles	Address Broadcast Timing Figure(s)	Data Transfer Timing Figure(s)
ADO	Address-Only	2 - 12	N/A
ADOH	Address-Only-With-Handshake	2 - 27	2 - 33
A40	A40 address broadcast	2 - 27	all
A64	A64 address broadcast	2 - 27	all
D08(E0)	Single-even-byte transfer		
	Byte(0) Read	2 - 12, 13, 27	2 - 16, 33
	Byte(2) Read	2 - 12, 13, 27	2 - 16, 33
	Byte(0) Write	2 - 12, 13, 27	2 - 18, 33
	Byte(2) Write	2 - 12, 13, 27	2 - 18, 33
D08(E0) or D08(O)	Single-odd-byte transfers		
	Byte(1) Read	2 - 12, 13, 27	2 - 16, 33
	Byte(3) Read	2 - 12, 13, 27	2 - 16, 33
	Byte(1) Write	2 - 12, 13, 27	2 - 18, 33
	Byte(3) Write	2 - 12, 13, 27	2 - 18, 33
D16	Double-byte transfers		
	Byte(0 - 1) Read	2 - 12, 13, 27	2 - 17, 33
	Byte(2 - 3) Read	2 - 12, 13, 27	2 - 17, 33
	Byte(0 - 1) Write	2 - 12, 13, 27	2 - 19, 33
	Byte(2 - 3) Write	2 - 12, 13, 27	2 - 19, 33
D32	Quad-byte transfers		
	Byte(0-3) Read	2 - 12, 13, 27	2 - 17, 33
	Byte(0-3) Write	2 - 12, 13, 27	2 - 19, 33
MD32	Multiplexed quad byte transfers		
	Byte(0-3) Read	2 - 27	2 - 28, 33
	Byte(0-3) Write	2 - 27	2 - 29, 33
D08(E0): BLT	Single-byte block transfers		
	Single-Byte Block Read	2 - 12, 14, 27	2 - 16, 33
	Single-Byte Block Write	2 - 12, 14, 27	2 - 18, 33
D16: BLT	Double-byte block transfers		
	Double-Byte Block Read	2 - 12, 17, 27	2 - 17, 33
	Double-Byte Block Write	2 - 12, 17, 27	2 - 19, 33

Table 2 -17 Timing Diagrams That Define Master, Slave, And Location-Monitor Operation (See Table 2-22 For Timing Values) (Continued)

Mnemonic	Type of Cycles	Address Broadcast Timing Figure(s)	Data Transfer Timing Figure(s)
D32: BLT	Quad-byte block transfers		
	Quad-Byte Block Read	2 - 12, 17, 27	2 - 17, 33
	Quad-Byte Block Write	2 - 12, 17, 27	2 - 22, 33
MD32:BLT	Multiplexed quad byte block transfers		
	Multiplexed Quad-Byte Block Read	2 - 27	2 - 28, 33
	Multiplexed Quad-Byte Block Write	2 - 27	2 - 29, 33
MBLT	Eight-byte block transfers		
	Eight-Byte Block Read	2 - 27	2 - 28, 33
	Eight-Byte Block Write	2 - 27	2 - 29, 33
D08(EO) RMW	Single-byte RMW transfers		
	Byte(0) Read-Modify-Write	2 - 12, 15, 27	2 - 20, 33
	Byte(1) Read-Modify-Write	2 - 12, 15, 27	2 - 20, 33
	Byte(2) Read-Modify-Write	2 - 12, 15, 27	2 - 20, 33
	Byte(3) Read-Modify-Write	2 - 12, 15, 27	2 - 20, 33
D16: RMW	Double-byte RMW transfers		
	Byte(0-1) Read-Modify-Write	2 - 12, 15, 27	2 - 18, 33
	Byte(2-3) Read-Modify-Write	2 - 12, 15, 27	2 - 18, 33
D32: RMW	Quad-Byte RMW transfers		
	Byte(0-3) Read-Modify-Write	2 - 12, 15, 27	2 - 18, 33
MD32: RMW	Multiplexed Quad-Byte RMW transfers		
	Byte(0-3) Read-Modify-Write	2 - 27	2 - 32, 33
D32: UAT	Unaligned transfers		
	Byte(0-2) Read	2 - 12, 13, 27	2 - 16, 33
	Byte(1-3) Read	2 - 12, 13, 27	2 - 16, 33
	Byte(1-2) Read	2 - 12, 13, 27	2 - 17, 33
	Byte(0-2) Write	2 - 12, 13, 27	2 - 18, 33
	Byte(1-3) Write	2 - 12, 13, 27	2 - 18, 33
	Byte(1-2) Write	2 - 12, 13, 27	2 - 19, 33
N/A	Retry request	N/A	2 - 30, 31

Table 2 - 18 Definitions Of Mnemonics Used In Tables 2 -19, 2 - 20, And 2 - 21

Mnemonic	Description	RULES and PERMISSIONS
DVBM	Driven Valid By Master	RULE 2.21: The Master MUST drive DVBM lines to a valid level.
DLBM	Driven Low By Master	RULE 2.22: The Master MUST drive DLBM lines to a low level.
DHBM	Driven High By Master	RULE 2.23: The Master MUST drive DHBM lines to a high level.
dhbm?	Driven High By Master?	PERMISSION 2.10: The Master MAY drive dhbm? lines high RULE 2.24: The Master MUST NOT drive dhbm? lines low.
dxbm?	Driven By Master?	PERMISSION 2.11: The Master MAY drive dxbm? lines, or it may leave these lines undriven. (When dxbm? lines are driven, they carry no valid information.)
DVBS	Driven Valid By Slave	RULE 2.25: The responding Slave MUST drive DVBS lines to a valid level.
dxbs?	Driven By Slave?	PERMISSION 2.12: The responding Slave MAY drive dxbs? lines, or it may leave these lines undriven. (When dxbs? lines are driven, they carry no valid information.)
DVBB	Driven Valid By Both Slave and Master	RULE 2.26: During the read portion of a Read-Modify-Write Cycle, the responding Slave MUST drive DVBB lines with valid data. During the write portion of a Read-Modify-Write Cycle, the Master MUST drive DVBB lines with valid data.
dxbb?	Driven By Both Slave and Master?	PERMISSION 2.13: During the read portion of a Read-Modify-Write Cycle, the responding Slave MAY drive dxbb? lines, or it MAY leave them undriven. During the write portion of a Read-Modify-Write Cycle, the Master MAY drive dxbb? lines, or it MAY leave them undriven. (When dxbb? lines are driven they carry no valid information.)

Table 2 - 19 Use Of The Address And Data Lines To Select A Byte Group

Addressing Mode	A[15..2] Note (1)	A[23..16]	A[31..24]	D[31..16] Note (2)	D[15..0] Note (2)	IACK*
A16	DVBM	dxbm?	dxbm?	dxbm?	dxbm?	dhbm?
A24	DVBM	DVBM	dxbm?	dxbm?	dxbm?	dhbm?
A32	DVBM	DVBM	DVBM	dxbm?	dxbm?	dhbm?
A40	DVBM	DVBM	dxbm?	dxbm?	DVBM	dhbm?
A64	DVBM	DVBM	DVBM	DVBM	DVBM	dhbm?

NOTES:

(1) A1 is used in conjunction with LWORD*, DS0* and DS1* to select which of the bytes within the byte group is accessed (see Table 2 - 20).

(2) During the A16, A24 and A32 addressing modes, D[31..0] are not used to select byte-locations. They are multiplexed to carry the upper address bits only when using the A40 and A64 addressing modes.

Table 2 - 20 Use of DS1*, DS0*, A1, A2, and LWORD* During the Address Phase of the Various Cycles

Mnemonic	Type of Cycles	DS1*	DS0*	A1	A2	LWORD*
ADO	Address-Only	dhbm?	dhbm?	dxbm?	DVBM	dxbm?
ADOH	Address-Only-With-Handshake	DLBM	DLBM	DVBM	DVBM	DVBM
D08(EO)	Single-even-byte transfers					
	Byte(0) Read or Write	DLBM	dhbm?	DLBM	DVBM	dhbm?
	Byte(2) Read or Write	DLBM	dhbm?	DHBM	DVBM	dhbm?
D08(EO) or D08(O)	Single-odd-byte transfers					
	Byte(1) Read or Write	dhbm?	DLBM	DLBM	DVBM	dhbm?
	Byte(3) Read or Write	dhbm?	DLBM	DHBM	DVBM	dhbm?
D16	Double-byte transfers					
	Byte(0-1) Read or Write	DLBM	DLBM	DLBM	DVBM	dhbm?
	Byte(2-3) Read or Write	DLBM	DLBM	DHBM	DVBM	dhbm?
D32, MD32	Quad-byte transfers					
	Byte(0-3) Read or Write	DLBM	DLBM	DLBM	DVBM	DLBM
D08(EO): BLT	Single-byte block transfers					
	Single-Byte Block Read or Write	<	Note (1)	>	DVBM	dhbm?
D16:BLT	Double-byte block transfers					
	Double-Byte Block Read or Write	DLBM	DLBM	Note 2	DVBM	dhbm?
D32:BLT, MD32:BLT	Quad-byte block transfers					
	Quad-Byte Block Read or Write	DLBM	DLBM	DLBM	DVBM	DLBM

Table 2 - 20 Use Of DS1*, DS0*, A1, A2, And LWORD*
During The Address Phase Of The Various Cycles (continued)

Mnemonic	Type of Cycles	DS1*	DS0*	A1	A2	LWORD*
MBLT	Eight-byte block transfers Eight-Byte Block Read or Write	DLBM	DLBM	DLBM	DLBM	DLBM
D08(EO): RMW	Single-byte Read-Modify-Write transfers					
	Byte(0) RMW	DLBM	dhbm?	DLBM	DVBM	dhbm?
	Byte(1) RMW	dhbm?	DLBM	DLBM	DVBM	dhbm?
	Byte(2) RMW	DLBM	dhbm?	DHBM	DVBM	dhbm?
	Byte(3) RMW	dhbm?	DLBM	DHBM	DVBM	dhbm?
D16: RMW	Double-byte read-modify-write transfers					
	Byte(0-1) RMW	DLBM	DLBM	DLBM	DVBM	dhbm?
	Byte(2-3) RMW	DLBM	DLBM	DHBM	DVBM	dhbm?
D32:RMW, MD32:RMW	Quad-byte read-modify-write transfers					
	Byte(0-3) RMW	DLBM	DLBM	DLBM	DVBM	DLBM
D32:UAT	Unaligned transfers					
	Byte(0-2) Read or Write	DLBM	dhbm?	DLBM	DVBM	DLBM
	Byte(1-3) Read or Write	dhbm?	DLBM	DLBM	DVBM	DLBM
	Byte(1-2) Read or Write	DLBM	DLBM	DHBM	DVBM	DLBM

NOTES:

(1) During single-byte block transfers, the two data strobes are alternately driven low. Either data strobe might be driven low on the first transfer. If the first accessed byte location is Byte(0) or Byte(2), then the Master drives DS1* low first. If the first accessed byte location is Byte(1) or Byte(3), then it drives DS0* low first. A1 is valid only on the first data transfer (that is, until the Slave drives DTACK* or BERR* low the first time) and might be either high or low depending upon which byte the single byte block transfer begins with. If the first byte location is Byte(0) or Byte(1), then the Master drives A1 to low. If the first byte location is Byte(2) or Byte(3), then the Master drives A1 to high.

An example of the use of DS0*, DS1*, A1, A2, and LWORD* during a single-byte block transfer cycle that starts with Byte(2) is given below.

		DS1*	DS0*	A1	A2	LWORD*
First data transfer	Byte(2)	DLBM	DHBM	DHBM	DVBM	dhbm?
	Byte(3)	DHBM	DLBM	dxbm?	DVBM	dxbm?
	Byte(0)	DLBM	DHBM	dxbm?	DVBM	dxbm?
	Byte(1)	DHBM	DLBM	dxbm?	DVBM	dxbm?
Last data transfer	Byte(2)	DLBM	DHBM	dxbm?	DVBM	dxbm?

(2) During a double-byte block transfer, A1 is valid only on the first data transfer (that is, until the Slave drives DTACK* or BERR* low the first time) and is driven either high or low depending upon what double-byte group the double-byte block transfer begins with. If the first double-byte group is Byte(0-1), then the Master drives A1 low. If the first double-byte group is Byte(2-3), then the Master drives A1 high.

Table 2 - 21 Use Of The Data Lines To Transfer Data

Mnemonic	Type of Cycles	A[31..1] LWORD*	D[31..24]	D[23..16]	D[15..8]	D[7..0]
ADO	Address-Only	dxbm? ²	dxbm?	dxbm?	dxbm?	dxbm?
ADOH	Address-Only With-Handshake	dxbm? ²	dxbm?	dxbm?	dxbm?	dxbm?
D08(EO)	Single-even-byte transfers					
	Byte(0) Read	dxbm? ²	dxbs?	dxbs?	DVBS	dxbs?
	Byte(2) Read	dxbm? ²	dxbs?	dxbs?	DVBS	dxbs?
	Byte(0) Write	dxbm? ²	dxbm?	dxbm?	DVBM	dxbm?
	Byte(2) Write	dxbm? ²	dxbm?	dxbm?	DVBM	dxbm?
D08(O)	Single-odd-byte transfers					
	or					
D08(O)	Byte(1) Read	dxbm? ²	dxbs?	dxbs?	dxbs?	DVBS
	Byte(3) Read	dxbm? ²	dxbs?	dxbs?	dxbs?	DVBS
	Byte(1) Write	dxbm? ²	dxbm?	dxbm?	dxbm?	DVBM
	Byte(3) Write	dxbm? ²	dxbm?	dxbm?	dxbm?	DVBM
D16	Double-byte transfers					
	Byte(0-1) Read	dxbm? ²	dxbs?	dxbs?	DVBS	DVBS
	Byte(2-3) Read	dxbm? ²	dxbs?	dxbs?	DVBS	DVBS
	Byte(0-1) Write	dxbm? ²	dxbm?	dxbm?	DVBM	DVBM
	Byte(2-3) Write	dxbm? ²	dxbm?	dxbm?	DVBM	DVBM
D32	Quad-byte transfers					
	Byte(0-3) Read	dxbm? ²	DVBS	DVBS	DVBS	DVBS
	Byte(0-3) Write	dxbm? ²	DVBM	DVBM	DVBM	DVBM
MD32	Multiplexed quad-byte transfers					
	Quad-Byte Read	DVBS ²	dxbs?	dxbs?	DVBS	DVBS
	Quad-Byte Write	DVBM ²	dxbm?	dxbm?	DVBM	DVBM

Table 2 - 21 (continued)
Use Of The Data Lines To Transfer Data

Mnemonic	Type of Cycles	A[31..1] LWORD*	D[31..24]	D[23..16]	D[15..8]	D[7..0]
D08(EO):	Single-byte block transfers					
BLT	Single-Byte Block Read	dxbm? ²	dxbs?	dxbs?	Note 1	Note 1
	Single-Byte Block Write	dxbm? ²	dxbm?	dxbm?	Note 1	Note 1
D16:BLT	Double-byte block transfers					
	Double-Byte Block Read	dxbm? ²	dxbs?	dxbs?	DVBS	DVBS
	Double-Byte Block Write	dxbm? ²	dxbm?	dxbm?	DVBM	DVBM
D32:BLT	Quad-byte block transfers					
	Quad-Byte Block Read	dxbm? ²	DVBS	DVBS	DVBS	DVBS
	Quad-Byte Block Write	dxbm? ²	DVBM	DVBM	DVBM	DVBM
MD32:BLT	Quad-byte multiplexed block transfers					
	Quad-Byte Block Read	DVBS ²	dxbs?	dxbs?	DVBS	DVBS
	Quad-Byte Block Write	DVBM ²	dxbm?	dxbm?	DVBM	DVBM
MBLT	Eight-byte block transfers					
	Eight-Byte Block Read	DVBS	DVBS	DVBS	DVBS	DVBS
	Eight-Byte Block Write	DVBM	DVBM	DVBM	DVBM	DVBM
D08(EO):	Single-byte Read-Modify-Write transfers					
RMW	Byte(0) RMW	dxbm? ²	dxbb?	dxbb?	DVBB	dxbb?
	Byte(1) RMW	dxbm? ²	dxbb?	dxbb?	dxbb?	DVBB
	Byte(2) RMW	dxbm? ²	dxbb?	dxbb?	DVBB	dxbb?
	Byte(3) RMW	dxbm? ²	dxbb?	dxbb?	dxbb?	DVBB

Table 2 - 21 (continued)
Use Of The Data Lines To Transfer Data

Mnemonic	Type of Cycles	A[31..1] LWORD*	D[31..24]	D[23..16]	D[15..8]	D[7..0]
D16:	Double-byte Read-Modify-Write transfers					
RMW	Byte(0-1) RMW	dxbm? ²	dxbb?	dxbb?	DVBB	DVBB
	Byte(2-3) RMW	dxbm? ²	dxbb?	dxbb?	DVBB	DVBB
D32:	Quad-byte Read-Modify-Write transfers					
RMW	Byte(0-3) RMW	dxbm? ²	DVBB	DVBB	DVBB	DVBB
MD32:	Multiplexed quad-byte Read-Modify-Write transfers					
RMW	Quad-Byte RMW	DVBB ²	dxbb?	dxbb?	DVBB	DVBB
D32:UAT	Unaligned transfers					
	Byte(0-2) Read	dxbm? ²	DVBS	DVBS	DVBS	dxbs?
	Byte(1-3) Read	dxbm? ²	dxbs?	DVBS	DVBS	DVBS
	Byte(1-2) Read	dxbm? ²	dxbs?	DVBS	DVBS	dxbs?
	Byte(0-2) Write	dxbm? ²	DVBM	DVBM	DVBM	dxbm?
	Byte(1-3) Write	dxbm? ²	dxbm?	DVBM	DVBM	DVBM
	Byte(1-2) Write	dxbm? ²	dxbm?	DVBM	DVBM	dxbm?

NOTES:

- (1) During single-byte block transfers, data is transferred 8 bits at a time over D07-D00 or D15-D08. A single-byte block read example is given below:

	D[15..8]	D[7..0]
First data transfer	DVBS	dxbs?
	dxbs?	DVBS
	DVBS	dxbs?
	dxbs?	DVBS
	DVBS	dxbs?
	dxbs?	DVBS
Last data transfer	DVBS	dxbs?

- (2) A[15..1] and LWORD* are used to carry data only during MBLT and MD32 cycles. In addition, A[31..16] are used to carry data only during MBLT cycles. A[31..16] may be driven by Slave during MD32 read cycles and by Master during MD32 write cycles; however, they carry no valid information.

Table 2 - 22 Master, Slave, And Location Monitor Timing Parameters

Timing Parameter Number	Master (See also Table 2-24)		Slave (See also Table 2-25)		Location Monitor (See also Table 2-26)	
	MIN	MAX	MIN	MAX	MIN	MAX
1	0					
2	0					
3	60					
4	35		10		10	
5	40		30		30	
6	0		0			
7	0		0			
8	35		10		10	
9	0		0			
10	0		-10		-10	
11	40		30		30	
12	35		10		10	
13		10		20		20
14	0		0			
15	0		0			
16	0		0			
17	40		30		30	
18	0		0			
18A	0		0			
19	40		30		30	
20	0		0			
21	0		0			
22	0		0			
23	10		0		0	
24A	0					
24B	0					
25		25				
26	0		0			
27	-25		0			
28	30	2T	30			
28A			30			
28B			0			
29	0		0			
30	0		0			
30A			0	30		
31	0		0			
32			10		10	
33			30		30	
44	30		30			
44A			30			
45	225	2T	225			
46	0		0	30		
46A			0	30		
47	0	200	0	225		
48	0		0			

NOTES:

- (1) All times are in nanoseconds
- (2) T = time-out value in microseconds

Table 2 - 23 Bus-Timer Timing Parameters
(See also Table 2 - 27)

Parameter Number	Bus Timer	
	MIN	MAX
28	T	2T
30	0	

NOTE: T = time-out value in microseconds.

Table 2 - 24 Master, Timing RULEs and OBSERVATIONS

Note: The numbers correspond to the timing parameters specified in Table 2-22.

1. **RULE 2.27:**
When taking control of the DTB, the Master MUST NOT drive any of the IACK*, AM[5..0], A[31..1], LWORD*, D[31..0], WRITE*, DS0*, DS1*, or AS* lines until after the previous Master or interrupt handler allows AS* to rise above the low level.

OBSERVATION 2.35:
Chapter 3 describes how a Master's requester is granted use of the DTB.
2. **RULE 2.28:**
When taking control of the DTB, the Master MUST NOT drive any of the IACK*, AM[5..0], A[31..1], LWORD*, D[31..0], WRITE*, DS0*, DS1*, or AS* lines until after its requester is granted the bus.

OBSERVATION 2.36:
Chapter 3 describes how a Master's requester is granted use of the DTB.
3. **RULE 2.29:**
When taking control of the DTB, the Master MUST NOT drive AS* low until this time after the previous Master or interrupt handler allows AS* to rise above the low level.

OBSERVATION 2.37:
Timing parameter 3 ensures that timing parameter 5 for Slaves is guaranteed when there is an interchange of the DTB Mastership.
4. **RULE 2.30:**
The Master MUST NOT drive AS* low until IACK* has been high, and the required lines among A[31..1], AM[5..0], and LWORD* have been valid for this minimum time.

OBSERVATION 2.38:
Table 19 specifies the lines among A[31..1] that the Master is required to drive. Table 3 specifies how it uses AM[5..0], and Table 20 how it uses LWORD*.
5. **RULE 2.31:**
When using the DTB for two consecutive cycles, the Master MUST NOT drive AS* low until it has been high for this minimum time.
6. **RULE 2.32:**
After all read cycles, the Master MUST NOT drive any of D[31..0] until both DTACK* and BERR* are high.
7. **RULE 2.33a:**
During all read cycles except during A64, A40 and MBLT, and except during the A40BLT address phase, the Master MUST NOT drive DSA* low until it has released all of the D[31..0] lines. In addition, during MBLT and MD32 read cycles and when a data transfer follows A64 and A40 addressing, the Master MUST NOT drive DSA* low until it has released all the D[31..0], A[31..1], and LWORD* lines.

Table 2 - 24. Master, Timing RULEs and OBSERVATIONS (continued)

8. **RULE 2.34a:**
During write cycles and when using A64 or A40 addressing, the Master MUST NOT drive DSA* low until the required lines among D[31..0] have been valid for this minimum time. In addition, during MBLT, A40BLT, and MD32 write cycles, the Master MUST NOT drive DSA* low until A[31..1] and LWORD* have been valid for this minimum time.
- OBSERVATION 2.39a:**
Table 2-21 specifies the lines among D[31..0] that a Master is required to drive. Table 2-19 specifies how A64 Master and A40 Master use D[31..0] for A64 and A40 addressing.
9. **RULE 2.35:**
The Master MUST NOT drive DSA* low until both DTACK* and BERR* are high.
10. **RULE 2.36:**
The Master MUST NOT drive DSA* low until it has driven AS* low.
11. **RULE 2.37:**
The Master MUST NOT drive DSA* low until DS0* and DS1* have both been simultaneously high for this minimum time.
12. **RULE 2.38:**
The Master MUST NOT drive DSA* low for the first time until WRITE* has been valid for this minimum time.
13. **RULE 2.39:**
During cycles where the Master drives both DS0* and DS1* low, it MUST drive DSB* low within this maximum time after it drives DSA* low.
- OBSERVATION 2.40:**
Timing parameter 13 does not apply to transfers where either DS0* is driven low or DS1* is driven low, but not both.
14. **RULE 2.40:**
During all data transfer cycles, excluding A16, A24 and A32 Read-Modify-Write and Retry cycles, the Master MUST drive A[31..1] and LWORD* with valid addressing information until it detects the first falling edge on DTACK* or BERR*.
- OBSERVATION 2.41:**
During all data transfer cycles, except block transfer cycles and all Read-Modify-Write cycles, there will be only one falling edge on DTACK* or BERR*.
15. **RULE 2.41:**
During all Read-Modify-Write cycles, the Master MUST hold A[31..1] valid and maintain the appropriate level on LWORD* until it detects the second falling edge on DTACK* or BERR*.

Table 2 - 24 Master, Timing RULEs and OBSERVATIONs (continued)

16. **RULE 2.42:**
During all data transfer cycles, the Master MUST maintain a valid AM code, and ensure that IACK* stays high until it detects the last falling edge on DTACK* or BERR*.
- OBSERVATION 2.42:**
During all data transfer cycles, except block transfer and Read-Modify-Write cycles, there will be only one falling edge on DTACK* or BERR*.
17. **RULE 2.43:**
The Master MUST NOT change the levels on IACK*, A[31..1], AM[5..0], or LWORD* for this minimum time after it drives AS* low.
18. **RULE 2.44a:**
During all data transfer cycles except MBLT, A40BLT, and MD32 read cycles, the Master MUST hold AS* low until it detects the last falling edge on DTACK* or BERR*.
- 18A. **RULE 2.94:**
During MBLT, A40BLT, and MD32 read cycles, the Master MUST hold AS* low until it detects the last rising edge on DTACK* or BERR*.
19. **RULE 2.45:**
The Master MUST hold AS* low for this minimum time.
20. **RULE 2.46a:**
Once a Master has driven DSA* low, it MUST maintain that line low until it detects the falling edge on DTACK*, RETRY*, or BERR*.
- OBSERVATION 2.102**
During a read cycle, data may not be valid at the bus Master for up to 25 nanoseconds after DTACK* is detected low. (See parameter 27). To ensure that valid data is read, the bus Master should maintain DSA* asserted for a least 25 nanoseconds after detecting DTACK* low.
21. **RULE 2.47a:**
Once a Master has driven DSB* low, it MUST maintain that line low until it detects the falling edge on DTACK*, RETRY*, or BERR*.
- See OBSERVATION 2.102
22. **RULE 2.48a:**
During all write cycles and during A64 or A40 addressing, once the Master has driven DSA* low, it MUST NOT change any of D[31..0] until it detects DTACK*, BERR* or RETRY* low. In addition, during MBLT and MD32 write data transfers, once the Master has driven DSA* low, it MUST NOT change any of A[31..1] and LWORD* until it detects DTACK* , BERR* or RETRY* low.

Table 2 - 24 Master, Timing RULEs and OBSERVATIONs (continued)

23. **RULE 2.49:**
 Once a Master has driven DSA* low, it MUST NOT change the level of WRITE* until this minimum time after DSB* is high.
- 24A. **RULE 2.50:**
 IF the Master drives or releases AS* high after its requester releases BBSY*,
 THEN it MUST release IACK*, AM[5..0], A[31..1], LWORD*, D[31..0], WRITE*, DS0*, and DS1* before allowing AS* to rise above the low level.
- OBSERVATION 2.43:**
 Section 3 describes how a Master's requester releases the BBSY* line.
- 24B. **RULE 2.51:**
 IF the Master drives or releases AS* high before its requester releases BBSY*,
 THEN it MUST release AS*, IACK*, AM[5..0], A[31..1], LWORD*, D[31..0], WRITE*, DS0*, and DS1* before allowing its requester to release BBSY*.
- OBSERVATION 2.44:**
 Section 3 describes how a Master's requester releases the BBSY* line.
25. **RULE 2.52:**
 When the Master drives or releases AS* high after its requester releases BBSY*, then it MUST release AS* within this time after allowing it to rise above the low level.
- OBSERVATION 2.45:**
 Section 3 describes how a Master's requester releases the BBSY* line.
26. **OBSERVATION 2.46a:**
 During all read cycles except A64, A40, MBLT and A40BLT address phase, the Master is guaranteed that the responding Slave will not drive D[31..0] until the Master drives DSA* low. In addition, during MBLT and A40BLT read cycles and all A64 and A40 read cycles, the Master is guaranteed that the responding Slave will not drive any of D[31..0], A[31..1], and LWORD* lines until the Master drives DSA* low for the second time.
27. **OBSERVATION 2.47a:**
 During read cycles, the Master is guaranteed that D[31..0] will be valid within 25 nanoseconds after DTACK* goes low. In addition, during MBLT, A40BLT, and MD32 read cycles, the Master is guaranteed that A[31..1] and LWORD* will be valid within 25 nanoseconds after DTACK* goes low. This time does not apply to cycles where the Slave drives BERR* low instead of DTACK*.
28. **OBSERVATION 2.48:**
 The Master is guaranteed that neither DTACK* nor BERR* will go low until this minimum time after it drives DSA* low. The bus timer guarantees the Master that if DTACK* is not detected low after its time-out period has elapsed, and within twice its time-out period, then the bus timer will drive BERR* low.

Table 2 - 24 Master, Timing RULEs and OBSERVATIONS (continued)

29. OBSERVATION 2.49a:
During read cycles, the Master is guaranteed that D[31..0] remain valid until it drives DSA* high. In addition, during MBLT, A40BLT, and MD32 read cycles, the Master is guaranteed that A[31..1] and LWORD* remain valid until it drives DSA* high.
30. OBSERVATION 2.50:
The Master is guaranteed that DTACK*, and BERR* will not go high until it drives both DS0* and DS1* high.
31. OBSERVATION 2.51a:
During all read cycles, the Master is guaranteed that D[31..0] lines are released by the time it detects the rising edge of DTACK*, RETRY* or BERR*. In addition, during MBLT, A40BLT, and MD32 read cycles, the Master is guaranteed that A[31..1] and LWORD* lines have been released by the time it detects the rising edge of DTACK*, BERR* or RETRY*.
44. OBSERVATION 2.103:
The Master is guaranteed that RETRY* will not go low until this minimum time after it drives DSA* low.
45. OBSERVATION 2.104:
IF the Slave drives RETRY* low,
THEN the Master is guaranteed that neither DTACK* nor BERR* will be driven low until this minimum time after RETRY* has been driven low.
46. OBSERVATION 2.105:
The Master is guaranteed that RETRY* will not go high until it drives both DS0* and DS1* high.
47. RULE 2.95:
IF the Master responds to the falling edge on RETRY*,
THEN it MUST drive the data strobes high within this specified time after detecting the falling edge on RETRY*.
48. RULE 2.96:
IF the Master responds to RETRY* by driving data strobes high within the time specified by Timing Parameter 47 after detecting the falling edge on RETRY*,
THEN it MUST NOT drive or release AS* high and MUST NOT change the levels on IACK* and AM[5..0] until it detects the rising edge on RETRY*. In addition, during address phase, it MUST NOT change the levels on LWORD* and the address lines, including D[31..0] for A40 and A64 addressing, until it detects the rising edge on RETRY*.
- OBSERVATION 2.106:
Since RETRY* may not be monitored by all Masters and Slaves in the system, the Master responding to RETRY* must guarantee that the data lines have been released prior to the transfer of bus mastership. This is guaranteed by timing parameters 31 and 48.

Table 2 - 25 Slave, Timing RULEs and OBSERVATIONS

Note: The numbers correspond to the timing parameters specified in Table 2-22.

4. OBSERVATION 2.52:
All Slaves are guaranteed that the IACK*, A[31..1], AM[5..0], and LWORD* lines have been valid for this minimum time when they detect a falling edge on AS*.
5. OBSERVATION 2.53:
All Slaves are guaranteed this minimum high time on AS* between DTB cycles.
6. OBSERVATION 2.54:
During read cycles, the responding Slave is guaranteed that none of the D[31..0] lines will be driven by any other module until it releases DTACK* and BERR* to high.
7. OBSERVATION 2.55a:
During all read cycles except A64, A40, MBLT and A40BLT address phase, the responding Slave is guaranteed that the D[31..0] lines are released by the time DSA* goes low. In addition, during MBLT and A40BLT read cycles and when a data transfer follows A64 and A40 addressing, the responding Slave is guaranteed that the D[31..0], A[31..1], and LWORD* lines are released by the time DSA* goes low.
8. OBSERVATION 2.56a:
During write cycles and during A64 or A40 addressing, all Slaves are guaranteed that the D[31..0] lines have been valid for this minimum time when they detect a falling edge on DSA*. In addition, during MBLT and MD32 write cycles, all Slaves are guaranteed that the A[31..1] and LWORD* lines have been valid for this minimum time when they detect a falling edge on DSA*.
9. OBSERVATION 2.57:
The responding Slave is guaranteed that neither DS0* nor DS1* will go low until it drives DTACK* and BERR* high.
10. OBSERVATION 2.58:
Due to bus skew, Slaves on the DTB might detect a falling edge on DSA* before detecting the falling edge on AS*. However, Slaves are guaranteed that a falling edge on DSA* will not precede the falling edge on AS* by more than this time.
11. OBSERVATION 2.59:
Slaves are guaranteed this minimum time during which both DS0* and DS1* are simultaneously high.
12. OBSERVATION 2.60:
Slaves are guaranteed that WRITE* has been valid for this minimum time before the first falling edge on DSA*.
13. OBSERVATION 2.61:
IF both data strobes are going to be driven low,
THEN the responding Slave is guaranteed that DSB* will go low within this maximum time after DSA* has gone low.

Table 2 - 25 Slave, Timing RULES and OBSERVATIONS (continued)

14. OBSERVATION 2.62:
During all DTB cycles, excluding A16, A24 and A32 Read-Modify-Write cycles, the responding Slave is guaranteed that A[31..1] and LWORD* carry valid addressing information until it drives DTACK* or BERR* low for the first time, provided that it does so within the bus time-out period.
15. OBSERVATION 2.63:
During all Read-Modify-Write cycles, the responding Slave is guaranteed that the A[31..1] and LWORD* lines remain valid until it drives DTACK* or BERR* low for the second time, provided that it does so within the bus time-out period.
16. OBSERVATION 2.64:
The responding Slave is guaranteed that IACK* and AM[5..0] remain valid until it drives DTACK* or BERR* low for the last time, provided that it does so within the bus time-out period.
17. OBSERVATION 2.65:
Slaves are guaranteed that IACK*, A[31..1], AM[5..0], and LWORD* remain valid for this minimum time after the falling edge of AS*. During address-only cycles this time is guaranteed by the Master. During all other cycle types, this time is derived from timing parameters 10, 14, 16, and 28.
18. OBSERVATION 2.66a:
During all data transfer cycles except MBLT, A40BLT, and MD32 read cycles, the responding Slave is guaranteed that AS* will remain low until it drives DTACK* or BERR* low for the last time, provided that it does so within the bus time-out period.
- 18A. OBSERVATION 2.107:
During MBLT, A40BLT, and MD32 read cycles, the responding Slave is guaranteed that AS* will remain low until it drives DTACK* or BERR* high for the last time, provided that it does so within the bus time-out period.
19. OBSERVATION 2.67:
Slaves are guaranteed that AS* remains low for this minimum time.
20. OBSERVATION 2.68a:
The responding Slave is guaranteed that once DSA* goes low, it remains low until it drives DTACK*, RETRY*, or BERR* low, provided that the Slave does so within the bus time-out period.
21. OBSERVATION 2.69a:
The responding Slave is guaranteed that once DSB* goes low, it remains low until it drives DTACK*, RETRY*, or BERR* low, provided that the Slave does so within the bus time-out period.
22. OBSERVATION 2.70a:
During all write cycles and during A64 or A40 addressing, the responding Slave is guaranteed that D[31..0] remain valid until it drives DTACK* or BERR* low, provided that it does so within the bus time-out period. In addition, during MBLT and MD32 write data transfers, the responding Slave is guaranteed that A[31..1] and LWORD* remain valid until it drives DTACK* or BERR* low, provided that it does so within the bus time-out period.

Table 2 - 25 Slave, Timing RULEs and OBSERVATIONs (continued)

23. OBSERVATION 2.71 :
The responding Slave is guaranteed that WRITE* remains valid until both DS0* and DS1* are high.
26. RULE 2.53a:
During all read cycles except during A64, A40, MBLT, or A40BLT address phase, the responding Slave MUST NOT drive the D[31..0] lines until DSA* goes low. In addition, during MBLT and A40BLT read cycles and all A64 and A40 read cycles, the responding Slave MUST NOT drive any of D[31..0], A[31..1], and LWORD* lines until DSA* goes low for the second time.
27. RULE 2.54a:
During all read cycles, the responding Slave MUST NOT drive DTACK* low before it drives D[31..0] with valid data, except as noted in Timing Parameter 26. In addition, during MBLT and A40BLT read cycles, the responding Slave MUST NOT drive DTACK* low before it also drives A[31..1] and LWORD* with valid data.
- OBSERVATION 2.72:
RULE 2.54a does not apply to cycles where the responding Slave drives BERR* low instead of DTACK*.
28. RULE 2.55:
The responding Slave MUST wait this minimum time after DSA* goes low before driving DTACK* or BERR* low.
- OBSERVATION 2.108:
All Slaves are guaranteed that those bus lines that carry data during write cycles remain valid for this minimum time after DSA* goes low. This time is guaranteed by the requirement that the Master maintain data valid until it receives DTACK* or BERR* low.
- 28A. RULE 2.97:
Responding Slaves MUST wait this minimum time after detecting the first falling edge of DSA* before driving DTACK* to any level.
- 28B. PERMISSION 2.19:
Except for the first DSA*, responding Slaves MAY drive DTACK* high immediately after detecting the falling edge of DSA*.
29. RULE 2.56a:
During all read cycles, once the responding Slave has driven DTACK* low, it MUST NOT change D[31..0] until DSA* goes high. In addition, during MBLT, A40BLT, and MD32 read cycles, once the responding Slave has driven DTACK* low, it MUST NOT change A[31..1] or LWORD* until DSA* goes high.
30. RULE 2.57:
Once the responding Slave has driven DTACK* or BERR* low, it MUST NOT release them or drive DTACK* high until it detects both DS0* and DS1* high.
- 30A. RULE 2.98:
When a Slave drives DTACK* high at the end of a cycle, it MUST release DTACK* within this time.

Table 2 - 25 Slave, Timing RULEs and OBSERVATIONs (continued)

31. **RULE 2.58a:**
During all read cycles, the responding Slave MUST release all of the D[31..0] lines before releasing DTACK* or BERR* high or driving RETRY* or DTACK* high. In addition, during MBLT, and MD32 read cycles, the responding Slave MUST release all of the A[31..1] and LWORD* lines before releasing DTACK* or BERR* high or driving RETRY* or DTACK* high.
32. **OBSERVATION 2.73:**
Slaves are guaranteed that IACK*, LWORD*, A[31..1], and AM[5..0] have been valid for this minimum time when they detect a falling edge on DSA*. This time is derived from timing parameters 4 and 10.
33. **OBSERVATION 2.74:**
During data transfer cycles, Slaves are guaranteed that either DS0*, or DS1*, or both remain low for at least this minimum time. This time is derived from timing parameter 28, where the responding Slave is required to wait a minimum time before driving BERR* or DTACK* low.
44. **RULE 2.99:**
The responding Slave MUST wait this minimum time after DSA* goes low before driving RETRY* low.
- 44A. **RULE 2.100:**
Responding Slaves MUST wait this minimum time after detecting the falling edge of DSA* before driving RETRY* to any level.
45. **RULE 2.101:**
Once the responding Slave drives RETRY* low, it MUST NOT drive either DTACK* or BERR* low until RETRY* has been low for this minimum time.
- RULE 2.102:**
IF a Slave has asserted RETRY* low and the Master has not driven DSB* high within the time specified by parameter 47
THEN the SLAVE MUST complete the bus cycle with either DTACK* or BERR*.
46. **RULE 2.103:**
Once the responding Slave drives RETRY* low it MUST NOT drive RETRY* high until it detects both DS0* and DS1* high.
- RULE 2.104:**
Once the responding Slave drives RETRY* low, it MUST drive RETRY* high within this time after it detects both DS0* and DS1* high.
- 46A. **RULE 2.105:**
The Slave MUST release RETRY* within this time after it has driven RETRY* high at the end of the cycle.
47. **OBSERVATION 2.109:**
IF the Master has not driven DSB* high within this time after the Slave has driven RETRY* low,
THEN the Slave is guaranteed that the Master will not respond to RETRY* and that DSA* will remain low until either DTACK* or BERR* is driven low.

Table 2 - 25 Slave, Timing RULES and OBSERVATIONS (continued)

48. OBSERVATION 2.110:
Slaves are guaranteed that AS* remains asserted and that IACK* and AM[5..0] remain valid until RETRY* is driven high. In addition, during address phase, Slaves are guaranteed that LWORD* and all address lines, including D[31..0] for A40 and A64 addressing, remain valid until RETRY* is driven high. This timing parameter guarantees that Slaves that do not monitor RETRY* will not see invalid addressing information while data strobes are low and neither DTACK* nor BERR* are low.

Table 2 - 26 Location Monitor, Timing OBSERVATIONS

Note: The numbers correspond to the timing parameters specified in Table 2-22.

4. OBSERVATION 2.75:
Location Monitors are guaranteed that IACK*, A[31..1], AM[5..0], and LWORD* have been valid for this minimum time when they detect a falling edge on AS*.
5. OBSERVATION 2.76:
Location Monitors are guaranteed this minimum high time on AS* between DTB cycles.
8. OBSERVATION 2.112:
A64 Location Monitors are guaranteed that D[31..0] have been valid for this minimum time when they detect a falling edge on DSA*.
10. OBSERVATION 2.77:
Due to bus skew, Location Monitors might detect a falling edge on DSA* before detecting the falling edge on AS*. However, they are guaranteed that the falling edge on DSA* will not precede the falling edge on AS* by more than this time.
11. OBSERVATION 2.78:
Location Monitors are guaranteed this minimum time during which both data strobes are simultaneously high between consecutive data transfers.
12. OBSERVATION 2.79:
Location Monitors are guaranteed that WRITE* has been valid for this minimum time when they detect the first falling edge on DSA*.
13. OBSERVATION 2.80:
IF both data strobes are going to be driven low,
THEN the Location Monitor is guaranteed that DSB* will go low within this maximum time after DSA* has gone low.

Table 2 - 26 Location Monitor, Timing OBSERVATIONS (continued)

17. OBSERVATION 2.81:
The Location Monitor is guaranteed that IACK*, A[31..1], AM[5..0], and LWORD* will remain valid for this minimum time after the falling edge of AS*. During Address-Only cycles this time is guaranteed by the Master. During all other cycle types, this time is derived from timing parameters 10, 14, 16, and 28.
19. OBSERVATION 2.82:
The Location Monitor is guaranteed that AS* will remain low for this minimum time.
23. OBSERVATION 2.83:
The Location Monitor is guaranteed that the WRITE* line will remain valid until both data strobes go high.
32. OBSERVATION 2.84:
The Location Monitor is guaranteed that IACK*, LWORD*, A[31..1], and AM[5..0] have been valid for this minimum time when it detects a falling edge on DSA*.
33. OBSERVATION 2.85:
During data transfer cycles, the Location Monitor is guaranteed that DS0* and/or DS1* will remain low for at least this minimum time. This time is derived from timing parameter 28, where the responding Slave is required to wait a minimum time before driving BERR* or DTACK* to low.

Table 2 - 27 BUS TIMER, Timing RULES

Note: The numbers correspond to the timing parameters specified in Table 2-22.

28. RULE 2.59:
The bus timer MUST wait at least its time-out time, but no longer than twice its time-out time, after the first data strobe goes low before driving BERR* low.
30. RULE 2.60:
Once it has driven BERR* low, the bus timer MUST NOT release BERR* until it detects both DS0* and DS1* high.
- RECOMMENDATION 2.13
To ensure predictable system operation this time-out value should be longer than longest Slave response time in the system.

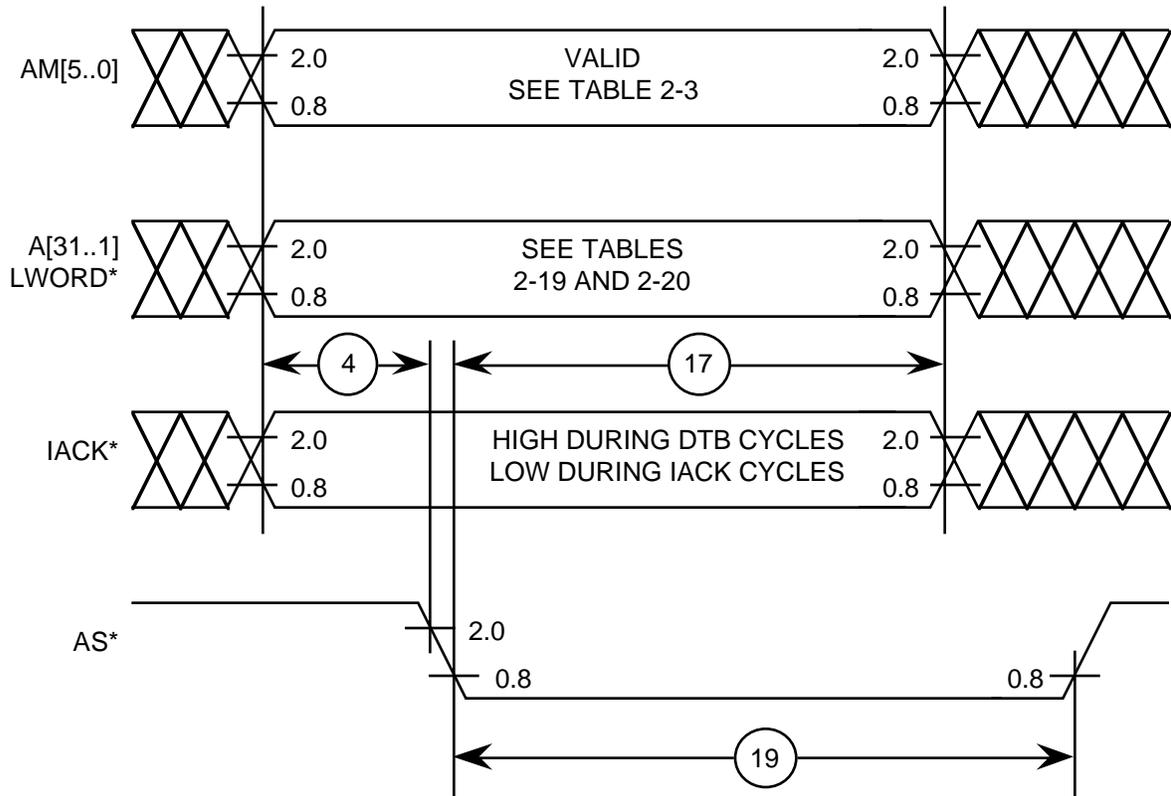


Figure 2 - 12 Address Broadcast Timing - All Cycles

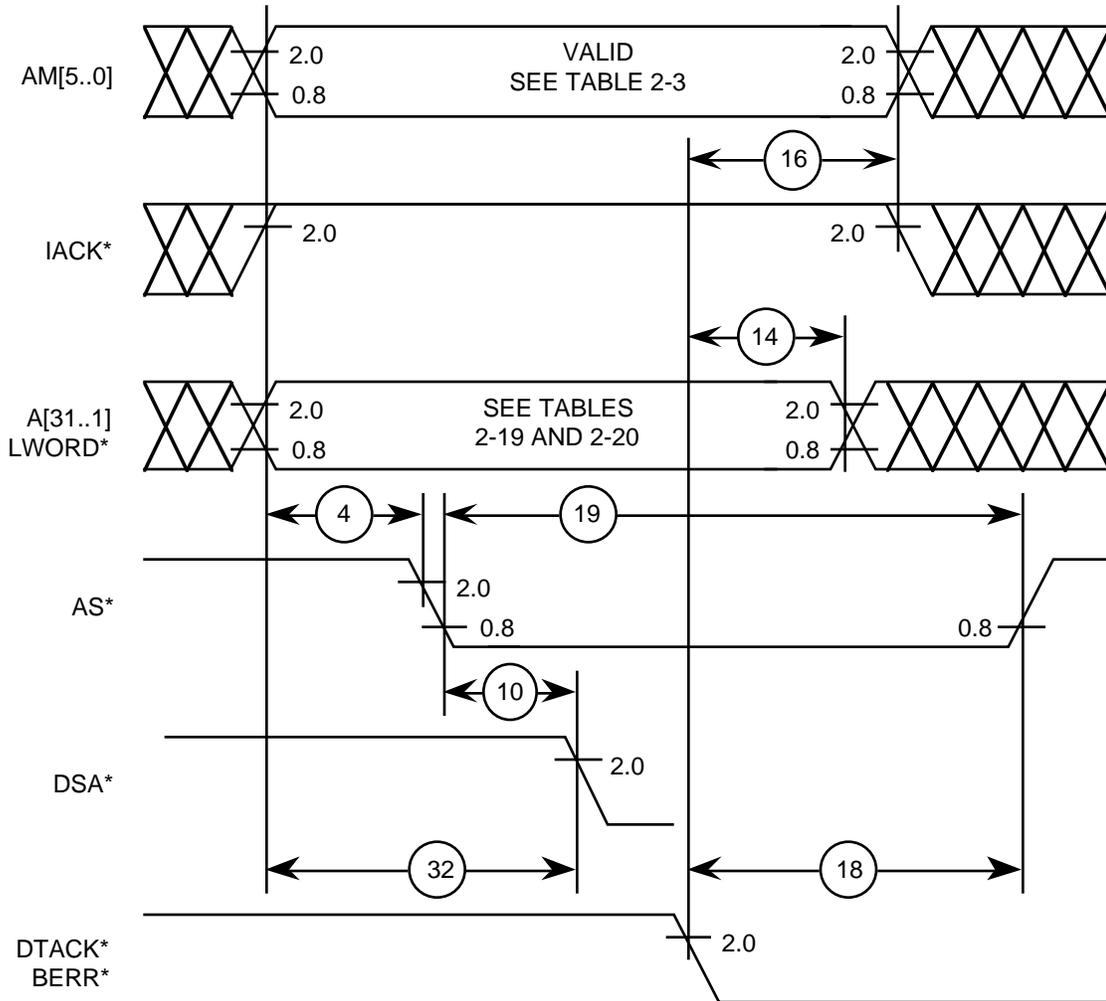


Figure 2 - 13 A16, A24, A32 Master, Responding Slave, And Location Monitor
 Address Broadcast Timing
 Single-Even-Byte Transfers
 Single-Odd-Byte Transfers
 Double-Byte Transfers
 Quad-Byte Transfers
 Unaligned Transfers

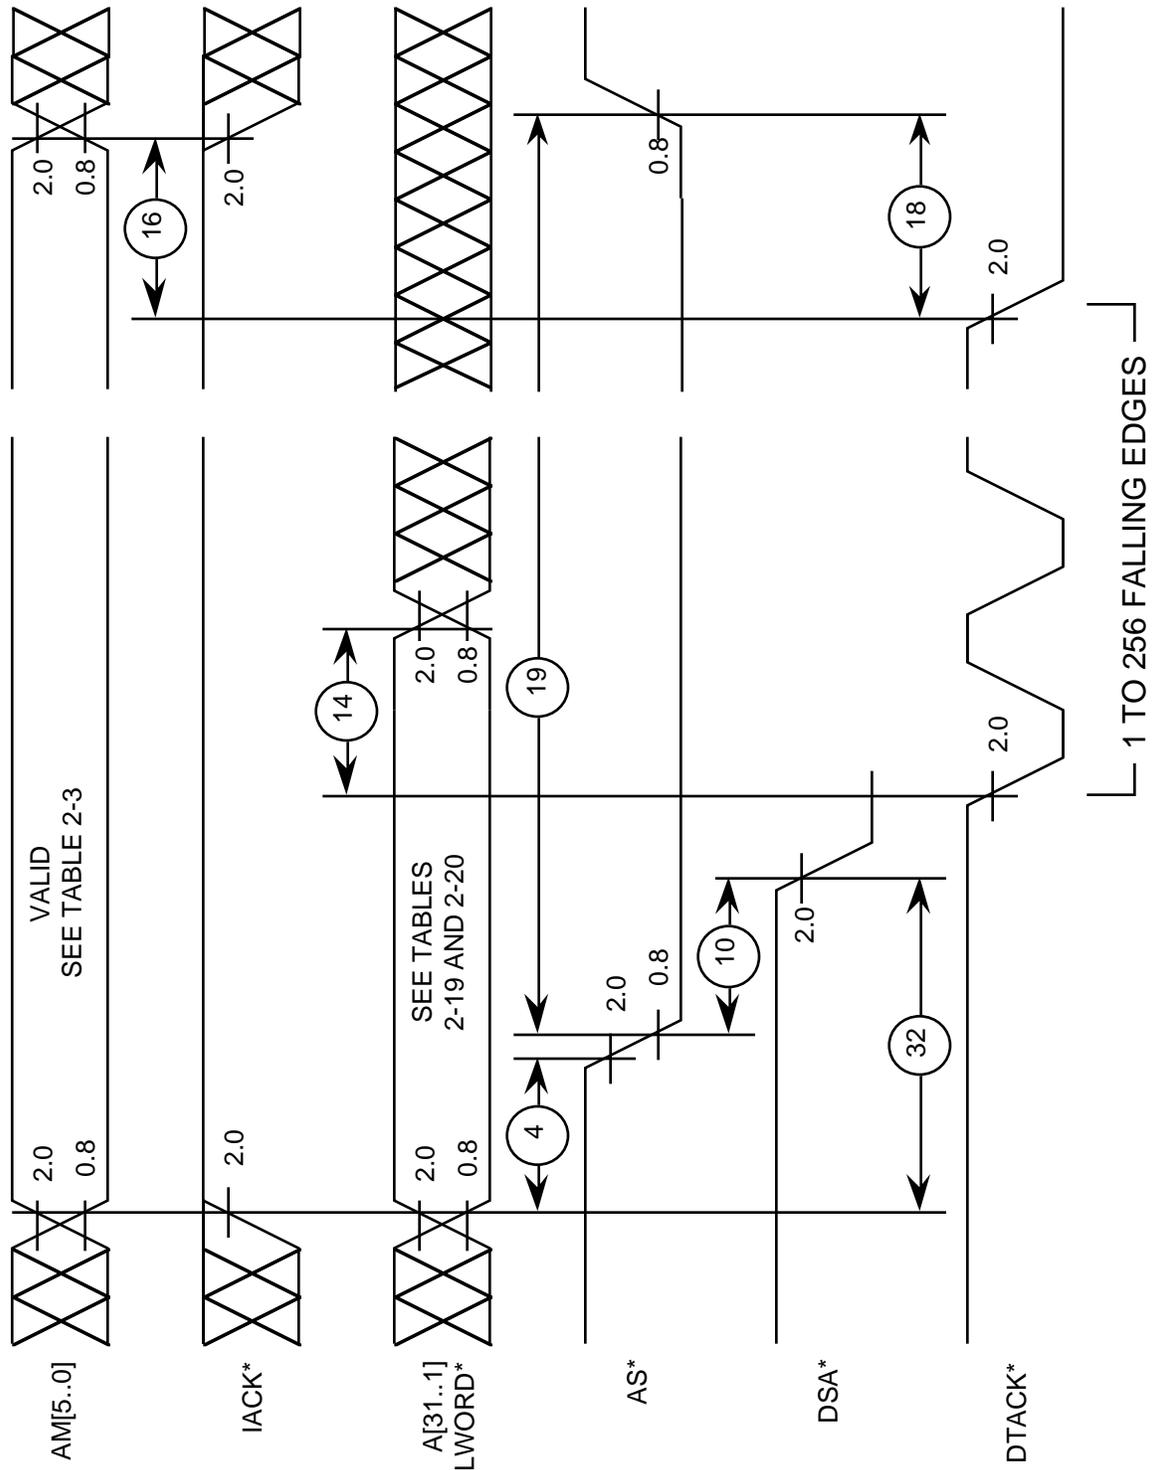


Figure 2 - 14 Master, Slave, and Location Monitor - A16, A24 and A32 Address Broadcast Timing
 Single-Byte Block Transfers
 Double-Byte Block Transfers
 Quad-Byte Block Transfers

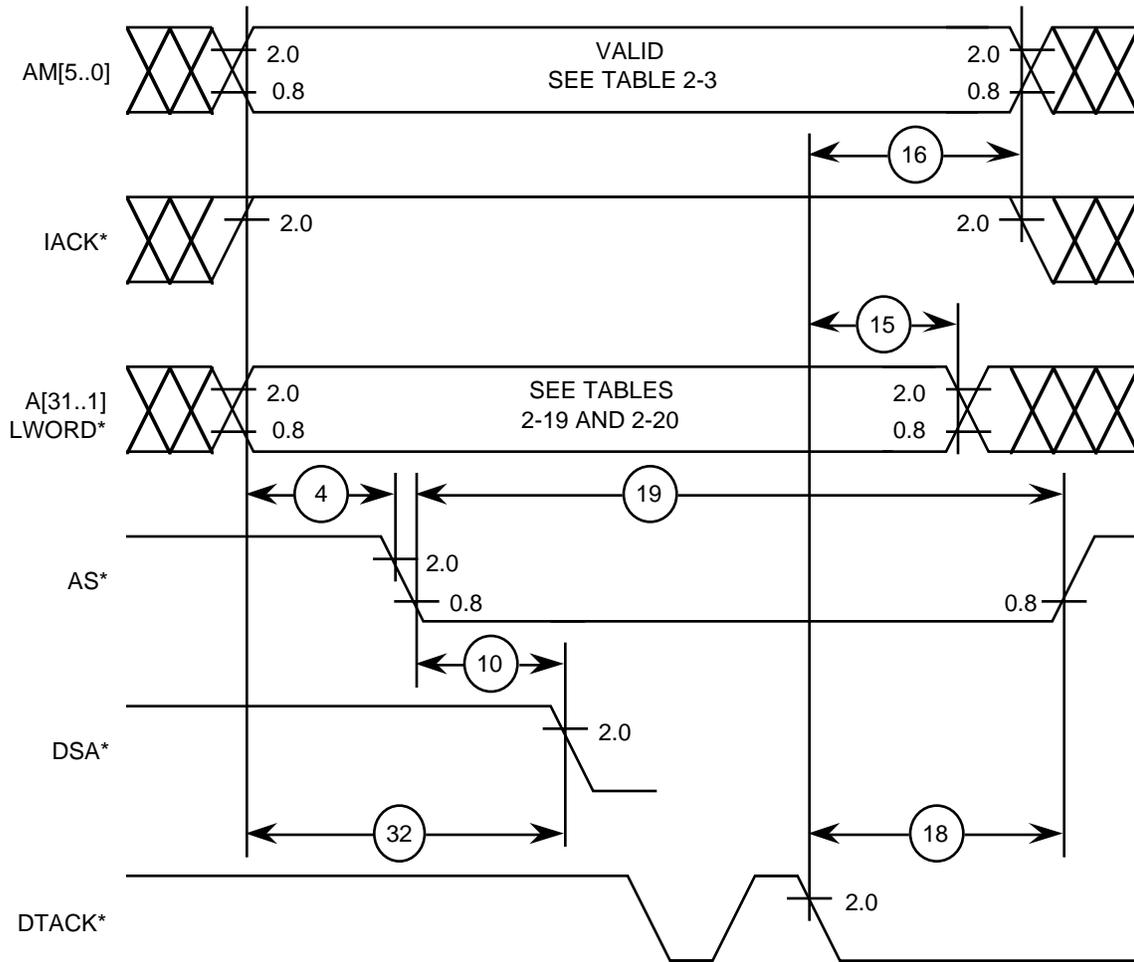


Figure 2 - 15 Master, Slave, And Location Monitor A16, A24, And A32 Address
 Broadcast Timing
 Single-Byte RMW Cycles
 Double-Byte RMW Cycles
 Quad-Byte RMW Cycles

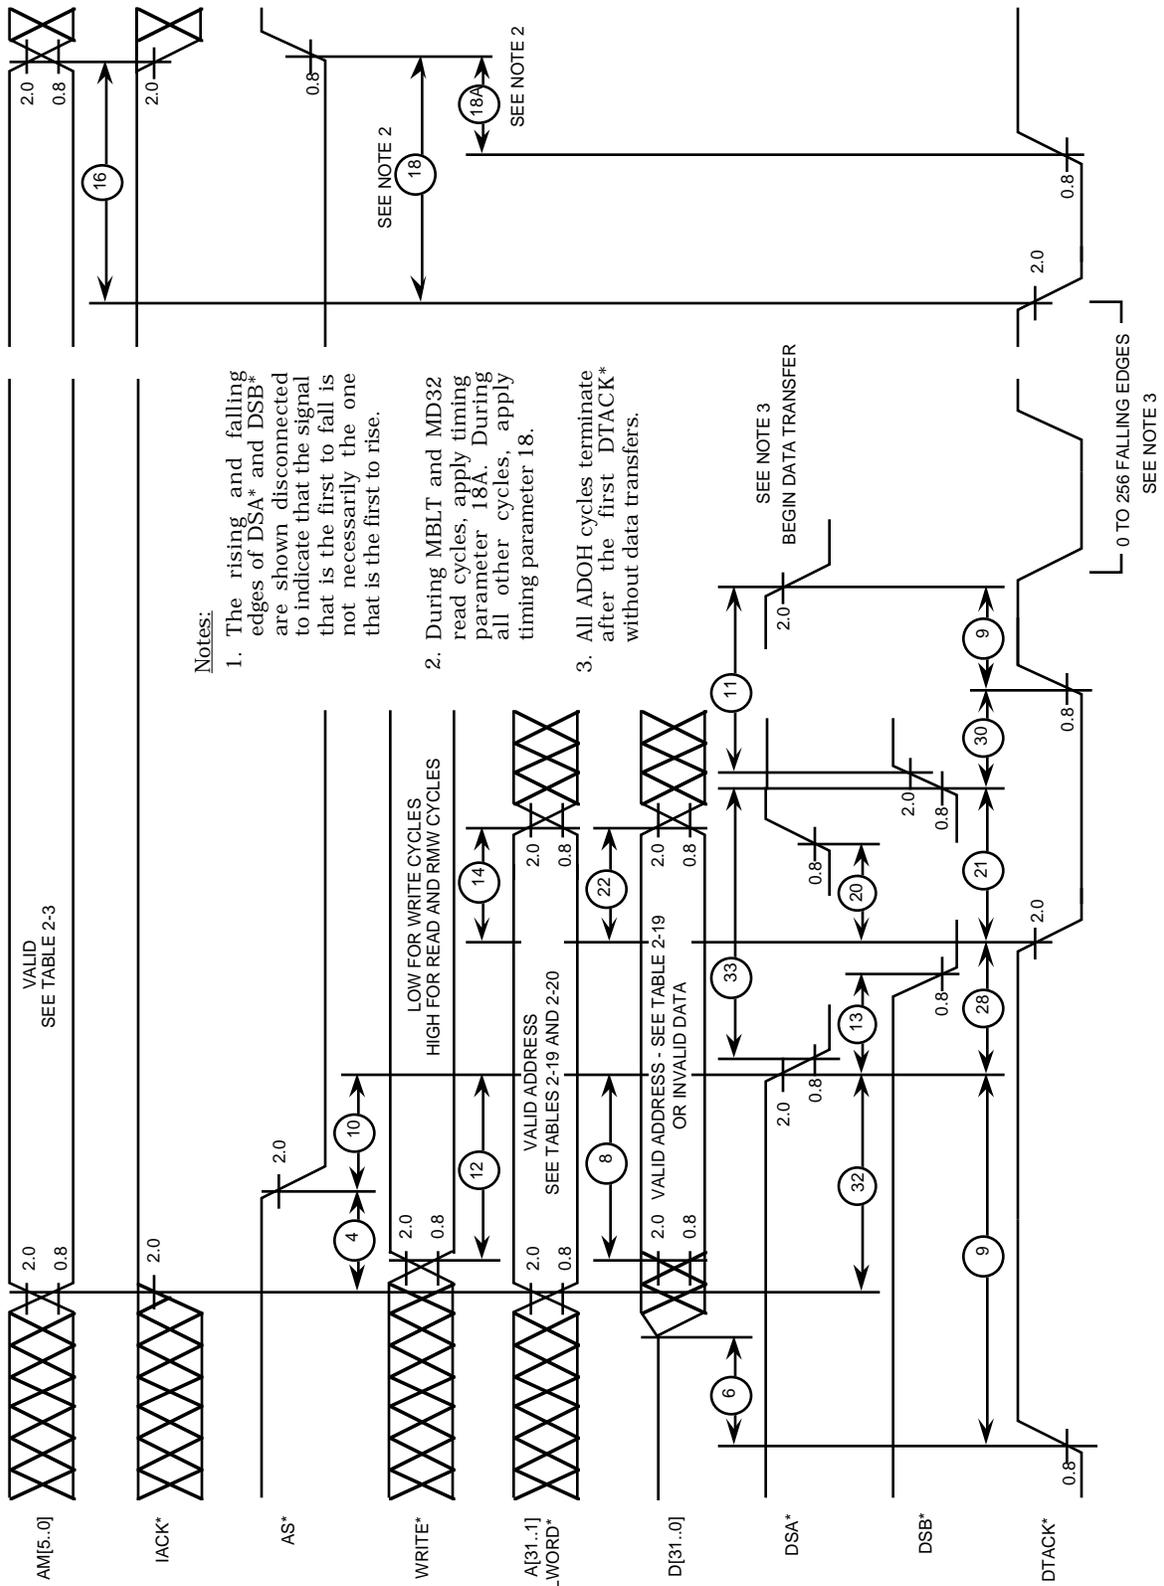


Figure 2 - 27 Master, Slave And Location Monitor - A64, A40, And ADOH Address Broadcast Timing All Data Transfer Cycles

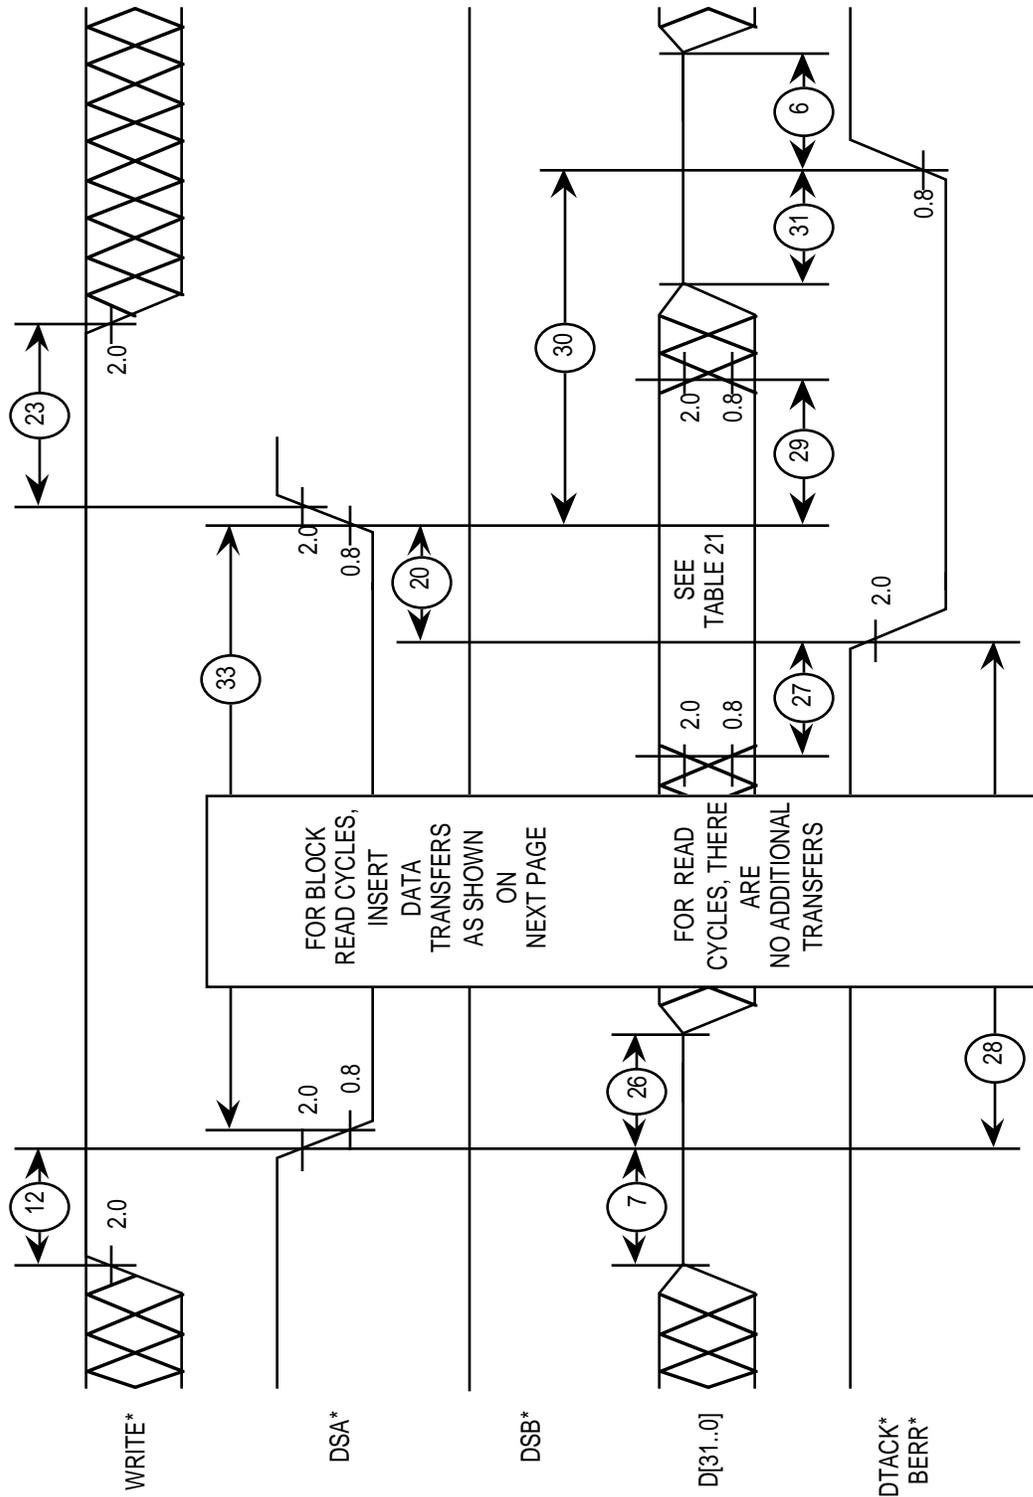


Figure 2 - 16 Master, Slave, And Location Monitor Data Transfer Timing
 Single Byte Read
 Byte(0-2) UAT Read
 Byte(1-3) UAT Read
 Single-Byte Block Read

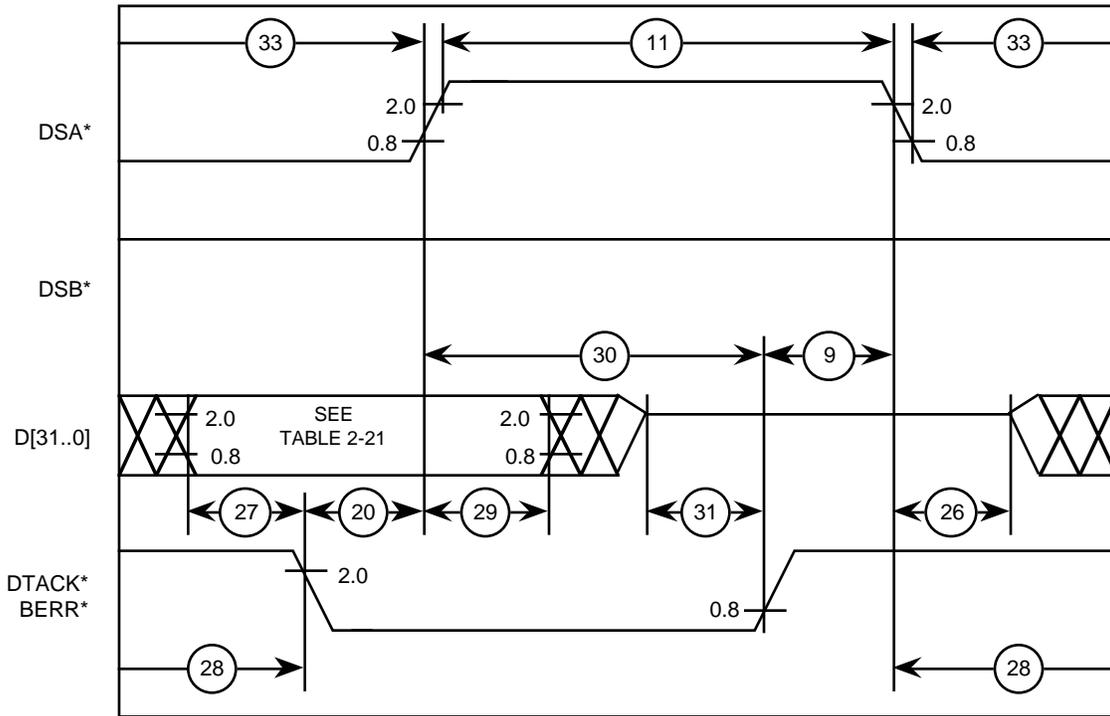
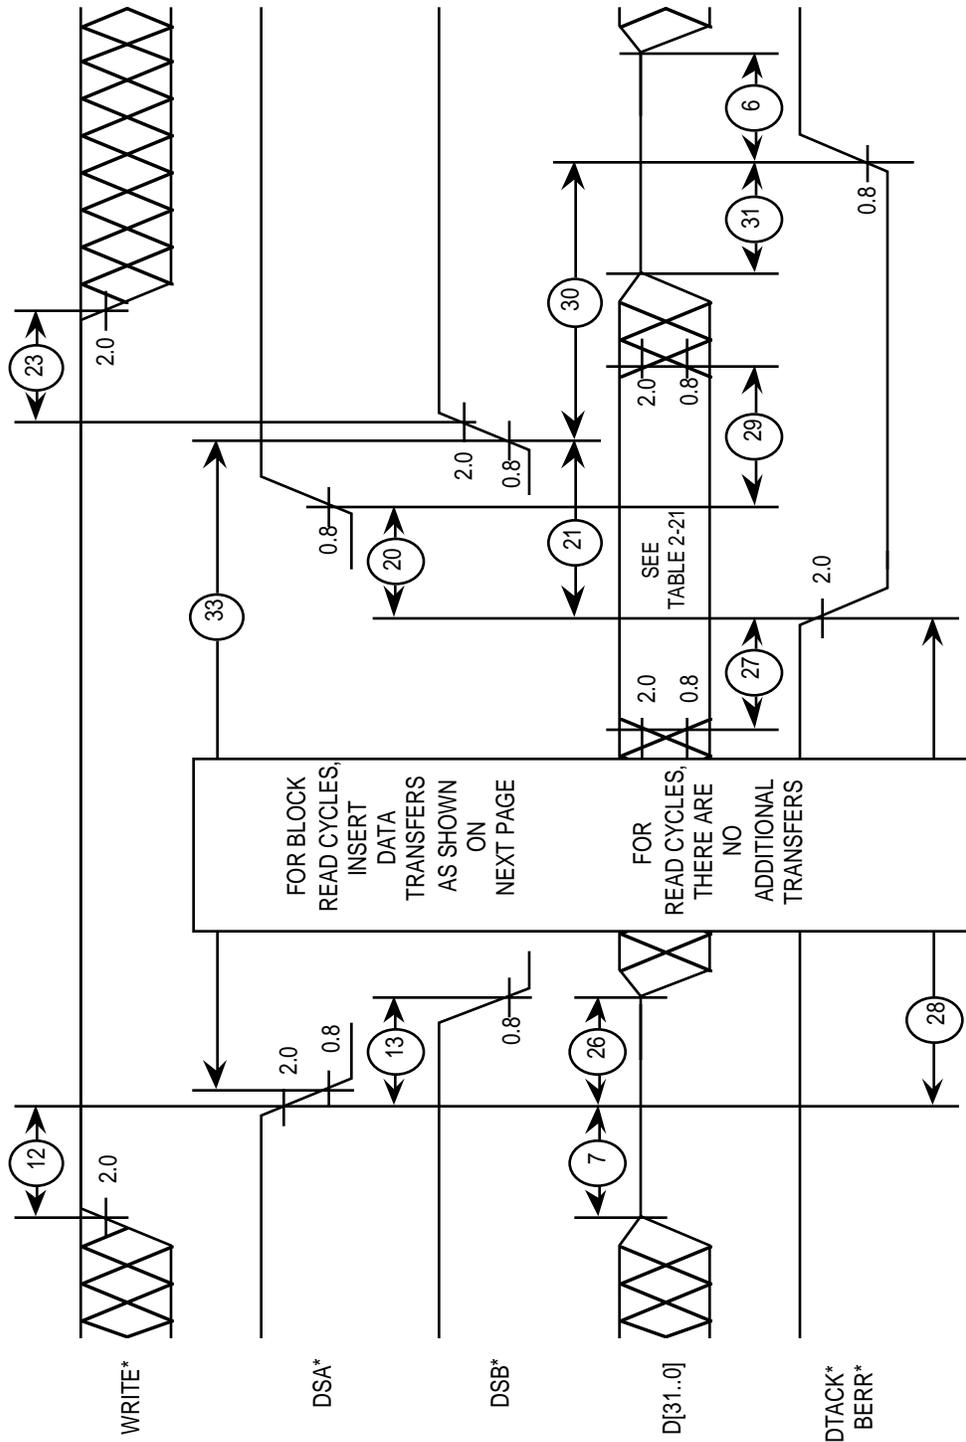
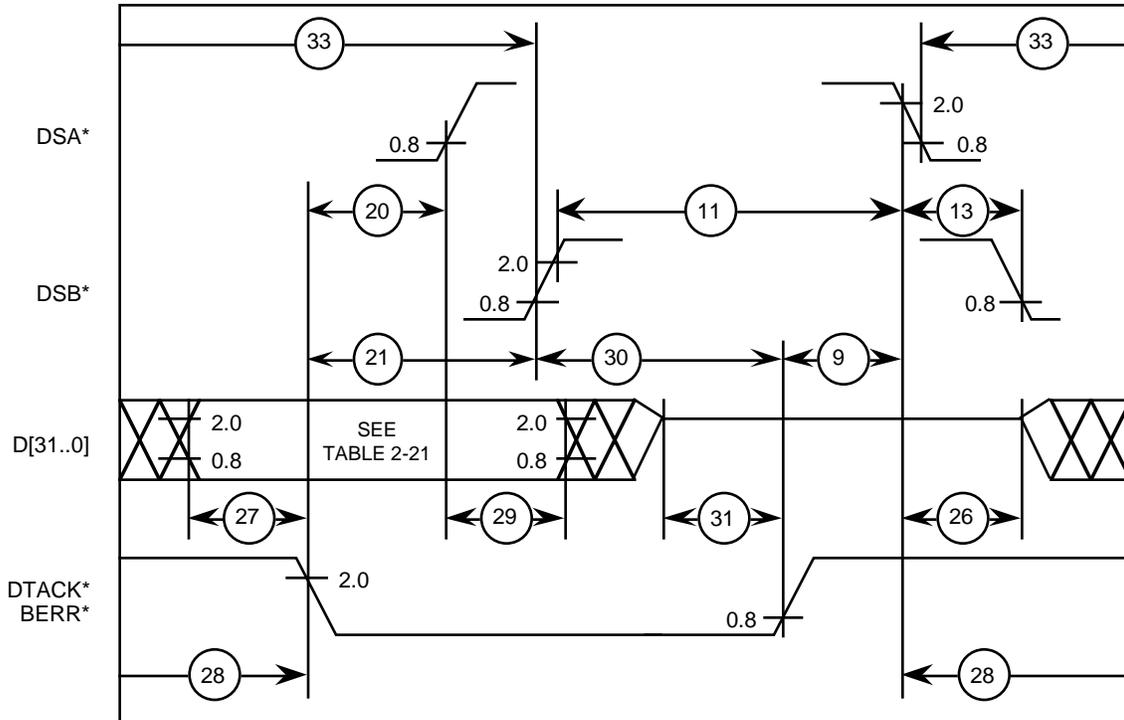


Figure 2 - 16 Master, Slave, And Location Monitor Data Transfer Timing (Continued)
 Single Byte Read
 Byte(0-2) UAT Read
 Byte(1-3) UAT Read
 Single-Byte Block Read



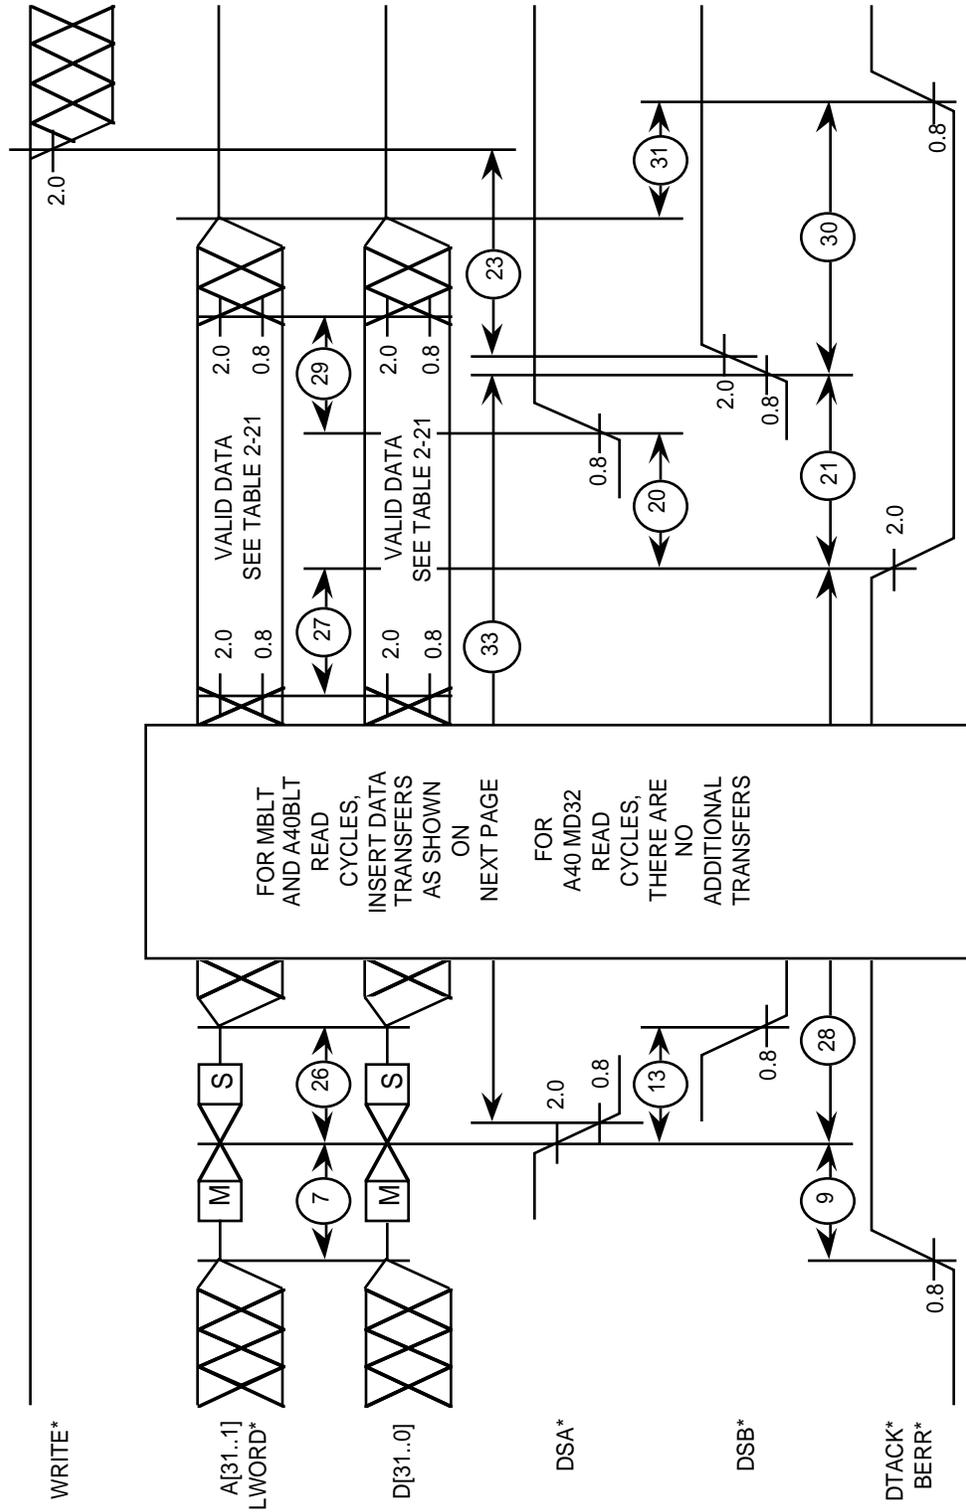
Note:
 The falling and rising edges of DSA* and DSB* are not connected to indicate that the signal that is the first to fall is not necessarily the one that is the first to rise.

Figure 2 - 17 Master, Slave, And Location Monitor Data Transfer Timing
 Byte(0-1) Read, Byte(2-3) Read
 Byte(0-3) Read, Byte(1-2) Read
 Double-Byte Block Read
 Quad-Byte Block Read



Note:
 The rising and falling edges of DSA* and DSB* are not connected to indicate that the signal that is the first to rise is not necessarily the one that is the first to fall

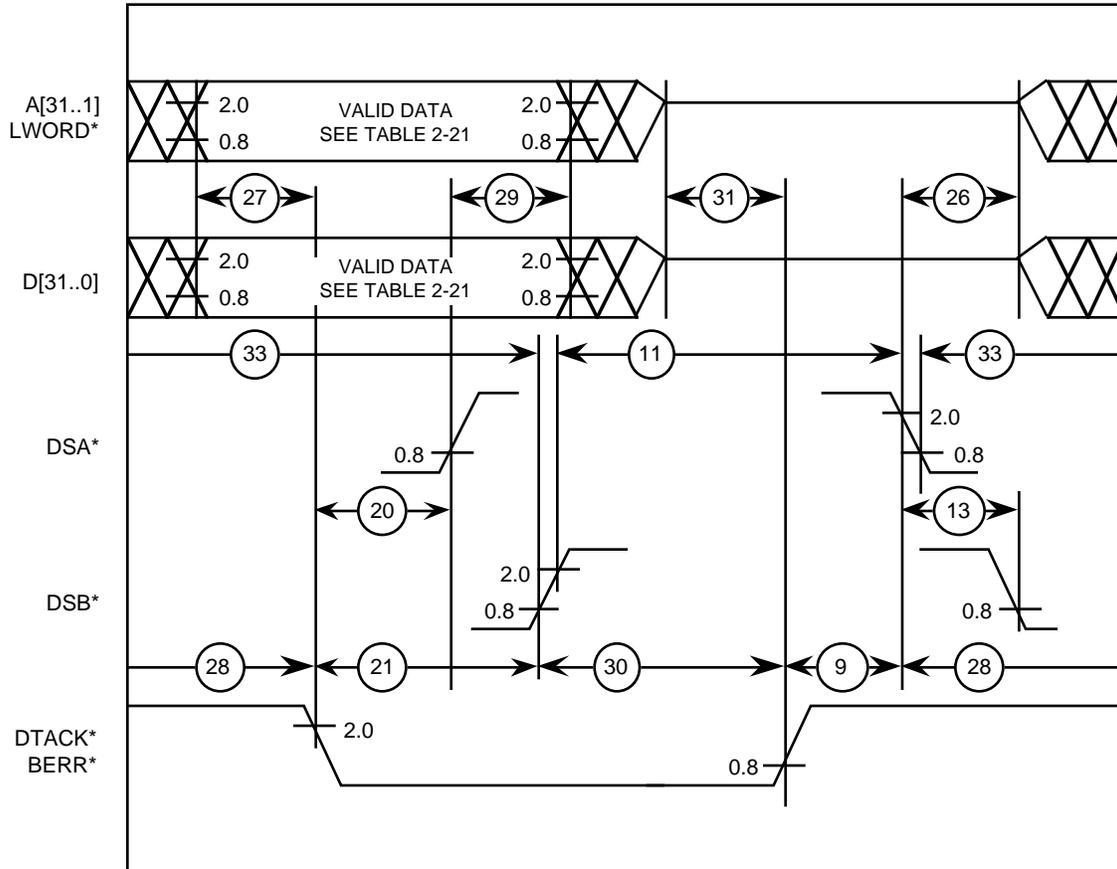
Figure 2 - 17 Master, Slave, And Location Monitor Data Transfer Timing (Continued)
 Byte(0-1) Read , Byte(2-3) Read
 Byte(0-3) Read, Byte(1-2) Read
 Double-Byte Block Read
 Quad-Byte Block Read



Notes:

1. The rising and falling edges of **DSA*** and **DSB*** are shown disconnected to indicate that the signal that is the first to fall is not necessarily the one that is the first to rise.
2. See Figure 2-27 for address broadcast timing.

Figure 2 - 28 Master, Slave And Location Monitor Data Transfer Timing
 A40 Multiplexed Quad Byte Read, A40BLT Multiplexed Quad Byte Block Read ,MBLT
 Eight Byte Block Read



Note:

The rising and falling edges of DSA* and DSB* are not connected to indicate that the signal that is the first to rise is not necessarily the one that is the first to fall

Figure 2 - 28 Master, Slave And Location Monitor Data Transfer Timing
 A40 Multiplexed Quad Byte Read, A40BLT Multiplexed Quad Byte Block Read ,MBLT
 Eight Byte Block Read (Continued)

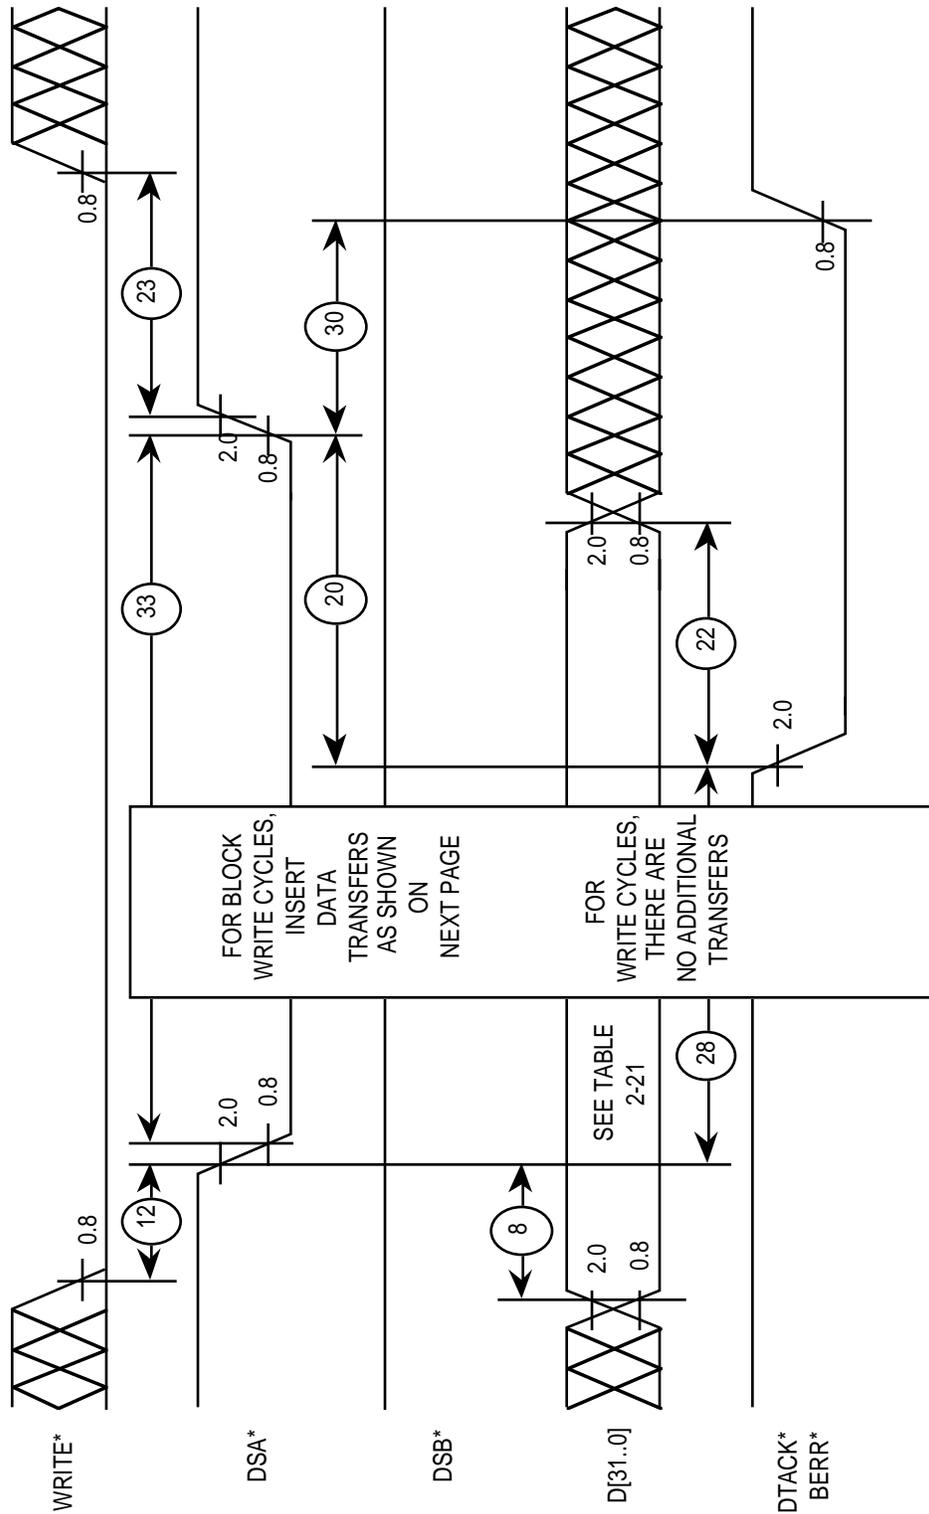


Figure 2 - 18 Master, Slave, And Location Monitor Data Transfer Timing
 Byte(0) Write, Byte(1) Write
 Byte(2) Write , Byte(3) Write
 Byte(0-2) Write, Byte(1-3) Write
 Single-Byte Block Write

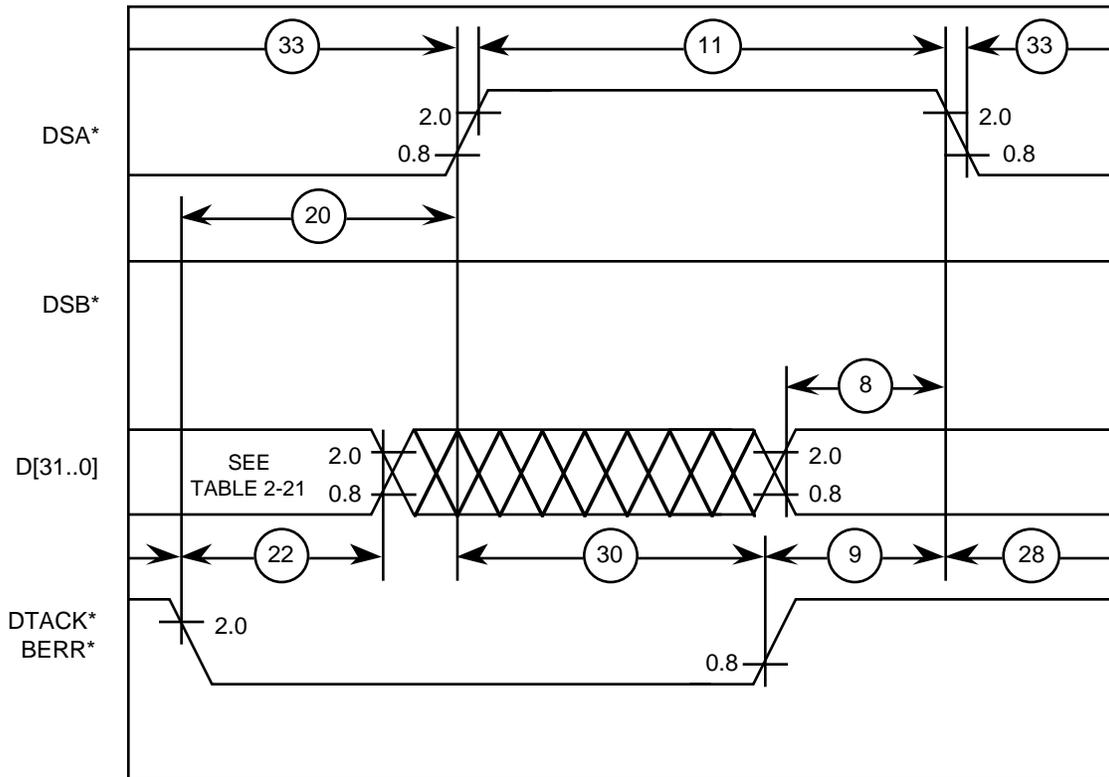
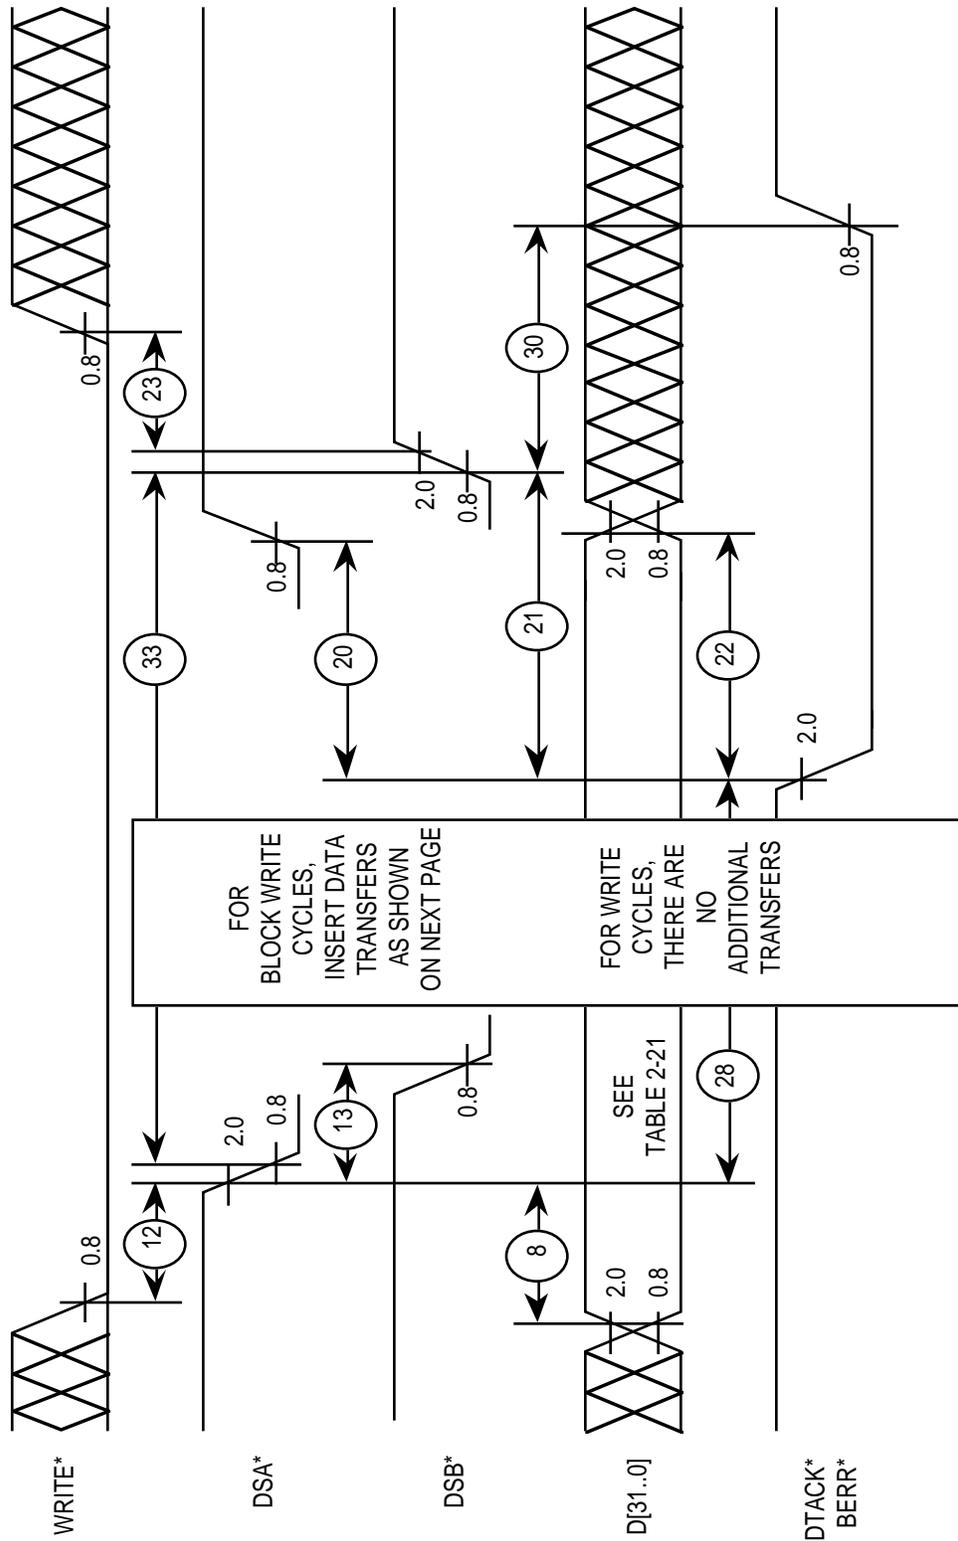


Figure 2 - 18 Master, Slave, And Location Monitor Data Transfer Timing (Continued)
 Byte(0) Write, Byte(1) Write
 Byte(2) Write, Byte(3) Write
 Byte(0-2) Write, Byte(1-3) Write
 Single-Byte Block Write



Note:
 The falling and rising edges of DSA* and DSB* are not connected to indicate that the signal that is the first to fall is not necessarily the one that is the first to rise.

Figure 2 - 19 Master, Slave, And Location Monitor Data Transfer Timing
 Byte(0-1) Write , Byte(2-3) Write
 Byte(0-3) Write, Byte(1-2) Write
 DOUBLE-Byte Block Write
 Quad-Byte Block Write

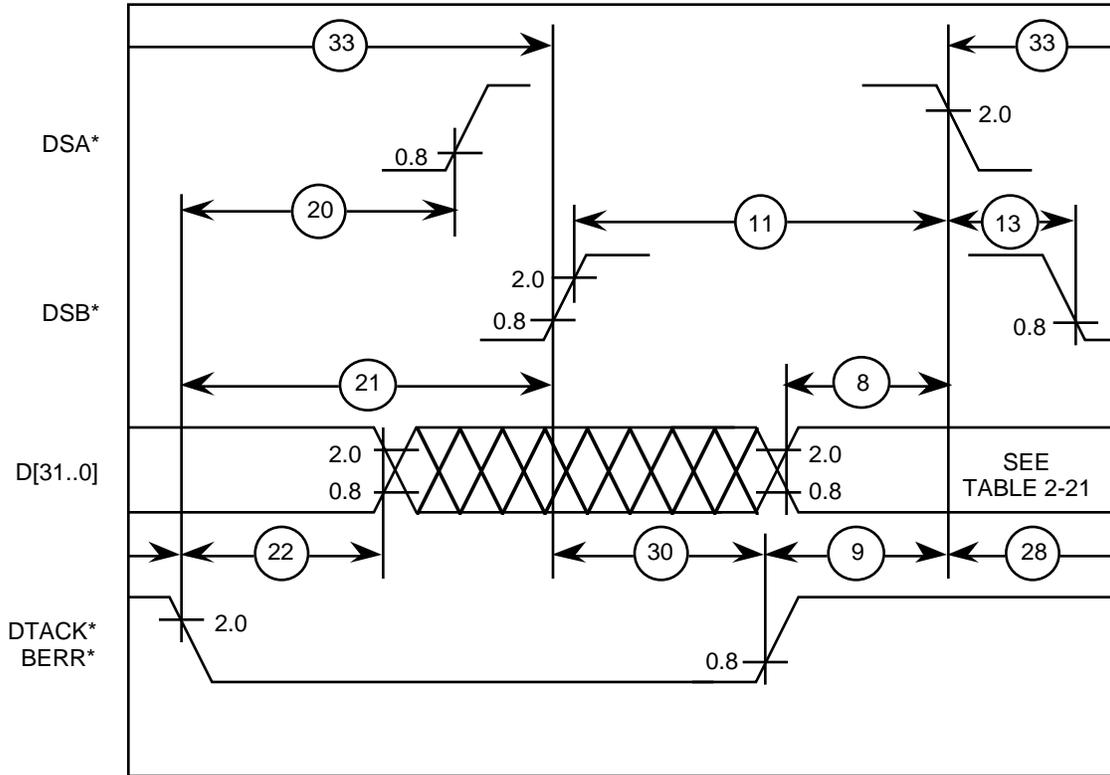
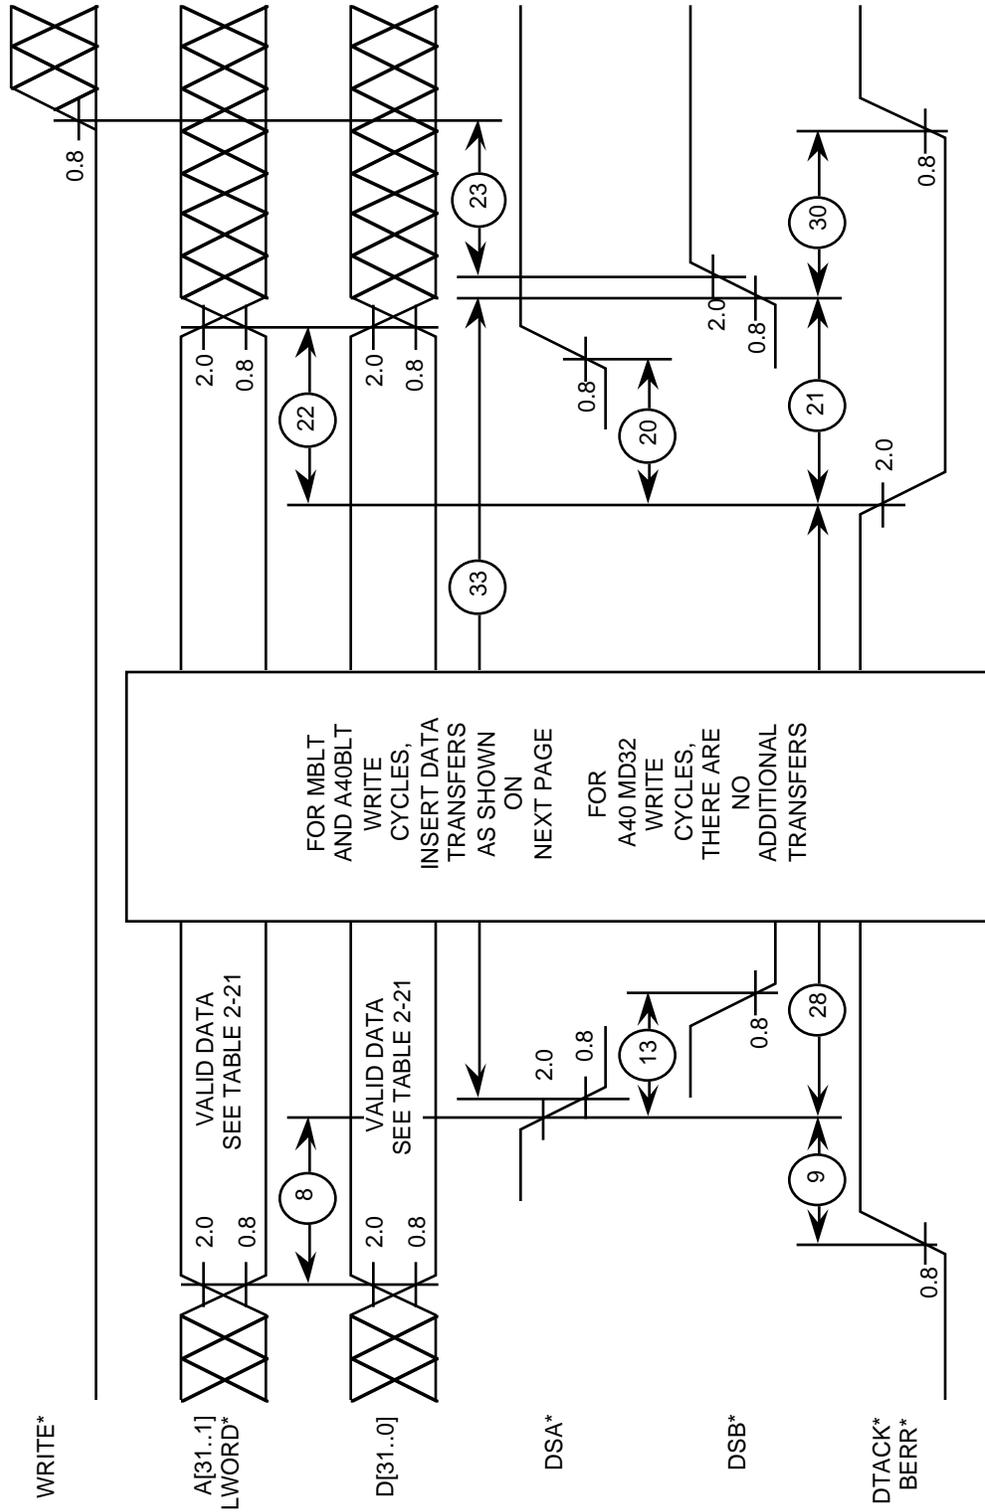


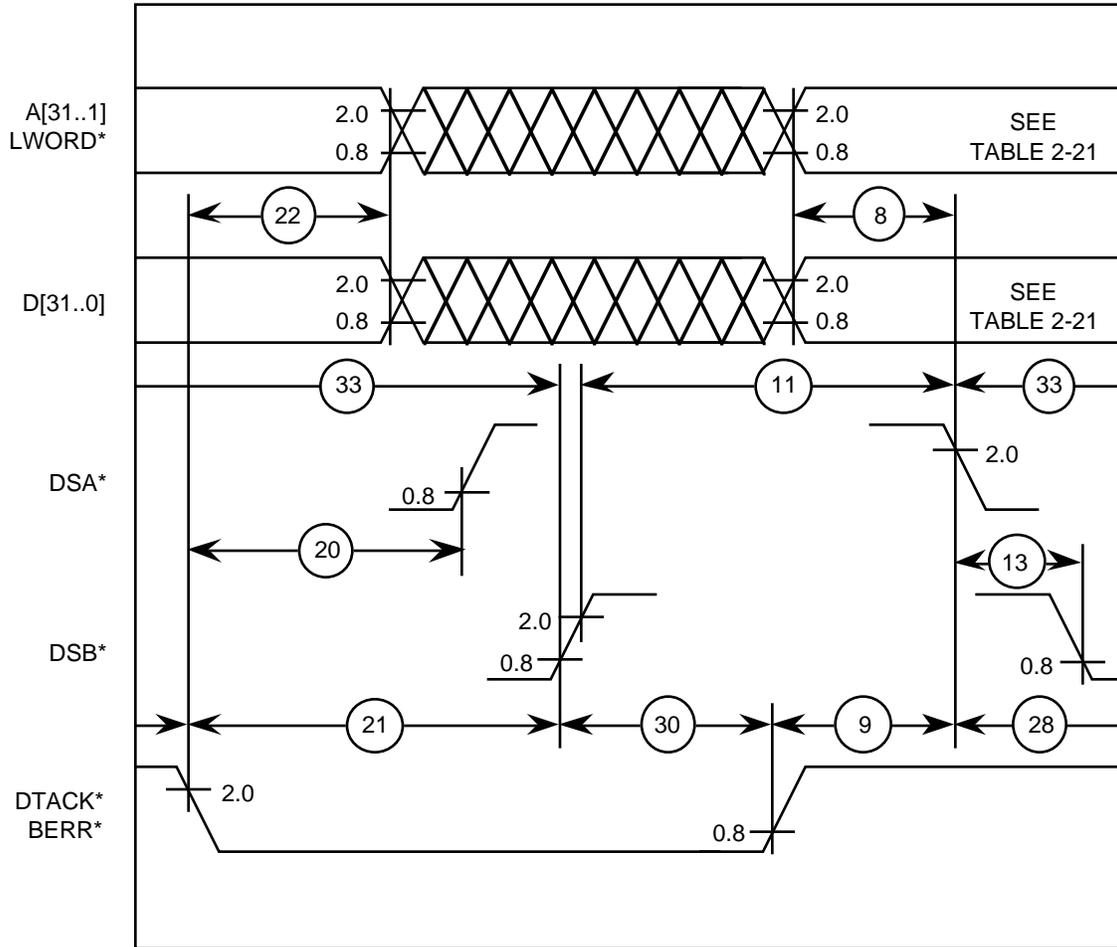
Figure 2 - 19 Master, Slave, And Location Monitor Data Transfer Timing (Continued)
 Byte(0-1) Write, Byte(2-3) Write
 Byte(0-3) Write, Byte(1-2) Write
 DOUBLE-Byte Block Write
 Quad-Byte Block Write



Note:

1. The rising and falling edges of DSA* and DSB* are shown disconnected to indicate that the signal that is the first to fall is not necessarily the one that is the first to rise.
2. See Figure 2-27 for address broadcast timing.

Figure 2 - 29 Master, Slave And Location Monitor Data Transfer Timing
 A40 Multiplexed Quad Byte Write, A40BLT Multiplexed Quad Byte Block Write ,MBLT
 Eight Byte Block Write



Note:

The rising and falling edges of DSA* and DSB* are not connected to indicate that the signal that is the first to rise is not necessarily the one that is the first to fall

Figure 2 - 29 Master, Slave And Location Monitor Data Transfer Timing
 A40 Multiplexed Quad Byte Write, A40BLT Multiplexed Quad Byte Block Write ,MBLT
 Eight Byte Block Write (Continued)

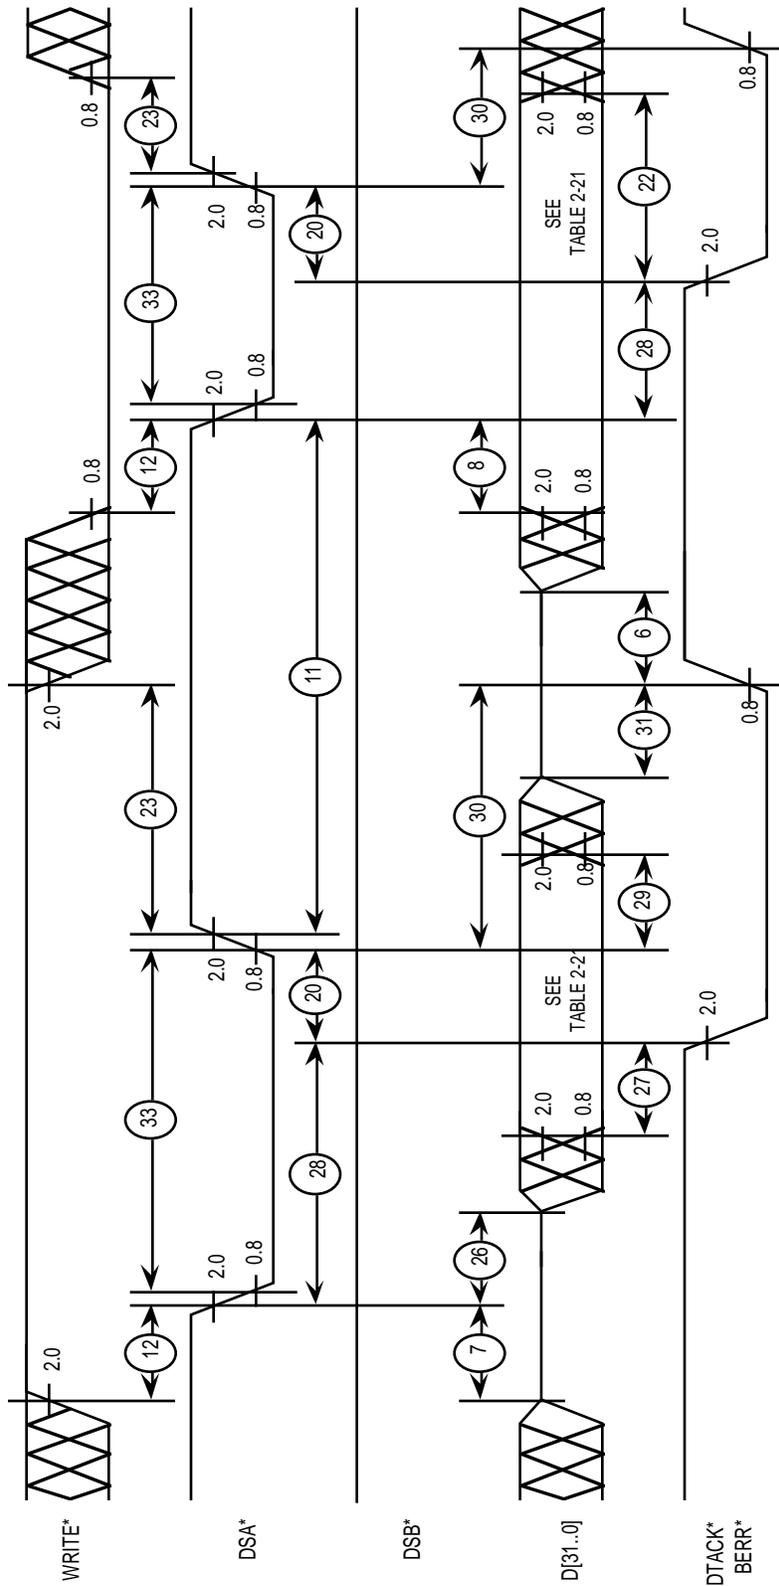
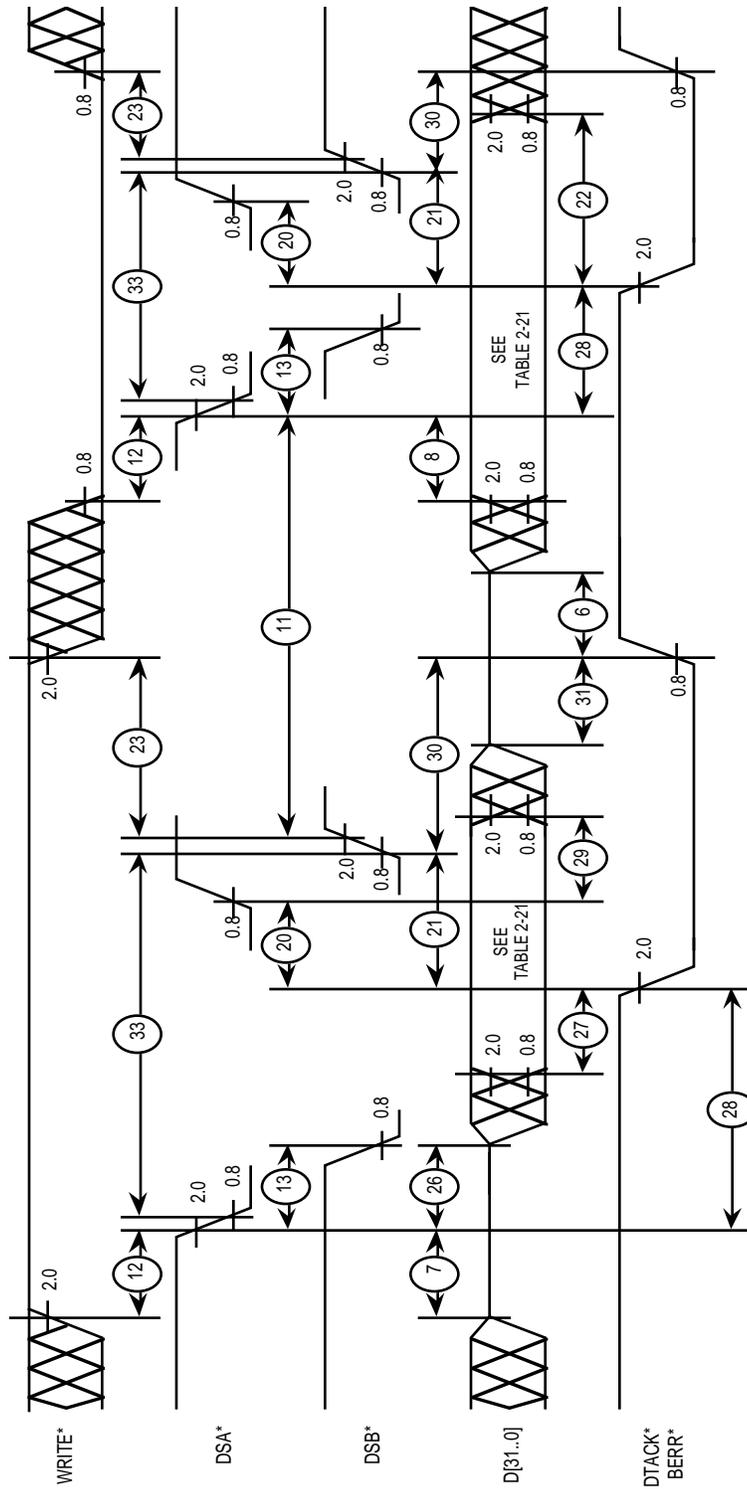


Figure 2 - 20 Master, Slave, And Location Monitor Data Transfer Timing Single-Byte RMW Cycles



Note:
The rising and falling edges of DSA* and DSB* are not connected to indicate that the signal that is the first to rise is not necessarily the one that is the first to fall

Figure 2 - 21 Master, Slave, and Location Monitor Data Transfer Timing Double-Byte RMW Cycles, Quad-Byte RMW Cycles

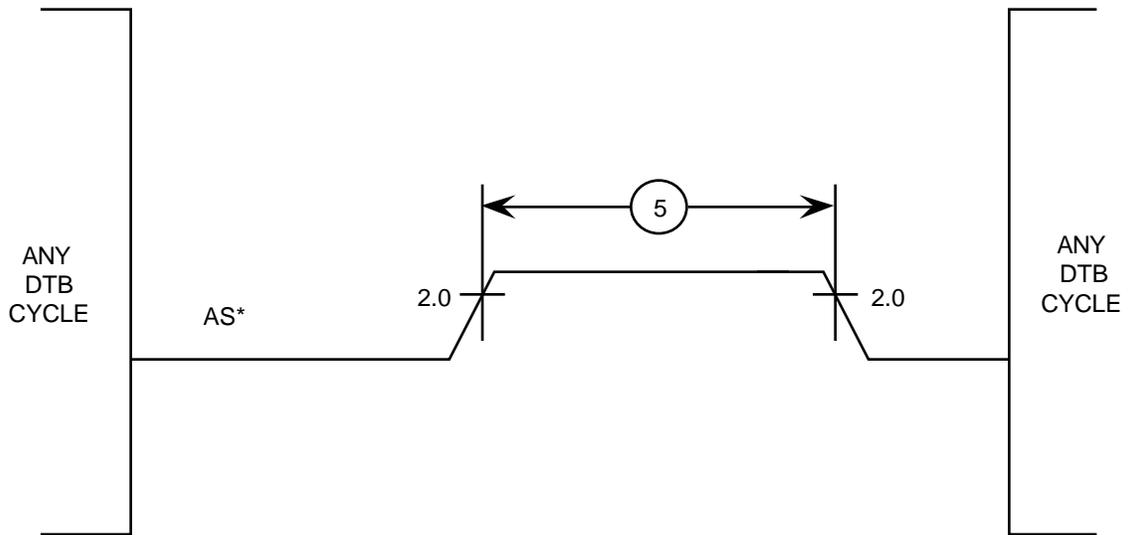
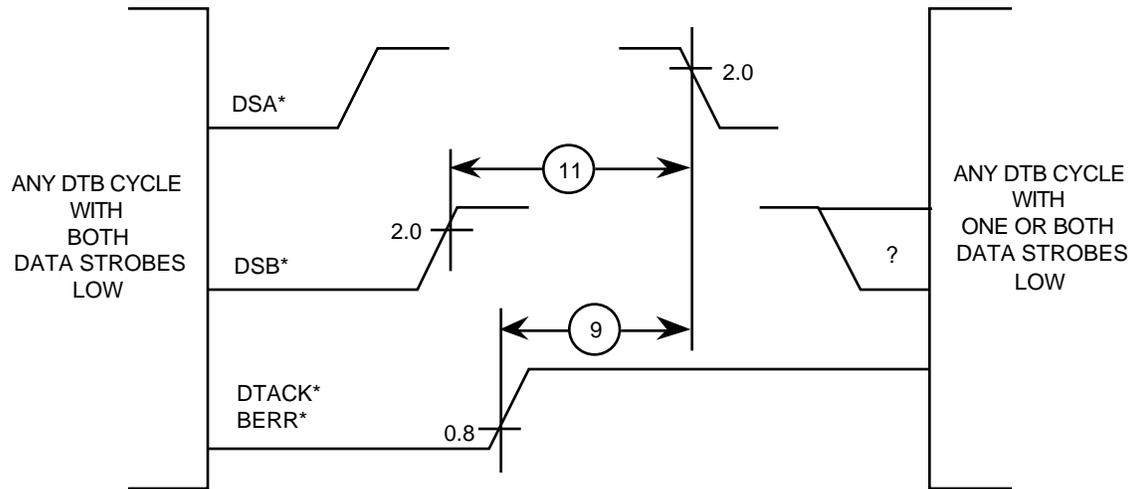
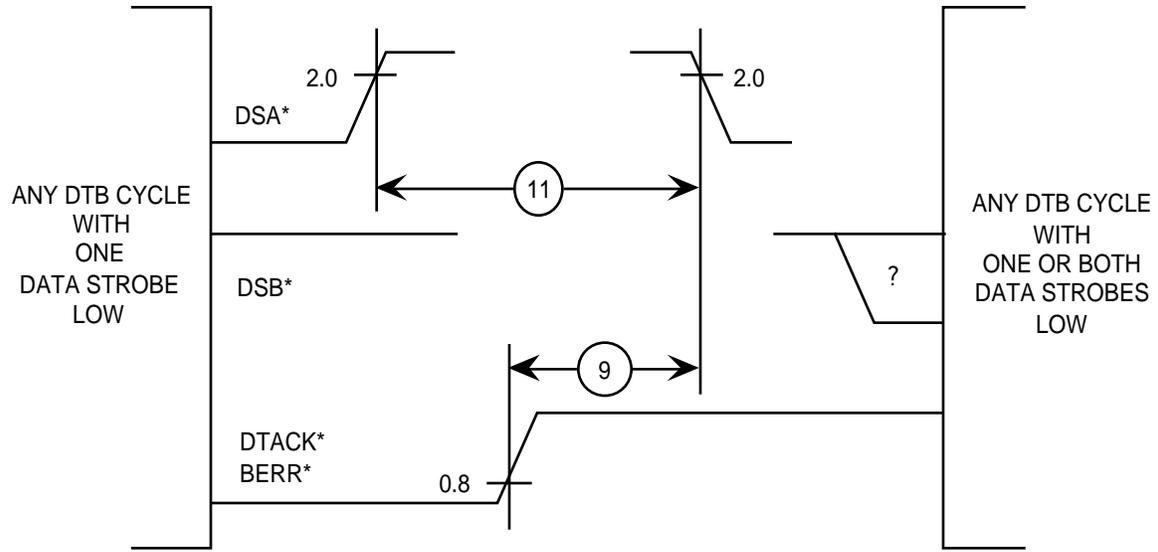


Figure 2 - 22 Address Strobe Inter-Cycle Timing



An example of a cycle where both data strobes go low followed by a cycle where one or both data strobes go low

Figure 2 - 23 Data Strobe Inter-Cycle Timing



An example of a cycle where one data strobe goes low followed by a cycle where one or both data strobes go low

Figure 2 - 24 Data Strobe Inter-Cycle Timing

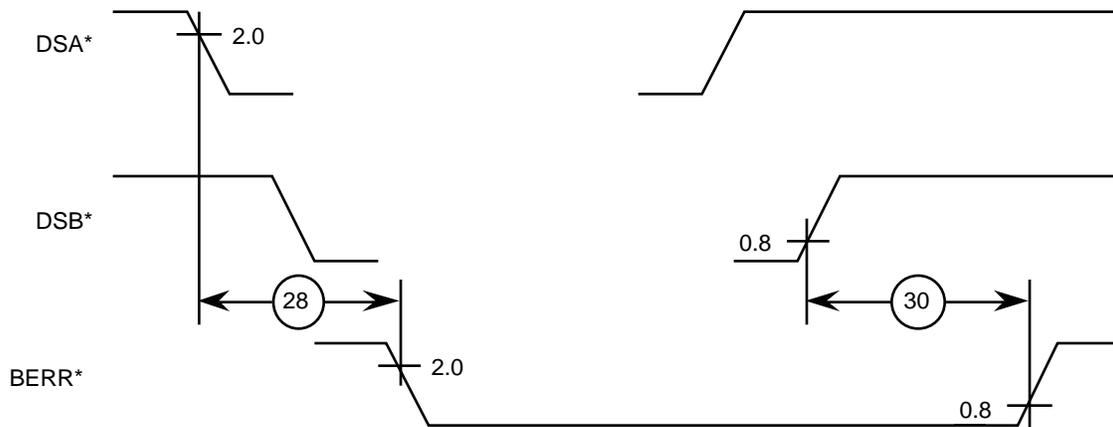


Figure 2 - 25 Master, Slave, And Bus Timer Data Transfer Timing Timed-Out Cycle

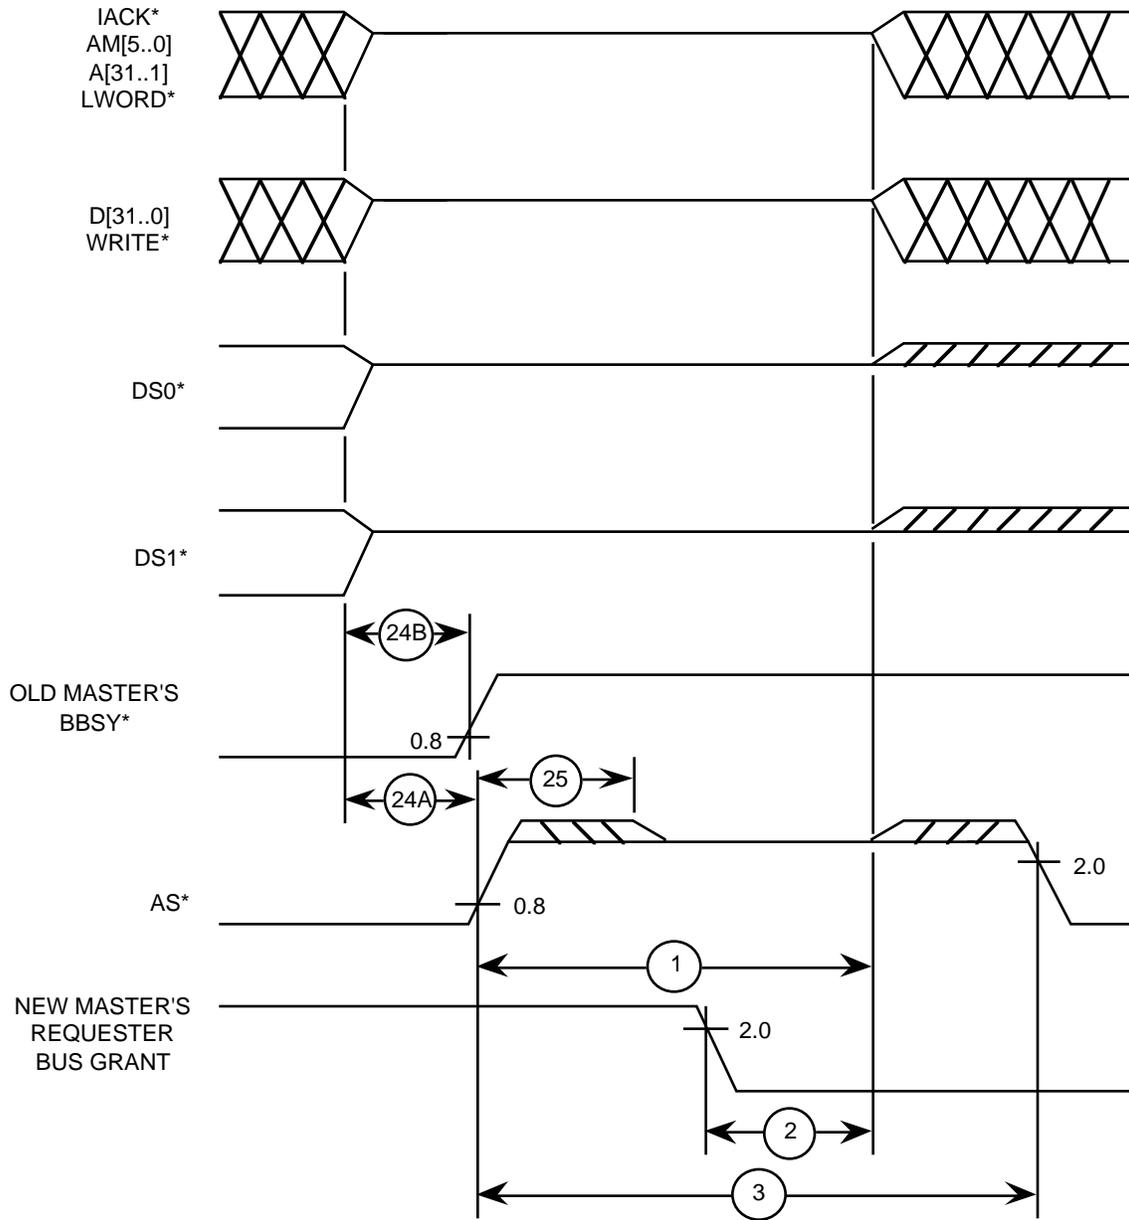
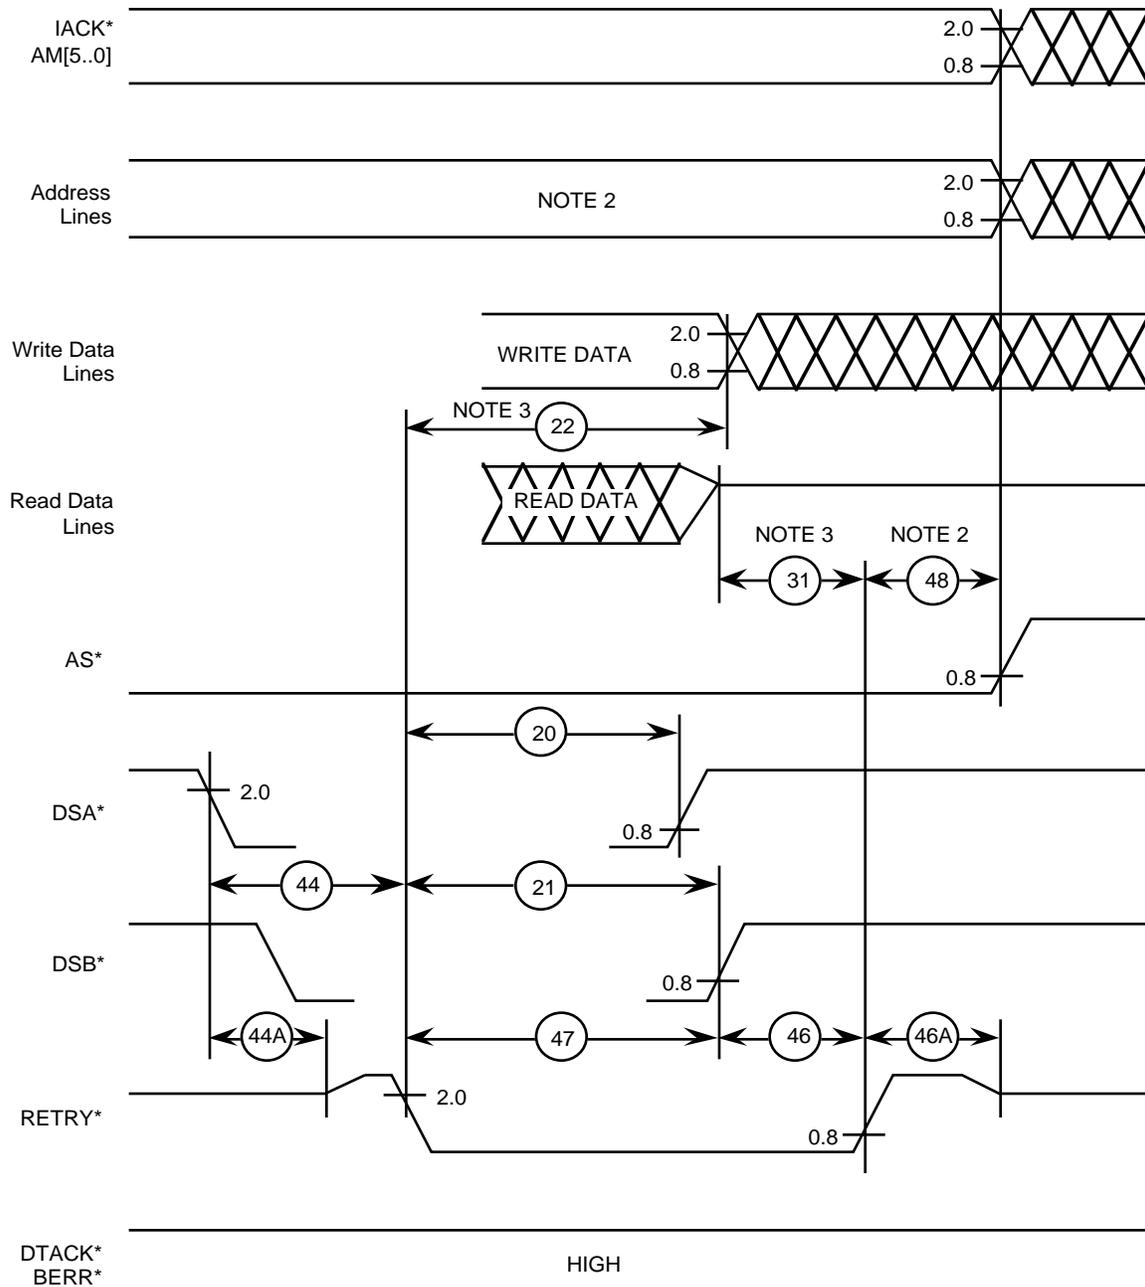


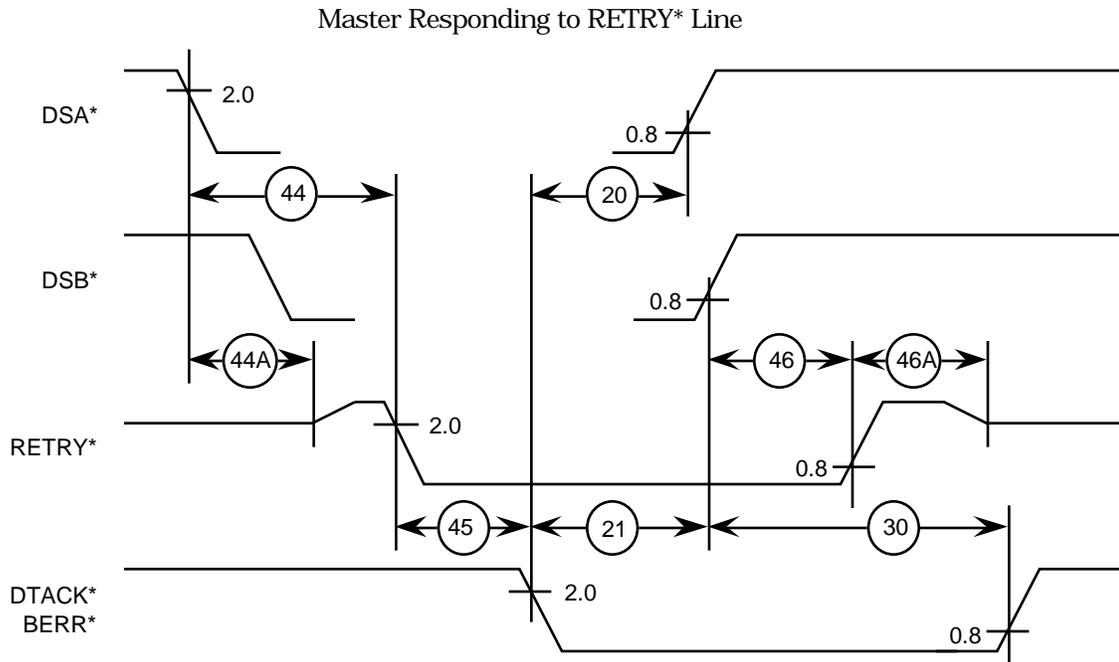
Figure 2 - 26 Master DTB Control Transfer Timing



Notes:

1. The rising and falling edges of DSA* and DSB* are not connected to indicate that the signal that is the first to rise is not necessarily the one that is the first to fall
2. Timing Parameter 48 does not apply to the address lines during the data phase of multiplexed-address or multiplexed-data cycles. Nor does it imply any timing relationships between AS* and address information.
3. Apply Timing Parameter 22 for write cycles and Parameter 31 for read cycles.

Figure 2 - 30 Master And Slave Data Transfer Timing



Note:

The rising and falling edges of DSA* and DSB* are not connected to indicate that the signal that is the first to rise is not necessarily the one that is the first to fall

Figure 2 - 31 Master And Slave Data Transfer Timing
Master Ignoring RETRY* Line

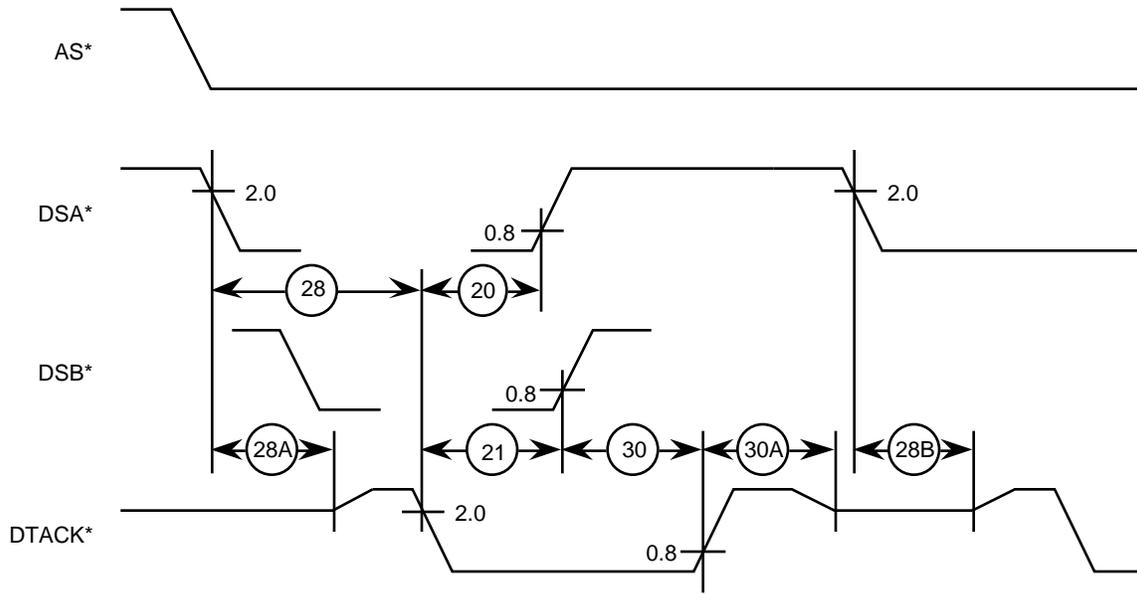


Figure 2 - 33 Rescinding DTACK Timing

Chapter 3

Data Transfer Bus Arbitration

3.1 Bus Arbitration Philosophy

As microprocessor costs decrease, it is becoming more cost effective to design systems with multiple processors sharing global resources.

The most fundamental of these global resources is the Data Transfer Bus through which all other global resources are accessed. Therefore, any system that supports multiprocessing needs to provide an efficient allocation method for the Data Transfer Bus. Because speed of allocation is vital, a hardware allocation scheme is the only practical alternative. The VMEbus meets this need with its Arbitration subsystem. (See Figure 3-1).

The VMEbus arbitration subsystem:

- a. Prevents simultaneous use of the bus by two Masters.
- b. Schedules requests from multiple Masters for optimum bus use.

3.1.1 Types Of Arbitration

When several boards request use of the DTB simultaneously, the arbitration subsystem detects these requests and grants the bus to one board at a time. The decision of which board is granted the bus first depends upon what scheduling algorithm is used. Many algorithms are possible. This specification describes four arbitration algorithms: prioritized, round-robin, single level and fair arbitration.

Prioritized arbitration assigns the bus according to a fixed priority scheme where each of four bus request lines has a priority from highest (BR3*) to lowest (BR0*).

Round-robin arbitration assigns the bus on a rotating priority basis. When the bus is granted to the Requester on bus request line BR(n)*, then the highest priority for the next arbitration is assigned to bus request line BR(n-1)*.

Single level arbitration only accepts requests on BR3*, and relies on BR3*'s bus grant daisy-chain to arbitrate the requests.

Fair arbitration ensures that all requesters on the same request level get equal access to the DTB. In the fair arbitration algorithm a requester will not issue a request for the DTB if a request has been issued by another requester on its own level.

PERMISSION 3.1:

Scheduling algorithms other than priority, round-robin, or single level MAY be used. For example, an Arbiter's algorithm might give highest priority to BR3*, but grant the bus to BR0* through BR2* on a round-robin basis.

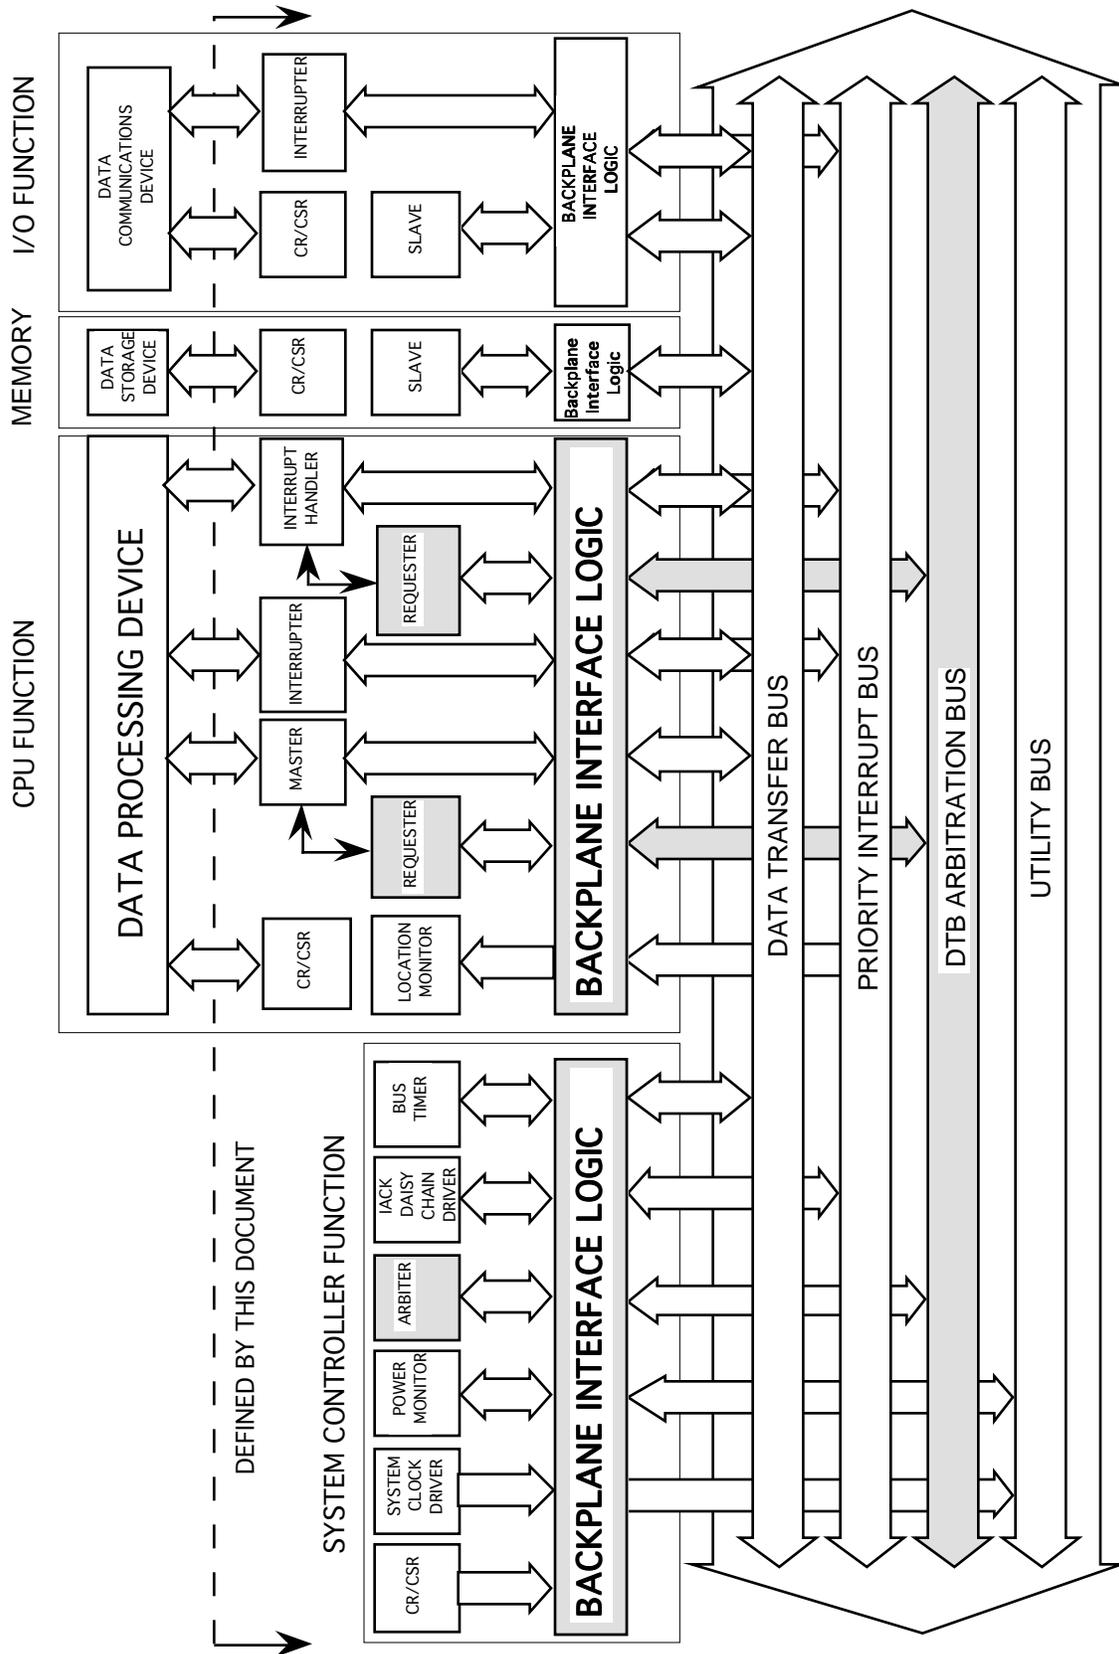


Figure 3 - 1 Arbitration Functional Block Diagram

3.2 Arbitration Bus Lines

The Arbitration Bus consists of six bused VMEbus lines and four daisy-chained lines. These daisy-chained lines require special signal names. The signals entering each board are called "Bus Grant IN" lines (BG[3..0]IN*), while the signals leaving each board are called "Bus Grant OUT" lines (BG[3..0]OUT*). The lines which leave slot n as BG[3..0]OUT* enter slot n+1 as BG[3..0]IN*. This is illustrated in Figure 3-2.

OBSERVATION 3.1:

In all descriptions in this chapter, the terms BRx*, BG[3..0]IN*, and BG[3..0]OUT* are used to describe the bus request and bus grant lines, where x takes on any value from zero to three.

In the VMEbus arbitration system, a Requester module drives the following lines:

- 1 bus request line (one of BR0* through BR3*)
- 1 bus grant out line (one of BG0OUT* through BG3OUT*)
- 1 bus busy line (BBSY*)

RULE 3.1

- IF a VMEbus board does not generate bus requests on some bus request levels,
- THEN it MUST propagate the daisy-chain signals for those levels from its BG[3..0]IN* lines to its respective BG[3..0]OUT* lines.

PERMISSION 3.2:

The propagation for the unused lines of the bus grant daisy-chain MAY be done using jumpers or active logic. The latter approach allows selection of the request level under software control, while the former results in faster propagation through the daisy-chain.

Three types of Arbiters are described in this standard. They are:

- PRIoritized (PRI)
- Round-Robin-Select (RRS)
- SinGLe level (SGL)

The operation of these three types of Arbiters is described in Section 3.3.

A PRI Arbiter drives the following:

- 1 bus clear line (BCLR*)
- 4 bus grant lines (Slot 1 BG0IN* through BG3IN*) if the Arbiter's board also has a Requester.
(Slot 1 BG0OUT* through BG3OUT*) if the Arbiter's board does not have a Requester.

An RRS Arbiter drives the four Slot 1 BG[3..0]IN* or BG[3..0]OUT* lines and, optionally, the BCLR* line.

A SGL Arbiter drives only BG3IN* or BG3OUT* at Slot 1, and, optionally the BCLR* line.

Two additional lines are connected (although indirectly) with the arbitration system during power-up and power-down sequencing: SYSRESET* and ACFAIL*. While their impact on the arbitration system is included in this chapter, these lines will be discussed further in Chapter 5.

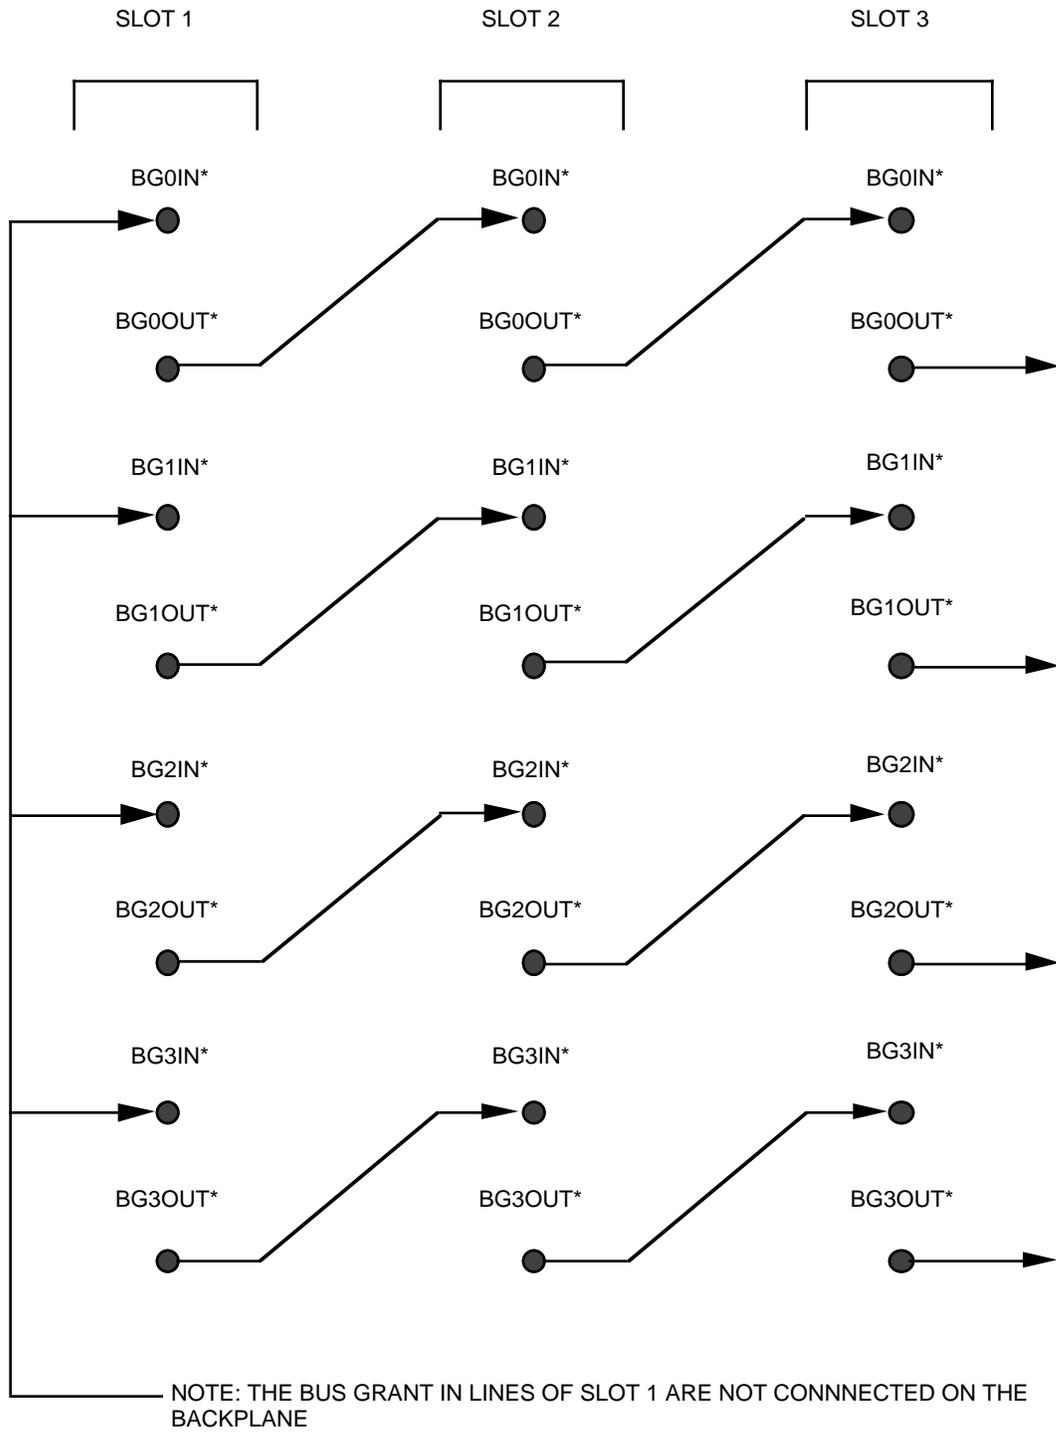


Figure 3 - 2 Illustration Of The Daisy Chain Bus Grant Lines

3.2.1 Bus Request And Bus Grant Lines

The bus request lines are used by each Requester to request use of the DTB. The bus grant lines allow the Arbiter to award use of the bus. It does this by driving a bus grant daisy-chain line low. This low level propagates down the daisy-chain, typically passing through several boards in the process. If a board never uses a particular request/grant level, the signal is passed through that board. When the board uses a request/grant level x , the corresponding signal $BG[3..0]IN^*$ is gated on board. If its on-board Requester is currently requesting the DTB on that level, it does not pass the low level on to its $BG[3..0]OUT^*$. Otherwise, it passes on the low level to the next slot in the chain.

RULE 3.2:

IF a VMEbus backplane slot is not occupied by a board, and if there are boards farther down the daisy-chain.

THEN jumpers MUST be installed at the empty slot to pass through all daisy-chain signals.

OBSERVATION 3.2a:

The backplane mechanical specification in Chapter 7 describes a provision for the installation of jumpers at each slot. It is possible to provide automatic mechanical or electronic jumpering on VMEbus backplanes.

RULE 3.3:

The Arbiter MUST always be located in Slot 1.

3.2.2 Bus Busy Line (BBSY*)

Once a Requester has been granted control of the Data Transfer Bus via the bus grant daisy-chain, it drives BBSY* low. It then has control of the DTB until it releases BBSY*, allowing the Arbiter to grant the DTB to some other Requester.

3.2.3 Bus Clear Line (BCLR*)

The PRI Arbiter drives BCLR* low to inform the Master, currently in control of the DTB, when a higher priority request is pending. The current Master does not have to relinquish the bus within any prescribed time limit. It can continue transferring data until it reaches an appropriate stopping point, and then allow its on-board Requester to release BBSY*.

PERMISSION 3.3:

Although RRS Arbiters are not required to drive the BCLR* line they MAY do so.

SUGGESTION 3.1:

IF an RRS Arbiter drives the BCLR* line low.

THEN design it to do so whenever there is a request pending on any of the non-granted bus request lines

3.3 Functional Modules

The arbitration subsystem is composed of several modules:

- a. One Arbiter
- b. One or more Requesters

Figures 3-3 and 3-4 provide block diagrams for the two types of Arbitration Bus modules.

RULE 3.4:

Output signal lines shown with solid lines in Figure 3-3 and 3-4 **MUST** be driven by the module, unless it would always drive them high.

RULE 3.5:

Input signal lines shown with solid lines in Figures 3-3 and 3-4 **MUST** be monitored and responded to in the appropriate fashion.

OBSERVATION 3.3:

RULEs and PERMISSIONs for driving and monitoring the signal lines shown with dotted lines in Figures 3-3 and 3-4 are given in Tables 3-1 and 3-2.

OBSERVATION 3.4:

IF a bused output signal line is not driven
 THEN terminators on the backplane ensure that it is high.

OBSERVATION 3.5:

Although SYSRESET* and ACFAIL* are not specified as part of the Arbitration Bus, they are important here because Masters, which are paired with Requesters, respond to these signal lines. (SYSRESET* and ACFAIL* are driven by the Power Monitor module which is discussed in Chapter 5.)

3.3.1 Arbiter

The Arbiter is a functional module that decides which Requester should be granted control of the DTB when several Requesters request it simultaneously. There are many possible algorithms that could be used to make this decision. Three types of Arbiters are described in this specification: a prioritized (PRI) Arbiter, a round-robin (RRS) Arbiter, and a single level (SGL) Arbiter.

An Arbiter responds to incoming bus requests and grants the DTB to the appropriate Requester with one of the bus grant lines.

When the Arbiter detects BBSY* high for at least 40 nS, and after it detects one or more bus requests, it issues a bus grant, corresponding to the highest priority bus request.

When the Requester receives the bus grant, it drives BBSY* low and signals to its on-board Master or Interrupt Handler that it has been granted the DTB. After its on-board Master or Interrupt Handler finishes using the DTB, the Requester releases BBSY*. The resulting rising edge of BBSY* signals the Arbiter to issue another bus grant, based upon the levels of the bus request lines at that time.

In addition to the arbitration provided by the Arbiter, a secondary level of arbitration is provided by the bus grant daisy-chains. Because of these daisy-chains, Requesters sharing a common request line are prioritized by slot position. The Requester closest to Slot 1 has the highest priority.

SGL Arbiters respond only to bus requests on BR3* and depend on the BG3IN*/BG3OUT* daisy-chain to do the priority arbitration.

The PRI Arbiter prioritizes the four bus request lines, from BR0* (the lowest) to BR3* (the highest), and responds with BG0IN* through BG3IN*, as appropriate. A PRI Arbiter also informs any Master, currently in control of the bus, when a higher level request is pending by driving BCLR* low.

To visualize an RRS Arbiter, consider a mechanical switch being driven by a stepping motor. Each position on the switch connects a bus request line to its corresponding bus grant line. When the bus is busy, the switch is stopped on the current level. Upon release of the bus, the switch steps one position lower (i.e., from BR(n)* to BR(n-1)*) and tests for a request. It continues this scanning operation until a request is found, sending a bus grant over the appropriate line.

PERMISSION 3.4:

An Arbiter MAY be designed with a built-in time-out feature that causes it to withdraw a bus grant if BBSY* is not driven low by a Requester within a prescribed time.

OBSERVATION 3.6:

The time used by the Arbiter allowed by PERMISSION 3.4 needs to be longer than the longest possible bus grant daisy-chain propagation delay time, plus the time the slowest Requester takes to generate BBSY*.

RULE 3.6:

Except for a time-out situation where no Requester responds, once the Arbiter grants the bus to a Requester it MUST NOT generate a new bus grant before that Requester generates a rising edge on the BBSY* line. (The Requester generates a rising edge by driving BBSY* low and then releasing it.)

OBSERVATION 3.7:

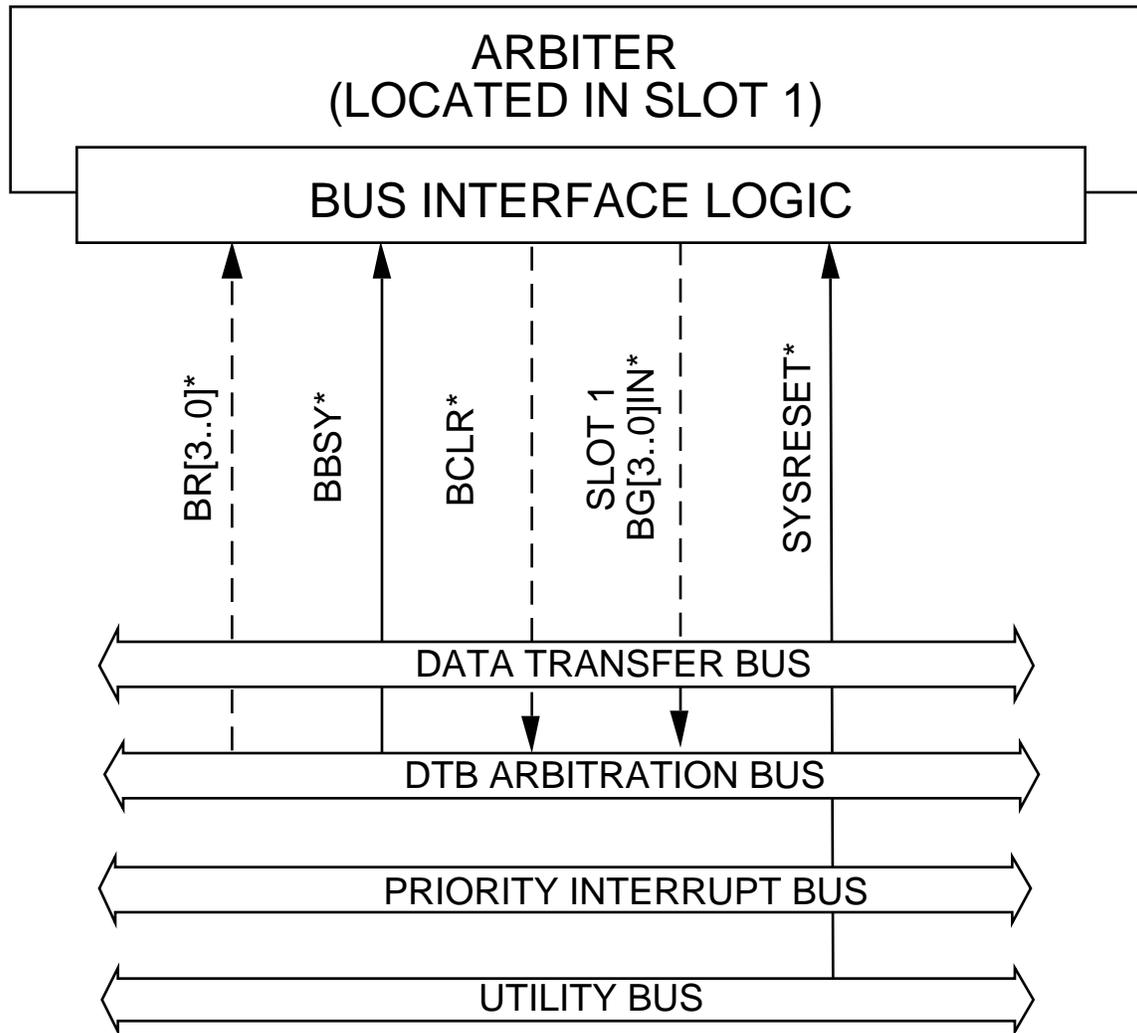
IF an Arbiter uses a "snapshot" of the request lines taken prior to the rising edge of BBSY*.
 THEN it might grant the bus to a Requester that has since removed its request.

RULE 3.15:

The Arbiter MUST wait for a minimum of 40nS after detecting the rising edge of BBSY* before driving any BG[0..3]IN* lines low.

OBSERVATION 3.19:

Rule 3.15 guarantees a 40nS minimum high time on BBSY* (30nS on the non-System Controller boards) to support future protocols that need to detect the negation of BBSY*.



Note : The RULES and PERMISSIONS for monitoring and driving the dotted lines are given in Table 3-1.

Figure 3 - 3 Block Diagram: Arbitrator

Table 3 - 1 RULEs and PERMISSIONs That Specify The Use Of The Dotted Lines By The Various Types Of Arbiters Defined By This Document

Type of Arbiter	Use of dotted line
SGL	MUST drive Slot 1 BG3IN*(see note). MUST ensure Slot 1 BG[2..0]IN* are high (see note). MUST monitor BR3*. MAY drive BCLR*, or Slot 1 BG[2..0]IN* (see note). MAY monitor BR[2..0]*.
RRS	MUST drive Slot 1 BG[3..0]IN* (see note). MUST monitor BR[3..0]*. MAY drive BCLR*.
PRI	MUST drive BCLR* and Slot1 BG[3..0]IN* (see note). MUST monitor BR[3..0]*.

Note: If there is no Requester on the board then as per Rule 3.1 the board propagates the daisy chain signals BG[3..0]IN to BG[3..0]OUT.

3.3.2 Requester

Each Requester in the system:

- Monitors the DEVICE WANTS BUS signal from its on-board Master or Interrupt Handler and generates a bus request when the DTB is needed.
- If it detects a low level on its BG[3..0]IN* line, and its on-board Master or Interrupt Handler does not need the DTB, then it passes on that low level to its BG[3..0]OUT*
- If it detects a low level on its BG[3..0]IN* line, and its on-board Master or Interrupt Handler needs the DTB, it generates an on-board DEVICE GRANTED BUS signal to indicate the DTB is available, and drives the BBSY* signal low.

Three types of Requesters are described in this specification: a Release When Done (RWD) Requester, a Release On Request (ROR) Requester and a FAIR Requester.

The RWD Requester releases BBSY* when the Master or Interrupt Handler drives the on-board DEVICE WANTS BUS signal false.

The ROR Requester does not release BBSY* when its on-board DEVICE WANTS BUS signal goes false unless some other Requester on the bus drives one of the bus request lines low. It monitors the four bus request lines and releases BBSY* only if another bus request is pending. ROR Requesters reduce the number of arbitrations initiated by a Master which is generating a large percentage of the bus traffic. See Figure 3-4.

FAIR Requesters are used to ensure guaranteed access to the bus when there are more than four Masters and Interrupt Handlers. Fair Requesters do not request the bus if there are other bus requests pending on the same level. When the current "batch" of requests are serviced, other requesters needing access to the bus issue requests. This approach guarantees each requester on a level will get access to the bus. It does not however guarantee that each requester get equal bandwidth or that the accesses are granted in a consistent order. These are implementation and system dependent.

OBSERVATION 3.17:

When configuring a system to provide fairness for more than four Masters or Interrupt Handlers all the requesters in the system should be Fair Requesters. If one or more of the Requesters are not FAIR Requesters, then fairness can still be guaranteed if the sum total of bus traffic generated by the non-Fair Requesters does not exceed the capacity of the bus.

OBSERVATION 3.18:

The RWD and ROR capabilities describe under what conditions a requester relinquishes control of the DTB. On the other hand, the FAIR capability describes under what conditions the requester requests control of DTB. Therefore, both the RWD and the ROR requesters can include the FAIR capability as well.

Assuming that the Requester's DEVICE WANTS BUS input is true, when it receives a bus grant it does three things:

- a. It drives BBSY* to low.
- b. It releases its low on BR[3..0]*.
- c. It drives the DEVICE GRANTED BUS on-board signal true, allowing the Master or Interrupt Handler to initiate bus transfers.

These events might occur in any order. It is even possible, although meaningless, that the Master or Interrupt Handler might not use the bus in response to this particular grant.

However, the following RULES apply:

RULE 3.7:

On detecting BGxIN going low, the Requester MUST drive BBSY* to low for at least 90 nanoseconds.

RULE 3.8:

On detecting BGxIN going low, the Requester MUST release Bus Request BR[3..0]* to high.

RULE 3.9:

On detecting BGxIN going low, the Requester MUST maintain BBSY* low for at least 30 nanoseconds after it releases its bus request.

OBSERVATION 3.8:

The 30 nanosecond delay between the bus request's rising edge and the BBSY* rising edge ensures that the Arbiter does not mistakenly interpret the old bus request as a new one and issue another grant.

RULE 3.10:

The Requester MUST hold BBSY* low until its bus grant goes high.

OBSERVATION 3.9:

RULE 3.10 ensures that the BBSY* transition to low has been seen by the Arbiter and that all segments of the bus grant daisy-chain have returned to high, in preparation for the next arbitration.

PERMISSION 3.5:

IF a Requester has a bus request pending and, if it sees some other Requester drive BBSY* low

THEN it MAY withdraw its request by releasing the bus request line.

RULE 3.11:

IF a Requester withdraws a bus request without having first been granted the bus,

THEN it MUST wait to do so until BBSY* goes low and it MUST do so within 50 nanoseconds after BBSY* goes low.

SUGGESTION 3.2:

Design Requesters so that they pass on the bus grant daisy-chain as fast as possible after receipt of a bus grant. This will improve system performance.

RULE 3.14:

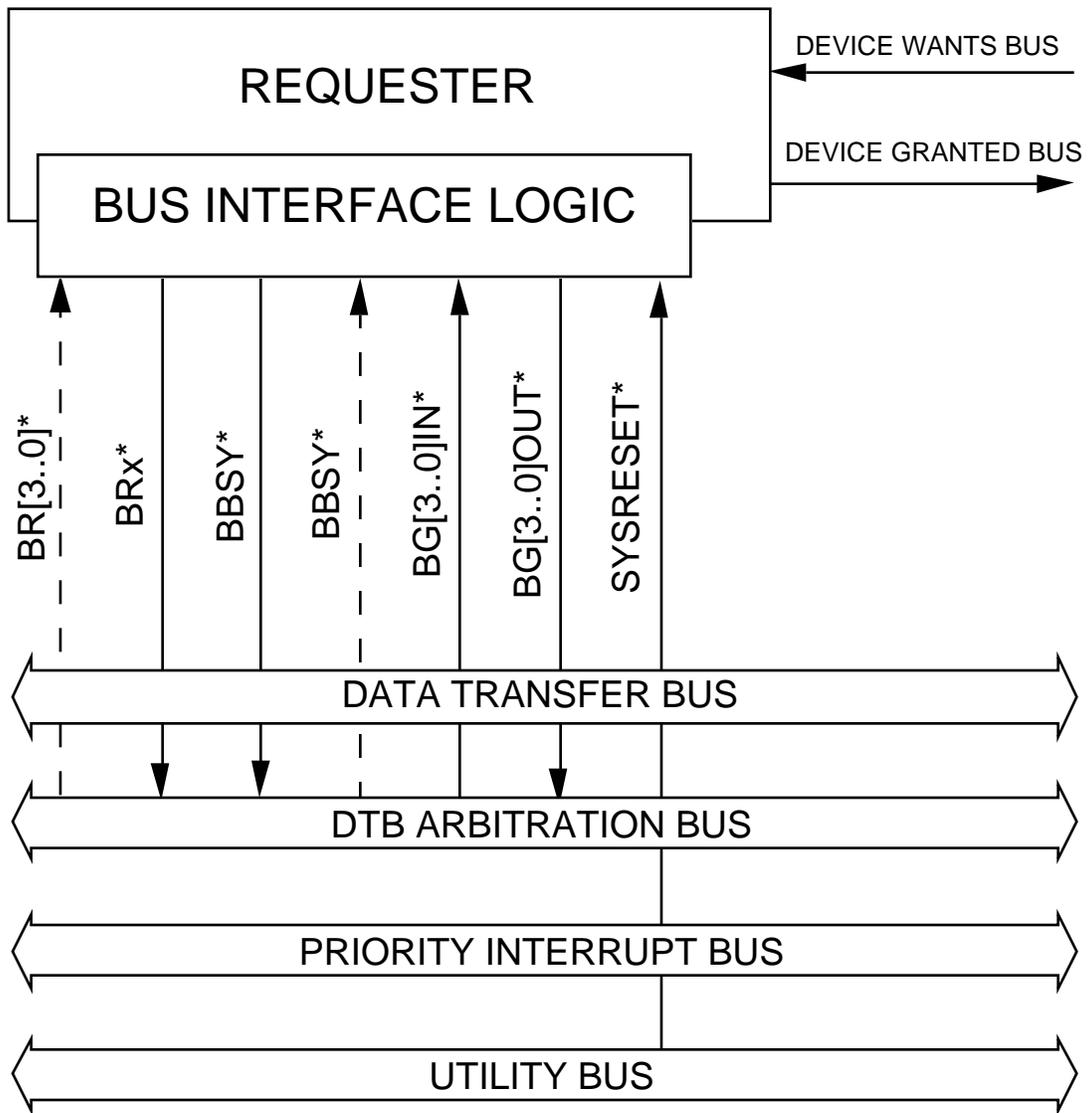
A FAIR Requester MUST NOT issue a bus request unless it detects a high on its bus request line.

RECOMMENDATION 3.3:

A FAIR Requester should wait at least 40nS after it detects a high on its bus request line before asserting its bus request.

OBSERVATION 3.20:

If a FAIR Requester asserts its request line immediately after detecting the line going high then other requesters may not recognize the line being high. This may result in requesters not receiving fair access to the bus.



Note: The RULES and PERMISSIONS for monitoring the dotted lines are given in Table 3-2.

Figure 3 - 4 Block Diagram: Requester

Table 3 - 2 RULES And PERMISSIONS That Specify The Use Of The Dotted Lines
By The Various Types Of Requesters Defined By This Document

Type of Requester	Use of dotted line
RWD	MAY monitor BR[3..0]*. MAY monitor BBSY*
ROR	MUST monitor BR[3..0]*. MAY monitor BBSY*
FAIR	MUST monitor the bus request line it is using MAY monitor BBSY*

3.3.3 Data Transfer Bus Master

3.3.3.1 Release Of The DTB

The bus arbitration protocol determines how and when the DTB is granted to the various Masters and Interrupt Handlers in the system. It does not, however, control when Masters and Interrupt Handlers release the DTB.

Masters and Interrupt Handlers use several criteria in deciding when to release the DTB. Interrupt Handlers give up the bus after their interrupt acknowledge cycle. Masters give up the bus when they finish their data transfers.

Some Masters also monitor the ACFAIL* and BCLR* VMEbus signals. Both of these signals inform the Master that the DTB is needed for some higher priority activity. In the case of BCLR*, the Master's design determines how long it takes to release the bus. For example, a Master on a disk controller board might not be able to relinquish the bus during a disk sector transfer without loss of data, so it might keep the bus until the sector transfer is finished. ACFAIL* informs the Master that an AC power loss has been detected, and the problems the Master will face in surrendering the bus are insignificant compared to the needs of the total system.

RECOMMENDATION 3.1:

Design Masters to release the DTB within 200 microseconds after ACFAIL* goes low, except to participate in the ensuing power failure activities.

OBSERVATION 3.10:

The 200 microsecond time specified in RECOMMENDATION 3.1 is intended to provide enough time for an orderly shut-down of the system.

Whatever criteria are used to decide when to release the DTB, arbitration is done before some other Master or Interrupt Handler begins using it. This arbitration takes place either during the last data transfer or after that transfer, depending on how much notice the Master or Interrupt Handler gives to its on-board Requester.

PERMISSION 3.6:

Masters and Interrupt Handlers **MAY** release the DTB either during or after their last data transfer.

For example, if the Master notifies its on-board Requester that it no longer wants the bus during its last data transfer, the Requester releases BBSY* and arbitration takes place during the last transfer. But if the Master waits until the last transfer has completed before signaling its on-board Requester, the DTB remains idle while the arbitration is done. (This was illustrated in Section 2.5)

Chapters 2 and 4 contain RULEs that pertain to the release of the DTB.

SUGGESTION 3.3:

Design block transfer Master boards so that they signal their Requester to release BBSY* during the last data strobe of the block transfer. If it is released at the beginning of the block transfer, high priority bus requests initiated during the block transfer might not be taken into account by the Arbiter until the next arbitration cycle.

3.3.3.2 Acquisition Of The DTB

To ensure that no DTB line is ever driven to opposite states by two Masters or Interrupt Handlers, when these modules take control of the DTB they are constrained by the following rule.

RULE 3.12:

When a Master or Interrupt Handler is given control of the DTB by its on-board Requester, it **MUST** wait until it detects AS* high before turning on its DTB drivers.

OBSERVATION 3.11:

IF the previous Master or Interrupt Handler releases the bus **DURING** its last data transfer,

THEN RULE 3.12 ensures that the data transfer will be finished before the new Master or Interrupt Handler starts using the DTB. (If the previous Master or Interrupt Handler waited until the data transfer was finished before releasing the bus, AS* would already be high.)

3.3.3.3 Other Information**RECOMMENDATION 3.2:**

To allow for prompt servicing of interrupt requests and for optimum use of the DTB, design Masters so that they release the DTB as soon as possible after they detect BCLR* low.

PERMISSION 3.7:

A Master or Interrupt Handler **MAY** have more than one Requester, where each Requester generates bus requests on a different bus request line.

OBSERVATION 3.12:

Where a Master or Interrupt Handler has two or more Requesters, it can do high priority data transfers using one Requester and low priority transfers using another Requester.

3.4 Typical Operation

3.4.1 Arbitration Of Two Different Levels Of Bus Request

Figures 3-5 and 3-6 illustrate the sequence of events that takes place when two Requesters send simultaneous bus requests to a PRI Arbiter on different bus request lines. When the sequence begins, each of the Requesters drives its respective bus request line low (Requester A drives BR1* and Requester B drives BR2*). The Arbiter detects BR1* and BR2* low simultaneously, and it drives BG2IN* low to its own slot (Slot 1). That BG2IN* signal is monitored by Requester B (also in Slot 1). When Requester B detects BG2IN* low, it responds by driving BBSY* low. Requester B then releases the BR2* line and informs its own Master (Master B) that the DTB is available. When BBSY* goes low, the Arbiter drives BG2IN* of Slot 1 high.

When Master B completes its data transfer(s), and signals that fact by driving DEVICE WANTS BUS false, Requester B releases BBSY*, provided that its BG2IN* has been received high and 30 nanoseconds have elapsed since it released BR2*.

The Arbiter interprets the release of BBSY* as a signal to arbitrate any current bus requests. Since BR1* is the only bus request being driven low, the Arbiter grants the DTB to Requester A by driving BG1IN* low. Requester A responds by driving BBSY* low. When Master A completes its data transfer(s), and signals that fact by driving DEVICE WANTS BUS false, Requester A releases BBSY*, provided that its BG1IN* has been received high and 30 nanoseconds have elapsed since it released BR1*.

In this example, since no bus request lines are low when Requester A releases BBSY*, the Arbiter waits until it detects a bus request.

OBSERVATION 3.13:

The description illustrated in Figures 3-5 and 3-6 would hold for both PRI and RRS Arbiters, unless we consider an RRS Arbiter where the last active request was level BR2*. In this case, the Arbiter would process the BR1* request first and then proceed to the BR2* request.

OBSERVATION 3.14:

BBSY* and the bus grants are fully interlocked as shown in Figure 3-6:

- 1) The Arbiter does not drive the bus grant high until it detects BBSY* low.
- 2) The Requester does not release BBSY* to high until it detects the bus grant high.
- 3) The Arbiter does not drive the next bus grant low until it detects BBSY* high for 40 ns.
- 4) The next Requester does not drive BBSY* low until it detects the bus grant low.

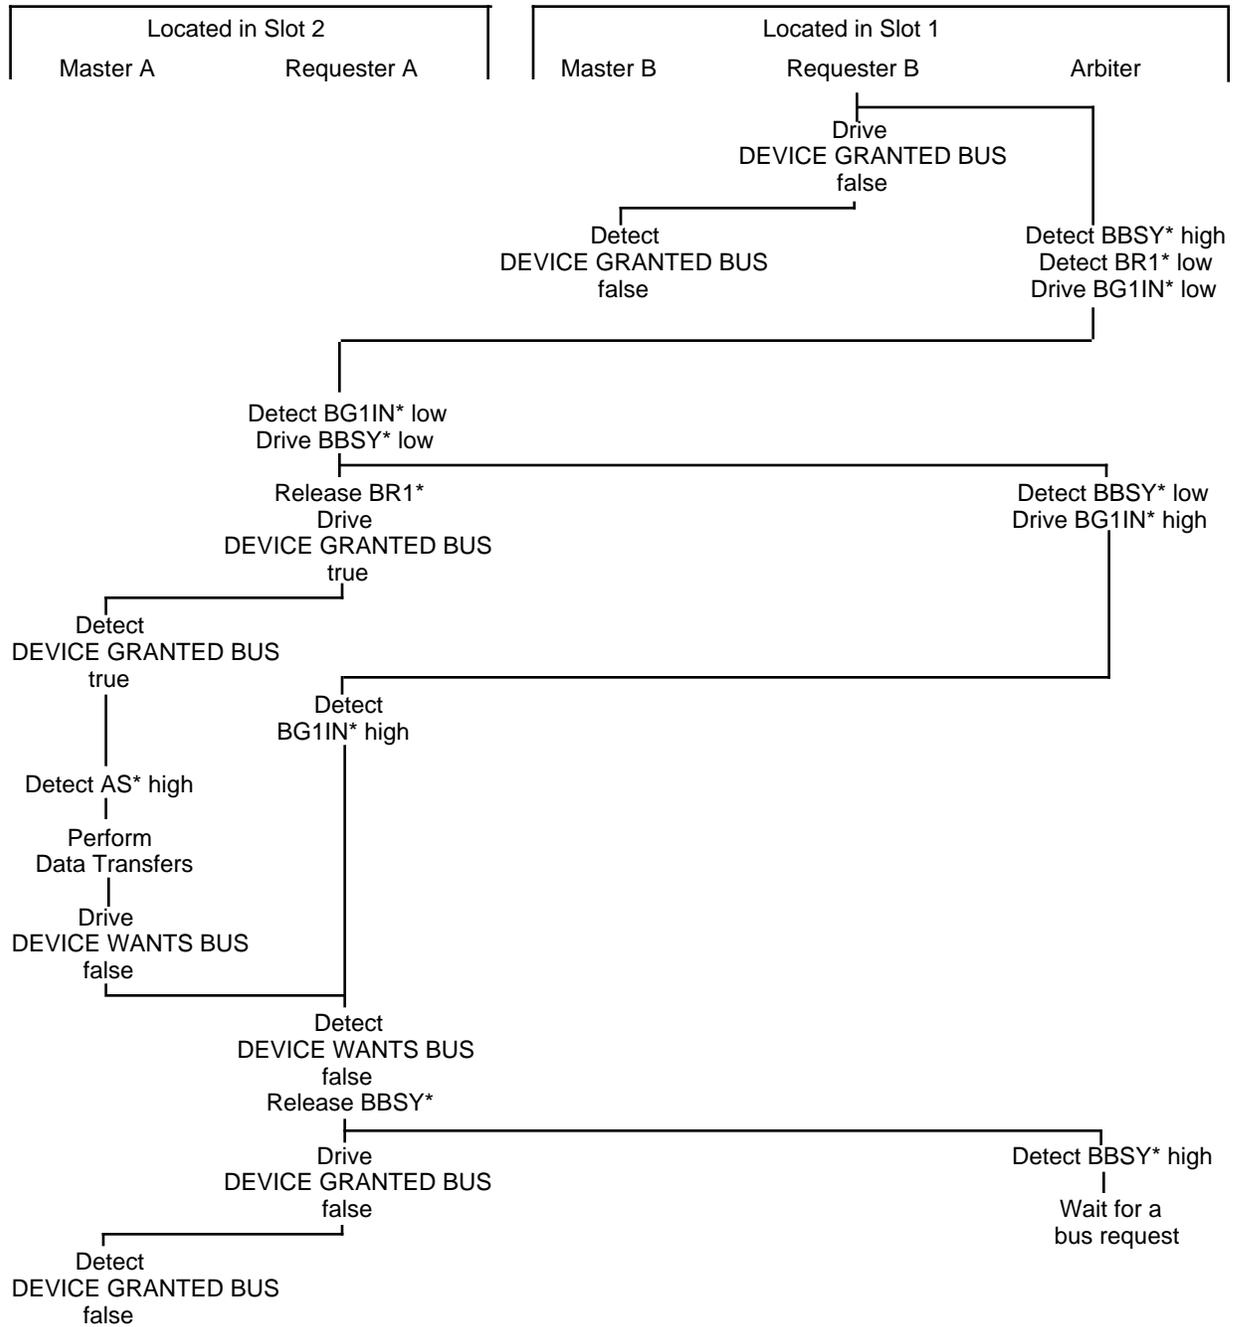
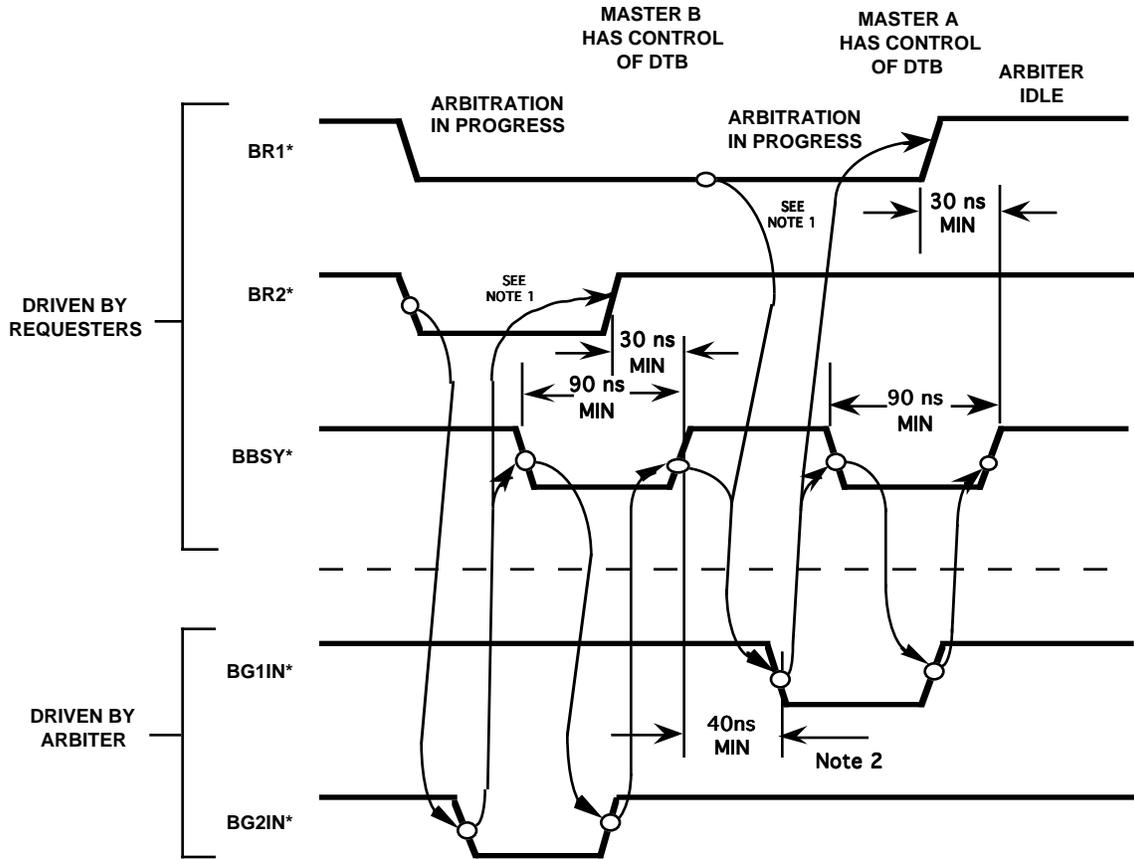


Figure 3 - 5 Arbitration Flow Diagram Two Requesters, Two Request Levels (Sheet 2 of 2)



Notes 1: In this example each Requester maintains its bus request line low until it is granted the DTB. In some cases a Requester might release its bus request line without receiving a bus grant (see Section 3.3.2).

2: This specification has been added to accommodate systems that require a minimum BBSY* high time such as VME64 LOCK protocols and Autobahn.

Figure 3 - 6 Arbitration Sequence Diagram Two Requesters, Two Request Levels

3.4.2 Arbitration Of Two Bus Requests On The Same Bus Request Line

Figures 3-7 and 3-8 illustrate the sequence of events which take place when an ROR Requester and an RWD Requester send simultaneous requests to a PRI Arbiter on a common bus request line. In this example, the Arbiter and RWD Requester are located on the system controller board in Slot 1, with the ROR Requester located in Slot 2. When the sequence begins, both of the Requesters drive BR1* low simultaneously. The Arbiter then drives BG1IN* low to its own slot (Slot 1). That BG1IN* signal is monitored by Requester A (also in Slot 1). When Requester A detects BG1IN* low, it responds by driving BBSY* low. Requester A then releases BR1* and informs Master A that the DTB is available.

OBSERVATION 3.15:

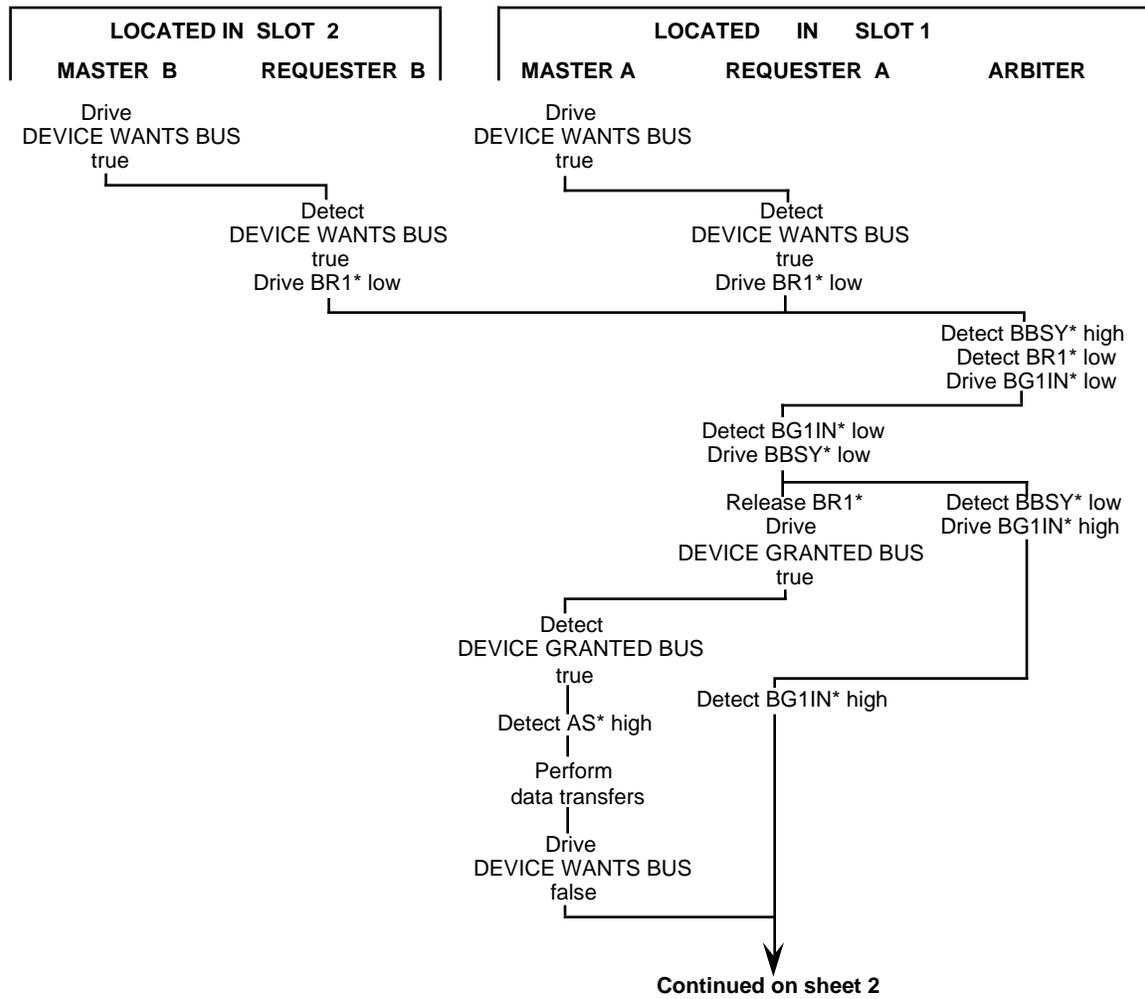
Even though Requester A releases BR1*, Requester B continues to drive it low (see Figures 3-7 and 3-8).

After detecting BBSY* low, the Arbiter drives BG1IN* high. When Master A has completed its data transfer(s), it drives DEVICE WANTS BUS false. When Requester A detects this, and when the 30 nanoseconds delay since the release of BR1* has been satisfied, Requester A releases BBSY*.

The Arbiter interprets the release of BBSY* as a signal to arbitrate any current bus requests. Since the BR1* line is still low, the Arbiter drives BG1IN* low again. When Requester A detects BG1IN* low, it drives its BG1OUT* low because it does not need the DTB. Requester B then detects the low on its BG1IN* and responds by driving BBSY* low. When the Arbiter detects the low on BBSY*, it drives BG1IN* high, which causes Requester A to drive its BG1OUT* high.

Some time later, when Master B has finished its data transfers, it drives DEVICE WANTS BUS false, indicating that it has finished using the DTB.

Since Requester B is an ROR Requester, it does not release BBSY*, but keeps it driven low. In the event that Master B needs to use the DTB again, no arbitration will be required. In this example, however, Requester A drives BR1* low, indicating a need to use the DTB, and Requester B (which is monitoring the bus request lines) releases the BBSY* line. The Arbiter then grants the DTB to Requester A.



Note: DEVICE WANTS BUS and DEVICE GRANTED BUS are on-board signals between the Master and its Requester (See Figure 3-4)

Figure 3 - 7 Arbitration Flow Diagram
Two Requesters, Same Request Level (Sheet 1 of 3)

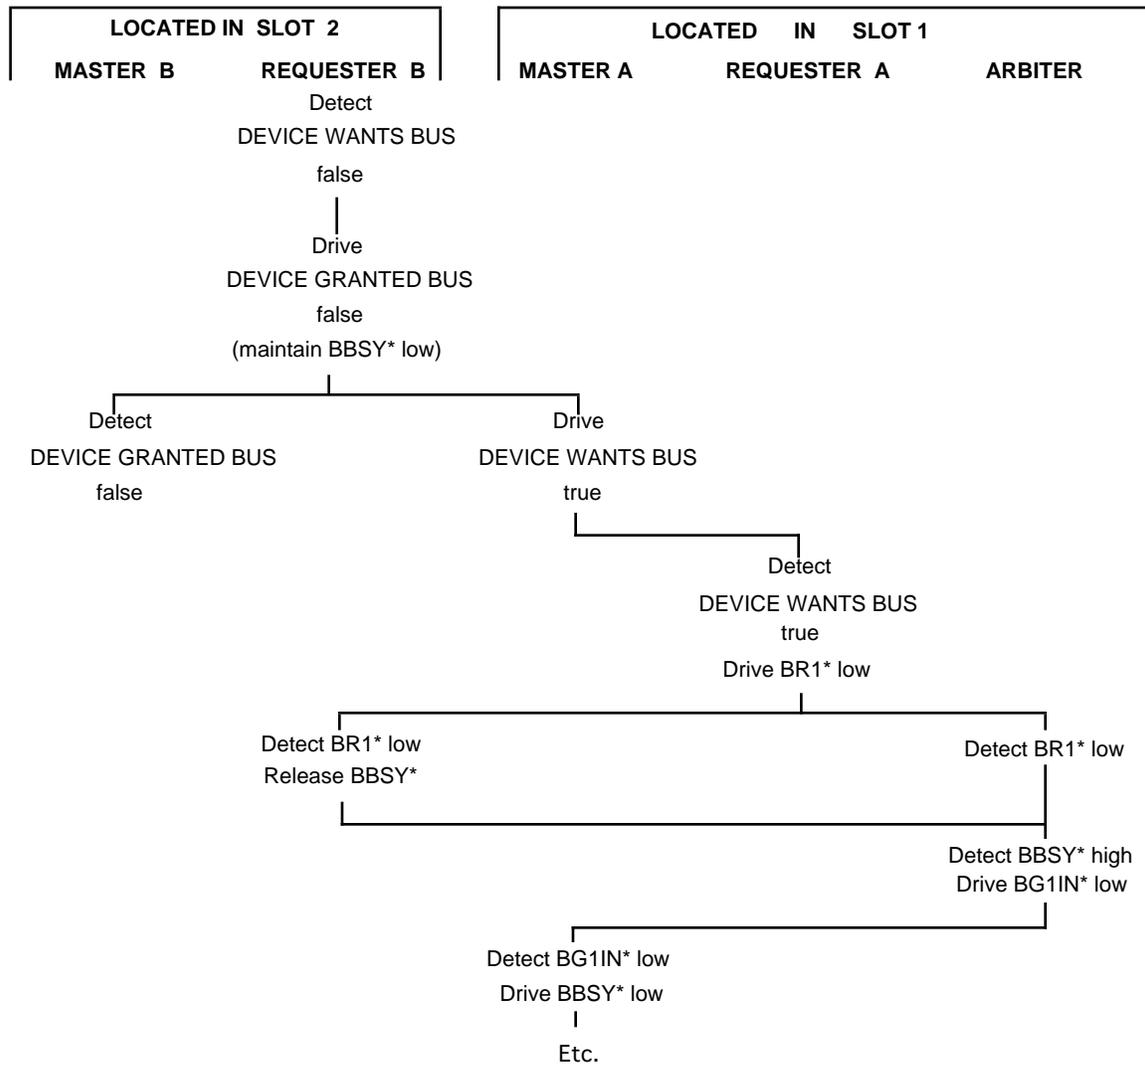
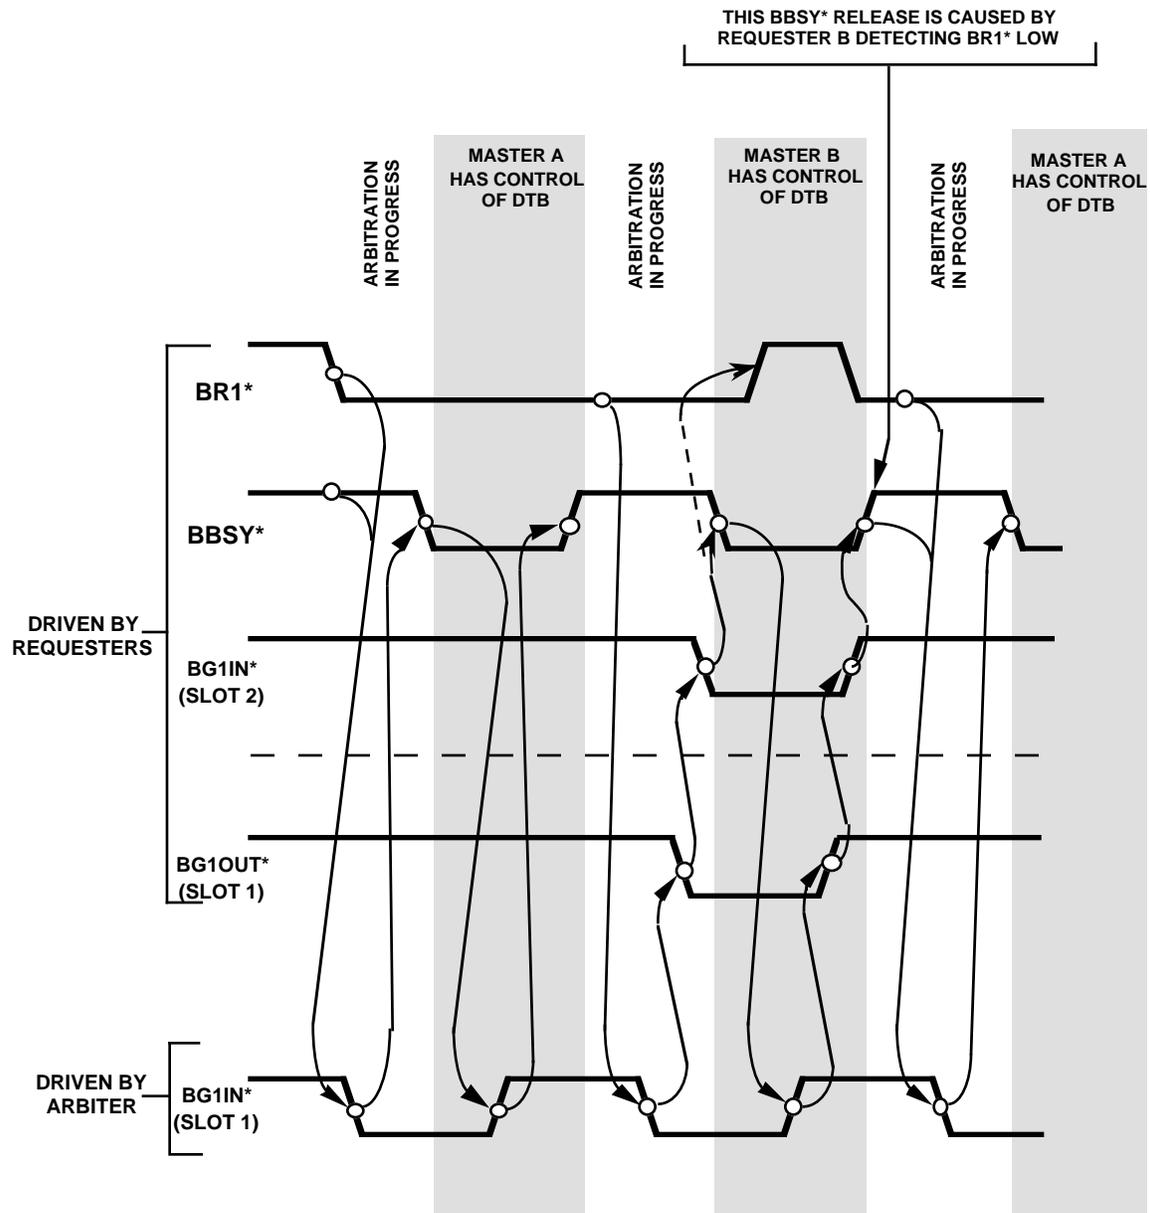


Figure 3 - 7 Arbitration Flow Diagram
Two Requesters, Same Request Level (Sheet 3 of 3)



Note: In this example each Requester maintains its bus request line low until it is granted the DTB. In some cases a Requester might release its bus request line without receiving a bus grant (see Section 3.3.2).

Figure 3 - 8 Arbitration Sequence Diagram
Two Requesters, Same Request Level

3.5 Race Conditions Between Master Requests and Arbiter Grants

Suppose that there are two Requesters, Requester A and Requester B, that share a common bus request line. Requester B, which is farther down the daisy-chain, requests the bus and the Arbiter drives the corresponding bus grant line low. This bus grant arrives at Requester A just as Master A signals that it wants the bus. If Requester A is improperly designed, this situation might cause it to momentarily drive its BG[3..0]OUT* line low and then high again resulting in a low-going transient.

RULE 3.13:

Requesters **MUST** be designed to ensure that no momentary low-going transients are generated on their BG[3..0]OUT* out line.

OBSERVATION 3.16:

If the Requester is designed such that it latches the state of the on-board DEVICE WANTS BUS line upon the falling edge of its bus grant in line, and if that signal is in transition when the falling edge occurs, the outputs of the latch will sometimes oscillate, or remain in the threshold region between the high and low levels, for a short time. Because of this, the VMEbus specification does not set a time limit for the Requester to pass along the bus grant. It only prohibits the Requester from generating a low-going transient on its BG[3..0]OUT* line which might be interpreted as a bus grant by a Requester further down the daisy-chain.

PERMISSION 3.8:

IF a Requester detects that its on-board Master needs the bus between the time that it receives a bus grant intended for another Requester and the time it would pass that bus grant on,

THEN it **MAY** treat the bus grant as its own. In this case the other Requester will maintain its bus request until another bus grant is issued.

Chapter 4

Priority Interrupt Bus

4.1 Introduction

The VMEbus includes a Priority Interrupt Bus which provides the signal lines needed to generate and service interrupts. Figure 4-1 shows a typical VMEbus system. Interrupters use the Priority Interrupt Bus to send interrupt requests to Interrupt Handlers which respond to these requests.

Any system which has interrupt capability includes software routines that are called interrupt service routines, which are invoked by the interrupts. Interrupt subsystems can be classified into two groups:

- a. Single handler systems, which have only one Interrupt Handler that receives and services all bus interrupts.
- b. Distributed systems, which have two or more Interrupt Handlers that receive and service bus interrupts.

4.1.1 Single Handler Systems

In a single handler system, all interrupts are received by one Interrupt Handler and all interrupt service routines are executed by one processor. Figure 4-2 shows the interrupt structure of a single handler system. This type of architecture is well suited to machine or process control applications, where a supervisory processor coordinates the activities of dedicated processors. The dedicated processors are typically interfaced to the machine or the process being controlled.

The supervisory processor is the destination for all bus interrupts, servicing them in a prioritized manner. The dedicated processors are not required to service interrupts from the bus and can give primary attention to controlling a machine or process.

4.1.2 Distributed Systems

Figure 4-3 shows the interrupt structure of a distributed system. This system includes two or more Interrupt Handlers, each servicing only a subset of the bus interrupts. In a typical implementation, each of the Interrupt Handlers resides on a different processor board. This type of architecture is well suited to distributed computing applications, where multiple, co-equal processors execute the application software. As each of the co-equal processors executes part of the system software, it might need to communicate with the other processors. In the distributed system, each processor services only those interrupts directed to it, establishing dedicated communication paths among all processors.

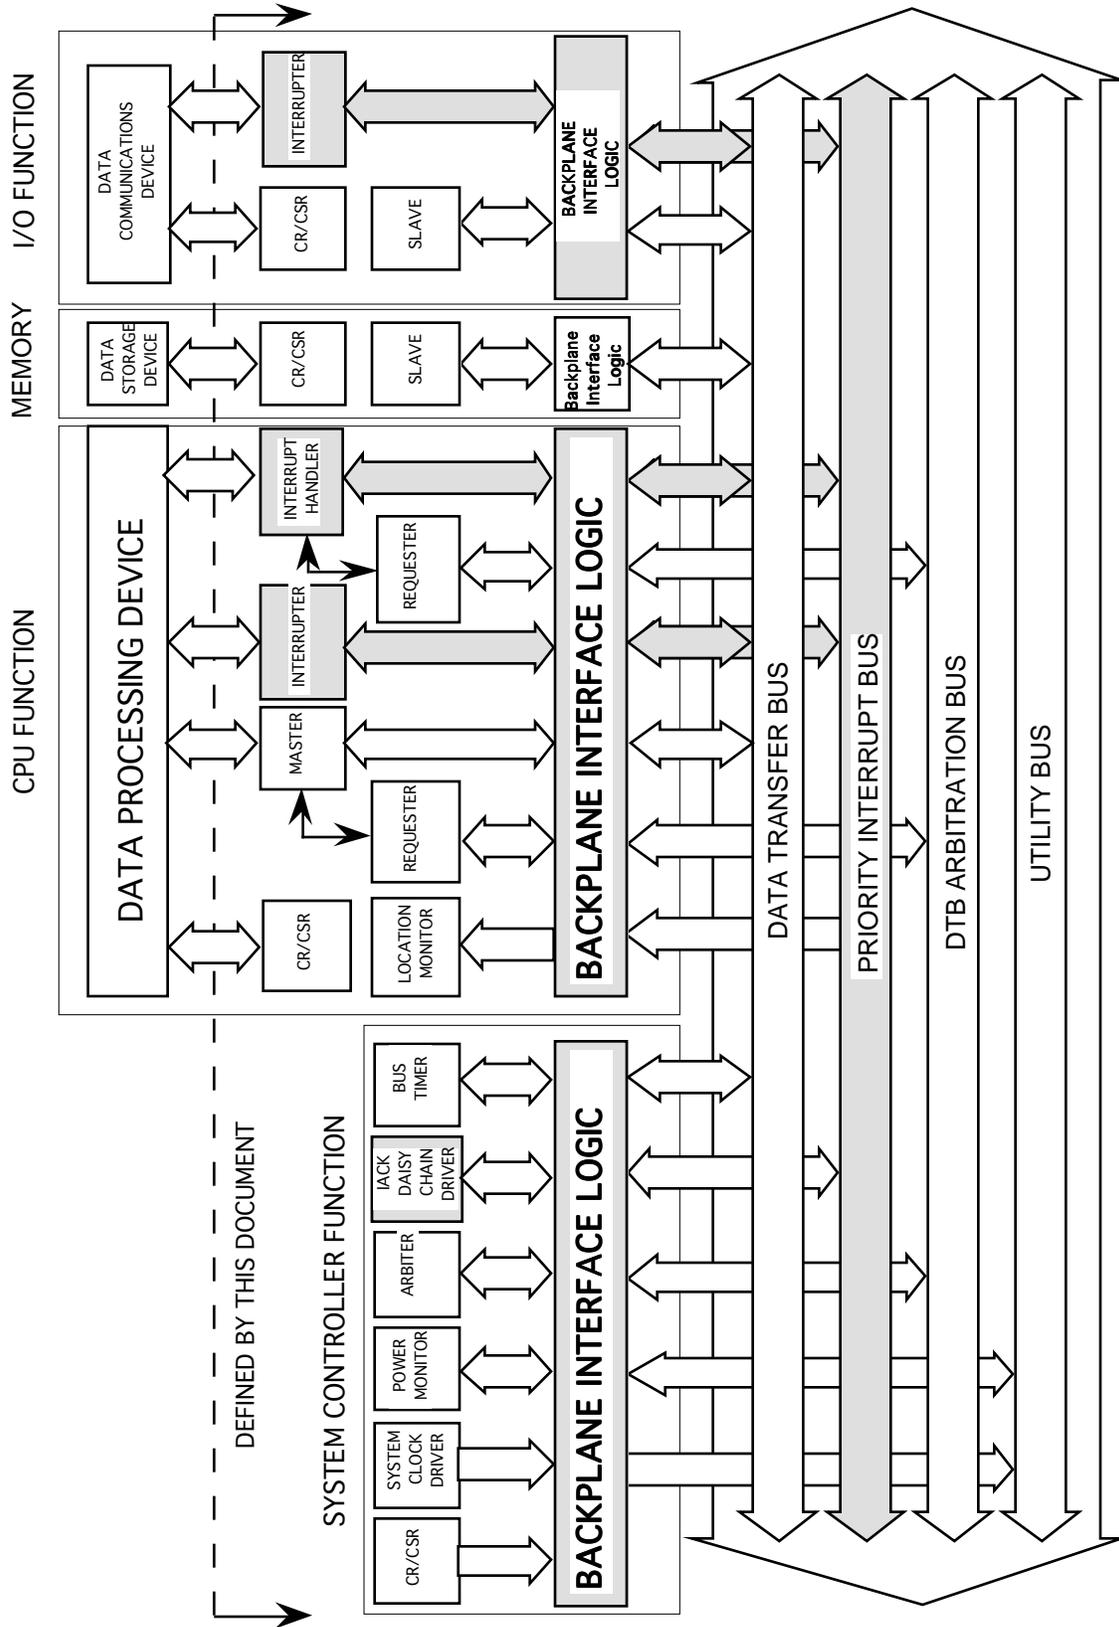


Figure 4 - 1 Priority Interrupt Bus Functional Diagram

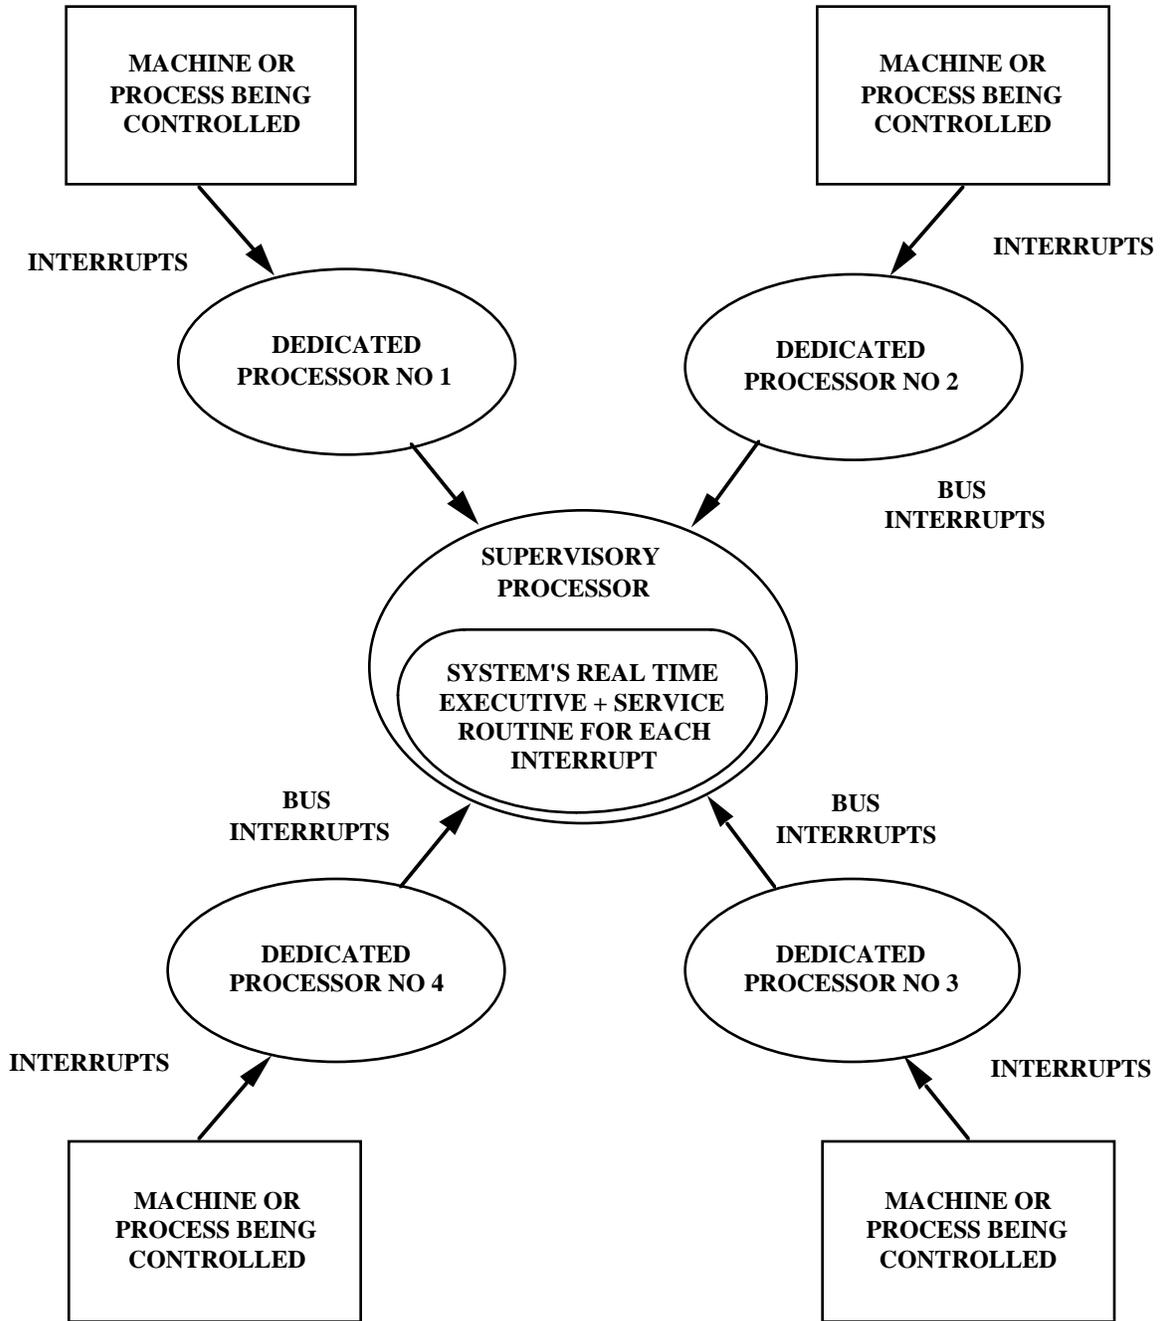


Figure 4 - 2 Interrupt Subsystem Structure: Single Handler System

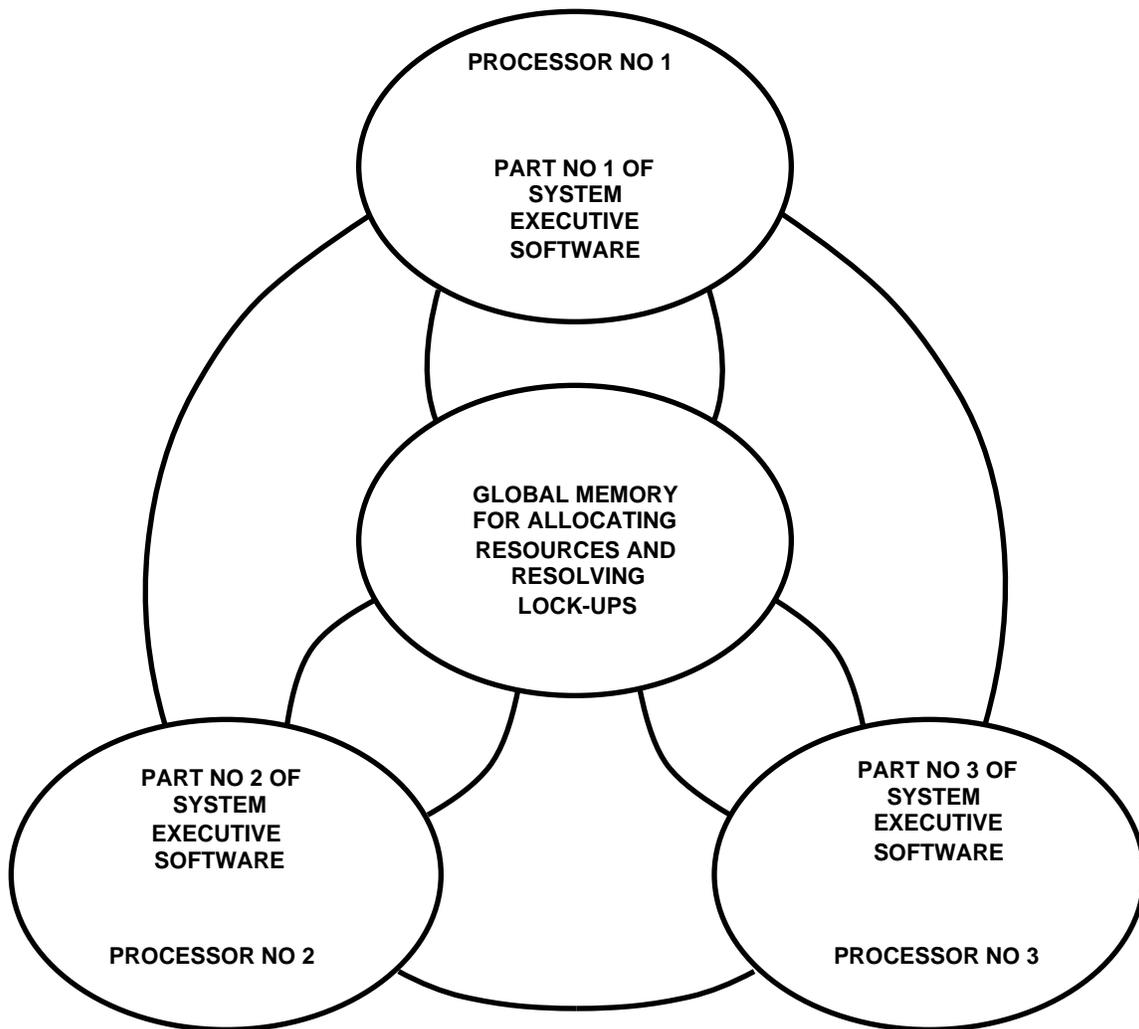


Figure 4 - 3 Interrupt Subsystem Structure: Distributed System

4.2 Priority Interrupt Bus Lines

The Data Transfer Bus, the Arbitration Bus, and the Priority Interrupt Bus are all used in the process of generating and handling bus interrupts.

The following discussion of the Priority Interrupt Bus assumes that the reader understands the operation of the Data Transfer Bus described in Chapter 2, and the Arbitration Bus described in Chapter 3.

The Priority Interrupt Bus consists of seven interrupt request signal lines, one interrupt acknowledge line, and one interrupt acknowledge daisy-chain:

IRQ1*	Interrupt Request 1
IRQ2*	Interrupt Request 2
IRQ3*	Interrupt Request 3
IRQ4*	Interrupt Request 4
IRQ5*	Interrupt Request 5
IRQ6*	Interrupt Request 6
IRQ7*	Interrupt Request 7
IACK*	Interrupt Acknowledge
IACKIN*/IACKOUT*	Interrupt Acknowledge Daisy-Chain

4.2.1 Interrupt Request Lines

Interrupters request interrupts by driving an interrupt request line low. In a single handler system, these interrupt request lines are prioritized, with IRQ7* having the highest priority.

4.2.2 Interrupt Acknowledge Line

The IACK* line runs the full length of the backplane and is connected to the IACKIN* pin of Slot 1 (see Figure 4-4). When driven low, the IACKIN* pin causes the IACK Daisy-Chain Driver, located in Slot 1, to propagate a falling edge down the interrupt acknowledge daisy-chain.

4.2.3 Interrupt Acknowledge Daisy-Chain - IACKIN*/IACKOUT*

Each of the seven interrupt request lines can be shared by two or more Interrupter modules. The interrupt acknowledge daisy-chain assures that only one Interrupter responds to the interrupt acknowledge cycle. This daisy-chain line passes through each board on the VMEbus. Each Interrupter that is driving an interrupt request line low waits for a falling edge to arrive at its IACKIN* daisy-chain input. Only upon receiving this falling edge does an Interrupter respond to an Interrupt Acknowledge Cycle. It does not pass the falling edge down the Daisy-Chain, preventing other Interrupters from responding to the Interrupt Acknowledge Cycle.

RULE 4.1:

IF	a VMEbus backplane slot is not occupied by a board, and if there are boards farther down the interrupt acknowledge daisy-chain,
THEN	jumpers or a functionally equivalent mechanism MUST be installed at the empty slot to pass-through the daisy-chain signal.

RULE 4.50:

IF	a VMEbus board does not generate interrupt requests
THEN	it MUST propagate the IACK daisy chain from its IACKIN* line to its IACKOUT* line.

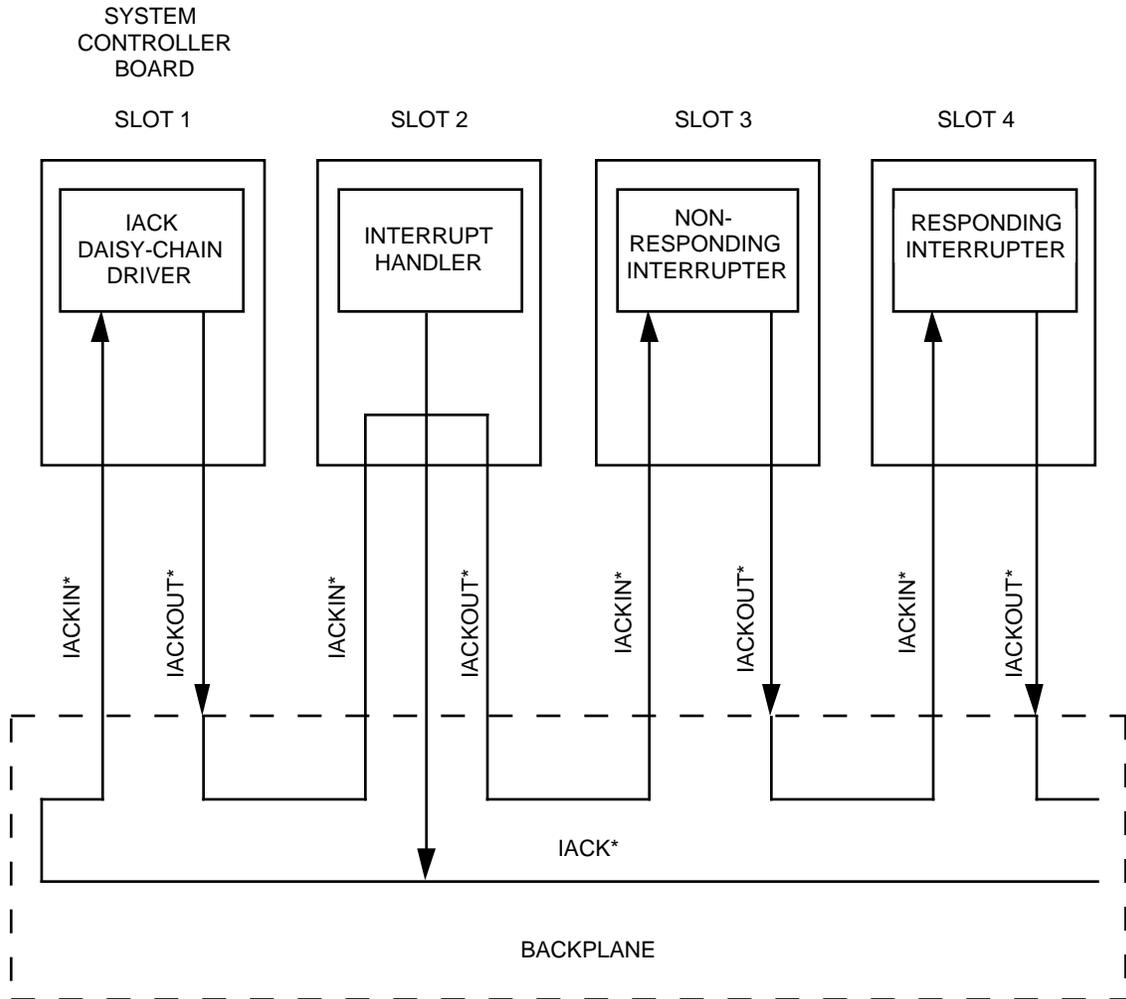


Figure 4 - 4 IACKIN*/IACKOUT* DAISY-CHAIN

4.3 Priority Interrupt Bus Modules - Basic Description

There are three types of functional modules associated with the Priority Interrupt Bus: Interrupters, Interrupt Handlers, and an IACK Daisy-Chain Driver. The capabilities of Interrupt Handlers and Interrupters are described by a list of mnemonics that show what interrupt acknowledge cycle types they can generate and accept.

Sections 4.3.1 through 4.3.3 provide block diagrams for the three Priority Interrupt Bus modules: Interrupt Handler, Interrupter, and IACK Daisy-Chain Driver.

RULE 4.2:

Output signal lines shown with solid lines in Figures 4-5 through 4-7 MUST be driven by the module, unless it would always drive them high.

OBSERVATION 4.1a:

IF a bused output line is not driven,
THEN terminators on the backplane ensure that it is high.

RULE 4.3:

Input signal lines shown with solid lines in Figures 4-5 through 4-7 MUST be monitored and responded to in the appropriate fashion.

OBSERVATION 4.2:

RULEs and PERMISSIONs for driving and monitoring signal lines shown with dotted lines in Figures 4-5 through 4-7, are given in Tables 4-1 and 4-2.

4.3.1 Interrupt Handler

The Interrupt Handler is used to accomplish several tasks:

- a. It prioritizes the incoming interrupt requests within its assigned group of interrupt request lines (highest of IRQ[7..1]*).
- b. It uses its on-board Requester to request the DTB and, when granted use of the DTB, initiates an interrupt acknowledge cycle, reading a Status/Id from the Interrupter being acknowledged.
- c. It initiates the appropriate interrupt servicing sequence based on the information received in the Status/Id.

OBSERVATION 4.3:

The VMEbus specification does not dictate what will happen during the interrupt servicing sequence. Servicing of the interrupt might or might not involve use of the VMEbus.

The Interrupt Handler uses the DTB to read a Status/Id from the Interrupter. In this respect, the Interrupt Handler acts like a Master and the Interrupter acts like a Slave. However, there are four important differences. The Interrupt Handler:

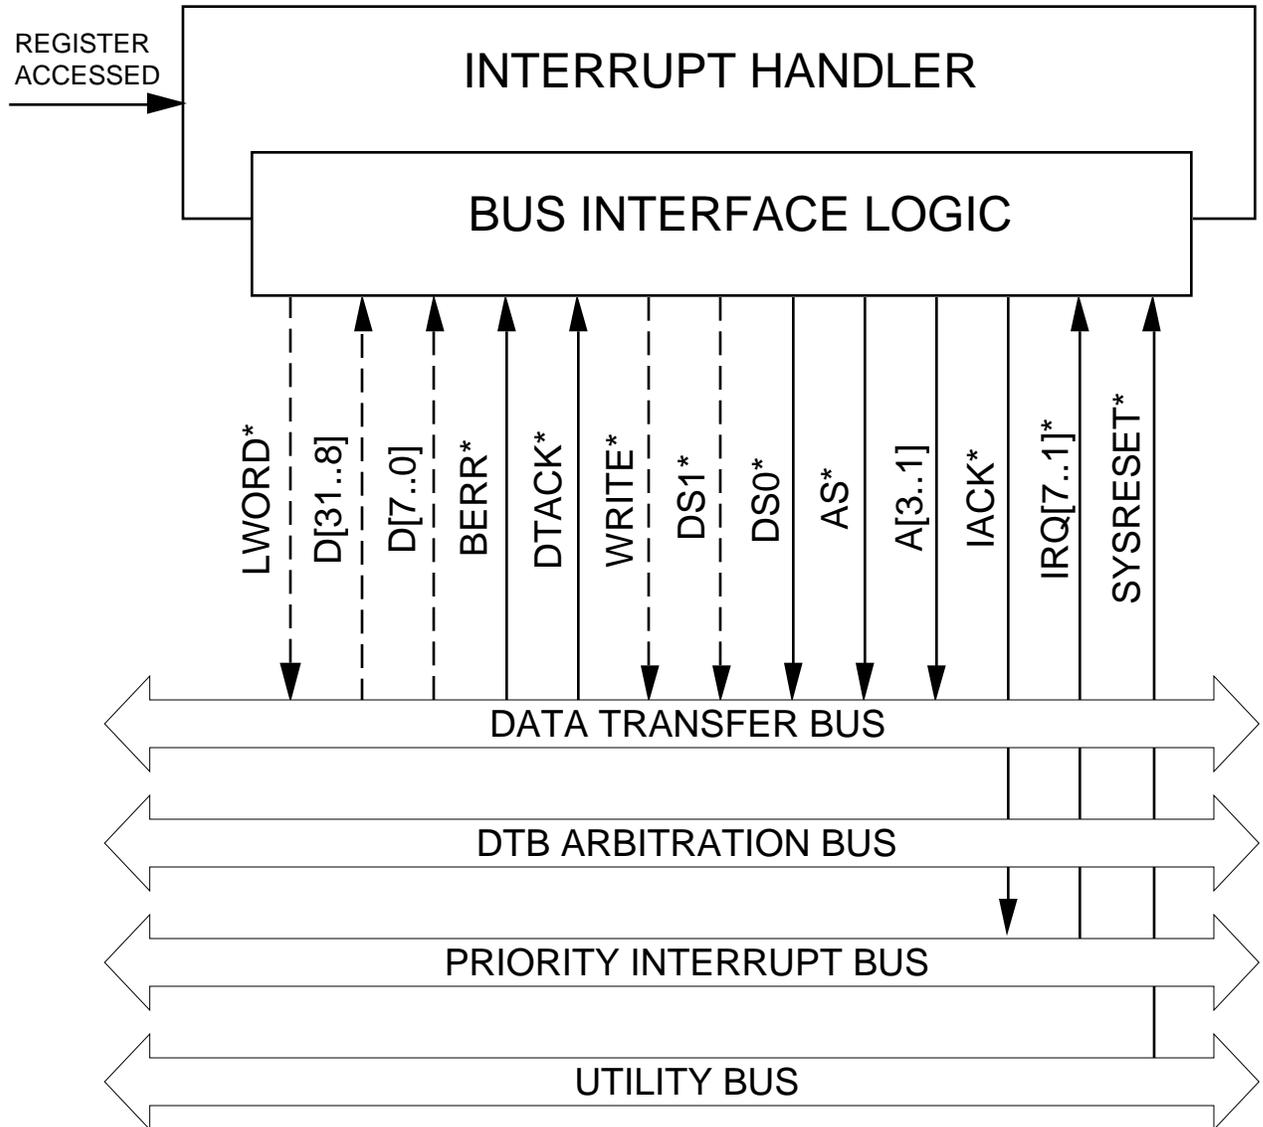
- a. always drives IACK* low.
- b. is not required to drive the address modifier lines.
- c. only uses the lowest three address lines A[3..1].
- d. never drives the data bus.

The Interrupt Handler always drives IACK* low when it accesses the bus. The Master either drives it high or does not drive it at all.

The Interrupt Handler does not have to drive the address modifier lines with a valid code and it only drives the lowest three address lines A[3..1], with valid information. The levels of these three address lines indicate which of the seven interrupt request lines is being acknowledged, as shown in Table 4-7. A Master drives the address bus with the address of the Slave being accessed and provides an address modifier code on the address modifier lines.

The Interrupt Handler does not drive the data lines (i.e., it does not “write” to the Interrupter) and does not have to drive the WRITE* line. A Master uses the data lines to a Slave bidirectionally and, during normal use, drives WRITE* low or high as required.

A block diagram of the Interrupt Handler is shown in Figure 4-5.



Note:

The RULES and PERMISSIONS for monitoring and driving the dotted lines are given in Table 4-1.

Figure 4 - 5 Block Diagram: Interrupt Handler

Table 4 - 1 RULES And PERMISSIONS That Specify The Use Of The Dotted Lines In Figure 4-5 By The Various Types Of Interrupt Handlers

Type of Interrupt Handler	Use of dotted lines
D08(O)	MUST monitor D[7..0]. MAY drive LWORD* and DS1*. MAY monitor D[31..0].
D16	MUST drive DS1*. MUST monitor D[15..0]. MAY drive LWORD*. MAY monitor D[31..16].
D32	MUST drive DS1* and LWORD*. MUST monitor D[31..0].
ALL	MUST NOT drive WRITE* low.

Note: The mnemonics D08(O), D16, and D32 are defined in Table 4-5.

4.3.2 Interrupter

The Interrupter functions as follows:

- a. It requests an interrupt from the Interrupt Handler which monitors its interrupt request line.
- b. IF it receives a falling edge on the interrupt acknowledge daisy-chain input, THEN
 - IF it is requesting an interrupt AND the levels on the three valid address lines correspond to the interrupt request line it is using, AND the width of the requested Status/Id is either equal to, or greater than the size it can supply,
 - THEN it supplies a Status/Id,
 - ELSE it passes the falling edge down the interrupt acknowledge daisy-chain.

Each Interrupter module drives only one interrupt request line. The VMEbus specification describes a board that generates interrupt requests on several interrupt lines as having several Interrupter modules.

PERMISSION 4.1:

Since the Interrupter is just a conceptual model, logic on a VMEbus board MAY be shared between several Interrupter modules.

The Interrupter uses one of seven lines to request an interrupt. It then monitors the lowest three lines of the address bus A[3..1], the IACKIN*/IACKOUT* daisy-chain, and optionally IACK*, to determine when its interrupt is being acknowledged. When acknowledged, it places its Status/Id on the data bus and signals the Interrupt Handler that the Status/Id is valid by driving DTACK* low.

There are five primary differences in the use of the DTB by the Interrupter and the Slave. The Interrupter:

- a. only responds when its IACKIN* is low.
- b. does not have to monitor the address modifier lines.
- c. only monitors the lowest three address lines.
- d. does not monitor the WRITE* line.
- e. is permitted to respond with data of a different size than that requested.

The Slave monitors AS*, and interprets a falling edge on AS* as the signal that a valid bus cycle is in progress. It then proceeds to decode the appropriate number of address lines and the address modifier lines and then, based on this information, it determines whether it was addressed. However, the Slave responds only if IACK* is high.

The Interrupter, on the other hand, interprets the falling edge on its IACKIN* line as a signal that it can respond to the interrupt acknowledge cycle in progress. It decodes only the lowest three address lines A[3..1], ignoring the address modifier lines.

The Interrupter does not need to monitor WRITE*, since it is never written to. Slaves need to monitor WRITE* so that they can distinguish read cycles from write cycles.

The Interrupter places a Status/Id on the bus and responds with DTACK*, even if the LWORD*, DS1*, and DS0* lines call for a Status/Id whose width is greater than the Interrupter is able to provide. For example, the Interrupt Handler might drive LWORD* and both data strobes low, indicating that it will read 32-bits of Status/Id from D[31..0], but a D08(O) Interrupter would still respond with its 8-bit Status/Id on D[7..0]. In contrast, when a Slave cannot provide the requested data width, it either responds with BERR* or does not respond at all, typically resulting in a bus time-out.

OBSERVATION 4.4:

When an Interrupter places a Status/Id on the data bus, any undriven data lines are read by the Interrupt Handler as high because of the bus terminators. For example, if a D16 Interrupt Handler initiates a double byte interrupt acknowledge cycle, a D08(O) Interrupter would place an 8-bit Status/Id on D[7..0]. The upper 8 bits, read by the Interrupt Handler from D[15..8], are read as ones (high), since they are not driven by the D08(O) Interrupter.

RULE 4.4:

Before responding to an interrupt acknowledge cycle, the Interrupter:

1. MUST have an interrupt request pending.
2. The level of that request MUST match the level indicated on A[3..1].
3. The width of the requested Status/Id MUST be equal to or greater than the size it can respond with.
4. It MUST have received an incoming falling edge on its IACKIN* daisy-chain input.

IF any of the four conditions above are not met.

THEN the Interrupter MUST NOT respond to the interrupt acknowledge cycle.

IF condition 4 is met, but either 1, 2, or 3 is not.

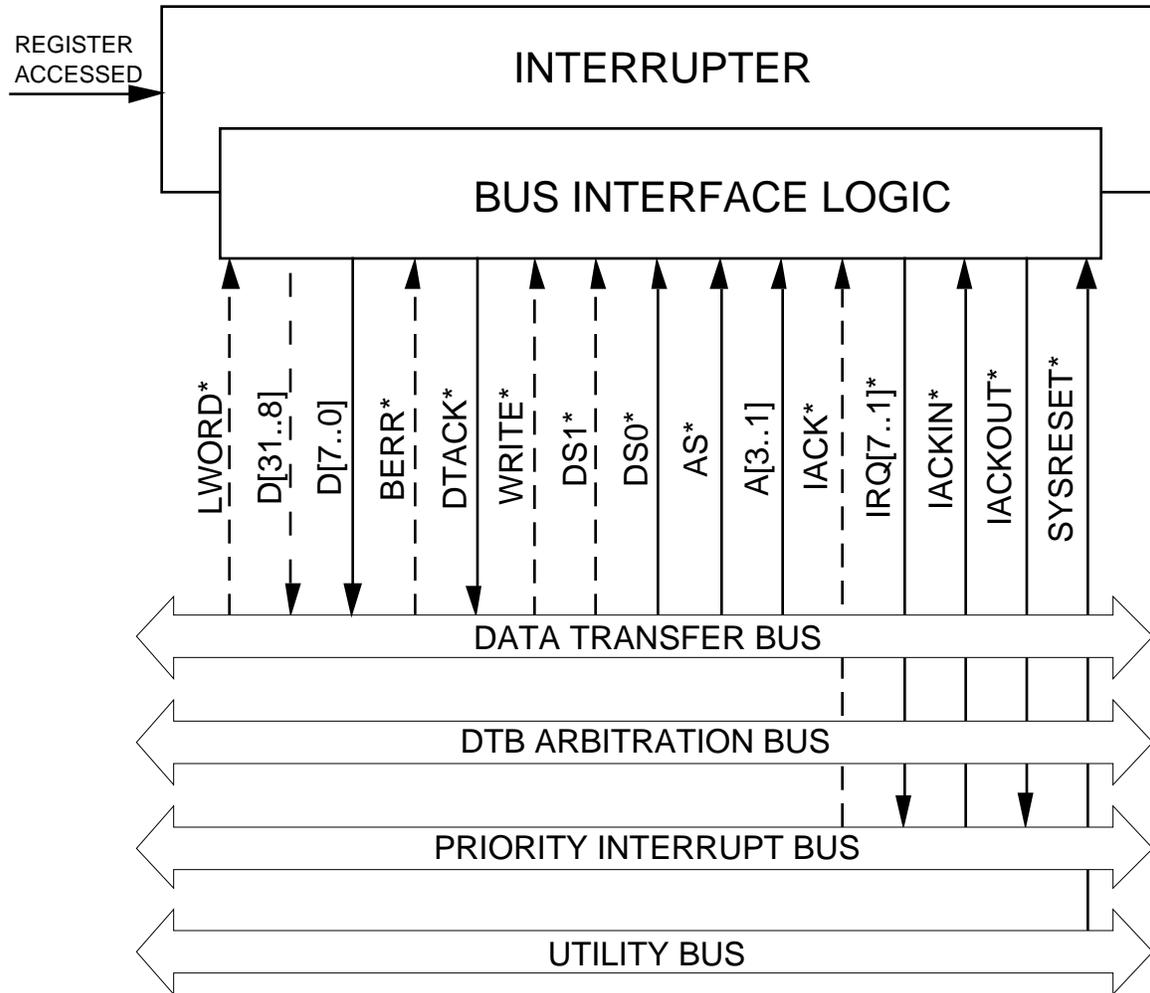
THEN the Interrupter MUST pass the falling edge of IACKIN* to the next Interrupter module in the daisy-chain.

A block diagram of the Interrupter is shown in Figure 4-6.

Table 4 - 2 RULES And PERMISSIONS That Specify The Use Of The Dotted Lines In Figure 4 - 6 By The Various Types Of Interrupters

Type of Interrupter	Use of dotted lines
D08(O)	MUST drive D[7..0]. MAY drive D[31..8]. MAY monitor LWORD* or DS1*.
D16	MUST monitor DS1*. MUST drive D[15..0]. MAY drive D[31..16] low. MAY monitor LWORD*.
D32	MUST monitor DS1* and LWORD*. MUST drive D[31..0].
ALL	MAY monitor WRITE* and IACK*. MAY drive BERR*.

Note: The mnemonics D08(O), D16, and D32 are defined in Table 4-5.



Notes :

1. The RULES and PERMISSIONS for driving and monitoring the dotted lines are given in Table 4-2.
2. The Register Accessed input signal is present on RORA Interrupters only.

Figure 4 - 6 Block Diagram: Interrupter

4.3.3 IACK Daisy-Chain Driver

The IACK Daisy-Chain Driver is another module that interacts with Interrupt Handlers and Interrupters to coordinate the servicing of interrupts. It generates a falling edge on the interrupt acknowledge daisy-chain each time an Interrupt Handler initiates an interrupt acknowledge cycle.

A block diagram of the IACK Daisy-Chain Driver is given in Figure 4-7.

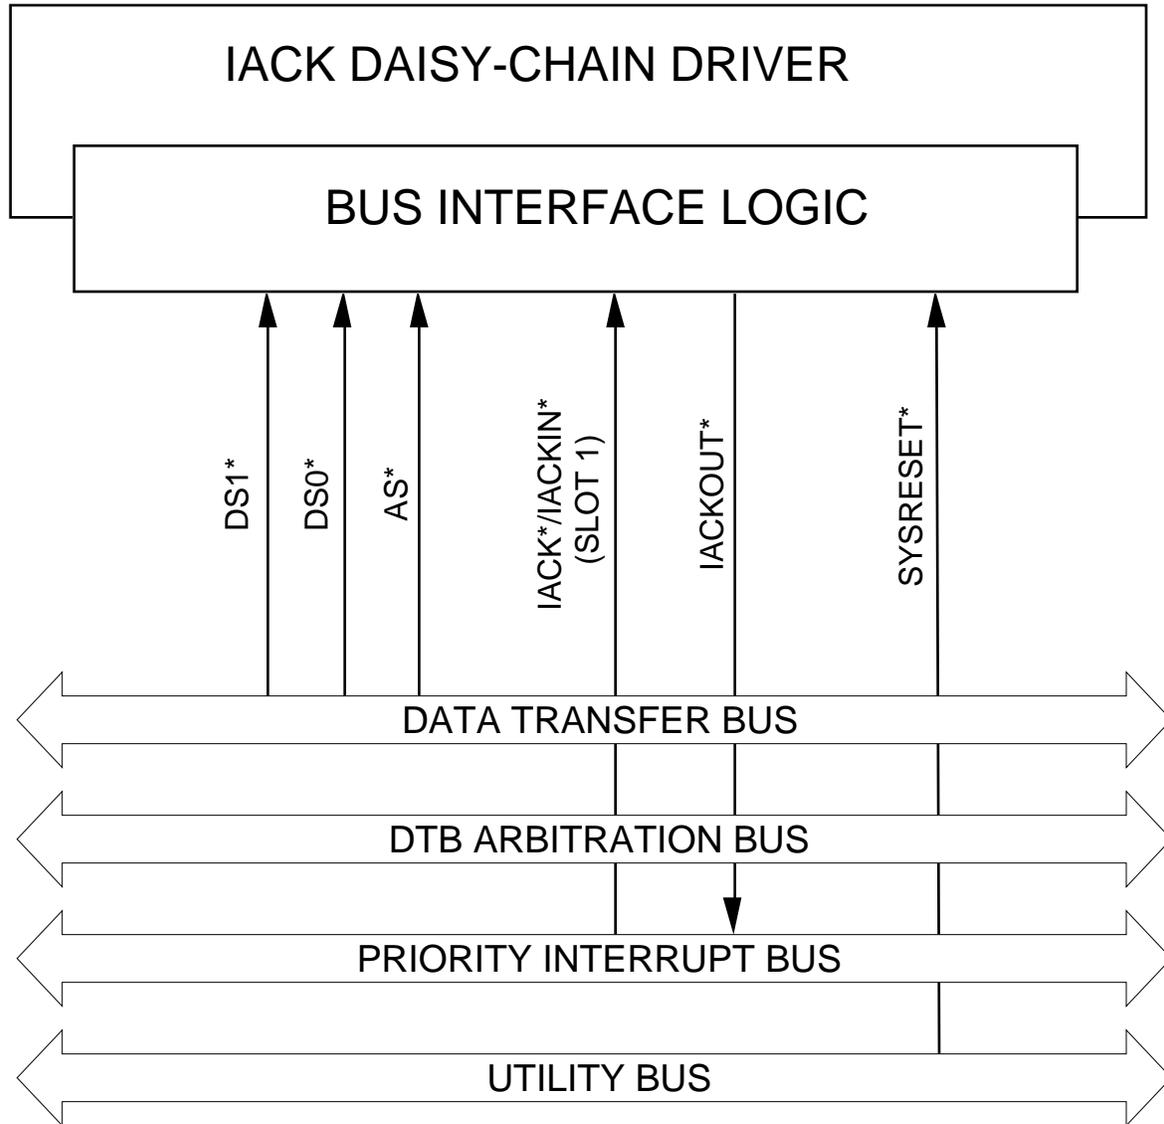


Figure 4 - 7 Block Diagram: IACK Daisy-Chain Driver

4.3.4 Interrupt Request Handling Capabilities

Interrupt Handlers can be designed to handle interrupt requests received on one to seven interrupt request lines. Table 4-3 shows how the IH() mnemonic is used to describe the interrupt handling capabilities of Interrupt Handlers.

Table 4 - 3 Use Of The IH() Mnemonic To Specify Interrupt Request Handling Capabilities

The Following Mnemonic	When Applied to an	Means that it
IH(x-y)	Interrupt Handler	can generate interrupt acknowledge cycles in response to interrupt requests on lines IRQx* through IRQy*.
IH(x)	Interrupt Handler	can generate interrupt acknowledge cycles in response to interrupt requests on line IRQx*.

4.3.5 Interrupt Request Generation Capabilities

Interrupters can be designed to generate an interrupt request on any of the seven interrupt request lines. Table 4-4 shows how the I() mnemonic is used to describe the interrupt request generation capabilities of Interrupters.

Table 4 - 4 Use Of The I() Mnemonic To Specify Interrupt Request Generation Capabilities

The Following Mnemonic	When Applied to an	Means that it
I(x)	Interrupter	can generate interrupt requests on line IRQx*

4.3.6 Status/Id Transfer Capabilities

There are three Status/Id transfer capabilities: D08(O), D16 and D32. Table 4-5 shows how these mnemonics are used to describe the interrupt handling capabilities of Interrupt Handlers and Interrupters.

Table 4 - 5 Mnemonics That Specify Status/Id Transfer Capabilities

The Following Mnemonic	When Applied to an	Means that it
D08(O)	Interrupter	responds to 8-bit, 16-bit, and 32-bit interrupt acknowledge cycles by providing an 8-bit Status/Id on D[7..0].
	Interrupt Handler	generates 8-bit interrupt acknowledge cycles in response to the requests on the interrupt request line(s) and reads an 8-bit Status/Id from D[7..0].
D16	Interrupter	responds to 16-bit and 32-bit interrupt acknowledge cycles by providing a 16-bit Status/Id on D[15..0].
	Interrupt Handler	generates 16-bit interrupt acknowledge cycles in response to the requests on the interrupt request line(s) and reads a 16-bit Status/Id from D[15..0].
D32	Interrupter	responds to 32-bit interrupt acknowledge cycles by providing a 32 bit Status/Id on D[31..0].
	Interrupt Handler	generates 32-bit interrupt acknowledge cycles in response to the requests on the interrupt request line(s) and reads a 32-bit Status/Id from D[31..0].

4.3.7 Interrupt Request Release Capabilities

Many widely used peripheral ICs generate interrupt requests. Unfortunately, there is no standard method for indicating to these ICs when it is time for them to remove their interrupt request from the bus. Three methods are used:

- a. When the relevant processor senses an interrupt request from a peripheral device, it enters an interrupt service routine, and Reads a status register in the device. The peripheral device interprets this read cycle on its status register as a signal to remove its interrupt request.
- b. When the relevant processor senses an interrupt request from a peripheral device, it enters an interrupt service routine, and Writes to a control register in the device. The peripheral device interprets this write cycle to its control register as a signal to remove its interrupt request.
- c. When the relevant processor senses an interrupt request from a peripheral device, it reads a Status/Id from the device. The peripheral device interprets this read cycle as a signal to remove its interrupt request.

The VMEbus specification calls Interrupters that use methods a and b Release On Register Access (RORA) Interrupters, and those that use method c Release On Acknowledge (ROAK) Interrupters. Figure 4-8 shows how an ROAK Interrupter releases its interrupt request line when the Interrupt Handler reads its Status/Id and how an RORA Interrupter releases its interrupt request upon an access to a control or status register.

OBSERVATION 4.5:

The Slave that provided the access to the Interrupter's control or status register is typically on the same board as the Interrupter, and it generates an on-board signal to the Interrupter when it has completed the register access.

RULE 4.5:

An RORA Interrupter **MUST NOT** release its interrupt request line before it detects a falling edge on DSA* during the register access cycle. It **MUST** release the interrupt request line within 2 microseconds after the last data strobe goes high at the end of the register access cycle.

RULE 4.6:

An ROAK Interrupter **MUST NOT** release its interrupt request line before it detects a falling edge on DSA* during the interrupt acknowledge cycle which acknowledges its interrupt, and it **MUST** release its interrupt request line within 500 nanoseconds after the last data strobe goes high at the end of the Status/Id read cycle.

RECOMMENDATION 4.2:

RULEs 4.5 and 4.6 have been retained for compatibility with previous versions of the VMEbus specification. The recommended practice for new designs is that the Interrupt REQUEST line be released as soon as possible but in all cases less than 50ns from data strobe high of the register access cycle (RULE 4.5) or the Status Id read cycle (RULE 4.6).

OBSERVATION 4.51:

The reason for speeding up the release time is to reduce the possibility of software race condition. 500ns can now be a large number of CPU instructions.

RULE 4.7:

Both RORA and ROAK Interrupters **MUST** provide a Status/Id during the interrupt acknowledge cycle that was initiated in response to their interrupt request.

RULE 4.8:

After an Interrupt Handler initiates an interrupt acknowledge cycle, and reads the Status/Id from an RORA Interrupter, it **MUST** ignore the low level on the interrupt request line for 2 microseconds after its on-board signal REGISTER ACCESSED goes true.

OBSERVATION 4.6:

RULE 4.8 prevents the Interrupt Handler from misinterpreting the low level on that line as a new interrupt request.

OBSERVATION 4.7:

The Master that accesses the Interrupter's control or status register is typically on the same board as the Interrupt Handler, and it generates an on-board signal to the Interrupt Handler when it has completed the register access.

PERMISSION 4.2:

IF a procedure is established to allow the Master to signal the Interrupt Handler that an access to the Interrupter's control or status registers has taken place.

THEN the Master and Interrupt Handler **MAY** reside on different boards.

Table 4-6 shows how the RORA and ROAK mnemonics are used to describe Interrupters.

Table 4 - 6 Mnemonics That Specify Interrupt Request Release Capabilities

The Following Mnemonic	When Applied to an	Means that it
RORA	Interrupter	releases its interrupt request line when a Master accesses an on-board status or control register.
ROAK	Interrupter	releases its interrupt request line when its Status/Id is read during an interrupt acknowledge cycle.

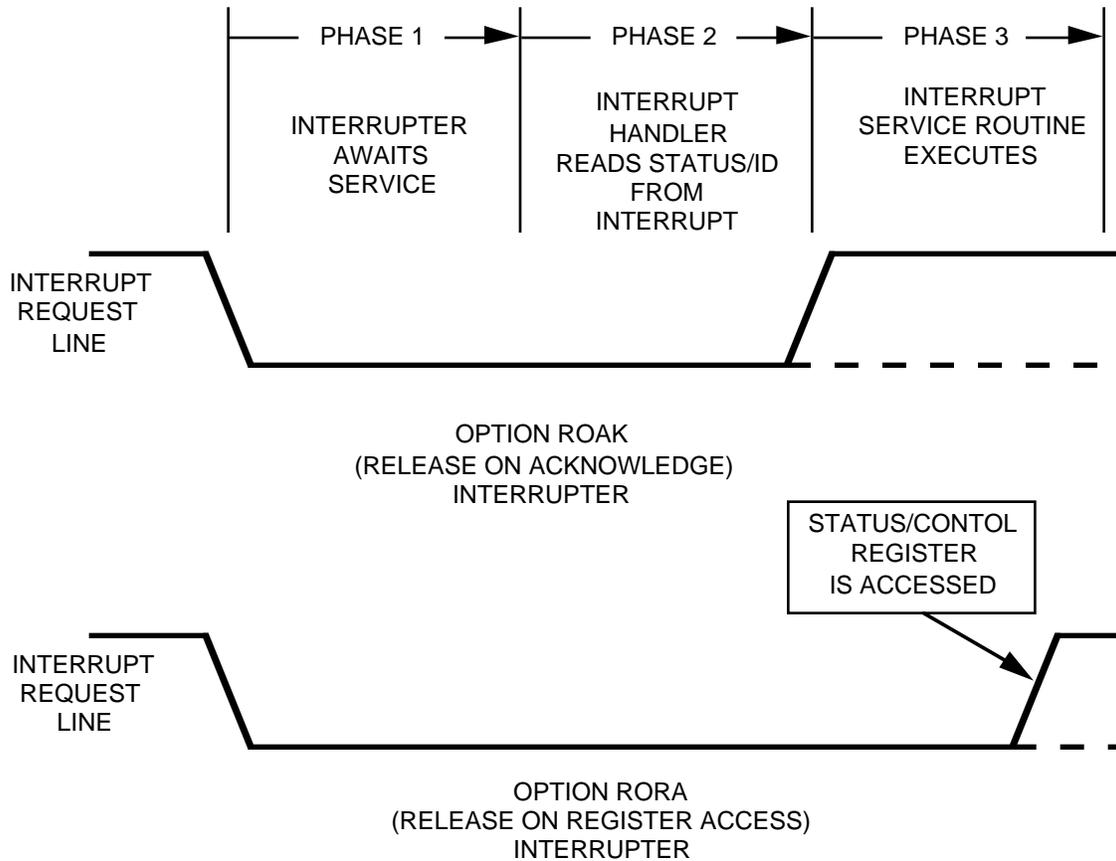


Figure 4 - 8 Release Of Interrupt Request Lines By ROAK And RORA Interrupters

4.3.8 Interaction Between Priority Interrupt Bus Modules

In the following discussions, several on-board signals are defined to describe the interaction between the Interrupter and Interrupt Handler modules and other on-board logic. These signals are only intended to illustrate the information which is passed to and from the modules, rather than to define their designs.

PERMISSION 4.3:

VMEbus boards MAY be designed with on-board signals that differ from those used in the following discussions.

Figure 4-4 shows how the IACKIN*/IACKOUT* daisy-chain is routed through a typical configuration of boards on the VMEbus. The IACK* line runs the full length of the backplane and can be driven by any Interrupt Handler that has control of the DTB. The backplane connects IACK* to the IACKIN* pin of Slot 1. The IACK Daisy-Chain Driver resides in Slot 1 and monitors the level of Slot 1's IACKIN* line.

When an Interrupt Handler drives IACK* (and Slot 1's IACKIN*) low, and then drives DSA* low, the IACK Daisy-Chain Driver generates a falling edge on its IACKOUT* pin. This pin is connected to the IACKIN* pin of Slot 2. A jumper on the board in Slot 2 routes the falling edge on the IACKIN* pin to the IACKOUT* pin, and through the backplane to the IACKIN* pin of the board in Slot 3. The Interrupter in Slot 3 does not have a pending interrupt request, so it passes on the falling edge to its IACKOUT* pin. The Interrupter in Slot 4 then detects the falling edge on its IACKIN* line and responds by placing its Status/Id on the data bus, and then driving DTACK* low.

PERMISSION 4.4:

An Interrupter MAY reside on the system controller board, installed in Slot 1, along with the IACK Daisy-Chain Driver. Figure 4-9 shows how the two modules would be connected.

PERMISSION 4.5:

More than one Interrupter MAY reside on a board. Figure 4-10 shows how this might be done.

OBSERVATION 4.8:

In some cases, board designers might not know whether or not the board they are designing will be installed in Slot 1, or in some other slot of a VMEbus system.

RECOMMENDATION 4.1:

IF a board includes both an IACK Daisy-Chain Driver and an Interrupter,
and might or might not be installed in Slot 1.

THEN design it as shown in Figure 4-9.

PERMISSION 4.6:

Several boards containing IACK Daisy-Chain Drivers MAY be installed in a VMEbus system.

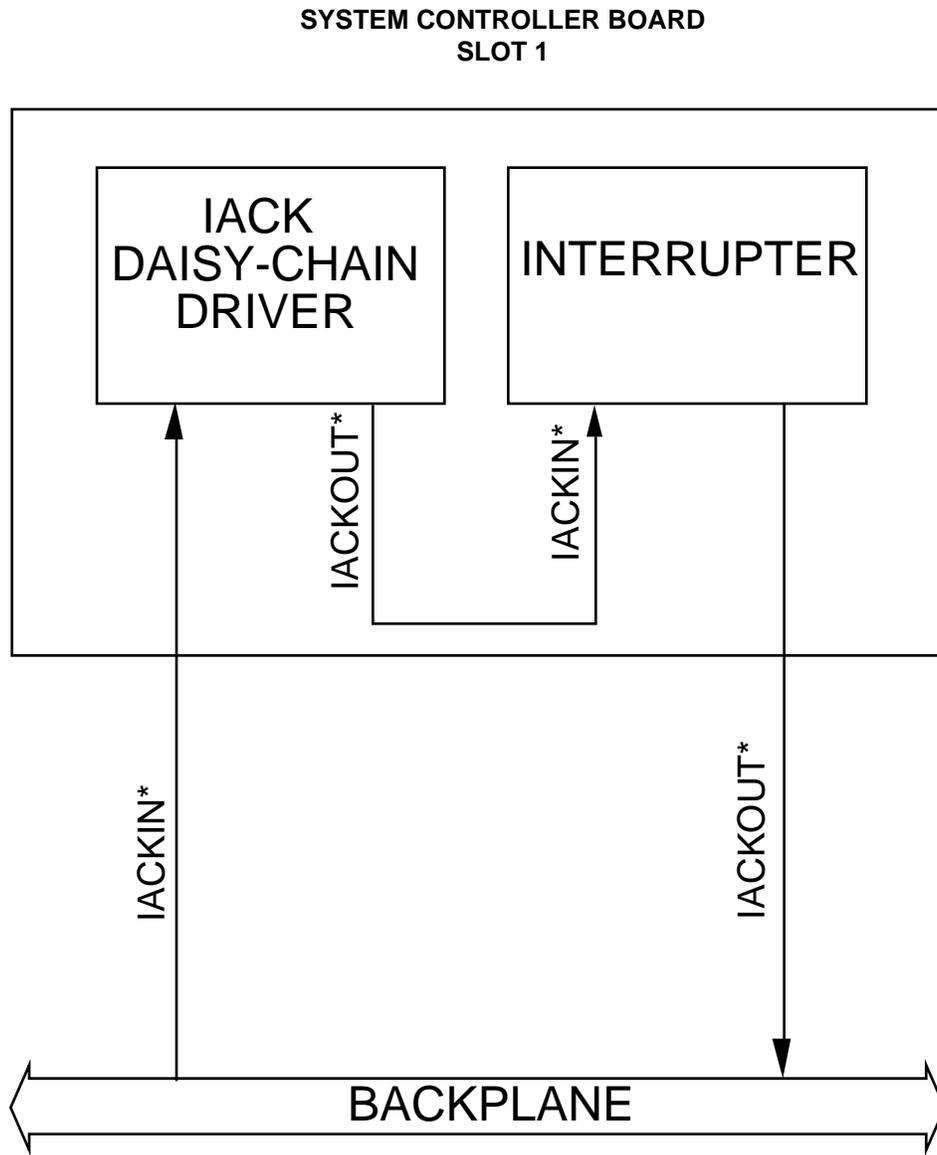


Figure 4 - 9 An IACK Daisy-Chain Driver And Interrupter On The Same Board

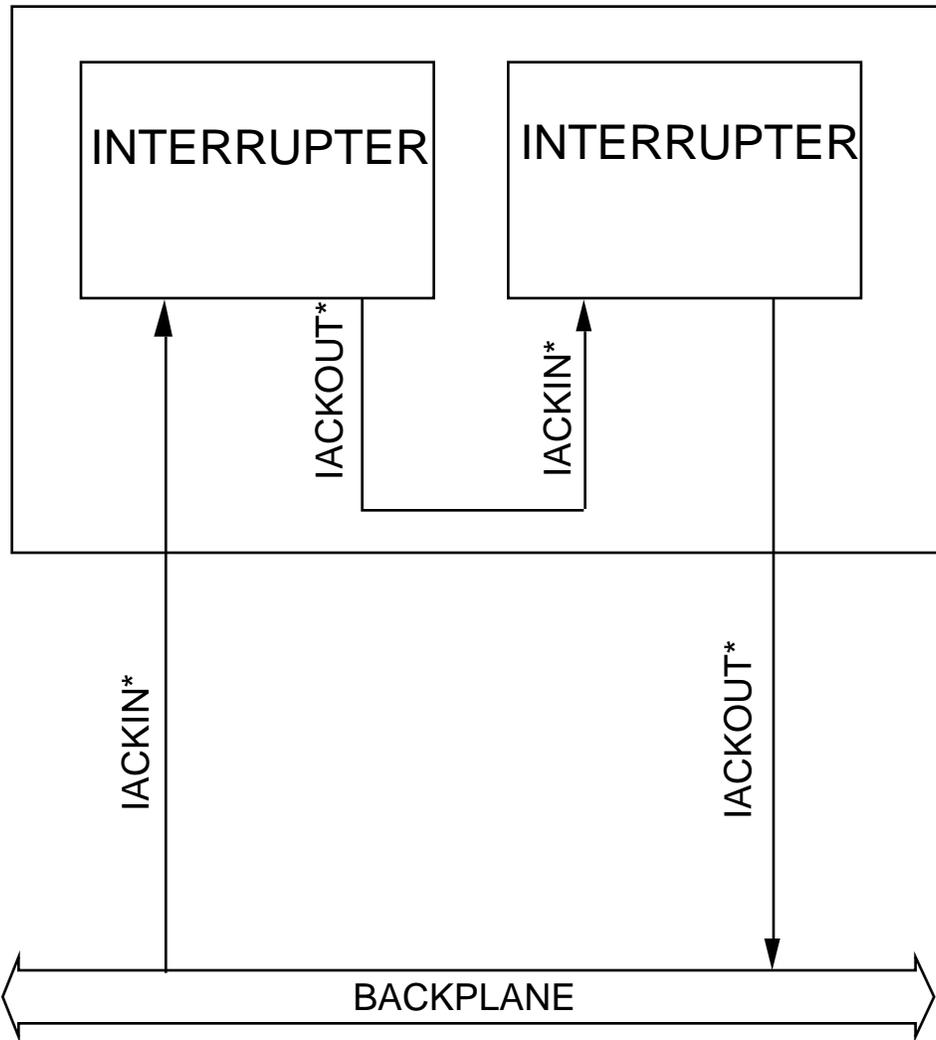


Figure 4 - 10 Two Interrupters On The Same Board

4.4 Typical Operation

A typical interrupt sequence can be divided into three phases:

- Phase 1: The interrupt request phase.
- Phase 2: The interrupt acknowledge phase.
- Phase 3: The interrupt servicing phase.

Figure 4-11 illustrates the timing relationships between the three phases.

Phase 1 starts when an Interrupter drives an interrupt request line low and ends when the Interrupt Handler gains control of the DTB. During phase 2 the Interrupt Handler uses the DTB to read the Interrupter's Status/Id. During phase 3 an interrupt servicing routine is executed. (This might or might not involve data transfers on the VMEbus.)

The protocol for the interrupt subsystem describes the module interaction required during phase 1 and phase 2. Any data transfers which take place during phase 3 will follow the Data Transfer Bus protocol described in Chapter 2.

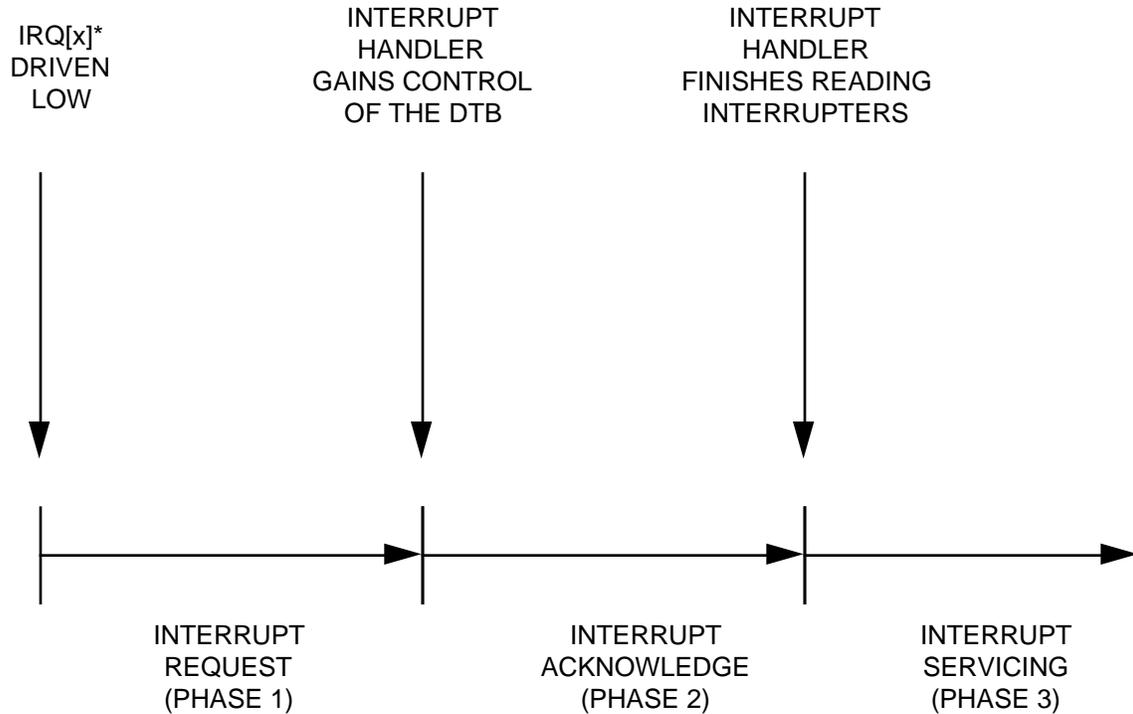


Figure 4 - 11 The Three Phases Of An Interrupt Sequence

4.4.1 Single Handler Interrupt Operation

In single handler interrupt systems, the seven interrupt request lines are all monitored by a single Interrupt Handler. The interrupt request lines are prioritized such that $IRQ7^*$ has the highest priority and $IRQ1^*$ has the lowest priority. When the Interrupt Handler detects simultaneous requests on two interrupt request lines, it acknowledges the highest priority request first.

4.4.2 Distributed Interrupt Operation

Distributed interrupt systems contain from two to seven Interrupt Handlers. For purposes of the following discussion, distributed interrupt systems will be considered in two groups:

- a. distributed interrupt systems with seven Interrupt Handlers,
- b. distributed interrupt systems with two to six Interrupt Handlers.

4.4.2.1 Distributed Interrupt Systems With Seven Interrupt Handlers

In distributed interrupt systems with seven Interrupt Handlers, each of the interrupt request lines is monitored by a separate Interrupt Handler. Each Interrupt Handler gains control of the DTB before it reads the Status/Id from Interrupters driving its interrupt request line.

OBSERVATION 4.9:

There is no specified relationship between the interrupt request line that an Interrupt Handler services and the bus request line used by its on-board Requester. For example, an Interrupt Handler that services IRQ7* might have a Requester that uses BR0*, and an Interrupt Handler that services IRQ1* might have a Requester that uses BR3*. It is clear from this that there is no implied interrupt priority between lines serviced by different Interrupt Handlers.

Figure 4-12 illustrates a distributed interrupt system where Interrupt Handler A monitors IRQ2* and has an on-board Requester which requests the DTB on BR2*. Interrupt Handler B monitors IRQ5* and has an on-board Requester which requests the DTB on BR3*. Two Interrupters simultaneously drive IRQ2* and IRQ5* low, and the two Interrupt Handlers cause their on-board Requesters to drive BR2* and BR3* low simultaneously. In this example, priority arbitration is used and, since both bus requests go low together, the Arbiter first grants control of the DTB to Interrupt Handler B's Requester, and Interrupt Handler A waits until B has finished using the DTB.

OBSERVATION 4.10:

If round-robin arbitration is used, either of the Interrupt Handlers described in Figure 4-12 might be granted the bus first.

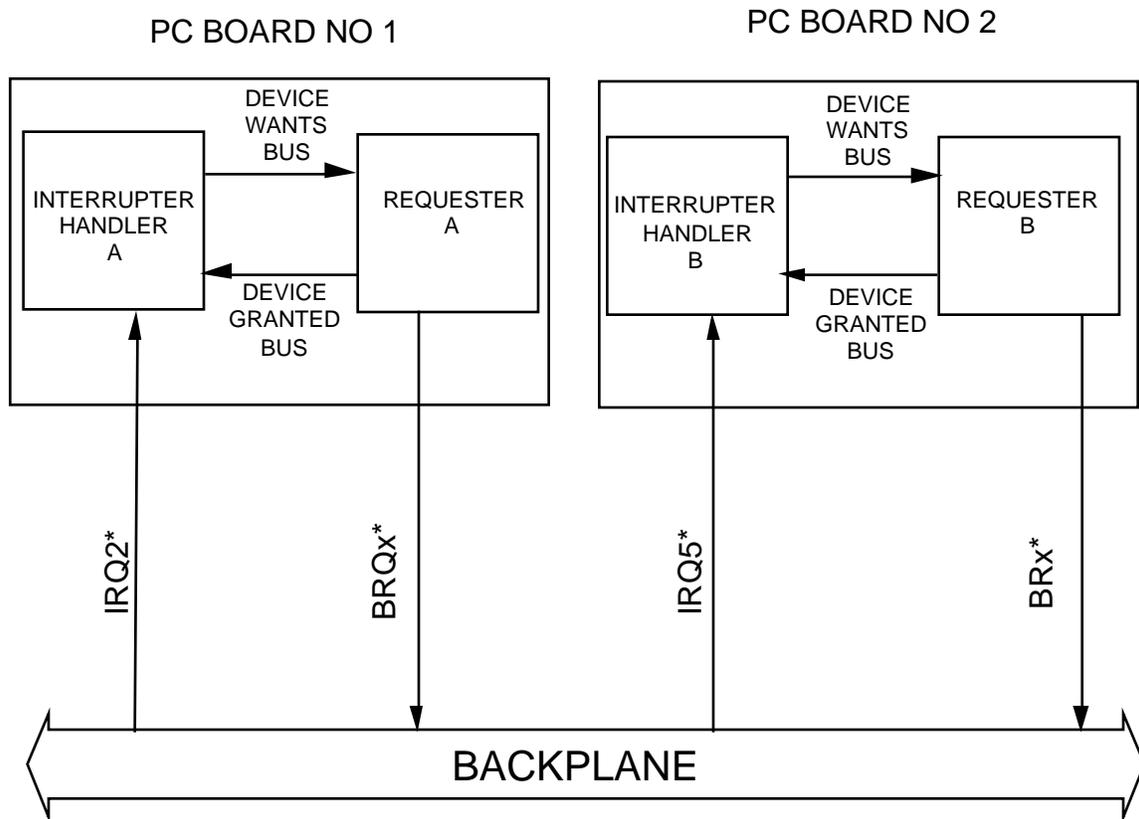


Figure 4 - 12 Two Interrupt Handlers, Each Monitoring One Interrupt Request Line

4.4.2.2 Distributed Interrupt Systems With Two To Six Interrupt Handlers.

It is also possible to configure a distributed interrupt system in which two or more of the interrupt request lines are monitored by a single Interrupt Handler. Figure 4-13 illustrates a system configured with two Interrupt Handlers in which Interrupt Handler A monitors IRQ[4..1]*, and Interrupt Handler B monitors IRQ[7..5]*. In this case, the IRQ[4..1]*, lines are prioritized; IRQ4* = highest priority for Interrupt Handler A, and the IRQ[7..5]*, lines are prioritized; IRQ7* = highest priority for Interrupt Handler B. The DTB arbitration still determines which Interrupt Handler is granted the use of the DTB first.

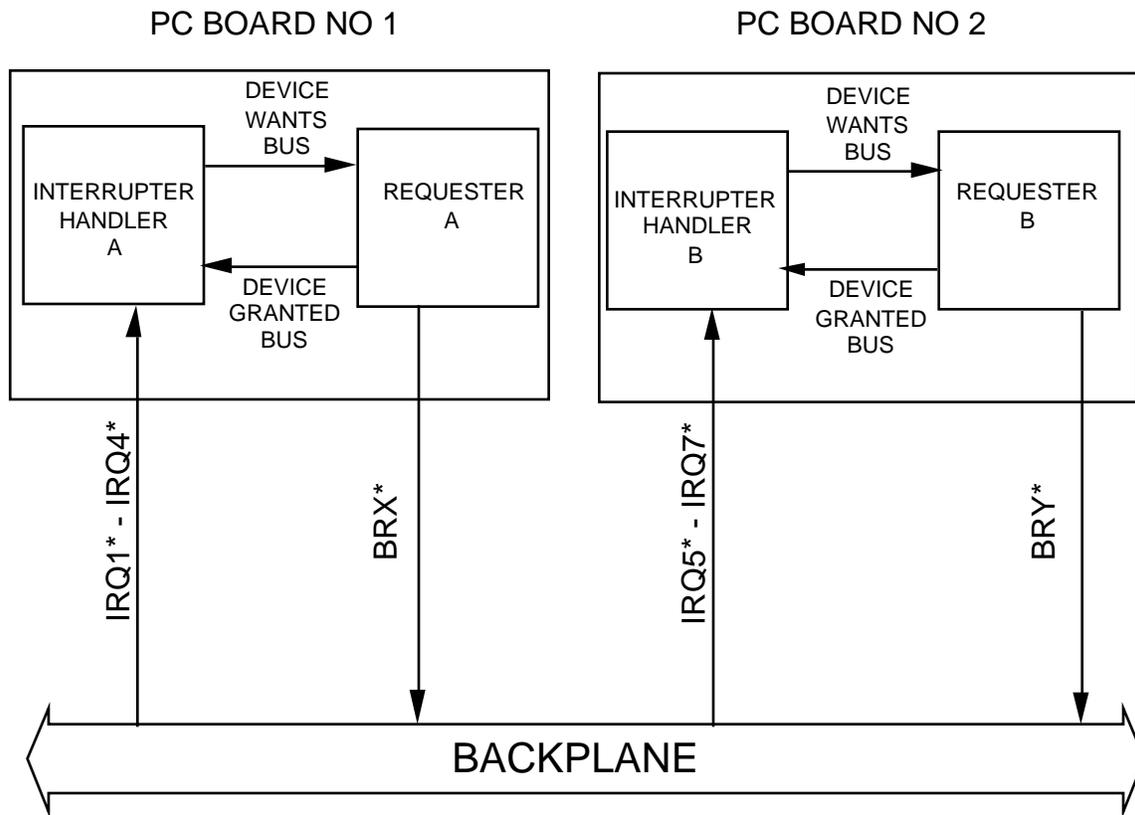


Figure 4 - 13 Two Interrupt Handlers, Each Monitoring Several Interrupt Request Lines

4.4.3 Example: Typical Single Handler Interrupt System Operation

Figure 4-14 illustrates the operation of a single handler interrupt system in which one Interrupt Handler monitors and prioritizes all seven interrupt lines. At the top of the diagram, a Master is using the DTB to move data within the system at a bus request level of 2. An Interrupter in Slot 3 requests an interrupt by driving IRQ4* low. When the Interrupt Handler detects the low level on IRQ4* it sends a signal to its on-board Requester, indicating that it needs the bus. This Requester then drives BR3* low. Upon detecting the bus request, the Arbiter drives BCLR* low, indicating that a higher priority Requester is waiting for the DTB. (This example assumes a PRI Arbiter). When Master A detects the low level on BCLR*, it stops moving data and allows its requester to relinquish control of the DTB, and release BBSY*.

OBSERVATION 4.11:

The active Master is not required to relinquish the DTB within any specified time, but a prompt response to the BCLR* line allows the interrupt to be serviced quicker.

When the Arbiter detects BBSY* high, it grants the DTB to Requester B, which informs its Interrupt Handler that the DTB is available (see Figures 2-26 and 2-27). The Interrupt Handler then puts a 3-bit code on the lower three address lines, indicating that it is acknowledging the interrupt request on the IRQ4* line (see Table 4-7). At the same time, it drives IACK* low, indicating that it is acknowledging an interrupt, and drives AS* low. The low level on IACK* is connected, by a signal trace in the backplane, to the IACKIN* pin of Slot 1 and causes the IACK Daisy-Chain Driver to generate a falling edge down the IACKIN*/IACKOUT* daisy-chain.

When the Interrupter detects a falling edge on its incoming daisy-chain line (IACKIN*) it checks the lower three address bits to see if they match the interrupt request line which it is driving low. Since the 3-bit code matches the line on which it is making its interrupt request, when the Interrupter detects the data strobe(s) low, it places its Status/Id on the data bus and drives the DTACK* line low. When the Interrupt Handler detects the low DTACK*, it reads the Status/Id and activates the appropriate interrupt service routine.

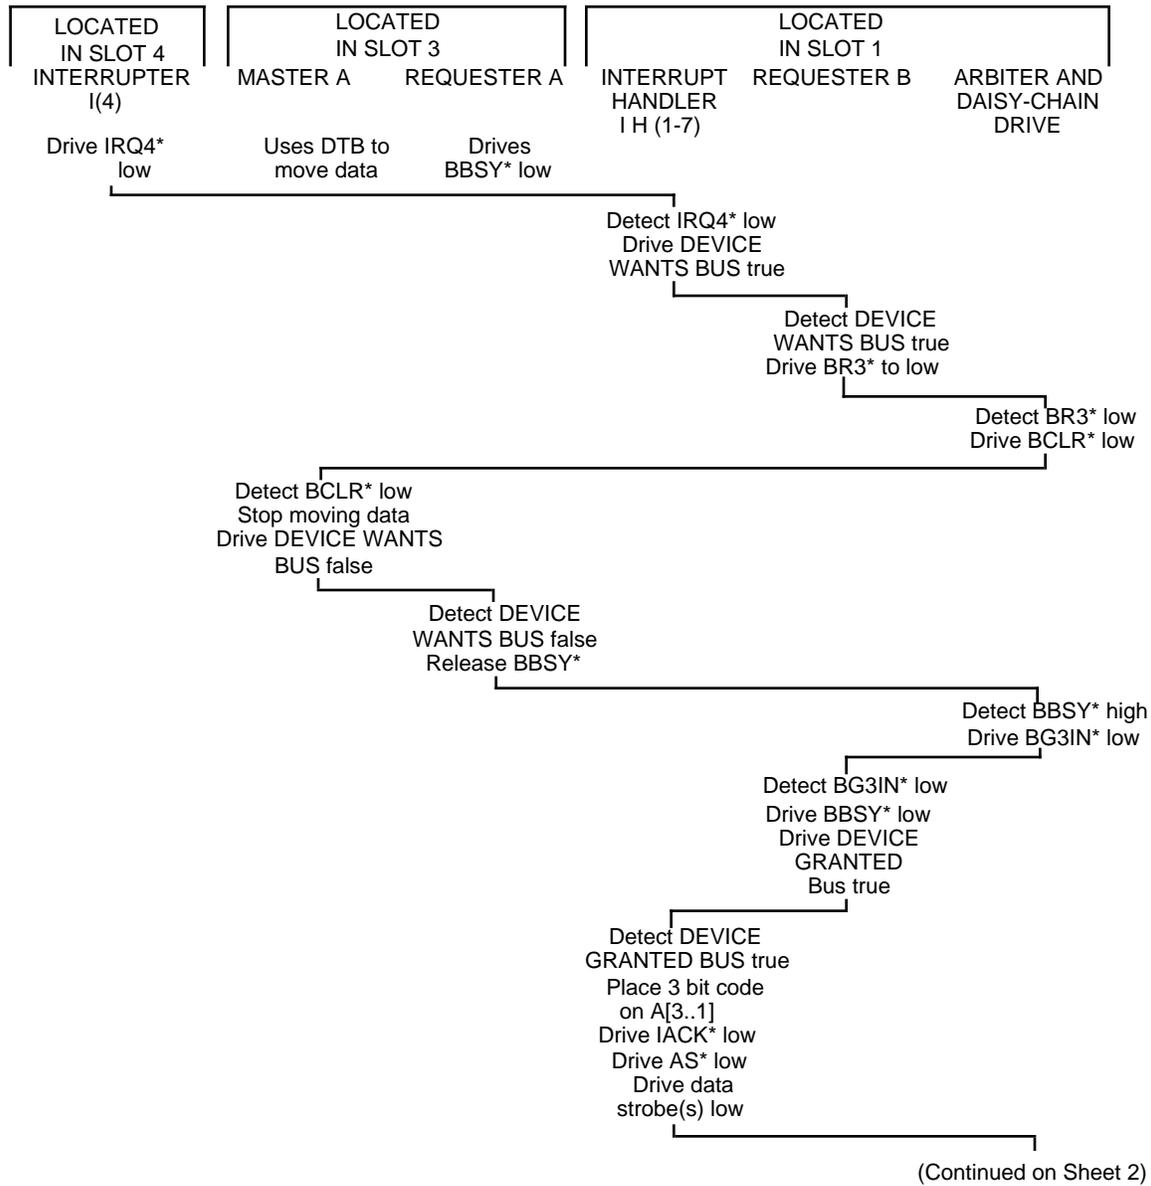


Figure 4 - 14 Typical Single Handler Interrupt System Operation Flow Diagram (Sheet 1 of 2)

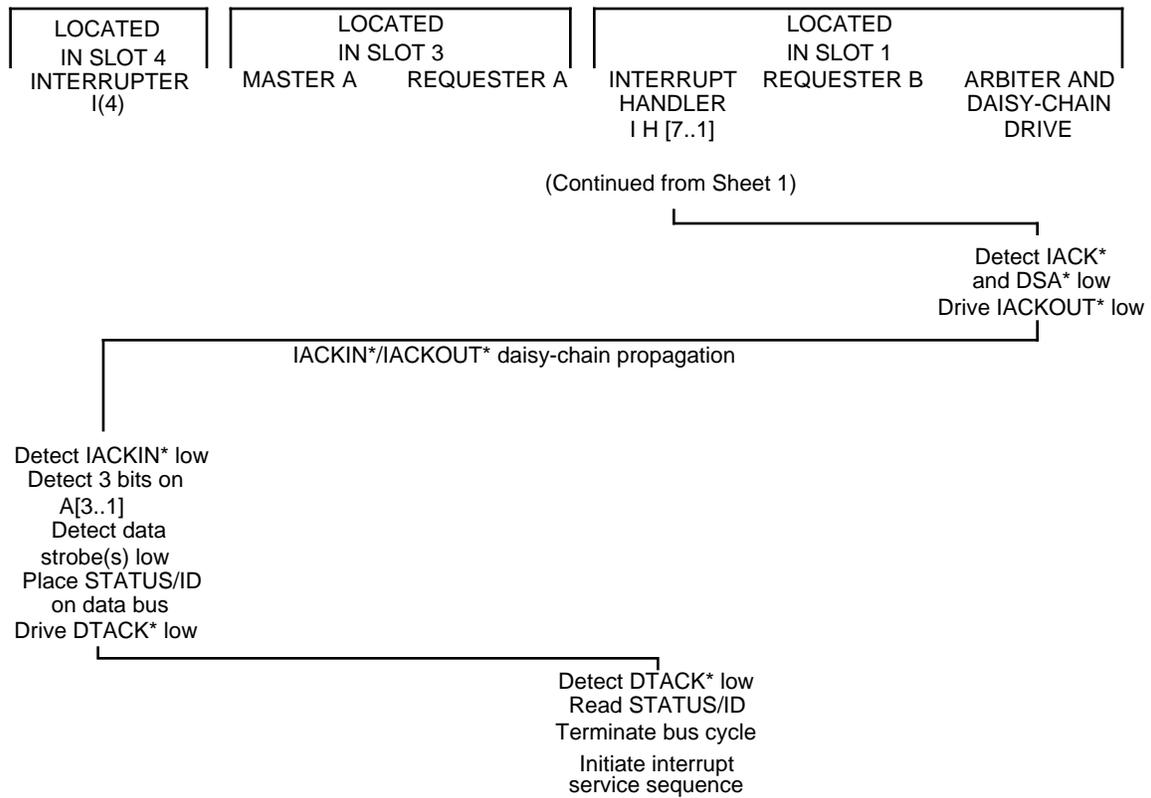


Figure 4 - 14 Typical Single Handler Interrupt System Operation Flow Diagram (Sheet 2 of 2)

Table 4 - 7 3-Bit Interrupt Acknowledge Code

Interrupt line being acknowledged	Use of the address bus to broadcast the 3-bit interrupt acknowledge code		
	A3	A2	A1
IRQ1*	L	L	H
IRQ2*	L	H	L
IRQ3*	L	H	H
IRQ4*	H	L	L
IRQ5*	H	L	H
IRQ6*	H	H	L
IRQ7*	H	H	H

H = High level

L = Low level

4.4.4 Example: Prioritization of Two Interrupts In A Distributed Interrupt System

Figure 4-15 illustrates the operation of a distributed interrupt system with two Interrupt Handlers. Interrupt Handler A monitors IRQ[4..1]*, while Interrupt Handler B monitors IRQ[7..5]*. Interrupt Handler A treats IRQ4* as its highest priority interrupt, while Interrupt Handler B treats IRQ7* as its highest priority interrupt. At the top of the diagram, Interrupter C drives IRQ3* low, and Interrupter D drives IRQ6* low. Both Interrupt Handlers detect their respective interrupt request lines low, and both simultaneously indicate to their on-board Requester that they need the DTB. Both the Requesters drive BR3* low. Upon detecting BR3* low, the DTB Arbiter drives BG3IN* low on Slot 1. This low signal is passed down the BG3IN*/BG3OUT* daisy-chain until it is detected by the Requester B in Slot 4. This Requester then signals its on-board Interrupt Handler B that the DTB is available. Interrupt Handler B then reads the Status/Id from Interrupter D.

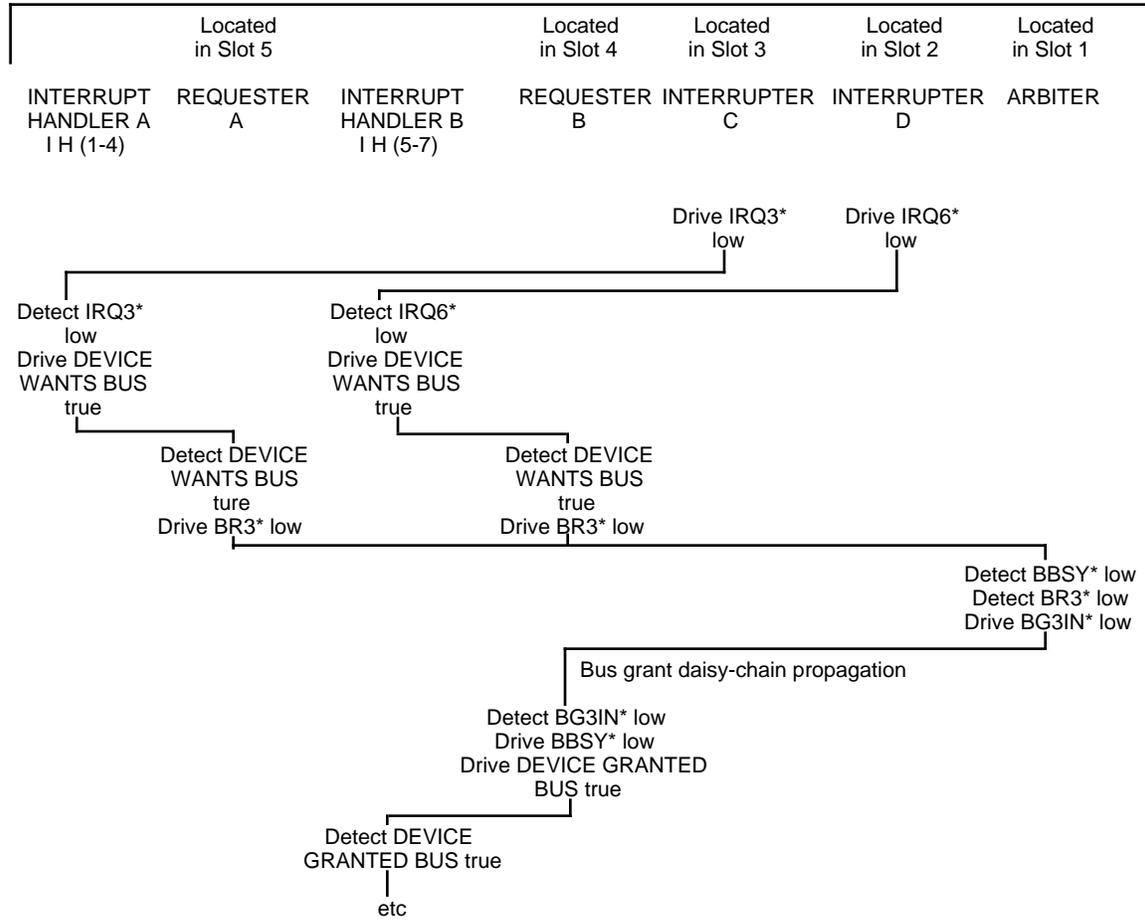


Figure 4 - 15 Typical Distributed Interrupt System With Two Interrupt Handlers, Flow Diagram

4.5 Race Conditions

Consider the case of two interrupters A and B. Interrupter B is down the interrupt acknowledge daisy-chain from A. Interrupter B requests an interrupt. After the relevant interrupt handler is granted the bus, it acknowledges the interrupt request by driving the IACK* line low. The resulting low going edge on the interrupt acknowledge daisy-chain arrives at interrupter A at about the time it drives its own interrupt request line low. If interrupter A is improperly designed, this situation might cause it to momentarily drive its IACKOUT* line low, and then high again. This would result in a low going transient on the interrupt acknowledge daisy-chain.

RULE 4.49

Interrupters **MUST** be designed to ensure that no momentary low-going transients are generated on their IACKOUT* line.

OBSERVATION 4.50

If the interrupter is designed so that it latches the state of an on-board generated interrupt request line upon the falling edge of its IACKIN* line, and if that signal is in transition when the falling edge occurs, the outputs of the latch will sometimes oscillate, or remain in the threshold region between the high and low levels for a short time. Because of this, no time limit is set for the interrupter to pass along the interrupt acknowledge. It is only prohibited from generating a low-going transient on its IACKOUT* line, which might be interpreted as an acknowledge by an interrupter further down the daisy chain.

PERMISSION 4.11

If an interrupter is about to drive an interrupt request line low between the time it receives an interrupt acknowledge intended for another interrupter and the time it would pass the interrupt acknowledge on, then it **MAY** treat the interrupt acknowledge as its own. In this case, the other interrupter maintains its interrupt request until another acknowledge is issued.

4.6 Priority Interrupt Bus Timing Rules and Observations

This section describes the timing RULES and OBSERVATIONS that govern the behavior of Interrupt Handlers, Interrupters, and IACK Daisy-Chain Driver during the selection of the responding Interrupter (i.e. the Interrupter that is to provide its Status/Id in response to the Interrupt Acknowledge Cycle). This timing information is in the form of figures and tables.

The Interrupt Acknowledge Cycle begins with the selection of the responding Interrupter. This is called the Interrupter selection portion of the cycle. Once an Interrupter responds, the Interrupt Handler reads the Status/Id from it. This is called the Status/Id transfer portion of the cycle.

When the Interrupt Handler initiates an interrupt acknowledge cycle, there might be several Interrupters between it and the Interrupter being acknowledged that either:

- a. do not have an interrupt pending,
- b. have an interrupt pending on a different interrupt request line than the one being acknowledged.

Although these Interrupters do not respond with a Status/Id, they do participate in the Interrupt Acknowledge Cycle by passing the falling edge from their IACKIN* line on to their IACKOUT* line. For this reason, these Interrupters are called participating Interrupters.

The first Interrupter in the daisy-chain that has an interrupt pending on the interrupt request line that is being acknowledged responds with a Status/Id. For this reason it is called the responding Interrupter. All other Interrupters are called non-participating Interrupters.

Table 4-8 lists timing tables and timing diagrams that specify Interrupt Handler and Interrupter operation.

Table 4-9 lists timing tables and timing diagrams that specify IACK Daisy-Chain Driver operation.

Table 4-10 lists timing tables and timing diagrams that specify participating Interrupter operation.

Table 4-11 lists timing tables and timing diagrams that specify responding Interrupter operation.

Table 4-12 defines the mnemonics that are used in Tables 4-13 through 4-15.

Tables 4-13 through 4-15 specify the use of bus signal lines by the Priority Interrupt Bus functional modules.

Tables 4-16 through 4-19 specify the timing parameters for the Priority Interrupt Bus functional modules. (The reference numbers used in Tables 4-17 through 4-19 correspond to the timing parameter numbers in Table 4-16.)

Figures 4-16 through 4-23 are timing diagrams that specify the timing during interrupt acknowledge cycles.

Figure 4-24 specifies additional intercycle timing for the IACKIN*/IACKOUT* daisy-chain.

All of the RULES and OBSERVATIONS associated with the figures listed below also apply to Interrupt Handlers, Interrupters, and the IACK Daisy-Chain Driver.

Figures 2-22 through 2-24 in Chapter 2 specify the timing for the address and data strobes between data transfer cycles.

Figure 2-25 specifies the timing of a timed-out cycle.

Figures 2-26 and 2-27 specify the timing during mastership transfer.

In order to meet the specified timing, board designers have to take into account the worst case propagation delays of the bus drivers and receivers used on their VMEbus boards. The propagation delay of the drivers depends on their output loads. However, manufacturer specifications do not always give enough information to calculate the propagation delays under various loads. To help the VMEbus board designer, some suggestions are offered in Chapter 6.

The OBSERVATIONS specify the timing of signal transitions on incoming lines. These times can be relied upon as long as the backplane loading RULEs in Chapter 6 are not violated. The RULEs for the bus terminators in Chapter 6 guarantee that the timing parameters for signal lines that are released after they have been driven, are met.

Typically, for each timing RULE there is a corresponding timing OBSERVATION. However, the time that is guaranteed in the OBSERVATION might differ from the time specified by the RULE. For example, a careful inspection of the timing diagrams shows that the Interrupt Handler is required to provide 35 nanoseconds of address set-up time, but the Interrupter is only guaranteed 10 nanoseconds. This is because the address drivers are not always able to drive the backplane's signal lines completely through the low to high threshold region, until the transition propagates to the end of the backplane and is reflected back. The falling edge of the address strobe, however, typically crosses the 0.8-volt threshold without waiting for a reflection. Therefore, the resulting set-up time at the Interrupter is the Interrupt Handler's set-up time less two bus propagation times.

A special notation has been used to describe the data strobe timing. The two data strobes (DS0* and DS1*) do not always make their transitions simultaneously. For purposes of these timing diagrams, DSA* represents the first data strobe to make its transition (whether that is DS0* or DS1*). DSB* represents the second data strobe to make its transition (whether that is DS1* or DS0*). The broken line shown while the data strobes are stable indicates that the first data strobe to make a falling transition might not be the first to make its rising transition, i.e., DSA* might represent DS0* on its falling edge and DS1* on its rising edge.

Table 4 - 8 Timing Diagrams That Define Interrupt Handler And Interrupter Operation

Mnemonic	Type of cycle	Interrupter Selection Timing Diag Figure	Status/Id Transfer Timing Diag Figure
D08(O)	Single Byte Status/Id Read	2-12 & 4-16	4-20
D16	Double Byte Status/Id Read	2-12 & 4-16	4-21
D32	Quad Byte Status/Id Read	2-12 & 4-16	4-21

Note: See Tables 4-16 and 4-17 for timing parameters

Table 4 - 9 Timing Diagrams That Define IACK Daisy-Chain Driver Operation

Type of cycle	Interrupter Selection Timing Diagram
Single Byte Status/Id Read	Figure 4-17
Double Byte Status/Id Read	Figure 4-17
Quad Byte Status/Id Read	Figure 4-17

Note: See Tables 4-16 and 4-19 for timing values

Table 4 - 10 Timing Diagrams That Define Participating Interrupter Operation

Type of cycle	Interrupter Selection Timing Diagram
Single Byte Status/Id Read	Figure 4-18
Double Byte Status/Id Read	Figure 4-18
Quad Byte Status/Id Read	Figure 4-18

Note: See Tables 4-16 and 4-18 for timing values

Table 4 - 11 Timing Diagrams That Define Responding Interrupter Operation

Mnemonic	Interrupter Type of cycle Timing Diag	Status/Id Selection Timing Diag	Transfer
D08(O)	Single Byte Status/Id Read	Figure 4-19	Figure 4-22
D16	Double Byte Status/Id Read	Figure 4-19	Figure 4-23
D32	Quad Byte Status/Id Read	Figure 4-19	Figure 4-23

Notes : See Tables 4-16 and 4-18 for timing parameters

Table 4 - 12 Definitions Of Mnemonics Used In Tables 4-13, 4-14, And 4-15

Mnemonic	Description	Comments
DVBIH	Driven Valid by Interrupt Handler	RULE 4.9: The Interrupt Handler MUST drive DVBIH lines to a valid level.
DLBIH	Driven Low by Interrupt Handler	RULE 4.10: The Interrupt Handler MUST drive DLBIH lines to a low level.
DHBIH	Driven High by Interrupt Handler	RULE 4.11: The Interrupt Handler MUST drive DHBIH lines to a high level.
dhbih?	Driven High by Interrupt Handler?	PERMISSION 4.7: The Interrupt Handler MAY drive dhbih? lines high. RULE 4.12: The Interrupt Handler MUST NOT drive dhbih? lines low during the cycle.
dxbih?	Driven by Interrupt Handler?	PERMISSION 4.8: The Interrupt Handler MAY drive dxbih? lines during the cycle, or it MAY leave dxbih? lines undriven. (When the line is driven it carries no valid information.)
DVBI	Driven Valid by Interrupter	RULE 4.13: The Interrupter MUST drive DVBI lines a to valid level.
dhbi?	Driven by Interrupter?	PERMISSION 4.9: The Interrupter MAY drive dhbi? lines high. RULE 4.14: The Interrupter MUST NOT drive dhbi? lines low.
dxbi?	Driven by Interrupter?	PERMISSION 4.10: The Interrupter MAY drive dxbi? lines during the cycle, or it MAY leave the line undriven. (When the line is driven, it carries no valid information.)

Table 4 - 13 Use Of Addressing Lines During Interrupt Acknowledge Cycles

Interrupt line being acknowledged	A03	A02	A01	IACK*
IRQ1*	DLBIH	DLBIH	DHBIH	DLBIH
IRQ2*	DLBIH	DHBIH	DLBIH	DLBIH
IRQ3*	DLBIH	DHBIH	DHBIH	DLBIH
IRQ4*	DHBIH	DLBIH	DLBIH	DLBIH
IRQ5*	DHBIH	DLBIH	DHBIH	DLBIH
IRQ6*	DHBIH	DHBIH	DLBIH	DLBIH
IRQ7*	DHBIH	DHBIH	DHBIH	DLBIH

Table 4 - 14 Use Of The DS1*, DS0*, LWORD*, And WRITE* Lines During Interrupt Acknowledge Cycles

Mnemonic	Type of cycles	DS1*	DS0*	LWORD*	WRITE*
D08(O)	Single byte interrupt acknowledge	dhbih?	DLBIH	dhbih?	dhbih?
D16	Double byte interrupt acknowledge	DLBIH	DLBIH	dhbih?	dhbih?
D32	Quad byte interrupt acknowledge	DLBIH	DLBIH	DLBIH	dhbih?

Table 4 - 15 Use of the data bus lines to transfer the Status/Id

Mnemonic	Type of cycles	D31-24	D23-16	D15-08	D07-00
D08(O)	Single, double, and quad byte interrupt acknowledge	dhbi?	dhbi?	dhbi?	DVBI
D16	Double and quad byte interrupt acknowledge	dhbi?	dhbi?	DVBI	DVBI
D32	Quad byte interrupt acknowledge	DVBI	DVBI	DVBI	DVBI

Table 4 - 16 Interrupt Handler, Interrupter, And IACK DAISY-CHAIN DRIVER Timing Parameters

Parameter Number	Interrupt Handler See Table 4-17		Interrupter See Table 4-18		IACK Daisy-Chain Driver See Table 4-19	
	MIN	MAX	MIN	MAX	MIN	MAX
1	0					
2	0					
3	60					
4	35		10			
5	40		30		30	
6			0			
7			0			
9	0		0			
10	0		-10			
11	40		30			
12	35		10			
13		10		20		
14	0		0			
16	0		0			
18	0		0			
19	40		30		30	
20	0		0			
21	0		0			
23	10		0			
24A	0					
24B	0					
25		25				
26	0		0			
27	-25		0			
28	30	2T	30			
29	0		0			
30	0		0			
31	0		0			
32			10		10	
34					40	
34A			30			
35				30	0	30
36			0			
37			0			
38A			0			
38B			0			
39				40		
40			30		30	
41			0			
42					30	
43			30			

- Notes
1. All times are in nanoseconds.
 2. T = the bus time-out value.
 3. Some entries in this table are referred to in the following tables only and are not used in the timing diagrams.

Table 4 - 17 Interrupt Handler, Timing RULEs And OBSERVATIONS

Note: The numbers correspond to the timing parameters specified in Table 4-16 .

1. RULE 4.15:
When taking control of the VMEbus, the Interrupt Handler MUST NOT drive any of IACK*, A[3..1], LWORD*, WRITE*, DS0*, DS1* or AS* until after the previous Master or Interrupt Handler allows AS* to rise above the low level.

OBSERVATION 4.12:
Chapter 3 describes how an Interrupt Handler's Requester is granted use of the DTB.
2. RULE 4.16:
When taking control of the DTB, the Interrupt Handler MUST NOT drive any of IACK*, A[3..1], LWORD*, WRITE*, DS0*, DS1*, or AS* until after it receives DEVICE GRANTED BUS true.

OBSERVATION 4.13:
Chapter 3 describes how an Interrupt Handler's Requester is granted use of the DTB.
3. RULE 4.17:
When taking control of the DTB, the Interrupt Handler MUST NOT drive AS* low until this time after the previous Master or Interrupt Handler allows AS* to rise above the low level.

OBSERVATION 4.14:
RULE 4.17 ensures that timing parameter 5 for Interrupters and Slaves is guaranteed when there is an interchange of the DTB mastership.
4. RULE 4.18:
The Interrupt Handler MUST NOT drive AS* low until IACK* has been low, and LWORD* and A[3..1] have been valid for this minimum time.
5. RULE 4.19:
When using the DTB for two consecutive cycles, the Interrupt Handler MUST NOT drive AS* low until it has been high for this minimum time.
9. RULE 4.20:
The Interrupt Handler MUST NOT drive DSA* low until both DTACK* and BERR* are high.
10. RULE 4.21:
The Interrupt Handler MUST NOT drive DSA* low before it has driven AS* low.

Table 4 - 17 (Continued)

11. RULE 4.22:
The Interrupt Handler MUST NOT drive DSA* low until DS0* and DS1* have been simultaneously high for this minimum time.
12. RULE 4.23:
The Interrupt Handler MUST NOT drive DSA* low until WRITE* has been high for this minimum time.
13. RULE 4.24:
During double byte or quad byte Interrupt Acknowledge read cycles, the Interrupt Handler MUST drive DSB* low within this maximum time after it drives DSA* low.

OBSERVATION 4.15:
Timing parameter 13 does not apply to single byte Interrupt Acknowledge reads.
14. RULE 4.25:
During all interrupt acknowledge cycles, the Interrupt Handler MUST hold the 3-bit interrupt acknowledge code valid and maintain the appropriate level on LWORD* until it detects a falling edge on DTACK* or BERR*.
16. RULE 4.26:
During all interrupt acknowledge cycles, the Interrupt Handler MUST maintain IACK* low until it detects a falling edge on DTACK* or BERR*.
18. RULE 4.27:
The Interrupt Handler MUST hold AS* low until it detects DTACK* or BERR* low.
19. RULE 4.28:
The Interrupt Handler MUST hold AS* low for this minimum time.
20. RULE 4.29:
Once an Interrupt Handler has driven DSA* low, it MUST maintain it low until it detects DTACK* or BERR* low.

OBSERVATION 4.52:
During a read cycle, data may not be valid at the bus Interrupt Handler for up to 25ns after DTACK* is detected low. (See parameter 27). To ensure that valid data is read, the bus Interrupt Handler should maintain DSA* asserted for a least 25ns after detecting DTACK* low.
21. RULE 4.30:
Once an Interrupt Handler has driven DSB* low, it MUST maintain it low until it detects DTACK* or BERR* low.
See OBSERVATION 4.52

Table 4 - 17 (Continued)

23. RULE 4.31:
Once an Interrupt Handler has driven DSA* low, it MUST maintain a high on the WRITE* line, until this minimum time after both data strobes are high.
- 24A. RULE 4.32:
IF the Interrupt Handler drives or releases AS* to high after its Requester releases BBSY*,
THEN it MUST release IACK*, A[3..1], LWORD*, WRITE*, DS0*, and DS1* before allowing AS* to rise above the low level.
- 24B. RULE 4.33:
IF the Interrupt Handler drives or releases AS* to high before its Requester releases BBSY*,
THEN it MUST release IACK*, A[3..1], LWORD*, WRITE*, DS0*, and DS1* before changing its DEVICE WANTS BUS signal from true to false.
25. RULE 4.34:
IF the Interrupt Handler drives or releases AS* to high after its Requester releases BBSY*,
THEN it MUST release AS* within this time after allowing it to rise above the low level.
26. OBSERVATION 4.16:
Timing parameter 26 guarantees that the data bus will not be driven until the Interrupt Handler drives DSA* low.
27. OBSERVATION 4.17:
The Interrupt Handler is guaranteed that the data bus will be valid within this time after DTACK* goes low. This time does not apply to cycles where the Interrupter drives BERR* low instead of DTACK*.
28. OBSERVATION 4.18:
The Interrupt Handler is guaranteed that neither DTACK* nor BERR* will go low until this minimum time after it drives DSA* low. The Bus Timer guarantees the Interrupt Handler that if DTACK* has not gone low after its time-out period has elapsed, and within twice its time-out period, then the Bus Timer will drive BERR* low.
29. OBSERVATION 4.19:
The Interrupt Handler is guaranteed that the data bus remains valid until it drives DSA* high.
30. OBSERVATION 4.20:
This guarantees that neither DTACK* nor BERR* goes high until the Interrupt Handler drives both DS0* and DS1* high.
31. OBSERVATION 4.21:
The Interrupt Handler is guaranteed that the data bus has been released by the time DTACK* and BERR* are high.

Table 4 - 18 Interrupter, Timing RULEs And OBSERVATIONS

Note: The numbers correspond to the timing parameters specified in Table 4-16.

4. OBSERVATION 4.22:
Interrupters are guaranteed that IACK*, LWORD*, and A[3..1] have been valid for this minimum time when they detect a falling edge on AS*.
 5. OBSERVATION 4.23:
All Interrupters are guaranteed this minimum high time on AS* between DTB cycles.
 6. OBSERVATION 4.24:
The responding Interrupter is guaranteed that none of D[31..0] will be driven by any other module until the responding Interrupter releases DTACK* and BERR* to high.
 7. OBSERVATION 4.25:
The responding Interrupter is guaranteed that the data bus will be released by all other modules by the time DSA* goes low.
 9. OBSERVATION 4.26:
The responding Interrupter is guaranteed that neither DS0* nor DS1* will go low until DTACK* and BERR* from the previous cycle have gone high.
 10. OBSERVATION 4.53:
Due to bus skew, Slaves on the DTB might detect a falling edge on DSA* before detecting the falling edge on AS*. However, Slaves are guaranteed that a falling edge on DSA* will not precede the falling edge on AS* by more than this time.
 11. OBSERVATION 4.27:
Interrupters are guaranteed this minimum time during which both data strobes are simultaneously high between cycles.
 12. OBSERVATION 4.28:
Interrupters are guaranteed that WRITE* has been high for this minimum time when they detect a falling edge on DSA*.
 13. OBSERVATION 4.29:
IF both data strobes are going to be driven low,
THEN the responding Interrupter is guaranteed that DSB* will go low within this maximum time after DSA*.
- And therefore:
- IF the DSB* does not go low within this maximum time,
THEN the responding Interrupter assumes that it is to respond with a single byte Status/Id.

Table 4 - 18 (Continued)

14. OBSERVATION 4.30:
The responding Interrupter is guaranteed that LWORD* and A[3..1] will remain valid until it drives DTACK* or BERR* low, provided it does so within the bus time-out period.
16. OBSERVATION 4.31:
The responding Interrupter is guaranteed that IACK* will remain low until it drives DTACK* or BERR* low, provided it does so within the bus time-out period.
18. OBSERVATION 4.32:
The responding Interrupter is guaranteed that AS* will remain low until it drives DTACK* or BERR* low, provided that it does so within the bus time-out period.
19. OBSERVATION 4.33:
Interrupters are guaranteed that the AS* will remain low for this minimum time.
20. OBSERVATION 4.34:
The responding Interrupter is guaranteed that once DSA* goes low, it will remain low until the Interrupter has driven DTACK* or BERR* low, provided that the Interrupter does so within the bus time-out period.
21. OBSERVATION 4.35:
The responding Interrupter is guaranteed that once DSB* goes low, it will remain low until the Interrupter has driven DTACK* or BERR* low, provided that the Interrupter does so within the bus time-out period.
23. OBSERVATION 4.36:
Interrupters are guaranteed that the WRITE* line remains high until both data strobes are high.
26. RULE 4.35:
The Interrupter MUST NOT drive the data bus until DSA* goes low.
27. RULE 4.36:
The responding Interrupter MUST NOT drive DTACK* low before it drives the data lines with a valid Status/Id.

OBSERVATION 4.37:
This time does not apply to cycles where the Interrupter drives BERR* low instead of DTACK*.
28. RULE 4.37:
The responding Interrupter MUST wait this minimum time after DSA* goes low before driving DTACK* or BERR* low.

Table 4 - 18 (Continued)

29. RULE 4.38:
Once the responding Interrupter has driven DTACK* low, it MUST NOT change D[31..0] until DSA* goes high.
30. RULE 4.39:
Once the responding Interrupter has driven DTACK* or BERR* to low, it MUST NOT release it until it detects both DS0* and DS1* high.
31. RULE 4.40:
The responding Interrupter MUST release all of D[31..0] before releasing DTACK* and BERR* to high.
32. OBSERVATION 4.38:
The responding Interrupter is guaranteed that IACK*, LWORD*, and A[3..1] have been valid for this minimum time when it detects a falling edge on DSA*. This time is derived from timing parameters 4 and 10.
33. OBSERVATION 4.54:
During data transfer cycles, Interrupters are guaranteed that either DS0*, or DS1*, or both remain low for at least this minimum time. This time is derived from timing parameter 28, where the responding Interrupter is required to wait a minimum time before driving BERR* or DTACK* low.
- 34A. OBSERVATION 4.39:
The Interrupter is guaranteed that AS* has been low for this minimum time, when it detects a falling edge on IACKIN*.
35. RULE 4.41:
A participating Interrupter MUST drive its IACKOUT* high within this maximum time after the rising edge on AS*.
36. RULE 4.42:
The Interrupter MUST NOT drive the data bus until its IACKIN* line goes low.
37. RULE 4.43:
IF a participating Interrupter drives the data bus,
THEN it MUST release it before driving its IACKOUT* line low.
- 38A. RULE 4.44:
A participating Interrupter MUST NOT drive its IACKOUT* line low, until it detects its IACKIN* line low.
- 38B. RULE 4.45:
The responding Interrupter MUST NOT drive its DTACK* line low, until it detects its IACKIN* line low.

Table 4-18. (Continued)

39. OBSERVATION 4.40:
This time guarantees that each Interrupter's IACKIN* line will go high within this time after the rising edge on AS*. This time is derived from timing parameter 35, where the IACK Daisy-Chain Driver and participating Interrupters are required to drive their IACKOUT* line high within a maximum time.
40. OBSERVATION 4.41:
All Interrupters are guaranteed that their IACKIN* line will stay high for this minimum time between consecutive DTB cycles.
41. OBSERVATION 4.42:
This time guarantees that A[3..1] and LWORD* remain valid until this time after the participating Interrupter drives its IACKOUT* low, provided it does so within the bus time-out period.
43. OBSERVATION 4.43:
This time guarantees that AS* remains low for this minimum time after the participating Interrupter drives its IACKOUT* low, provided it does so within the bus time-out period.

Table 4 - 19 IACK Daisy-Chain Driver, Timing RULES And OBSERVATIONS

Note: The numbers correspond to the timing parameters specified in Table 4-16.

OBSERVATION 4.44:

Since the backplane connects IACK* to the Slot 1 IACKIN*, these two signals are equivalent. Therefore, all RULEs and OBSERVATIONS that apply to one, are applicable to the other as well.

5. OBSERVATION 4.45:

The IACK Daisy-Chain Driver is guaranteed this minimum high time on AS* between DTB cycles.

19. OBSERVATION 4.46:

The IACK Daisy-Chain Driver is guaranteed that the AS* will remain low for this minimum time. This time is derived from timing parameters 8, 16, and 27 of the Interrupter.

32. OBSERVATION 4.47:

The IACK Daisy-Chain Driver is guaranteed that IACK* (and the Slot 1 IACKIN*) has been valid for this minimum time when it detects a falling edge on DSA*.

34. RULE 4.46:

IF the IACKIN* line is low when the IACK Daisy-Chain Driver detects a falling edge on DSA*,
THEN it MUST drive its IACKOUT* line low, but it MUST NOT do so until this time after the falling edge on DSA*.

OBSERVATION 4.48:

The IACK Daisy-Chain Driver does not drive its IACKOUT* line low every time DSA* goes low. It only does so when the IACK* line is low as well, indicating that an interrupt acknowledge cycle is in progress.

35. RULE 4.47:

IF the IACK Daisy-Chain Driver drives its IACKOUT* line low,
THEN it MUST drive its IACKOUT* high within this time after the rising edge of AS*.

40. RULE 4.48:

The IACK Daisy-Chain Driver MUST NOT drive IACKOUT* low until it has been high for this minimum time.

42. OBSERVATION 4.49:

IF the IACK Daisy-Chain Driver drives its IACKOUT* line low within the bus time-out period,
THEN this time guarantees that IACK* (and the Slot 1 IACKIN*) remains valid for this minimum time.

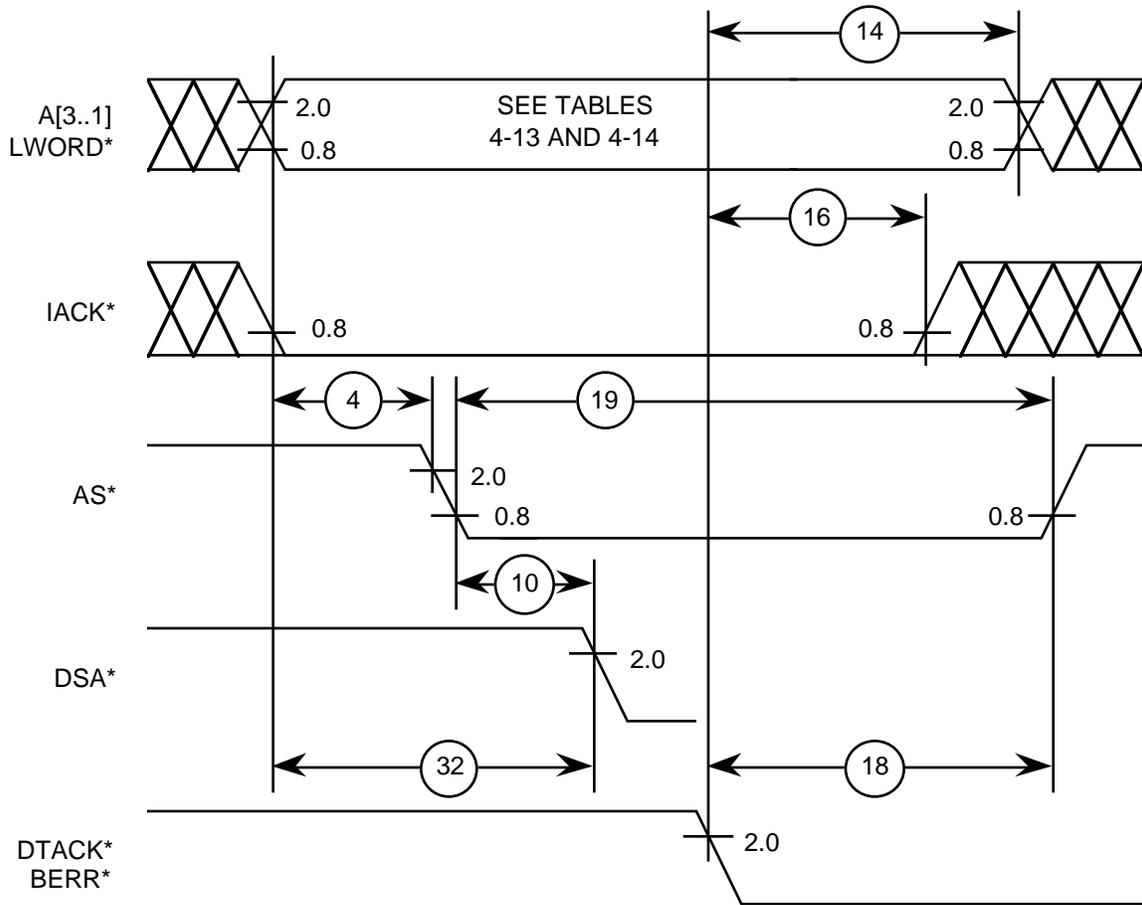


Figure 4 - 16 Interrupt Handler And Interrupter--Interrupter Selection Timing Single-Byte, Double-Byte, And Quad-Byte Interrupt Acknowledge Cycles

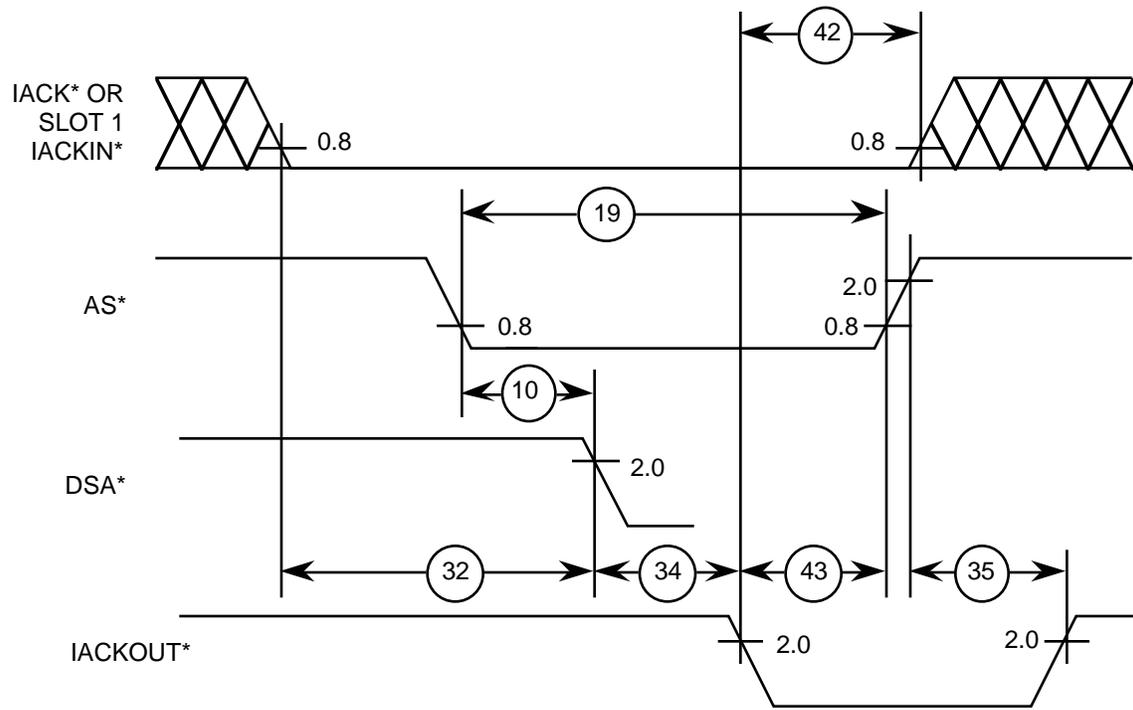


Figure 4 - 17 IACK Daisy-Chain Driver--Interrupter Selection Timing Single-Byte, Double-Byte, And Quad-Byte Interrupt Acknowledge Cycles

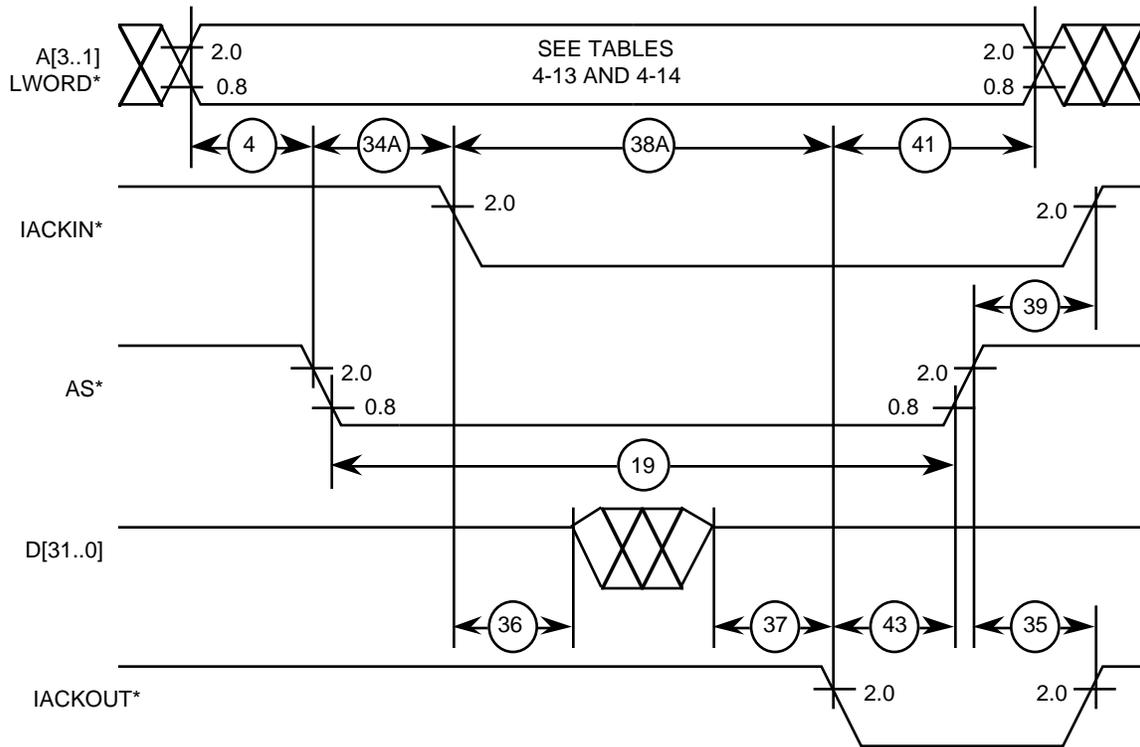


Figure 4 - 18 Participating Interrupter--Interrupter Selection Timing
Single-Byte, Double-Byte, And Quad-Byte Interrupt Acknowledge Cycles

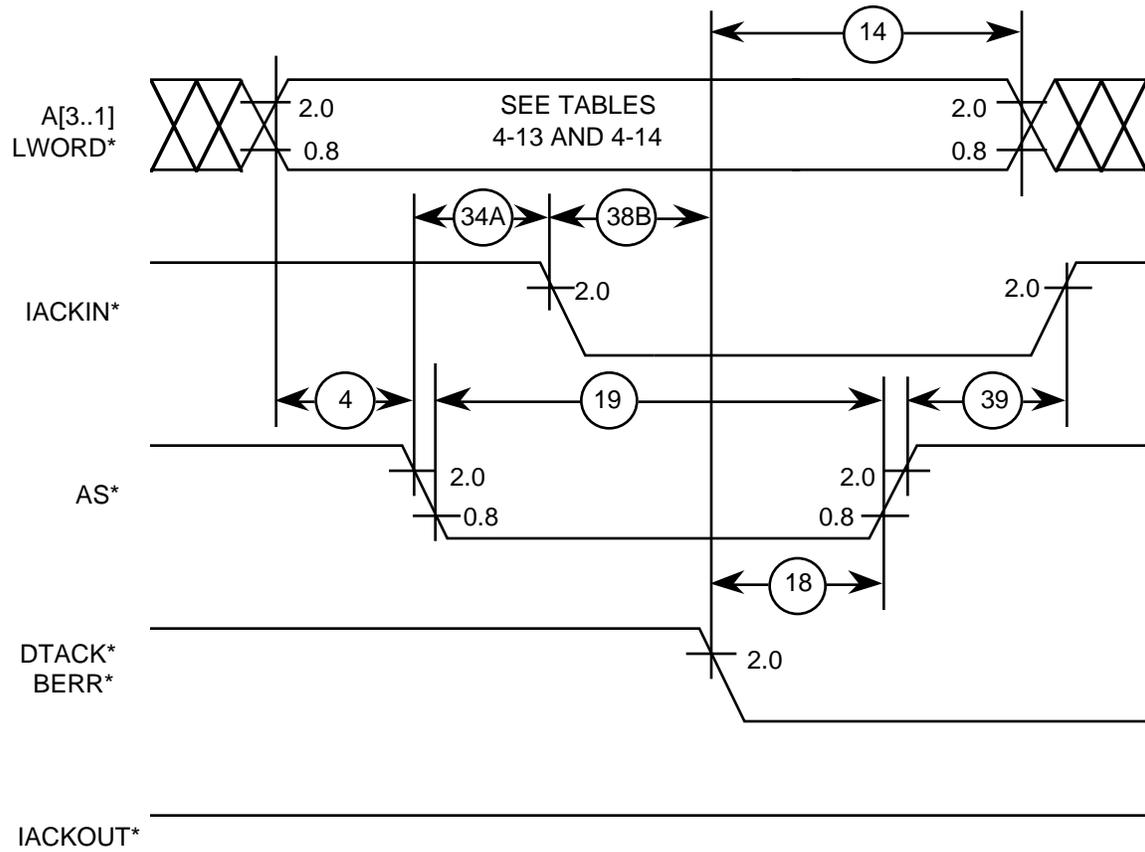


Figure 4 - 19 Responding Interrupter--Interrupter Selection Timing
Single-Byte, Double-Byte, And Quad-Byte Interrupt Acknowledge Cycles

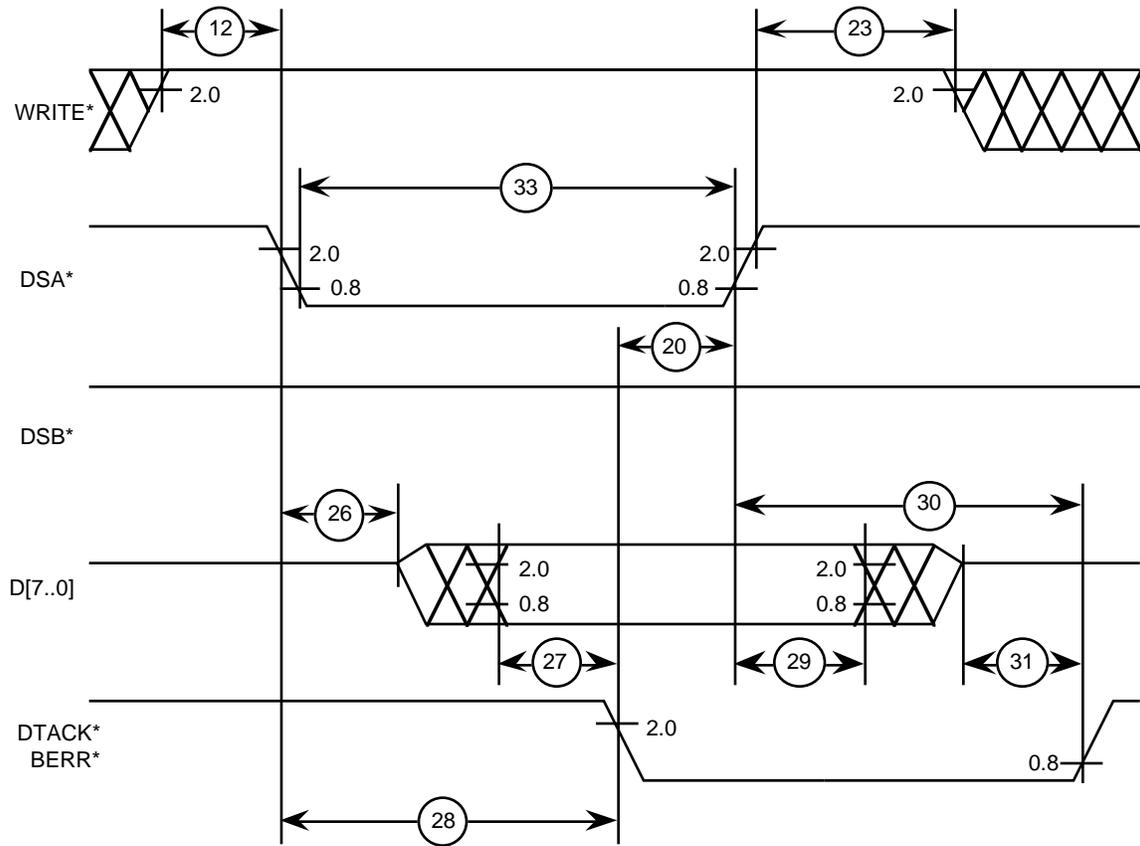


Figure 4 - 20 Interrupt Handler--Status/Id Transfer Timing
Single-Byte Interrupt Acknowledge Cycle

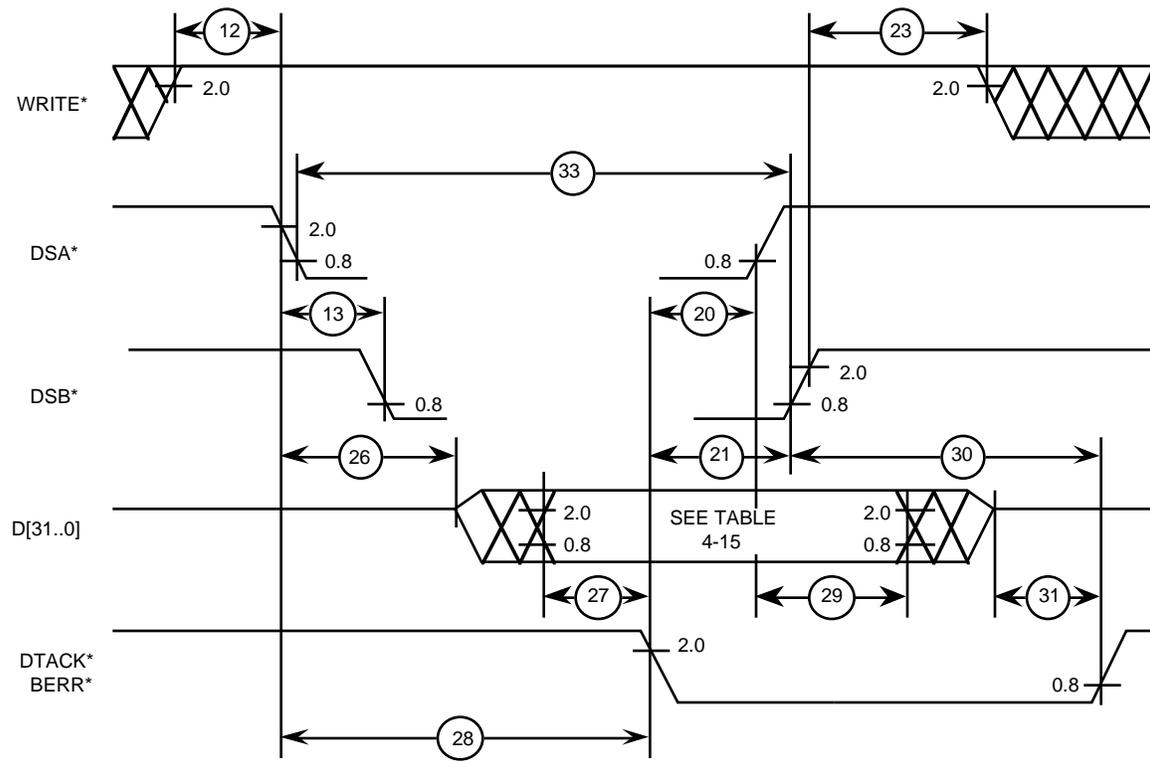


Figure 4 - 21 Interrupt Handler--Status/Id Transfer Timing
Double-Byte And Quad-Byte Interrupt Acknowledge Cycles

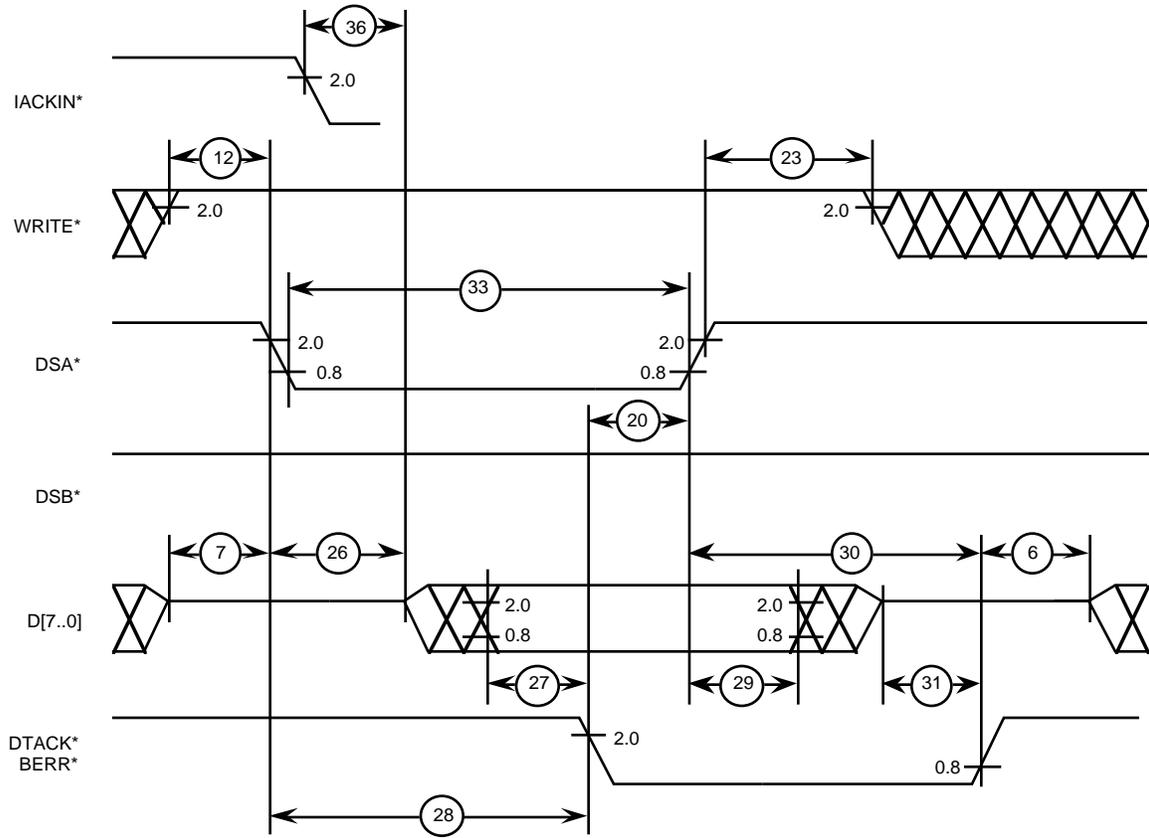


Figure 4 - 22 Responding Interrupter--Status/Id Transfer Timing
Single-Byte Interrupt Acknowledge Cycle

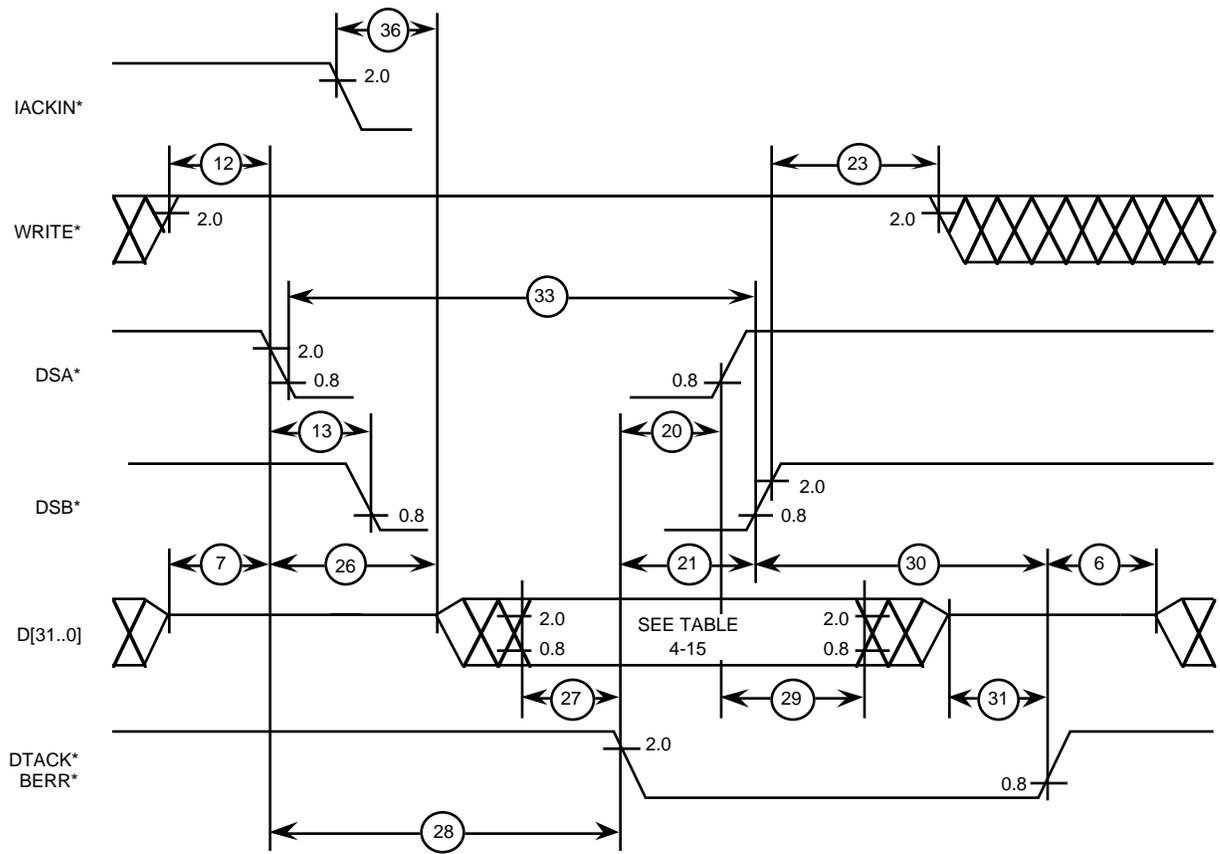


Figure 4 - 23 Responding Interrupter--Status/Id Transfer Timing
 Double-Byte Interrupt Acknowledge Cycle
 Quad-Byte Interrupt Acknowledge Cycle

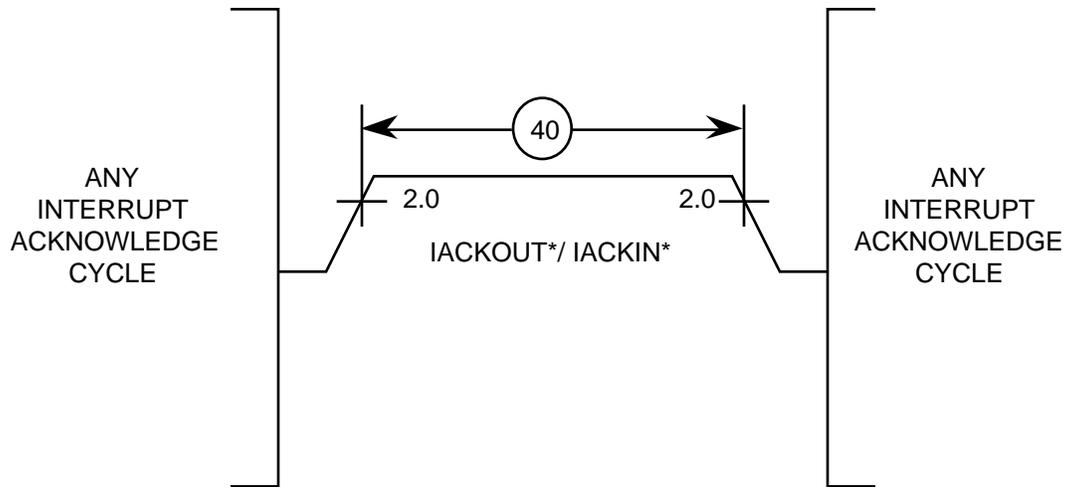


Figure 4 - 24 IACK Daisy-Chain Driver, Responding Interrupter, And Participating Interrupter IACK Daisy-Chain Inter-Cycle Timing

PAGE LEFT BLANK

Chapter 5

Utility Bus

5.1 Introduction

This chapter identifies and defines the signal lines and modules which provide utility functions for the VMEbus. The Utility Bus supplies periodic timing, initialization and diagnostic capability for the VMEbus (See Figure 5-1).

5.2 Utility Bus Signal Lines.

The Utility Bus signal lines are listed below:

System Clock	(SYSCLK)
AC Fail	(ACFAIL*)
System Reset	(SYSRESET*)
System Failure	(SYSFAIL*)
Serial Bus A	(SERA)
Serial Bus B	(SERB)

5.3 Utility Bus Modules

5.3.1 System Clock Driver.

The system clock is an independent, nongated, fixed frequency, 16 MHz, 50% (nominal) duty cycle signal. The SYSCLK driver is located on the System Controller located in board slot one (see Chapter 1). It provides a known time base that is useful for counting off time delays. Figure 5-2 shows the SYSTEM CLOCK DRIVER timing diagram. For additional information see 6.5.3.

OBSERVATION 5.1:

SYSCLK has no fixed phase relationships with other VMEbus timing.

5.3.2 The Serial Bus Lines

RULE 5.11:

Connector pins 'b21' and 'b22' of P1/J1 are reserved for communication via a serial bus. These pins MUST be bussed across the backplane and terminated in accordance with the specific implementation. Specific logical and electrical definitions relating to the serial bus are beyond the scope of this document. Serial bus implementations may include but are not limited to IEEE 1394 and AUTOBAHN technology.

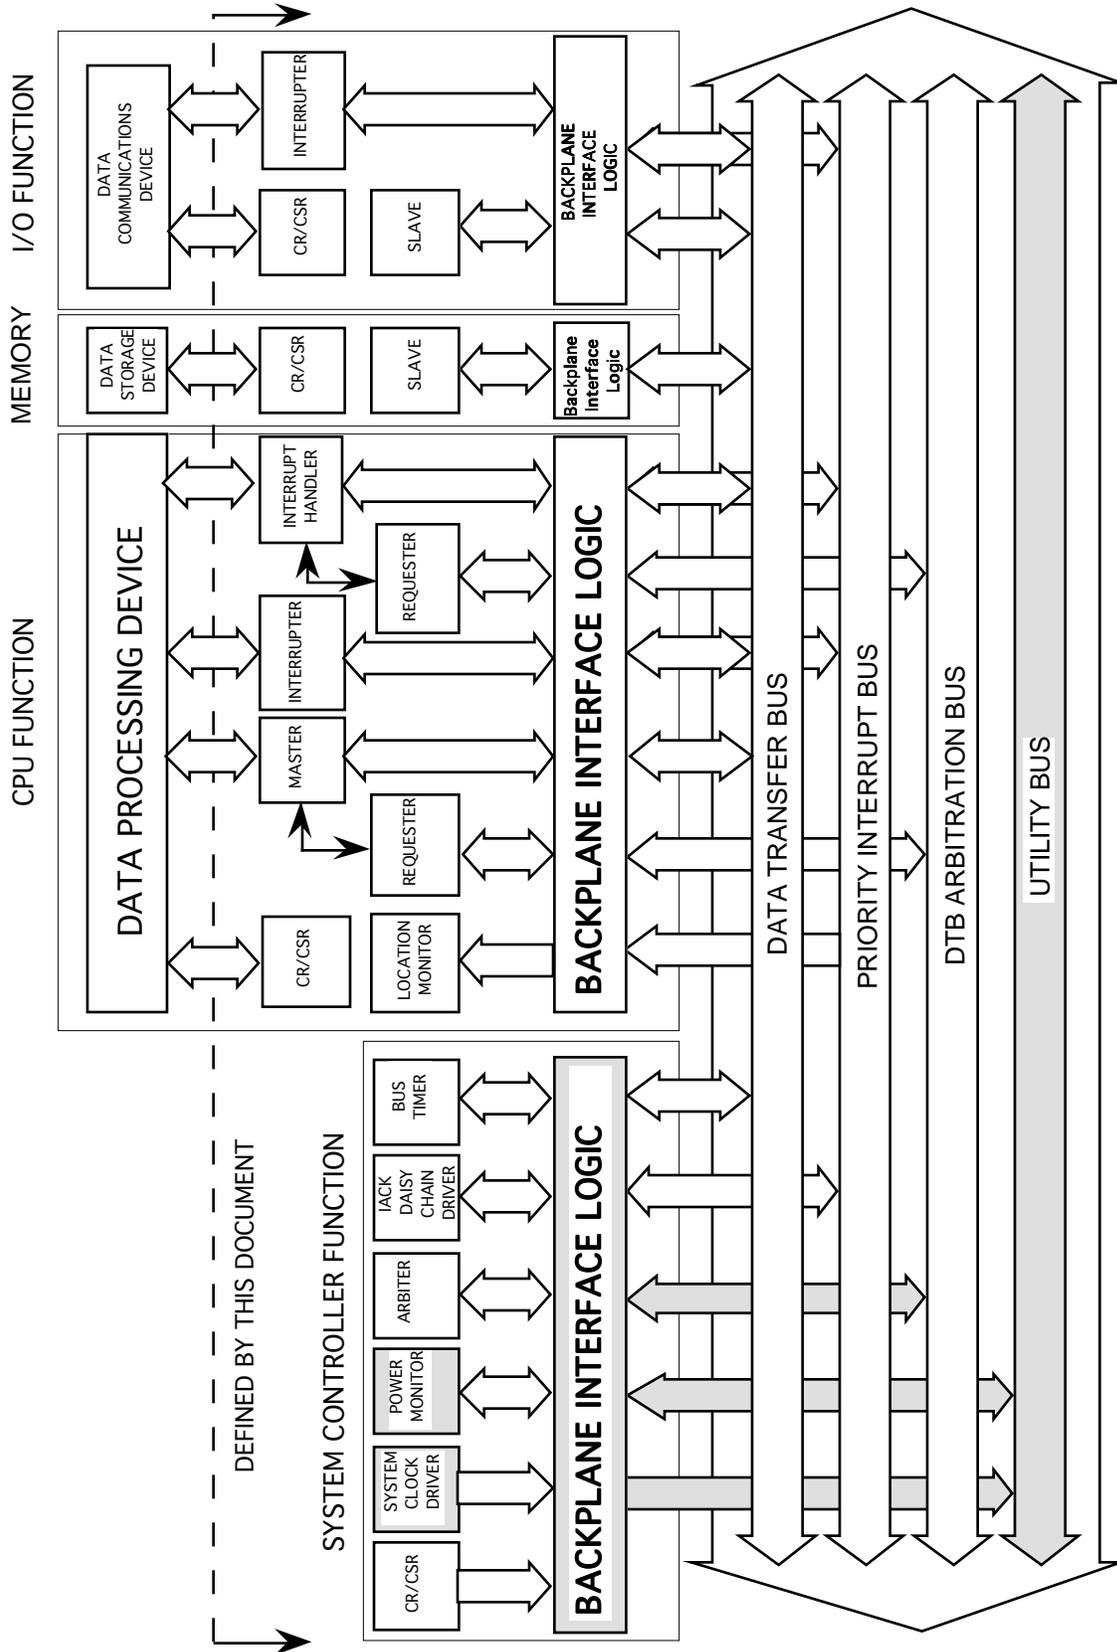


Figure 5 - 1 Utility Bus Block Diagram

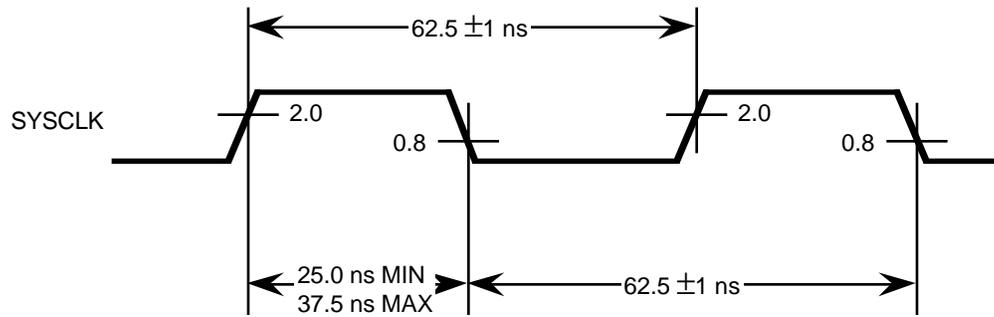


Figure 5 - 2 System Clock Driver Timing

5.3.3 Power Monitor

Figure 5-3 is the block diagram for the Power Monitor module. This module detects power failures and signals the VMEbus system in time to effect orderly shut-down. When power is then reapplied to the system the Power Monitor signals the other VMEbus modules for proper initialization.

The Power Monitor might also monitor a manually operated push button and initialize the VMEbus system whenever that button is depressed by the operator.

The ACFAIL* and SYSRESET* transitions, and the point at which the system DC voltages violate specification, have certain timing relationships. These relationships are shown in Figures 5-4 and 5-5.

PERMISSION 5.1:

VMEbus systems MAY be built with or without a Power Monitor module.

RULE 5.1:

Power Monitors MUST comply with the timing given in Figures 5-4 and 5-5.

PERMISSION 5.2:

The SYSRESET* line MAY be driven low by any VMEbus board to initialize the system from a manual push-button. Where a board drives SYSRESET*, but does not drive ACFAIL*, the timing in Figure 5-4 and Figure 5-5 does not apply.

RULE 5.2:

Whenever any board drives SYSRESET* low, it MUST hold SYSRESET* low for a minimum period of 200 milliseconds.

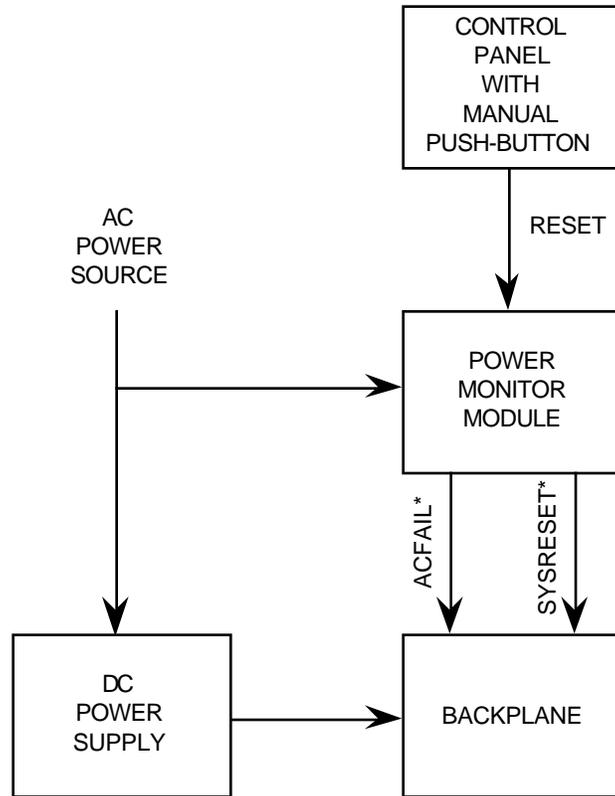


Figure 5 - 3 Block Diagram Of Power Monitor Module

5.4 System Initialization and Diagnostics

The VMEbus provides protocols which allow the system to be shut down and powered up in an orderly manner. Two signal lines are used in the power-down and power-up sequence: ACFAIL* and SYSRESET*. Another signal line is used in the power-up sequence: SYSFAIL*.

The following specifies the behavior of VMEbus functional modules during the power-down sequence:

RECOMMENDATION 5.1

Design Masters so that they will not request the bus for any purpose except power-fail activity, after ACFAIL* has been low for 200 microseconds.

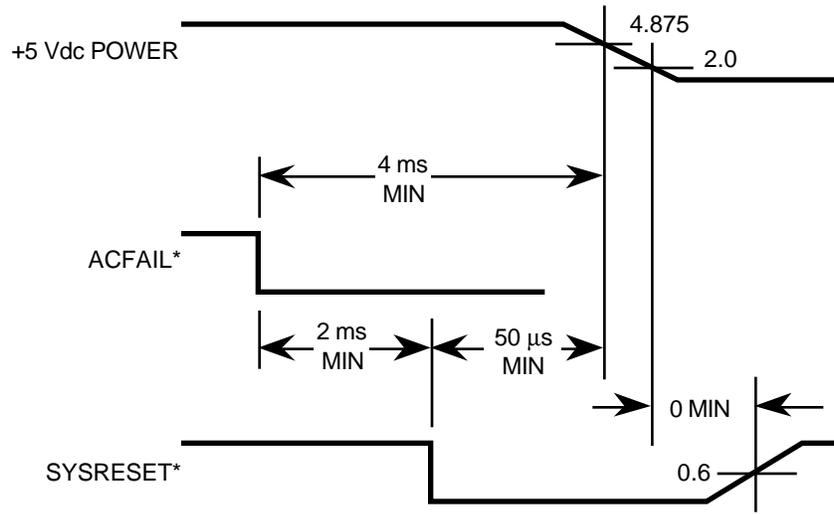


Figure 5 - 4 Power Monitor Power Failure Timing

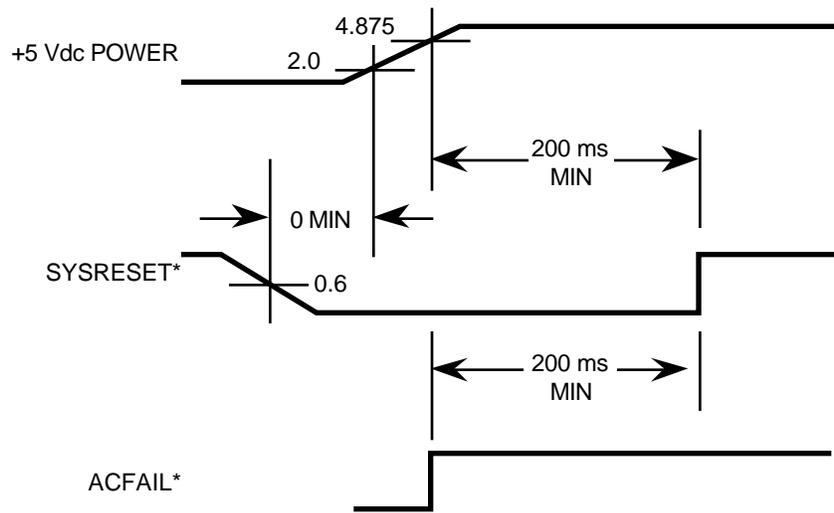


Figure 5 - 5 Power Monitor System Restart Timing

RECOMMENDATION 5.2:

IF Masters and Interrupt Handlers have a bus request pending prior to detecting ACFAIL* low,
 THEN they should limit their subsequent non-power-fail activity to 200 microseconds.

OBSERVATION 5.2:

Bus accesses required to save and restore system data to global memory depend upon the application, and are not specified here. (The operating system has to ensure that data saved during the shutdown process is restored prior to system operation.) In the case of a multiprocessing system, this might require some interprocessor communication.

System Reset (SYSRESET*) is an open-collector line driven by the Power Monitor module, or by any board in response to a push-button switch closure.

OBSERVATION 5.3:

Special circuitry is needed where push-button reset switches are used, to ensure that switch bounce does not cause the board to violate the 200 ms minimum SYSRESET* low time.

RULE 5.3

The System Clock driver MUST continue to provide the specified SYSCLK waveform regardless of the state of the SYSRESET* line.

PERMISSION 5.3 not used

PERMISSION 5.4

When SYSRESET* goes low, any board that requires more than 200 ms to complete its initialization may turn on its SYSRESET* driver low to maintain SYSRESET* low for the required period.

RULE 5.4

IF the +5 Vdc power source is within its specified range when SYSRESET* goes low,
 THEN functional modules MUST satisfy the timing RULEs given in Table 5-1 within the specified time after SYSRESET* goes low.

RULE 5.5

IF SYSRESET* is low when the +5 Vdc enters its specified range,
 THEN functional modules MUST satisfy the timing RULEs given in Table 5-1 and then MUST refrain from driving specified lines until SYSRESET* goes high.

RULE 5.6

After satisfying the RULEs in Table 5-1, functional modules MUST NOT change the state of their drivers until SYSRESET* goes high, unless the +5 V dc power source exits its specified range.

RULE 5.7:
IF the +5 Vdc power source is within the specified range when SYSRESET* goes low, and a Master or Interrupt Handler is driving AS*, DS0*, or DS1* low.
THEN it **MUST** maintain these strobes low long enough to satisfy the minimum low times given in Chapters 2 and 4.

Table 5 - 1 Module Drive During Power-Up and Power-Down Sequences

MODULE	MUST REFRAIN FROM DRIVING	After SYSRESET* HAS BEEN LOW FOR
Masters and Interrupt Handlers	AS*, DS0*, or DS1* from high to low	5 μ s
Masters and Interrupt Handlers	IACK*, LWORD*, AS*, DS0*, DS1*, AM[5..0], A[31..1], WRITE*, or D[31..0]	20 μ s
Slaves and Interrupters	D[31..0], DTACK*, or BERR*	30 μ s
Interrupters	IRQ[7..1]*	30 μ s
Bus Timer	BERR*	30 μ s
Arbiter	BG[3..0]OUT* from high to low	5 μ s
Arbiter	BG[3..0]IN* low	30 μ s
Requesters	BBSY*	30 μ s
Requesters	BG[3..0]OUT* from high to low	5 μ s
Requesters	BG[3..0]OUT* low	30 μ s

SYSFAIL* is an open collector line that is held low when the system is powered-up and remains low until system self-tests are complete (see Figure 5 - 6). The following applies:

SUGGESTION 5.1

On intelligent Master boards, include a locally accessible control register bit that is initialized to drive SYSFAIL* low when power is first applied. This permits the board's local intelligence to do a self-test and release SYSFAIL* only if the self-test passes.

SUGGESTION 5.2

Design non-intelligent boards with a globally accessible control register bit that is initialized to drive SYSFAIL* low. This allows a Master on the VMEbus to run a test on the non-intelligent board, and then write to the global control register bit, releasing the board's SYSFAIL* driver.

SUGGESTION 5.3

Where a SYSFAIL* control register bit is included on VMEbus boards, provide a status LED on the board's front panel to indicate the status of the control register bit. Then, if a system failure is indicated by the SYSFAIL* signal line, a visual inspection will help determine which board has failed.

RULE 5.8:

IF SYSFAIL* control register bits are included on VMEbus boards,
 THEN they MUST drive SYSFAIL* low within 50 milliseconds after SYSRESET* goes low, as shown in Figure 5-6 and hold SYSFAIL* asserted during the assertion of SYSRESET*.

PERMISSION 5.4

A VMEbus board MAY also drive SYSFAIL* low at any time during normal operation to indicate that it has detected some kind of failure.

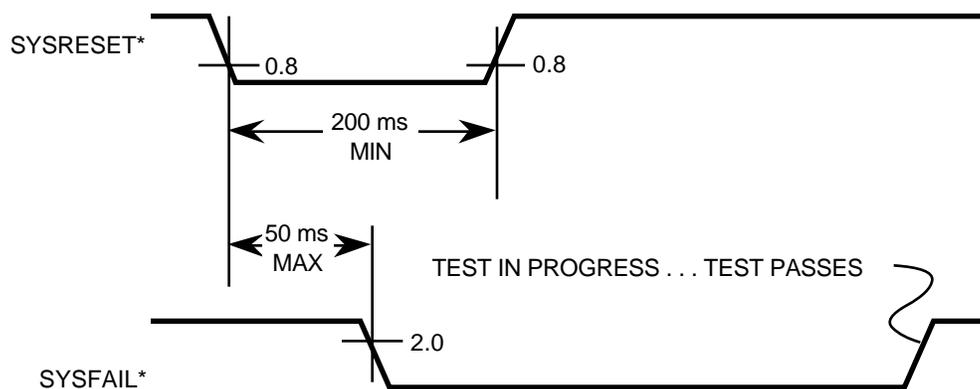


Figure 5 - 6 SYSRESET* and SYSFAIL* Timing Diagram

5.5 Power and Ground pins

Figure 5-7 gives the current rating for the VMEbus power pins at various temperatures.

OBSERVATION 5.4

Some connector pins have a slightly higher contact resistance than others when plugged into the backplane. This produces unbalanced current flow in pins which are paralleled. Suppose that two pins are paralleled and are carrying a total of two amperes of current. If the contact resistance on one is 1 milliohm and the other is 2 milliohms, then one pin will be carrying only 0.67 amperes while the other carries 1.33 amperes.

RULE 5.9

VMEbus connector pins **MUST** be capable of carrying the currents shown by the solid line in Figure 5-7.

OBSERVATION 5.5

If one or more power pins fail completely, all of the load current flows through the remaining pins. For example, if half of the pins fail, the remaining pins carry twice the normal current. Depending upon the load current, this might cause damage to these remaining good pins.

SUGGESTION 5.4

When designing a VMEbus board with a high current load, divide the board's area into zones which are each powered by a separate power grid. Don't connect these grids to each other on the VMEbus board. Instead, connect each to its own VMEbus power pin.

OBSERVATION 5.6:

IF a double height VMEbus board which draws more power than its P1 connector can provide when it is plugged into a subrack which contains only a J1 backplane,

THEN its P1 power pins will overheat, and might be damaged.

SUGGESTION 5.5:

Boards that support modes that drive all the address and data lines plus control signals at the same time should be designed with DIN connectors that provide 16 signal grounds on the lower outer shield.

OBSERVATION 5.8:

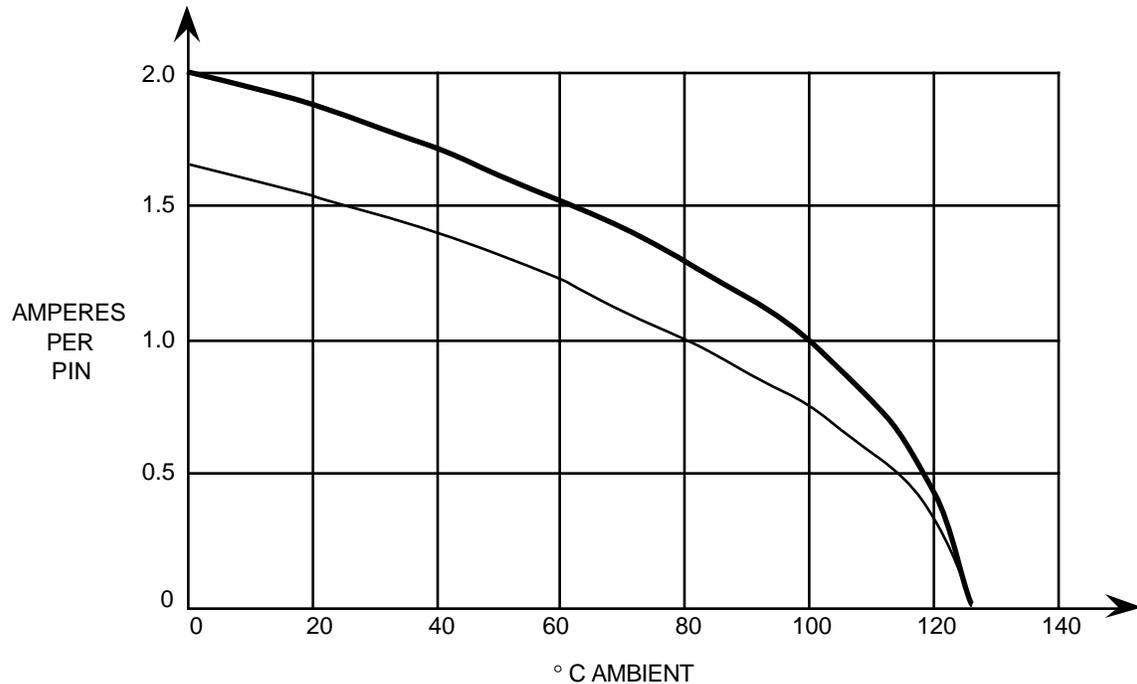
Use of DIN connectors with signal grounds on the outer shell will improve signal integrity.

5.6 Reserved Line

The Reserved line originally defined in VME Specification Revision C.1 is now used for the **RETRY*** function. Refer to Chapter 2.

OBSERVATION 5.7: not used

RULE 5.10: not used



NOTES:

1. The solid bold line shows the current that can be drawn per power pin where each pin is connected to a separate power grid.
2. The thin line shows the current that can be drawn per power pin where two or more pins are connected to a common on-board power grid.

Figure 5 - 7 Current Rating For Power Pins

5.7 Auto Slot Id

Auto ID is an optional method of assigning the CR/CSR base address to each VMEbus board. The Auto ID Slave uses a level 2, D08(O) Interrupter along with additional board-specific hardware to obtain the CR/CSR base address. An Auto ID Master, called the Monarch, uses a level 2, D08(O) Interrupt Handler to acknowledge the Auto ID interrupt and assigns a base address to each board's CR/CSR.

Auto ID Slaves have the defined CSR register set and a CR/CSR Base Address Register located in the CR/CSR address space. Accesses to the board in the CR/CSR space are inhibited when a system reset occurs and also when the board powers up. Accesses to the board in the CR/CSR space remain inhibited until the board participates in the identification process. After power up or a system reset the Auto ID Slave drives SYSFAIL* low and inhibits CR/CSR accesses. It may also clear its CR/CSR Base Address Register to 0x00 at this time. When SYSRESET* goes high, the Auto ID Slave will prepare itself to enter into the Auto ID process. When ready to enter into the Auto ID process, the Auto ID Slave

- 1) signals its Interrupter to generate a level 2 interrupt, and

2) quits driving SYSFAIL* low.

When the Auto ID Slave's Interrupter responds to its level 2 IACK cycle, it

- 1) enables Slave accesses in its CR/CSR space,
- 2) presents a Status/ID byte that identifies the interrupt as an Auto ID request,
- 3) drives DTACK* low, and
- 4) releases IRQ2* (if a ROAK Interrupter).

The Monarch may now perform CR/CSR reads and writes to this board at A[23..19] = 0x00 in the CR/CSR space as defined by the Status/ID byte. Before responding to another IRQ2*, the Monarch reassigns the base address of this board's CR/CSR by writing a new value to this board's CR/CSR Base Address Register.

RULE 5.12:

Within 5 μ s after SYSRESET* goes low, each Auto ID board MUST inhibit accesses to its CR/CSR. It MUST continue to inhibit CR/CSR accesses until after it receives IACKIN* low, but before it drives DTACK* low during the interrupt acknowledge cycle in which it will provide the Status/ID byte for an Auto ID request.

RULE 5.13:

Each Auto ID Slave with a CR/CSR MUST have a writable CR/CSR Base Address Register accessible within the board's CR/CSR address space.

RULE 5.14:

An Auto ID Slave MUST initialize its CR/CSR Base Address Register to 0x00 before its Interrupter acknowledges the Auto ID interrupt.

RULE 5.15:

A write to the CR/CSR Base Address Register MUST move the board's CR/CSR base address to the value written to this register before allowing DTACK* to go high at the end of the write cycle.

RULE 5.16:

The Auto ID Slave MUST release IRQ2* before terminating the data transfer which remaps the CR/CSR or Auto ID register unless IRQ2* is being driven low for another purpose.

OBSERVATION 5.9:

A ROAK Interrupter will easily guarantee RULE 5.15 in an Auto ID system. However, if the remainder of the board's functions require an RORA Interrupter, it may be desirable to use an RORA Interrupter for Auto ID requests.

PERMISSION 5.5:

Either ROAK or RORA Interrupters MAY be used for Auto ID interrupt.

RECOMMENDATION 5.3:

To provide interoperability between boards from different vendors, the Status/ID byte for Auto ID requests should be 0xFE.

The Monarch requires a level 2 Interrupt Handler and a Master capable of initiating data transfer cycles in the CR/CSR space. After powering up or after a system reset completes, the Monarch typically waits to acknowledge interrupt requests on IRQ2* until after SYSFAIL* goes high. If the Monarch then detects IRQ2* low its Interrupt

Handler initiates an Interrupt Acknowledge cycle on level 2. When the Monarch receives the Status/ID byte for an Auto ID request it

- 1 masks IRQ2*,
- 2 performs accesses at 0x00 in the CR/CSR space to gain information about the Auto ID Slave board,
- 3 moves the CR/CSR to its new location, and
- 4 unmask IRQ2*.

Auto ID cycles do not start until the last Auto ID Slave releases the SYSFAIL* line. All Auto ID requests utilize the level 2 Interrupt Acknowledge cycle. Interrupt Acknowledge Cycles on the same level are resolved in slot-sequential order using the IACK Daisy-Chain. Therefore, if the Monarch waits until it detects SYSFAIL* high before initiating level 2 Interrupt Acknowledge cycles, then each Auto ID Slave will respond in slot-sequential order. The Monarch can use this knowledge to assign CR/CSR base addresses in sequential order to each Auto ID Slave board so that each Auto ID Slave board's CR/CSR base address relates to its relative slot position.

OBSERVATION 5.10:

It is possible for a failed board to keep SYSFAIL* asserted after all Auto ID boards have quit driving SYSFAIL* low. The low SYSFAIL* will keep the Monarch from assigning CR/CSR base addresses to Auto ID Slaves.

RECOMMENDATION 5.4:

Design Monarchs such that if they do not detect SYSFAIL* high within a certain amount of time, they will assume a system error has occurred and may optionally begin the Auto ID process.

OBSERVATION 5.11:

Some Auto ID systems may not require knowledge of relative slot positions.

PERMISSION 5.6:

In those systems which do not require knowledge of relative slot positions, the Monarch does not have to wait for SYSFAIL* to be high before responding to a low level on IRQ2*. In such a system, each Auto ID Slave MAY receive its final CR/CSR base address based on some prior agreement between the Monarch and the Auto ID Slaves.

PERMISSION 5.7:

In those systems which do not require knowledge of relative slot positions, the Auto ID Slave MAY quit driving SYSFAIL* low before driving IRQ2* low to request an Auto ID.

OBSERVATION 5.12:

Some systems may have a mix of Auto ID Slaves and boards with their CR/CSR base addresses at fixed or default locations in the CR/CSR space.

RECOMMENDATION 5.5:

To allow for mixed Auto ID and non-Auto ID systems, the Monarch should search the system CR/CSR space for responding CR/CSR boards and either record their base addresses, or reassign their base addresses to some benign values before beginning the Auto ID process. This helps keep the Monarch from assigning the same base address to more than one CR/CSR.

RULE 5.17:

The Monarch MUST NOT assign the same CR/CSR base address to more than one CR/CSR.

RULE 5.18:

In mixed Auto ID and non-Auto ID systems, non-Auto ID boards with CR/CSR's MUST NOT be accessible at base address 0x00 in the CR/CSR space.

OBSERVATION 5.13:

Rule 5.18 keeps non-Auto ID boards from interfering with Auto ID boards during the Auto ID sequence.

Figure 5-8 shows a typical Auto ID sequence for a three-slot system.

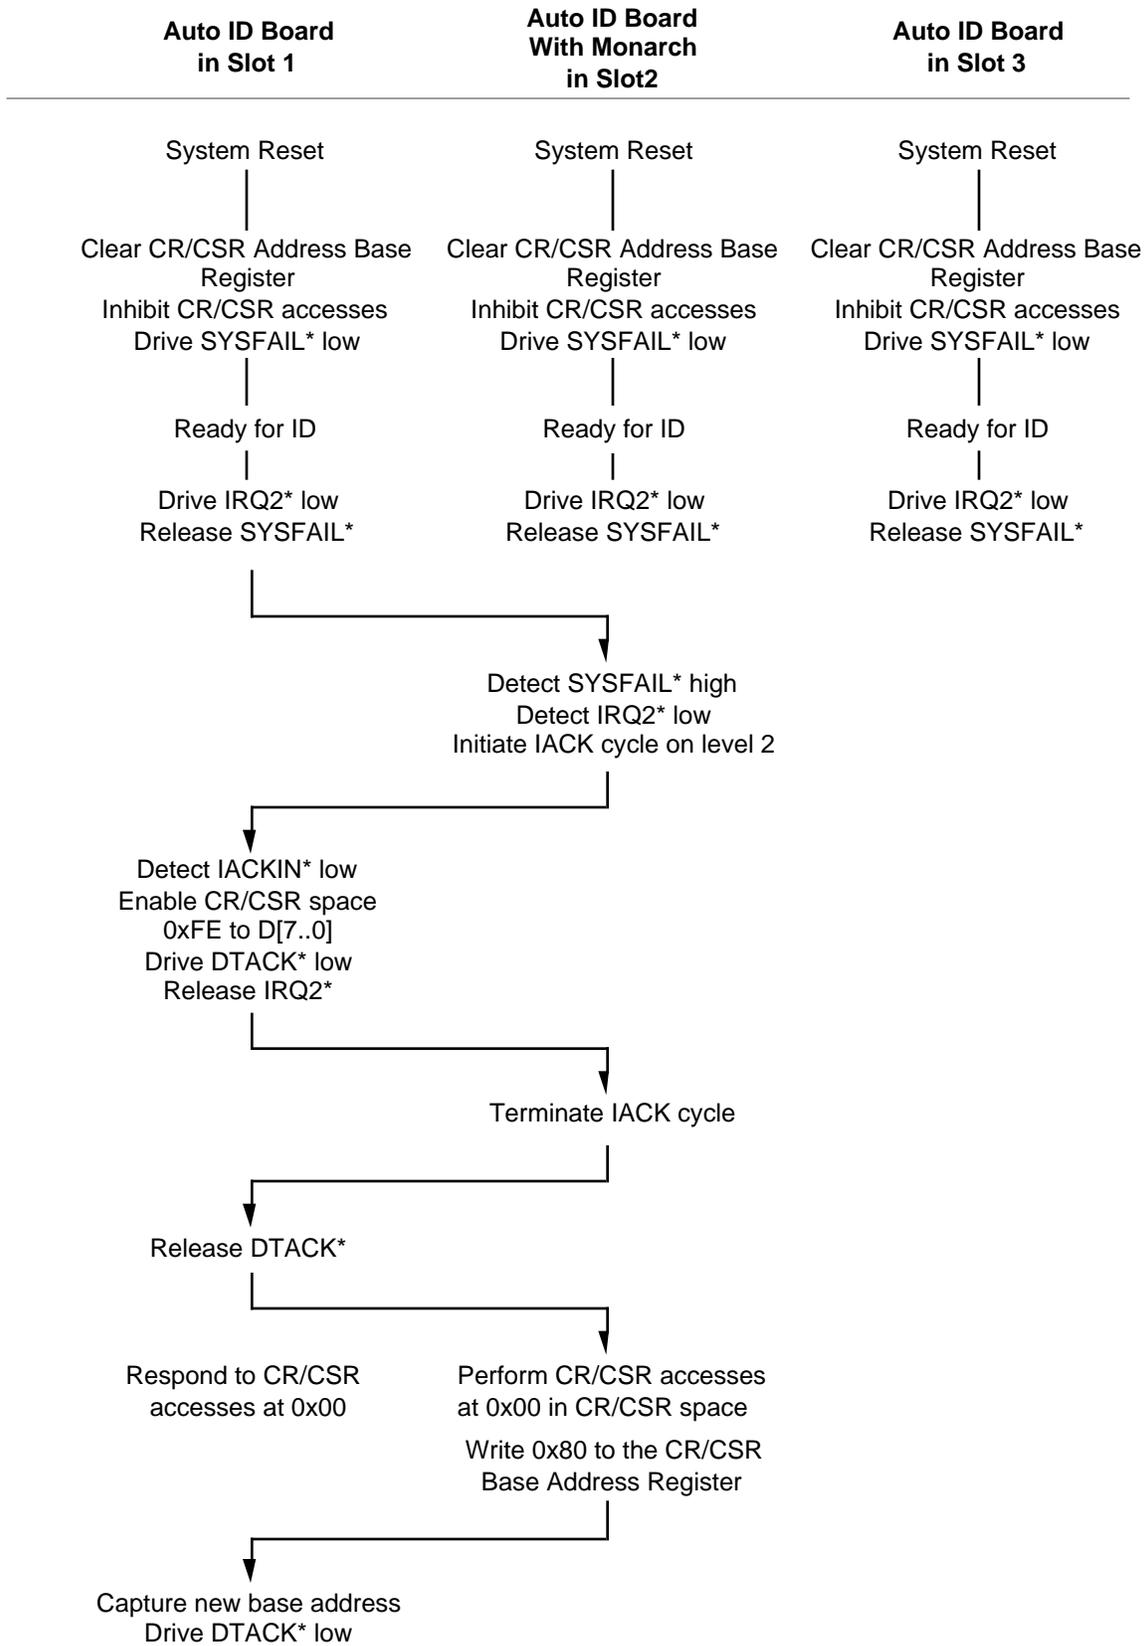


Figure 5 - 8 CR/CSR Auto ID Slave Initialization Algorithm, 1 of 3 pages

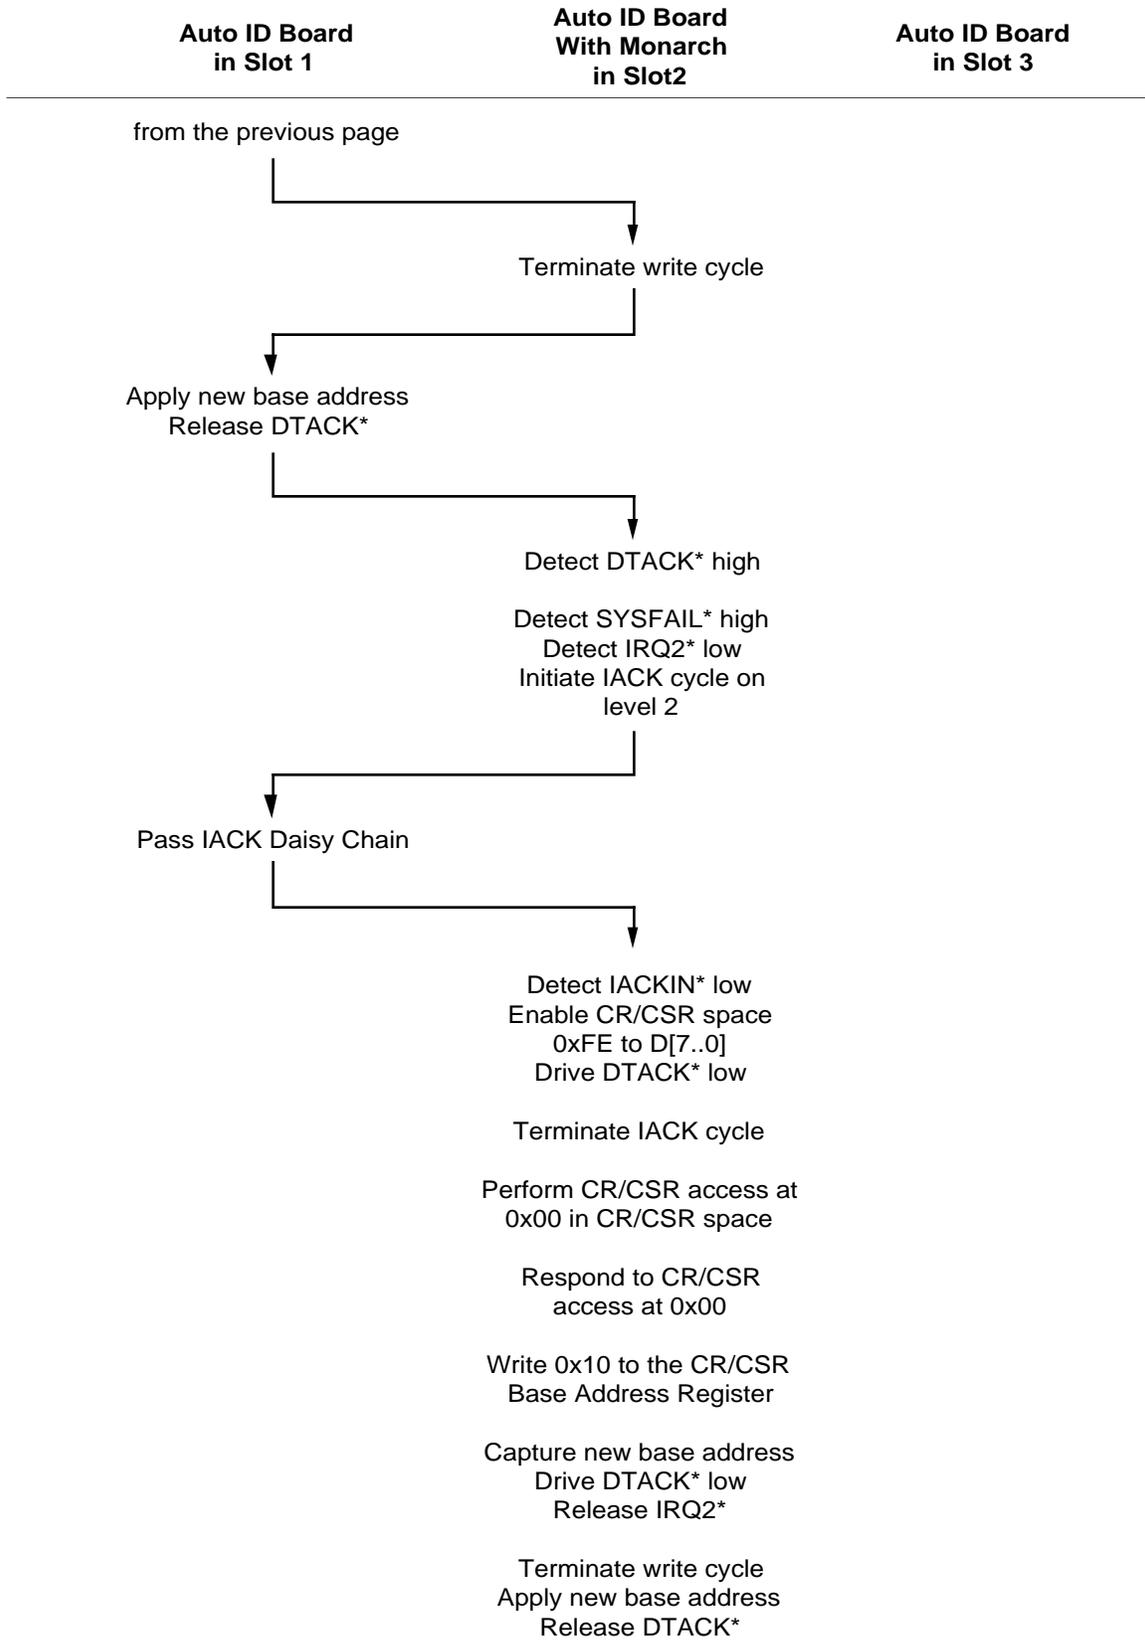


Figure 5 - 8 CR/CSR Auto ID Slave Initialization Algorithm, 2 of 3 pages

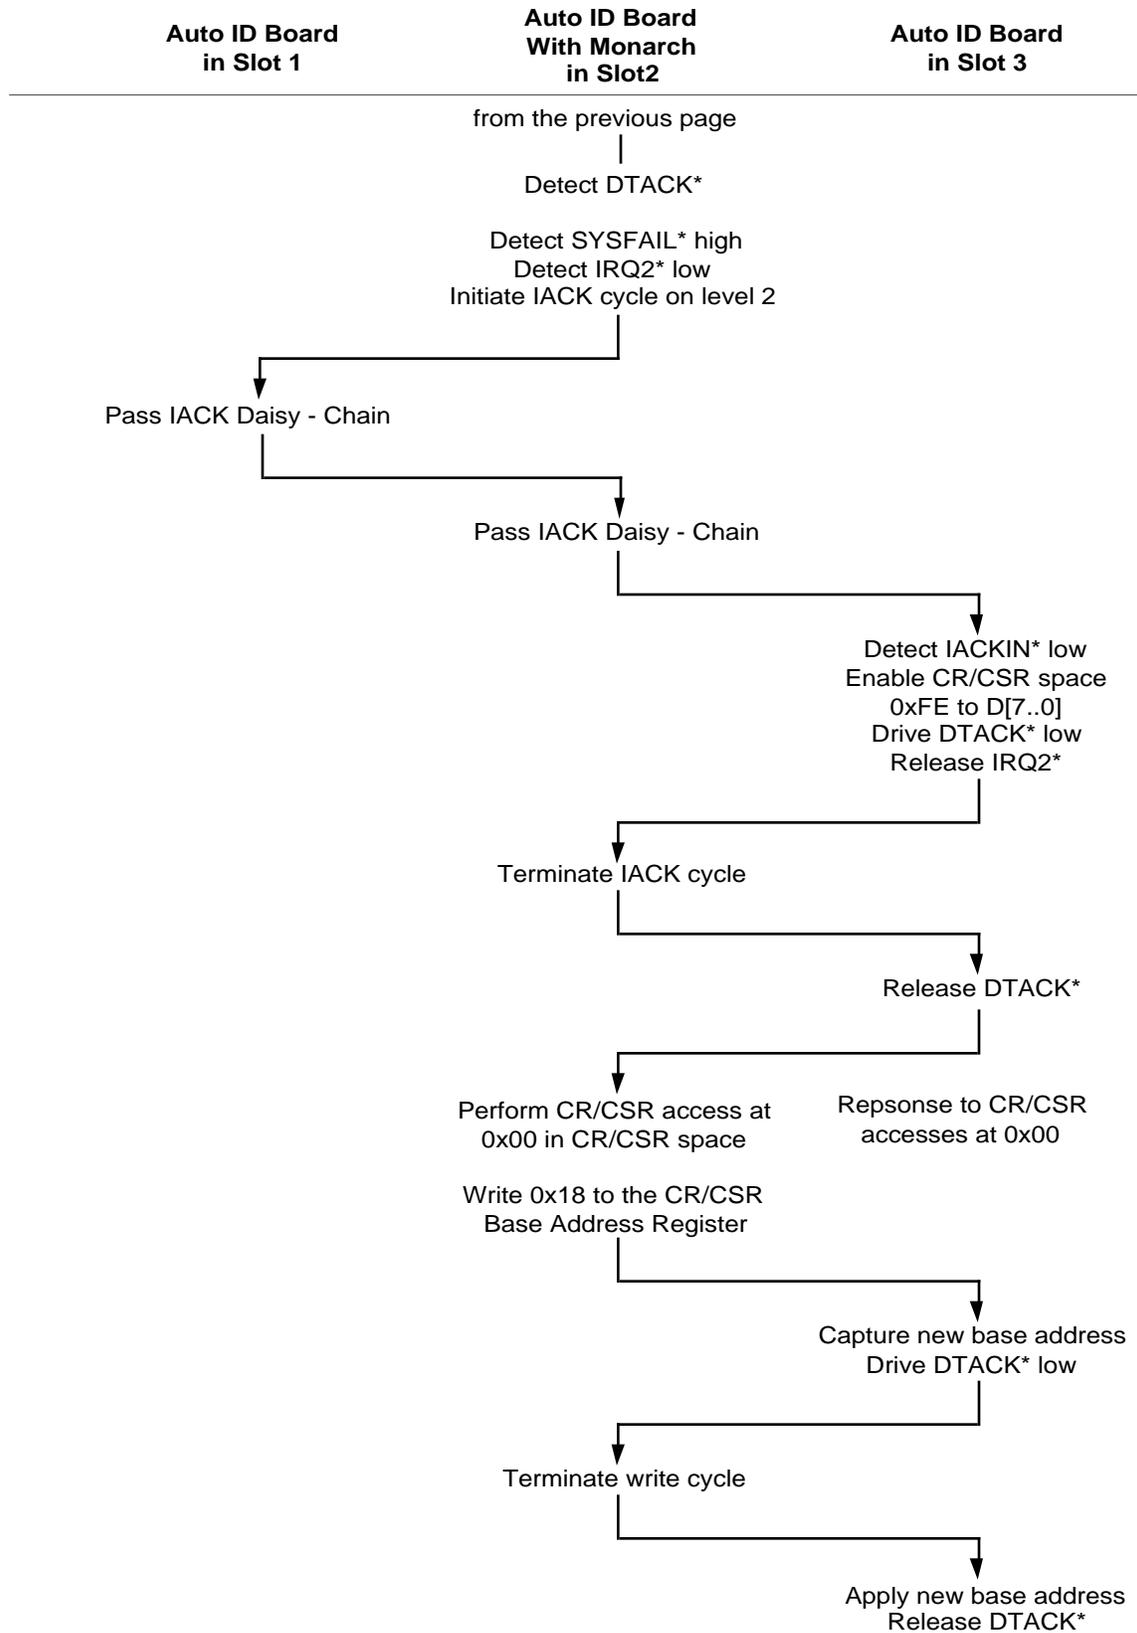


Figure 5 - 8 CR/CSR Auto ID Slave Initialization Algorithm, 3 of 3 pages

5.8 Auto System Controller

VMEbus requires that the first board in the system be configured as the System Controller. In some systems, the System Controller functionality is supplied by a separate board. In other systems, the System Controller functionality is supplied as part of the CPU boards. When included as part of a CPU board, the System Controller functions are usually enabled or inhibited by means of an on-board jumper. Configuring jumpers is a source of error in a system. The First Slot Detector (FSD) module allows a board to determine if it should enable its System Controller functions without the need for jumpers.

The typical Backplane Interface Logic on an Auto System Controller board connects a 10Kohm resistor between BG3IN* and ground. During power-up, the FSD keeps SCON deasserted while the system voltages ramp up towards their operating potential. When the on-board logic power reaches 4.75 volts, the on-board Power Monitor asserts POWER INSPEC. The FSD delays POWER INSPEC for forty milliseconds to allow for board-to-board Power Monitor variations, then latches the state of BG3IN*. If BG3IN* is low, the FSD asserts SCON to enable the System Controller functions. If BG3IN* is high, the FSD keeps SCON deasserted to maintain the on-board System Controller in the inhibited state.

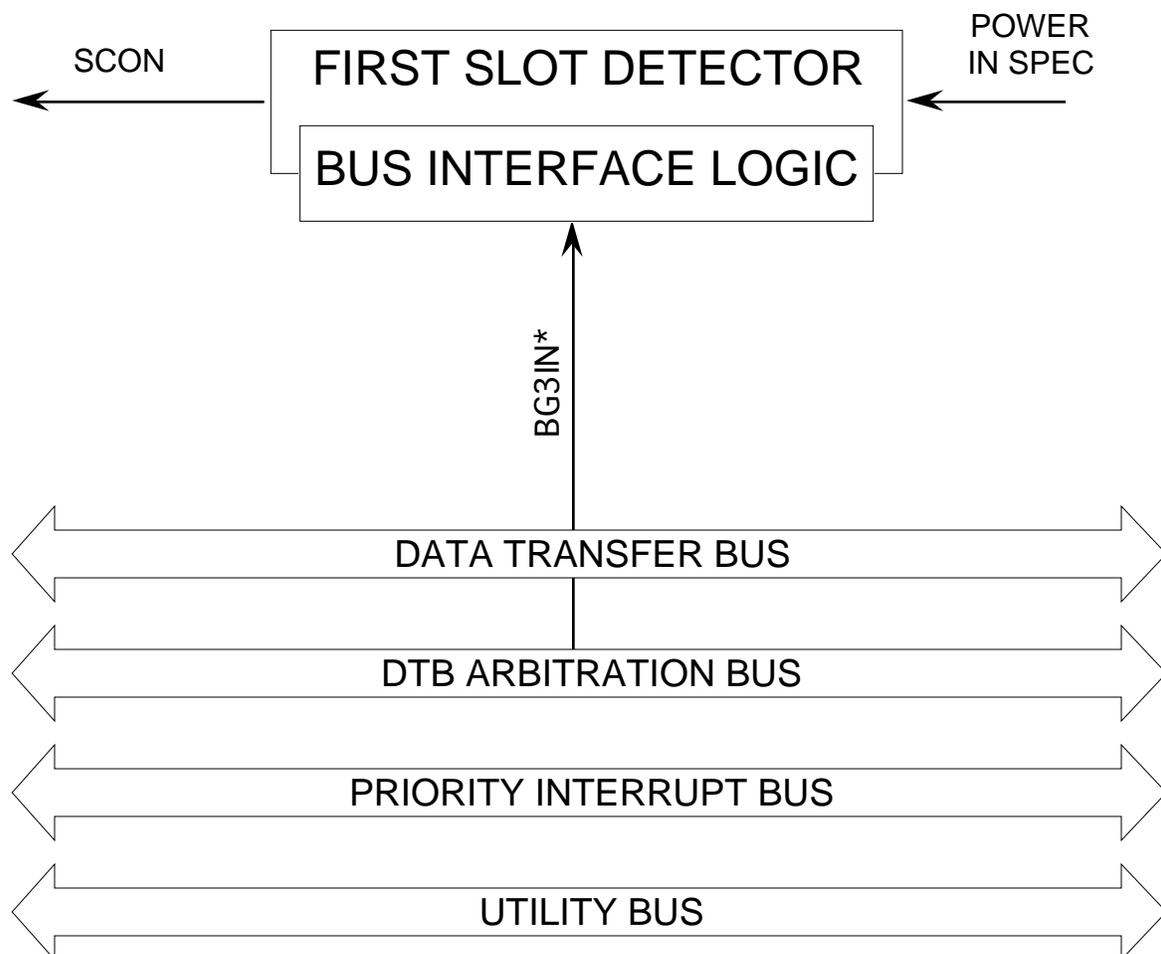


Figure 5 - 9 FIRST SLOT DETECTOR (FSD)

RULE 5.19:

Autoconfigured System Controller **MUST NOT** drive BG[3..0]IN*, or SYSCLK until the First Slot Detector (FSD) asserts SCON indicating that this board is the System Controller.

RULE 5.20:

An Autoconfigured System Controller **MUST** follow the power-up sequence requirements shown in Table 5 - 1.

RULE 5.21:

An Autoconfigured System Controller's Backplane Interface Logic **MUST** have a pull down on BG3IN*. (See Suggestion 6.9).

OBSERVATION 5.14:

It is necessary to use BG3IN* for the First Slot Detector because BG3IN* is the only bus grant input signal associated with a minimum function Arbiter.

RULE 5.22:

Once an Autoconfigured System Controller becomes the System Controller, it **MUST** remain the System Controller until after SYSRESET* is detected low or the system is powered down. If an Autoconfigured System Controller does not become the System Controller after powerup, it **MUST** remain in the non-System Controller state until after SYSRESET* is detected low or the system is powered down.

Chapter 6

Electrical Specifications

6.1 Introduction

The transmission of data between VMEbus boards such as processors, memories and I/O devices takes place over one or two backplanes, depending on the design. The rules in this chapter ensure proper timing, minimal noise and minimal cross talk problems on the backplane signal lines. The design of VMEbus backplanes is governed by the following RULEs:

RULE 6.1:

VMEbus backplanes **MUST NOT** have any signal conductors longer than 500 mm (19.68 inches).

RULE 6.2:

VMEbus backplanes **MUST NOT** have more than 21 slots.

RULE 6.3:

For those lines requiring termination (see Section 6.7) the backplane **MUST** provide some means for terminating them at both ends of the signal line.

RULE 6.4:

The backplane **MUST** provide power conductors for distribution of +5V, +5 STDBY, +12V, and -12V to all of the power pins specified in Section 7.7.

RULE 6.5:

The backplane **MUST** provide ground connections to all of the ground pins specified in Section 7.7.

PERMISSION 6.1:

VMEbus signal lines are normally driven by bipolar drivers, but any technology which complies with this specification **MAY** be used.

6.2 Power Distribution

Power in a VMEbus system is distributed on the backplane(s) as regulated direct current (DC) voltages. The available voltages are:

+5 VDC

This is the main power source for most VMEbus systems. Most of the system circuitry, including TTL logic, MOS microprocessors, and memories, requires this voltage.

- ± 12 VDC These are often used for powering RS232C drivers. They are also sometimes used for powering MOS and analog devices. In some cases -5 VDC bias voltage or -5.2 VDC ECL voltages are also derived from the -12VDC source using on-board regulators. These supplies normally don't supply as much power to the VMEbus system as the +5VDC source.
- +5 VDC STDBY This is used to sustain memory, time-of-day clocks, etc., when the +5VDC power is lost.

6.2.1 DC Voltage Specifications

Table 6-1 summarizes the DC voltage specifications. The listed specifications are the maximum allowed variance as measured at the connector pins of any card plugged into the backplane.

RECOMMENDATION 6.1:

Design and connect backplanes so that the power supply sense point is located somewhere near the center of the backplane, and as close as possible to the point where power is introduced into the backplane.

OBSERVATION 6.1:

Placing the power supply sense point near the power input point prevents boards near the power input point from receiving too high a voltage.

Table 6 - 1 Bus Voltage Specification

MNEMONIC	DESCRIPTION	ALLOWED VARIATION (see OBSERVATION 6.2)	RIPPLE/NOISE BELOW 10 MHz (Peak-to-Peak)
+5V	+5 VDC	+0.25V / -0.125V	50 mV
+12V	+12 VDC power	+0.60V / -0.36V	50 mV
-12V	-12 VDC power	-0.60V / +0.36V	50 mV
+5V STDBY	+5 VDC standby	+0.25V / -0.125V	50 mV
GND	Ground	REFERENCE	

OBSERVATION 6.2:

The non-symmetric variation given in Table 6-1 ensures that the DC power remains within the tolerance required by most ICs despite the typical voltage drops that occur in the power distribution network.

OBSERVATION 6.3:

The power consumed by some systems fluctuates over a wide range during normal system operation. For example, dynamic memory refreshing might cause significant fluctuations if large amounts of memory are refreshed at one time. In this case the response time of the voltage distribution system becomes important.

RECOMMENDATION 6.2:

Use bypass capacitors on VMEbus boards to minimize the effects of power transients.

6.2.2 Pin and Socket Connector Electrical Ratings

RULE 6.6:

The 96 pin connector used by the VMEbus MUST provide the following:

Voltage rating:	>	100 volts DC, isolation pin to pin
Contact resistance:	<	50 milliohms, at rated current
Insulation resistance:	>	100 Megaohms, pin to pin

6.3 Electrical Signal Characteristics

RULE 6.7:

VMEbus boards MUST NOT drive any backplane signal line to a higher steady-state voltage than the highest voltage on any of its +5V power pins, or to a lower steady-state voltage than the lowest voltage on any of its GND pins.

RULE 6.8:

VMEbus boards MUST use drivers and receivers that meet the following characteristics:

Steady-state driver	low output level	< 0.6 V
Steady-state receiver	low input level	< 0.8 V
Steady-state driver	high output level	> 2.4 V
Steady-state receiver	high input level	> 2.0 V

Figure 6-1 gives a simple graphic representation of these levels.

VMEbus boards drive the backplane lines with three-state, open collector, and totem-pole drivers. Section 6.4 specifies the drive and loading requirements for the various signal lines. Section 6.7 provides a summary, showing which types of drivers are used to drive each signal line.

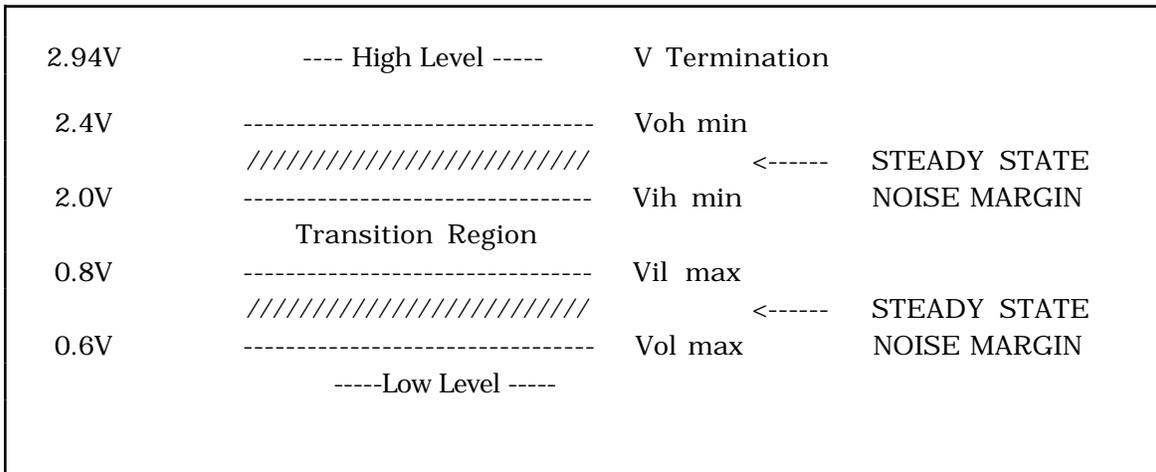


Figure 6 - 1 VMEbus Signal Levels

RULE 6.9:

When making voltage threshold measurements on a VMEbus board to verify compliance with timing specifications, the ground reference **MUST** be taken from the board's ground pin nearest the signal pin being measured, and the signal voltage **MUST** be measured on the board's connector pin.

6.4 Bus Driving and Receiving Requirements

This section defines the driver and receiver specifications for all VMEbus signal lines. Table 6-2 lists all of the signals and shows which of the following sub-sections discuss it.

6.4.1 Bus Driver Definitions

Totem-pole, three-state, and open-collector drivers are defined as follows:

Totem-pole — an active driver in both states which sinks current in the low state and sources current in the high state. Totem-pole drivers are used on signals having only a single driver per line (e.g., daisy-chain lines).

Three-state — similar to a totem-pole driver except that it can go to a high impedance state (drivers turned off) in addition to the low and high logic states. Three-state drivers are used for lines that can be driven by several devices at different points on the bus (e.g., address or data lines). Only one of these drivers can be active at any one time.

Rescinding -- a three state driver that is driven high before being released to a high impedance.

Open-collector — sinks current in the low state but sources no significant current in the high state. Terminating resistors on the backplane ensure that the signal line voltage rises to a high level whenever it is not driven low. Open-collector drivers are used for signal lines which can be driven by several devices simultaneously (e.g., interrupt and bus request lines).

Table 6 - 2 Bus Driving And Receiving Requirements

SIGNAL NAME	SUB-SECTION 6.4.2.x
A[31..1]	2
ACFAIL*	5
AM[5..0]	2
AS*	1
BBSY*	5
BCLR*	3
BERR*	5
BG[3..0]OUT*	4
BG[3..0]IN*	4
BR[3..0]*	5
D[31..0]	2
DS0*	1
DS1*	1
DTACK*	1, 5
IACK*	2, 5
IACKOUT*	4
IRQ[7..1]*	5
LWORD*	2
RETRY*	1
SYSCLK	3
SYSFAIL*	5
SYSRESET*	5
WRITE*	2

6.4.2 Driving And Loading RULEs For All VMEbus Lines

RULE 6.10:

All VMEbus boards **MUST** provide clamping on each VMEbus signal line that they monitor to prevent negative excursions below -1.5 V.

OBSERVATION 6.4:

Standard 74LSxxx and 74Fxxx devices have internal clamping diodes on their inputs that will satisfy the clamping requirement specified in RULE 6.10.

RULE 6.11:

VMEbus receivers **MUST** guarantee detection of a high logic level above a threshold of 2.0 volts, as shown in Figure 6 - 1.

RULE 6.12:

VMEbus receivers **MUST** guarantee detection of a low logic level below a threshold of 0.8 volts, as shown in Figure 6 - 1.

PERMISSION 6.2:

A three-state driver **MAY** be used as a totem-pole driver if its output is permanently enabled.

6.4.2.1 Driving And Loading RULEs For High Current Three-State Lines (AS*, DS0*, DS1*, Rescinding DTACK*, RETRY*)

RULE 6.13:

IF a VMEbus board drives AS*, DS0*, DS1*, or Rescinding DTACK*

THEN its drivers for these lines MUST meet the following specifications:

Low state sink current:	IOL	>	64 mA
Low state voltage:	VOL	<	0.6V at IOL = 64 mA
High state source current:	IOH	>	3 mA
High state voltage:	VOH	>	2.4V at IOH = 3 mA
Minimum source current with board pin grounded:	IOS	>	50 mA at 0V
Maximum source current with board pin grounded:	IOS	<	225 mA at 0V

RULE 6.14:

When drivers are turned off, VMEbus boards MUST limit their loading of AS*, DS0*, and DS1* to the following values:

Current sourced by board at 0.6 V, including leakage current:	IOZL+IIL	<	450 uA
Current sunk by board at 2.4 V, including leakage current:	IOZH+IIH	<	100 uA
Total capacitive load on signal, including signal trace:	CT	<	20 pF

OBSERVATION 6.17

The rescinding DTACK* has the drive characteristics specified in RULE 6.13 and the loading characteristics specified in RULE 6.22. The control of the 3 state enable for the rescinding DTACK* is specified by RULE 2.96.

OBSERVATION 6.5:

The source and sink currents listed in RULEs 6.13 and 6.14 include both driver and receiver currents sourced and sunk on the board.

6.4.2.2 Driving And Loading RULEs For Standard Three-State Lines (A[31..1], D[31..0], AM[5..0], IACK*, LWORD*, WRITE*)

RULE 6.15:

IF a VMEbus board drives the lines A[31..1], D[31..0], AM0[5..0], IACK*, LWORD*, or WRITE*,

THEN its drivers for these lines MUST meet the following specifications:

Low state Sink current:	IOL	>	48 mA
Low state voltage:	VOL	<	0.6V at IOL = 48 mA
High state source current:	IOH	>	3 mA
High state voltage:	VOH	>	2.4V at IOH = 3 mA
Minimum source current with board pin grounded:	IOS	>	50 mA at 0V
Maximum source current with board pin grounded:	IOS	<	225 mA at 0V

RULE 6.16:

When drivers are turned off, VMEbus boards MUST limit their loading of the lines A[31..1], D[31..0], AM[5..0], IACK*, LWORD*, and WRITE* to the following values:

Current sourced by board at 0.6 V, including leakage current:	IOZL+IIL	<	700	uA
Current sunk by board at 2.4 V, including leakage current:	IOZH+IIH	<	150	uA
Total capacitive load on signal, including signal trace:	CT	<	20	pF

OBSERVATION 6.6:

The source and sink currents specified in RULEs 6.15 and 6.16 include both driver and receiver currents sourced and sunk on the board.

6.4.2.3 Driving And Loading RULEs For High Current Totem-Pole Lines (SYSCLK, BCLR*)

RULE 6.17:

VMEbus systems MUST have no more than one board driving each of the lines SYSCLK, or BCLR*. Its drivers for these lines MUST meet the following specifications:

Low state sink current:	IOL	>	64	mA
Low state voltage:	VOL	<	0.6V at IOL = 64	mA
High state source current:	IOH	>	3	mA
High state voltage:	VOH	>	2.4V at IOH = 3	mA
Minimum source current with board pin grounded	IOS	>	50	mA at 0 V
Maximum source current with board pin grounded	IOS	<	255	mA at 0V

RULE 6.18:

All VMEbus boards MUST limit their loading of the lines SYSCLK, and BCLR* to the following values:

Current sourced by board at 0.6 V, including leakage current:	IOZL+IIL	<	600	uA
Current sunk by board at 2.4 V, including leakage current:	IOZH+IIH	<	50	uA
Total capacitive load on signal, including signal trace, for system controllers (which have drivers).	CT	<	20	pF
Total capacitive load on signal, including signal trace, for other boards (which have no drivers).	CT	<	12	pF

OBSERVATION 6.7:

The source and sink currents specified in RULEs 6.17 and 6.18 include both driver and receiver currents sourced and sunk on the board.

6.4.2.4 Driving And Loading RULEs For Standard Totem-Pole Lines.

(BG[3..0]OUT*/BG[3..0]IN*, IACKOUT*/IACKIN*)

RULE 6.19:

IF a VMEbus board drives the lines BG[3..0]OUT*/BG[3..0]IN*, or IACKOUT*/IACKIN*,

THEN its drivers for these lines MUST meet the following specifications:

Low state sink current:	IOL	>	8 mA
Low state voltage:	VOL	<	0.6V at IOL = 8 mA
High state source current:	IOH	>	400 uA
High state voltage:	VOH	>	2.7V at IOH = 400 uA

RULE 6.20:

All VMEbus boards MUST limit their loading of each of the lines BG[3..0]OUT*/BG[3..0]IN* , and IACKOUT*/IACKIN* to the following values:

Current sourced by board at 0.6 V, including leakage current:	IOZL+IIL	<	600 uA
Current sunk by board at 2.4 V, including leakage current:	IOZH+IIH	<	300 uA
Total capacitive load on signal, including signal trace:	CT	<	20 pF

OBSERVATION 6.8:

The source and sink currents specified in RULEs 6.19 and 6.20 include both driver and receiver currents sourced and sunk on the board.

OBSERVATION 6.18:

BG[x]IN* sink current has been increased to 300uA to accommodate the Auto System Controller function (see Section 5.8). Standard VMEbus BGOUT drivers which provide 400uA of source current are fully compliant to this new specification.

SUGGESTION 6.9:

For Auto System Controller boards use a $\geq 10K$ ohm to tie an open BG3IN* to logic low.

Use $\leq 40uA$ Iih receivers to receive BG3IN*

6.4.2.5 Driving And Loading RULEs For Open-Collector Lines

(BR[3..0]*, BBSY*, IRQ[7..1]*, DTACK*, BERR*, SYSFAIL*, SYSRESET*, ACFAIL*, IACK*)

RULE 6.21:

IF a VMEbus board drives the lines BR[3..0]*, BBSY*, IRQ[7..1]*, DTACK*, BERR*, SYSFAIL*, SYSRESET*, ACFAIL*, IACK*,

THEN its drivers for these lines MUST meet the following specifications:

Low state sink current:	IOL	>	48 mA
Low state voltage:	VOL	<	0.6V at IOL = 48 mA

OBSERVATION 6.19:

The driver specification for rescinding DTACK* is specified in RULE 6.13.

RULE 6.22:

All VMEbus boards MUST limit their loading of the lines BR[3..0]*, BBSY*, IRQ[7..1]*, DTACK*, BERR*, SYSFAIL*, SYSRESET*, ACFAIL*, IACK* to the following values:

Current sourced by board at 0.6 V, including leakage current:	IOZL+IIL	<	400 uA (DTACK* and BERR*) < 600 uA (all others)
Current sunk by board at 2.4 V, including leakage current:	IOZH+IIH	<	50 uA
Total capacitive load on signal, including signal trace:	CT	<	20 pF

OBSERVATION 6.9:

The sink current specified in RULEs 6.21 and 6.22 includes both driver and receiver currents sourced and sunk on the board.

SUGGESTION 6.6:

Since most TTL drivers do not work reliably when the +5 VDC power source is out of its specified range, drive SYSRESET* on a POWER MONITOR module with a driver built from a discrete high-gain small signal transistor.

6.5 Backplane Signal Line Interconnections

The VMEbus is a high performance interface system. Its design takes into account transmission line effects on the backplane. The address and data set-up times specified in Chapters 2 and 4 take into account the fact that most drivers available today do not reliably drive backplane signal lines from the low to the high level until there is a reflection from the end of the bus. Although these reflections serve a useful purpose, they cannot be excessive or ringing will result. The following paragraphs specify the backplane characteristics that achieve the desired result.

6.5.1 Termination Networks

RULE 6.23a:

Termination networks **MUST** be used on each end of all VMEbus signal lines except the daisy-chain lines and the Serial Bus lines.

OBSERVATION 6.10:

The terminations in the VMEbus serve four purposes:

They reduce reflections from the ends of the backplanes.

They provide a high state pull-up for open-collector drivers.

They restore the signal lines to the high level when three-state devices are disabled.

They provide a standing current for the driver sink transistor to switch off, causing the signal line to rise more swiftly on positive transitions.

The Thevenin equivalent of the termination is shown in Figure 6-2. The voltage divider also shown provides this termination value.

OBSERVATION 6.11:

IF a maximum tolerance of $\pm 5\%$ is maintained on the resistor values and source voltage used in the resistor network shown in Figure 6-2.

THEN the circuit shown will meet the tolerances shown for the Thevenin equivalent.

OBSERVATION 6.12:

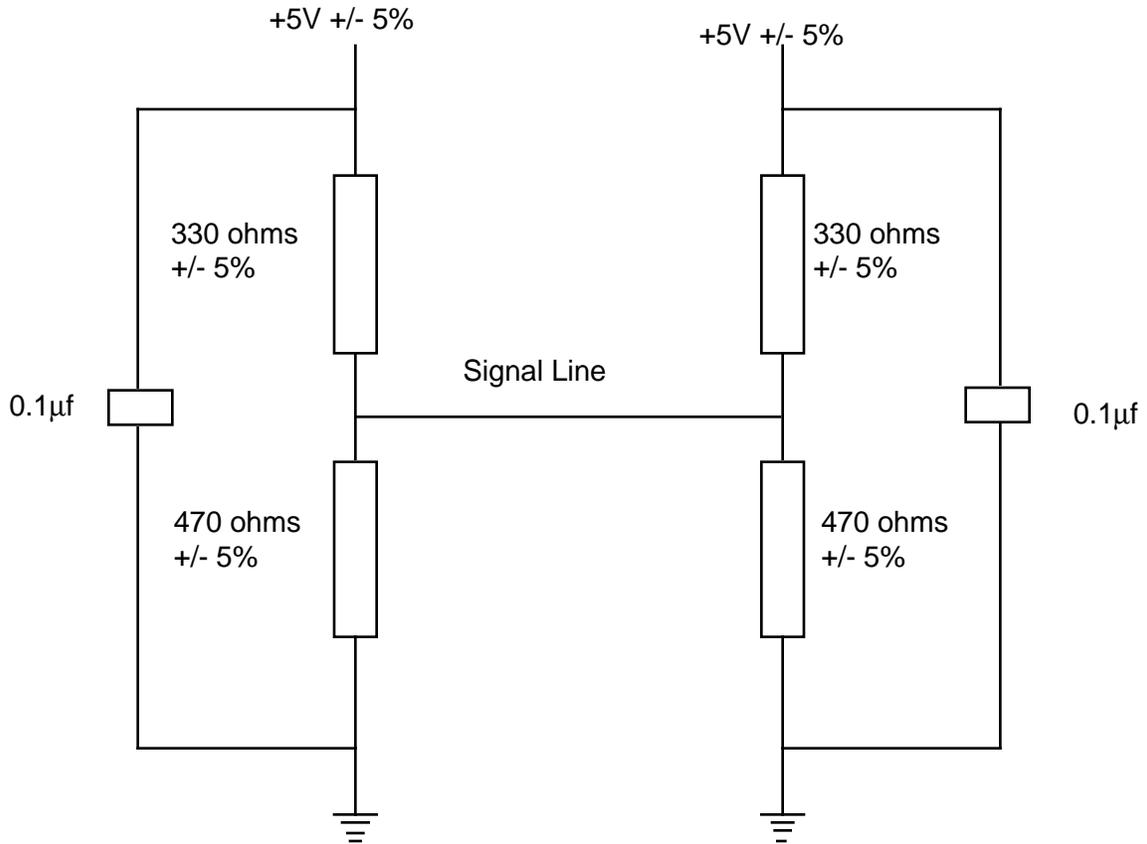
The resistor network shown in Figure 6-2 presents its Thevenin equivalent impedance only when its +5V source is adequately decoupled to ground by a bypass capacitor.

RECOMMENDATION 6.3:

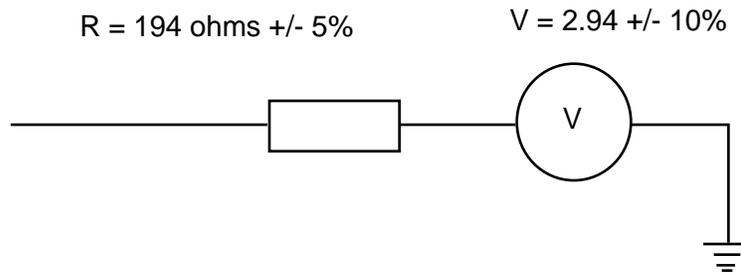
Provide a bypass capacitor with a value in the range of 0.01 to 0.1 μF as close as possible to the Vcc pin of each resistor termination package.

PERMISSION 6.3:

Any resistor network and voltage source **MAY** be used to provide the termination, as long as they provide the Thevenin equivalent shown in Figure 6-2.



RESISTOR NETWORKS THAT PROVIDE THE REQUIRED TERMINATION



THEVENIN EQUIVALENT FOR EACH NETWORK

Figure 6 - 2 Standard Bus Termination

6.5.2 Characteristic Impedance

Each signal line in the backplane has an associated characteristic impedance Z_0 . This characteristic impedance is important because discontinuities in Z_0 (due to capacitive effect and loads on the bus) and mismatches between Z_0 and the terminations can cause distortions of signal waveforms.

More information on characteristic impedance and backplane design can be found in the MECL System Design Handbook, Motorola, 1983.

The terminations on the VMEbus signal lines reduce distortion of their signal waveforms. Although a perfect impedance match (which totally eliminates distortions due to reflections) is not maintained between the termination networks and the signal lines, it is important not to allow too great a mismatch, as might be the case if a signal line's Z_0 value is too low.

The actual characteristic impedance of a backplane signal line is called the effective characteristic impedance (Z_0'), and will be lower than Z_0 , due to the capacitance of plated through holes and connector pins. This additional capacitance makes Z_0' go below Z_0 . Although plated-through holes are necessary to accommodate connectors, other holes should be kept to a minimum.

The backplane signal line impedance (without any boards plugged into the backplane) can be calculated with the following equation:

$$Z_0' = \frac{Z_0}{\sqrt{1 + Cd/Co}}$$

where:

Z_0 = The impedance of the microstrip line, ignoring the loading effects of plug-in pc boards, connectors, and plated-through holes.

Cd = The distributed capacitance, per unit of distance, of the plated-through holes, and backplane connectors.

Co = The intrinsic line capacitance, per unit of distance, of the microstrip line, ignoring the loading effects of plug-in pc boards, connectors, and plated-through holes.

Z_0' = The backplane signal line impedance, including the loading effects of connectors, and plated-through holes but excluding the loading effects of plug-in pc boards.

OBSERVATION 6.13:

Typical Z_0' values for a VMEbus backplane, with no boards inserted, range from 50 to 60 ohms. If this impedance is 50 ohms, or higher, it will provide satisfactory operation.

6.5.3 Additional Information

RULE 6.24:

All circuit traces from the 96 pin connectors to the on-board circuitry except for signals ACFAIL*, SYSRESET*, SYSFAIL*, IRQ[7..0] MUST NOT have a length of greater than 50.8 mm (2 inches) . Circuit traces from the 96 pin connectors to the on-board circuitry for signals ACFAIL*, SYSRESET*, SYSFAIL*, IRQ[7..0] MUST NOT have a length of greater than 101.6mm (4 inches) .

OBSERVATION 6.14:

IF the trace from the 96 pin connector to on-board circuitry branches,
THEN the length of each branch is added to get the total length specified in
RULE 6.24.

RULE 6.25:

There MUST NOT be more than one driver driving the SYSCLK line.

RULE 6.26:

The SYSTEM CLOCK DRIVER module MUST be installed in Slot 1 of the backplane.

OBSERVATION 6.15:

Locating the SYSTEM CLOCK DRIVER on the board in Slot 1 minimizes the distortion of their waveforms reflected from the terminated ends of the backplane.

SUGGESTION 6.7:

If actual capacitance loading values cannot be obtained from manufacturer specifications sheets, the following values can be used to estimate the total capacitive loading of a VMEbus board:

The typical capacitance of a receiver	3	-	5 pF
The typical capacitance of a driver	10	-	12 pF
The typical capacitance of a transceiver	15	-	18 pF
The typical capacitance of a 50.8 mm (2 inch) PC trace	2	-	3 pF

OBSERVATION 6.16:

Circuit traces which run parallel to each other, such as in a backplane, sometimes induce signal transitions in each other. This phenomenon is commonly known as cross talk. When designing VMEbus backplanes, the spacings of lines and their position relative to ground and power planes have a large effect on the amount of cross talk observed.

SUGGESTION 6.8:

Propagation delays through bus drivers depend on how heavily they are loaded and VMEbus signal lines typically represent heavy loads. This has to be taken into account when calculating worst case timing. If the manufacturer's data sheet for the driver gives a propagation delay for a 300 pF load, use that to do the worst case calculations. If the only propagation delay values are for a 30 pF load, add 10 ns to the propagation delay and 15 ns to the turn-on delay.

6.6 User Defined Signals

RECOMMENDATION 6.5:

If a board has a 96 pin connector in its P2 location, do not allow any of the P2 pins to be driven to a voltage greater than $\pm 15V$. This reduces the likelihood of serious damage to the VMEbus system in the event that a signal trace from one of these pins is accidentally shorted to some other signal line.

6.7 Signal Line Drivers and Terminations

This section summarizes the types of drivers which have to be used for each of the signal lines on the VMEbus.

In order to simplify Table 6-3, an abbreviated notation is used to describe the various types of drivers. The notations used are shown below:

Totem-pole (high current)	-	TP HC
Totem-pole (standard)	-	TP STD
Three-state (high current)	-	3 HC
Three-state (standard)	-	3 STD
Open-collector	-	OC

For detailed specifications, see Section 6.4.

Table 6 - 3 Bus Driver Summary

SIGNAL MNEMONIC	SIGNAL NAME	DRIVER TYPE	BUSED AND TERMINATED
A[31..1] (31 lines)	ADDRESS BUS	3 STD	YES
ACFAIL*	AC POWER FAILURE	OC	YES
AM[5..0] (6 lines)	ADDRESS MODIFIER	3 STD	YES
AS*	ADDRESS STROBE	3 HC	YES
BBSY*	BUS BUSY	OC	YES
BCLR*	BUS CLEAR	TP HC	YES
BERR*	BUS ERROR	OC	YES
BG[3..0]IN* BG[3..0]OUT* (Daisy-chain)	BUS GRANT DAISY- CHAIN	TP STD	NO

Table 6 - 3 Bus Driver Summary (Continued)

SIGNAL MNEMONIC	SIGNAL NAME	DRIVER TYPE	BUSED AND TERMINATED
BR[3..0]* (4 lines)	BUS REQUEST	OC	YES
D[31..0] (32 lines)	DATA BUS	3 STD	YES
DS0*-DS1* (2 lines)	DATA STROBES	3 HC	YES
DTACK*	DATA TRANSFER ACKNOWLEDGE	OC 3 HC	YES
IACK*	INTERRUPT ACKNOWLEDGE	3 STD or OC	YES
IACKIN*/IACKOUT* (Daisy-chain)	INTERRUPT ACKNOWLEDGE DAISY-CHAIN	TP STD	NO
IRQ[7..1]* (7 lines)	INTERRUPT REQUEST	OC	YES
LWORD*	LONGWORD	3 STD	YES
RETRY*	RETRY	3 HC	YES
SER_A	SERIAL BUS	Note 1	Note 2
SER_B	SERIAL BUS	Note 1	Note 2
SYSCLK	SYSTEM CLOCK	TP HC	YES
SYSFAIL*	SYSTEM FAILURE	OC	YES
SYSRESET*	SYSTEM RESET	OC	YES
WRITE*	WRITE	3 STD	YES

Notes:

- 1) Driver types are defined by the serial bus specification
- 2) Terminations are defined by the serial bus specification

PAGE LEFT BLANK

Chapter 7

Mechanical Specifications

7.1 Introduction

Information is provided in this chapter to ensure that the VMEbus board assemblies, backplanes, subracks, and associated mechanical accessories are dimensionally compatible.

The mechanical dimensions given in this chapter conform to IEC publications 297-1, 297-3, 297-4 and 603-2. The electrical characteristics for VMEbus connectors, as specified in Chapters 5 and 6 supersede publication 603-2 where they differ.

Two connector types are permitted in this specification, one for 96 pins configurations and one for 160 pin¹ configurations. For 96 pin configurations use the IEC publication 603-2 which describes a family of connector types which are identified by labels of the form: 603-2-IEC-xxxxxx-xxx. All of the P1/J1 and P2/J2 connectors used on VMEbus boards and backplanes are members of the 603-2-IEC family. In this chapter, the label 603-2-IEC-xxxxxx-xxx is used when referring to all of these connector types as a group. The label 603-2-IEC-C096Mx-xxx is used when referring to the male connector types within these families which are used on VMEbus boards. 603-2-IEC-C096Fx-xxx is used when referring to the female connector types which are used on VMEbus backplanes.

The drawings in this specification show the 160 pin family of connectors. These connectors have 5 rows of contacts. Use the center three rows when applying the drawings to 96 pin family of connectors.

Figure 7-1 is a front view of a 19-inch wide subrack and shows how single height and double height VMEbus boards can be mixed in a single subrack. Boards are inserted into the subrack from the front, in a vertical plane with the component face of the board on the right.

PERMISSION 7.1:

A VMEbus system MAY be composed of single height boards, double height boards, or a mixture of both.

RULE 7.1:

Single height VMEbus subracks MUST have a single J1 backplane.

RULE 7.2:

Double height VMEbus subracks MUST have

- 1) a J1 backplane mounted in the upper portion of the subrack,
- OR
- 2) a J1 and a J2 backplane, with the J1 backplane mounted in the upper portion and the J2 backplane mounted in the lower portion.
- OR
- 3) a double height backplane which provides both J1 and J2 connectors.

¹At the time this document was printed the 160 pin connector was being standardized in the IEC 1076-4 family of connectors rather than the 603-2 family. In the rest of this chapter when the 160 pin connector is referenced, the reader should refer to the IEC 1076 standard rather than the IEC 603-2 standard. Additionally, when the text refers to the 603 style connector in a general way the reader should assume a reference to the 1076 connector is intended also.

RULE 7.3:

VMEbus backplanes MUST NOT have more than 21 slots.

PERMISSION 7.2:

When using backplanes of fewer than 21 slots, the VMEbus subrack MAY be less than the standard 482.6 mm (19 inches) rack size.

RULE 7.4:

Except for the rack width, which varies with the number of slots, all subrack dimensions MUST agree with those given in this chapter to ensure mechanical compatibility between the boards and the subrack.

7.2 VMEbus Boards

RECOMMENDATION 7.1:

Make VMEbus boards 1.6 ± 0.2 mm (0.063 ± 0.008 inch) thick.

OBSERVATION 7.1:

The thickness of VMEbus boards is important because the subrack's guide rails are designed to accommodate boards of this thickness. Thicker boards might not fit into the guides and thinner boards might not be guided properly into the backplane's J1 and J2 connectors.

OBSERVATION 7.2:

The dimensioning of the 603-2-IEC-xxxxxx-xxx connectors provides a certain distance between the connector's mounting face and the center line of each of the connector's pins. This ensures that the P1 and P2 connector pin centers will align properly with the J1 and J2 connectors of the backplane.

PERMISSION 7.3:

VMEbus boards MAY be designed with board thicknesses greater than 1.6 mm (0.063 inch) if:

1. the thickness of the top and bottom edges of the board, which fit into the guide rails, is reduced to 1.6 mm (0.63 inch) for distance of 2.5 mm (0.098 inch) from the top and bottom edges of the board (See Figure 7-2 and Figure 7-3) and,
2. the mounting surface provided by the board for the IEC 603-2 connector(s) is 4.07 mm (0.160 inch) from the interboard separation plane. (See Figure 7-5)

Two board sizes are defined as standard VMEbus boards: single height and double height (see Figure 7-2 and Figure 7-3).

7.2.1 Single Height Boards

OBSERVATION 7.3:

A single height VMEbus board is 100 mm (3.937 inch) high and 160 mm (6.299 inch) deep with an area of approximately 160.0 sq. cm (24.8 sq. inch).

RULE 7.5:

All single height boards MUST be designed according to the dimensions given in Figure 7-2.

RULE 7.6:

The hole pattern for the 160 pin P1 connector **MUST** be as shown in Figure 7-2. The hole pattern for the 96 pin P1 connector **MUST** be the center three rows of the pattern shown in Figure 7-2.

SUGGESTION 7.1:

Use the PC layout grid shown in Figure 7-2.

OBSERVATION 7.4:

Modular front panel hardware is available from several manufacturers which, when mounted on the grid shown in Figure 7-2, properly aligns with the front panel grid shown in Figure 7-7.

PERMISSION 7.4:

Components, other than the 603-2-IEC-CO96Mx-xxx connector, **MAY** be placed in such a way that they do not align with the grid pattern.

7.2.2 Double Height Boards**OBSERVATION 7.5:**

A double height VMEbus board is 233.35 mm (9.187 inch) high and 160 mm (6.299 inch) deep with an area of approximately 373.4 sq. cm (57.9 sq. inch).

RULE 7.7:

All double height boards **MUST** be designed according to the dimensions given in Figure 7-3.

RULE 7.8:

IF a 96 pin or 160 pin 603-2-IEC-C096Mx-xxx connector is used for P2
THEN the hole pattern for the 160 pin P2 connector **MUST** be as shown in Figure 7-3 and the hole pattern for the 96 pin P2 connector **MUST** be the center three rows of the pattern shown in Figure 7-3.

RULE 7.9:

The hole pattern for the 160 pin P1 connector **MUST** be as shown in Figure 7-3. The hole pattern for the 96 pin P1 connector **MUST** be the center three rows of the pattern shown in Figure 7-3.

OBSERVATION 7.6:

As in the case of the single height boards above, the grid shown in Figure 7-3 properly aligns modular front panel components with the front panel grid shown in Figure 7-8.

PERMISSION 7.5:

Components, other than the 603-2-IEC-xxxxxx-xxx connectors, **MAY** be placed in such a way that they do not align with the grid pattern.

OBSERVATION 7.7:

There is a discontinuity of 1.27 mm (0.05 inch) between the 2.54 mm (0.10 inch) grid patterns for the upper and lower half of the board.

RULE 7.37

When a board is fitted with a connector that provides additional current paths to ground via the backplane, the electrically mating surfaces **MUST** be as defined in Figures 7.5 and 7.6. Mechanically compatibility with 603-2-IEC-C096Fx-xxx connectors **MUST** be maintained. The inductance of the ground path **MUST NOT** exceed 10 nanoHenries. The DC Resistance of the ground path **MUST NOT** exceed 10 milliohms. Minimum overall conformance **MUST** meet DIN41612; Class 2.

7.2.3 Board Connectors

The single height board has only one connector on its back edge. It is called the P1 connector. A double height board has either one or two connectors on its back edge. If it has one connector, that connector is called the P1 connector and is located on the upper half of the back edge. If it has two connectors the upper one is called the P1 connector and the lower one is called the P2 connector.

RULE 7.10:

The P1 and P2 connectors of all VMEbus boards **MUST** meet or exceed the IEC mechanical specifications of a 603-2-IEC-C096Mx-xxx class 2 connector, and **MUST** be mounted as shown in Figure 7-4.

OBSERVATION 7.8:

603-2-IEC class 2 connectors have a minimum mechanical endurance of 400 insertion/extraction cycles.

OBSERVATION 7.9:

The symmetry symbol in the box below each board outline in Figure 7-4 sets an upper limit on how much the connector's center line can be tilted with respect to the board's lower edge. This limit is described in the following rule.

RULE 7.11:

The perpendicular distance (d1) from the board's lower edge to the point A in Figure 7-4 **MUST NOT** differ from its perpendicular distance (d2) to point B by more than 0.3 mm (0.012 inch).

RULE 7.12:

IF a VMEbus board is designed to use the center row of P2 for address or data bus expansion, or the board requires more power than the P1 connector can provide,

THEN a 603-2-IEC-C096Mx-xxx P2 connector **MUST** be provided and it **MUST** be mounted as shown in Figure 7-4.

PERMISSION 7.6:

Where neither address nor data bus expansion is required, and where the board does not require more power than the P1 connector can provide, any 603-2-IEC-xxxxxx-xxx connector **MAY** be used for P2 on a double height VMEbus board or the board **MAY** be designed without a P2 connector.

PERMISSION 7.7:

On double height boards, the pins of rows a and c on the P2 connector **MAY** be used to provide user defined connections (See Section 7.7.2).

PERMISSION 7.8:

I/O cables **MAY** be connected to the front edge of a VMEbus board. No connectors are prescribed for these cable connections.

SUGGESTION 7.2:

Where possible, avoid the use of cable connections to the front edge of VMEbus boards. This makes it much easier to install and remove boards from the subrack during maintenance.

RULE 7.38:

IF a 96 pin 603-2-IEC-C096Mx-xxx is used for P1
THEN the pin assignments **MUST** be as shown in Table 7-1 using rows a, b and c.

RULE 7.39:

IF a 160 pin 603-2-IEC-C096Mx-xxx is used for P1
THEN the pin assignments **MUST** be as shown in Table 7-1 using rows z, a, b, c and d.

RULE 7.40:

IF a 96 pin 603-2-IEC-C096Mx-xxx is used for P2
THEN the pin assignments **MUST** be as shown in Table 7-2 using rows a, b and c.

RULE 7.41:

IF a 160 pin 603-2-IEC-C096Mx-xxx is used for P2
THEN the pin assignments **MUST** be as shown in Table 7-2 using rows z, a, b, c and d.

7.2.4 Board Assemblies

The board assembly typically consists of a PC board, as defined above with either one or two 603-2-IEC-xxxxxx-xxx connectors affixed to the board's back edge, electronic components, and an optional front panel with handles. For more detail on the front panels see Section 7.3.

RULE 7.13:

Solder filets, tracking, and components on VMEbus boards **MUST NOT** be closer than 2.5 mm (0.098 inch) from the top and bottom edges of the board to guarantee clearance between them and the board guides. Figure 7-2 and Figure 7-3 show these dimensions.

Figure 7-5 shows a cross sectional view of a PC board, its front panel, its connector and the backplane. The dimensions given are nominal values and are based upon the dimensions given in the other drawings in this chapter.

7.2.5 Board Widths

VMEbus boards designed to occupy a single slot of the subrack are called single width boards.

PERMISSION 7.9:

VMEbus boards **MAY** be designed to occupy more than one slot.

7.2.6 VMEbus Board Warpage, Lead Length and Component Height.

During the manufacturing process boards sometimes become warped.

RULE 7.14:

The sum of warpage and component lead length MUST be less than 2.47mm (0.097 inch) from where the solder side of an ideal (unwarped) board would be, and the sum of component height and warpage (in the other direction) MUST be less than 13.71mm (0.54 inch) plus an integral multiple (N) of 20.32mm (0.8 inch), from where the component side of an ideal (unwarped) board would be. (Where N = the number of slots that the board occupies minus 1.)

OBSERVATION 7.10:

During insertion into the subrack, the component leads of a board might contact the right edge of the front panel to its left. For this reason, the board should be inserted carefully to avoid bending the component leads.

SUGGESTION 7.3:

Where possible, trim the component lead lengths to less than 1.52 mm (0.06 inch). This makes installation and removal of boards more convenient. Since the maximum allowable warpage is affected by the length of these component leads, trimming them allows boards with larger warpage to meet the specifications.

OBSERVATION 7.11:

The dimensions given in SUGGESTION 7.3 ensure that there will be a clearance of at least 2.54 mm (0.1 inch) between components of each board and the component leads of the board to its right. This space allows adequate air flow and prevents additional warpage and vibration from causing interboard contact.

RULE 7.15:

All VMEbus boards MUST be measured, after they are assembled, to ensure that the combination of board warpage, component lead length, and component height do not exceed the specified limits when the board is inserted into a subrack. In order to properly make these measurements, the board MUST be placed in a subrack (or a similar test fixture) while the measurements are taken.

Figure 7-7 shows a board (single or double height) in a subrack, and shows how the combination of board warpage, component lead length, and component height are to be measured. The interboard separation planes defined in Figure 7-7 provide the reference from which the measurements are made.

OBSERVATION 7.12:

A special test fixture, which simulates a subrack, is helpful in speeding up the measurements shown in Figure 7-7.

7.3 Front Panels

This section provides the mechanical specifications for the single and double height board front panels and associated hardware.

PERMISSION 7.10:

VMEbus boards MAY be manufactured with or without front panels.

RECOMMENDATION 7.2:

Use front panels and associated hardware to prevent VMEbus boards from vibrating out of the subrack, and to guide airflow through the subrack.

RULE 7.16:

IF front panels are used,
THEN screws **MUST** be provided on those panels to secure the top and bottom of the panels to the subrack, and their screw threads **MUST** be M2.5 X 0.45 pitch (see Figure 7-8).

Figure 7-8 shows a single height, single width front panel. Figure 7-9 shows a double height, single width front panel.

SUGGESTION 7.4: not used

PERMISSION 7.20: not used.

7.3.1 Handles

PERMISSION 7.11:

The VMEbus board designer **MAY** design VMEbus board front panels with or without handles.

RECOMMENDATION 7.3:

Provide handles to make VMEbus boards easier to remove from the subrack.

OBSERVATION 7.13:

The handles available from various manufacturers vary somewhat in overall shape.

SUGGESTION 7.5:

Choose handles whose depth and height conform to the dimensions shown in Figure 7-8 and Figure 7-9.

RECOMMENDATION 7.4:

When mounting handles on VMEbus board front panels, choose one or more of the locations shown in Figure 7-8, Figure 7-9, Figure 7-12 and Figure 7-13.

PERMISSION 7.12:

On single height boards, handles **MAY** be installed in any of the following combinations

1. Top only
2. Bottom only
3. Top and bottom

PERMISSION 7.13:

On double height boards handles MAY be installed in any of the following combinations:

1. Top only
2. Middle only
3. Bottom only
4. Top and middle
5. Bottom and middle
6. Top and bottom

OBSERVATION 7.14:

When double and single height boards share the same subrack, handles in the center of the double height front panels align with the handles of single height boards, forming an unbroken line and giving a more consistent appearance.

OBSERVATION 7.15

Removal of double height boards that have both a P1 and a 96-pin P2 connector requires up to 180 (N) Newtons (40.5 lbf) of extraction force. Placing handles at the top and bottom of double height boards makes removal of these boards easiest.

7.3.2 Front Panel Mounting**RECOMMENDATION 7.5:**

IF front panels are used,
THEN keep the reserved areas shown in Figure 7-8 and Figure 7-9 free of components to allow for installation of front panel mounting brackets. Locate the holes used to mount these brackets as shown in Figure 7-2 and Figure 7-3.

RECOMMENDATION 7.6:

IF front panels are used on double height boards,
THEN provide at least one center mounting bracket at one of the two locations shown in Figure 7-9 and Figure 7-11.

RULE 7.17:

IF front panels are used,
THEN the test dimension shown in Figure 7-10 and Figure 7-11, from the rear face of the front panel to the rear face of the connector, MUST be maintained.

OBSERVATION 7.16:

The test dimensions from the rear face of the front panel to the front face of the backplane guarantee that the P1 and P2 connectors will be fully engaged and that the front panel fasteners will be able to secure the board into the subrack.

7.3.3 Front Panel Dimensions

All dimensioning of the front panel is done from a datum point 0.15 mm (0.006 inch) to the left of its top left corner as viewed from the front.

RULE 7.18:

Single width front panels **MUST** be designed to the dimensions shown in Figure 7-8 and Figure 7-9.

RECOMMENDATION 7.7:

Make front panels 2.5 mm (0.098 inch) nominal thickness.

OBSERVATION 7.17:

The 20.02 mm (0.788 inch) width of the single slot front panels is 0.30 mm (0.012 inch) narrower than the 20.32 mm (0.8 inch) slot spacing. This prevents mechanical interference between adjacent front panels due to tolerances in the board assemblies and subracks.

RULE 7.19:

IF a board occupies more than one slot and has a front panel,
THEN the width of that front panel **MUST** be 20.02 mm (0.788 inch) plus an integral multiple (N-1) of 20.32 mm (0.8 inch), where N is the number of slots that the board occupies.

RULE 7.20:

Single slot front panels **MUST** be equipped with one fastener at the top and another at the bottom, located as shown in Figure 7-8 and Figure 7-9.

7.3.4 Filler Panels

Filler panels are sometimes used where the backplane positioning leaves a gap on the left or right end of the subrack's front panel or where there are empty slots. These filler panels require no mounting brackets because they are not attached to printed circuit boards. They are secured to the subrack by screws or quarter-turn fasteners on their top and bottom ends like the front panels of VMEbus boards.

RULE 7.21:

Filler panels **MUST** be designed to conform to the dimensions given in Figure 7-12 and Figure 7-13.

RULE 7.22:

Single slot filler panels **MUST** be equipped with one fastener at the top and another at the bottom, located as shown in Figure 7-12 and Figure 7-13.

RECOMMENDATION 7.8:

Use filler panels on VMEbus based systems to maintain proper air flow within the subrack and to improve the appearance of the assembled VMEbus system.

SUGGESTION 7.6:

Equip filler panels with handles. This will give a more consistent appearance when they are installed in the same subrack with VMEbus boards that have front panels and handles.

SUGGESTION 7.7:

Provide additional mounting holes on filler panels wider than 101.60 mm (4.0 inch) to provide better attachment to the subrack.

7.3.5 Board Ejectors/Injectors

OBSERVATION 7.18:

Several vendors offer different types of board ejectors/injectors that make insertion and removal of VMEbus boards easier. Unfortunately, ejectors/injectors designed for use with one vendor's products do not always work with other vendor's products. Because no agreement has been reached between the various manufacturers, the VMEbus specification does not specify any as the preferred type.

OBSERVATION 7.19:

The insertion force for a single 96 pin 603-2-IEC-CO96Mx-xxx connector can be as high as 90 N (Newtons) (20.23 lbf).

PERMISSION 7.14:

VMEbus boards MAY be equipped with any type of ejectors, injectors, or retainers as long as they do not make the products they are used on incompatible with boards or subracks designed to the VMEbus specifications.

7.4 Backplanes

The primary backplane is designated as the J1 backplane. In some cases this is the only backplane in the VMEbus system. When a double height subrack is used, this backplane is mounted in the upper portion of that subrack. When the expanded VMEbus is used, a second backplane, designated the J2 backplane, is installed below the J1 backplane in the lower portion of the subrack. This J2 backplane buses only the center row (row b) of the P2 connector pins, allowing the two outer rows (rows a and c) to be used to implement the VSBbus or for any other user defined functions.

Board slots are designated 1, 2, 3, ... 21, with the slot numbering starting at the left end of the subrack, as viewed from the front. The daisy-chain propagation starts with slot 1 and goes to 21.

RULE 7.23:

J1 VMEbus backplanes MUST bus all signals in rows a, b and c in all slots except for the daisy-chained signals. (See Section 7.7.1)

RULE 7.24:

When a J2 backplane is used to provide for 32-bit wide address and data transfers, it MUST bus all pins of the row b of those slots for which it has connectors. (See Section 7.7.2)

RULE 7.25:

The 603-2-IEC-CO96Fx-xxx connectors MUST be used on all J1 VMEbus backplanes and on all J2 expansion backplanes.

RULE 7.42:

IF a 96 pin 603-2-IEC-CO96Fx-xxx is used for J1
THEN the pin assignments MUST be as shown in Table 7-1 using rows a, b and c.

RULE 7.43:

IF a 160 pin 603-2-IEC-CO96Fx-xxx is used for J1
THEN the pin assignments MUST be as shown in Table 7-1 using rows z, a, b, c and d.

RULE 7.44:

IF a 96 pin 603-2-IEC-C096Fx-xxx is used for J2
THEN the pin assignments MUST be as shown in Table 7-2 using rows a, b and c.

RULE 7.45:

IF a 160 pin 603-2-IEC-C096Fx-xxx is used for J2
THEN the pin assignments MUST be as shown in Table 7-2 using rows z, a, b, c and d.

RULE 7.26:

All VMEbus J1 backplanes MUST have some provision for jumpering the interrupt acknowledge and bus grant daisy-chains when boards are not plugged into a slot.

PERMISSION 7.25:

Jumpering can be implemented with either an automatic mechanical or electronic mechanism.

SUGGESTION 7.8:

To provide for the jumpering of the J1 backplane daisy-chains, use 603-2-IEC-C096Fx-xxx connectors that have wire-wrap pins.

SUGGESTION 7.9:

IF J1 connectors without wire-wrap pins are used,
THEN place all of the jumper pins for daisy-chains next to the J1 connector that they are jumpering.

RECOMMENDATION 7.15:

To improve system ground integrity use a J1/J2 backplane.

PERMISSION 7.15: not used.

SUGGESTION 7.10:

Use 603-2-IEC-C096Fx-xxx connectors with wire-wrap pins for the connectors on J2 backplanes. This allows attachment of ribbon cables and secondary backplanes to those pins.

7.4.1 Backplane Dimensional Requirements

Figures 7-14 and 7-15 depict a 21-slot backplane.

PERMISSION 7.16:

VMEbus backplanes MAY be designed with up to 21 slots.

RECOMMENDATION 7.9:

When designing a backplane with fewer than 21 slots, make the width (N x 20.32 mm)-1.44 mm, +0/-0.3 mm where N is the number of slots. This allows two backplanes to be installed next to each other without wasting a slot position in the subrack.

RECOMMENDATION 7.10:

Do not design backplanes with widths greater than 425.28 mm (16.743 inches). Backplanes longer than this might not fit into widely available subracks.

PERMISSION 7.17: not used.

RULE 7.27:

Except for the width dimension which varies with the number of slots, VMEbus backplanes **MUST** be designed to the dimensions given in Figure 7-14 and Figure 7-15.

OBSERVATION 7.20:

The backplane dimensions shown in Figure 7-14 and Figure 7-15 repeat every 20.32 mm (0.8 inch).

7.4.2 Signal Line Termination Networks

RULE 7.28:

VMEbus backplanes **MUST** either have built-in terminations or they **MUST** provide some way to connect plug-on terminator boards for terminating all of the signal lines indicated in Section 6.7.

PERMISSION 7.18:

The termination circuitry **MAY** either be built onto the ends of the backplane or it **MAY** be provided by separate modules that plug onto either the front or the back of the backplane at the two end slots.

RULE 7.29:

Backplane signal trace lengths, including any plug-on terminator boards, **MUST NOT** exceed 508 mm (20.0 inches).

OBSERVATION 7.21:

The entire board width of a 21-slot VMEbus backplane is required to accommodate the 603-2-IEC-xxxxxx-xxx connectors.

7.5 Assembly of VMEbus Subracks

This section shows how VMEbus subracks are assembled. All horizontal dimensions are from the left hand edge of the front opening of the subrack.

7.5.1 Subracks And Slot Widths

Figure 7-18 shows a typical double height 21-slot subrack.

PERMISSION 7.19:

Double height subracks **MAY** be used to house double height boards or they **MAY** be subdivided by the installation of a bracket, into two single height sections, one above the other.

OBSERVATION 7.22:

The bracket allowed in **PERMISSION 7.19** provides two board guides: the lower guide for the board above it and the upper guide for the board below it.

SUGGESTION 7.11:

Where possible, install the left-most (Slot 1) board guides so that their center line is 3.27 mm (0.129 inches) from the left end of the rack opening. This provides sufficient clearance for the component leads of the board installed in that slot, and doesn't waste any horizontal space. (If this board guide is installed farther to the right there will not be room for 21 slots.)

7.5.2 Subrack Dimensions**RULE 7.30:**

All double height VMEbus subracks **MUST** meet all of the dimensional requirements given in Figure 7-18, except for the width dimension, which varies according to how many slots are in the subrack.

RULE 7.31:

All single height VMEbus subracks **MUST** meet all of the dimensional requirements given in Figure 7-18, except for the width dimension, which varies according to how many slots are in the subrack, and the vertical distance between the lower and upper board guides, which is 100.2 mm, +0.4/-0.0 mm (3.94 inches +0.015/-0.0 inches) instead of 233.55 mm (9.19 inches).

OBSERVATION 7.23:

The dimension from the front panel mounting surface to the front face of the backplane is particularly critical, since it guarantees correct connector engagement.

7.6 Conduction Cooled VMEbus Systems**RULE 7.46:**

Conduction cooled VMEbus boards, backplanes and chassis **MUST** conform to IEEE 1101.2.

7.7 VMEbus Backplane Connectors and VMEbus Board Connectors**RULE 7.47:**

Reserved pins **MUST** be either not connected or not populated.

7.7.1 Pin Assignments For The J1/P1 Connector

Table 7-1 provides signal names for the J1/P1 connector pins. (The 96 pin connector consists of three rows of pins labeled rows a, b, and c. The 160 pin connector has five rows labeled z, a, b, c and d.)

Table 7 - 1 J1/P1 Pin Assignments

Pin Number	Row z Signal Mnemonics	Row a Signal Mnemonics	Row b Signal Mnemonics	Row c Signal Mnemonics	Row d Signal Mnemonics
1	Reserved	D00	BBSY*	D08	Reserved
2	GND	D01	BCLR*	D09	Reserved
3	Reserved	D02	ACFAIL*	D10	Reserved
4	GND	D03	BG0IN*	D11	Reserved
5	Reserved	D04	BG0OUT*	D12	Reserved
6	GND	D05	BG1IN*	D13	Reserved
7	Reserved	D06	BG1OUT*	D14	Reserved
8	GND	D07	BG2IN*	D15	Reserved
9	Reserved	GND	BG2OUT*	GND	Reserved
10	GND	SYSCLK	BG3IN*	SYSFAIL*	Reserved
11	Reserved	GND	BG3OUT*	BERR*	Reserved
12	GND	DS1*	BR0*	SYSRESET*	Reserved
13	Reserved	DS0*	BR1*	LWORD*	Reserved
14	GND	WRITE*	BR2*	AM5	Reserved
15	Reserved	GND	BR3*	A23	Reserved
16	GND	DTACK*	AM0	A22	Reserved
17	Reserved	GND	AM1	A21	Reserved
18	GND	AS*	AM2	A20	Reserved
19	Reserved	GND	AM3	A19	Reserved
20	GND	IACK*	GND	A18	Reserved
21	Reserved	IACKIN*	SERA	A17	Reserved
22	GND	IACKOUT*	SERB	A16	Reserved
23	Reserved	AM4	GND	A15	Reserved
24	GND	A07	IRQ7*	A14	Reserved
25	Reserved	A06	IRQ6*	A13	Reserved
26	GND	A05	IRQ5*	A12	Reserved
27	Reserved	A04	IRQ4*	A11	Reserved
28	GND	A03	IRQ3*	A10	Reserved
29	Reserved	A02	IRQ2*	A09	Reserved
30	GND	A01	IRQ1*	A08	Reserved
31	Reserved	-12V	+5VSTDBY	+12V	Reserved
32	GND	+5V	+5V	+5V	Reserved

7.7.2 Pin Assignments For The J2/P2 Connector

Table 7-2 provides signal names for the J2/P2 connector pins. (The 96 pin connector consists of three rows of pins labeled rows a, b, and c. The 160 pin connector has five rows labeled z, a, b, c and d.)

Table 7 - 2 J2/P2 Pin Assignments

Pin Number	Row z Signal Mnemonics	Row a Signal Mnemonics	Row b Signal Mnemonics	Row c Signal Mnemonics	Row d Signal Mnemonics
1	Reserved	User Defined	+5V	User Defined	Reserved
2	GND	User Defined	GND	User Defined	Reserved
3	Reserved	User Defined	RETRY*	User Defined	Reserved
4	GND	User Defined	A24	User Defined	Reserved
5	Reserved	User Defined	A25	User Defined	Reserved
6	GND	User Defined	A26	User Defined	Reserved
7	Reserved	User Defined	A27	User Defined	Reserved
8	GND	User Defined	A28	User Defined	Reserved
9	Reserved	User Defined	A29	User Defined	Reserved
10	GND	User Defined	A30	User Defined	Reserved
11	Reserved	User Defined	A31	User Defined	Reserved
12	GND	User Defined	GND	User Defined	Reserved
13	Reserved	User Defined	+5V	User Defined	Reserved
14	GND	User Defined	D16	User Defined	Reserved
15	Reserved	User Defined	D17	User Defined	Reserved
16	GND	User Defined	D18	User Defined	Reserved
17	Reserved	User Defined	D19	User Defined	Reserved
18	GND	User Defined	D20	User Defined	Reserved
19	Reserved	User Defined	D21	User Defined	Reserved
20	GND	User Defined	D22	User Defined	Reserved
21	Reserved	User Defined	D23	User Defined	Reserved
22	GND	User Defined	GND	User Defined	Reserved
23	Reserved	User Defined	D24	User Defined	Reserved
24	GND	User Defined	D25	User Defined	Reserved
25	Reserved	User Defined	D26	User Defined	Reserved
26	GND	User Defined	D27	User Defined	Reserved
27	Reserved	User Defined	D28	User Defined	Reserved
28	GND	User Defined	D29	User Defined	Reserved
29	Reserved	User Defined	D30	User Defined	Reserved
30	GND	User Defined	D31	User Defined	Reserved
31	Reserved	User Defined	GND	User Defined	Reserved
32	GND	User Defined	+5V	User Defined	Reserved

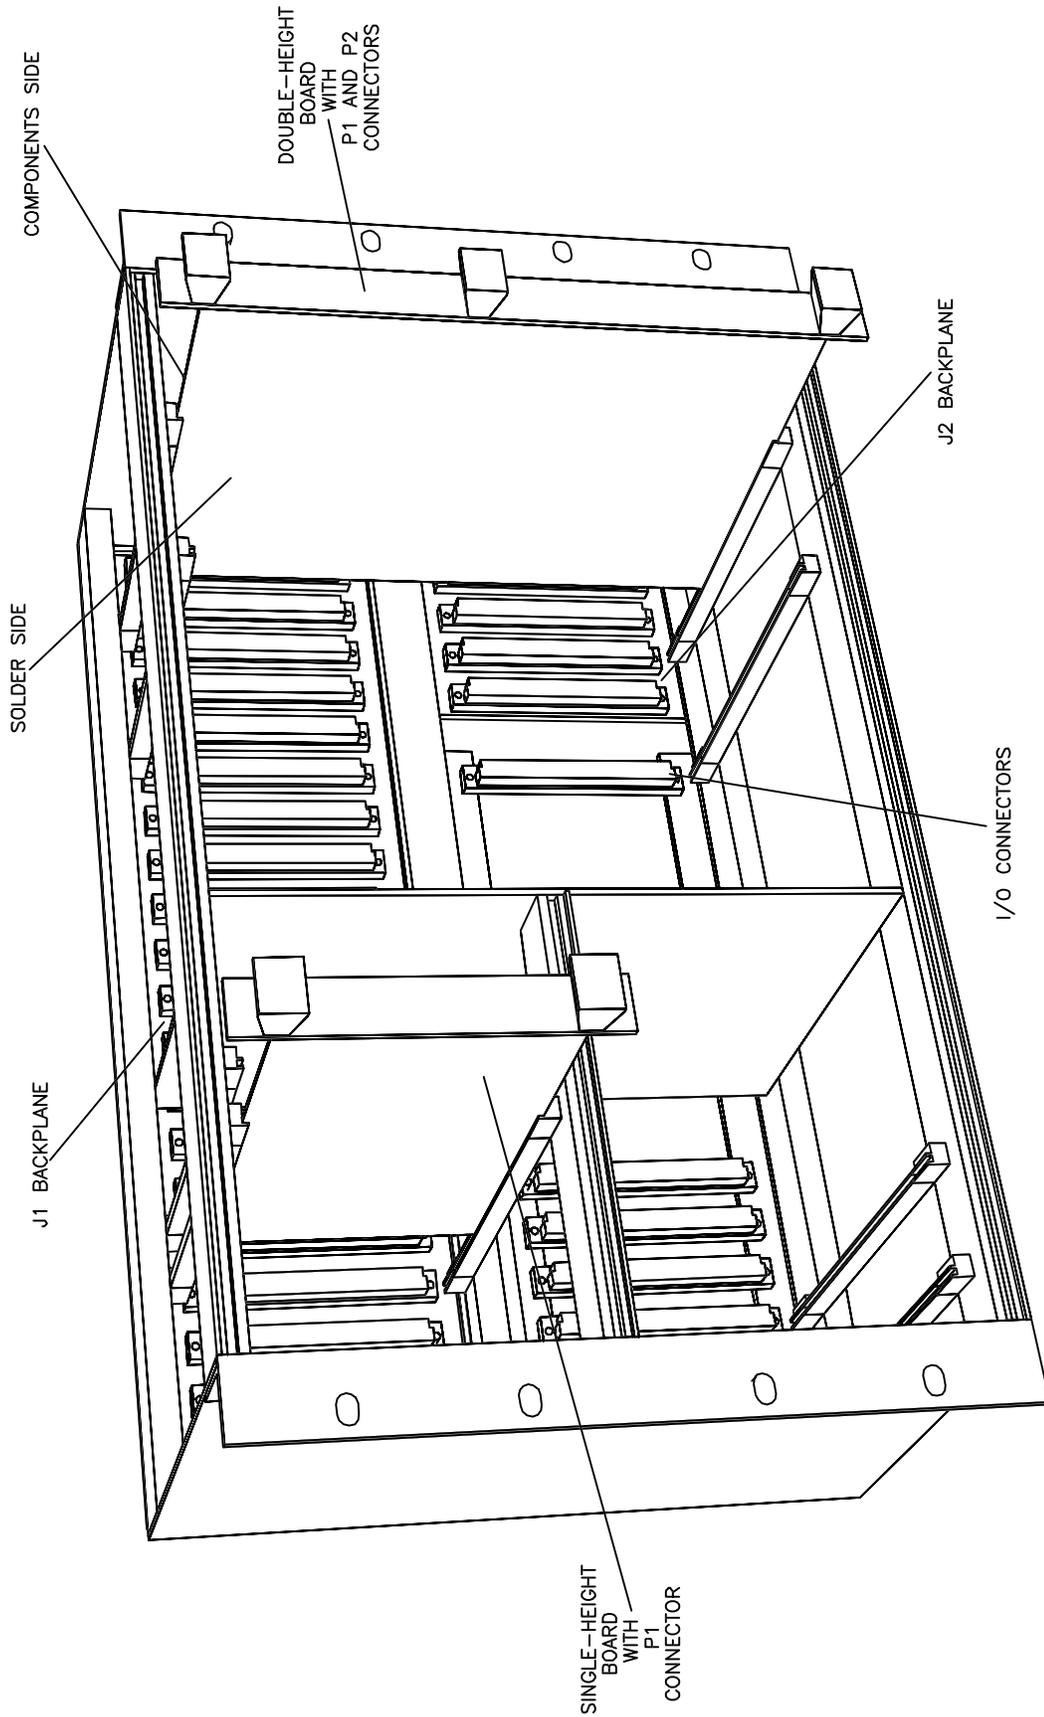
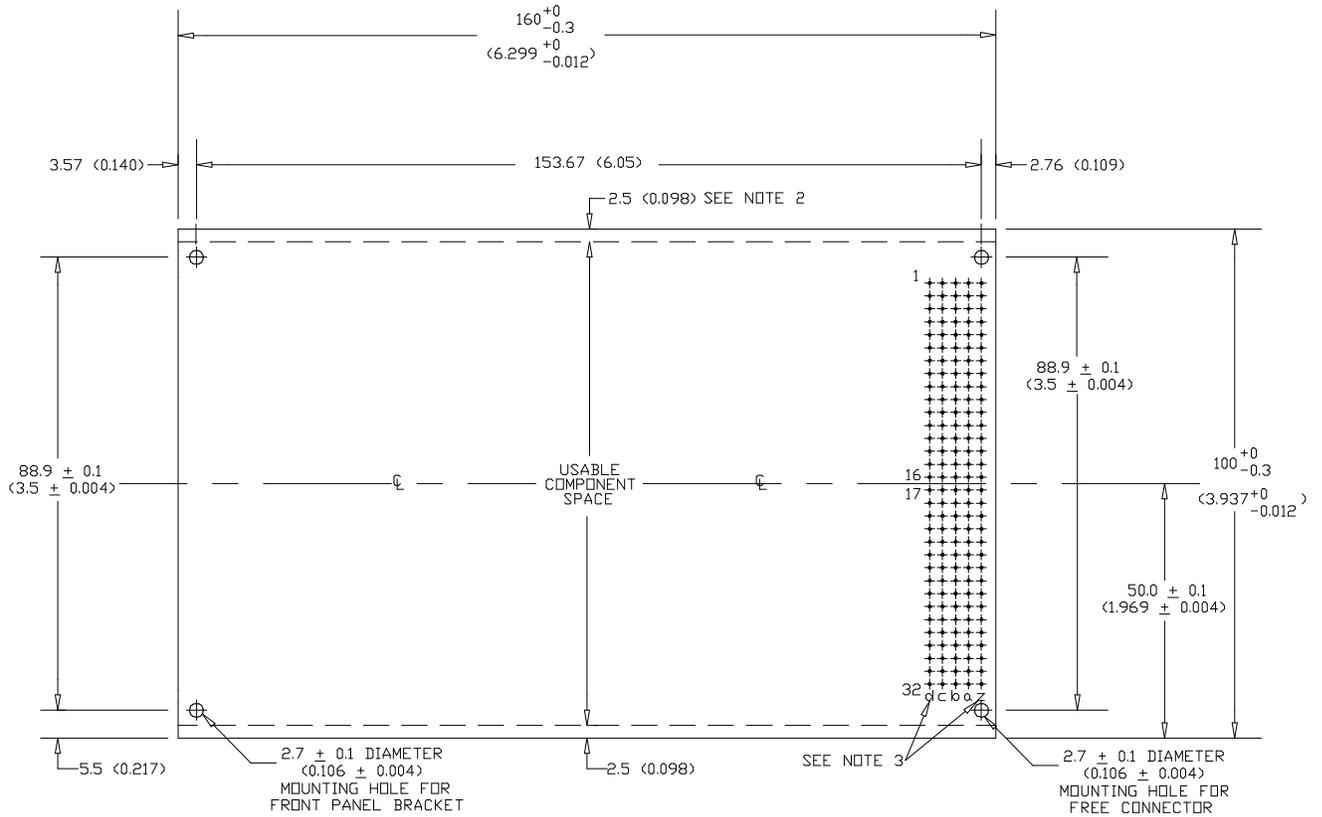


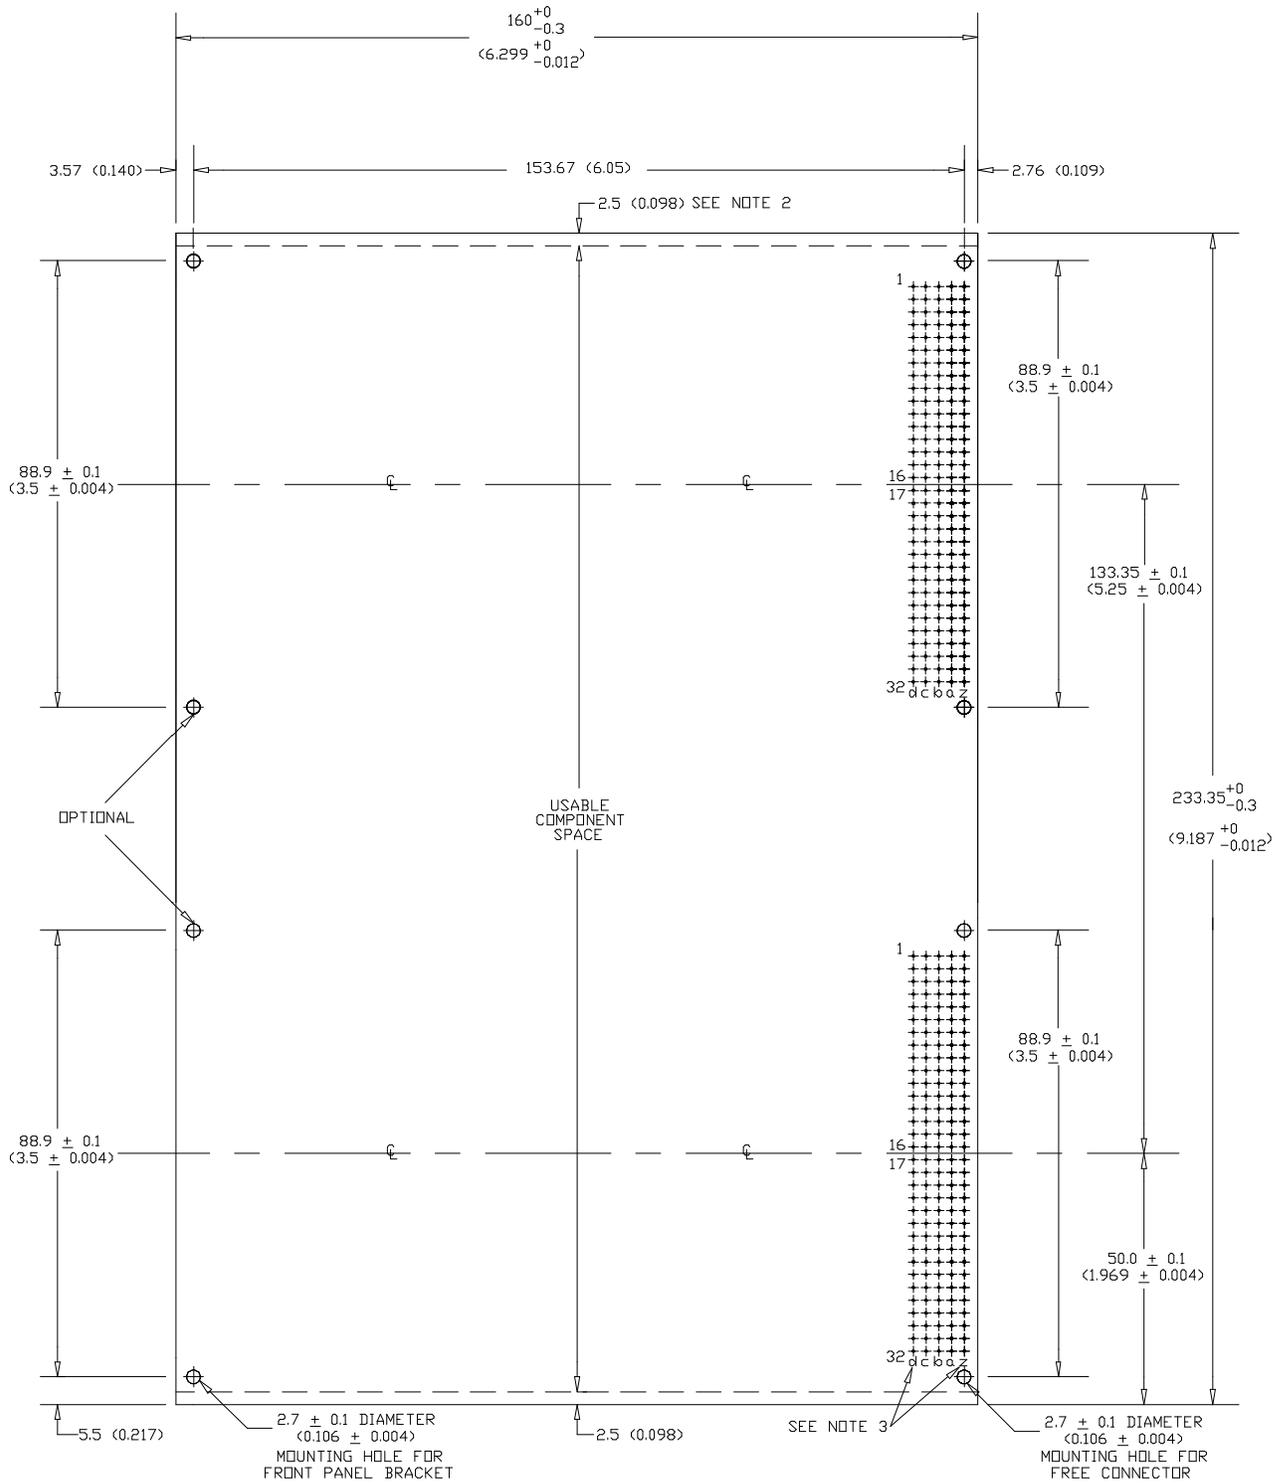
Figure 7 - 1 Subrack with Mixed Board Sizes



NOTES:

1. All dimensions are shown in millimeters. Inch dimensions are shown in parentheses.
2. **Rule 7.32:** Boards **MUST** be 1.6±0.2mm (0.063±0.008 inch) thick in the guide area.
3. "d" and "z" row pin positions used for optional 5 row DIN connector.

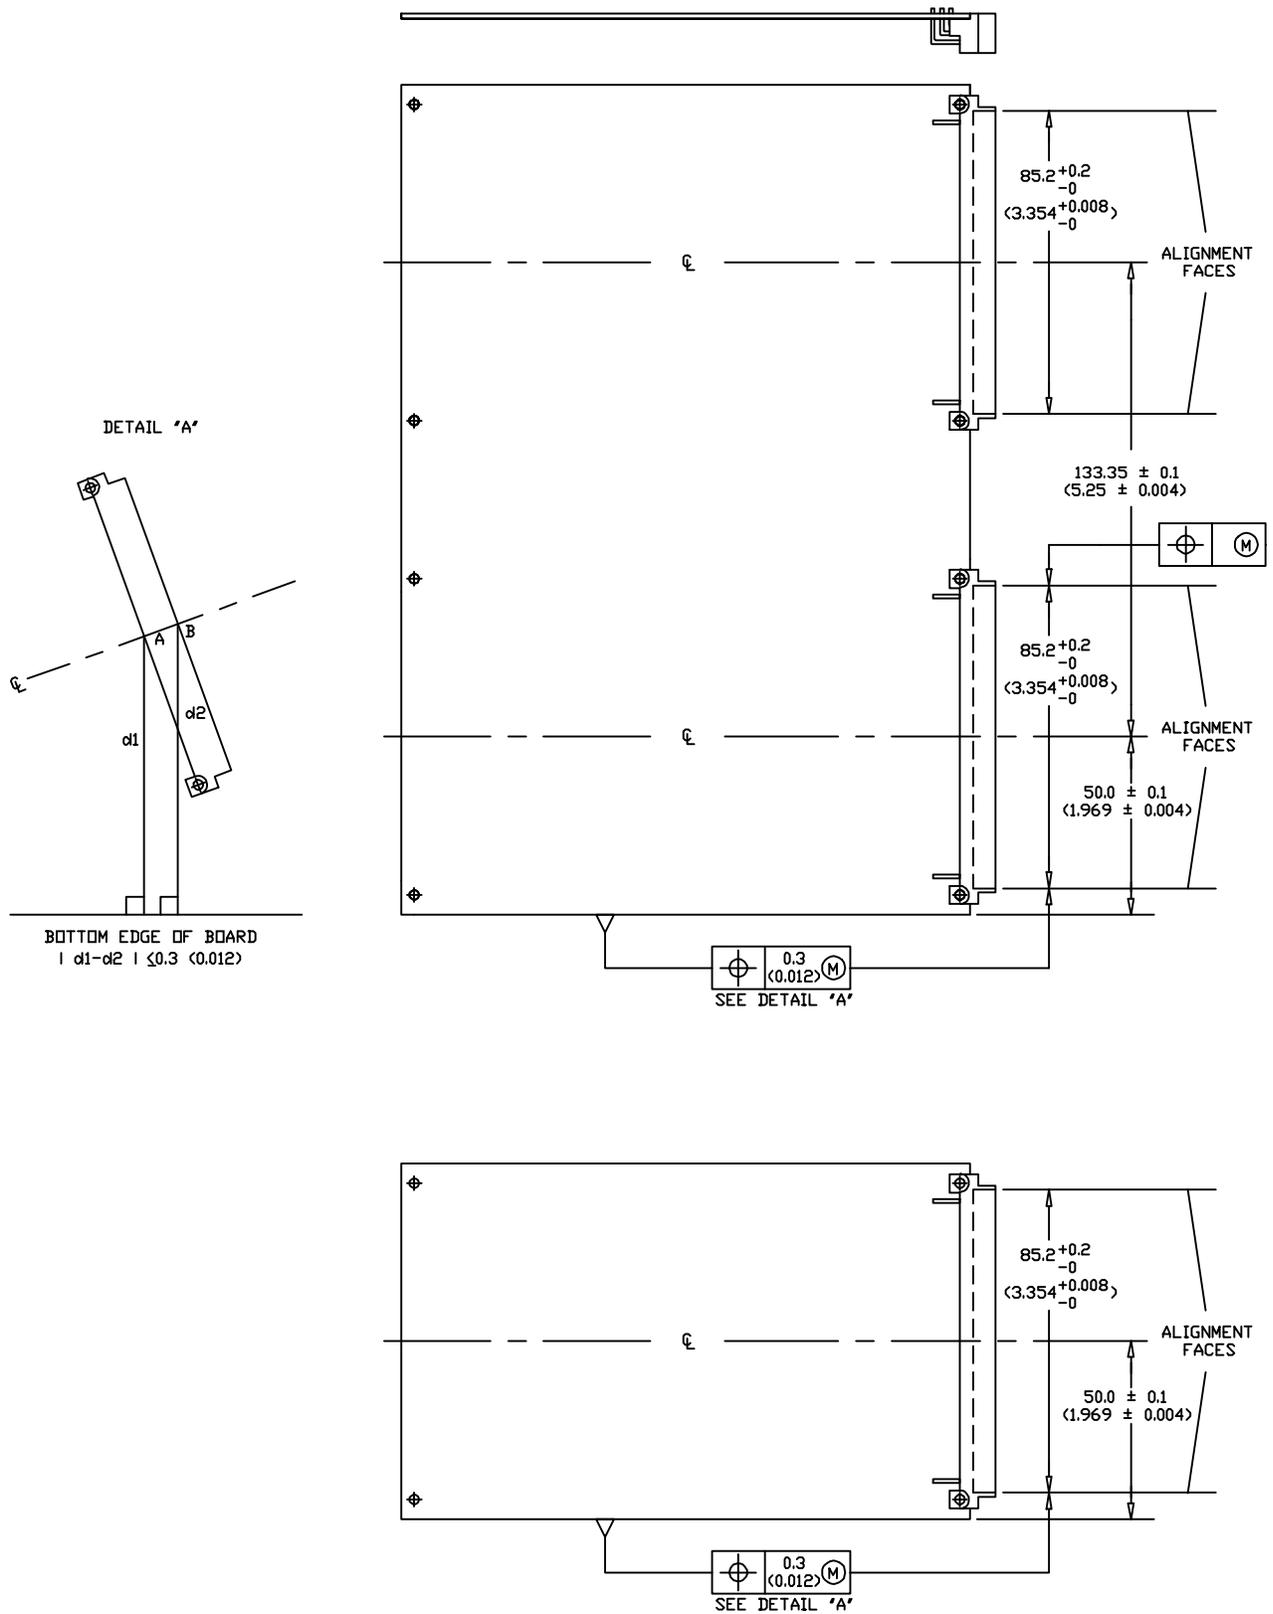
Figure 7 - 2 Single Height Board: Basic Dimensions



NOTES:

1. All dimensions are shown in millimeters. Inch dimensions are shown in parentheses.
2. **Rule 7.33:**
Boards **MUST** be 1.6±0.2mm (0.063±0.008 inch) thick in the guide area.
3. "d" and "z" row pin positions used for optional 5 row DIN connector.

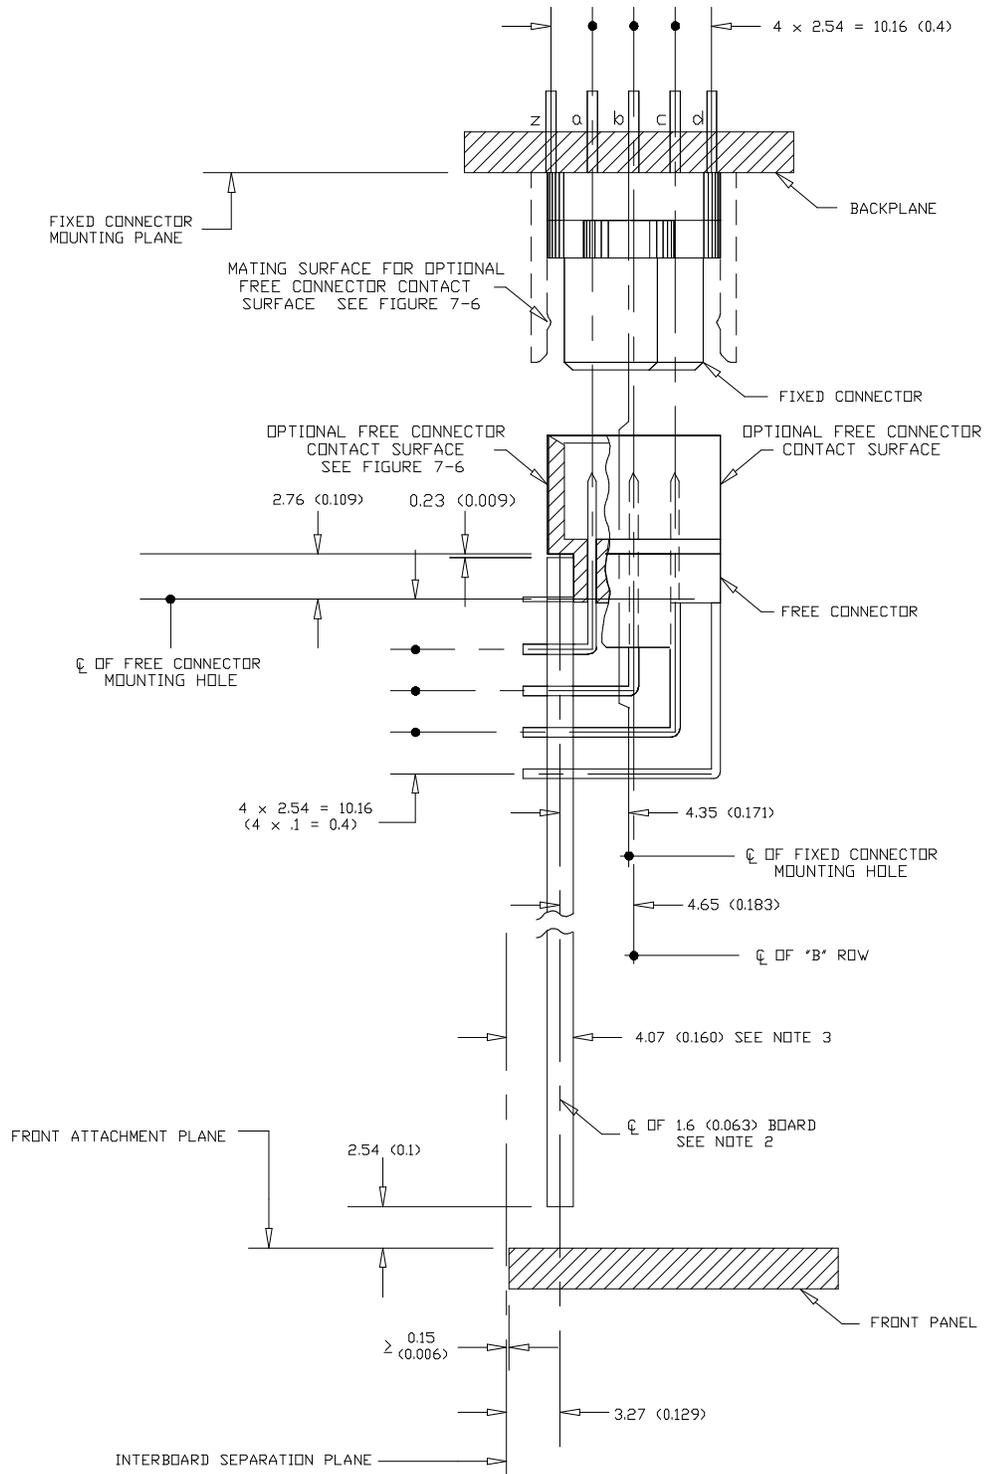
Figure 7-3 Double Height Board: Basic Dimensions



NOTE:

1. All dimensions are shown in millimeters. Inch dimensions are shown in parentheses.

Figure 7-4 Connector Positions On Single and Double Height Boards



NOTES:

1. All dimensions are shown in millimeters. Inch dimensions are shown in parentheses.
2. For additional information regarding the allowed thickness of boards, see Section 7.2.
3. This 4.07mm (0.160 in) dimension is the same regardless of the thickness of the board.

Figure 7-5 Cross Sectional View of Board, Connector, Backplane, And Front Panel

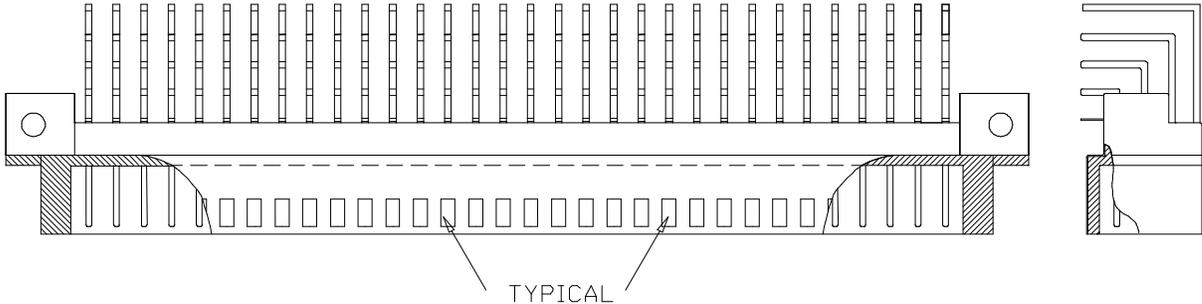
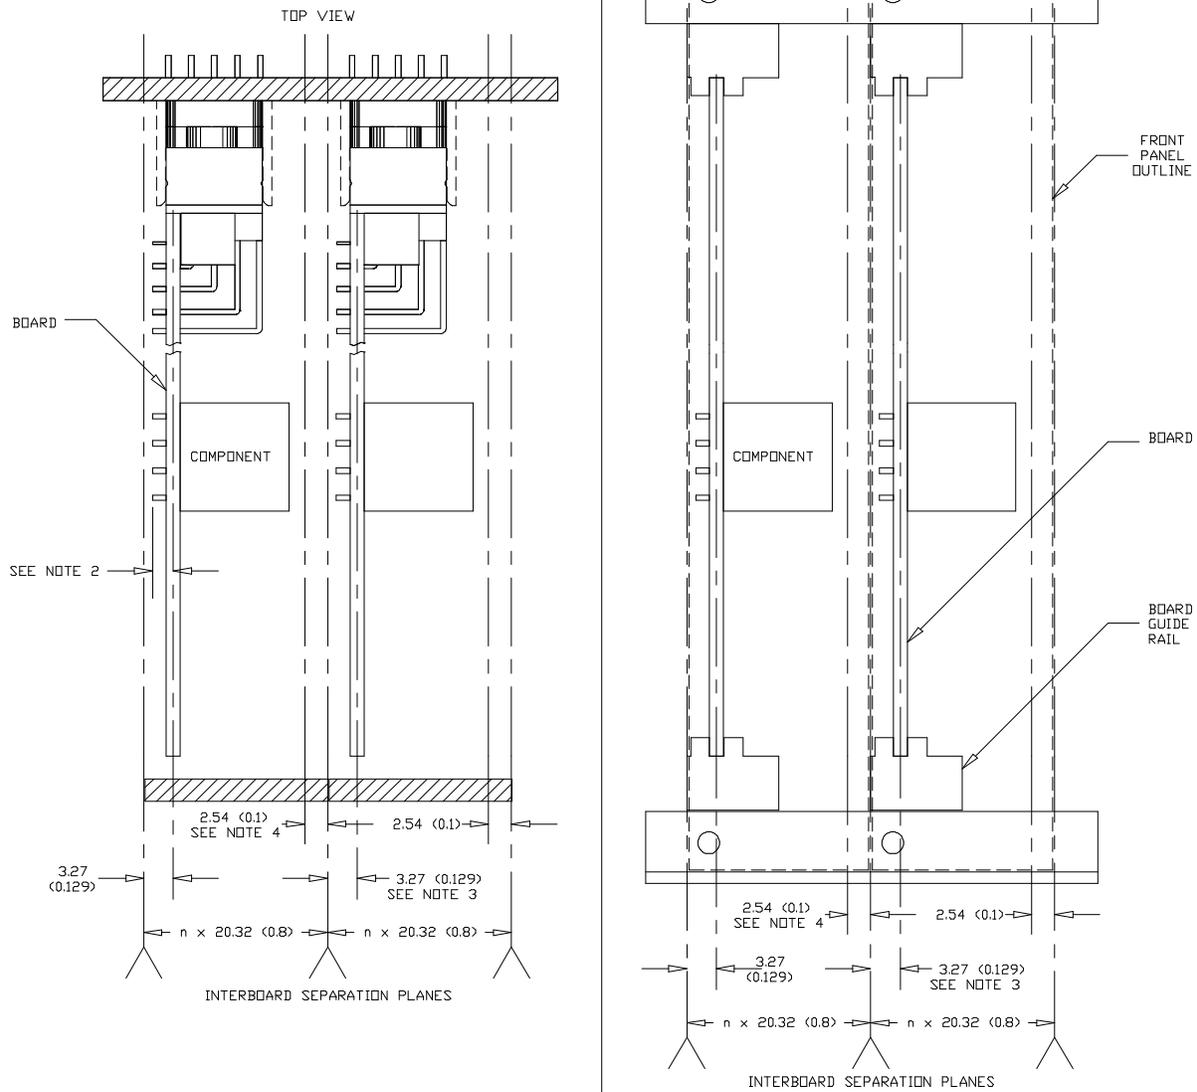


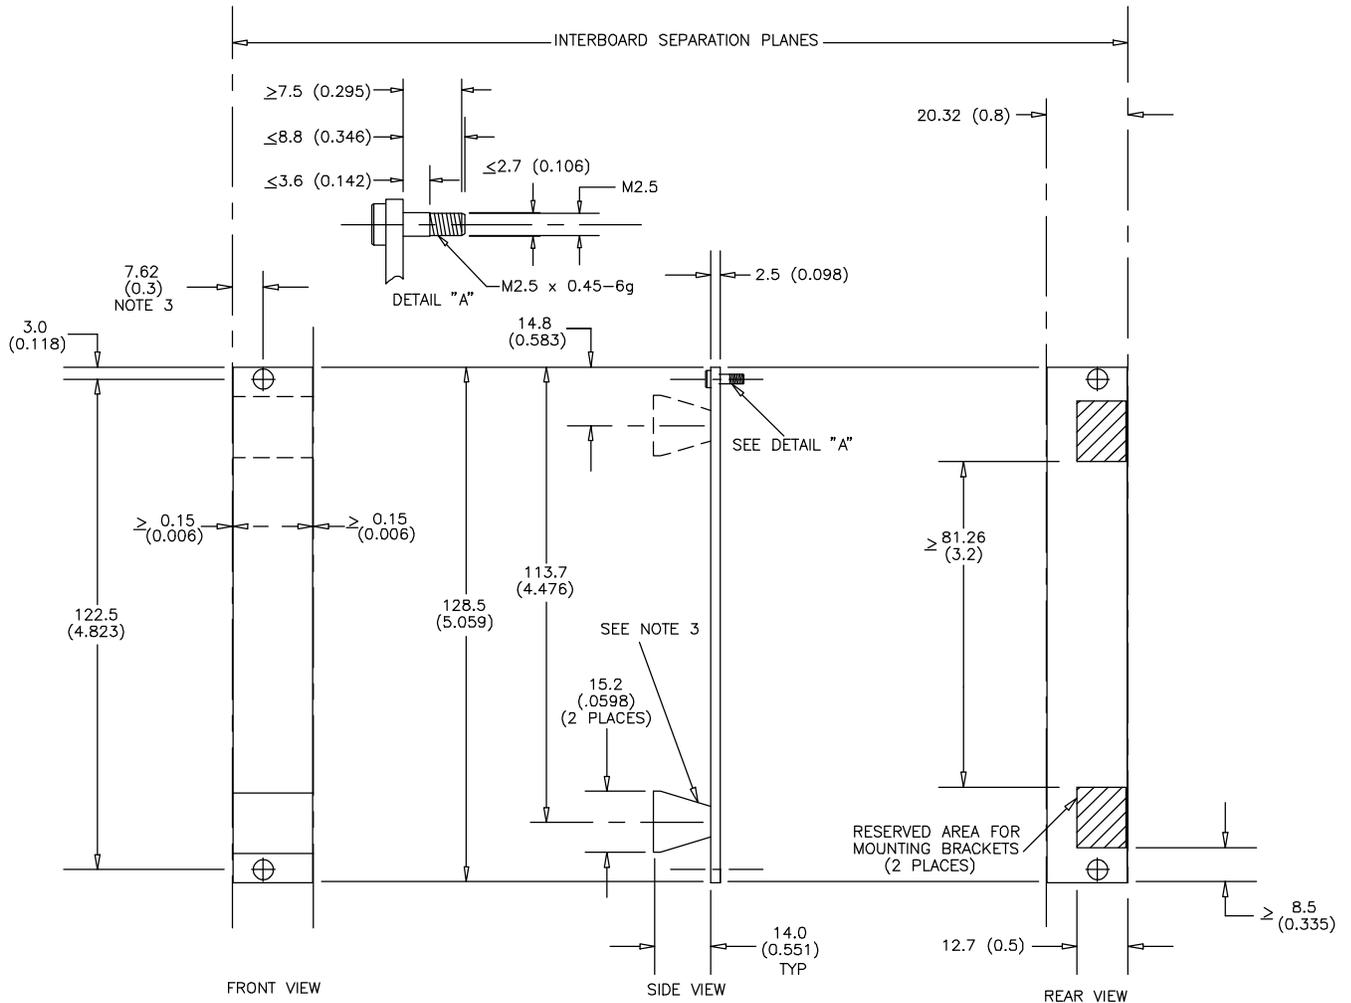
Figure 7-6 Optional Enhanced DIN Connector



NOTES:

1. All dimensions are shown in millimeters. Inch dimensions are shown in parentheses.
2. **SUGGESTION 7.12:**
Trim component lead lengths to no more than 1.52 mm (0.06 in).
3. **RULE 7.34:**
Components leads and components mounted on the solder side of the board **MUST NOT** protrude through the interboard separation plane after it has been completely seated in the backplane.
4. **RULE 7.35:**
Components mounted on the component side of the board **MUST NOT** be any closer than 2.54 mm (0.1 in) to the interboard separation plane after it has been completely seated in the backplane.

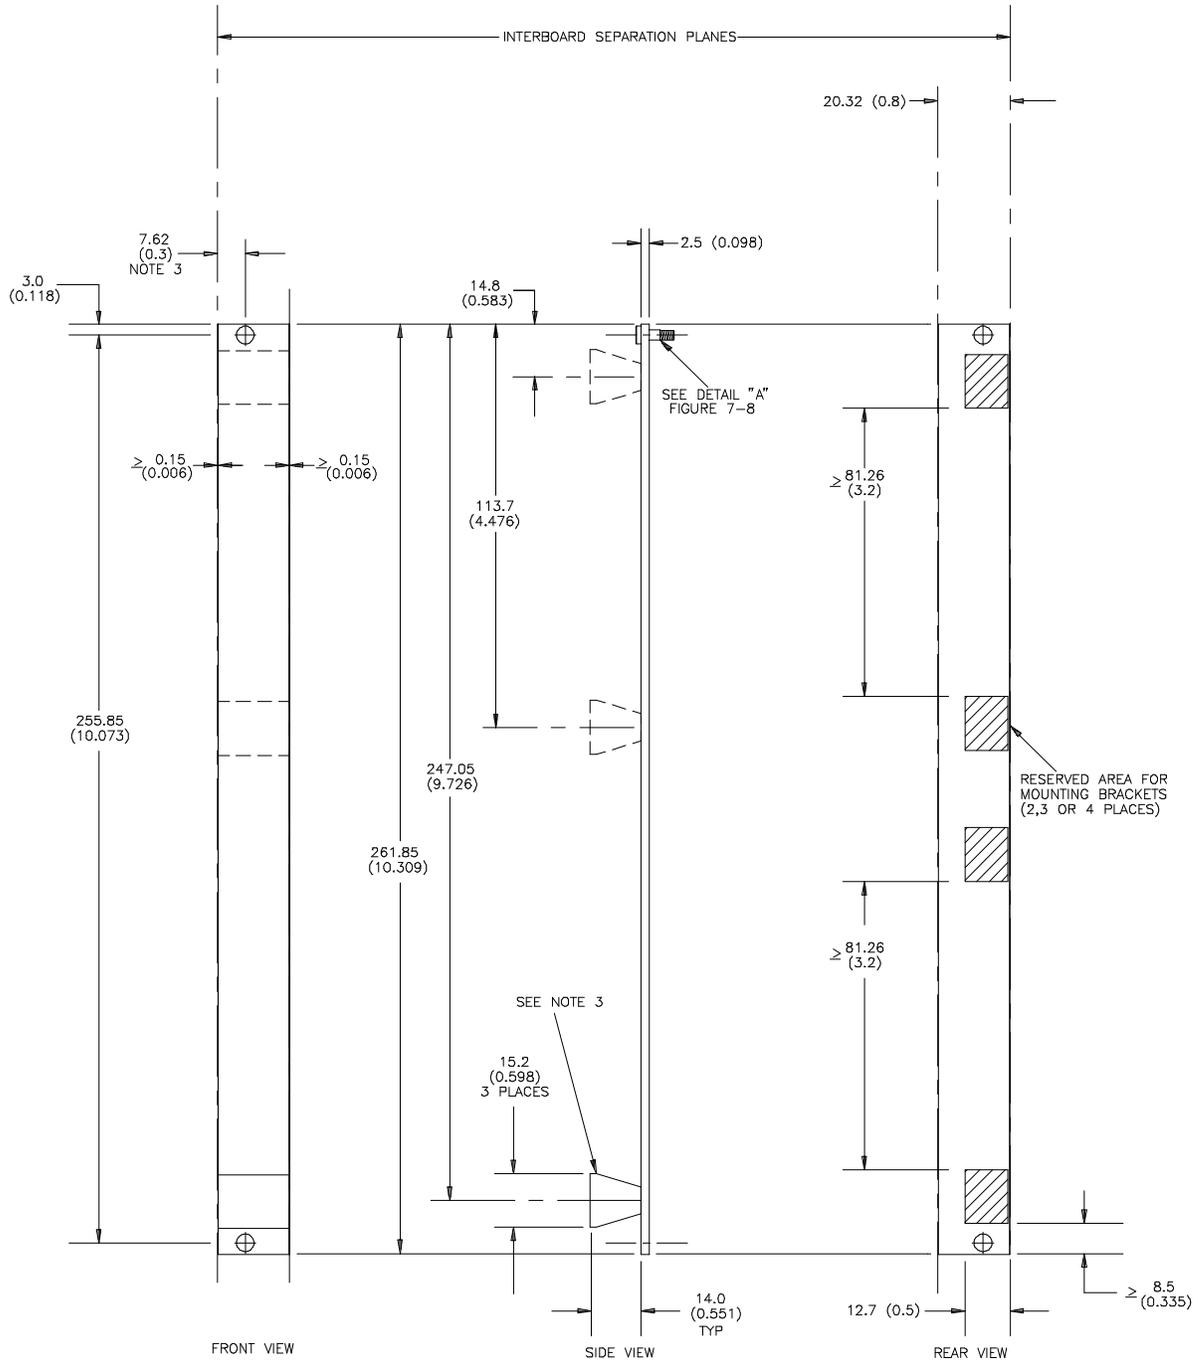
Figure 7-7 Component Height, Lead Length And Board Warpage



NOTES:

1. All dimensions are shown in millimeters. Inch dimensions are shown in parentheses.
2. Dimensions given for height and depth of handles are suggestions only.
3. **RECOMMENDATION 7.11:**
Locate the mounting hole 7.62 mm (0.3 in) from the interboard separation plane.
- PERMISSION 7.20:**
The mounting hole **MAY** be located 12.7 mm (0.5 in) from the interboard separation plane.

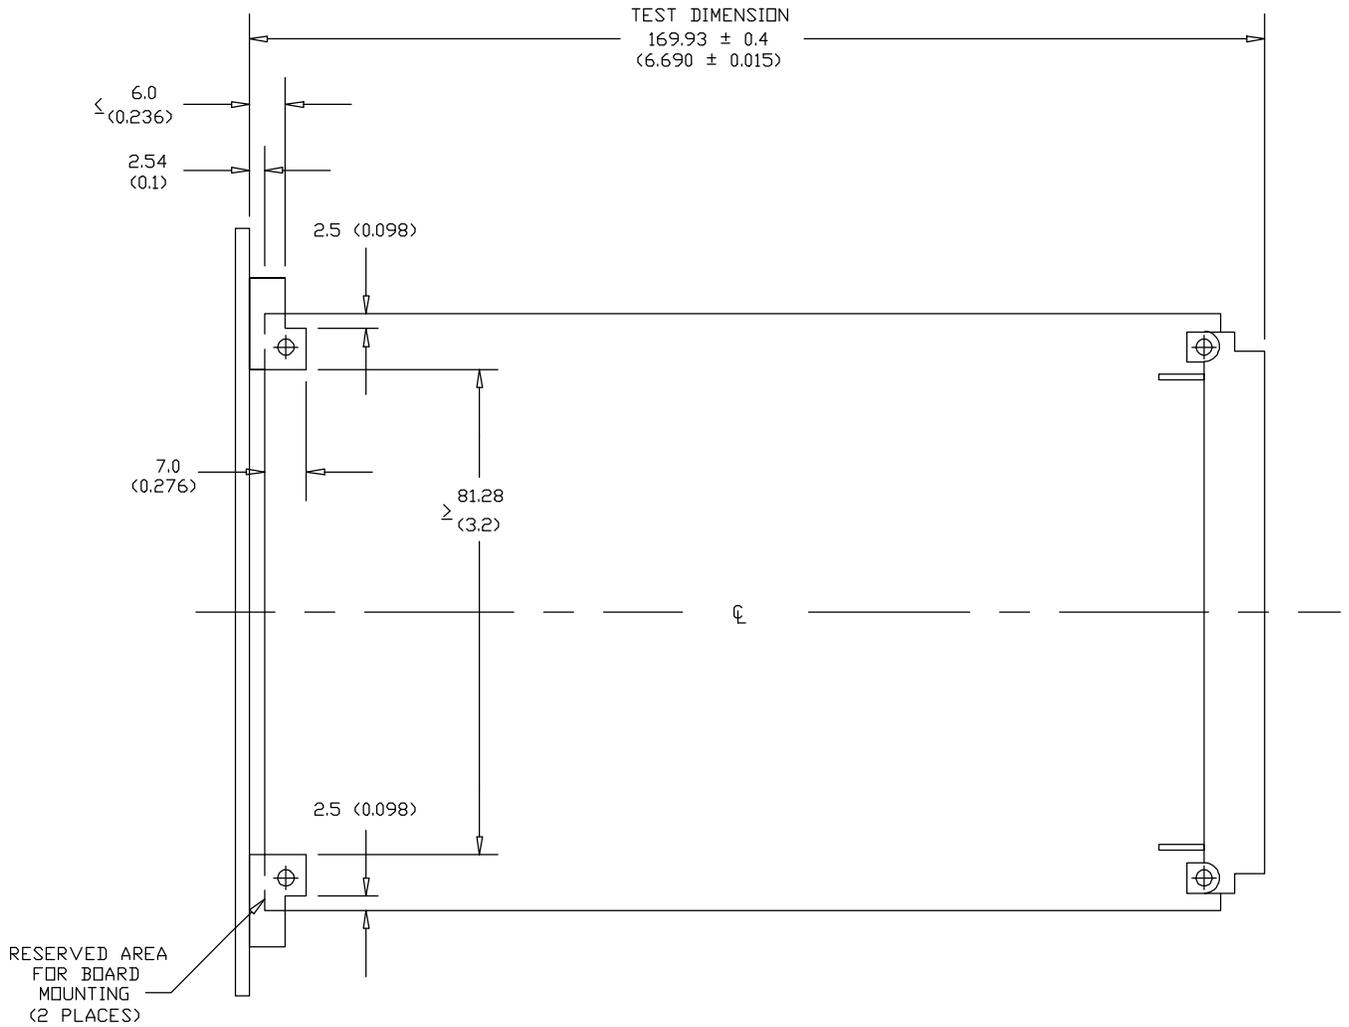
Figure 7 - 8 Single Height, Single Width Front Panel



NOTES:

1. All dimensions are shown in millimeters. Inch dimensions are shown in parentheses.
2. Dimensions given for height and depth of handles are suggestions only.
3. **RECOMMENDATION 7.12:**
Locate the mounting hole 7.62 mm (0.3 in) from the interboard separation plane.
PERMISSION 7.21:
The mounting hole **MAY** be located 12.7 mm (0.5 in) from the interboard separation plane.

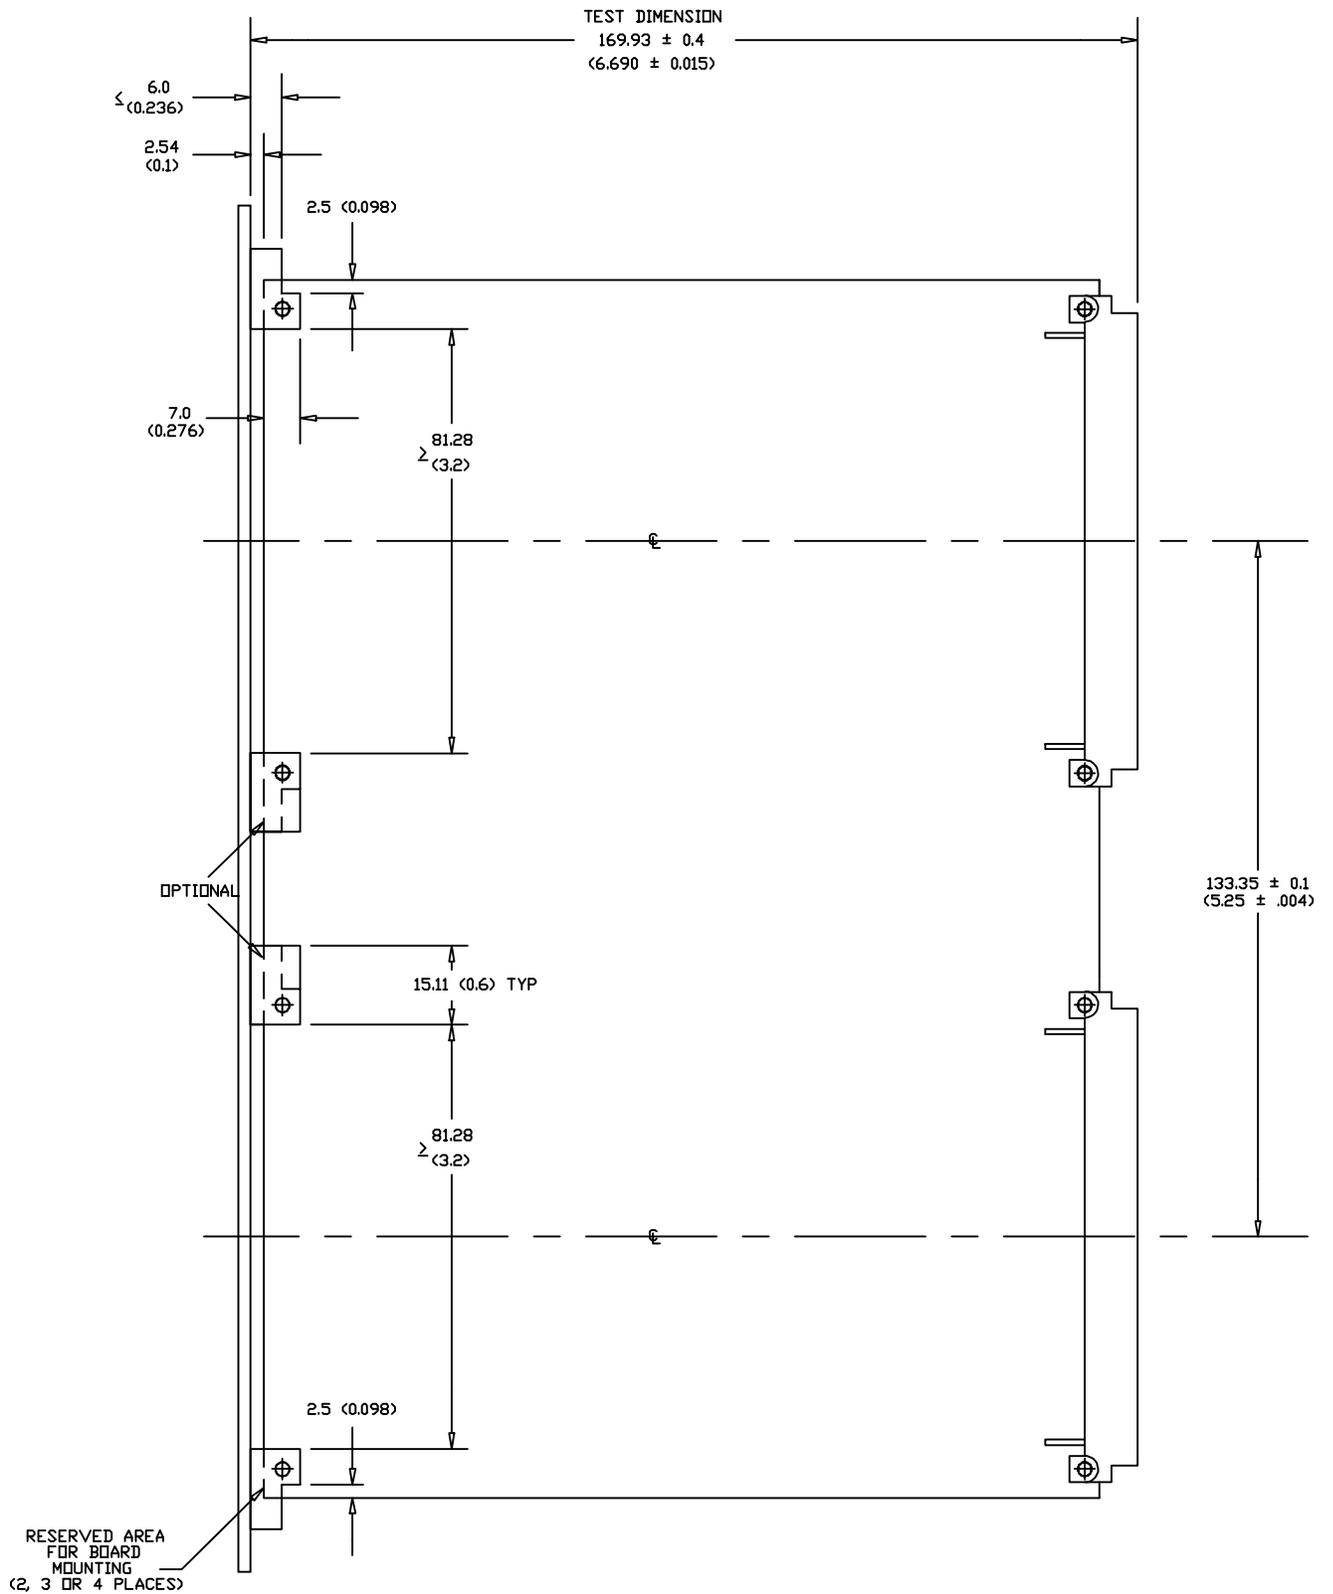
Figure 7 - 9 Double Height, Single Width Front Panel



NOTES:

1. All dimensions are shown in millimeters. Inch dimensions are shown in parentheses.

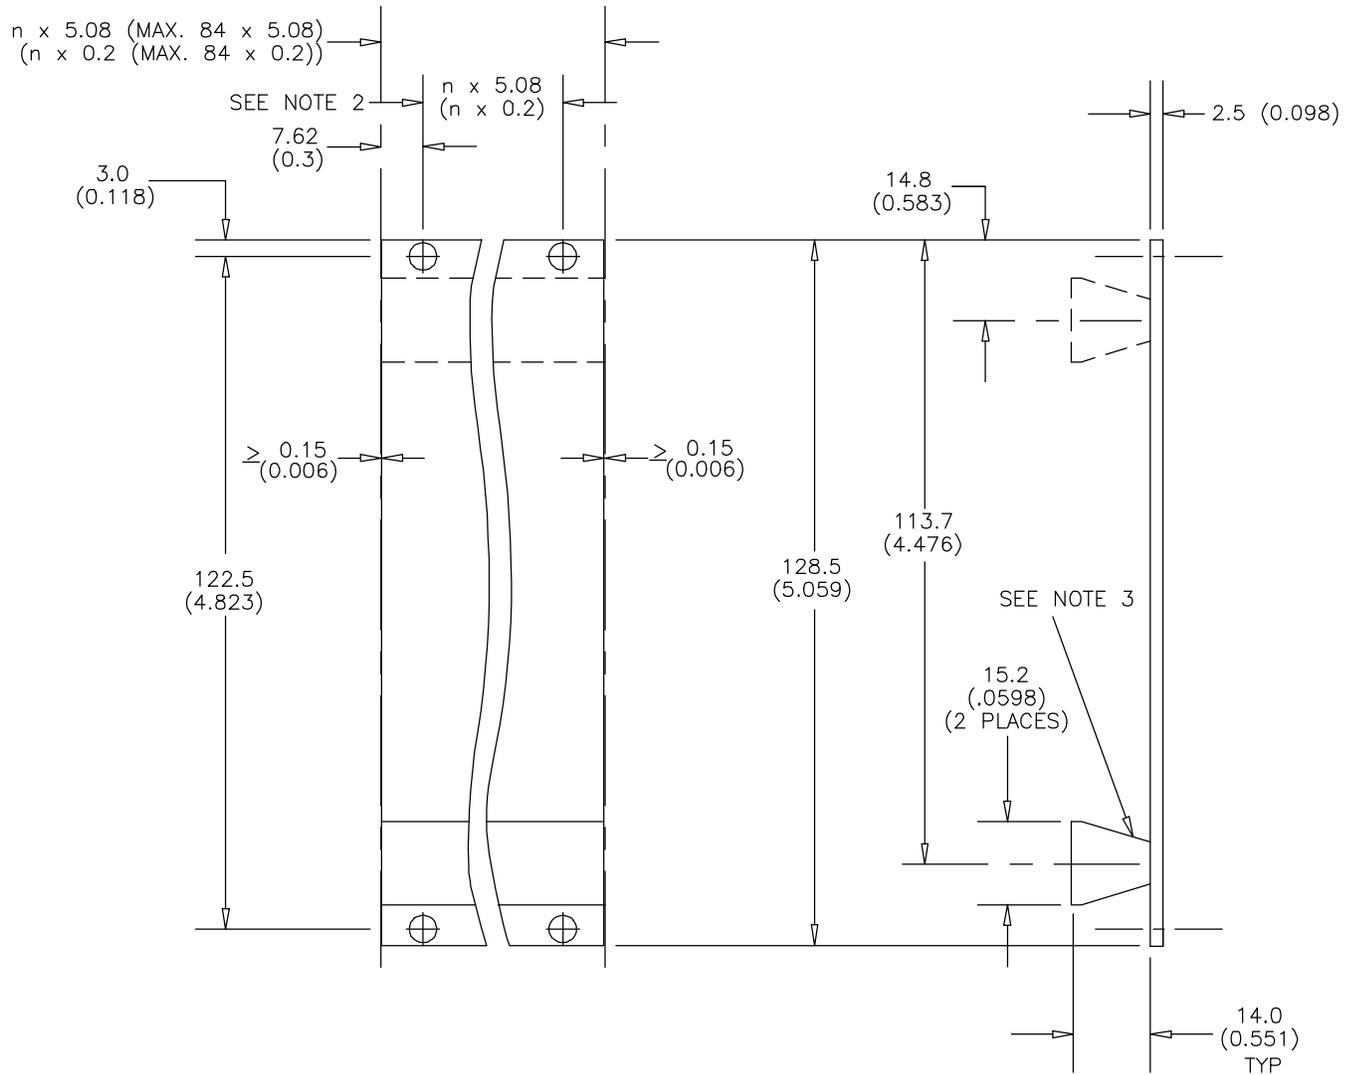
Figure 7 - 10 Front Panel Mounting Brackets And Dimension Of Single Height Boards.



NOTES:

1. All dimensions are shown in millimeters. Inch dimensions are shown in parentheses.

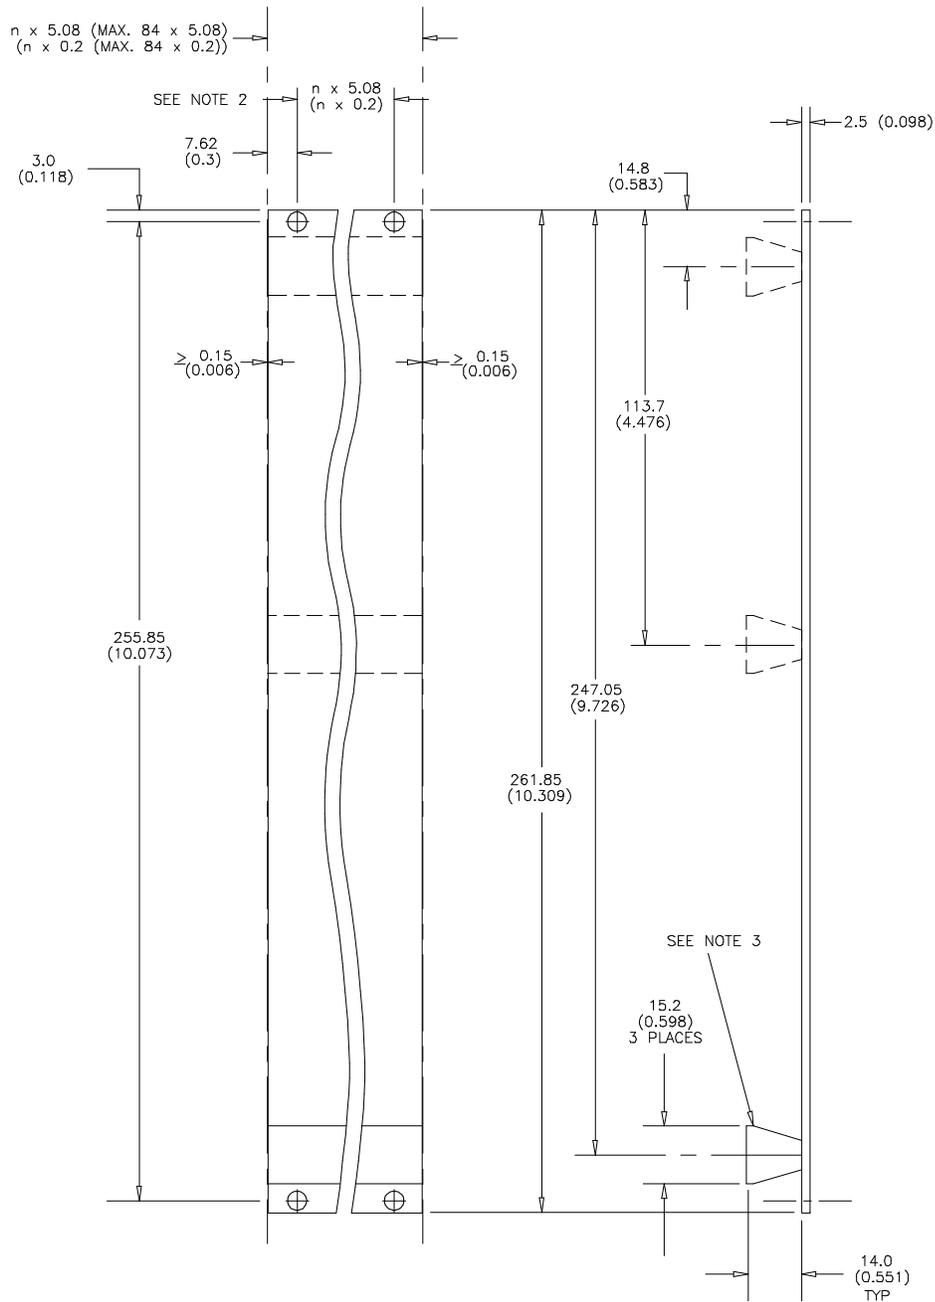
Figure 7 - 11 Front Panel Mounting Brackets And Dimension Of Double Height Boards.



NOTES:

1. All dimensions are shown in millimeters. Inch dimensions are shown in parentheses.
2. **RECOMMENDATION 7.13:**
Where a panel is more than 50.8 mm (2.0 in) wide use at least 4 mounting holes; two at the top and two at the bottom.
3. Dimensions given for height and depth of handles are suggestions only.

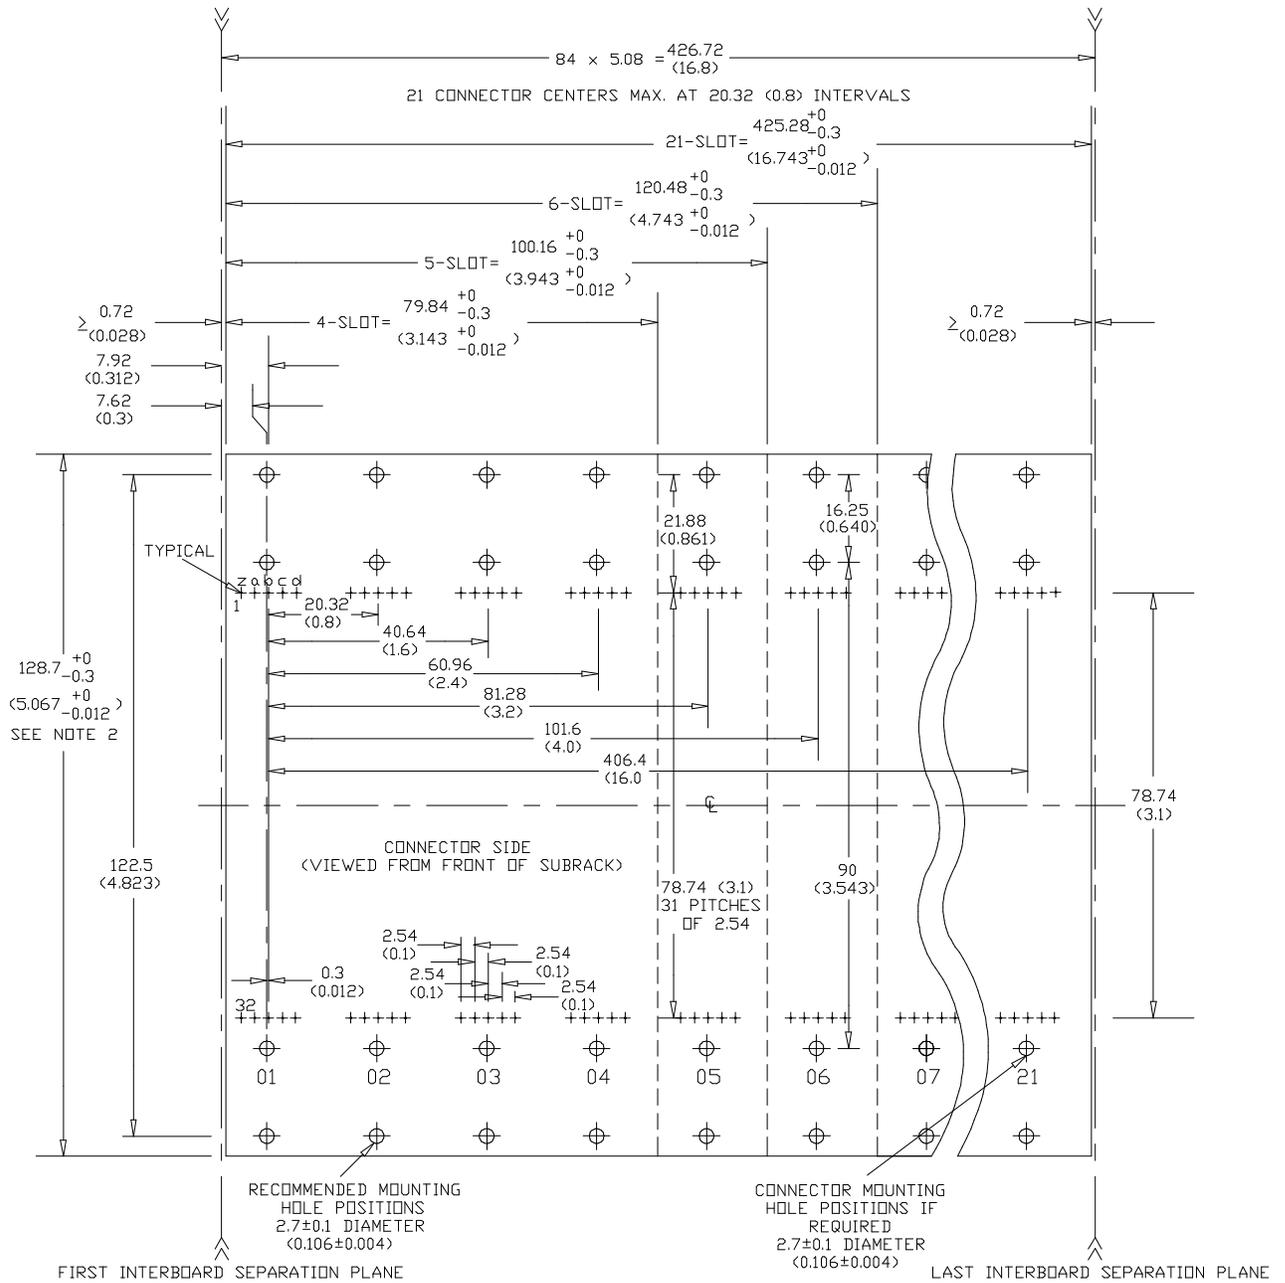
Figure 7 - 12 Single Height Filler Panel



NOTES:

1. All dimensions are shown in millimeters. Inch dimensions are shown in parentheses.
2. **RECOMMENDATION 7.14:**
Where a panel is more than 50.8 mm (2.0 in) wide use at least 4 mounting holes; two at the top and two at the bottom.
3. Dimensions given for height and depth of handles are suggestions only.

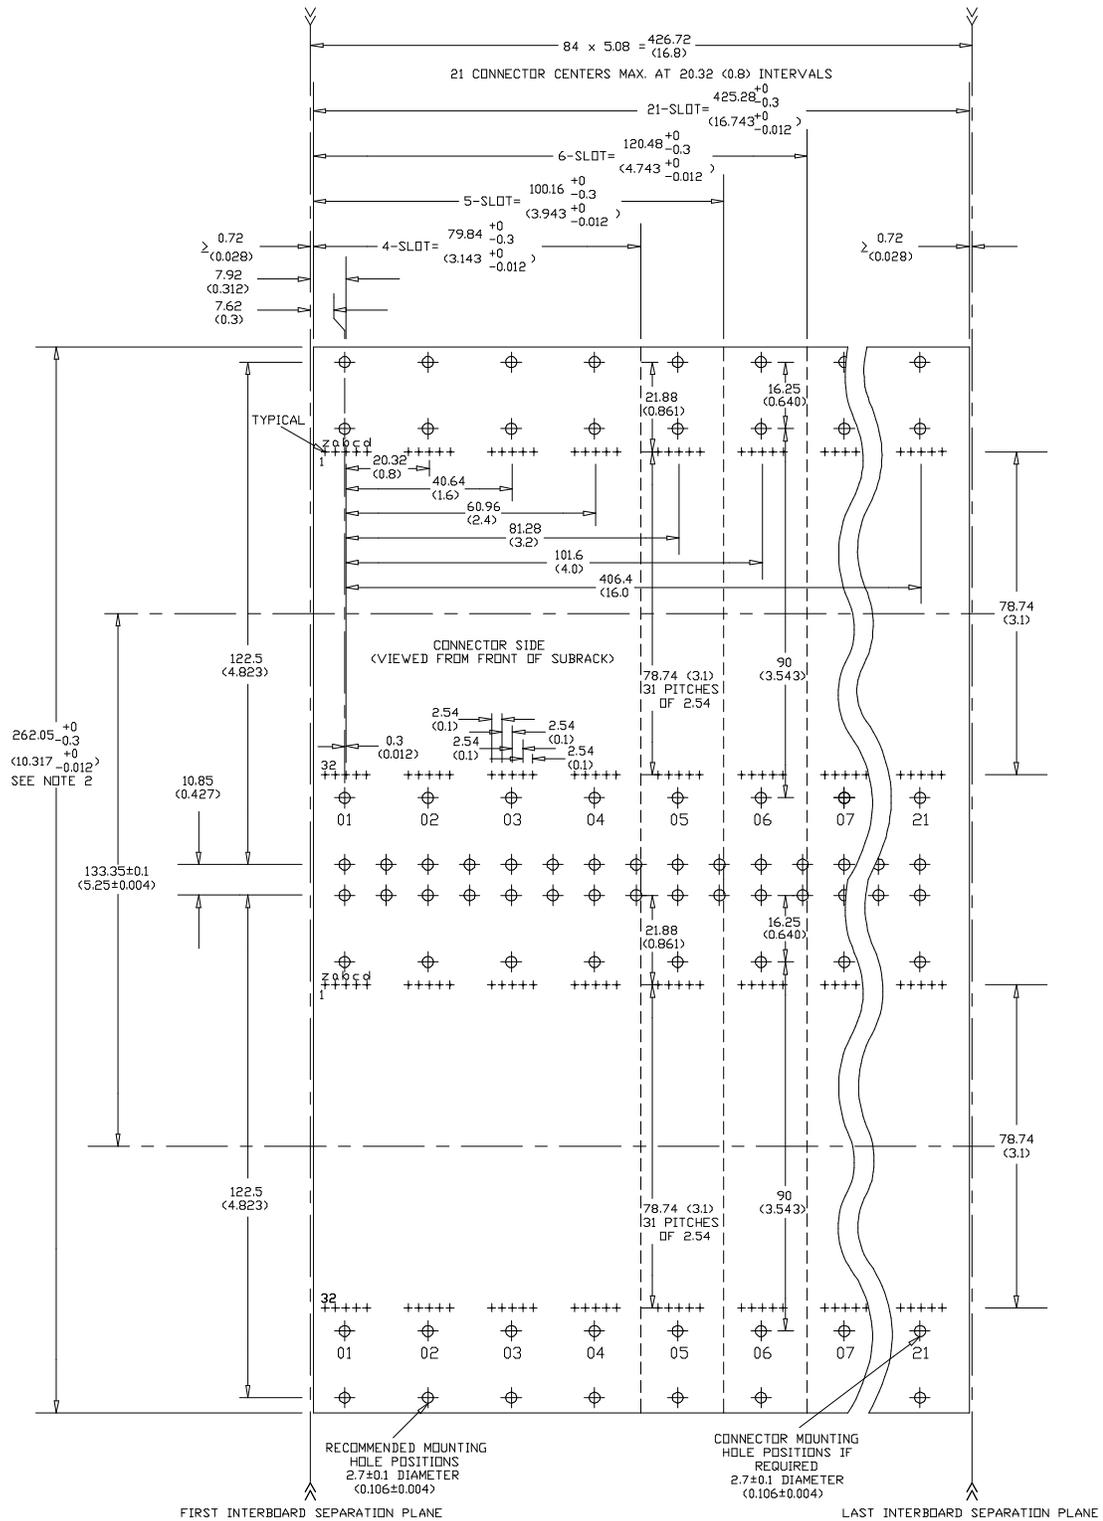
Figure 7 - 13 Double Height Filler Panel



NOTE:

1. All dimensions are shown in millimeters. Inch dimensions are shown in parentheses.
2. **PERMISSION 7.23**
The overall height of a J1 or a J2 backplane may be 130.0 ± 0.3 (5.118 ± 0.012)

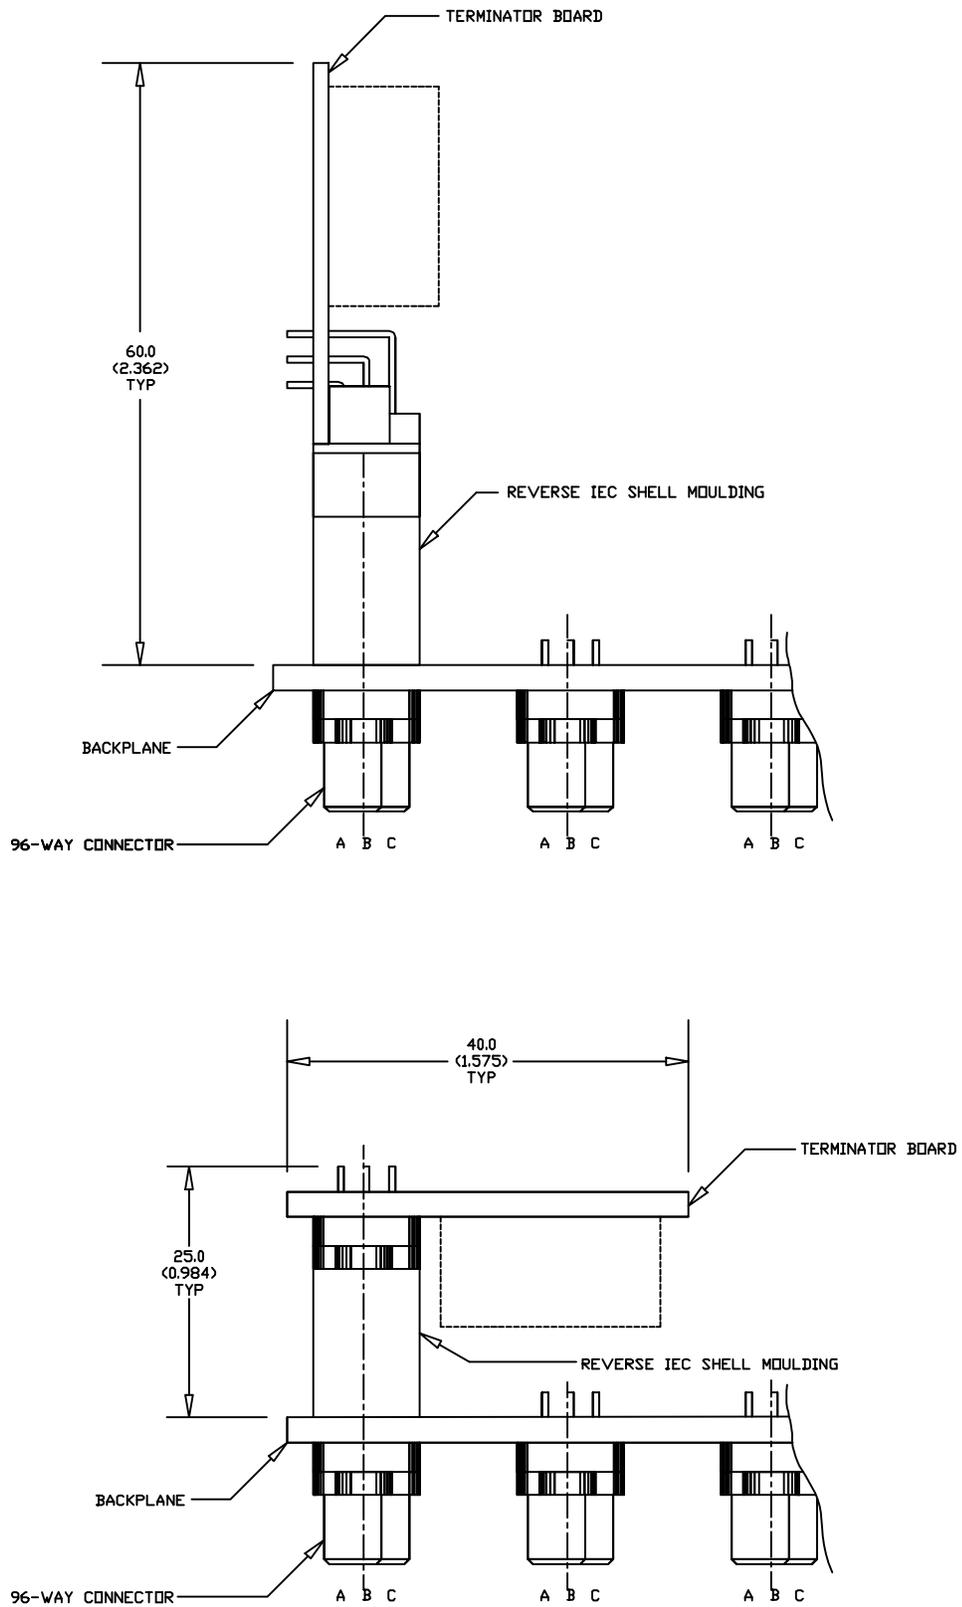
Figure 7 - 14 Detailed Dimensions of a J1 and a J2 Backplane



NOTE:

1. All dimensions are shown in millimeters. Inch dimensions are shown in parentheses.
2. **PERMISSION 7.24**
The overall height of a J1/J2 backplane may be 260.0 ± 0.3 (10.236 ± 0.012).

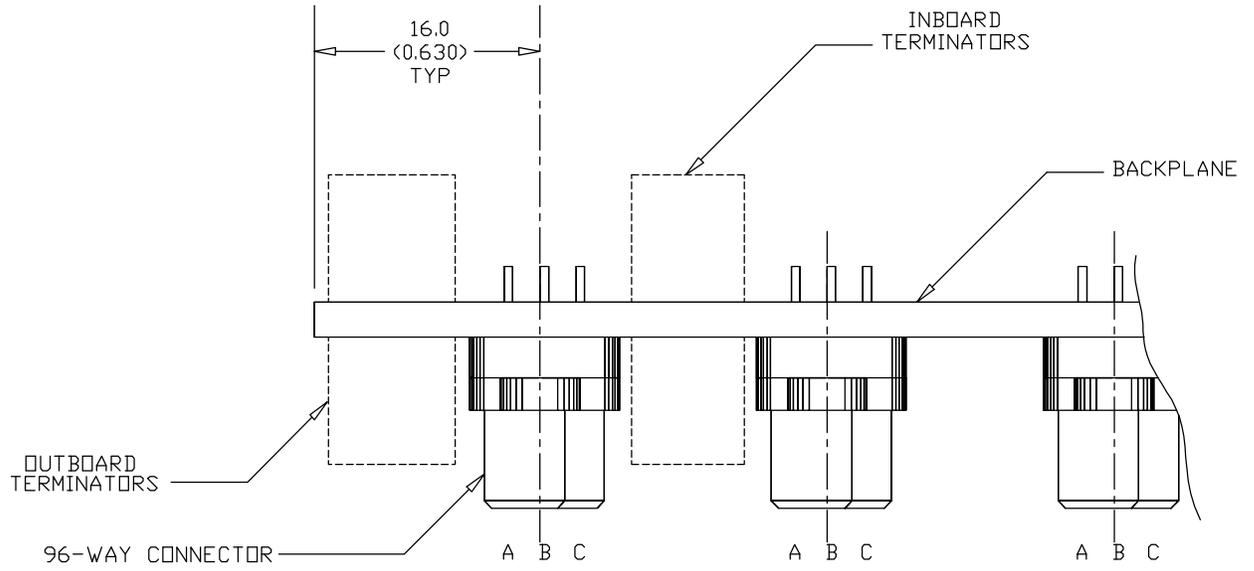
Figure 7 - 15 Detailed Dimensions of a J1/J2 Backplane



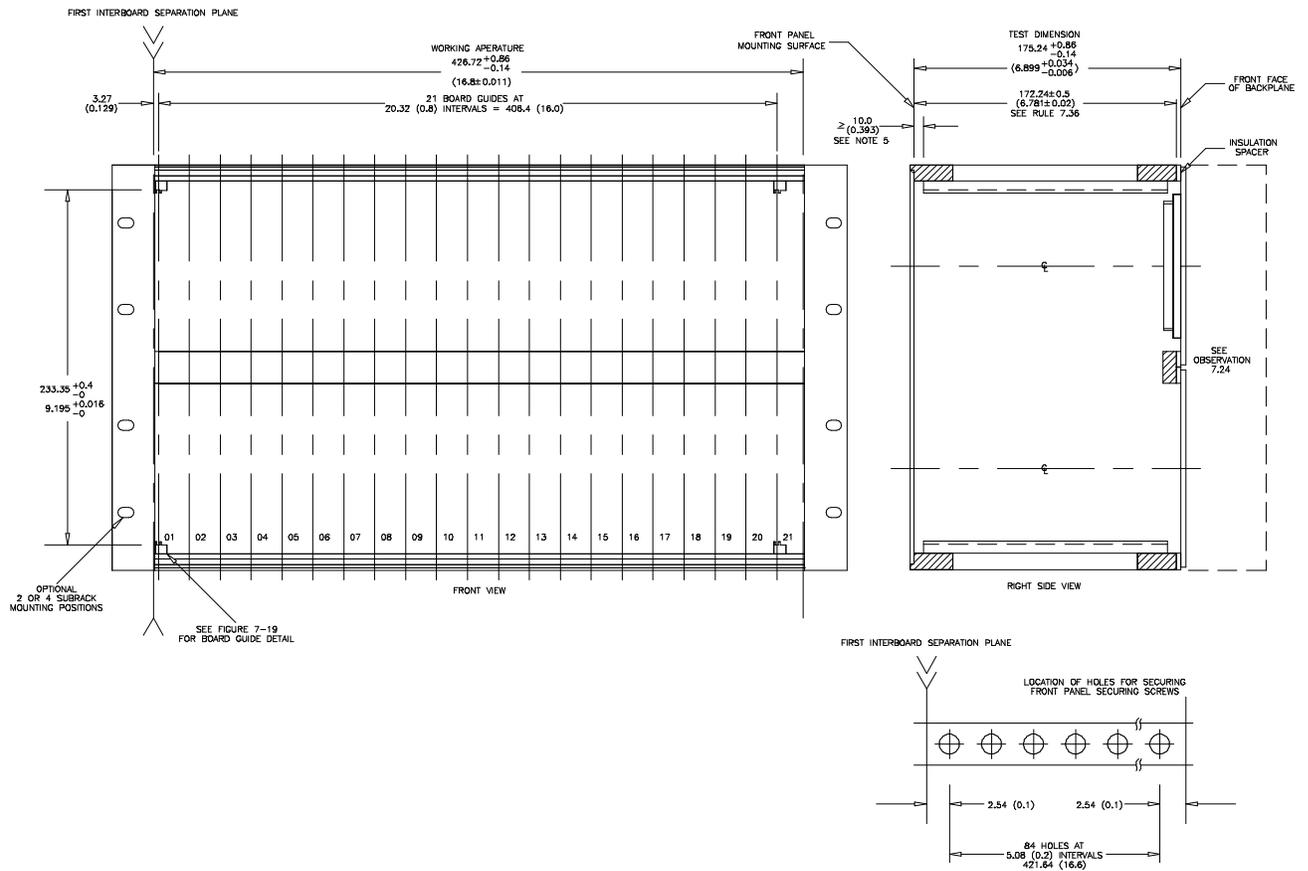
NOTE:

1. All dimensions are shown in millimeters. Inch dimensions are shown in parentheses.

**Figure 7 - 16 "Off-Board Type" Backplane Terminations
(Viewed From Top Of Backplane)**



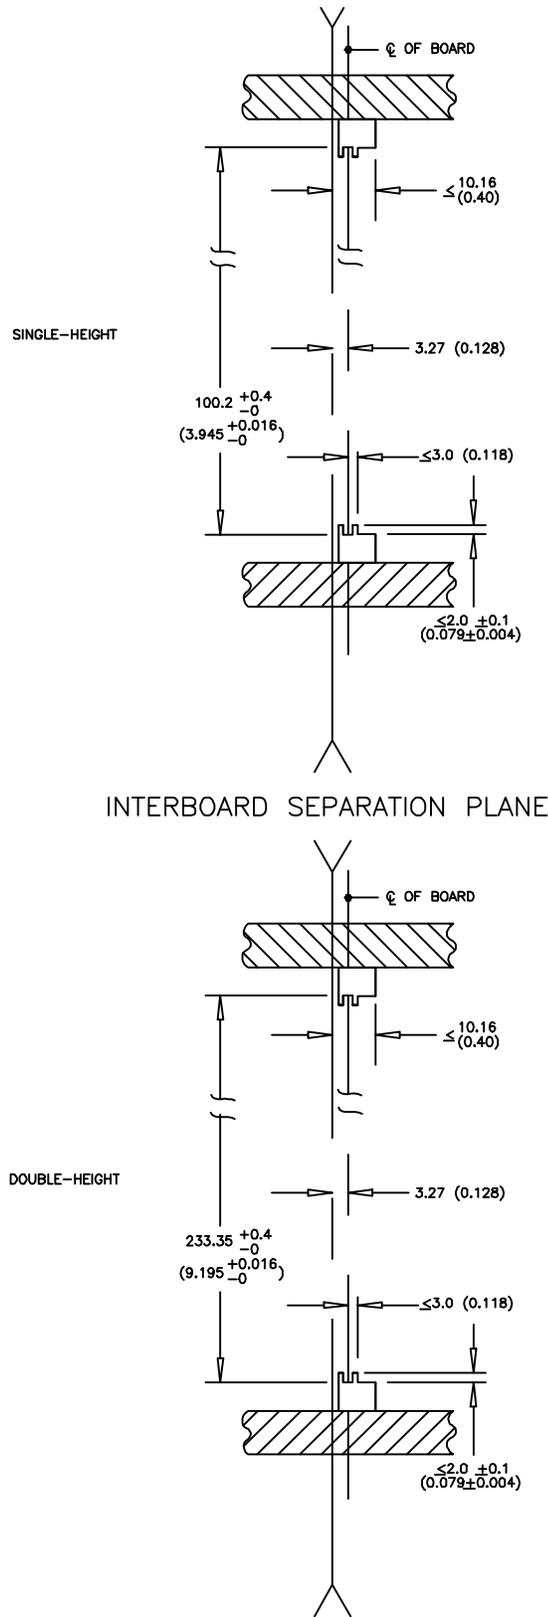
**Figure 7 - 17 "On-Board Type" Backplane Terminations
(Viewed From Top Of Backplane)**



NOTE:

1. All dimensions are shown in millimeters. Inch dimensions are shown in parentheses.
2. **RULE 7.36:**
The thickness of the insulation spacer (shown in the right side view) **MUST** be chosen to assure that the face of the backplane is the correct distance from the front panel mounting surface.
OBSERVATION 7.24:
The thickness of the spacer specified in **RULE 7.36** may vary from one vendor to another.
3. **PERMISSION 7.22:**
Side plates **MAY** be extended to the rear as required.
4. For additional information, see IEC 297-1, IEC 297-3 and IEC 297-4.
5. Changed from ≥ 6 (0.236) to ≥ 10 (0.393) in conformance to IEEE 1101.1 and IEC 297-4.

Figure 7 - 18 21 Slot Sub Rack



NOTE:

1. All dimensions are shown in millimeters. Inch dimensions are shown in parentheses.

Figure 7 - 19 Board Guide Detail

Appendix A

Glossary of VMEBus Terms

- A16 A type of module that provides or decodes an address on address lines A[15..01].
- A24 A type of module that provides or decodes an address on address lines A[23..01].
- A32 A type of module that provides or decodes an address on address lines A[31..01].
- A40 A type of module that provides or decodes an address on address lines A[23..01] and data lines D[15..00] for a total of 40 address bits.
- A64 A type of module that provides or decodes an address on address lines A[31..01] and data lines D[31..00] for a total of 64 address bits.

ARBITRATION

The process of assigning control of the DTB to a Requester.

ADDRESS-ONLY CYCLE

A DTB cycle that consists of an address broadcast, but no data transfer. Slaves do not acknowledge Address-Only cycles and Masters terminate the cycle without waiting for an acknowledgment. No data or acknowledge strobes are asserted in an Address-Only Cycle.

ADDRESS-ONLY-WITH-HANDSHAKE CYCLE

A DTB cycle that consists of an address broadcast, but no data transfer. The addressed Slave responds in the same manner as a standard access cycle.

ARBITER

A functional module that accepts bus requests from Requester modules and grants control of the DTB to one Requester at a time.

ARBITRATION BUS

One of the four buses provided by the VMEbus backplane. This bus allows an Arbiter module and several Requester modules to coordinate use of the DTB.

ARBITRATION CYCLE

An Arbitration Cycle begins when the Arbiter senses a bus request. The Arbiter grants the bus to a Requester, which signals that the DTB is busy. The Requester terminates the cycle by releasing the bus busy signal which causes the Arbiter to sample the bus requests again.

BACKPLANE (VMEbus)

A printed circuit (PC) board with 96-pin connectors and signal paths that bus the connector pins. Some VMEbus systems have a single PC board, called the J1 backplane. It provides the signal paths needed for basic operation. Other VMEbus systems also have an optional second PC board, called a J2 backplane. It provides the additional 96 pin connectors and signal paths needed for wider data and address transfers. Still others have a single monolithic PC board that provides the signal conductors and connectors of both the J1 and J2 backplanes.

BACKPLANE INTERFACE LOGIC

Special interface logic that takes into account the characteristics of the backplane: its signal line impedance, propagation time, termination values, etc. The VMEbus specification prescribes certain rules for the design of this logic based on the maximum length of the backplane and its maximum number of board slots.

BLOCK READ CYCLE

A DTB cycle used to transfer a block of 1 to 256 bytes from a Slave to a Master. This transfer is done using a string of 1, 2, or 4 byte data transfers. Once the block transfer is started, the Master does not release the DTB until all of the bytes have been transferred. It differs from a string of read cycles in that the Master broadcasts only one address and address modifier (at the beginning of the cycle). Then the Slave increments this address on each transfer so that the data for the next cycle is retrieved from the next higher location.

BLOCK WRITE CYCLE

A DTB cycle used to transfer a block of 1 to 256 bytes from a Master to a Slave. The block write cycle is very similar to the block read cycle. It uses a string of 1, 2, or 4 byte data transfers and the Master does not release the DTB until all of the bytes have been transferred. It differs from a string of write cycles in that the Master broadcasts only one address and address modifier (at the beginning of the cycle). Then the Slave increments this address on each transfer so that the next transfer is stored in the next higher location.

BOARD

A printed circuit (PC) board, its collection of electronic components, and either one or two 96-pin connectors that can be plugged into VMEbus backplane connectors.

BUS TIMER

A functional module that measures how long each data transfer takes on the DTB and terminates the DTB cycle if a transfer takes too long. Without this module, if the Master tries to transfer data to or from a non-existent Slave location it could wait forever for a Slave to respond. The Bus Timer prevents this by terminating the cycle.

CR/CSR

A functional module that provides configuration ROM information and Control and Status Registers. The module provides manufacturing and board ID and other important board information. The CSRs are used for software configuration of a VMEbus system.

D08(O)

A Slave that sends and receives data 8-bits at a time over D[7..0],
OR
an Interrupt Handler that receives 8-bit Status/IDs over D[7..0],
OR
an Interrupter that sends 8 bit Status/IDs over D[7..0].

D08(E0)

A Master that sends or receives data 8 bits at a time over either D[7..0] or D[15..8],
OR
a Slave that sends and receives data 8 bits at a time over either D[7..0] or D[15..8].

D16

A Master that sends and receives data 16 bits at a time over D[15..0],
OR
a Slave that sends and receives data 16 bits at a time over D[15..0],
OR
an Interrupt Handler that receives 16-bit Status/IDs over D[15..0],
OR
an Interrupter that sends 16-bit Status/IDs over D[15..0].

D32

A Master that sends and receives data 32 bits at a time over D[31..0],
OR
a Slave that sends and receives data 32 bits at a time over D[31..0],
OR
an Interrupt Handler that receives 32-bit Status/IDs over D[31..0],
OR
an Interrupter that sends 32-bit Status/IDs over D[31..0].

DAISY-CHAIN

A special type of VMEbus signal line that is used to propagate a signal level from board to board, starting with the first slot and ending with the last slot. There are four bus grant daisy-chains and one interrupt acknowledge daisy-chain on the VMEbus backplane.

DATA TRANSFER BUS

One of the four buses provided by the VMEbus backplane. The DATA TRANSFER BUS allows Masters to direct the transfer of binary data between themselves and Slaves. (DATA TRANSFER BUS is often abbreviated DTB.)

DATA TRANSFER BUS CYCLE

A sequence of level transitions on the signal lines of the DTB that result in the transfer of an address or an address and data between a Master and a Slave. There are 34 types of data transfer bus cycles.

DTB

An acronym for DATA TRANSFER BUS.

FUNCTIONAL MODULE

A collection of electronic circuitry that resides on one VMEbus board and works together to accomplish a task.

IACK DAISY-CHAIN DRIVER

A functional module which activates the interrupt acknowledge daisy-chain whenever an Interrupt Handler acknowledges an interrupt request. This daisy-chain ensures that only one Interrupter will respond with its Status/ID when more than one has generated an interrupt request on the same level.

INTERRUPT ACKNOWLEDGE CYCLE

A DTB cycle, initiated by an Interrupt Handler, that reads a "Status/ID" from an Interrupter. An Interrupt Handler generates this cycle when it detects an interrupt request from an Interrupter and it has control of the DTB.

PRIORITY INTERRUPT BUS

One of the four buses provided by the VMEbus backplane. The Priority Interrupt Bus allows Interrupter modules to send interrupt requests to Interrupt Handler modules, and Interrupt Handler modules to acknowledge these interrupt requests.

INTERRUPTER

A functional module that generates an interrupt request on the Interrupt Bus and then provides Status/ID information when the Interrupt Handler requests it.

INTERRUPT HANDLER

A functional module that detects interrupt requests generated by Interrupters and responds to those requests by asking for Status/ID information.

LOCATION MONITOR

A functional module that monitors data transfers over the DTB in order to detect accesses to the locations it has been assigned to watch. When an access occurs to one of these assigned locations, the Location Monitor generates an on-board signal.

MASTER

A functional module that initiates DTB cycles in order to transfer data between itself and a Slave module.

MD32

A Master that sends and receives data 32 bits at a time using D[15..0], A[15..1] and LWORD*. Commonly used by 3U boards.

OR

a Slave that sends and receives data 32 bits at a time using D[15..0], A[15..1] and LWORD*. Commonly used by 3U boards.

MONARCH

The Monarch is the board that controls the assignment of CR/CSR base addresses during the Auto-Slot ID process.

MULTIPLEXED ADDRESS CYCLE

A Multiplexed Address Cycle is a cycle in which both the Address Bus and the Data Bus are used to transfer address information.

MULTIPLEXED DATA CYCLE

A Multiplexed Data Cycle is a cycle in which both the Address Bus and the Data Bus are used to transfer data.

POWER MONITOR MODULE

A functional module that monitors the status of the primary power source to the VMEbus system and signals when the power has strayed outside the limits required for reliable system operation. Since most systems are powered by an AC source, the power monitor is typically designed to detect drop-out or brown-out conditions on AC lines.

READ CYCLE

A DTB cycle used to transfer 1, 2, 3, 4 or 8 bytes from a Slave to a Master. The cycle begins when the Master broadcasts an address and an address modifier. Each Slave captures this address and address modifier, and checks to see if it is to respond to the cycle. If so, it retrieves the data from its internal storage, places it on the data bus, and acknowledges the transfer. Then the Master terminates the cycle.

READ-MODIFY-WRITE CYCLE

A DTB cycle that is used to first read from, and then write to, a Slave location without permitting any other Master to access that location during that cycle. This cycle is most useful in multiprocessing systems where certain memory locations are used to control access to certain systems resources.(i.e., implementation of semaphores).

REQUESTER

A functional module that resides on the same board as a Master or Interrupt Handler and requests use of the DTB whenever its Master or Interrupt Handler needs it.

SERIAL BUS

A functional module that provides a bussed 2-wire interface between cards in the backplane, independent of the other VMEbus modules.

SLAVE

A functional module that detects DTB cycles initiated by a Master and, when those cycles specify its participation, transfers data between itself and the Master.

SLOT

A position where a board can be inserted into a VMEbus backplane. If the VMEbus system has both a J1 and a J2 backplane (or a combination J1/J2 backplane) each slot provides a pair of 96-pin connectors. If the system has only a J1 backplane, then each slot provides a single 96-pin connector.

STATUS/ID

An eight, sixteen, or thirty two bit value returned by an interrupter to an interrupt handler during an interrupt acknowledge cycle.

SUBRACK

A rigid framework that provides mechanical support for boards inserted into the backplane, ensuring that the connectors mate properly and that adjacent boards do not contact each other. It also guides the cooling airflow through the system, and ensures that inserted boards do not disengage themselves from the backplane due to vibration or shock.

SYSTEM CLOCK DRIVER

A functional module that provides a 16 MHz timing signal on the Utility Bus.

SYSTEM CONTROLLER BOARD

A board which resides in slot 1 of a VMEbus backplane and has a System Clock Driver, a DTB Arbiter, an IACK Daisy-Chain Driver, and a Bus Timer. Some also have a Power Monitor.

UNALIGNED TRANSFER (UAT)

A Master that sends or receives data in an unaligned fashion,

OR

A Slave that sends and receives data in an unaligned fashion.

UTILITY BUS

One of the four buses provided by the VMEbus backplane. This bus includes signals that provide periodic timing and coordinates the power-up and power-down of VMEbus systems.

WRITE CYCLE

A DTB cycle used to transfer 1, 2, 3, 4 or 8 bytes from a Master to a Slave. The cycle begins when the Master broadcasts an address and address modifier and places data on the DTB. Each Slave captures this address and address modifier, and checks to see if it is to respond to the cycle. If so, it stores the data and then acknowledges the transfer. The Master then terminates the cycle.

PAGE LEFT BLANK

Appendix B

VMEBus Connector/Pin Description

Introduction

This appendix describes the VMEbus signal lines. The following table identifies the VMEbus signals by signal mnemonic, and describes the signal characteristics.

VMEbus Signal Identification

SIGNAL

MNEMONIC

SIGNAL NAME AND DESCRIPTION

A[15..1]	Address bus (bits 15-1) - Three-state driven address lines that are used to broadcast an A16, A24, A32, A40 or A64 address and may be used to carry data for MBLT and MD32.
A[23..16]	Address bus (bits 23-16) - Three-state driven address lines that are used in conjunction with A[15..1] to broadcast an A24, A32, A40 or A64 address and may be used to carry data for MBLT.
A[31..24]	Address bus (bits 31-24) - Three-state driven address lines that are used in conjunction with A[23..1] to broadcast an A32 or A64 address and may be used carry data for MBLT.
ACFAIL*	AC Failure - An open-collector driven signal which indicates that the AC input to the power supply is no longer being provided or that the required AC input voltage levels are not being met.
AM[5..0]	Address Modifier (bits 5-0) - Three-state driven lines that are used to broadcast information such as address size, cycle type, and/or Master identification.
AS*	Address Strobe - A three-state driven signal that indicates when a valid address has been placed on the address bus.
BBSY*	Bus Busy - An open-collector driven signal driven low by the current Master to indicate that it is using the bus. When the Master releases this line, the resultant rising edge causes the Arbiter to sample the bus grant lines and grant the bus to the highest priority requester.
BCLR*	Bus Clear - A totem-pole driven signal, generated by an Arbiter to indicate when there is a higher priority request for the bus. This signal requests the current Master to release the DTB .
BERR*	Bus Error - An open-collector driven signal generated by a Slave or Bus Timer. This signal indicates to the Master that the data transfer was not completed or the data contained an error.

SIGNAL

MNEMONIC

SIGNAL NAME AND DESCRIPTION

BG[3..0]IN*	Bus Grant [3..0]In - Totem-pole driven signals generated by the Arbiter and Requesters. "Bus Grant In" and "Bus Grant Out" signals form bus grant daisy chains. The "Bus Grant In" signal indicates, to the board receiving it, that it may use the DTB.
BG[3..0]OUT*	Bus Grant (0-3) Out - Totem-pole driven signals generated by Requesters. The Bus Grant Out signal indicates to the next board in the daisy-chain that it may use the DTB.
BR[3..0]*	Bus Request (0-3) - Open-collector driven signal generated by RequestERS. A low level on one of these lines indicates that some Master needs to use the DTB.
D[31..0]	Data Bus - Three-state driven bidirectional data lines used to transfer data between Masters and Slaves, and Status/ID information between Interrupters and Interrupt Handlers. Some or all of these lines may also be used to broadcast addresses in the A40 and A64 modes.
DS0*, DS1*	Data Strobe Zero, One - Three-state driven signals used in conjunction with LWORD* and A1 to indicate how many data bytes are being transferred (1, 2, 3, or 4). During a write cycle, the falling edge of the first data strobe indicates that valid data is available on the data bus. On a read cycle, the rising edge of the first data strobe indicates that data has been accepted from the data bus.
DTACK*	Data Transfer Acknowledge - An open-collector driven signal generated by a Slave. The falling edge of this signal indicates that valid data is available on the data bus during a read cycle, or that data has been accepted from the data bus during a write cycle. The rising edge indicates when the Slave has released the data bus at the end of a Read Cycle.
GND	The DC voltage reference for the VMEbus system.
IACK*	Interrupt Acknowledge - An open-collector or three-state driven signal used by an Interrupt Handler acknowledging an interrupt request. It is routed, via a backplane signal trace, to the IACKIN* pin of Slot 1, where it is monitored by the IACK Daisy-Chain Driver.
IACKIN*	Interrupt Acknowledge In - A totem-pole driven signal. The IACKIN* and IACKOUT* signals form a daisy-chain. The IACKIN* signal indicates to the VMEbus board receiving it that it is allowed to respond to the Interrupt Acknowledge Cycle that is in progress.
IACKOUT*	Interrupt Acknowledge Out - A totem-pole driven signal. The IACKIN* and IACKOUT* signals form a daisy-chain. The IACKOUT* signal is sent by a board to indicate to the next board in the daisy-chain that it is allowed to respond to the Interrupt Acknowledge Cycle that is in progress.

SIGNAL MNEMONIC	SIGNAL NAME AND DESCRIPTION
IRQ[7..1]*	Interrupt Request (7-1) - Open-collector driven signals, generated by an Interrupter, which carry interrupt requests. When several lines are monitored by a single Interrupt Handler the highest numbered line is given the highest priority.
LWORD*	Longword - A three-state driven signal used in conjunction with DS0*, DS1*, and A01 to select which byte location(s) within the 4-byte group are accessed during the data transfer. This line may also be used to transfer data during the data phase of MBLT and MD32 cycles.
RETRY*	A rescinding line driven by a Slave to inform the Master that it cannot complete the cycle at the time and that the Master should try again later.
SERA	Used by Serial Bus modules to transfer encoded data.
SERB	Used by Serial Bus modules to transfer encoded data.
SYSCLK	System Clock - A totem-pole driven signal which provides a constant 16-MHz clock signal that is independent of any other bus timing.
SYSFAIL*	System Fail - An open-collector driven signal that indicates that a failure has occurred in the system. This signal may be generated by any board on the VMEbus.
SYSRESET*	System Reset - An open-collector driven signal which, when low, causes the system to be reset.
WRITE*	Write - A three-state driven signal generated by the Master to indicate whether the data transfer cycle is a read or a write. A high level indicates a read operation; a low level indicates a write operation.
+5V STDBY	+5 VDC Standby - This line supplies +5 VDC to devices requiring battery backup.
+5V	+5 VDC Power - Used by system logic circuits.
+12V	+12 VDC Power - Used by system logic circuits.
-12V	-12 VDC Power - Used by system logic circuits.

PAGE LEFT BLANK

Appendix C

Manufacturer's Board Identification

The manufacturer's board ID is a three octet (24 bit) identifier assigned to manufacturers by the IEEE Registration Authority. This identifier is referred to as the Organizationally Unique Identifier by the IEEE. The IEEE has been designated by the ISO Council to act as the world wide registration authority. Identifiers are available from the IEEE for a cost determined by the IEEE.

For information and to obtain an identifier contact the

IEEE Registration Authority
IEEE Standards Department
445 Hoes Lane, P.O. Box 1331
Piscataway NJ 08855-1331
Phone: 908 562 3813
Fax: 908 562 1571

The 24 bit identifier is used in the Configuration ROM. The most significant byte is placed at location 0x27. The second byte is placed at location 0x2B and the least significant byte is placed at location 0x2F.

Octet 1		Octet 2		Octet 3	
MSB	LSB	MSB	LSB	MSB	LSB
23	16	15	8	7	0

Configuration ROM	
0x27	Octet 1
0x2B	Octet 2
0x2F	Octet 3

PAGE LEFT BLANK

Appendix D

RULE INDEX

Rule 2.1a.....	17	Rule 2.42.....	86
Rule 2.2.....	20	Rule 2.43.....	86
Rule 2.3.....	20	Rule 2.44a.....	86
Rule 2.4.....	40	Rule 2.45.....	86
Rule 2.5.....	40	Rule 2.46a.....	86
Rule 2.6.....	49	Rule 2.47a.....	86
Rule 2.7.....	28	Rule 2.48a.....	86
Rule 2.8.....	28	Rule 2.49.....	87
Rule 2.9.....	37	Rule 2.50.....	87
Rule 2.10.....	38	Rule 2.51.....	87
Rule 2.11.....	38	Rule 2.52.....	87
Rule 2.12a.....	42	Rule 2.53a.....	91
Rule 2.13a.....	58	Rule 2.54a.....	91
Rule 2.14a.....	58	Rule 2.55.....	91
Rule 2.15a.....	58	Rule 2.56a.....	91
Rule 2.16.....	58	Rule 2.57.....	91
Rule 2.17a.....	59	Rule 2.58a.....	92
Rule 2.18.....	67	Rule 2.59.....	94
Rule 2.19.....	68	Rule 2.60.....	94
Rule 2.20.....	68	Rule 2.61.....	Not Used *
Rule 2.21.....	74	Rule 2.62.....	Not Used *
Rule 2.22.....	74	Rule 2.63.....	NotUsed *
Rule 2.23.....	74	Rule 2.64.....	Not Used *
Rule 2.24.....	74	Rule 2.65.....	Not Used *
Rule 2.25.....	74	Rule 2.66.....	43
Rule 2.26.....	74	Rule 2.67.....	Not Used *
Rule 2.27.....	84	Rule 2.68.....	51
Rule 2.28.....	84	Rule 2.69.....	17
Rule 2.29.....	84	Rule 2.70.....	23
Rule 2.30.....	84	Rule 2.71.....	37
Rule 2.31.....	84	Rule 2.72.....	39
Rule 2.32.....	84	Rule 2.73.....	39
Rule 2.33a.....	84	Rule 2.74.....	39
Rule 2.34a.....	85	Rule 2.75.....	39
Rule 2.35.....	85	Rule 2.76.....	39
Rule 2.36.....	85	Rule 2.77.....	39
Rule 2.37.....	85	Rule 2.78.....	42
Rule 2.38.....	85	Rule 2.79.....	43
Rule 2.39.....	85	Rule 2.80.....	43
Rule 2.40.....	85	Rule 2.81.....	51
Rule 2.41.....	85	Rule 2.82.....	52
		Rule 2.83.....	52
		Rule 2.84.....	52
		Rule 2.85.....	53
		Rule 2.86.....	53
		Rule 2.87.....	54

Rule 2.88.....	56	Rule 4.12.....	178
Rule 2.89.....	57	Rule 4.13.....	178
Rule 2.90.....	57	Rule 4.14.....	178
Rule 2.91.....	57	Rule 4.15.....	181
Rule 2.92.....	59	Rule 4.16.....	181
Rule 2.93.....	67	Rule 4.17.....	181
Rule 2.94.....	86	Rule 4.18.....	181
Rule 2.95.....	88	Rule 4.19.....	181
Rule 2.96.....	88	Rule 4.20.....	181
Rule 2.97.....	91	Rule 4.21.....	181
Rule 2.98.....	91	Rule 4.22.....	182
Rule 2.99.....	92	Rule 4.23.....	182
Rule 2.100.....	92	Rule 4.24.....	182
Rule 2.101.....	92	Rule 4.25.....	182
Rule 2.102.....	92	Rule 4.26.....	182
Rule 2.103.....	92	Rule 4.27.....	182
Rule 2.104.....	92	Rule 4.28.....	182
Rule 2.105.....	92	Rule 4.29.....	182
		Rule 4.30.....	182
Rule 3.1.....	123	Rule 4.31.....	183
Rule 3.2.....	125	Rule 4.32.....	183
Rule 3.3.....	135	Rule 4.33.....	183
Rule 3.4.....	126	Rule 4.34.....	183
Rule 3.5.....	126	Rule 4.35.....	185
Rule 3.6.....	127	Rule 4.36.....	185
Rule 3.7.....	130	Rule 4.37.....	185
Rule 3.8.....	130	Rule 4.38.....	186
Rule 3.9.....	130	Rule 4.39.....	186
Rule 3.10.....	130	Rule 4.40.....	186
Rule 3.11.....	131	Rule 4.41.....	186
Rule 3.12.....	134	Rule 4.42.....	186
Rule 3.13.....	144	Rule 4.43.....	186
Rule 3.14.....	131	Rule 4.44.....	186
Rule 3.15.....	127	Rule 4.45.....	186
		Rule 4.46.....	188
Rule 4.1.....	149	Rule 4.47.....	188
Rule 4.2.....	150	Rule 4.48.....	188
Rule 4.3.....	151	Rule 4.49.....	174
Rule 4.4.....	154	Rule 4.50.....	149
Rule 4.5.....	160		
Rule 4.6.....	160	Rule 5.1.....	201
Rule 4.7.....	161	Rule 5.2.....	201
Rule 4.8.....	161	Rule 5.3.....	204
Rule 4.9.....	178	Rule 5.4.....	204
Rule 4.10.....	178	Rule 5.5.....	204
Rule 4.11.....	178	Rule 5.6.....	204

Rule 5.7.....	205	Rule 7.3.....	233
Rule 5.8.....	206	Rule 7.4.....	233
Rule 5.9.....	207	Rule 7.5.....	233
Rule 5.10.....	207	Rule 7.6.....	234
Rule 5.11.....	199	Rule 7.7.....	234
Rule 5.12.....	209	Rule 7.8.....	234
Rule 5.13.....	209	Rule 7.9.....	234
Rule 5.14.....	209	Rule 7.10.....	236
Rule 5.15.....	209	Rule 7.11.....	236
Rule 5.16.....	209	Rule 7.12.....	236
Rule 5.17.....	211	Rule 7.13.....	237
Rule 5.18.....	211	Rule 7.14.....	238
Rule 5.19.....	216	Rule 7.15.....	238
Rule 5.20.....	216	Rule 7.16.....	239
Rule 5.21.....	216	Rule 7.17.....	240
Rule 5.22.....	216	Rule 7.18.....	241
		Rule 7.19.....	241
Rule 6.1.....	217	Rule 7.20.....	241
Rule 6.2.....	217	Rule 7.21.....	241
Rule 6.3.....	217	Rule 7.22.....	241
Rule 6.4.....	217	Rule 7.23.....	242
Rule 6.5.....	217	Rule 7.24.....	242
Rule 6.6.....	219	Rule 7.25.....	242
Rule 6.7.....	219	Rule 7.26.....	243
Rule 6.8.....	219	Rule 7.27.....	244
Rule 6.9.....	220	Rule 7.28.....	244
Rule 6.10.....	221	Rule 7.29.....	244
Rule 6.11.....	221	Rule 7.30.....	245
Rule 6.12.....	221	Rule 7.31.....	245
Rule 6.13.....	222	Rule 7.32.....	249
Rule 6.14.....	222	Rule 7.33.....	250
Rule 6.15.....	222	Rule 7.34.....	254
Rule 6.16.....	223	Rule 7.35.....	254
Rule 6.17.....	223	Rule 7.36.....	265
Rule 6.18.....	223	Rule 7.37.....	236
Rule 6.19.....	224	Rule 7.38.....	237
Rule 6.20.....	224	Rule 7.39.....	237
Rule 6.21.....	225	Rule 7.40.....	237
Rule 6.22.....	225	Rule 7.41.....	237
Rule 6.23a.....	226	Rule 7.42.....	242
Rule 6.24.....	229	Rule 7.43.....	242
Rule 6.25.....	229	Rule 7.44.....	243
Rule 6.26.....	229	Rule 7.45.....	243
		Rule 7.46.....	245
Rule 7.1.....	233	Rule 7.47.....	245
Rule 7.2.....	233		

Observation 1.1.....	11	Observation 2.45.....	87
Observation 2.2.....	11	Observation 2.46a.....	87
Observation 2.1.....	14	Observation 2.47a.....	87
Observation 2.2.....	17	Observation 2.48.....	87
Observation 2.3.....	17	Observation 2.49a.....	88
Observation 2.4.....	17	Observation 2.50.....	88
Observation 2.5.....	20	Observation 2.51a.....	88
Observation 2.6.....	Not Used	Observation 2.52.....	89
Observation 2.7.....	26	Observation 2.53.....	89
Observation 2.8.....	26	Observation 2.54.....	89
Observation 2.9a.....	26	Observation 2.55a.....	89
Observation 2.10.....	Not Used *	Observation 2.56a.....	89
Observation 2.11.....	28	Observation 2.57.....	89
Observation 2.12.....	28	Observation 2.58.....	89
Observation 2.13.....	34	Observation 2.59.....	89
Observation 2.14a.....	37	Observation 2.60.....	89
Observation 2.15a.....	38	Observation 2.61.....	89
Observation 2.16.....	Not Used *	Observation 2.62.....	90
Observation 2.17.....	Not Used *	Observation 2.63.....	90
Observation 2.18a.....	42	Observation 2.64.....	90
Observation 2.19a.....	43	Observation 2.65.....	90
Observation 2.20a.....	43	Observation 2.66a.....	90
Observation 2.21.....	43	Observation 2.67.....	90
Observation 2.22.....	47	Observation 2.68a.....	90
Observation 2.23.....	48	Observation 2.69a.....	90
Observation 2.24.....	49	Observation 2.70a.....	90
Observation 2.25.....	50	Observation 2.71.....	91
Observation 2.26a.....	58	Observation 2.72.....	91
Observation 2.27.....	58	Observation 2.73.....	92
Observation 2.28.....	Not Used *	Observation 2.74.....	92
Observation 2.29.....	59	Observation 2.75.....	93
Observation 2.30a.....	59	Observation 2.76.....	93
Observation 2.31.....	60	Observation 2.77.....	93
Observation 2.32.....	67	Observation 2.78.....	93
Observation 2.33.....	67	Observation 2.79.....	93
Observation 2.34.....	68	Observation 2.80.....	93
Observation 2.35.....	84	Observation 2.81.....	94
Observation 2.36.....	84	Observation 2.82.....	94
Observation 2.37.....	84	Observation 2.83.....	94
Observation 2.38.....	84	Observation 2.84.....	94
Observation 2.39a.....	85	Observation 2.85.....	94
Observation 2.40.....	85	Observation 2.86.....	14
Observation 2.41.....	85	Observation 2.87.....	14
Observation 2.42.....	86	Observation 2.88.....	38
Observation 2.43.....	87	Observation 2.89.....	39
Observation 2.44.....	87	Observation 2.90.....	42

Observation 2.91.....	43	Observation 4.4.....	154
Observation 2.92.....	43	Observation 4.5.....	160
Observation 2.93.....	43	Observation 4.6.....	161
Observation 2.94.....	51	Observation 4.7.....	161
Observation 2.95.....	52	Observation 4.8.....	163
Observation 2.96.....	52	Observation 4.9.....	167
Observation 2.97.....	56	Observation 4.10.....	167
Observation 2.98.....	56	Observation 4.11.....	169
Observation 2.99.....	57	Observation 4.12.....	181
Observation 2.100.....	67	Observation 4.13.....	181
Observation 2.101.....	67	Observation 4.14.....	181
Observation 2.102.....	86	Observation 4.15.....	182
Observation 2.103.....	88	Observation 4.16.....	183
Observation 2.104.....	88	Observation 4.17.....	183
Observation 2.105.....	88	Observation 4.18.....	183
Observation 2.106.....	88	Observation 4.19.....	183
Observation 2.107.....	90	Observation 4.20.....	183
Observation 2.108.....	91	Observation 4.21.....	183
Observation 2.109.....	92	Observation 4.22.....	184
Observation 2.110.....	93	Observation 4.23.....	184
Observation 2.111.....	93	Observation 4.24.....	184
		Observation 4.25.....	184
Observation 3.1.....	123	Observation 4.26.....	184
Observation 3.2a.....	125	Observation 4.27.....	184
Observation 3.3.....	126	Observation 4.28.....	184
Observation 3.4.....	126	Observation 4.29.....	184
Observation 3.5.....	126	Observation 4.30.....	185
Observation 3.6.....	127	Observation 4.31.....	185
Observation 3.7.....	127	Observation 4.32.....	185
Observation 3.8.....	130	Observation 4.33.....	185
Observation 3.9.....	130	Observation 4.34.....	185
Observation 3.10.....	133	Observation 4.35.....	185
Observation 3.11.....	134	Observation 4.36.....	185
Observation 3.12.....	134	Observation 4.37.....	185
Observation 3.13.....	135	Observation 4.38.....	186
Observation 3.14.....	135	Observation 4.39.....	186
Observation 3.15.....	139	Observation 4.40.....	187
Observation 3.16.....	144	Observation 4.41.....	187
Observation 3.17.....	130	Observation 4.42.....	187
Observation 3.18.....	130	Observation 4.43.....	187
Observation 3.19.....	127	Observation 4.44.....	188
Observation 3.20.....	131	Observation 4.45.....	188
		Observation 4.46.....	188
Observation 4.1a.....	150	Observation 4.47.....	188
Observation 4.2.....	151	Observation 4.48.....	188
Observation 4.3.....	151	Observation 4.49.....	188

Observation 4.50.....	174	Observation 7.6.....	235
Observation 4.51.....	160	Observation 7.7.....	235
Observation 4.52.....	182	Observation 7.8.....	236
Observation 4.53.....	184	Observation 7.9.....	236
Observation 4.54.....	186	Observation 7.10.....	238
Observation 5.1.....	199	Observation 7.11.....	238
Observation 5.2.....	204	Observation 7.12.....	238
Observation 5.3.....	204	Observation 7.13.....	239
Observation 5.4.....	207	Observation 7.14.....	240
Observation 5.5.....	207	Observation 7.15.....	240
Observation 5.6.....	207	Observation 7.16.....	240
Observation 5.7.....	Not Used *	Observation 7.17.....	241
Observation 5.8.....	207	Observation 7.18.....	242
Observation 5.9.....	209	Observation 7.19.....	242
Observation 5.10.....	210	Observation 7.20.....	244
Observation 5.11.....	210	Observation 7.21.....	244
Observation 5.12.....	210	Observation 7.22.....	244
Observation 5.13.....	211	Observation 7.23.....	245
Observation 5.14.....	216	Observation 7.24.....	265
Observation 6.1.....	218	Permission 2.1.....	17
Observation 6.2.....	219	Permission 2.2.....	17
Observation 6.3.....	219	Permission 2.3.....	20
Observation 6.4.....	221	Permission 2.4.....	23
Observation 6.5.....	222	Permission 2.5a.....	26
Observation 6.6.....	223	Permission 2.6.....	27
Observation 6.7.....	223	Permission 2.7.....	38
Observation 6.8.....	224	Permission 2.8a.....	67
Observation 6.9.....	225	Permission 2.9.....	68
Observation 6.10.....	226	Permission 2.10.....	74
Observation 6.11.....	226	Permission 2.11.....	74
Observation 6.12.....	226	Permission 2.12.....	74
Observation 6.13.....	228	Permission 2.13.....	74
Observation 6.14.....	229	Permission 2.14.....	27
Observation 6.15.....	229	Permission 2.15.....	51
Observation 6.16.....	229	Permission 2.16.....	54
Observation 6.17.....	222	Permission 2.17.....	57
Observation 6.18.....	224	Permission 2.18.....	57
Observation 6.19.....	225	Permission 2.19.....	91
Observation 7.1.....	234	Permission 3.1.....	121
Observation 7.2.....	234	Permission 3.2.....	123
Observation 7.3.....	234	Permission 3.3.....	125
Observation 7.4.....	235	Permission 3.4.....	127
Observation 7.5.....	235	Permission 3.5.....	131
		Permission 3.6.....	134

Permission 3.7.....	134	Permission 7.20.....	255
Permission 3.8.....	144	Permission 7.21.....	256
Permission 4.1.....	153	Permission 7.22.....	265
Permission 4.2.....	161	Permission 7.23.....	261
Permission 4.3.....	162	Permission 7.24.....	262
Permission 4.4.....	163	Suggestion 2.1.....	Not Used *
Permission 4.5.....	163	Suggestion 2.2.....	Not Used *
Permission 4.6.....	163	Suggestion 2.3a.....	58
Permission 4.7.....	178	Suggestion 2.4.....	67
Permission 4.8.....	178	Suggestion 2.5a.....	68
Permission 4.9.....	178	Suggestion 2.6a.....	39
Permission 4.10.....	178	Suggestion 2.7.....	Not Used *
Permission 4.11.....	174	Suggestion 2.8a.....	40
Permission 5.1.....	201	Suggestion 2.9.....	40
Permission 5.2.....	201	Suggestion 5.1.....	206
Permission 5.3.....	204	Suggestion 5.2.....	206
Permission 5.4.....	204/206 **	Suggestion 5.3.....	206
Permission 5.5.....	209	Suggestion 5.4.....	207
Permission 5.6.....	210	Suggestion 5.5.....	207
Permission 5.7.....	210	Suggestion 6.1.....	Not Used *
Permission 6.1.....	217	Suggestion 6.2.....	Not Used *
Permission 6.2.....	221	Suggestion 6.3.....	Not Used *
Permission 6.3.....	226	Suggestion 6.4.....	Not Used *
Permission 7.1.....	233	Suggestion 6.5.....	Not Used *
Permission 7.2.....	234	Suggestion 6.6.....	Not Used *
Permission 7.3.....	234	Suggestion 6.6.....	225
Permission 7.4.....	235	Suggestion 6.7.....	229
Permission 7.5.....	235	Suggestion 6.8.....	229
Permission 7.6.....	236	Suggestion 6.9.....	224
Permission 7.7.....	236	Suggestion 7.1.....	235
Permission 7.8.....	237	Suggestion 7.2.....	237
Permission 7.9.....	237	Suggestion 7.3.....	238
Permission 7.10.....	238	Suggestion 7.4.....	239
Permission 7.11.....	239	Suggestion 7.5.....	239
Permission 7.12.....	239	Suggestion 7.6.....	241
Permission 7.13.....	240	Suggestion 7.7.....	241
Permission 7.14.....	242	Suggestion 7.8.....	243
Permission 7.15.....	243	Suggestion 7.9.....	243
Permission 7.16.....	243	Suggestion 7.10.....	243
Permission 7.17.....	244	Suggestion 7.11.....	245
Permission 7.18.....	244	Suggestion 7.12.....	254
Permission 7.19.....	244	Recommendation 2.1.....	23
		Recommendation 2.2.....	27
		Recommendation 2.3.....	39

Recommendation 2.4.....	39
Recommendation 2.5.....	39
Recommendation 2.6.....	49
Recommendation 2.7.....	52
Recommendation 2.8.....	52
Recommendation 2.9.....	52
Recommendation 2.10.....	57
Recommendation 2.11.....	58
Recommendation 2.12.....	59
Recommendation 2.13.....	94
Recommendation 3.1.....	133
Recommendation 3.2.....	134
Recommendation 3.3.....	131
Recommendation 4.1.....	163
Recommendation 4.2.....	160
Recommendation 5.1.....	202
Recommendation 5.2.....	204
Recommendation 5.3.....	209
Recommendation 5.4.....	210
Recommendation 5.5.....	210
Recommendation 6.1.....	218
Recommendation 6.2.....	219
Recommendation 6.3.....	226
Recommendation 6.4.....	Not Used *
Recommendation 6.5.....	230
Recommendation 7.1.....	234
Recommendation 7.2.....	239
Recommendation 7.3.....	239
Recommendation 7.4.....	239
Recommendation 7.5.....	240
Recommendation 7.6.....	240
Recommendation 7.7.....	241
Recommendation 7.8.....	241
Recommendation 7.9.....	243
Recommendation 7.10.....	244
Recommendation 7.11.....	255
Recommendation 7.12.....	256
Recommendation 7.13.....	259
Recommendation 7.14.....	260
Recommendation 7.15.....	243

* Obsolete and not used in this specification

** Error in number from previous editions has been left unchanged in this edition,