

CRAY T3D™ Hardware Reference Manual

Volume 2 of 2

CMM-0602-0A0

Cray Research Proprietary

Cray Research, Inc.

Any shipment to a country outside of the United States requires a letter of assurance from Cray Research, Inc.

This document is the property of Cray Research, Inc. The use of this document is subject to specific license rights extended by Cray Research, Inc. to the owner or lessee of a Cray Research, Inc. computer system or other licensed party according to the terms and conditions of the license and for no other purpose.

Cray Research, Inc. Unpublished Proprietary Information — All Rights Reserved.

Autotasking, CF77, CRAY, Cray Ada, CRAY Y-MP, CRAY-1, HSX, SSD, UniChem, UNICOS, and X-MPEA are federally registered trademarks and CCI, CF90, CFT, CFT2, CFT77, COS, CRAY APP, CRAY C90, CRAY C90D, Cray C++ Compiling System, CRAY EL, Cray NQS, CRAY S-MP, CRAY T3D, CRAY X-MP, CRAY XMS, CRAY-2, Cray/REELibrarian, CRInform, CRI/TurboKiva, CSIM, CVT, Delivering the power . . ., DGauss, Docview, EMDS, IOS, ND Series Network Disk Array, Network Queuing Environment, Network Queuing Tools, OLNETH, RQS, SEGLDR, SMARTE, SUPERCLUSTER, SUPERLINK, Trusted UNICOS, and UNICOS MAX are trademarks of Cray Research, Inc.

Requests for copies of Cray Research, Inc. publications should be directed to:

CRAY RESEARCH, INC.
Logistics
6251 South Prairie View Road
Chippewa Falls, WI 54729

Comments about this publication should be directed to:

CRAY RESEARCH, INC.
Hardware Publications and Training
890 Industrial Blvd.
Chippewa Falls, WI 54729

CONTENTS – VOLUMES 1 AND 2

VOLUME 1

1 ARCHITECTURE OVERVIEW

2 INTERCONNECT NETWORK

3 PROCESSING ELEMENT NODE

4 I/O GATEWAY

5 STAND-ALONE DIAGNOSTIC SYSTEM INITIALIZATION

6 ADDRESSING

7 OPERATIONS

VOLUME 2

8 BLOCK TRANSFER ENGINE

9 BARRIER SYNCHRONIZATION

10 CONTROL AND STATUS

11 LOSP AND HISP CHANNELS

12 REFERENCE MATERIAL

13 CONTROL PANEL SWITCHES

GLOSSARY

INDEX

CONTENTS

8 BLOCK TRANSFER ENGINE

Functional Description	8-1
Register Groups	8-2
Remote Addressing	8-2
Local Addressing	8-2
Control and Status	8-2
BLT Register Functions	8-3
Remote Address Generation	8-3
Local Address Generation	8-14
Control and Status Information	8-16
Register Mapping	8-17
Remote Index Register	8-18
Remote Stride Register	8-21
Vector Length Register	8-22
Index Vector Register	8-23
Remote Mask Register	8-24
Remote Base Register	8-26
Remote Limit Register	8-27
Local Address Register	8-28
Local Stride Register	8-29
Control Register	8-30
Status Register	8-36
Hardware Characteristics	8-46
BLT Complete Interrupt	8-46
BLT Memory-mapped Register Write Ordering	8-47
Hardware Description	8-48
Initializing the BLT or EE Option	8-49

8 BLOCK TRANSFER ENGINE (continued)

Transferring Blocks of Data	8-57
Remote Read	8-57
Remote Write	8-69

9 BARRIER SYNCHRONIZATION

Functional Description	9-1
Logical Barrier Synchronization Circuits	9-2
Barrier Synchronization	9-2
Eureka Synchronization	9-6
Logical Partitions	9-7
Physical Barrier Synchronization Circuits	9-9
Physical Partitions	9-9
Timing	9-16
Register Mapping	9-16
Barrier Synchronization Register 0	9-18
Barrier Synchronization Register 1	9-19
Barrier Synchronization Function Register	9-20
Barrier Synchronization Mask and Interrupt Register	9-21
Barrier Timing and Refresh Register	9-23
Network Mode Register	9-25
Hardware Description	9-34
Physical Barrier Synchronization Circuitry	9-34
AR Option	9-34
ED Option	9-36
BC Option	9-37
Fanin Example of Barrier Bits 0 through 3 from PCB Bypass Point 000 ₁₆	9-49

10 CONTROL AND STATUS

Functional Description	10-1
Control Signals	10-1
Hardware Interrupt Pins	10-2
Instruction Cache Test Mode Pins	10-6
Cycle Request Pins	10-7
Cycle Acknowledge Pins	10-8
Cache Line Invalidate Pins	10-8
Status Information	10-9
Register Mapping	10-9
System Control Register	10-10
System Status Register	10-13
User Status Register	10-18
Invalidate Control Register	10-19
DRAM Control Register	10-21
Barrier Timing and Refresh Register	10-29
Network Enable Register	10-31
Network Performance Register	10-32
Network Mode Register	10-34
No-operation Register	10-37
Hardware Description	10-38
AR Option	10-38
System Control Register	10-39
Barrier Timing and Refresh Register	10-40
Hardware Interrupts	10-42
AM Option	10-43
System Status Register	10-43
User Status Register	10-46
Invalidate Control Register	10-46
DRAM Control Register	10-47
EE Option	10-52
Network Enable Register	10-52
Network Performance Register	10-53
Network Mode Register	10-54

11 LOSP AND HISP CHANNELS

LOSP Channel	11-2
Input Channel	11-3
Initialization	11-4
Interrupt Driven Versus Polled	11-4
Receiving Data	11-5
Errors	11-8
Clearing the Channel	11-10
Master Clear	11-10
Output Channel	11-12
Initialization	11-13
Interrupt Driven Versus Polled	11-13
Sending Data	11-13
Errors	11-16
Clearing the Channel	11-19
HISP Channel	11-20
Master Input Channel	11-22
Requesting a Transfer	11-23
Initialization	11-23
Interrupt Driven Versus Polling	11-24
BLT Parameters	11-24
I/O Registers	11-26
Sending the Address and Block Length	11-27
Receiving Data	11-27
Errors	11-29
Disable SECDED	11-33
Clearing the Channel	11-33
Slave Input Channel	11-34
Initialization	11-34
Interrupt Driven Versus Polled	11-34
Receiving the Address and Block Length	11-34
I/O Registers and BLT Parameters	11-36
Receiving Data	11-36
Errors	11-37
Disable SECDED	11-39

11 LOSP AND HISP CHANNELS (continued)

Clearing the Channel	11-39
Master Output Channel	11-40
Requesting a Transfer	11-41
Initialization	11-41
Interrupt Driven Versus Polled	11-42
BLT Parameters	11-42
I/O registers	11-44
Sending the Address and Block Length	11-44
Sending Data	11-45
Errors	11-47
Disable SECDDED	11-49
Clearing the Channel	11-49
Slave Output Transfers	11-50
Initialization	11-50
Interrupt Driven Versus Polled	11-50
Receiving the Address and Block Length	11-51
I/O Registers and BLT Parameters	11-52
Sending Data	11-52
Errors	11-53
Disable SECDDED	11-54
Clearing the Channel	11-54
Register Mapping	11-55
Interrupt Flags Register	11-56
HISP Error (Bit 0)	11-57
Microprocessor Single-bit Error (Bit 1)	11-58
Microprocessor Multiple-bit Error (Bit 2)	11-58
HISP Single-bit Error (Bit 3)	11-59
HISP Multiple-bit Error (Bit 4)	11-59
HISP Address Transfer Complete (Bit 5)	11-59
HISP Data Transfer Complete (Bit 6)	11-60
LOSP Input Data Available (Bit 7)	11-60
LOSP Input Disconnect (Bit 8)	11-60
LOSP Input Error (Bit 9)	11-61
LOSP Output Complete (Bit 10)	11-61
LOSP Output Error (Bit 11)	11-61

11 LOSP AND HISP CHANNELS (continued)

Master Clear In (Bit 12)	11-62
Revision Level (Bits 13 and 14)	11-62
Master or Slave I/O Node (Bit 15)	11-62
Interrupt Mask Register	11-63
I/O Channel Error Register	11-64
Master I/O Gateway	11-64
Slave I/O Gateway	11-69
Clear I/O Register	11-71
Clear LOSP Input Control (Bit 0)	11-72
Clear LOSP Output Control (Bit 1)	11-72
Clear HISP Address Control (Bit 2)	11-73
Clear HISP Data Control (Bit 3)	11-74
Clear HISP ROA Address (Bit 4)	11-74
Clear LOSP In Ready Waiting Bit (Bit 5)	11-75
Clear LOSP Buffer and Priority (Bit 6)	11-75
Clear LOSP Input Error Flags (Bit 7)	11-75
Clear LOSP Output Error Flags (Bit 8)	11-76
Clear LOSP Input Disconnect (Bit 9)	11-76
Clear LOSP Input Data Available (Bit 10)	11-76
LOSP Channel Configuration Register	11-77
Input Node	11-77
Output Node	11-78
LOSP Channel Data In Register	11-79
LOSP Channel Data Out Register	11-80
LOSP Output Channel Disconnect Register	11-80
HISP Channel Address Configuration Register	11-81
Master I/O Gateway	11-81
Slave I/O Gateway	11-85
HISP Data Configuration Register	11-87
Input Node	11-87
Output Node	11-88
HISP Channel Lower Block Length Register	11-90
HISP Channel Upper Block Length Register	11-91

11 LOSP AND HISP CHANNELS (continued)

HISP Channel Address 0 Register	11-92
100 Mbytes/s Mode	11-93
200 Mbytes/s Mode	11-94
HISP Channel Address 1 Register	11-95
HISP Buffer Registers	11-96
HISP Channel ROA Address Register	11-97
Syndrome Register	11-98
HISP Syndrome	11-99
Microprocessor Syndrome	11-99
Clear Microprocessor Syndrome Register	11-100
Clear HISP Syndrome Register	11-101
Heartbeat Generation Register	11-102

12 REFERENCE MATERIAL

13 CONTROL PANEL SWITCHES

Multiple-cabinet Control Panel Switches (6002 – 6015)	13-1
LOSP ENABLE Switches	13-3
CLOCK SELECT Switches	13-4
SYSTEM SIZE Switches	13-5
CHASSIS SELECT Switches	13-6
HEARTBEAT INTERVAL Switches	13-7
HEARTBEAT ENABLE Switch	13-8
HEARTBEAT SELECT Switch	13-8
REMOTE ENABLE Switch	13-9
Multiple-cabinet Control Panel Switches (6001, 6016 – 6099)	13-9
LOSP ENABLE Switches	13-11
CLOCK SELECT Switches	13-12
SYSTEM SIZE Switches	13-13
CHASSIS SELECT Switches	13-14

13 CONTROL PANEL SWITCHES (continued)

HEARTBEAT INTERVAL Switches	13-15
HEARTBEAT ENABLE Switch	13-16
HEARTBEAT SELECT Switch	13-16
REMOTE ENABLE Switch	13-17
Single-cabinet Control Panel Switches	13-17
LOSP ENABLE Switches	13-19
CLOCK SELECT Switches	13-20
HEARTBEAT INTERVAL Switches	13-21
HEARTBEAT ENABLE Switch	13-22
HEARTBEAT SELECT Switch	13-22
CPU IDLE Switches	13-23
CLOCK SELECT Switches	13-24
SECTION SELECT Switches	13-25
MASTER CPU Switches	13-26
MAINTENANCE MODE Switch	13-27
MCU SELECT Switches	13-27
CONTROL DUMP Switch	13-28
REMOTE ENABLE Switch	13-28
RTI ENABLE Switch	13-29

GLOSSARY

INDEX

FIGURES

Figure 8-1.	Remote Address Generation	8-3
Figure 8-2.	Remote Index Generation	8-4
Figure 8-3.	Centrifuge Functional Blocks	8-6
Figure 8-4.	Separation of Remote Index into PE Number and Index Offset	8-10

FIGURES (continued)

Figure 8-5.	Address Offset Generation	8-11
Figure 8-6.	Virtual PE Number Range Check	8-12
Figure 8-7.	PE Node Mask for a 32-node Partition	8-13
Figure 8-8.	Local Address Generation	8-14
Figure 8-9.	Remote Index Register Address Bit Format	8-18
Figure 8-10.	BLT_RSR Address Bit Assignments	8-21
Figure 8-11.	BLT_VLR Address Bit Assignments	8-22
Figure 8-12.	BLT_IVR Address Bit Assignment	8-23
Figure 8-13.	BLT_RMR Address Bit Assignments	8-24
Figure 8-14.	BLT_RMR Bit Format	8-25
Figure 8-15.	BLT_RBR Address Bit Assignments	8-26
Figure 8-16.	BLT_RLR Address Bit Assignments	8-27
Figure 8-17.	BLT_LAR Address Bit Assignments	8-28
Figure 8-18.	BLT_LSR Address Bit Assignments	8-29
Figure 8-19.	BLT_CR Address Bit Assignments	8-30
Figure 8-20.	Sample BLT_RMR Value	8-33
Figure 8-21.	BLT_SR Address Bit Assignments	8-36
Figure 8-22.	BLT Software Queue	8-43
Figure 8-23.	Same PE Request	8-45
Figure 8-24.	Other PE Request	8-45
Figure 8-25.	BLT Complete Interrupt	8-46
Figure 8-26.	EE Option	8-48
Figure 8-27.	Control and Address for Loading the BLT Memory-mapped Registers	8-51
Figure 8-28.	AM Option Circuitry that Determines BLT Memory-map Reference	8-52
Figure 8-29.	BLT Memory-map Reference Control Generated by the EB Option	8-55
Figure 8-30.	BLT Generation of Address and Control	8-58
Figure 8-31.	EA Option, EB Option, and Routing Tag Look-up Table	8-59

FIGURES (continued)

Figure 8-32.	Request Packet Assembly	8-60
Figure 8-33.	Input of Request Packet	8-62
Figure 8-34.	Generation of DRAM Control and Response Packet Header Phits	8-64
Figure 8-35.	Reading Data from the DRAM	8-66
Figure 8-36.	Response Packet Assembly	8-67
Figure 8-37.	Control to Read the Remote Write Data from Local Memory	8-70
Figure 9-1.	Barrier Register 0 and Barrier Register 1	9-2
Figure 9-2.	Simplified Barrier Synchronization Circuit	9-3
Figure 9-3.	BSFR Format	9-6
Figure 9-4.	BSMI Register Format	9-7
Figure 9-5.	Bypass Points	9-10
Figure 9-6.	Partitioned Barrier Synchronization Circuit	9-11
Figure 9-7.	PE Node Bypass Point Partitions in a 512-PE CRAY T3D System	9-14
Figure 9-8.	Barrier Synchronization Register 0 Address Bit Format	9-18
Figure 9-9.	Barrier Synchronization Register 1 Address Bit Format	9-19
Figure 9-10.	Barrier Synchronization Function Register Address Bit Format	9-20
Figure 9-11.	Barrier Synchronization Mask and Interrupt Register Address Bit Format	9-21
Figure 9-12.	Barrier Timing and Refresh Register Address Bit Format	9-23
Figure 9-13.	Network Mode Register Address Bit Format	9-25
Figure 9-14.	Four-PE Barrier Partition in Physical Barrier Synchronization Circuit 0	9-28
Figure 9-15.	Spare PE Node Barrier Partition in Physical Barrier Synchronization Circuit 0	9-29
Figure 9-16.	I/O Gateway Barrier Partition in Physical Barrier Synchronization Circuit 0	9-31

FIGURES (continued)

Figure 9-17.	Bypass 2A	9-33
Figure 9-18.	Barrier Synchronization Circuitry within the AR Option	9-35
Figure 9-19.	Barrier Synchronization Circuitry within the ED Options	9-36
Figure 9-20.	Fanin and Fanout of Barrier Bits within the BC Option	9-37
Figure 9-21.	Fanin and Fanout of Barrier Bits n+0 and n+1 in a Two Chassis System	9-39
Figure 9-22.	Fanin and Fanout of Barrier Bits n+2 and n+3 in a Two Chassis System	9-40
Figure 9-23.	Fanin of Barrier Bits n+0 or n+1 in a Four Chassis System	9-41
Figure 9-24.	Fanout of Barrier Bits n+0 or n+1 in a Four Chassis System	9-43
Figure 9-25.	Fanin of Barrier Bits n+2 or n+3 in a Four Chassis System	9-45
Figure 9-26.	Fanout of Barrier Bits n+2 or n+3 in a Four Chassis System	9-47
Figure 9-27.	Fanin of Barrier Bits n+0 through n+3 from PCB Bypass Point 000 ₁₆ (page 1 of 2)	9-51
Figure 9-28.	Fanin of Barrier Bits n+0 through n+3 from PCB Bypass Point 000 ₁₆ (page 2 of 2)	9-53
Figure 10-1.	Microprocessor and Support Circuitry Control Signals	10-1
Figure 10-2.	System Control Register Address Bit Format	10-10
Figure 10-3.	System Status Register Address Bit Format	10-13
Figure 10-4.	User Control and Status Register Address Bit Format	10-18
Figure 10-5.	Invalidate Control Register Address Bit Format ...	10-19
Figure 10-6.	DRAM Control Register Address Bit Format	10-21
Figure 10-7.	Transfer 0	10-23
Figure 10-8.	Transfer 1	10-23
Figure 10-9.	Transfer 2	10-23

FIGURES (continued)

Figure 10-10. Transfers 3 through 6	10-24
Figure 10-11. Transfer 7	10-27
Figure 10-12. Barrier Timing and Refresh Register Address Bit Format	10-29
Figure 10-13. Network Enable Register Address Bit Format	10-31
Figure 10-14. Network Performance Register Address Bit Format	10-32
Figure 10-15. Network Mode Register Address Bit Format	10-34
Figure 10-16. Invalidate Control Register Address Bit Format	10-37
Figure 10-17. AR Option	10-38
Figure 10-18. Barrier Timing Circuitry	10-41
Figure 10-19. AM Option	10-43
Figure 10-20. EE Option	10-52
Figure 10-21. Performance Monitor Circuitry in the SR Option	10-53
Figure 10-22. Performance Monitor Circuitry in the ED Option	10-54
Figure 11-1. I/O Gateway	11-1
Figure 11-2. LOSP Channel Signals	11-2
Figure 11-3. Receiving LOSP Input Data	11-5
Figure 11-4. LOSP Input Errors	11-8
Figure 11-5. Master Clear	11-11
Figure 11-6. Sending LOSP Output Data	11-14
Figure 11-7. LOSP Output Errors	11-16
Figure 11-8. HISP Channel Signals	11-20
Figure 11-9. Master HISP Input Transfer	11-26
Figure 11-10. Master HISP Input Errors	11-30
Figure 11-11. Slave HISP Input Transfer	11-35
Figure 11-12. Slave HISP Input Errors	11-37
Figure 11-13. Master HISP Output Channel	11-44
Figure 11-14. Master HISP Output Errors	11-47

FIGURES (continued)

Figure 11-15. Slave HISP Output Transfer	11-51
Figure 11-16. Slave HISP Output Errors	11-53
Figure 11-17. IO_IFLAG Register Address Bit Format	11-56
Figure 11-18. Interrupt Mask Register Address Bit Format	11-63
Figure 11-19. IO_ERR Register Address Bit Format	11-64
Figure 11-20. Clear IO Register Address Bit Format	11-71
Figure 11-21. LOSP Channel Configuration Register Address Bit Format	11-77
Figure 11-22. LOSP_DI Register Address Bit Format	11-79
Figure 11-23. LOSP_DO Register Address Bit Format	11-80
Figure 11-24. LOSP_DISC Register Address Bit Format	11-80
Figure 11-25. HISP Channel Address Configuration Register Address Bit Format	11-81
Figure 11-26. HISP Data Configuration Register Address Bit Format	11-87
Figure 11-27. HISP Lower Block Length Register Address Bit Format	11-90
Figure 11-28. HISP Upper Block Length Register Bit Format ...	11-91
Figure 11-29. HISP Channel Address 0 Register Address Bit Format	11-92
Figure 11-30. HISP Channel Address 1 Register Address Bit Format	11-95
Figure 11-31. HISP Buffer Register Address Bit Format	11-96
Figure 11-32. BLT and Microprocessor HISP_BUFFER Addresses	11-96
Figure 11-33. HISP Channel ROA Register Address Bit Format	11-97
Figure 11-34. Syndrome Register Address Bit Format	11-98
Figure 11-35. Clear Microprocessor Syndrome Register Address Bit Format	11-100
Figure 11-36. Clear HISP Syndrome Register Address Bit Format	11-101

FIGURES (continued)

Figure 11-37.	Heartbeat Generation Register Address Bit Format	11-102
Figure 12-1.	Processing Element Node Memory-mapped Registers	12-7
Figure 12-2.	I/O Memory-mapped Registers	12-9
Figure 12-3.	PE Module	12-11
Figure 12-4.	Master Input/Output Module	12-13
Figure 12-5.	Slave Input/Output Module	12-15
Figure 12-6.	Processing Element Module with Daughter Cards	12-17
Figure 12-7.	Processing Element Module	12-19
Figure 12-8.	I/O Module	12-21
Figure 12-9.	System Printed Circuit Board	12-23
Figure 13-1.	Multiple-cabinet Control Panel Switches (p2)	13-2
Figure 13-2.	Multiple-cabinet (p2) LOSP ENABLE Switches ..	13-3
Figure 13-3.	Multiple-cabinet (p2) CLOCK SELECT Switches	13-4
Figure 13-4.	Multiple-cabinet (p2) SYSTEM SIZE Switches ...	13-5
Figure 13-5.	Multiple-cabinet (p2) CHASSIS SELECT Switches	13-6
Figure 13-6.	Multiple-cabinet (p2) HEARTBEAT INTERVAL Values	13-7
Figure 13-7.	Multiple-cabinet (p2) HEARTBEAT ENABLE Switch	13-8
Figure 13-8.	Multiple-cabinet (p2) HEARTBEAT SELECT Switch	13-8
Figure 13-9.	Multiple-cabinet (p2) REMOTE ENABLE Switch	13-9
Figure 13-10.	Multiple-cabinet Control Panel Switches (np2)	13-10
Figure 13-11.	Multiple-cabinet (np2) LOSP ENABLE Switches ..	13-11
Figure 13-12.	Multiple-cabinet (np2) CLOCK SELECT Switches	13-12
Figure 13-13.	Multiple-cabinet (np2) SYSTEM SIZE Switches ..	13-13

FIGURES (continued)

Figure 13-14. Multiple-cabinet (np2) CHASSIS SELECT Switches	13-14
Figure 13-15. Multiple-cabinet (np2) HEARTBEAT INTERVAL Values	13-15
Figure 13-16. Multiple-cabinet (np2) HEARTBEAT ENABLE Switch	13-16
Figure 13-17. Multiple-cabinet (np2) HEARTBEAT SELECT Switch	13-16
Figure 13-18. Multiple-cabinet (np2) REMOTE ENABLE Switch	13-17
Figure 13-19. Single-cabinet Control Panel Switches	13-18
Figure 13-20. Single-cabinet LOSP ENABLE Switches	13-19
Figure 13-21. Single-cabinet CLOCK SELECT Switches	13-20
Figure 13-22. Single-cabinet HEARTBEAT INTERVAL Values	13-21
Figure 13-23. Single-cabinet HEARTBEAT ENABLE Switch ...	13-22
Figure 13-24. Single-cabinet HEARTBEAT SELECT Switch ...	13-22
Figure 13-25. CRAY Y-MP CPU IDLE Switches	13-23
Figure 13-26. CRAY Y-MP CLOCK SELECT Switches	13-24
Figure 13-27. CRAY Y-MP SECTION SELECT Switches	13-25
Figure 13-28. CRAY Y-MP MASTER CPU Switches	13-26
Figure 13-29. CRAY Y-MP MAINTENANCE MODE Switch ...	13-27
Figure 13-30. CRAY Y-MP MCU SELECT Switches	13-27
Figure 13-31. CRAY Y-MP CONTROL DUMP Switch	13-28
Figure 13-32. Single-cabinet REMOTE ENABLE Switch	13-28
Figure 13-33. CRAY Y-MP RTI ENABLE Switch	13-29

TABLES

Table 8-1. Remote Index Format Examples	8-8
Table 8-2. BLT Registers	8-17
Table 8-3. BLT_RIR	8-19
Table 8-4. BLT_RSR Bit Format	8-21

TABLES (continued)

Table 8-5.	BLT_VLR Bit Format	8-22
Table 8-6.	BLT_IVR Bit Format	8-24
Table 8-7.	BLT_RBR Bit Format	8-26
Table 8-8.	BLT_RLR Bit Format	8-27
Table 8-9.	BLT_LAR Bit Format	8-28
Table 8-10.	BLT_LSR Bit Format	8-29
Table 8-11.	BLT_CR Bit Format	8-30
Table 8-12.	Transfer Types	8-31
Table 8-13.	PE Range Mask Values	8-35
Table 8-14.	BLT_SR Bit Format	8-36
Table 8-15.	Status Select Codes	8-38
Table 8-16.	Interrupt Status	8-38
Table 8-17.	BLT Error Status	8-39
Table 8-18.	BLT_VLR Bits 0 through 3 Status	8-41
Table 8-19.	BLT_VLR Bits 4 through 7 Status	8-41
Table 8-20.	BLT_VLR Bits 8 through 11 Status	8-42
Table 8-21.	BLT_VLR Bits 12 through 15 Status	8-42
Table 8-22.	Address Offset Bits 5 through 28	8-53
Table 8-23.	Decode of Address Offset	8-53
Table 8-24.	Decode of Register Select Bits	8-55
Table 9-1.	Effect of BSMI register and BSFR Bits	9-8
Table 9-2.	Physical Barrier Synchronization Circuit Inputs	9-9
Table 9-3.	Levels in a 512-PE Physical Barrier Synchronization Circuit	9-13
Table 9-4.	Barrier Synchronization Registers	9-17
Table 9-5.	BSR0 Format	9-18
Table 9-6.	BSR1 Format	9-19
Table 9-7.	BSFR Format	9-20
Table 9-8.	BSMI Register Write Format	9-21
Table 9-9.	BSMI Register Read Format	9-22
Table 9-10.	BAR_TMG Bit Format	9-23

TABLES (continued)

Table 9-11.	Timing Procedure Values	9-25
Table 9-12.	NODE_CSR Bit Format	9-26
Table 9-13.	Bit 11 of NODE_CSR in a Processing Element Node	9-27
Table 9-14.	Bit 11 of NODE_CSR in a Spare Processing Element Node	9-29
Table 9-15.	Bit 11 of NODE_CSR in an I/O Gateway	9-30
Table 9-16.	Bit 12 of NODE_CSR in a PE Node or Spare PE Node	9-32
Table 9-17.	Bit 12 of NODE_CSR in an I/O Gateway	9-32
Table 10-1.	Hardware Interrupt Signals	10-2
Table 10-2.	Instruction Cache Test Mode Pins	10-6
Table 10-3.	Cycle Request Codes	10-7
Table 10-4.	Cycle Acknowledge Codes	10-8
Table 10-5.	Control and Status Registers	10-9
Table 10-6.	SCR Format	10-10
Table 10-7.	SSR Format	10-13
Table 10-8.	USR Format	10-18
Table 10-9.	INV_CR Format	10-20
Table 10-10.	Enable Invalidation and Invalidation Filter Bits	10-21
Table 10-11.	DRAM_CR Format	10-22
Table 10-12.	Timing Information Transfers	10-25
Table 10-13.	CAS Cycle Timing Bit Values	10-27
Table 10-14.	DRAM Readout Sample Timing Values	10-28
Table 10-15.	BAR_TMG Bit Format	10-29
Table 10-16.	DRAM Refresh Timing Values	10-30
Table 10-17.	NET_ENA Bit Format	10-31
Table 10-18.	NET_PFM Bit Format	10-32
Table 10-19.	Network Switch Performance Parameters	10-33
Table 10-20.	Switch or BLT Select Bits	10-34
Table 10-21.	NODE_CSR Bit Format	10-35

TABLES (continued)

Table 10-22. Configuration Data	10-50
Table 11-1. LOSP Input Registers	11-3
Table 11-2. LOSP Output Registers	11-12
Table 11-3. HISP Channel Signals	11-21
Table 11-4. HISP Input Registers	11-22
Table 11-5. Sample BLT Input Parameters	11-25
Table 11-6. HISP Output Registers	11-40
Table 11-7. Sample BLT Output Parameters	11-43
Table 11-8. Register Assignments for I/O Node	11-55
Table 11-9. IO_IFLAG Register Bit Format	11-56
Table 11-10. IO_IMASK Register Bit Format	11-63
Table 11-11. Master IO_ERR Register Bit Format	11-64
Table 11-12. Slave IO_ERR Register Bit Format	11-69
Table 11-13. IO_CLR Register Bit Format	11-71
Table 11-14. LOSP_CFIG Register Format	11-77
Table 11-15. LOSP_CFIG Register Bit Format in Output Node ..	11-78
Table 11-16. LOSP_DI Register Bit Format	11-79
Table 11-17. HISP_A_CFIG Register Bit Format	11-82
Table 11-18. Slave HISP_A_CFIG Register Bit Format	11-85
Table 11-19. HISP_D_CFIG Register Format in the Input Node	11-87
Table 11-20. HISP_D_CFIG Register Bit Format in the Output Node	11-88
Table 11-21. HISP_L_BL Register Bit Format	11-90
Table 11-22. HISP_U_BL Register Bit Format	11-91
Table 11-23. HISP_ADDR0 Register Bit Format in 100 Mbytes/s Mode	11-93
Table 11-24. HISP_ADDR0 Register Bit Format in 200 Mbytes/s Mode	11-94
Table 11-25. HISP_ADDR1 Register Bit Format	11-95
Table 11-26. HISP_ROA Register Bit Format	11-97
Table 11-27. IO_SYN Register Bit Format	11-98

TABLES (continued)

Table 11-28. HB_GEN Register Bit Format	11-102
Table 12-1. Page Reference for the Memory-mapped Registers	12-1
Table 12-2. A-series Options	12-3
Table 12-3. E-series Options	12-4
Table 12-4. I-series Options	12-5
Table 13-1. Multiple-cabinet System Control Panel Locations (p2)	13-1
Table 13-2. Multiple-cabinet Switch Label and P_WHOAMI Values	13-3
Table 13-3. Multiple-cabinet System Control Panel Locations (np2)	13-9

8 BLOCK TRANSFER ENGINE

The block transfer engine (BLT) is an asynchronous direct memory access device. The BLT redistributes system data between remote memory and local memory in either of the PEs in a processing element node. The BLT can create up to 65,536 packets with a 1-word or 4-word body without interruption from the PE.

This section describes the function and hardware of the BLT.

Functional Description

The BLT initiates four types of data transfer operations: constant stride read, constant stride write, gather, and scatter. A constant stride read operation transfers data from fixed increment address locations in remote memory to fixed increment address locations in local memory. A constant stride write operation transfers data from fixed increment address locations in local memory to fixed increment address locations in remote memory. A gather operation transfers data from nonsequential memory locations in remote memory to fixed increment address locations in local memory. A scatter operation transfers data from fixed increment address locations in local memory to nonsequential memory locations in remote memory.

The BLT receives initial control information from one of the PEs in a processing element node and then functions independently from the PEs. The PEs send control information to the BLT by writing information into BLT memory-mapped registers.

The following subsections describe the format and use of the memory-mapped registers (hereafter referred to as registers) that control the BLT. This information describes how the registers are grouped, the overall function of the registers, the register addressing, and the register bit formats.

Register Groups

The BLT registers are divided into three groups of registers: remote addressing registers, local addressing registers, and control and status registers. The following subsections briefly describe each group of registers.

Remote Addressing

The BLT uses the remote addressing registers when generating remote memory addresses. The remote memory addresses are used to generate information used in the header of request packets. The remote addressing registers include the following registers:

- Remote index register (BLT_RIR)
- Remote stride register (BLT_RSR)
- Index vector register (BLT_IVR)
- Vector length register (BLT_VLR)
- Remote mask register (BLT_RMR)
- Remote base register (BLT_RBR)
- Remote limit register (BLT_RLR)

Local Addressing

The BLT uses the local addressing registers when generating local memory addresses. The local memory addresses may also be used to generate information used in the header of request packets. The local addressing registers include the following registers:

- Local address register (BLT_LAR)
- Local stride register (BLT_LSR)

Control and Status

The PEs use the control and status registers to control the BLT and check the progress of a block transfer. The control and status registers include the following registers:


- Control register (BLT_CR)
- Status register (BLT_SR)

BLT Register Functions

The BLT uses the registers to perform three functions: remote address generation, local address generation, and status generation.

Remote Address Generation

Remote addressing circuitry in the BLT generates information that is placed in request packet headers and is used to reference remote memory. Figure 8-1 shows the functional blocks of remote address generation.



A-11663

Figure 8-1. Remote Address Generation

Remote Index Generation

The first step performed by the BLT during remote address generation is the generation of the remote index. The remote index is a memory address that contains a PE number and an index offset in a software defined bit format. The remote index is incremented during a constant stride read or write operation and held constant during a scatter or gather operation. Figure 8-2 shows the functional blocks of the remote index generation.




Figure 8-2. Remote Index Generation

Before a constant stride read or write operation, the PE that is using the BLT loads information into the BLT remote addressing registers. These registers include the remote index register, the remote stride register, and the vector length register.

Initially, the remote index register is loaded with the first remote index that will be converted into a remote address. The remote stride is loaded with a value that the BLT repeatedly adds to the remote index to generate successive remote indexes. The vector length register is loaded with a value equal to the total number of request packets the PE is requesting the BLT to generate. If the vector length register is set to 0, the BLT generates the maximum number of packets (65,536).

During the constant stride operation, the BLT reads the first remote index from the remote index register and sends the remote index to the centrifuge. When reading the first remote index, the remote stride register is disabled and its value is 0. As the BLT sends the first remote index to the centrifuge, it decrements the vector length register.

The BLT then adds the remote stride to the first remote index. This value is the second remote index value. As the BLT sends the second remote index value to the centrifuge, it decrements the vector length register.

The BLT then adds the remote stride to the second remote index. This sum is the third remote index value. As the BLT sends the third remote index value to the centrifuge, it decrements the vector length register. The BLT continues this process until the value of the vector length register is 0.

During a scatter or gather operation, the remote index value is routed through the remote stride register and the feedback to the remote index register is disabled. Before loading the BLT registers with information, the PE loads contiguous 64-bit local memory locations with remote indexes. Each memory location contains one remote index.

Before starting the scatter or gather operation, the PE loads the index vector register with a 24-bit address offset that points to the first remote index stored in local memory. The 24-bit address offset is equivalent to bits 3 through 26 of a byte-oriented partial physical address generated by a microprocessor. During the scatter or gather operation, the BLT reads the value of the index vector register and retrieves four remote indexes from local memory. (These four remote indexes are addressed by bits 2 through 23 of the index vector register.) After reading the contents of the index vector register, the BLT automatically increments the value of the index vector register by four.

Each remote index read from local memory is added to the remote index register. Because of this characteristic, the value of the remote index register can be zero or can be used as a base value for remote indexes. As the sum of the remote index and the value of the remote index register is sent to the centrifuge, the BLT decrements the value of the vector length register.

The BLT then reads the value of the index vector register again and retrieves the second set of four remote indexes from local memory. After reading the contents of the index vector register, the BLT automatically increments the value of the index vector register by four.

Each remote index in the second set of 4 remote indexes is also added to the value stored in the remote index register. As the sum of the remote index and the value of the remote index register is sent to the centrifuge, the BLT decrements the value of the vector length register. The BLT continues to read the remote indexes from local memory until the value of the vector length register is zero.

Because the BLT reads remote indexes from local memory four remote indexes at a time (cache-line size), the first remote index stored in memory must be cache-line aligned (bits 0 through 4 of the address set to 0). The last remote index stored in memory does not have to be cache-line aligned.

Centrifuge

After receiving the remote index, the centrifuge separates the remote index into a PE number and an index offset. Figure 8-3 shows the functional blocks of the centrifuge.

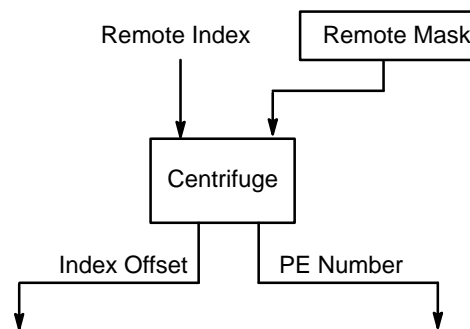


Figure 8-3. Centrifuge Functional Blocks

The PE number is either a logical PE number or a virtual PE number. When the BLT transfers information for the operating system, the BLT interprets the PE number as a logical PE number. When the BLT transfers information for a user, the BLT interprets the PE number as a virtual PE number.

Before a transfer begins, the PE sets a bit in the BLT control register that indicates whether the PE number should be interpreted as a virtual or logical PE number. When the BLT interprets the PE number as a virtual PE number, the BLT sends the virtual PE number through the virtual-to-logical translation circuitry before the PE number is used in a request packet. When the BLT interprets the PE number as a logical PE number, no translation of the PE number takes place.

The index offset is a word-oriented address offset in remote memory with respect to a base address offset. The BLT adds the index offset to a remote base offset before using the information in a request packet header.

Software defines the bit positions of the PE number and index offset in the remote index. This process gives a programmer the ability to redistribute system data in almost any type of simple or complex pattern.

As a simple example, the following paragraphs describe a remote index that contains 5 bits; 2 of the bits are assigned to the PE number, and 3 of the bits are assigned to the index offset. This example demonstrates the generation of remote indexes during a constant stride read or write operation. In this operation, the remote stride is set to 1. Also, the first remote index is set to 0.

Table 8-1 shows three examples of how the bits of the remote index may be defined. As the remote index increments sequentially from 0 to 31, each of the formats increments through PE numbers and index offsets differently.

In the first format, the 2 most significant bits of the remote index are assigned as the PE number. As the remote index increments from 0 to 7, the index offset increments from 0 to 7 in PE number 0. As the remote index increments from 8 to 17, the index offset increments from 0 to 7 in PE number 1. In this format, all of the index offsets increment from 0 to 7 before the PE number increments.

In the second format, the 2 least significant bits of the remote index are assigned as the PE number. As the remote index increments from 0 to 3, the PE number increments from 0 to 3 and the index offset is 0. As the remote index increments from 4 to 7, the PE number increments from 0 to 3 and the index offset is 1. In this format, all of the PE numbers increment from 0 to 3 before the index offset increments.

In the third format, bits 1 through 3 of the remote index are assigned as the PE number. As the remote index increments from 0 to 31, the PE number and index offset increment in a set pattern. In this format, all of the PE number and index offset combinations are incremented through but are not incremented through sequentially.

In the CRAY T3D system, the remote index is actually a 36-bit value; 12 of the bits are assigned to the PE number, and 24 of the bits are assigned to the index offset.

The centrifuge uses the remote mask register to determine which bits of the remote index are PE number bits and which bits of the remote index are index offset bits. Before the transfer begins, the PE loads the remote mask register with the appropriate values.

Table 8-1. Remote Index Format Examples

Remote Index Bits 4 through 0	Remote Index Bit Formats					
	Bits 4, 3 = PE Number Bits 2, 1, 0 = Index Offset		Bits 1, 0 = PE Number Bits 4, 3, 2 = Index Offset		Bits 3, 1 = PE Number Bits 4, 2, 0 = Index Offset	
	PE Number	Index Offset	PE Number	Index Offset	PE Number	Index Offset
00000	0	0	0	0	0	0
00001	0	1	1	0	0	1
00010	0	2	2	0	1	0
00011	0	3	3	0	1	1
00100	0	4	0	1	0	2
00101	0	5	1	1	0	3
00110	0	6	2	1	1	2
00111	0	7	3	1	1	3
01000	1	0	0	2	2	0
01001	1	1	1	2	2	1
01010	1	2	2	2	3	0
01011	1	3	3	2	3	1
01100	1	4	0	3	2	2
01101	1	5	1	3	2	3
01110	1	6	2	3	3	2
01111	1	7	3	3	3	3
10000	2	0	0	4	0	4
10001	2	1	1	4	0	5
10010	2	2	2	4	1	4
10011	2	3	3	4	1	5
10100	2	4	0	5	0	6


Table 8-1. Remote Index Format Examples (continued)

Remote Index Bits 4 through 0	Remote Index Bit Formats					
	Bits 4, 3 = PE Number Bits 2, 1, 0 = Index Offset		Bits 1, 0 = PE Number Bits 4, 3, 2 = Index Offset		Bits 3, 1 = PE Number Bits 4, 2, 0 = Index Offset	
	PE Number	Index Offset	PE Number	Index Offset	PE Number	Index Offset
10101	2	5	1	5	0	7
10110	2	6	2	5	1	6
10111	2	7	3	5	1	7
11000	3	0	0	6	2	4
11001	3	1	1	6	2	5
11010	3	2	2	6	3	4
11011	3	3	3	6	3	5
11100	3	4	0	7	2	6
11101	3	5	1	7	2	7
11110	3	6	2	7	3	6
11111	3	7	3	7	3	7

The remote mask is a 36-bit number. If a bit of the remote mask is set to 1, the corresponding bit of the remote index is assigned to the PE number. If a bit of the remote mask is set to 0, the corresponding bit of the remote index is assigned to the index offset.

Twelve bits of the remote mask register must be set to 1 to define the centrifuge output. Hardware in the BLT does not check the value of the remote mask register to ensure that 12 bits are set to 1.

Figure 8-4 illustrates the separation of the remote index into the PE number and index offset. The numbers shown in Figure 8-4 were arbitrarily chosen for this example.



A-11664

Figure 8-4. Separation of Remote Index into PE Number and Index Offset

Address Offset Generation

After the centrifuge separates the index offset, the BLT converts the index offset into a word-oriented address offset. Figure 8-5 shows the functional blocks of the address offset generation.

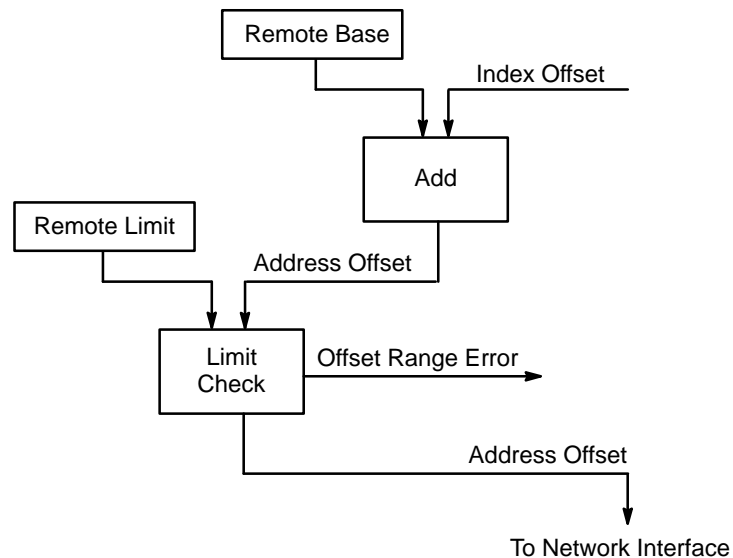


Figure 8-5. Address Offset Generation

To convert the index offset into an address offset, the BLT adds the contents of the remote base register to the index offset. The remote base register contains the starting word-oriented address offset of a remote data structure.

The BLT sends the address offset to the network interface for use in a request packet header. Before sending the address offset to the network interface, the BLT checks the value of the address offset.

The BLT checks the value of the address offset by comparing it to the value of the remote limit register. If the value of the address offset is greater than the value of the remote limit register, an offset range error occurs. More information on the offset range error is provided in “Register Mapping” in this section.

PE Range Check

After the centrifuge obtains the PE number, the BLT sends the PE number to the network interface. If the BLT interprets the PE number as a virtual PE number, the BLT checks the value of the virtual PE number.

Figure 8-6 shows the functional blocks of checking the virtual PE number.

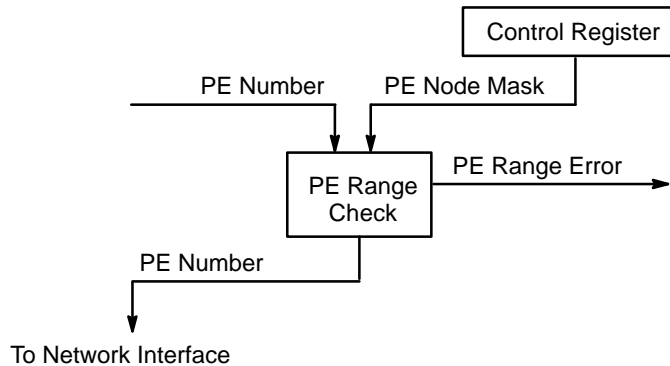


Figure 8-6. Virtual PE Number Range Check

To check the value of the virtual PE number, the BLT reads the PE node mask from the control register and compares this to the virtual node number. The virtual node number is the same as the virtual PE number minus the least significant bit of the virtual PE number. The PE node mask is an 11-bit number that contains a contiguous, right-justified field of 0's that indicate the size of a partition.

For example, Figure 8-7 shows the PE node mask for a 32-node partition. If the generated virtual node number is greater than 31, a PE range error occurs. (More information on the PE range error is provided in "Register Mapping" later in this section.) After checking the virtual PE number, the BLT sends the virtual PE number to the network interface where it is converted into a logical PE number and used in a request packet header.




Figure 8-7. PE Node Mask for a 32-node Partition

If the BLT interprets the PE number as a logical PE number, the BLT still compares the logical PE number to the PE node range mask in the control register. Because of this characteristic, the value of the PE node mask should be set to all 0's when operating the BLT in logical mode. The BLT then sends the logical PE number directly to the network interface for use in a request packet header. The BLT does not send the PE number through the virtual channel to logical conversion circuitry.

Outstanding Request Counter

Each time the BLT sends a request packet to a PE, the BLT increments a counter. This counter, the outstanding request counter, contains the number of request packets generated by the BLT that have not received responses.

The BLT contains two outstanding request counters, one for each PE in the processing element node. The counters operate independently and keep track of the number of request packets generated for each individual PE that has not received responses.

When the support circuitry in a processing element node receives a BLT response packet, the support circuitry signals the BLT that a response packet has arrived. After examining the response packet information, the BLT decrements the value of the appropriate outstanding request counter. When the value of the outstanding request counter reaches 0, a response packet has been received for every request packet the BLT generated. When this occurs, the block transfer operation for that PE is complete.

When enabled, the BLT sends a hardware interrupt to the PE to signal the PE that the transfer is complete. More information on the BLT hardware interrupt is provided in "Register Mapping" in this section and in Section 10, "Control and Status."

The outstanding request counter can be run in two modes. In the first mode, the BLT stalls when the request counter reaches 256 and does not continue until the outstanding request counter decrements. In the second mode, the BLT stalls when the outstanding request counter reaches 32 and does not continue until the counter decrements.

The outstanding request counter can be reset only by writing a code into the status register. The reset counter operation resets only the outstanding request counter associated with the PE that initiated the reset code. More information on resetting the counter is provided in “Register Mapping” later in this section.

Local Address Generation

Local addressing circuitry in the BLT may also generate information used in the header of request packets (packet type 2). Figure 8-8 shows the functional blocks of local address generation.

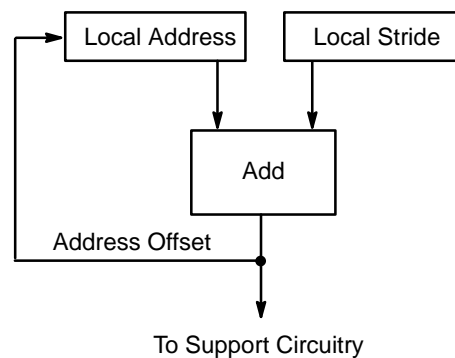


Figure 8-8. Local Address Generation

The BLT generates local addresses when the BLT is creating request packets. The BLT creates information for write request packets during a constant stride write or a scatter operation. The BLT creates information for read request packets during a constant stride read or a gather operation.

When the BLT creates information for write request packets, the BLT first reads the value of the local address register. This value is a word-oriented address offset in local memory where write data for a remote PE is stored. After retrieving 1 or 4 words of data from local memory, the BLT sends the data to the network interface for use in the body of a BLT write request packet.

The BLT then adds the local stride to the first local address offset. The local stride is a value that the BLT repeatedly adds to the local address offset to generate successive address offsets. During single-word BLT transfers, the local stride may be set to one. During cache-line size (hereafter referred to as a block) unit stride BLT transfers, the local stride must be set to a multiple of four.

The sum of the first local address offset and local stride is the second address offset in local memory where write data for a remote PE is stored. After retrieving 1 or 4 words of data from local memory, the BLT uses the data in the body of another write request packet. The BLT continues to use the local addressing circuitry to retrieve data from local memory until the BLT has generated all of the write request packets.

When the BLT creates read request packets, the BLT reads a value stored in the local address register. This value is a word-oriented address offset in local memory where read data from a remote PE will be stored. The BLT uses this address offset in the header of a type 2 read request packet. When a remote PE receives the read request packet, the remote PE uses this address offset in the header of a read response packet.

The BLT then adds the local stride to the first address offset. This sum is the second address offset in local memory where read data from a remote PE will be stored. The BLT uses this second address offset in the header of another read request packet.

During single-word BLT transfers, the value of the local stride may be one. During cache-line size unit stride BLT transfers, the value of the local stride must be a multiple of four.

The BLT continues to read address offsets from memory until the BLT has generated all of the read request packets.

Control and Status Information

The PEs in a processing element node control and monitor the BLT by using the BLT control and status registers. The BLT contains one control register and two status registers (one for each PE in a processing element node). These registers control and contain information on the BLT interrupt bits, the BLT error bits, the vector length register, the outstanding request counter, and the block transfer parameters.

There are 3 BLT interrupt bits. These bits indicate when the BLT is free and ready to accept information for another transfer, indicate when a transfer is complete, and indicate whether a BLT error occurred. If any of these bits are set to 1, the BLT sets the BLT hardware interrupt to the appropriate PE in a processing element node.

There are also 3 BLT error bits that signal the BLT when an error occurs. These errors include the offset range error, PE range error, and index memory error. As described in the “Address Offset Generation” subsection, the offset range error indicates the address offset is larger than the remote limit. The PE range error indicates the PE number is larger than the PE node mask in the BLT control register. The index memory error indicates whether a single- or double-bit error was detected when the BLT read a remote index value from local memory. If any of the three errors occur, circuitry in the BLT sets the BLT error interrupt bit. More information on the BLT error interrupt bit is provided in “Register Mapping” later in this section.

If an error occurs during a transfer, the PE can read the value of the vector length register using the status register. This value indicates how far the transfer progressed before encountering an error.

If an error occurs during a transfer, the outstanding request counter may have to be reset. The PE resets the appropriate outstanding request counter using the BLT status register.

Before initiating a transfer, the PE writes the final transfer parameters into the control register. As soon as the PE writes the final transfer parameters into the control register, the transfer begins. These parameters include the type of transfer operation, the size of the body of the packets, the PE number portion of the remote index (virtual or logical PE number), and the size of the user partition (PE range mask).

Register Mapping

The following subsections describe the addressing and bit assignments for the registers used by the BLT in the CRAY T3D system. Each subsection also provides a brief summary of the function of the register.

Table 8-2 provides a summary of the BLT registers and also lists the partial physical address of each register as it appears on the address pins of the microprocessor.

Table 8-2. BLT Registers

Address	Register Name	Direction	Description
10600000 ₁₆	BLT_RIR	Write	Remote index register
10610000 ₁₆	BLT_RSR	Write	Remote stride register
10640000 ₁₆	BLT_VLR	Write or read	Vector length register
10680000 ₁₆	BLT_IVR	Write	Index vector register
10650000 ₁₆	BLT_RMR	Write	Remote mask register
10660000 ₁₆	BLT_RBR	Write	Remote base register
10670000 ₁₆	BLT_RLR	Write	Remote limit register
10620000 ₁₆	BLT_LAR	Write	Local address register
10630000 ₁₆	BLT_LSR	Write	Local stride register
10690000 ₁₆	BLT_CR	Write	Control register
10414000 ₁₆	BLT_SR	Read or write	Status register

The BLT_VLR may be read using the BLT_SR.

NOTE: Because of multiplexed data paths, when one PE in a processing element node is modifying the contents of the BLT registers, the other PE in the node must not attempt to modify any of the shared registers at the same time. The shared registers include the X_WHOAMI, LPE_XLATE, ROUTE_LO, ROUTE_HI, NET_ENA, NET_PFM, NODE_CSR, and BLT registers.

Because software defines the virtual address, the addresses for each of the registers are given according to the partial physical address as it appears on the pins of the microprocessor.

Remote Index Register**Address 10600000₁₆**

The remote index register (BLT_RIR) is a 36-bit, write-only, system-privileged register that contains the remote index. The remote index is a remote address that points to an aligned 64-bit word in memory. Byte address bits 0 through 2 are not represented in the remote index. Figure 8-9 shows the bit assignments for the BLT_RIR address as they appear on the address pins of the microprocessor.

HEX	1	0				6				0				0				0				0							
Binary	2 ²⁸	2 ²⁷	2 ²⁶	2 ²⁵	2 ²⁴	2 ²³	2 ²²	2 ²¹	2 ²⁰	2 ¹⁹	2 ¹⁸	2 ¹⁷	2 ¹⁶	2 ¹⁵	2 ¹⁴	2 ¹³	2 ¹²	2 ¹¹	2 ¹⁰	2 ⁹	2 ⁸	2 ⁷	2 ⁶	2 ⁵	2 ⁴	2 ³	2 ²	2 ¹	2 ⁰
	1	x	x	x	x	x	1	1	0	0	0	0	0	x	x	x	x	x	x	x	x	x	x	x	x	x	x	x	x

NOTE: x = Don't Care

Figure 8-9. Remote Index Register Address Bit Format

The remote index contains a PE number and an index offset in a software defined bit format. The remote index must have 12 bits assigned to the PE number and 24 bits assigned to the index offset. Table 8-3 shows an example of a remote index in the BLT_RIR. The remote mask register defines the bit positions of the PE number and index offset.

The PE number is either a virtual PE number or a logical PE number. When transferring data for the operating system, the BLT interprets the PE number as a logical PE number. When transferring data for a user program, the BLT interprets the PE number as a virtual PE number.

The index offset is a word-oriented address offset with respect to a remote base address. The BLT adds the index offset to the remote base address to obtain the address offset for a remote address.

During a constant stride read or write operation, the BLT repeatedly adds the value of the remote stride to the remote index in the BLT_RIR to generate successive remote indexes. The first remote index is the value stored in the BLT_RIR. The second remote index is the value stored in the BLT_RIR plus the value of the remote stride. The BLT repeats this process until all the remote indexes have been generated.

During a gather or scatter operation, the BLT adds the value of a remote index read from local memory to the BLT_RIR (and the value of the BLT_RIR is held constant). Because of this characteristic, the BLT_RIR may be set to 0 or may be used as a base value during a gather or scatter operation.

Table 8-3. BLT_RIR

Bits	Description
0	Bit 0 of the PE number
1	Bit 0 of the index offset
2	Bit 1 of the index offset
3	Bit 2 of the index offset
4	Bit 1 of the PE number
5	Bit 2 of the PE number
6	Bit 3 of the index offset
7	Bit 4 of the index offset
8	Bit 5 of the index offset
9	Bit 6 of the index offset
10	Bit 3 of the PE number
11	Bit 4 of the PE number
12	Bit 5 of the PE number
13	Bit 6 of the PE number
14	Bit 7 of the index offset
15	Bit 8 of the index offset
16	Bit 9 of the index offset
17	Bit 7 of the PE number
18	Bit 8 of the PE number
19	Bit 9 of the PE number
20	Bit 10 of the index offset
21	Bit 11 of the index offset
22	Bit 12 of the index offset
23	Bit 10 of the PE number
24	Bit 11 of the PE number
25	Bit 13 of the index offset
26	Bit 14 of the index offset
27	Bit 15 of the index offset
28	Bit 16 of the index offset

Table 8-3. BLT_RIR (continued)

Bits	Description
29	Bit 17 of the index offset
30	Bit 18 of the index offset
31	Bit 19 of the index offset
32	Bit 20 of the index offset
33	Bit 21 of the index offset
34	Bit 22 of the index offset
35	Bit 23 of the index offset
63 – 36	Not used

Remote Stride Register
Address 10610000₁₆

The remote stride register (BLT_RSR) is a 36-bit, write-only, system-privileged register that contains the remote stride. The remote stride is a value that the BLT repeatedly adds to the remote index to generate successive remote indexes during a constant stride read or constant stride write operation.

Figure 8-10 shows the bit assignments for the BLT_RSR address as they appear on the address pins of the microprocessor.

HEX	1	0				6				1				0				0				0							
Binary	2 ²⁸	2 ²⁷	2 ²⁶	2 ²⁵	2 ²⁴	2 ²³	2 ²²	2 ²¹	2 ²⁰	2 ¹⁹	2 ¹⁸	2 ¹⁷	2 ¹⁶	2 ¹⁵	2 ¹⁴	2 ¹³	2 ¹²	2 ¹¹	2 ¹⁰	2 ⁹	2 ⁸	2 ⁷	2 ⁶	2 ⁵	2 ⁴	2 ³	2 ²	2 ¹	2 ⁰
	1	x	x	x	x	x	1	1	0	0	0	0	1	x	x	x	x	x	x	x	x	x	x	x	x	x	x	x	x

NOTE: x = Don't Care

Figure 8-10. BLT_RSR Address Bit Assignments

The remote stride represents a stride by 64-bit data words. Byte address bits 0 through 2 are not represented in the remote stride. Table 8-4 shows the bit format of the BLT_RSR.

Table 8-4. BLT_RSR Bit Format

Bits	Description
35 – 0	Contains the remote stride
63 – 36	Not used

When the BLT performs a block unit stride operation, restrictions apply to the value of the BLT_RSR. More information on the restrictions is provided in the description of the enable large packet size bit in “Control Register” later in this section.

Vector Length Register**Address 10640000₁₆**

The vector length register (BLT_VLR) is a 16-bit, write-only, system-privileged register that contains the vector length. The vector length indicates how many request packets the BLT generates. When the value of the vector length register is set to 0, the BLT generates the maximum number of packets (65,536).

Figure 8-11 shows the bit assignments for the BLT_VLR address as they appear on the address pins of the microprocessor.

HEX	1	0				6				4				0				0				0				0			
Binary	2 ²⁸	2 ²⁷	2 ²⁶	2 ²⁵	2 ²⁴	2 ²³	2 ²²	2 ²¹	2 ²⁰	2 ¹⁹	2 ¹⁸	2 ¹⁷	2 ¹⁶	2 ¹⁵	2 ¹⁴	2 ¹³	2 ¹²	2 ¹¹	2 ¹⁰	2 ⁹	2 ⁸	2 ⁷	2 ⁶	2 ⁵	2 ⁴	2 ³	2 ²	2 ¹	2 ⁰
	1	x	x	x	x	x	1	1	0	0	1	0	0	x	x	x	x	x	x	x	x	x	x	x	x	x	x	x	x

NOTE: x = Don't Care

Figure 8-11. BLT_VLR Address Bit Assignments

The BLT decrements the vector length register each time it creates a request packet. When the vector length is 0, the block transfer is complete.

If an error occurs during a transfer, the PE that requested the transfer can indirectly read the value of the BLT_VLR to determine when the error occurred. A constant correction value may need to be added to the BLT_VLR contents after an error to compensate for the number of pipeline stages in the BLT between the vector length counter and the limit address checking circuits. More information on reading the value of the BLT_VLR is provided in “Status Register” later in this section.

Table 8-5 shows the bit format of the BLT_VLR.

Table 8-5. BLT_VLR Bit Format

Bits	Description
15 – 0	Bits 15 – 0 contain the vector length
63 – 16	Not used

Index Vector Register
Address 10680000₁₆

The index vector register (BLT_IVR) is a 24-bit, write-only, system-privileged register that contains the index vector. The index vector is a word-oriented address offset that points to a 64-bit word in local memory where a remote index is stored. The BLT uses the remote index when creating information for the header of a request packet during a gather or scatter operation.

Figure 8-12 shows the bit assignments for the BLT_IVR address as they appear on the address pins of the microprocessor.

HEX	1	0				6				8				0				0				0				0			
Binary	2 ²⁸	2 ²⁷	2 ²⁶	2 ²⁵	2 ²⁴	2 ²³	2 ²²	2 ²¹	2 ²⁰	2 ¹⁹	2 ¹⁸	2 ¹⁷	2 ¹⁶	2 ¹⁵	2 ¹⁴	2 ¹³	2 ¹²	2 ¹¹	2 ¹⁰	2 ⁹	2 ⁸	2 ⁷	2 ⁶	2 ⁵	2 ⁴	2 ³	2 ²	2 ¹	2 ⁰
	1	x	x	x	x	x	1	1	0	1	0	0	0	x	x	x	x	x	x	x	x	x	x	x	x	x	x	x	x

NOTE: x = Don't Care

Figure 8-12. BLT_IVR Address Bit Assignment

Software must check the value of the index vector to ensure that it represents a valid local address offset. Software must also ensure that the user has read permission for the contiguous local memory locations, starting at the first index vector and ending with the last index vector. The value of the last index vector is calculated using the following equation:

$$\text{First Index Vector} + \text{Vector Length} - 1$$

The BLT reads remote indexes from local memory in cache-line-size blocks (4 remote indexes at a time). Because of this characteristic, the first index vector stored in the BLT_IVR must be cache-line aligned (bits 0 through 4 of the partial physical address set to 0). The last index vector does not need to be cache-line aligned.

Table 8-6 shows the bit format of the BLT_IVR.

Table 8-6. BLT_IVR Bit Format

Bits	Description
23 – 0	Contains the index vector
63 – 24	Not used

NOTE: During a scatter or gather operation, the remote indexes can conflict with updating the shared registers. The shared registers include the X_WHOAMI, LPE_XLATE, ROUTE_LO, ROUTE_HI, NET_ENA, NET_PFM, NODE_CSR, and BLT registers.

Remote Mask Register
Address 10650000₁₆

The remote mask register (BLT_RMR) is a 36-bit, write-only, system-privileged register that contains the remote mask. The remote mask contains 12 bits that are set to 1 to indicate that the corresponding bits of the remote index are used for the PE number. The BLT_RMR also contains 24 bits set to 0 to indicate that the corresponding bits of the remote index are used for the index offset.

Figure 8-13 shows the bit assignments for the BLT_RMR address as they appear on the address pins of the microprocessor.

HEX	1	0				6				5				0				0				0							
Binary	2 ²⁸	2 ²⁷	2 ²⁶	2 ²⁵	2 ²⁴	2 ²³	2 ²²	2 ²¹	2 ²⁰	2 ¹⁹	2 ¹⁸	2 ¹⁷	2 ¹⁶	2 ¹⁵	2 ¹⁴	2 ¹³	2 ¹²	2 ¹¹	2 ¹⁰	2 ⁹	2 ⁸	2 ⁷	2 ⁶	2 ⁵	2 ⁴	2 ³	2 ²	2 ¹	2 ⁰
	1	x	x	x	x	x	1	1	0	0	1	0	1	x	x	x	x	x	x	x	x	x	x	x	x	x	x	x	x

NOTE: x = Don't Care

Figure 8-13. BLT_RMR Address Bit Assignments

The remote mask placed in the BLT_RMR must have exactly 12 bits set to 1 or the centrifuge in the BLT will not operate correctly. Hardware in the BLT does not check to see if there are 12 bits set to 1 in the BLT_RMR. Figure 8-14 shows an example of a remote mask bit format.

Remote Base Register Address 10660000₁₆

The remote base register (BLT_RBR) is a 24-bit, write-only, system-privileged register that contains the remote base. The remote base is the starting word-oriented address offset of a remote data structure that will be transferred by the BLT.

Figure 8-15 shows the bit assignments for the BLT_RBR address as they appear on the address pins of the microprocessor.

HEX	1	0				6				6				0				0				0				0			
Binary	2 ²⁸	2 ²⁷	2 ²⁶	2 ²⁵	2 ²⁴	2 ²³	2 ²²	2 ²¹	2 ²⁰	2 ¹⁹	2 ¹⁸	2 ¹⁷	2 ¹⁶	2 ¹⁵	2 ¹⁴	2 ¹³	2 ¹²	2 ¹¹	2 ¹⁰	2 ⁹	2 ⁸	2 ⁷	2 ⁶	2 ⁵	2 ⁴	2 ³	2 ²	2 ¹	2 ⁰
	1	x	x	x	x	x	1	1	0	0	1	1	0	x	x	x	x	x	x	x	x	x	x	x	x	x	x	x	x

NOTE: x = Don't Care

Figure 8-15. BLT_RBR Address Bit Assignments

The address offset for a remote address generated by the BLT is equal to the remote base plus the index offset. The centrifuge in the BLT separates the remote index into the index offset and PE number.

Software must ensure that correct ownership and access permission exists for the remote memory referenced by the value in the BLT_RBR.

When the BLT performs a block unit stride operation, restrictions apply to the value of the BLT_RBR. More information on the restrictions is provided in the description of the enable large packet size bit in “Control Register” later in this section.

Table 8-7 shows the bit format of the BLT_RBR.

Table 8-7. BLT_RBR Bit Format

Bits	Description
23 – 0	Contains the remote base
63 – 24	Not used

Remote Limit Register
Address 10670000₁₆

The remote limit register (BLT_RLR) is a 24-bit, write-only, system-privileged register that contains the remote limit. The remote limit is the maximum remote word-oriented address offset that a user can access.

Figure 8-16 shows the bit assignments for the BLT_RLR address as they appear on the address pins of the microprocessor.

HEX	1	0				6				7				0				0				0				0			
Binary	2 ²⁸	2 ²⁷	2 ²⁶	2 ²⁵	2 ²⁴	2 ²³	2 ²²	2 ²¹	2 ²⁰	2 ¹⁹	2 ¹⁸	2 ¹⁷	2 ¹⁶	2 ¹⁵	2 ¹⁴	2 ¹³	2 ¹²	2 ¹¹	2 ¹⁰	2 ⁹	2 ⁸	2 ⁷	2 ⁶	2 ⁵	2 ⁴	2 ³	2 ²	2 ¹	2 ⁰
	1	x	x	x	x	x	1	1	0	0	1	1	1	x	x	x	x	x	x	x	x	x	x	x	x	x	x	x	x

NOTE: x = Don't Care

Figure 8-16. BLT_RLR Address Bit Assignments

The BLT compares the contents of the BLT_RLR to each address offset that is generated during a transfer. The address offset must be less than or equal to the value in the BLT_RLR or an address range error occurs.

When an address range error occurs, the BLT operation is aborted. The address offset that is out-of-range is not used in a request packet and a BLT hardware interrupt is sent to the PE that requested the transfer. More information on the BLT errors is provided in “Status Register” later in this section.

Table 8-8 shows the bit format of the BLT_RLR.

Table 8-8. BLT_RLR Bit Format

Bits	Description
23 – 0	Contains the remote limit
63 – 24	Not used

Local Address Register**Address 10620000₁₆**

The local address register (BLT_LAR) is a 24-bit, write-only, system-privileged register that contains the local address. The local address is either the word-oriented address in local memory where data for a BLT write request pack resides or the address in local memory where data from a BLT read response packet will be placed.

Figure 8-17 shows the bit assignments for the BLT_LAR address as they appear on the address pins of the microprocessor.

HEX	1	0				6				2				0				0				0				0			
Binary	2 ²⁸	2 ²⁷	2 ²⁶	2 ²⁵	2 ²⁴	2 ²³	2 ²²	2 ²¹	2 ²⁰	2 ¹⁹	2 ¹⁸	2 ¹⁷	2 ¹⁶	2 ¹⁵	2 ¹⁴	2 ¹³	2 ¹²	2 ¹¹	2 ¹⁰	2 ⁹	2 ⁸	2 ⁷	2 ⁶	2 ⁵	2 ⁴	2 ³	2 ²	2 ¹	2 ⁰
	1	x	x	x	x	x	1	1	0	0	0	1	0	x	x	x	x	x	x	x	x	x	x	x	x	x	x	x	x

NOTE: x = Don't Care

Figure 8-17. BLT_LAR Address Bit Assignments

Software must guarantee that the user has the appropriate read or write permission for the contiguous local memory block, starting at the address pointed to by the BLT_LAR and ending at the physical address represented by the following equation:

$$\text{Local Base Address} + (\text{Local Stride} * (\text{Vector Length} - 1)).$$

The BLT_LAR acts as the local address adder during a transfer operation. The BLT increments the value of the BLT_LAR by the local stride for each request packet the BLT generates. Table 8-9 shows the bit format of the BLT_LAR.

Table 8-9. BLT_LAR Bit Format

Bits	Description
23 – 0	Contains the local address
63 – 24	Not used

When the BLT performs a block unit stride operation, restrictions apply to the value of the BLT_LAR. More information on the restrictions is provided in the description of the enable large packet size bit in “Control Register” later in this section.

Local Stride Register
Address 10630000₁₆

The local stride register (BLT_LSR) is a 24-bit, write-only, system-privileged register that contains the local stride. The local stride is a positive value that the BLT repeatedly adds to the local address to generate successive local addresses during a transfer operation.

Figure 8-18 shows the bit assignments for the BLT_LSR address as they appear on the address pins of the microprocessor.

HEX	1	0				6				3				0				0				0							
Binary	2 ²⁸	2 ²⁷	2 ²⁶	2 ²⁵	2 ²⁴	2 ²³	2 ²²	2 ²¹	2 ²⁰	2 ¹⁹	2 ¹⁸	2 ¹⁷	2 ¹⁶	2 ¹⁵	2 ¹⁴	2 ¹³	2 ¹²	2 ¹¹	2 ¹⁰	2 ⁹	2 ⁸	2 ⁷	2 ⁶	2 ⁵	2 ⁴	2 ³	2 ²	2 ¹	2 ⁰
	1	x	x	x	x	x	1	1	0	0	0	1	1	x	x	x	x	x	x	x	x	x	x	x	x	x	x	x	x

NOTE: x = Don't Care

Figure 8-18. BLT_LSR Address Bit Assignments

When the BLT performs a block unit stride operation, restrictions apply to the value of the BLT_LSR. More information on the restrictions is provided in the description of the enable large packet size bit in “Control Register” later in this section.

Table 8-10 shows the bit format of the BLT_LSR.

Table 8-10. BLT_LSR Bit Format

Bits	Description
23 – 0	Contains the local stride
63 – 24	Not used

Control Register
Address 10690000₁₆

The control register (BLT_CR) is a 27-bit, write-only, system-privileged register. The BLT_CR contains bits that control the type of operation the BLT performs, define the range of legal PE numbers, and enable a transfer.

Figure 8-19 shows the bit assignments for the BLT_CR address as they appear on the address pins of the microprocessor.

HEX	1	0				6				9				0				0			0			0					
Binary	2 ²⁸	2 ²⁷	2 ²⁶	2 ²⁵	2 ²⁴	2 ²³	2 ²²	2 ²¹	2 ²⁰	2 ¹⁹	2 ¹⁸	2 ¹⁷	2 ¹⁶	2 ¹⁵	2 ¹⁴	2 ¹³	2 ¹²	2 ¹¹	2 ¹⁰	2 ⁹	2 ⁸	2 ⁷	2 ⁶	2 ⁵	2 ⁴	2 ³	2 ²	2 ¹	2 ⁰
	1	x	x	x	x	x	1	1	0	1	0	0	1	x	x	x	x	x	x	x	x	x	x	x	x	x	x	x	x

NOTE: x = Don't Care

Figure 8-19. BLT_CR Address Bit Assignments

The BLT_CR is the last register that should be written to before starting a block transfer (except for the BLT_SR). As soon as information is written to the BLT_CR, the BLT starts a transfer that uses the information stored in the other BLT registers. Table 8-11 shows the bit format of the BLT_CR, and the following paragraphs describe each bit.

Table 8-11. BLT_CR Bit Format

Bits	Description
0	Transfer type bit 0
1	Transfer type bit 1
2	Enable large packet size
3	Virtual or logical PE numbers
4	Enable I/O mode
5	Outstanding requests select
15 – 6	Reserved
26 – 16	PE range mask
63 – 27	Not used

Transfer Type Bits 0 and 1

These bits control the type of transfer the BLT performs. Table 8-12 lists the different types of transfers and the corresponding values of the transfer type bits.

Table 8-12. Transfer Types

Bit 1	Bit 0	Description
0	0	Constant stride read transfer
0	1	Constant stride write transfer
1	0	Gather transfer
1	1	Scatter transfer

Enable Large Packet Size Bit 2

This bit controls the body size of packets generated by the BLT. The body of a packet may be either a single 64-bit word or four 64-bit words (a block). When set to 0, this bit signals the BLT to write or read single words to or from local memory and generate packets with 1-word bodies.

When set to 1, this bit signals the BLT to write or read blocks of data to or from local memory and generate packets that have a 4-word body. When this bit is set to 1, several restrictions apply to other BLT registers:

- Bits 0 and 1 of the BLT_LSR must be set to 0 so the local stride is a multiple of four.
- Bits 0 and 1 of the BLT_LAR must be set to 0 so the local address offset is cache-line aligned.
- Bits 0 and 1 of the BLT_RBR must be set to 0 so the remote base address offset is cache-line aligned.
- The value of the BLT_RIR, BLT_RSR (constant stride operations only), and BLT_RMR must be set so that the index offset generated by the centrifuge has bits 0 and 1 set to 0. Bits 0 and 1 of the index offset must be set to 0 so the index offset is cache-line aligned.

The values of the BLT_RIR and BLT_RSR depend on the value set in the BLT_RMR. The following paragraphs provide examples of possible BLT_RMR values and the corresponding register restrictions that apply if the enable large packet size bit is set to 1.

Software may set bits 24 through 35 of the BLT_RMR to 1. In this case, bits 24 through 35 of the remote index are the PE number and bits 0 through 23 of the remote index are the index offset. If the enable large packet size bit is set to 1, the following restrictions apply:

- Bits 0 and 1 of the BLT_LSR must be set to 0 so the local stride is a multiple of four.
- Bits 0 and 1 of the BLT_LAR must be set to 0 so the local address offset is cache-line aligned.
- Bits 0 and 1 of the BLT_RBR must be set to 0 so the remote base address offset is cache-line aligned.
- Bits 0 and 1 of the BLT_RIR must be set to 0 so the remote index generated during a constant stride operation or the base remote index used during a gather or scatter operation has the index offset portion of the remote index cache-line aligned.
- Bits 0 and 1 of the BLT_RSR must be set to 0 so the remote stride used during a constant stride operation increments the index offset portion of the remote index by a multiple of four. The BLT_RSR is not used during a gather or scatter operation.

Software may set bits 0 through 11 of the BLT_RMR to 1. In this case, bits 0 through 11 of the remote index are the PE number and bits 12 through 35 of the remote index are the index offset. If the enable large packet size bit is set to 1, the following restrictions apply:

- Bits 0 and 1 of the BLT_LSR must be set to 0 so the local stride is a multiple of four.
- Bits 0 and 1 of the BLT_LAR must be set to 0 so the local address offset is cache-line aligned.
- Bits 0 and 1 of the BLT_RBR must be set to 0 so the remote base address offset is cache-line aligned.

- Bits 12 and 13 of the BLT_RIR must be set to 0 so the remote index generated during a constant stride operation or the base remote index used during a gather or scatter operation has the index offset portion of the remote index cache-line aligned.
- Bits 0 through 13 of the BLT_RSR must be set to 0 so the remote stride used during a constant stride operation increments the index offset portion of the remote index by a multiple of four. The BLT_RSR is not used during a gather or scatter operation.

Software may set the bits of the BLT_RMR as shown in Figure 8-20. In this case, the bits of the remote index that are used for the PE number are placed sporadically throughout the remote index.




Figure 8-20. Sample BLT_RMR Value

If the enable large packet size bit is set to 1, the following restrictions apply:

- Bits 0 and 1 of the BLT_LSR must be set to 0 so the local stride is a multiple of four.
- Bits 0 and 1 of the BLT_LAR must be set to 0 so the local address offset is cache-line aligned.
- Bits 0 and 1 of the BLT_RBR must be set to 0 so the remote base address offset is cache-line aligned.

- Bits 2 and 6 of the BLT_RIR must be set to 0 so the remote index generated during a constant stride operation or the base remote index used during a gather or scatter operation has the index offset portion of the remote index cache-line aligned.
- Bits 0 through 6 of the BLT_RSR must be set to 0 so the remote stride used during a constant stride operation increments the index offset portion of the remote index by a multiple of four. The BLT_RSR is not used during a gather or scatter operation.

Virtual or Logical PE Number
Bit 3

This bit signals the BLT to interpret the PE number portion of the remote index as a virtual PE number or a logical PE number. When set to 0, the BLT interprets the PE number as a virtual PE number and sends the PE number through the virtual-to-logical translation circuitry. This mode is used when the BLT transfers data for a user.

When set to 1, the BLT interprets the PE number as a logical PE number and does not send the PE number through the virtual-to-logical translation circuitry. This mode is used when the BLT transfers data for the operating system. Also, when bit 3 is set to 1, bits 16 through 26 must be set to 0.

Enable I/O mode
Bit 4

This bit is used only in a BLT that resides in the input or output node of an I/O gateway. When set to 1, this bit enables flow control between the BLT and the HISP channel buffers in the I/O gateway. In a PE node, this bit must be set to 0.

Outstanding Requests Select
Bit 5

This bit controls the maximum value that the outstanding request counter can obtain before the BLT stalls. When set to 0, the BLT stalls if the request counter reaches 256, and the BLT does not continue until the outstanding request counter decrements. When set to 1, the BLT stalls if the outstanding request counter reaches 32 and does not continue until the counter decrements.

PE Range Mask
 Bits 16 through 26

The PE range mask is initialized at the start of a transfer operation with a contiguous, right-justified field of 0's in the legal bit positions for the size of the user partition. The remaining bits are set to 1's. Table 8-13 shows the legal values of the PE range mask for all possible partition sizes in a CRAY T3D system. If the BLT interprets the PE number as a logical PE number, bits 16 through 26 of the BLT_CR must be set to 0.

Table 8-13. PE Range Mask Values

Bit 26	Bit 25	Bit 24	Bit 23	Bit 22	Bit 21	Bit 20	Bit 19	Bit 18	Bit 17	Bit 16	Number of PEs in Partition
1	1	1	1	1	1	1	1	1	1	1	2
1	1	1	1	1	1	1	1	1	1	0	4
1	1	1	1	1	1	1	1	1	0	0	8
1	1	1	1	1	1	1	1	0	0	0	16
1	1	1	1	1	1	1	0	0	0	0	32
1	1	1	1	1	1	0	0	0	0	0	64
1	1	1	1	1	0	0	0	0	0	0	128
1	1	1	1	0	0	0	0	0	0	0	256
1	1	1	0	0	0	0	0	0	0	0	512
1	1	0	0	0	0	0	0	0	0	0	1,024
1	0	0	0	0	0	0	0	0	0	0	2,048
0	0	0	0	0	0	0	0	0	0	0	2,048 + I/O

Status Register
Address 10414000₁₆

The status register (BLT_SR) is a 12-bit, read and write, system-privileged register that a PE may use to check the progress of a transfer and to initiate certain BLT control operations. Unlike the other BLT registers, there are two physical BLT_SR registers, one for each PE in a processing element node. The BLT_SR enables interrupts from the BLT to a PE and contains BLT status information.

Figure 8-21 shows the bit assignments for the BLT_SR address as they appear on the address pins of the microprocessor.

HEX	1	0				4				1				4				0				0				0			
Binary	2 ²⁸	2 ²⁷	2 ²⁶	2 ²⁵	2 ²⁴	2 ²³	2 ²²	2 ²¹	2 ²⁰	2 ¹⁹	2 ¹⁸	2 ¹⁷	2 ¹⁶	2 ¹⁵	2 ¹⁴	2 ¹³	2 ¹²	2 ¹¹	2 ¹⁰	2 ⁹	2 ⁸	2 ⁷	2 ⁶	2 ⁵	2 ⁴	2 ³	2 ²	2 ¹	2 ⁰
	1	x	x	x	x	x	1	0	0	0	0	0	1	0	1	x	x	x	x	x	x	x	x	x	x	x	x	x	x

NOTE: x = Don't Care

Figure 8-21. BLT_SR Address Bit Assignments

Either PE in a processing element node may write information into its corresponding BLT_SR at any time. Because writing information into the BLT_SR may abort BLT code, both PEs in the processing element node must be able to abort the BLT operation.

Table 8-14 shows the bit format of the BLT_SR, and the following paragraphs describe each bit.

Table 8-14. BLT_SR Bit Format

Bits	Read or Write	Description
0	Write only	Enable BLT complete interrupt
1	Write only	Enable BLT free interrupt
2	Write only	Enable BLT error interrupt
5 – 3	Write only	Status select for bits 9 – 6
9 – 6	Read only	Status bits
10	Write only	Enable remote PE BLT free interrupt

Table 8-14. BLT_SR Bit Format (continued)

Bits	Read or Write	Description
11	Write only	Clear enable BLT free interrupt latch set by remote PE
63 – 12	Not used	Not used

Enable BLT Complete Interrupt
Bit 0

When set to 1, this bit indicates that the PE will receive the BLT hardware interrupt when a BLT transfer is complete. The BLT transfer is complete when all responses from all outstanding transfer requests from a PE have returned (the outstanding request counter is 0).

Enable BLT Free Interrupt
Bit 1

When set to 1, this bit indicates that the PE will receive the BLT hardware interrupt when the BLT unit is free. The BLT is free when it transmits the last request element of the last initiated transfer. The BLT may perform another transfer as soon as the BLT is free. More information on the BLT free interrupt is provided in “Enable Remote PE BLT Free Interrupt” later in this section.

Enable BLT Error Interrupt
Bit 2

When set to 1, this bit indicates that the PE will receive the BLT hardware interrupt when a BLT error occurs. There are three possible BLT errors. More information on the BLT errors is provided in “Status Select” later in this section.

Status Select Bits 3 through 5

These bits select the status or error information presented on bits 6 through 9 of the BLT_SR. Table 8-15 lists the status or error information for each status select, and the following subsections describe each function.

Table 8-15. Status Select Codes

Bit5	Bit4	Bit 3	Function
0	0	0	Interrupt status
0	0	1	BLT error status
0	1	0	BLT abort
0	1	1	Reset request counter
1	0	0	BLT_VLR bits 3 – 0
1	0	1	BLT_VLR bits 7 – 4
1	1	0	BLT_VLR bits 11 – 8
1	1	1	BLT_VLR bits 15 – 12

When the status select bits are set to 0, the interrupt status is presented in bits 6 through 9 of the BLT_SR. The interrupt status indicates the present value of the BLT interrupt bits. The state of the interrupt bits is not affected by enabling or disabling the BLT interrupts using bits 0 through 2 of the BLT status register. Table 8-16 shows the bit format of the interrupt bits status in the BLT_SR, and the following paragraphs describe each bit.

NOTE: The status select bits must be set to 0 in order for the BLT to send a BLT hardware interrupt to the support circuitry in a PE.

Table 8-16. Interrupt Status

BLT_SR Bit	Description
6	BLT complete interrupt
7	BLT free interrupt
8	BLT error interrupt
9	Not used

The BLT complete interrupt is different for each PE in the node. When set to 1, this bit indicates that the outstanding request counter for that PE has reached 0 and the BLT operation is complete for that PE.

The BLT free interrupt is identical for each of the two PEs in a processing element node. When set to 1, this bit indicates that the BLT is free and another transfer operation can be initiated. This bit may be used to trigger arbitration between the two PEs for use of the BLT. The BLT free bit remains set to 1 until the next block transfer sequence is initiated by either PE. More information on the BLT free interrupt is provided in “Enable Remote PE BLT Free Interrupt” later in this section.

The BLT error interrupt is different for each PE in the node. When set to 1, this bit indicates that one of the three BLT errors occurred. The BLT sends one hardware interrupt signal to a PE to indicate a BLT interrupt. If any of the three interrupt signals, BLT free, BLT complete, or BLT error, are set to 1, the BLT sets the hardware interrupt to 1. The hardware interrupt can be enabled or disabled by mask bits located in the microprocessor. More information on the BLT hardware interrupt is provided in Section 10, “Control and Status.”

When the status select bits are set to 1, the BLT error status is presented in bits 6 through 9 of the BLT_SR. The BLT error status indicates the present value of the BLT error bits. Table 8-17 shows the bit format of the BLT error status in the BLT_SR, and the following paragraphs describe each bit.

Table 8-17. BLT Error Status

BLT_SR Bit	Description
6	Offset range error
7	Index memory error
8	PE range error
9	Not used

The offset range error may occur when the BLT is generating a remote address offset. When set to 1, this bit indicates that the final remote offset generated by the BLT was greater than the remote limit stored in the BLT_RLR.

The index memory error may occur when the BLT is performing a scatter or gather operation. When set to 1, this bit indicates that a single- or double-bit error was detected when the BLT read a remote index value from local memory. The BLT can detect single- or double-bit errors on the remote index, but it cannot correct the bits.

The PE range error may occur when the BLT is generating the PE number. When set to 1, this bit indicates that the PE number obtained from the remote index is greater than the range of PE numbers set by the PE range mask in the BLT_CR.

When a PE range error occurs, the BLT stops the transfer, however, a packet is still created and sent for the address that was out of range. For example, the PE range mask may indicate a virtual PE range of 0 to 7. When the virtual PE number from the centrifuge was 9, the BLT would indicate a PE range error and stop the BLT transfer, however, a packet would still be created and sent to virtual PE 1. When the packet was a write request packet, data would be written into the memory of virtual PE 1.

The BLT error bits are reset to 0 when a microprocessor writes to the BLT_CR register, when the microprocessor issues the BLT abort command, or when the node is reset.

Setting the status select bits to 2 initiates the BLT abort code. The BLT abort code aborts a transfer prematurely. This code may be used to stop a transfer when there is an error in the PE or when the user wants to terminate an undesirable transfer.

When the BLT abort code is initiated, the BLT stops generating and transferring request packets. Because remote PEs generate response packets, the BLT abort code cannot stop the generation of remote response packets.

When a transfer has been interrupted by the BLT abort code, the value of the BLT_VLR may be read to determine the exact point at which the BLT was interrupted.

The BLT abort code does not clear the outstanding request counter of a PE. If the counter were cleared, it would cause an underflow of the counter if any outstanding BLT request packets existed for that PE when the BLT abort was issued.

Setting the status select bits to 3 initiates the reset request counter code. The reset request counter code resets the outstanding request counter. The counter must be reset if the counter enters a confused state. Asserting the reset request counter code affects only the request counter for the PE that asserts the code.

When the status select bits are set to 4, the status of BLT_VLR bits 0 through 3 is presented in bits 6 through 9 of the BLT_SR. The BLT_VR 0 through 3 status shows the last value of bits 0 through 3 of the BLT_VLR (refer to Table 8-18).

Table 8-18. BLT_VLR Bits 0 through 3 Status

BLT_SR Bit	Description
6	Bit 0 of the BLT_VLR
7	Bit 1 of the BLT_VLR
8	Bit 2 of the BLT_VLR
9	Bit 3 of the BLT_VLR

When the status select bits are set to 5, the status of BLT_VLR bits 4 through 7 is presented in bits 6 through 9 of the BLT_SR. The BLT_VR 4 through 7 status shows the last value of bits 4 through 7 of the BLT_VLR (refer to Table 8-19).

Table 8-19. BLT_VLR Bits 4 through 7 Status

BLT_SR Bit	Description
6	Bit 4 of the BLT_VLR
7	Bit 5 of the BLT_VLR
8	Bit 6 of the BLT_VLR
9	Bit 7 of the BLT_VLR

When the status select bits are set to 6, the status of BLT_VLR bits 8 through 11 is presented in bits 6 through 9 of the BLT_SR. The BLT_VR 8 through 11 status shows the last value of bits 8 through 11 of the BLT_VLR (refer to Table 8-20).

Table 8-20. BLT_VLR Bits 8 through 11 Status

BLT_SR Bit	Description
6	Bit 8 of the BLT_VLR
7	Bit 9 of the BLT_VLR
8	Bit 10 of the BLT_VLR
9	Bit 11 of the BLT_VLR

When the status select bits are set to 7, the status of BLT_VLR bits 12 through 15 is presented in bits 6 through 9 of the BLT_SR. The BLT_VR 12 through 15 status shows the last value of bits 12 through 15 of the BLT_VLR (refer to Table 8-21).

Table 8-21. BLT_VLR Bits 12 through 15 Status

BLT_SR Bit	Description
6	Bit 12 of the BLT_VLR
7	Bit 13 of the BLT_VLR
8	Bit 14 of the BLT_VLR
9	Bit 15 of the BLT_VLR

Enable Remote PE BLT Free Interrupt Bit 10

The enable remote PE BLT free interrupt bit 10 sets a latch in the opposing PE of a processing element node so that the opposing PE will receive the BLT interrupt when the BLT is free. Writing a 0 to this bit has no effect and does not reset the latch in the other PE.

If PE 0 writes a 1 to bit 10 of its BLT_SR, PE 1 will receive a BLT hardware interrupt when the BLT is free, regardless of the state of bit 1 (enable BLT free interrupt bit) in the BLT_SR of PE 1. In order to clear the enable BLT free interrupt latch set by PE 0, PE 1 must write a 1 to bit 11 of its BLT_SR.

NOTE: Bits 3 through 5 of the BLT_SR in PE 1 must still be set to 0 (to select BLT interrupt status) in order for PE 1 to receive the BLT free interrupt.

Bits 10 and 11 of the BLT_SR together with bit 1 of the BLT_SR provide a mechanism that can simplify the software interaction between the two PEs in a processing element node while the PEs contend for access to the shared BLT. The following paragraphs describe the general nature of an algorithm using this mechanism.

A software shared circular queue may be used to arbitrate between and sequence several contending requests to use the shared BLT. Each queue entry represents a request to use the BLT, and both PEs may have multiple entries in the queue.

Both PEs have access to the queue. Because of this characteristic, shared lock variables (implemented using Atomic Swap operations) are used to ensure mutually exclusive access to the input and output ports of the queue. The input and output ports of the queue are accessed with a pair of pointers: the head pointer and the tail pointer.

Figure 8-22 shows a sample BLT software queue. The head pointer points to the entry that will be transferred to the BLT next. The tail pointer points to the entry in the software queue that will receive information next.




Figure 8-22. BLT Software Queue

Using both the local (bit 1) and remote (bit 10) BLT free interrupt enables, it is possible to direct the BLT free interrupt to the PE whose request will be sent to the BLT next. A PE servicing a request that has reached the head of the queue examines the next entry to identify which PE requested

that transfer. The PE servicing the request at the head of the queue then directs the BLT free interrupt to the PE specified in the next entry of the queue.

The PE that is not specified in the next entry of the queue does not receive and does not have to process the next BLT free interrupt. In addition, there is no messaging or other interaction required between the two PEs in order to reliably share the BLT (unless exceptional events, such as errors, occur).

The following paragraphs contain a detailed description of transactions using bits 1 and 10 of the BLT_SR. When a PE needs to use the BLT, it transfers the parameters for the BLT registers into the entry of the queue indicated by the tail pointer. The PE then compares the value indicated by the tail pointer to the value indicated by the head pointer.

When the value indicated by the tail pointer is not equal to the value indicated by the head pointer, the queue contains other entries. In this case, the PE increments the tail pointer after placing information in the queue. For example, in Figure 8-22, PE 1 placed an entry in entry 2 of the queue and then incremented the tail pointer. After placing an entry in the queue, the PE continues with other program instructions and sets bits 3 through 5 (status select bits) and bit 1 (enable BLT free interrupt bit) of its BLT_SR to 0.

When the value of the tail pointer is equal to the value of the head pointer, the queue is empty. In this case, after placing information in an entry of the queue, the PE increments the tail pointer and then sets bit 1 of its BLT_SR to 1 so it will receive the BLT free interrupt. If the BLT is not busy, the PE will immediately receive the BLT free interrupt. If the BLT is busy, the PE will receive the BLT free interrupt after the BLT has created the last packet of the currently active transfer.

When a PE receives the BLT free interrupt, it transfers information from the head of the queue to the BLT. This action starts a transfer by the BLT. After sending information to the BLT, the PE disables its BLT free interrupt by setting bit 1 of its BLT_SR to 0 and setting bit 11 of its BLT_SR to 1.

The PE then checks the next entry in the queue. If this entry is empty, the queue is empty and no more BLT transfers were requested. In this case, the PE just increments the head pointer. If this entry contains information, the PE increments the head pointer and checks the parameters stored in that entry.

If the parameters indicate the transfer is for the same PE, the PE enables the BLT free interrupt for itself by setting bit 1 of its BLT_SR to 1. For example, in Figure 8-23, after transferring information from entry 0 in the queue and incrementing the head pointer, PE 0 sets bit 1 of its BLT_SR to 1 so it will receive the next BLT free interrupt.




Figure 8-23. Same PE Request

If the parameters indicate the transfer is for the other PE, the PE enables the BLT free interrupt for the other PE by setting bit 10 of its own BLT_SR to 1. For example, in Figure 8-24, after transferring information from slot 1 in the queue and incrementing the head pointer, PE 0 sets bit 10 of its BLT_SR to 1 so PE 1 will receive the next BLT free interrupt. PE 0 then continues with other program instructions.




Figure 8-24. Other PE Request

Clear Enable BLT Free Interrupt Latch Set by Remote PE
Bit 11

When set to 1, this bit clears the enable BLT free interrupt latch set by the other PE in a processing element node. Writing 0 to this bit has no effect.

Hardware Characteristics

The following subsections describe characteristics, such as timing and signal flow through the CRAY T3D hardware, that may have an effect on software.

BLT Complete Interrupt

The BLT complete interrupt sets to 1 when the BLT outstanding request counter for a PE decrements to 0. In the CRAY T3D system, there are two BLT outstanding request counters per processing element node. These counters reside in the E-series options that make up the network interface and BLT (refer to Figure 8-25).

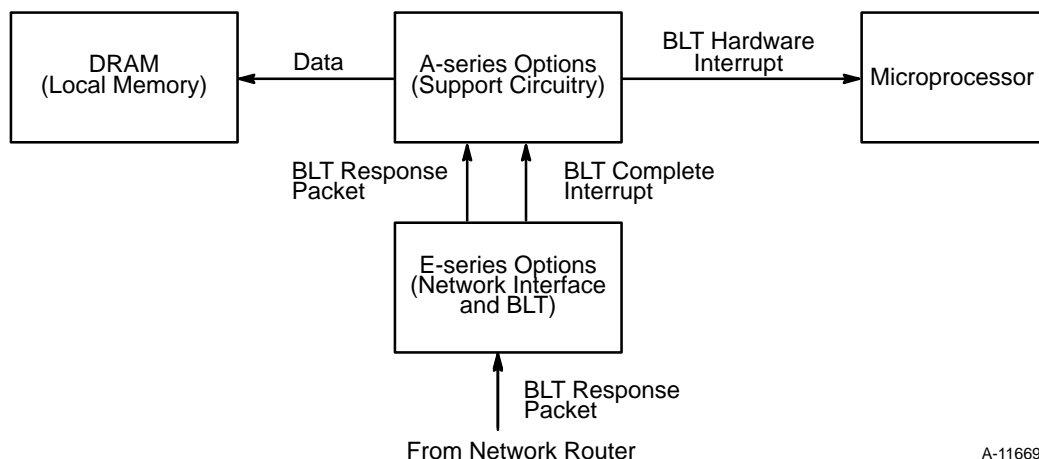


Figure 8-25. BLT Complete Interrupt

Each time the E-series options create a BLT request packet, the BLT outstanding request counter for the appropriate PE is incremented by 1. Each time the E-series options receive a BLT response packet from the network, the BLT outstanding request counter for the appropriate PE is decremented by 1.

When the BLT outstanding request counter for a PE decrements to 0, the E-series options send the BLT complete interrupt to the support circuitry in the appropriate PE. In addition, the E-series options send the BLT response packet to the support circuitry (refer to Figure 8-25).

After receiving a BLT response packet and the BLT complete interrupt, the support circuitry sets the BLT hardware interrupt for the microprocessor. When the BLT response packet is a read response, the support circuitry attempts to write the read data into local memory.

The support circuitry arbitrates access to local memory between microprocessor read and write requests, network interface read and write requests, BLT read requests, and refresh requests. Because of this characteristic, the support circuitry may not write the data from a BLT read response packet into local memory immediately after receiving the packet from the network interface. The support circuitry sends the BLT hardware interrupt to the microprocessor as soon as the write occurs. (refer to Figure 8-25).

BLT Memory-mapped Register Write Ordering

The following list is a suggested ordering for writing to the BLT memory-mapped registers:

- 1 . Write to the following registers (in any order):
 - BLT_RIR
 - BLT_RSR
 - BLT_VLR
 - BLT_IVR
 - BLT_RMR
 - BLT_RBR
 - BLT_RLR
 - BLT_LAR
 - BLT_LSR
- 2 . Write to the BLT_CR. As soon as information is written into the BLT_CR, the BLT starts the transfer.
- 3 . Write a value of 0 to the BLT_SR to reset all of the bits in the register, and select the BLT interrupt bits to be displayed in bits 6 through 9 of the BLT_SR.
- 4 . Write to the BLT_SR again, and set bits 0, 1, and/or 2 to 1 to enable the appropriate BLT interrupts for the transfer. These bits were not set in Step 3 to ensure that the write to the BLT_CR (Step 2), which also resets the BLT interrupt bits, occurs before the BLT_SR is written to enable the BLT interrupts.

Hardware Description

The hardware that makes up the BLT is contained in one option: the EE option. The EE option contains the BLT memory-mapped registers, the BLT centrifuge circuitry, the remote address generation circuitry, the local address generation circuitry, the PE range check circuitry, and the offset limit check circuitry. Refer to Figure 8-26.

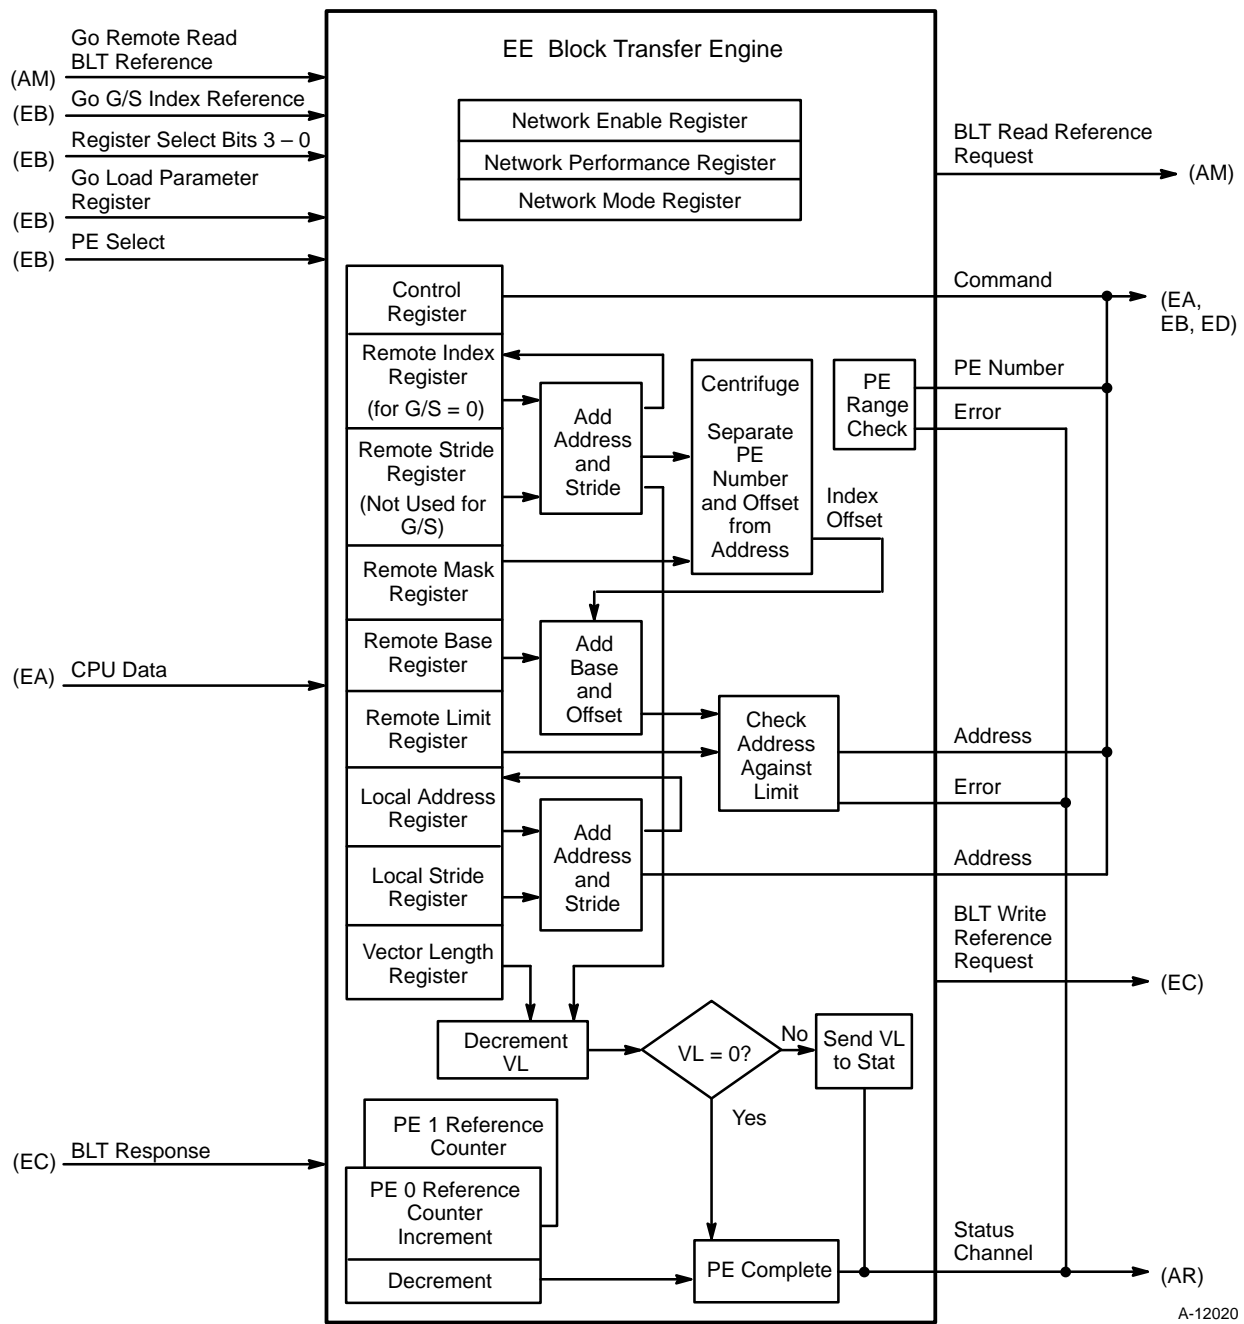


Figure 8-26. EE Option

The microprocessor initializes a block transfer by loading the BLT memory-mapped registers with parameter information. Once this is done, the BLT generates the necessary addresses and control to transfer the blocks of data. When the block transfer is complete, the BLT notifies the microprocessor. The following subsections describe how the BLT is initialized and how the BLT transfers blocks of data.

Initializing the BLT or EE Option

The microprocessor initializes the BLT by loading parameter information into the BLT memory-mapped registers. The BLT memory-mapped registers reside in the EE option, with the exception of the BLT status register and the index vector register. The BLT status register resides in the AR option, and the index vector register resides in the EC option.

To load the BLT memory-mapped registers, the microprocessor supplies parameter information, address, and control to the PE support circuitry. The PE support circuitry interprets this address and control and steers the parameter information to the proper BLT memory-mapped register.

The following parameter information is generated by the microprocessor:

- The remote index register is loaded with the remote PE number and index offset during a constant stride read or constant stride write operation. During a gather or scatter operation, the remote index register is loaded with a 0.
- The remote stride register is loaded with a value that is added to the remote index to generate successive remote indexes.
- The local address register is loaded with the local memory address in the source PE where remote write data resides or where remote read data will be stored.
- The local stride register is loaded with a value that is added to the local address. The sum of the local stride register and the local address register produces successive local addresses of the write data or read data storage locations.
- The vector length register is loaded with a value that indicates how many request packets the BLT needs to generate. This value is decremented each time the BLT creates a request packet. When the vector length equals 0, the request portion of the block transfer is complete.

- The remote mask register is loaded with 12 bits set to 1 and 24 bits set to 0. The 12 bits set to 1 indicate which bits in the remote index are the PE number. The 24 bits set to 0 indicate which bits are the address offset.
- The remote base register is loaded with an address offset that is added to the value of the remote index offset. The sum of the remote base register and the remote index offset produces the remote offset.
- The remote limit register is loaded with an address offset that specifies the highest remote offset that can be referenced. This address offset is compared to the sum of the remote base and remote index. If the sum of the remote base and remote index is greater than the remote limit, an address range error occurs. When this happens, the EE option sends the error information to the AR option by way of the status channel. The AR option sets the error interrupt, which informs the microprocessor of the error.
- The index vector register is loaded with an address offset that points to an address in local memory where a remote index is stored.
- The BLT control register is loaded with 16 bits of information that specify the type of BLT operation, define the range of legal PE numbers, and enable the transfer.

The microprocessor sends the parameter information to the AJ and AK options. The parameter information is held there until the AJ and AK options are instructed to transfer the parameter information to the memory-mapped registers.

The address and control used by the PE to steer the parameter information to the correct memory-mapped registers consists of an address offset and a PE cycle request code. The microprocessor generates the address offset and PE cycle request code for the memory-mapped register reference like it does for a memory reference. The microprocessor does not differentiate between loading a memory-mapped register and loading a memory reference.

The microprocessor sends bits 5 through 28 of the address offset to the AM option and sends bits 29 through 33 to the AR option. The microprocessor also sends the PE cycle request code to the AM option.

The AR option uses bits 29 through 33 of the address offset to address a location in the DTB annex. This location stores a logical or virtual node number and a memory function code. The node number and the memory

function code are sent to the AM option. For a register load, the node number equals the logical or virtual node number of the local PE and the memory function code equals 000 for a normal load.

At this point, the AM option has the virtual or logical PE number, a memory function code, address offset bits 5 through 28, and the PE cycle request code. The AJ and AK options are holding the parameter information. Refer to Figure 8-27.

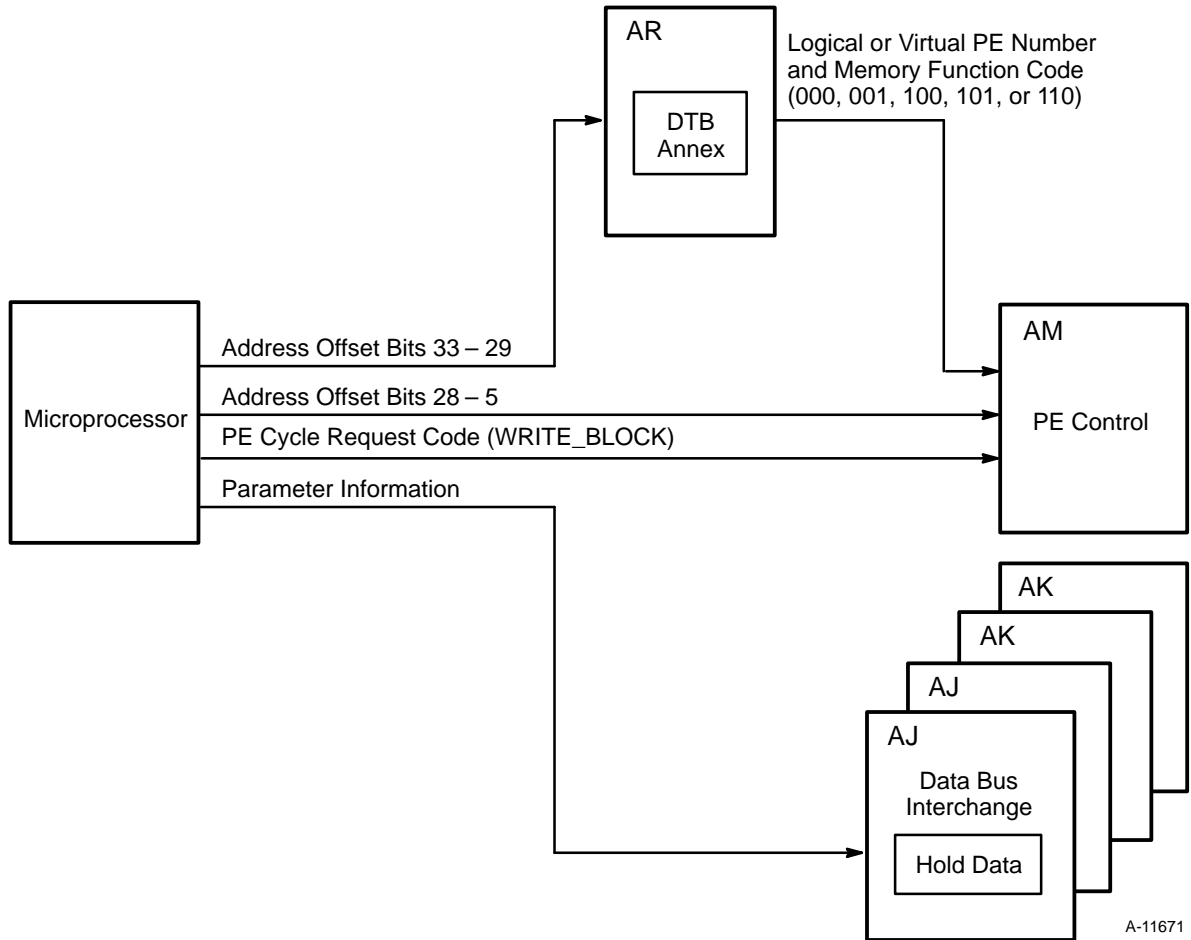


Figure 8-27. Control and Address for Loading the BLT Memory-mapped Registers

The AM option uses the logical or virtual PE number from the DTB annex to determine whether the destination of the parameter information is local or remote (refer to Figure 8-28). The AM option compares the PE number to the contents of the logical PE number register and the virtual PE number register. Because loading a BLT memory-mapped register is a local reference, the node number will equal one of the two PE number registers.

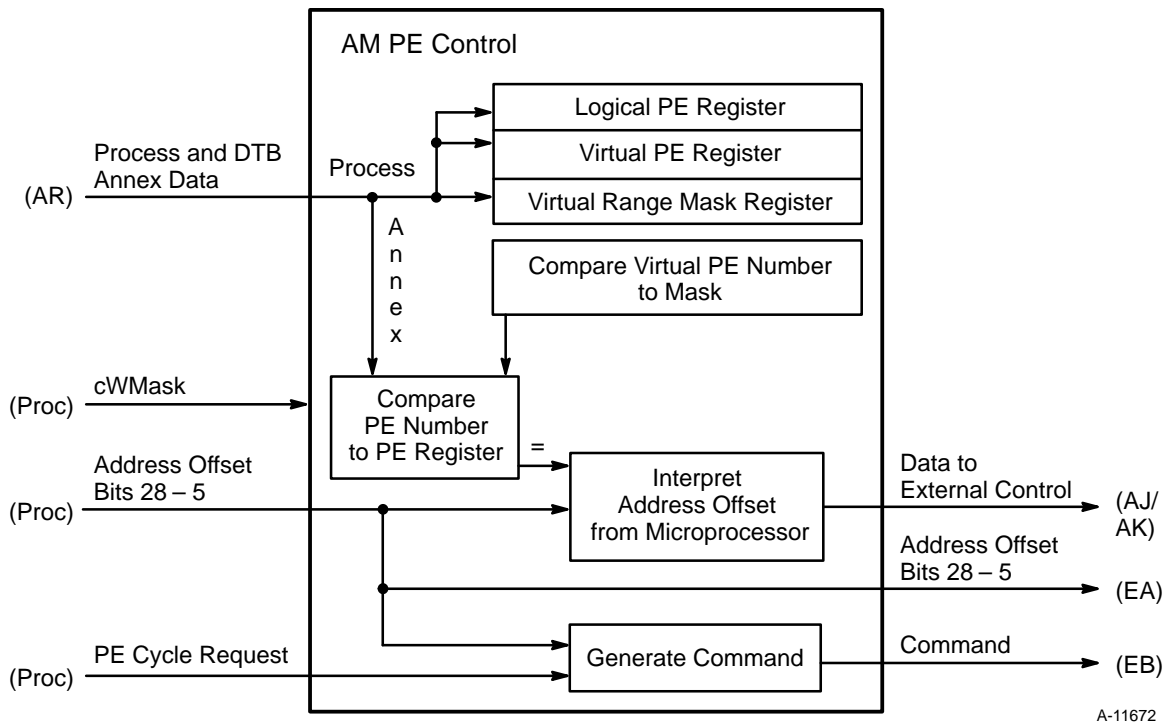


Figure 8-28. AM Option Circuitry that Determines BLT Memory-map Reference

When the AM option determines that a local reference is being requested, it interprets bits 5 through 28 of the address offset to determine whether the request is a local memory reference or a register reference. Table 8-22 lists the values of bits 5 through 28 of the address offset that specify a reference to one of the BLT memory-mapped registers. Table 8-23 shows the decode of the address offset.

Table 8-22. Address Offset Bits 5 through 28

BLT Memory-mapped Register	Address Offset Bits																													
	28	27	26	25	24	23	22	21	20	19	18	17	16	15	14	13	12	11	10	9	8	7	6	5						
BLT_RIR	1	0	0	0	0	0	1	1	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0
BLT_RSR	1	0	0	0	0	0	1	1	0	0	0	0	1	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0
BLT_LAR	1	0	0	0	0	0	1	1	0	0	0	1	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0
BLT_LSR	1	0	0	0	0	0	1	1	0	0	0	1	1	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0
BLT_VLR	1	0	0	0	0	0	1	1	0	0	1	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0
BLT_RMR	1	0	0	0	0	0	1	1	0	0	1	0	1	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0
BLT_RBR	1	0	0	0	0	0	1	1	0	0	1	1	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0
BLT_RLR	1	0	0	0	0	0	1	1	0	0	1	1	1	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0
BLT_IVR	1	0	0	0	0	0	1	1	0	1	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0
BLT_CR	1	0	0	0	0	0	1	1	0	1	0	0	1	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0
BLT_SR	1	0	0	0	0	0	1	0	0	0	0	0	1	0	1	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0


Table 8-23. Decode of Address Offset

Bits	Description
13 – 5	Specifies the address when loading the ROUTE_LO and ROUTE_HI registers
17 – 14	Specifies memory-mapped register
18	1 = AM Option
19	1 = Prefetch Queue
20	1 = Network Upper
22 – 21	00 = AM or AR User Register 01 = AJ or AK Write 10 = AM or AR Privileged 11 = External Write Command
23	1 = Upper Barrier Register
24	0
25	1 = HISP Buffer Data
26	1 = I/O 0 = DRAM
27	1 = Disable Error Correction
28	1 = Memory-mapped register

When the address offset translates into a register load, the AM option generates a Data to External Control signal that is sent to the AJ and AK options. This control instructs the AJ and AK options to steer the parameter information from the microprocessor to the EA option, using 16-bit transfers. The AM option also sends the address offset to the EB option.

The AM option uses the PE cycle request code along with address offset bits 21 and 28 to generate a command. The command for loading a BLT register is 100 0 010 xxxxx. The first 5 bits of the command (xxxxx) are not used for this type of reference. The next 3 bits specify what type of reference is being made (For example, 010 = local register write). Bit 8 specifies that there is no response needed for register writes. Bits 9 through 11 specify the packet type. The command is sent to the EB option in the network interface.

The EB option uses the address and command it receives from the AM option to generate the Go Scatter and Gather Index Reference, Register Select 0 through 3, PE Select, and Go Load Parameter Register signals. These signals are sent to the EA, EC, and EE options so that the BLT memory-mapped registers are loaded properly (refer to Figure 8-29). Table 8-24 lists the decode of the register select.



A-11674

Figure 8-29. BLT Memory-map Reference Control Generated by the EB Option

Table 8-24. Decode of Register Select Bits

BLT Register	Register Select Bits 3 – 0
BLT_RIR	0000
BLT_RSR	0001
BLT_LAR	0010
BLT_LSR	0011
BLT_VLR	0100
BLT_RMR	0101
BLT_RBR	0110
BLT_RLR	0111

Table 8-24. Decode of Register Select Bits (continued)

BLT Register	Register Select Bits 3 – 0
BLT_IVR	1000
BLT_CR	1001

The EA option uses the PE Select signal to input the parameter information for the correct PE. The EA option passes this information to the EE and EC options.

The EE option uses the Go Load Parameter Register signal to input the Register Select signals and parameter information. The EE option uses the Register Select signals to steer the parameter information to the correct memory-mapped register. The EE option uses the PE Select signal to initiate the block transfer for the correct PE.

The EC option also uses the Go Load Parameter Register signal to input the Register Select signals and the scatter and gather index parameter information. The EC option uses the Register Select signals to steer the scatter and gather index information to the index vector memory-mapped register.

The EC option also receives the Go Scatter and Gather Index Reference signal. The EC option uses this signal to enable a scatter and gather index counter.

At this point, the BLT memory-mapped registers have all been initialized. The last memory-mapped register that should be loaded with parameter information is the BLT control register. Once this register is loaded, the BLT begins the block transfer operation.

Transferring Blocks of Data

After the initialization of the memory-mapped registers, the BLT generates all of the necessary addresses and control needed to retrieve or store blocks or single-words of data. For the CRAY T3D system, a block of data is usually defined as four 64-bit words of data.

To retrieve data from the local memory of other PEs, a remote read operation is performed. To store data in the local memory of other PEs, a remote write operation is performed. The following subsections describe how the hardware performs the remote read operation and the remote write operation.

Remote Read

For a remote read operation, the BLT in the source node generates the addresses and control that are needed for the request. The addresses consist of a remote address and a local address. The control consists of a command and a destination PE number. The BLT also notifies the AM option in the source PE of the remote read operation (refer to Figure 8-30). For more information on how the BLT or EE option generates the address and control, refer to the “BLT Register Functions” subsection at the beginning of this section.

After receiving notification of the remote read operation, the AM option arbitrates for a request buffer in the EA option and the outgoing external channel. Once the BLT has priority to use a request buffer and the outgoing external channel, the AM option notifies the EE option.

When the EE option has permission to access a request buffer, it sends the local and remote addresses to the EA option and the command and destination PE number to the EB option. The EA option buffers the address in one of the request buffers. As the addresses are written into the request buffer, parity is generated for the address.

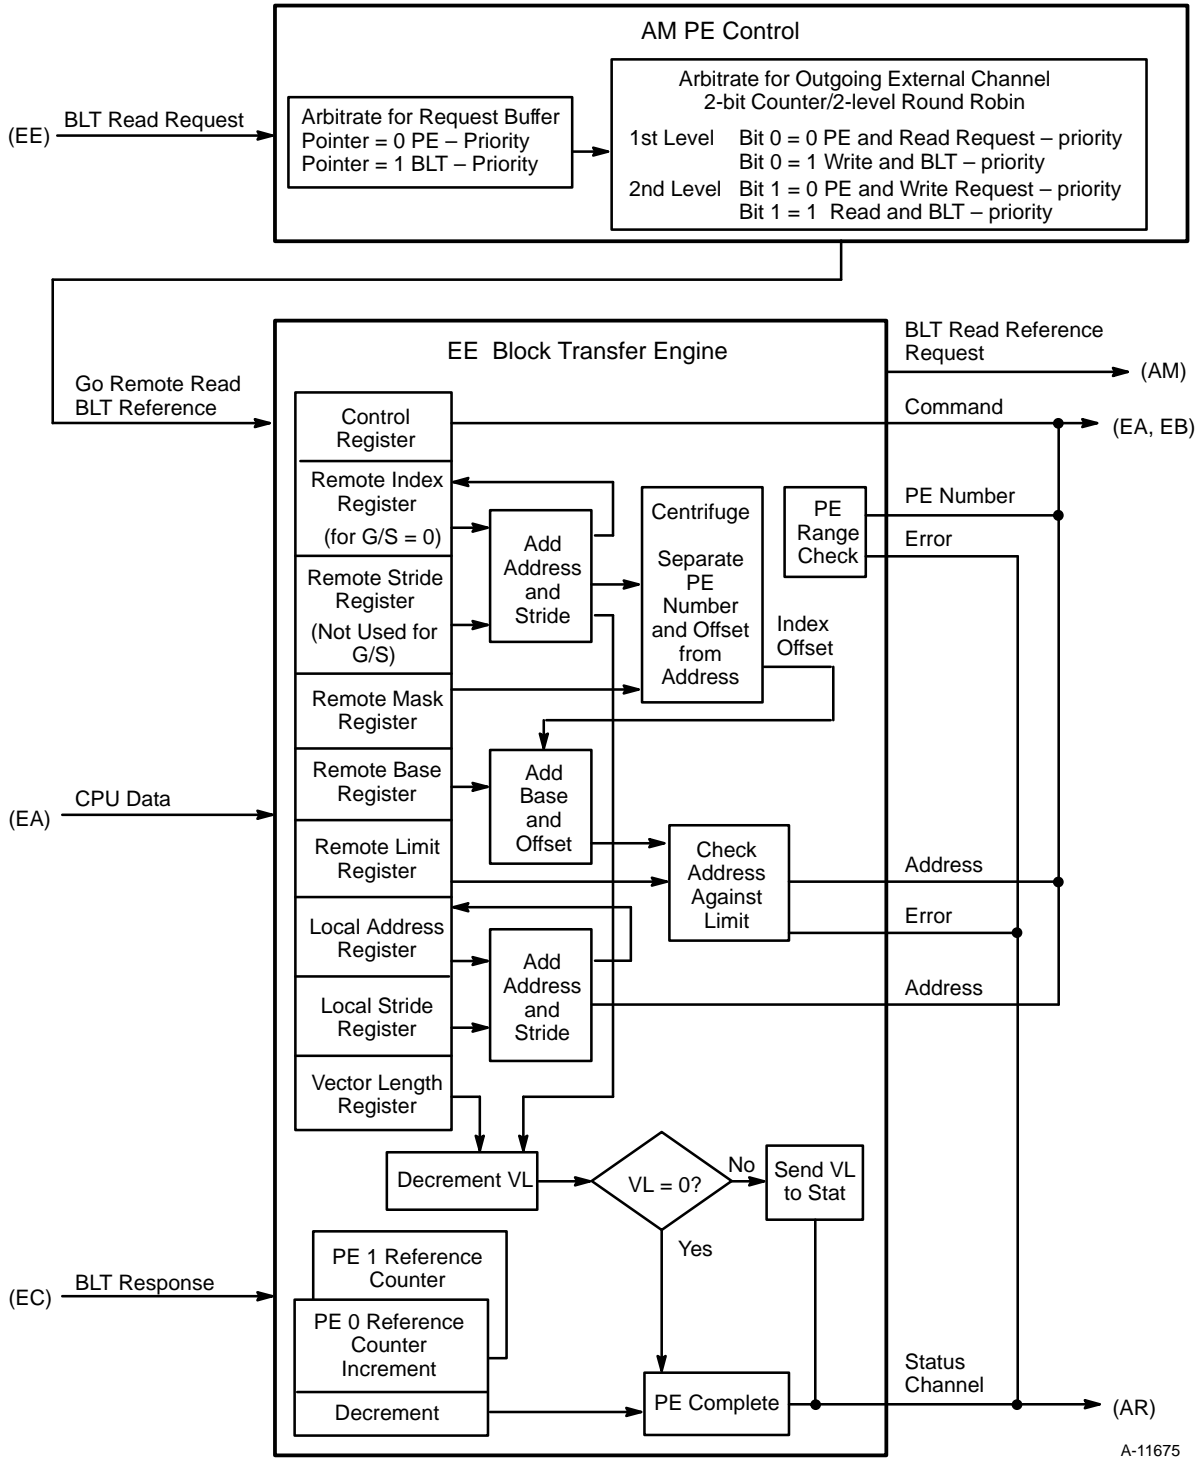
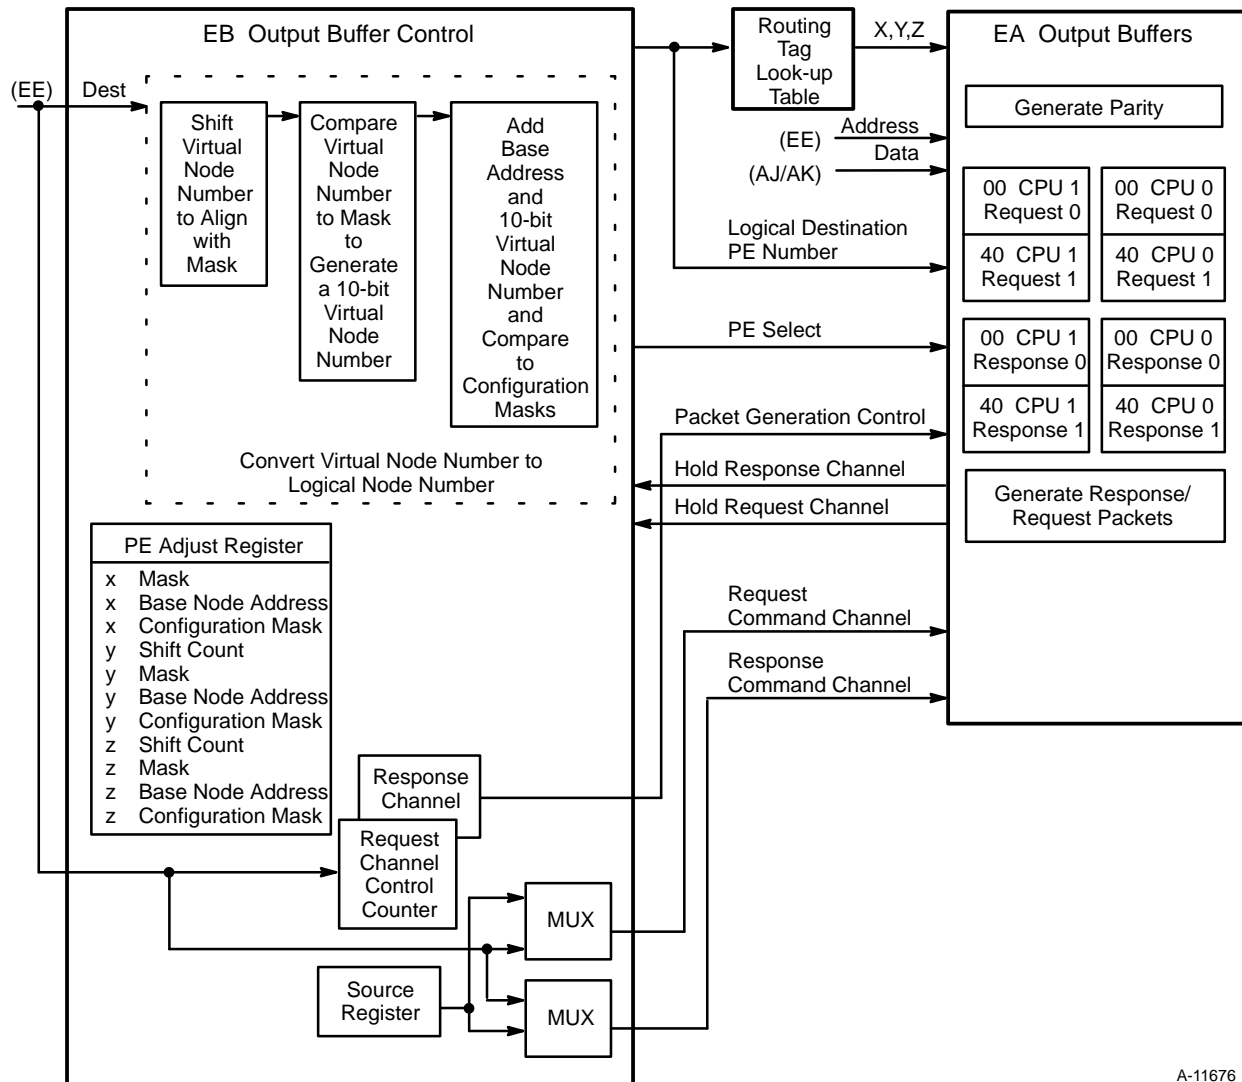


Figure 8-30. BLT Generation of Address and Control

The EB option inputs the command and destination PE number and analyzes this information. From the command, the EB option derives the information needed to generate signals that control the packet assembly on

the EA option (refer to Figure 8-31). The EB option also determines whether the destination PE number is virtual or logical, using bit 12 of the destination PE number. If the destination PE number is virtual, the EB option converts it into a logical PE number.



A-11676

Figure 8-31. EA Option, EB Option, and Routing Tag Look-up Table

The EB option sends the command, destination PE number, and packet generation control to the EA option. The EB option also sends the contents of its source PE number register to the EA option, and it sends the logical PE number to the routing tag look-up table.

The logical PE number is used to address a location in the routing tag look-up table. This location contains the routing tag used to route the request packet through the network. The routing tag is also sent to the EA option.

At this point, the EA option has the local and remote addresses, the command, the destination PE number, the source PE number, and the routing tag. All of this information will be assembled into a request packet. The EA option uses the packet generation control signals received from the EB option to read the addresses out of the request buffer; generate parity for the command, destination PE number, and the source PE number; and assemble this information into a type 2 packet (refer to Figure 8-32).





Figure 8-32. Request Packet Assembly

Under the control of the EB option, the EA option sends the request packets to the SR option that handles the X dimension of the interconnect network. From the SR option, the response packet is transferred through the network to the destination node.

NOTE: The number of request packets generated for the block transfer will vary depending upon how many blocks of data are transferred. For the remote read, each request packet will retrieve 1 block of data.

When the request packet reaches the destination node, the EC and ED options input the request packet. The ED option buffers the entire request packet, except for the routing tag phit. The SR option continues to send the packets to the EC and ED options as long as the EC option responds to the SR option with an acknowledge signal. The EC option will not respond with an acknowledge signal when the buffers on the ED option are full (refer to Figure 8-33).



A-11678

Figure 8-33. Input of Request Packet

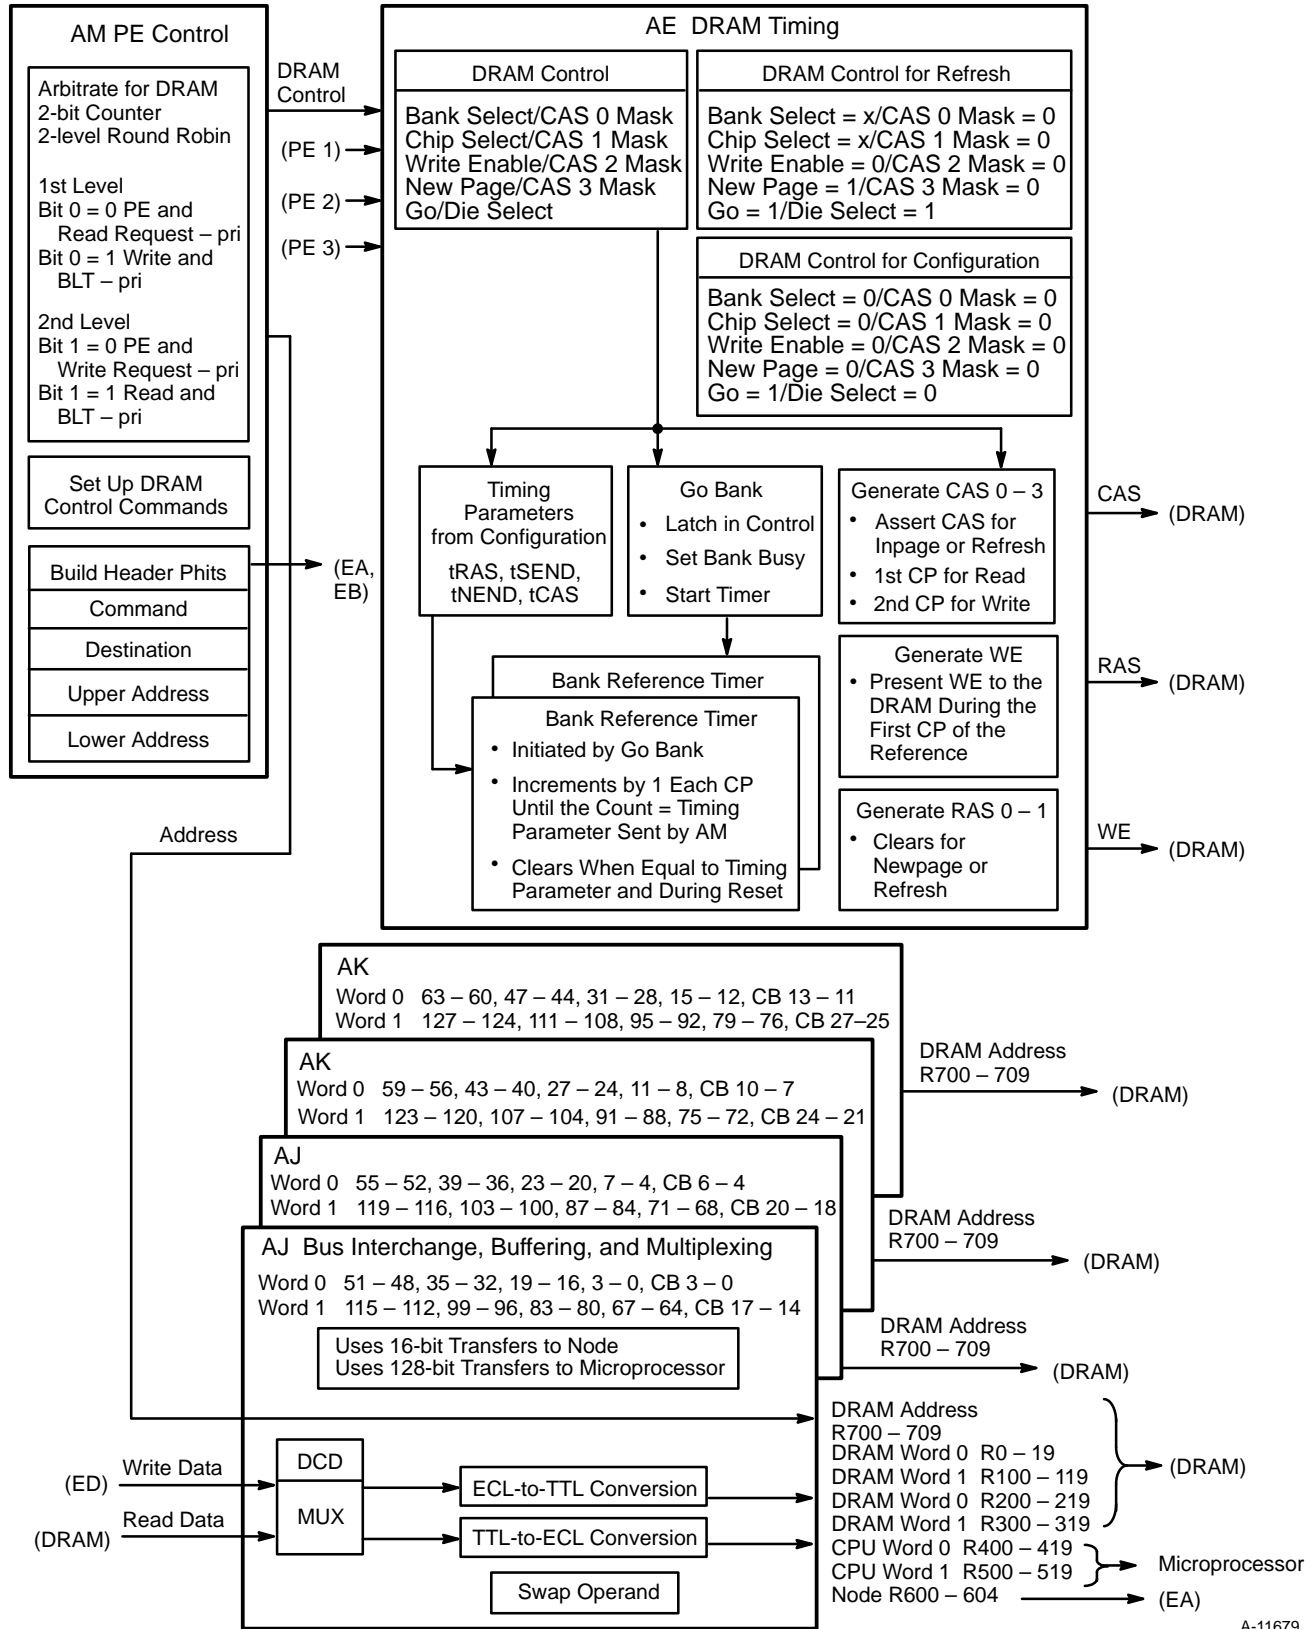
The EC option receives the packet and examines it to determine the status of the request. The EC option determines which PE of the node is destined for the request; whether the request is a read or write request; whether there are any parity errors on the command, address, and source; and whether the packet arrived at the correct destination. When there is a

network parity error or a destination error, the entire packet is written into a message queue so that it can be examined to determine the cause of the error.

The EC option services requests for both PEs and the BLT. If both PEs or a PE and the BLT require servicing simultaneously, the EC option decides which one to service first using round robin (rotating priorities) arbitration. The EC option also generates the control that the ED option uses to read the packet from its buffers.

As stated, the ED option buffers the incoming packets. As soon as the EC option determines that one of the ED option's buffers is full, it sends the ED option the PE Active signal. This signal enables the ED option to read the packet out of the buffer. Once the packet is read from the buffer, the ED option checks the packet for network buffer parity errors. Next, the ED option sends the command, remote address, local address, and source PE number of the packet to the AM option.

From the command, the addresses, and the source PE number, the AM option yields DRAM control signals and response header phits. The DRAM control signals, along with the remote address, are sent to the AJ, AK, and AE options (refer to Figure 8-34).



A-11679

Figure 8-34. Generation of DRAM Control and Response Packet Header Phits

The AJ and AK options, along with the AE option, address the DRAM. The AE option supplies the RAS, CAS, and WE to the DRAM. The AJ and AK options supply the internal address. After the DRAM is addressed, the AJ and AK options read the data from the DRAM.

The AM option generates a new command and destination PE number using the command and source PE number phits of the request packet. This information is sent to the EB option, and the local address is sent to the EA option. The EA option buffers the address in the response buffer until it is instructed to build a response packet by the EB option.

The AM option also instructs the AJ and AK options to send the data to the EA option. The EA option buffers the data in the same response buffer as the address. The EA option also receives the command and destination PE number and the source PE number information from the EB option (refer to Figure 8-35).





Figure 8-35. Reading Data from the DRAM

The EB option instructs the EA option to read the address and data out of the response buffer; generate parity for the command, destination PE number, and the source PE number; and assemble this information into a type 6 response packet (refer to Figure 8-36). The EA option sends the response packet to the SR option that handles the X dimension of the interconnect network. From this option, the response packet is transferred through the network, back to the source node.



A-11681

Figure 8-36. Response Packet Assembly

When the response packet reaches the source node, the EC and ED options input the response packet. The ED option buffers the entire response packet except for the routing tag phit. The SR option continues to send the packets to the EC and ED options as long as the buffer on the ED option is not full (refer again to Figure 8-33).

The EC option examines the response packet to determine which PE of the node will receive the response; whether the response is a read or write reference; whether there are any parity errors on the command, address, and source; and whether the response packet arrived at the correct destination. When there is a parity error or a destination error, the entire packet is written into a message queue so that it can be examined to determine the cause of the error.

As soon as the buffer on the ED option is full, the EC option sends the ED option a PE Active signal, which enables the ED option to read the packet from the buffer. Once the packet is read from the buffer, the ED option checks the packet for network buffer parity errors. Next, the ED option sends the command, remote address, and source PE number of the packet to the AM option. The data is sent to the AJ and AK options.

The EC option also notifies the EE option of the response. This decrements the outstanding reference counter. When the counter equals zero, the EE option sends a signal to the AR option to signify that the BLT reference is complete. This signal prompts the AR option to set the BLT interrupt. The microprocessor polls this interrupt and responds accordingly.

The AM option uses the command and address it receives from the ED option to generate the memory control signals needed to write the data into the DRAM. The DRAM control signals, along with the remote address, are sent to the AJ, AK, and AE options (refer again to Figure 8-34).

The AJ and AK options, along with the AE option, address the DRAM. The AE option supplies the RAS, CAS, and WE to the DRAM. The AJ and AK options supply the internal address. After the DRAM is addressed, the AJ and AK options write the data to the DRAM.

Remote Write

For a remote write operation, the BLT in the source node generates the addresses and control that are needed for the request. The addresses consist of a remote address and a local address. The control consists of a command and a destination PE number. For more information on how the BLT or EE option generates the address and control, refer to the “BLT Register Functions” subsection at the beginning of this section.

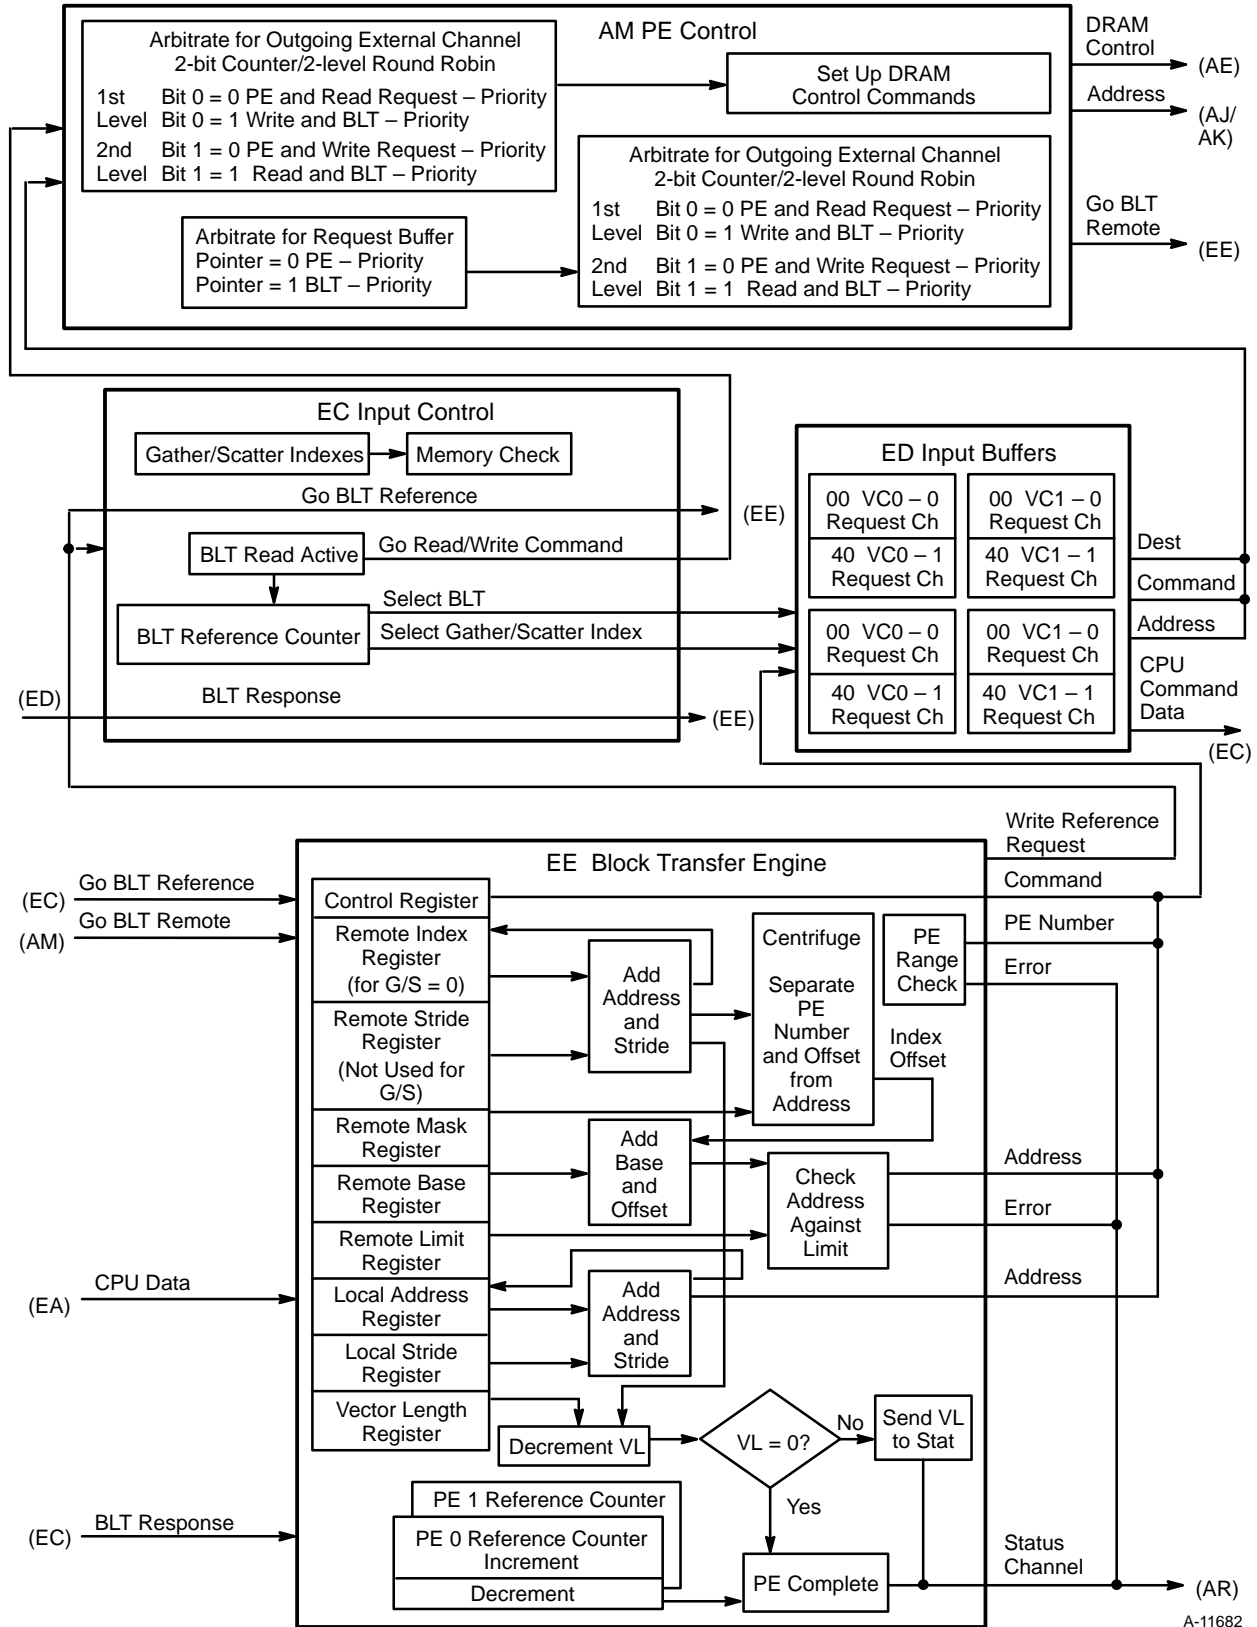
The BLT also generates the BLT Write Reference Request control signal using bit 0 of the BLT control register. The BLT sends this signal to the EC option. The EC option uses this signal to turn the BLT remote write request into a BLT local read request (to do a BLT remote write, a local read must take place first to retrieve the data). The EC option signals the ED option to input the command, destination PE number, and local address from the EE option (refer to Figure 8-37).

The EC option also activates the control signals Go Write Command and Go Read Command that are sent to the AM option. These signals enable the AM option to input the command, destination PE number, and local address from the ED option.

The AM option uses the command, destination PE number, and local address to set up for the local read and to set up for the BLT remote write. To set up for the local read, the AM option arbitrates for the DRAM and generates the DRAM control signals. The DRAM control signals are sent to the AE option, and the local address is sent to the AJ and AK options. To set up for the BLT remote write, the AM option arbitrates for a request buffer in the EA option and for the outgoing channel. After a successful arbitration, the AM option reserves the request buffer and the outgoing channel for the BLT.

The AJ and AK options, along with the AE option, address the DRAM. The AE option supplies the RAS, CAS, and WE to the DRAM. The AJ and AK options supply the internal address. After the DRAM is addressed, the AJ and AK options read the data from the DRAM. Next, the AM option instructs the AJ and AK options to send the data to the EA option where it is held in a request buffer.

When the EE option has permission to access a request buffer, it sends the remote addresses to the EA option and sends the command and destination PE number to the EB option. The EA option buffers the address in the request buffer.



A-11682

Figure 8-37. Control to Read the Remote Write Data from Local Memory

The EB option inputs the command and destination PE number and analyzes this information. From the command, the EB option derives the information needed to generate signals that control the packet assembly on the EA option. The EB option also determines whether the destination PE number is virtual or logical using bit 12 of the destination PE number. If the destination PE number is virtual, the EB option converts it into a logical PE number.

The EB option sends the command, destination PE number, and packet generation control to the EA option. The EB option also sends the EA option the contents of its source PE number register. The logical PE number is sent to the routing tag look-up table.

The logical PE number is used to address a location in the routing tag look-up table. This location contains the routing tag used to route the request packet through the network. The routing tag is also sent to the EA option.

At this point, the EA option has the remote addresses, the command, the destination PE number, the source PE number, and the routing tag. All of this information will be assembled into a request packet. The EA option uses the packet generation control signals it received from the EB option to read the remote address from the request buffer; generate parity for the command, destination PE number, and source PE number; and assemble this information into a type 6 packet.

Under the control of the EB option, the EA option sends the request packets to the SR option that handles the X dimension of the interconnect network. From this option, the response packet is transferred through the network to the destination node.

NOTE: The number of request packets that are generated for the block transfer will vary, depending upon how many blocks of data are being transferred. For the remote write, each request packet will transfer one block of data to the destination PE.

When the request packet reaches the destination node, the EC and ED options input the request packet. The ED option buffers the entire request packet except for the routing tag phit. The SR option continues to send the packets to the EC and ED options until the buffer on the ED option is full.

The EC option examines the request packet to determine which PE of the node is destined for the request; whether the request is a read or write request; whether there are any parity errors on the command, address, and source; and whether the packet arrived at the correct destination.

As soon as the EC option determines that the buffer on the ED option is full, it sends the ED option a PE Active signal, which enables the ED option to read the packet from the buffer. The ED option checks the packet for network buffer parity errors and sends the command, remote address, and source PE number of the packet to the AM option. The data is sent to the AJ and AK options.

From the command, remote address, and source PE number, the AM option yields DRAM control signals and response header phits. The DRAM control signals, along with the remote address, are sent to the AJ, AK, and AE options. Using the command and source PE number phits, the AM option generates a new command and destination and sends this information to the EB option.

The AJ and AK options, along with the AE option, address the DRAM. The AE option supplies the RAS, CAS, and WE to the DRAM. The AJ and AK options supply the internal address. After the DRAM is addressed, the AJ and AK options write the data to the DRAM.

The EB option inputs the command and destination PE number from the AM option and analyzes this information. From the command, the EB option derives the information needed to generate signals that control the packet assembly on the EA option. The EB option also determines whether the destination PE number is virtual or logical, using bit 12 of the destination PE number. If the destination PE number is virtual, the EB option converts it into a logical PE number.

The EB option sends the command, destination PE number, and packet generation control to the EA option. The EB option sends the logical PE number to the routing tag look-up table. The logical PE number is used to address a location in the routing tag look-up table. This location contains the routing tag used to route the request packet through the network. The routing tag is also sent to the EA option.

The EB option instructs the EA option to generate parity for the command and destination PE number and assemble this information into a type 0 response packet. The EA option sends the response packet to the SR option that handles the X dimension of the interconnect network. From this option, the response packet is transferred through the network, back to the source node.

When the response packet reaches the source node, the EC and ED options input the response packet. The ED option buffers the entire response packet, except for the routing tag phit. The SR option continues to send the response packets to the EC and ED options as long as the buffers on the ED option are not full.

The EC option examines the response packet to determine which PE of the node will receive the response; whether the response is a read or write reference; whether there are any parity errors on the command; and whether the response packet arrived at the correct destination.

The EC option notifies the EE option of the response. This causes a decrement of the outstanding reference counter. When the counter equals 0, the EE option sends a signal to the AR option to signify that the BLT reference is complete. This signal prompts the AR option to set the BLT interrupt. The microprocessor polls this interrupt and responds accordingly.

9 Barrier Synchronization

Barrier synchronization provides a low-latency method of synchronizing all or part of the processing elements (PEs) in the CRAY T3D system. This section describes the functions and hardware of the registers, circuitry, and signals used to perform barrier synchronization.

Functional Description

Barrier synchronization circuitry may perform two types of synchronization: barrier and eureka. A barrier is an event initiated by software that indicates when all of the PEs in the partition have reached the same point in a program. A eureka is an event initiated by software that indicates when at least one of the PEs in a partition has reached a specified point in a program.

For example, during barrier synchronization, all of the PEs in a partition may be performing identical matrix multiply computations in parallel, using different segments of data. As each PE completes the computations, it sets a signal that indicates it has reached the barrier. When the last PE completes the computations and sets the barrier signal, all of the PEs in the partition have their barrier signals set. This causes a barrier event, which interrupts all of the PEs in the partition and enables them to begin the next computational task.

During eureka synchronization, all of the PEs in a partition may be performing a distributed data base search on different segments of data. As soon as one PE finds the requested data, it sets a signal that indicates it has reached the eureka. This causes a eureka event, which interrupts all of the PEs in the partition so they may begin the next computational task.

Logical Barrier Synchronization Circuits

The CRAY T3D system has 16 logical barrier synchronization circuits. Each PE in the system has an input to each of the logical barrier synchronization circuits.

Each PE contains two 8-bit registers called barrier register 0 (BSR0) and barrier register 1 (BSR1). Figure 9-1 shows the format of BSR0 and BSR1. Each bit in the barrier registers is an input to one of the 16 logical barrier synchronization circuits.





Figure 9-1. Barrier Register 0 and Barrier Register 1

All of the logical barrier synchronization circuits function identically and are independent. As an example, the following subsections describe the function of the logical barrier synchronization circuit connected to bit 2 of BSR0 when it is used for barrier synchronization and for eureka synchronization.

Barrier Synchronization

A logical barrier synchronization circuit is actually an AND-tree and fan-out-tree circuit. Figure 9-2 shows a barrier synchronization circuit in a simplified CRAY T3D system. This simplified system contains 16 PEs in a $Z = 1$, $Y = 4$, $X = 2$ node shape.

Before the barrier synchronization begins, bit 2 of BSR0 is reset to 0 in all of the PEs. When a microprocessor reaches the barrier, the microprocessor sets bit 2 of its BSR0 to 1. This action sends a 1 to an AND gate in the first layer of the AND tree.



A-12021

Figure 9-2. Simplified Barrier Synchronization Circuit

The first layer of the simplified AND tree contains four AND gates. Each AND gate receives signals from four PEs. For example, one AND gate receives signals from the following four PEs which are located on the same printed circuit board.

- PE 000₁₆
- PE 001₁₆
- PE 010₁₆
- PE 011₁₆

When all of the microprocessors set bit 2 of BSR0 to 1, the output of each of the four AND gates is 1.

The second layer of the simplified AND tree contains two AND gates. Each AND gate receives signals from two of the AND gates in the first layer of the AND tree. When the output of all AND gates in the first layer of the AND tree is 1, the output of both AND gates in the second layer of the AND tree is also 1.

The third layer of the simplified AND tree contains the final AND gate. This AND gate receives signals from both AND gates in the second layer of the AND tree. When the output of both AND gates in the second layer of the AND tree are 1, the output of the final AND gate is 1. The output of the final AND gate sends an input to the fan-out-tree circuit.

The first fan-out block in the simplified fan-out tree receives a 1 from the final AND gate. After creating two copies of the 1, the first fan-out block sends the 1's to the two fan-out blocks in the second layer of the fan-out tree.

The two fan-out blocks in the second layer of the simplified fan-out tree each create two copies of the 1. The two fan-out blocks in the second layer of the fan-out tree then send the 1's to four fan-out blocks in the third layer of the fan-out tree.

The four fan-out blocks in the third layer of the simplified fan-out tree each create two copies of the 1. The fan-out blocks in the third layer of the fan-out tree then send the 1's to the sixteen PEs. This signals all of the PEs in the system that all of the microprocessors have reached the barrier and that the microprocessors may continue with other program instructions.

The microprocessor monitors the output of the fan-out circuitry using one of two methods: continuously looping on the barrier bit or setting the barrier interrupt.

In the first method, after the microprocessor sets bit 2 of BSR0 to 1, the microprocessor enters a loop that continuously checks the value of bit 2 of BSR0. After receiving a 1 from the fan-out circuitry, the support circuitry in the PE resets bit 2 of BSR0 to 0. Because the PE constantly checks the value of bit 2 of BSR0, the PE resumes its computational task as soon as bit 2 of BSR0 is reset to 0.

In the second method, after the microprocessor sets bit 2 of BSR0 to 1, the microprocessor enables the barrier hardware interrupt. The microprocessor may then issue program instructions that are not associated with the barrier. After receiving a 1 from the fan-out circuitry, the support circuitry in the PE resets bit 2 of BSR0 to 0 and sets the barrier interrupt. The barrier interrupt indicates to the microprocessor that all of the microprocessors have reached the barrier.

The microprocessor enables the barrier hardware interrupt using the hardware interrupt enable register (HIER) in the microprocessor and using bit 3 of the system control register. More information on the system control register is provided later in this section and in Section 10, “Control and Status.”

When the PE is looping on bit 2 of BSR0, the PE terminates its current task as soon as the support circuitry resets bit 2 of BSR0 to 0. When the barrier interrupt is enabled, the support circuitry sets the barrier interrupt when the support circuitry resets bit 2 of BSR0 to 0.

Logical Partitions

Not all of the PEs in the system need to be part of a barrier or eureka synchronization process. Each PE also contains a barrier synchronization mask and interrupt (BSMI) register, which is used to enable or disable a logical barrier synchronization circuit for a PE. The BSMI register contains a mask that indicates which bits of BSR0 and BSR1 are enabled for the PE.

Figure 9-4 shows the format of the BSMI. If a bit of the BSMI register is set to 1, the corresponding bit of BSR0 or BSR1 is enabled. If a bit of the BSMI register is set to 0, the corresponding bit of BSR0 or BSR1 is disabled.




Figure 9-4. BSMI Register Format

If a barrier register bit is disabled, the support circuitry automatically sends a 1 to the input of the corresponding barrier synchronization circuit. This prevents the disabled barrier bit from interrupting the function of the barrier synchronization circuit for the other PEs in the partition.

Software may use the BSMI register to allow only a subset of PEs to use one of the logical barrier synchronization circuits. For example, software may set bit 2 of the BSMI register to 1 in only four PEs. In this case, only these four PEs may use the logical barrier synchronization circuit connected to bit 2 of BSR0. This creates a logical barrier partition among the four PEs.

Table 9-1 shows the effect of a bit from the BSFR and the BSMI register on a bit in BSR0 or BSR1. When the contents of BSR0 are read, the value returned represents bits 0 through 7 of BSR0. When the contents of BSR1 are read, the value returned represents bits 8 through 15 of BSR1 and bits 0 through 7 of BSR0.

Table 9-1. Effect of BSMI register and BSFR Bits

BSMI Bit Value	BSFR Bit Value	Writing to BSR0 or BSR1	Reading from BSR0 or BSR1	Synchronization Type
0	0	No effect	Returns a 1	Disabled
0	1	No effect	Returns a 1	Disabled
1	0	Writing 1 indicates the microprocessor has reached the barrier. Writing 0 has no effect.	Returns a 1 if waiting for barrier to complete. Returns a 0 when barrier is complete.	Barrier
1	1	Writing 1 indicates the microprocessor is ready for eureka synchronization. Writing 0 indicates the microprocessor has completed the eureka process.	Returns a 1 if waiting for eureka to occur. Returns a 0 when eureka occurs.	Eureka

Physical Barrier Synchronization Circuits

Although each of the 16 bits in BSR0 and BSR1 represents an input to a logical barrier synchronization circuit, the CRAY T3D system does not contain 16 physical barrier synchronization circuits. Instead, the system contains 4 physical barrier synchronization circuits, and the 16 bits of BSR0 and BSR1 are time multiplexed.

Table 9-2 shows the input to each physical barrier synchronization circuit during each clock period (CP). Four CPs are required for the physical barrier synchronization circuits to receive input from all 16 bits in BSR0 and BSR1.

Table 9-2. Physical Barrier Synchronization Circuit Inputs


Input Type	First CP	Second CP	Third CP	Fourth CP
Input to physical barrier synchronization circuit 0	Bit 0 of BSR0	Bit 4 of BSR0	Bit 8 of BSR1	Bit 12 of BSR1
Input to physical barrier synchronization circuit 1	Bit 1 of BSR0	Bit 5 of BSR0	Bit 9 of BSR1	Bit 13 of BSR1
Input to physical barrier synchronization circuit 2	Bit 2 of BSR0	Bit 6 of BSR0	Bit 10 of BSR1	Bit 14 of BSR1
Input to physical barrier synchronization circuit 3	Bit 3 of BSR0	Bit 7 of BSR0	Bit 11 of BSR1	Bit 15 of BSR1

Physical Partitions

In a physical barrier synchronization circuit, each AND gate in the AND tree is paired with a fan-out block in the fan-out tree. An AND gate and fan-out block pair is called a bypass point. For example, Figure 9-5 shows the bypass points in a simplified barrier synchronization circuit.

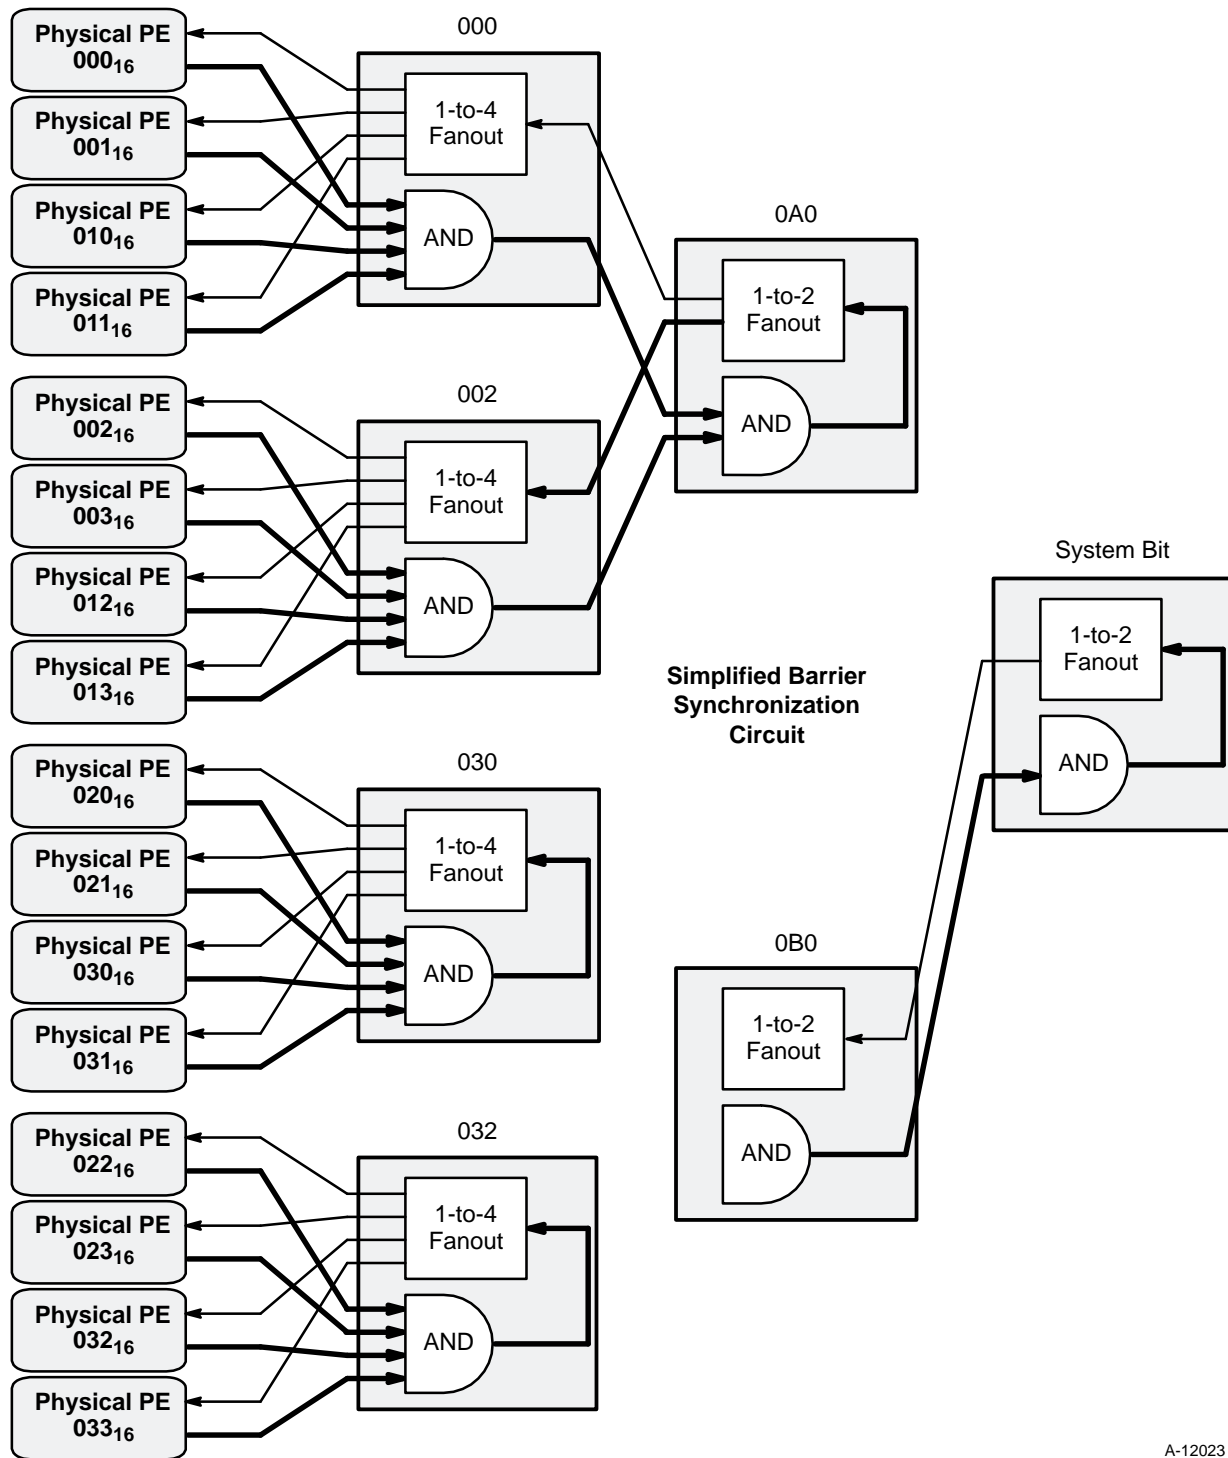
Software running in the host system can redirect the output of an AND gate in a bypass point so that the output of the AND gate connects to the fan-out block in the bypass point. Another memory-mapped register, called the network mode register (NODE_CSR), controls which AND gates in the AND tree have their outputs redirected to the corresponding fan-out block in a bypass point. More information on the NODE_CSR is provided later in this section and in Section 10, “Control and Status.”

Figure 9-6 shows the same barrier synchronization circuit as Figure 9-5; however, the output of the AND gate in some of the bypass points is redirected to the fan-out block in the bypass point. This partitions the physical barrier synchronization circuit into three smaller circuits.



A-12022

Figure 9-5. Bypass Points



A-12023

Figure 9-6. Partitioned Barrier Synchronization Circuit

The first smaller barrier synchronization circuit contains a two-level AND tree and a two-level fan-out tree. This circuit is created by redirecting the output of the AND gate in bypass point 0A0 to the fan-out block in bypass

point 0A0. This smaller circuit operates identically to the barrier synchronization circuit in Figure 9-5; however, this circuit receives an input from and sends an output to only 8 of the 16 PEs in the system.

The second small barrier synchronization circuit contains a one-level AND tree and fan-out tree. This circuit is created by redirecting the output of the AND gate in bypass point 030 to the fan-out block in bypass point 030. This smaller circuit operates identically to the barrier synchronization circuit in Figure 9-5; however, this circuit receives input from and sends an output to only 4 of the 16 PEs in the system.

The third small barrier synchronization circuit also contains a one-level AND tree and fan-out tree. This circuit is created by redirecting the output of the AND gate in bypass point 032 to the fan-out block in bypass point 032. This smaller circuit operates identically to the barrier synchronization circuit in Figure 9-5; however, this circuit receives input from and sends an output to only 4 of the 16 PEs in the system.

In an actual CRAY T3D system, each physical barrier synchronization circuit connects to all of the PEs in the system. The signals from the PEs in the processing element nodes travel through several layers of bypass points before arriving at the system bit. The signals from the PEs in the I/O gateways travel through one bypass point and then to the system bit. The signals from the PEs in the spare processing element nodes travel directly to the final system bit.


For an example of system bypass points, in a 512-PE CRAY T3D system, each physical barrier synchronization circuit contains a six-level AND tree and six-level fan-out tree (refer to Table 9-3). The first level contains bypass points that connect to the four PEs on a printed circuit board (PCB). The remaining levels contain bypass points that connect to other bypass points through the wiremat. More information on the actual physical barrier synchronization circuits is provided in “Register Mapping” later in this section.

Table 9-3. Levels in a 512-PE Physical Barrier Synchronization Circuit

Level	Name	PE Nodes		I/O Gateways	
		Bypass Points	Each Bypass Point Receives an Input From:	Bypass Points	Each Bypass Point Receives an Input From:
1	PCB bypass points	128	The four PEs on a PCB	8	The two PEs in an I/O gateway
2	Wiremat bypass points	32	Four PCB bypass points	NA	NA
3	Wiremat bypass points	8	Four level 2 bypass points	NA	NA
4	Wiremat bypass points	4	Two level 3 bypass points	NA	NA
5	Box bit	1	Four level 4 bypass points	NA	NA
6	System bit	The system bit receives an input from the box bit, the eight I/O gateway PCB bypass points, and the eight spare PEs. (Although the two spare PE PCBs each contain a bypass point, they may not be used.)			

The system administrator can redirect the output of the AND gate to the fan-out block in any of the bypass points listed in Table 9-3. This creates several types of barrier partitions. Figure 9-7 shows how the PE nodes in a 512-PE CRAY T3D system are partitioned at each level of a physical barrier synchronization circuit. The following paragraphs describe each level of the physical barrier synchronization circuit.

Each bypass point in level 1 of the physical barrier synchronization circuit connects to the 4 PE on a PE PCB. If all 128 of the PCB bypass points have the output of the AND gate redirected to the fan-out block, the PEs in the system are divided into 128 4-PE groups (refer to “1” in Figure 9-7).



A-11701

Figure 9-7. PE Node Bypass Point Partitions in a 512-PE CRAY T3D System

Each bypass point in level 2 of the physical barrier synchronization circuit connects to 4 PCB level 1 bypass points. If all 32 of these level 2 bypass points have the output of the AND gate redirected to the fan-out block, the PEs in the system are divided into 32 16-PE groups (refer to “2” in Figure 9-7). More information on the bypass point numbering is provided in “Register Mapping” later in this section.

Each bypass point in level 3 of a physical barrier synchronization circuit connects to four level 2 bypass points. If all eight of these bypass points have the output of the AND gate redirected to the fan-out block, the PEs in the system are divided into eight 64-PE groups (refer to “3” in Figure 9-7).

Each bypass point in level 4 of a physical barrier synchronization circuit connects to two level 3 bypass points. If all four of these bypass points have the output of the AND gate redirected to the fan-out block, the PEs in the system are divided into four 128-PE groups (refer to “4” in Figure 9-7).

The box bit bypass point in a physical barrier synchronization circuit connects to the four level 4 bypass points. If the box bit bypass point has the output of the AND gate redirected to the fan-out block, the PEs in the system are in one 512-PE group (refer to “5” in Figure 9-7).

When a bypass point has the output of the AND gate redirected to the fan-out block, the PEs in the barrier partition can still use all 16 bits in BSR0 and BSR1. However, because there are only 4 physical barrier synchronization circuits, creating a barrier partition affects 4 of the 16 bits in BSR0 and BSR1.

For example, if a level 4 bypass point in physical barrier synchronization circuit 2 has the output of the AND gate redirected to the fan-out block, the barrier partition contains 128-PEs. Because of this barrier partition, bits 2, 6, 10, and 14 of BSR0 and BSR1 in each of these 128 PEs affects only the 128 PEs in the barrier partition. In the 128 PEs, these bits cannot be used for system-level barrier or eureka synchronization.

By redirecting the output of AND gates in different level bypass points, the system administrator may divide a physical barrier synchronization circuit into a combination of the previously described barrier partitions. Also, as described in Section 6, “Addressing,” the system administrator may divide the PE nodes in the system into different user partitions. The number and shape of the PEs in a barrier partition may not exactly match the number and shape of the PEs in a user partition.

For example, in a 512-PE CRAY T3D system, when the output of an AND gate in a level 3 bypass point is redirected to the fan-out block, the barrier partition contains 64 PEs (32 nodes). These 64 PEs may actually be two 32-PE (16 node) user partitions. In this case, the operating system must use the BSMI register to disable some of the bits in BSR0 and BSR1 for half of the PEs and enable these bits for the remaining PEs in the barrier partition.

Timing

When a physical barrier synchronization circuit is divided into barrier partitions, the time needed for a bit to propagate through each physical barrier synchronization circuit is not consistent. A memory-mapped register called the barrier timing (BAR_TMGM) register controls the timing of signals for each physical barrier synchronization circuit. More information on the BAR_TMGM register is provided in the following subsection.

Register Mapping

The following subsections describe the addressing and bit assignments for the memory-mapped registers (hereafter referred to as registers) used for barrier and eureka synchronization in the CRAY T3D system. Each subsection also provides a brief summary of the function of the register.

Table 9-4 is a summary of the barrier synchronization registers and their names. Table 9-4 also lists the partial physical address of each register as it appears on the address pins of the microprocessor.

Table 9-4. Barrier Synchronization Registers

Address	Register Name	Direction	Description
10108000 ₁₆	BSR0	Read or write	Barrier synchronization register 0
1090C000 ₁₆	BSR1	Read or write	Barrier synchronization register 1
1041C000 ₁₆	BSFR	Write	Barrier synchronization function register
10418000 ₁₆	BSMI	Read or write	Barrier synchronization mask and interrupt
10404000 ₁₆	BAR_TMG	Write	Barrier timing and refresh register
106E0000 ₁₆	NODE_CSR	Write	Network mode register

NOTE: Because of multiplexed data paths, when one PE in a processing element node is modifying the contents of the NODE_CSR register, the other PE in the node must not attempt to modify any of the shared registers at the same time. The shared registers include the X_WHOAMI, LPE_XLATE, ROUTE_LO, ROUTE_HI, NET_ENA, NET_PFM, NODE_CSR, and BLT registers.

Because the virtual address is defined by software, the addresses for each of the registers are given according to the partial physical address as it appears on the pins of the microprocessor.

Bit 3 (enable BSR0 interrupt bit) of the system control register, the BSFR, and the BSMR are independent and may be written to in any order. When a PE is reset, the value of bit 3 of the SCR is 0, and bits 0 through 15 in BSR0 and BSR1 are set to 1's.

Barrier Synchronization Register 0**Address 10108000₁₆**

Barrier synchronization register 0 (BSR0) is an 8-bit, readable and writable, general access register. BSR0 contains the 8 least significant barrier bits for a PE. When read from, the value of BSR0 represents the value of bits 0 through 7 of BSR0. The remaining bits are not valid.

Figure 9-8 shows the bit assignments for the BSR0 address as they appear on the address pins of the microprocessor.

HEX	1	0				1				0				8				0				0				0			
Binary	2 ²⁸	2 ²⁷	2 ²⁶	2 ²⁵	2 ²⁴	2 ²³	2 ²²	2 ²¹	2 ²⁰	2 ¹⁹	2 ¹⁸	2 ¹⁷	2 ¹⁶	2 ¹⁵	2 ¹⁴	2 ¹³	2 ¹²	2 ¹¹	2 ¹⁰	2 ⁹	2 ⁸	2 ⁷	2 ⁶	2 ⁵	2 ⁴	2 ³	2 ²	2 ¹	2 ⁰
	1	x	x	x	x	x	x	0	1	0	0	0	0	1	0	x	x	x	x	x	x	x	x	x	x	x	x	x	x

NOTE: x = Don't Care

Figure 9-8. Barrier Synchronization Register 0 Address Bit Format

Table 9-5 shows the bit format of BSR0 and describes each bit of the register.

Table 9-5. BSR0 Format

Bits	Name
0	Barrier bit 0
1	Barrier bit 1
2	Barrier bit 2
3	Barrier bit 3
4	Barrier bit 4
5	Barrier bit 5
6	Barrier bit 6
7	Barrier bit 7
63 – 8	These bits are not used.

Barrier Synchronization Register 1
Address 1090C000₁₆

Barrier synchronization register 1 (BSR1) is an 8-bit, readable and writable, general access register. BSR1 contains the 8 most significant barrier bits for a PE. When read from, the value of BSR1 represents the value of bits 8 through 15 of BSR1 and 0 through 7 of BSR0. The remaining bits are not valid.

Figure 9-9 shows the bit assignments for the BSR1 address as they appear on the address pins of the microprocessor.

HEX	1	0				9				0				C				0				0				0			
Binary	2 ²⁸	2 ²⁷	2 ²⁶	2 ²⁵	2 ²⁴	2 ²³	2 ²²	2 ²¹	2 ²⁰	2 ¹⁹	2 ¹⁸	2 ¹⁷	2 ¹⁶	2 ¹⁵	2 ¹⁴	2 ¹³	2 ¹²	2 ¹¹	2 ¹⁰	2 ⁹	2 ⁸	2 ⁷	2 ⁶	2 ⁵	2 ⁴	2 ³	2 ²	2 ¹	2 ⁰
	1	x	x	x	x	1	x	0	1	0	0	0	0	1	1	x	x	x	x	x	x	x	x	x	x	x	x	x	x

NOTE: x = Don't Care

Figure 9-9. Barrier Synchronization Register 1 Address Bit Format

Table 9-6 shows the bit format of BSR1 and describes each bit of the register.

Table 9-6. BSR1 Format

Bits	Name
7 – 0	These bits are not used; however, when BSR1 is read, these bits contain the value of bits 7 through 0 of BSR0.
8	Barrier bit 8
9	Barrier bit 9
10	Barrier bit 10
11	Barrier bit 11
12	Barrier bit 12
13	Barrier bit 13
14	Barrier bit 14
15	Barrier bit 15
63 – 16	These bits are not used.

Barrier Synchronization Function Register**Address 1041C000₁₆**

The barrier synchronization function register (BSFR) is a 16-bit, write-only, system privileged register. The BSFR contains a 16-bit mask that indicates which bits of BSR0 and BSR1 are used for barrier synchronization and which bits are used for eureka synchronization.

Figure 9-10 shows the bit assignments for the BSR1 address as they appear on the address pins of the microprocessor.

HEX	1	0				4				1				C				0				0				0			
Binary	2 ²⁸	2 ²⁷	2 ²⁶	2 ²⁵	2 ²⁴	2 ²³	2 ²²	2 ²¹	2 ²⁰	2 ¹⁹	2 ¹⁸	2 ¹⁷	2 ¹⁶	2 ¹⁵	2 ¹⁴	2 ¹³	2 ¹²	2 ¹¹	2 ¹⁰	2 ⁹	2 ⁸	2 ⁷	2 ⁶	2 ⁵	2 ⁴	2 ³	2 ²	2 ¹	2 ⁰
	1	x	x	x	x	x	1	0	0	0	0	0	1	1	1	x	x	x	x	x	x	x	x	x	x	x	x	x	x

NOTE: x = Don't Care

Figure 9-10. Barrier Synchronization Function Register Address Bit Format

Table 9-7 shows the bit format of the BSFR and describes each bit of the register.

Table 9-7. BSFR Format

Bits	Description
7 – 0	These bits indicate which bits of BSR0 are used for eureka or barrier synchronization. For example, when bit 2 of the BSFR is set to 1, bit 2 of BSR0 is used for eureka synchronization. When bit 2 of the BSFR is set to 0, bit 2 of BSR0 is used for barrier synchronization.
15 – 8	These bits indicate which bits of BSR1 are used for eureka or barrier synchronization. For example, when bit 12 of the BSFR is set to 1, bit 12 of BSR1 is used for eureka synchronization. When bit 12 of the BSFR is set to 0, bit 12 of BSR1 is used for barrier synchronization.
63 – 16	These bits are not used.

Barrier Synchronization Mask and Interrupt Register

Address 10418000₁₆

The barrier synchronization mask and interrupt (BSMI) register is a 16-bit, readable and writable, system privileged register. When written to, the BSMI register indicates which bits in BSR0 and BSR1 are enabled for use by the PE. When read from, the BSMI register provides the current state of the barrier interrupts from BSR0 and BSR1 and then clears the interrupts.

Figure 9-11 shows the bit assignments for the BSMI register address as they appear on the address pins of the microprocessor.

HEX	1	0				4				1				8				0				0				0			
Binary	2 ²⁸	2 ²⁷	2 ²⁶	2 ²⁵	2 ²⁴	2 ²³	2 ²²	2 ²¹	2 ²⁰	2 ¹⁹	2 ¹⁸	2 ¹⁷	2 ¹⁶	2 ¹⁵	2 ¹⁴	2 ¹³	2 ¹²	2 ¹¹	2 ¹⁰	2 ⁹	2 ⁸	2 ⁷	2 ⁶	2 ⁵	2 ⁴	2 ³	2 ²	2 ¹	2 ⁰
	1	x	x	x	x	x	1	0	0	0	0	0	1	1	0	x	x	x	x	x	x	x	x	x	x	x	x	x	x

NOTE: x = Don't Care

Figure 9-11. Barrier Synchronization Mask and Interrupt Register Address Bit Format

The BSMI register has a different bit format when written to than it does when read from. Table 9-8 shows the bit format of the BSMI register when it is written to and describes each bit of the register. The BSMI does not mask a write to the BSFR. All 16 barrier bits for a barrier synchronization memory-mapped register are always transferred from the microprocessor to a register.

Table 9-8. BSMI Register Write Format

Bits	Name
7 – 0	These bits indicate which bits of BSR0 are enabled for use by the PE. For example, when bit 2 of the BSMI register is set to 1, bit 2 of BSR0 is enabled for use by the PE. When bit 2 of the BSMI register is set to 0, bit 2 of BSR0 is disabled and cannot be used by the PE.
15 – 8	These bits indicate which bits of BSR1 are enabled for use by the PE. For example, when bit 12 of the BSMI register is set to 1, bit 12 of BSR1 is enabled for use by the PE. When bit 12 of the BSMI register is set to 0, bit 12 of BSR1 is disabled and cannot be used by the PE.
63 – 16	These bits are not used.

Table 9-9 shows the bit format of the BSMI register when it is read from and describes each bit of the register.

Table 9-9. BSMI Register Read Format

Bits	Description
13 – 0	These bits are not valid.
14	This bit reflects the current state of the barrier interrupt associated with bits 7 through 0 of BSR0.
15	This bit reflects the current state of the barrier interrupt associated with bits 15 through 8 of BSR1.
63 – 16	These bits are not valid.

The support circuitry sets the barrier hardware interrupt when a barrier or eureka synchronization operation completes. The value of bit 3 (enable BSR0 interrupt bit) in the system control register (SCR) controls which barrier bits trigger the barrier hardware interrupt.

When set to 1, bit 3 of the SCR signals the support circuitry to set the barrier hardware interrupt if a barrier associated with BSR0 or BSR1 occurs. When set to 0, bit 3 of the SCR signals the support circuitry to set the barrier hardware interrupt only if a barrier associated with BSR1 occurs. In this case, the barrier interrupt triggered by the barrier bits in BSR0 is disabled.

After the support circuitry sets the barrier hardware interrupt, the microprocessor must read the BSMI register to determine whether the interrupt was associated with BSR0 or BSR1. When the microprocessor reads the value of the BSMI register, the support circuitry clears both the interrupt associated with BSR0 and the interrupt associated with BSR1.

If a barrier interrupt occurs while the microprocessor is reading the BSMI register, the interrupt still occurs and is not cleared. The microprocessor must then read the value of the BSMI register again to determine whether the interrupt was associated with BSR0 or BSR1 and to clear the interrupt.

More information on the system control register and the barrier hardware interrupt is provided in Section 10, “Control and Status.”

Barrier Timing and Refresh Register

Address 10404000₁₆

The barrier timing and refresh (BAR_TMGM) register is a 12-bit, write-only, system privileged register. The BAR_TMGM register controls the timing of each physical barrier synchronization circuit.

Figure 9-12 shows the bit assignments for the BAR_TMGM register address as they appear on the address pins of the microprocessor.

HEX	1	0				4				0				4				0				0							
Binary	2 ²⁸	2 ²⁷	2 ²⁶	2 ²⁵	2 ²⁴	2 ²³	2 ²²	2 ²¹	2 ²⁰	2 ¹⁹	2 ¹⁸	2 ¹⁷	2 ¹⁶	2 ¹⁵	2 ¹⁴	2 ¹³	2 ¹²	2 ¹¹	2 ¹⁰	2 ⁹	2 ⁸	2 ⁷	2 ⁶	2 ⁵	2 ⁴	2 ³	2 ²	2 ¹	2 ⁰
	1	x	x	x	x	x	1	0	0	0	0	0	0	0	1	x	x	x	x	x	x	x	x	x	x	x	x	x	x

NOTE: x = Don't Care

Figure 9-12. Barrier Timing and Refresh Register Address Bit Format

Table 9-10 shows the bit format of the BAR_TMGM register and the following paragraphs describe the function of bits 0 through 7 of the BAR_TMGM register. More information on bits 8 through 11 of the BAR_TMGM register is provided in Section 10, "Control and Status."

Table 9-10. BAR_TMGM Bit Format

Bits	Description
1 – 0	These bits control the timing for physical barrier synchronization circuit 0.
3 – 2	These bits control the timing for physical barrier synchronization circuit 1.
5 – 4	These bits control the timing for physical barrier synchronization circuit 2.
7 – 6	These bits control the timing for physical barrier synchronization circuit 3.
11 – 8	These bits control the DRAM refresh timing.
63 – 12	These bits are not used.

Because a physical barrier synchronization circuit may be divided into partitions, the total time for a bit to propagate through the circuit is not constant. This timing inconsistency may cause a bit to be in the wrong position when BSR0 or BSR1 is read. For example, if the timing is not set up correctly for physical barrier synchronization circuit 0, a bit that was originally written to the 0 bit location in BSR0 may appear in the 4 bit location of BSR0 when BSR0 is read.

The following procedure sets the BAR_TMG timing value for physical barrier synchronization circuit 0 in the PEs of a barrier partition. Before performing the procedure, disable the barrier hardware interrupt to the microprocessors.

- 1 . Write a value of 1111_{16} to the BSMI register in all of the PEs to enable barrier bits 0, 4, 8, and 12.
- 2 . Write a value of 1111_{16} to the BSFR in all of the PEs to set barrier bits 0, 4, 8, and 12 to eureka mode.
- 3 . Write a value of 1111_{16} to BSR0 and BSR1 in all of the PEs to start the eureka process.
- 4 . In one of the PEs, write a value of 1110_{16} to BSR0 to indicate that bit 0 has completed a eureka.
- 5 . In each of the PEs, read the value of BSR1, which contains the value of all 16 barrier bits. Also apply a software mask so that the only bits read are 0, 4, 8, and 12 and the remaining bits are set to 0. The value read from BSR1 may be 1110_{16} , 1101_{16} , 1011_{16} , or 0111_{16} . If the value is 1110_{16} , the timing is set up correctly. If the value is not 1110_{16} , increment bits 0 and 1 of the BAR_TMG register by 1 and read the value of BSR1 again.

If the value read from BSR1 is now 1110_{16} , the timing is set up correctly. If the value is not 1110_{16} , increment bits 0 and 1 of the BAR_TMG register by 1 and read the value of BSR1 again. This process must be repeated until the value read from BSR1 is 1110_{16} but should not need to be performed more than four times.

This procedure may be used to set the timing for any of the physical barrier synchronization circuits in a barrier partition. Table 9-11 lists the barrier bits affected and the write pattern for each physical barrier synchronization circuit.

Table 9-11. Timing Procedure Values

Physical Barrier Synchronization Circuit	Barrier Bits	Write Pattern
0	12, 8, 4, and 0	1110 ₁₆
1	13, 9, 5, and 1	2220 ₁₆
2	14, 10, 6, and 2	4440 ₁₆
3	15, 11, 7, and 3	8880 ₁₆

Network Mode Register
Address 106E0000₁₆

The network mode (NODE_CSR) register is a 14-bit, write only, system privileged register. The NODE_CSR enables or disables several types of error checking in the network interface and the BLT. The NODE_CSR is also used to partition the physical barrier synchronization circuits.

Figure 9-13 shows the bit assignments for the NODE_CSR address as they appear on the address pins of the microprocessor.

HEX	1	0				6				E				0				0				0				0			
Binary	2 ²⁸	2 ²⁷	2 ²⁶	2 ²⁵	2 ²⁴	2 ²³	2 ²²	2 ²¹	2 ²⁰	2 ¹⁹	2 ¹⁸	2 ¹⁷	2 ¹⁶	2 ¹⁵	2 ¹⁴	2 ¹³	2 ¹²	2 ¹¹	2 ¹⁰	2 ⁹	2 ⁸	2 ⁷	2 ⁶	2 ⁵	2 ⁴	2 ³	2 ²	2 ¹	2 ⁰
	1	x	x	x	x	x	1	1	0	1	1	1	0	x	x	x	x	x	x	x	x	x	x	x	x	x	x	x	x

NOTE: x = Don't Care

Figure 9-13. Network Mode Register Address Bit Format

Either PE in a node can write to the NODE_CSR. Software must determine which PE in the node sets the parameters for the NODE_CSR.

Table 9-12 shows the bit format of the NODE_CSR register, and the following subsections describe bits 11 through 13 of the register. For more information on the NODE_CSR register, refer to Section 10, “Control and Status.”

Table 9-12. NODE_CSR Bit Format

Bits	Description
10 – 0	These bits are used for control and status.
11	Barrier synchronization partition bit (PCB bypass n)
12	Barrier synchronization partition bit (PCB bypass n+1)
13	Barrier synchronization partition bit (wiremat bypass)
63 – 14	These bits are not used.

Barrier Synchronization Partition Bit 11

Bit 11 of the NODE_CSR controls a level 1 bypass point in physical barrier synchronization circuit 0 or 2. The physical circuit controlled depends on which node the NODE_CSR register is located in and on which PCB the node resides. Bit 11 of the NODE_CSR also controls a different type of bypass in a PE node or spare PE node than it does in an I/O gateway.

The output signals for each physical barrier synchronization circuit are located on the same side of a module even though the Y and Z sides of PCB 1 (b) are opposite the Y and Z sides of PCB 0 (a) when placed on the module. For example, the output signals for physical barrier synchronization circuit 0 from both PCBs are located on the Y side of a module. (For more information on the PCBs and modules, refer again to “Physical PE Number Register” in Section 6, “Addressing.”)

Because the signals are always on the same side of a module, the node on a PCB (node 0 or node 1) that controls a level 1 bypass point changes depending on whether the node is on PCB 0 or PCB 1. Table 9-13 lists the different values of node numbers and shows the corresponding physical barrier synchronization circuit that is controlled by bit 11 of the NODE_CSR in a PE node.

Table 9-13. Bit 11 of NODE_CSR in a Processing Element Node

Bits 7 – 4 of the Processing Element Node Number (from the P_WHOAMI register)	Bit 11 of the NODE_CSR Control Signals for:
xx00	Physical barrier synchronization circuit 0
xx01	Physical barrier synchronization circuit 2
xx10	Physical barrier synchronization circuit 2
xx11	Physical barrier synchronization circuit 0

When bit 11 of the NODE_CSR is set to 0, the output of an AND gate in a level 1 bypass point is redirected to the fan-out block in the bypass point. When bit 11 of the NODE_CSR is set to 1, the output of an AND gate in a level 1 bypass point is redirected to a second level bypass point.

Each level 1 bypass point for the processing element nodes receives four inputs. Each input is from one of the four PEs on a PCB.

Figure 9-14 shows a four-PE barrier partition in physical barrier synchronization circuit 0. To create this four-PE barrier partition, bit 11 of the NODE_CSR must be set to 0 in the node with a physical node number of 402_{16} (which corresponds to physical PE 402_{16} or 403_{16}). If bit 11 of the NODE_CSR in this node is set to 1, the output of the AND gate shown in Figure 9-14 is directed to a level 2 bypass point.




Figure 9-14. Four-PE Barrier Partition in Physical Barrier Synchronization Circuit 0

Every PE PCB contains four level 1 bypass points (one for each physical barrier synchronization circuit). Each bypass point is given a name that corresponds to the physical node that controls the bypass point.

For example, Figure 9-14 shows the bypass point 402. This indicates that the bypass is on the PCB that contains physical PEs 402, 403, 412, and 413. The bypass is controlled by setting bit 11 of the `NODE_CSR` in physical PE 402 or physical PE 403 to a 0.

Like in the PE nodes, bit 11 of the `NODE_CSR` in a spare PE node controls a level 1 bypass point in physical barrier synchronization circuit 0 or 2. Also like the PE nodes, the output signals for a physical barrier synchronization circuit are on the same side of a spare PE module. Table 9-14 lists the different values of node numbers and shows the corresponding physical barrier synchronization circuit that is controlled by bit 11 of the `NODE_CSR` in a spare PE node.

Table 9-14. Bit 11 of NODE_CSR in a Spare Processing Element Node

Bits 7 – 4 of the Spare Processing Element Node Number (from the P_WHOAMI register)	Bit 11 of the NODE_CSR Control Signals for:
xx00	Physical barrier synchronization circuit 0
xx01	Physical barrier synchronization circuit 2
xx10	Physical barrier synchronization circuit 2
xx11	Physical barrier synchronization circuit 0

Figure 9-15 shows a spare-PE barrier partition in physical barrier synchronization circuit 0. To create this spare-PE barrier partition, bit 11 of the NODE_CSR must be set to 0 in the node with a physical node number $90C_{16}$. If bit 11 of the NODE_CSR in this node is set to 1, the output of the AND gate shown in Figure 9-15 is directed to the system bit.

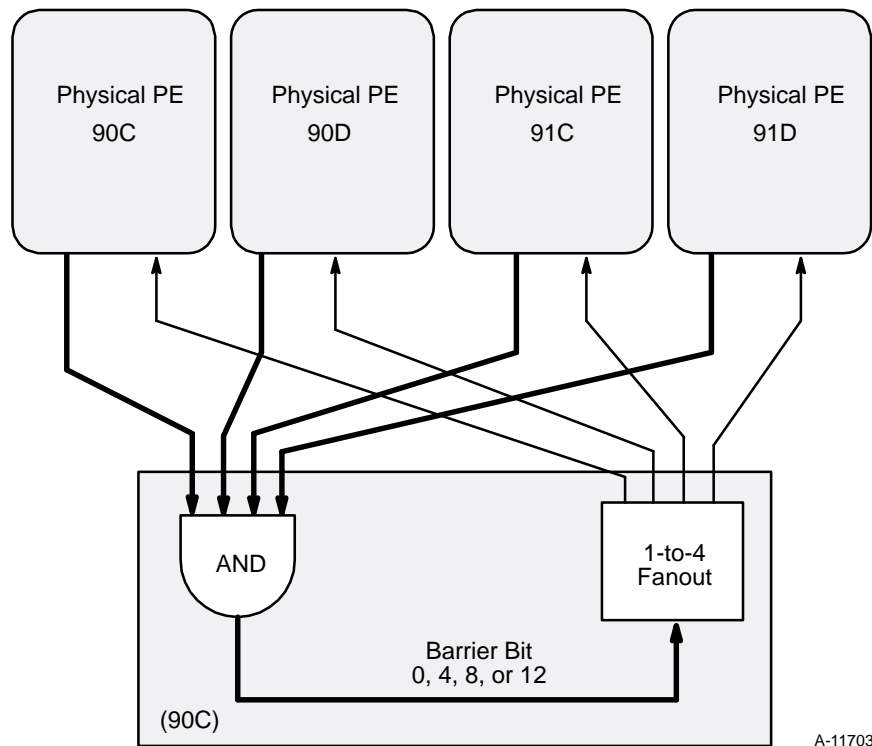


Figure 9-15. Spare PE Node Barrier Partition in Physical Barrier Synchronization Circuit 0


Like in the PE nodes, bit 11 of the NODE_CSR controls a level 1 bypass point in physical barrier synchronization circuit 0 or 2; however, the level 1 bypass point in an I/O gateway only receives input from two PEs (one in the input node and one in the output node).

Also like the PE nodes, the output signals for a physical barrier synchronization circuit are always on the same side of an I/O module. Table 9-15 lists the different values of node numbers and shows the corresponding physical barrier synchronization circuit that is controlled by bit 11 of the NODE_CSR in an input node or output node of an I/O gateway.

Table 9-15. Bit 11 of NODE_CSR in an I/O Gateway

Bits 7 – 4 of the I/O Node Number (from the P_WHOAMI register)	Bit 11 of the NODE_CSR Control Signals for:
xx00 (Input)	Physical barrier synchronization circuit 0
xx01 (Output)	Physical barrier synchronization circuit 2
xx10 (Input)	Physical barrier synchronization circuit 2
xx11 (Output)	Physical barrier synchronization circuit 0

Figure 9-16 shows an I/O gateway barrier partition in physical barrier synchronization circuit 0. To create this partition, bit 11 of the NODE_CSR must be set to 0 in the output physical node C30₁₆. If bit 11 of the NODE_CSR in this node is set to 1, the output of the AND gate, shown in Figure 9-16, is directed to the final system-level AND gate (system bit).



A-11704

Figure 9-16. I/O Gateway Barrier Partition in Physical Barrier Synchronization Circuit 0

Barrier Synchronization Partition

Bit 12

Bit 12 of the NODE_CSR operates identically to bit 11 of the NODE_CSR; however, bit 12 of the NODE_CSR controls the level 1 bypass for physical barrier synchronization circuit 1 or 3. Table 9-16 lists the different values of node numbers and shows the corresponding physical barrier synchronization circuit that is controlled by bit 12 of the NODE_CSR in a processing element node or a spare processing element node.

Table 9-16. Bit 12 of NODE_CSR in a PE Node or Spare PE Node

Bits 7 – 4 of the Processing Element Node Number (from the P_WHOAMI register)	Bit 12 of the NODE_CSR Control Signals for:
xx00	Physical barrier synchronization circuit 1
xx01	Physical barrier synchronization circuit 3
xx10	Physical barrier synchronization circuit 3
xx11	Physical barrier synchronization circuit 1

Table 9-17 lists the different values of node numbers and shows the corresponding physical barrier synchronization circuit that is controlled by bit 12 of the NODE_CSR in an input or output node of an I/O gateway.

Table 9-17. Bit 12 of NODE_CSR in an I/O Gateway


Bits 7 – 4 of the I/O Node Number (from the P_WHOAMI register)	Bit 12 of the NODE_CSR Control Signals for:
xx00 (Input)	Physical barrier synchronization circuit 1
xx01 (Output)	Physical barrier synchronization circuit 3
xx10 (Input)	Physical barrier synchronization circuit 3
xx11 (Output)	Physical barrier synchronization circuit 1

Barrier Synchronization Partition
Bit 13

Bit 13 of the NODE_CSR controls one bypass point in one of the physical barrier synchronization circuits. This bit is only used in specific PE nodes and has no effect in an I/O gateway or a spare PE node.

When set to 0, bit 13 redirects the output of the AND gate in the bypass point to the fan-out block in the bypass point. When set to 1, bit 13 directs the output of the AND gate in the bypass point to the next level bypass point.

As an example, Figure 9-17 shows a few of the bypass points for physical barrier synchronization circuit 0 in a CRAY T3D system. When bit 13 of the NODE_CSR is set to 0 in physical node 400₁₆, the output of the AND gate in bypass point 2A is redirected to the fan-out block in bypass point 2A. This creates a 64-PE barrier partition.



* The PCB bypass points are controlled by bit 11 or bit 12 of the NODE_CSR. For an example of this PCB bypass, refer again to Figure 9-14.

Figure 9-17. Bypass 2A

Information on the locations of the bypass points for each of the four physical barrier synchronization circuits in each configuration of the CRAY T3D system is provided in the manual, “CRAY T3D Configurations.”

Hardware Description

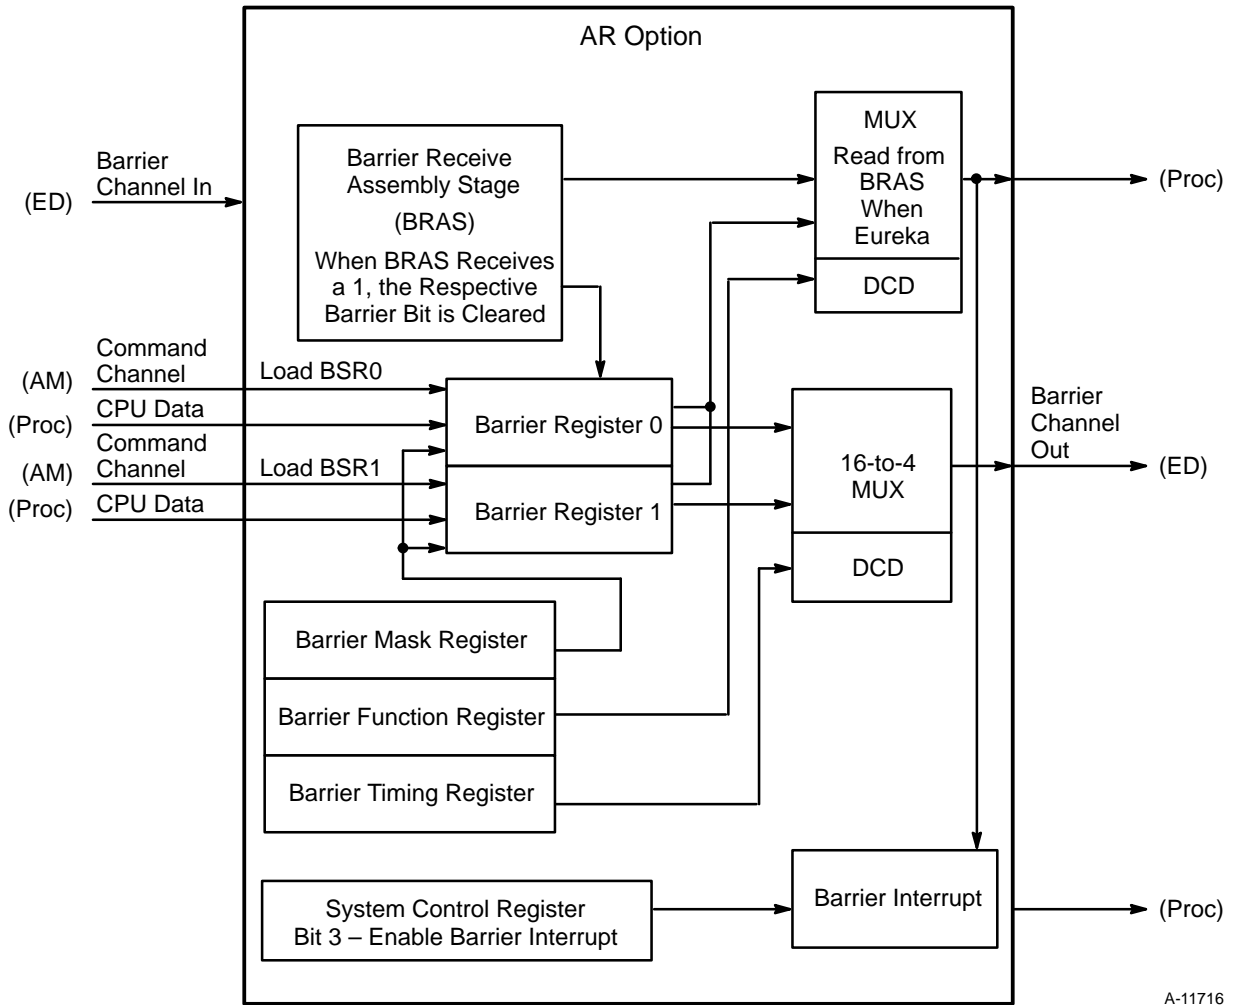
The hardware that makes up the barrier synchronization circuitry sets and clears the appropriate barrier register bits as specified by the microprocessor and performs the fanin and fanout of the barrier bits. The following subsections describe this hardware.

Physical Barrier Synchronization Circuitry

The options that contain the barrier synchronization circuitry are the AR option, the ED option, and the BC option.

AR Option

The AR option contains the two barrier registers, BSR0 and BSR1. Refer to Figure 9-18. When a microprocessor arrives at a barrier in a program, the microprocessor signals the AR option to set a specific barrier bit. This barrier bit remains set until the AR option receives a 1 on the appropriate barrier channel input. This 1 indicates that the barrier has been reached by all of the other microprocessors. After receiving this signal, the AR option clears the appropriate barrier bit and, if enabled, sets the barrier interrupt. When the microprocessor reads the barrier interrupt and finds it set, the microprocessor knows that a barrier synchronization operation is complete. The microprocessor then reads the BSR0 and BSR1 registers to determine which barrier completed. After determining which barrier completed, the microprocessor continues with the appropriate program code.



A-11716

Figure 9-18. Barrier Synchronization Circuitry within the AR Option

The AR option also contains the barrier mask register, the barrier function register, and the barrier timing register. For more information about these registers, refer to the “Register Mapping” subsection within this section and to the subsection, “AR Option” in Section 10, “Control and Status.”

ED Option

There are two ED options on each PCB. These ED options contain the AND (fan-in) tree and the fan-out tree circuitry. These options are ED0 and ED10. Refer to Figure 9-19. For the AND tree, the ED options perform a 4-to-1 fanin of the barrier bits. The 4-to-1 fanin is a 5-level AND tree that makes it possible to AND all of the microprocessors' barrier bits together in a reasonable amount of time.

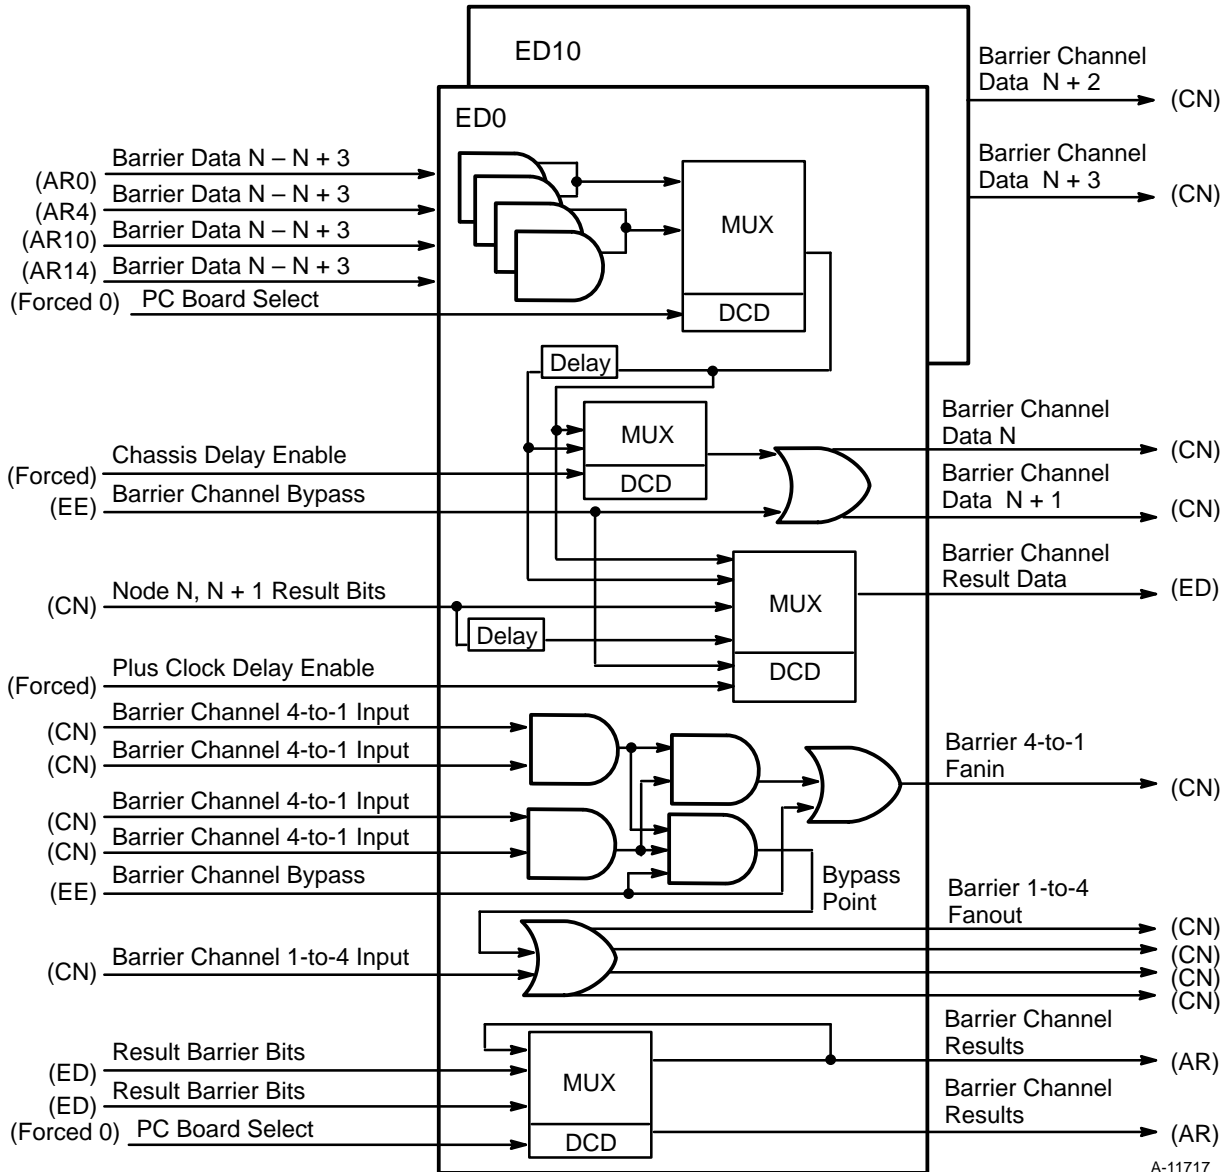


Figure 9-19. Barrier Synchronization Circuitry within the ED Options

For the fan-out tree, the ED options perform a 1-to-4 fanout of the barrier bits. The fanout also requires 5 levels.

As stated in the “Physical Partitions” subsection, each AND gate in the AND tree is paired with fan-out circuitry in the fan-out tree. This AND gate and fan-out circuitry is called a bypass point. Refer again to Figure 9-19. The bypass point is enabled when the Barrier Channel Bypass signal is set to a 1. The bypass point enables the output of the 4-to-1 fan-in circuitry to be redirected to the 1-to-4 fan-out circuitry.

BC Option

The BC options are located on the system module. These options receive the final barrier bits from the AND trees of the processing element modules, the I/O modules, and the redundant processing element module. Refer to Figure 9-20. The BC options perform the final AND of the barrier bits and fan out the result back to the processing element modules, the I/O modules, and the redundant processing element module.





Figure 9-20. Fanin and Fanout of Barrier Bits within the BC Option

The BC options also receive the final barrier bits from the AND trees from the other chassis within the CRAY T3D system. The BC options perform the final AND of all of the barrier bits within the system (or within a partition if the partition extends to one or more chassis) and fan out the result back to all of the chassis. When there are more than one chassis in the CRAY T3D system, the BC options of all of the chassis share in the fan-in and fan-out process. Refer to Figure 9-21 through Figure 9-26.

For a two chassis system, the processing element modules, the I/O modules, and the redundant PE module of each chassis fan in the barrier bits to the system PCB within their chassis. The BC options of the two system PCBs perform the AND of all of the barrier bits within their chassis and send the result of the AND to each other. The BC options then perform one more AND. This final AND gate ANDs the barrier bits of the two chassis together. The result of the final AND is then fanned out to the modules within the chassis.

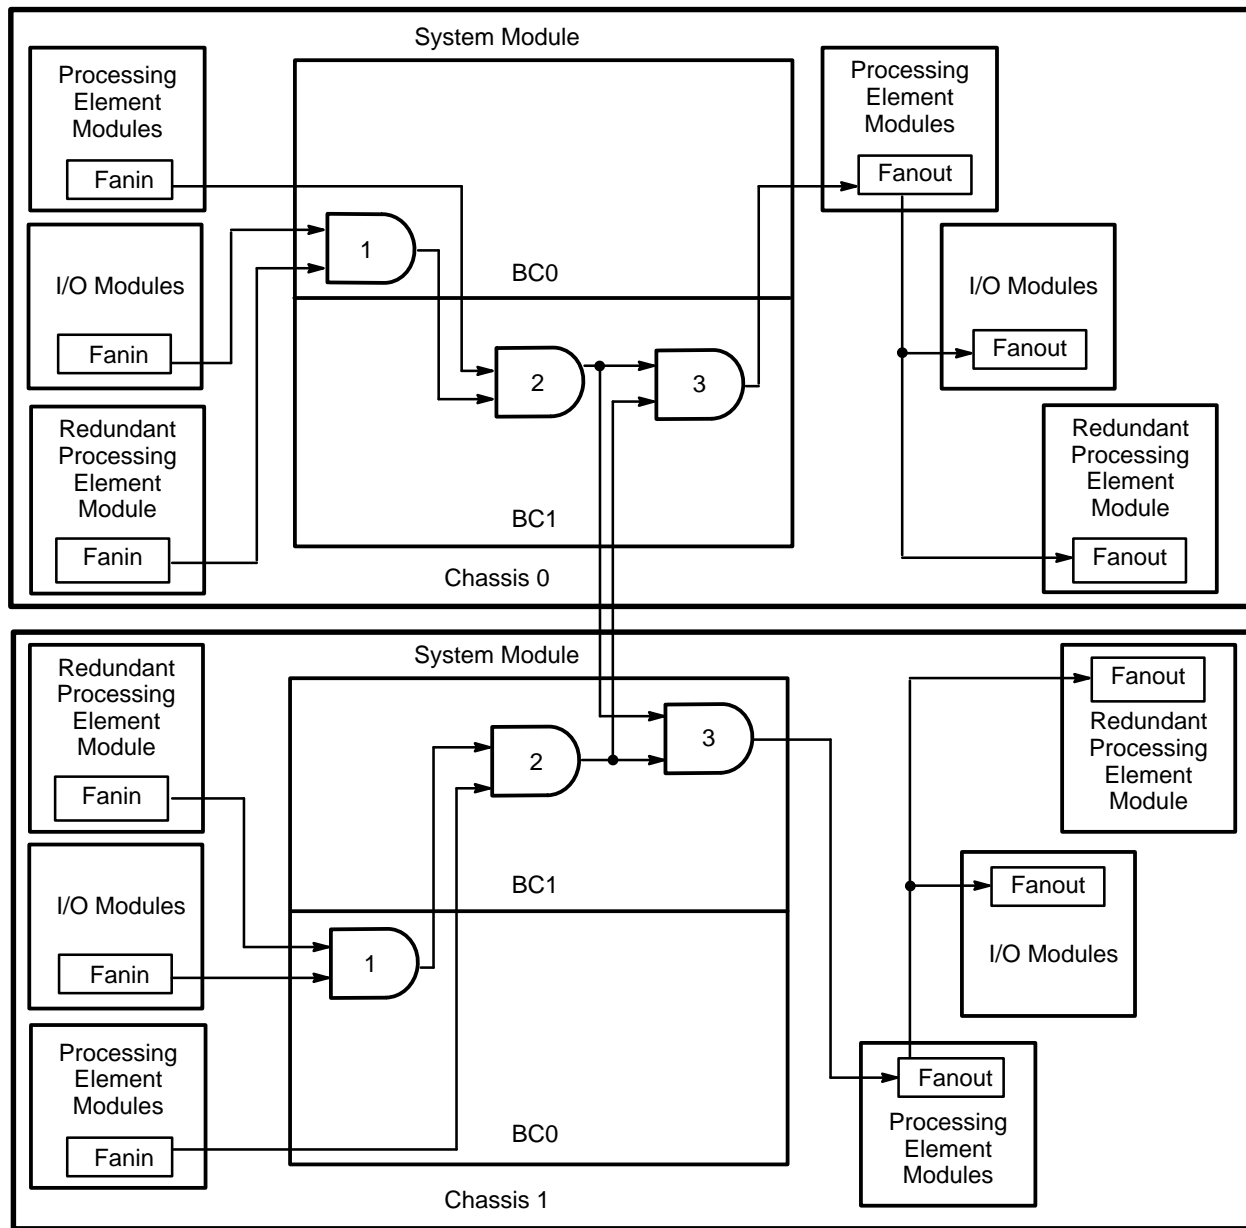
For a four chassis system, the processing element modules, the I/O modules, and the redundant processing element module of each chassis fan in the barrier bits to the system PCB within their chassis. The BC options of the four system PCBs perform the AND of all of the barrier bits within their individual chassis. The BC option of Chassis 0 sends the result of its AND to Chassis 1, and the BC option of Chassis 3 sends the result of its AND to Chassis 2. The BC option of Chassis 1 then ANDs its own results with the results from Chassis 0. The BC option of Chassis 2 ANDs its own results with the results from Chassis 3. Chassis 1 then sends this result to Chassis 2, and Chassis 2 sends its result to Chassis 1. Chassis 1 and 2 perform another AND, which ANDs the barrier bits of all four chassis together. Chassis 1 and Chassis 2 fan out the result to Chassis 0 and Chassis 3 respectively. Within each chassis, the barrier result is fanned out to the processing element modules, the I/O modules, and the redundant PE module.



1. AND barrier bits of I/O modules and redundant processing element module.
2. AND results from 1 to the barrier bits of the processing element node.
3. AND results from 2 to the barrier bits from other chassis.

A-11720

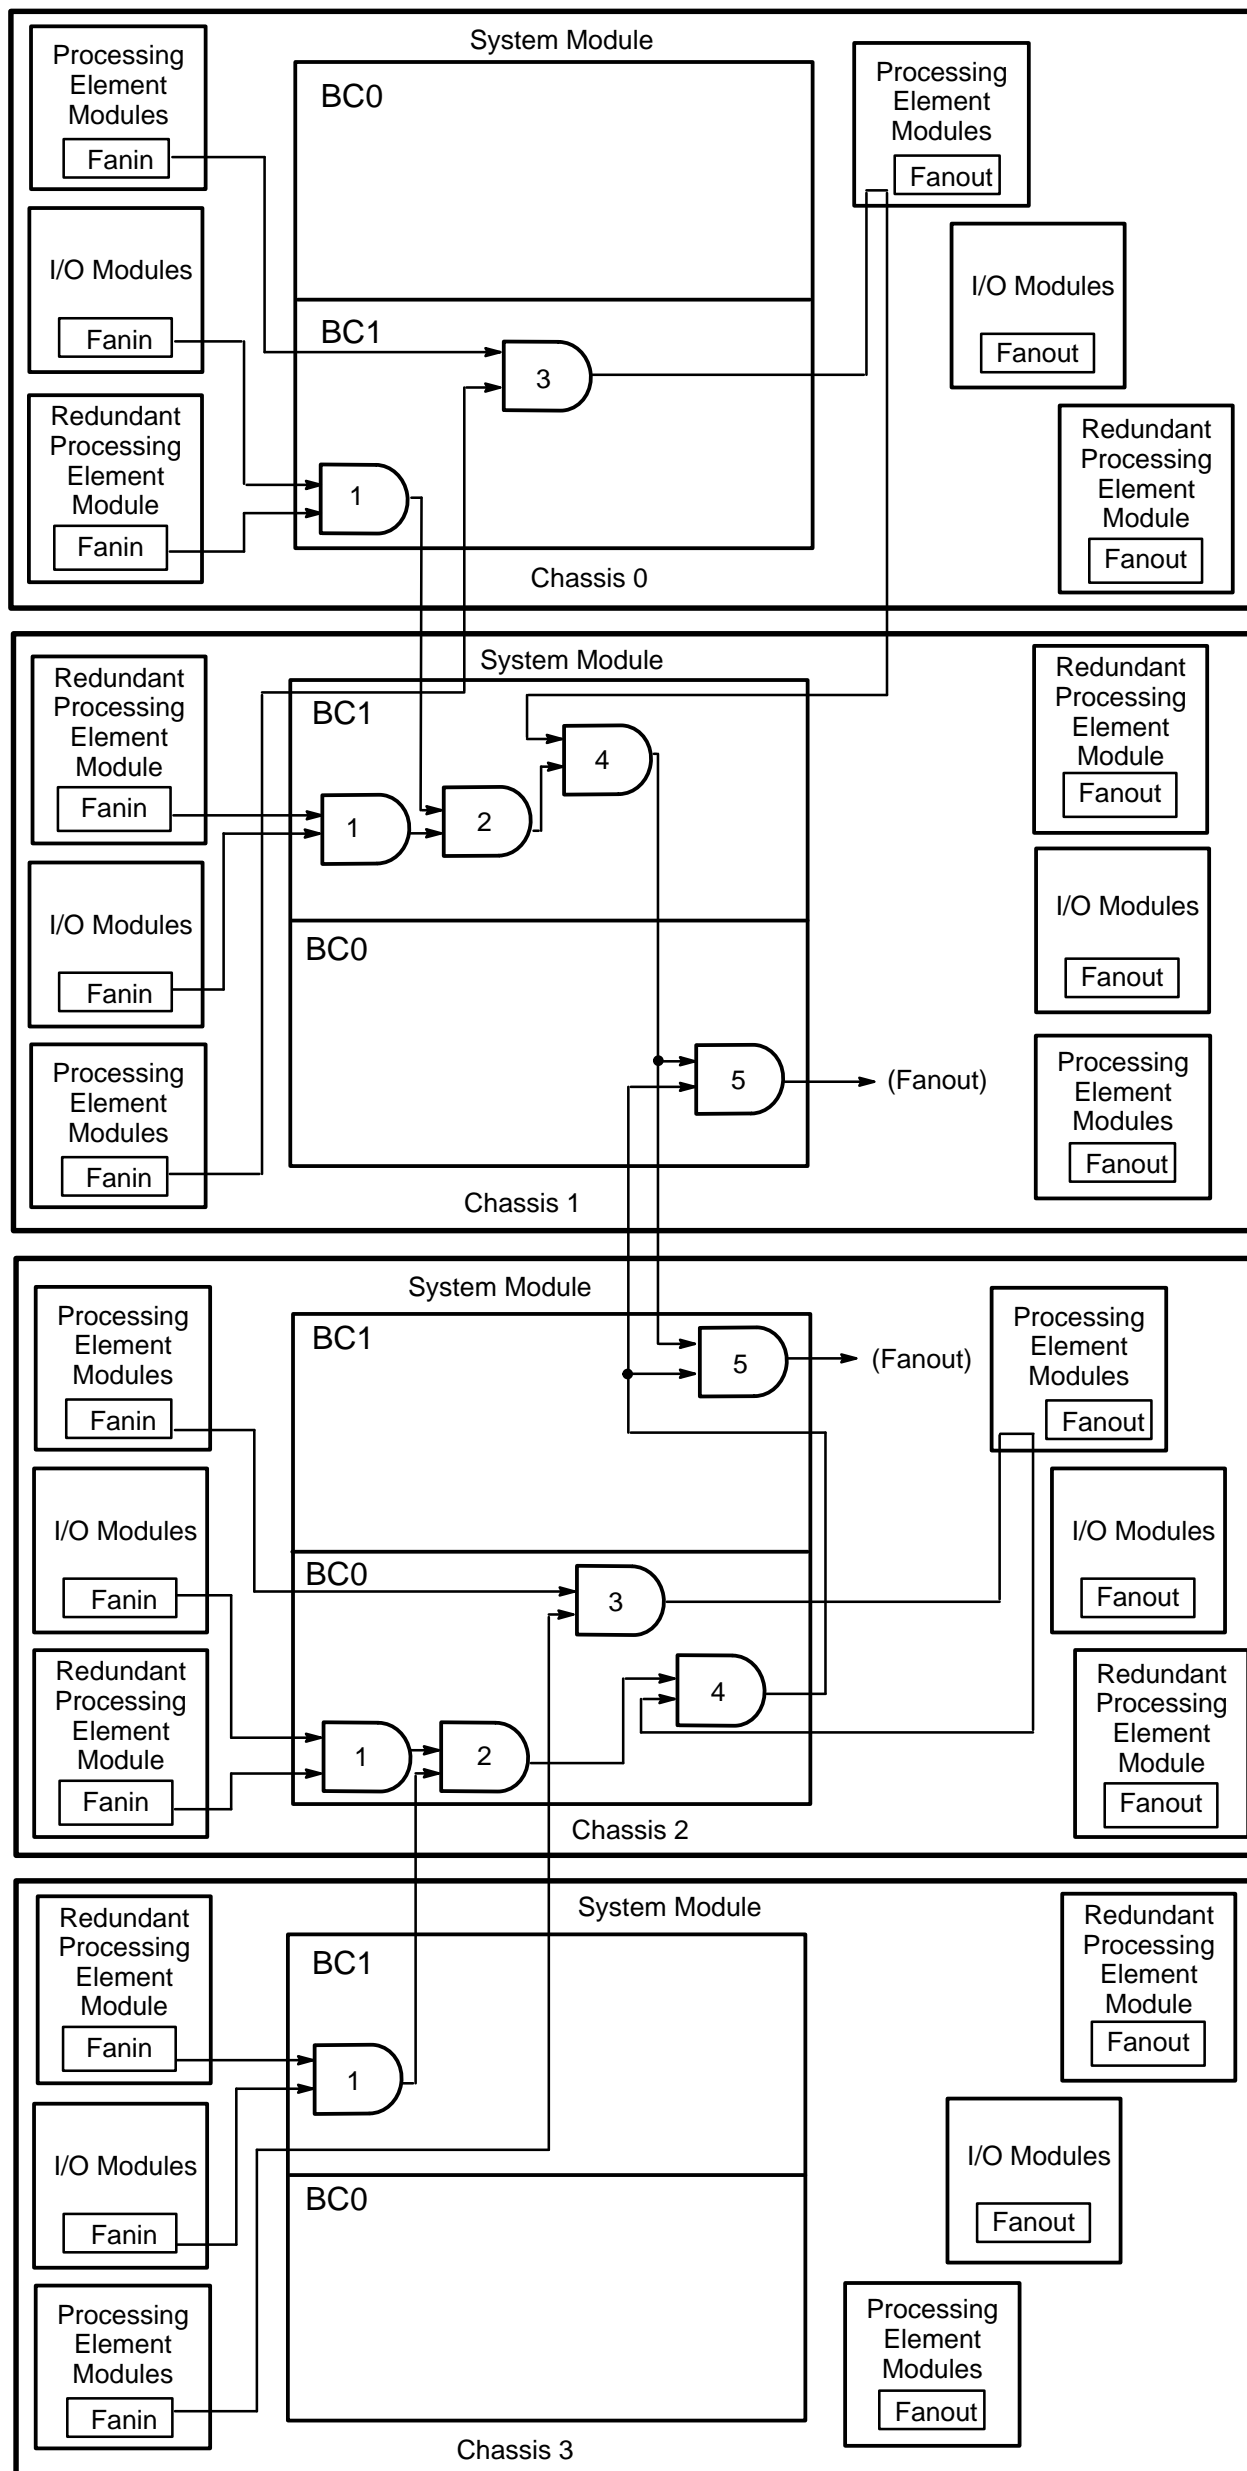
Figure 9-21. Fanin and Fanout of Barrier Bits n+0 and n+1 in a Two Chassis System



1. AND barrier bits of I/O modules and redundant processing element module.
2. AND results from 1 to the barrier bits of the processing element node.
3. AND results from 2 to the barrier bits from other chassis.

A-11721


Figure 9-22. Fanin and Fanout of Barrier Bits n+2 and n+3 in a Two Chassis System



1. AND barrier bits of I/O modules and redundant processing element module.
2. AND results of "AND 1" of chassis 0 and chassis 1 and of chassis 2 and chassis 3.
3. AND barrier bits of the processing element modules from chassis 0 and chassis 1 and from chassis 2 and chassis 3.
4. AND results of "AND 2" with the results of "AND 3."
5. AND results of "AND 4" from chassis 0 and chassis 1 with results of "AND 4" from chassis 2 and chassis 3.


A-11722

Figure 9-23. Fanin of Barrier Bits n+0 or n+1 in a Four Chassis System



A-11723

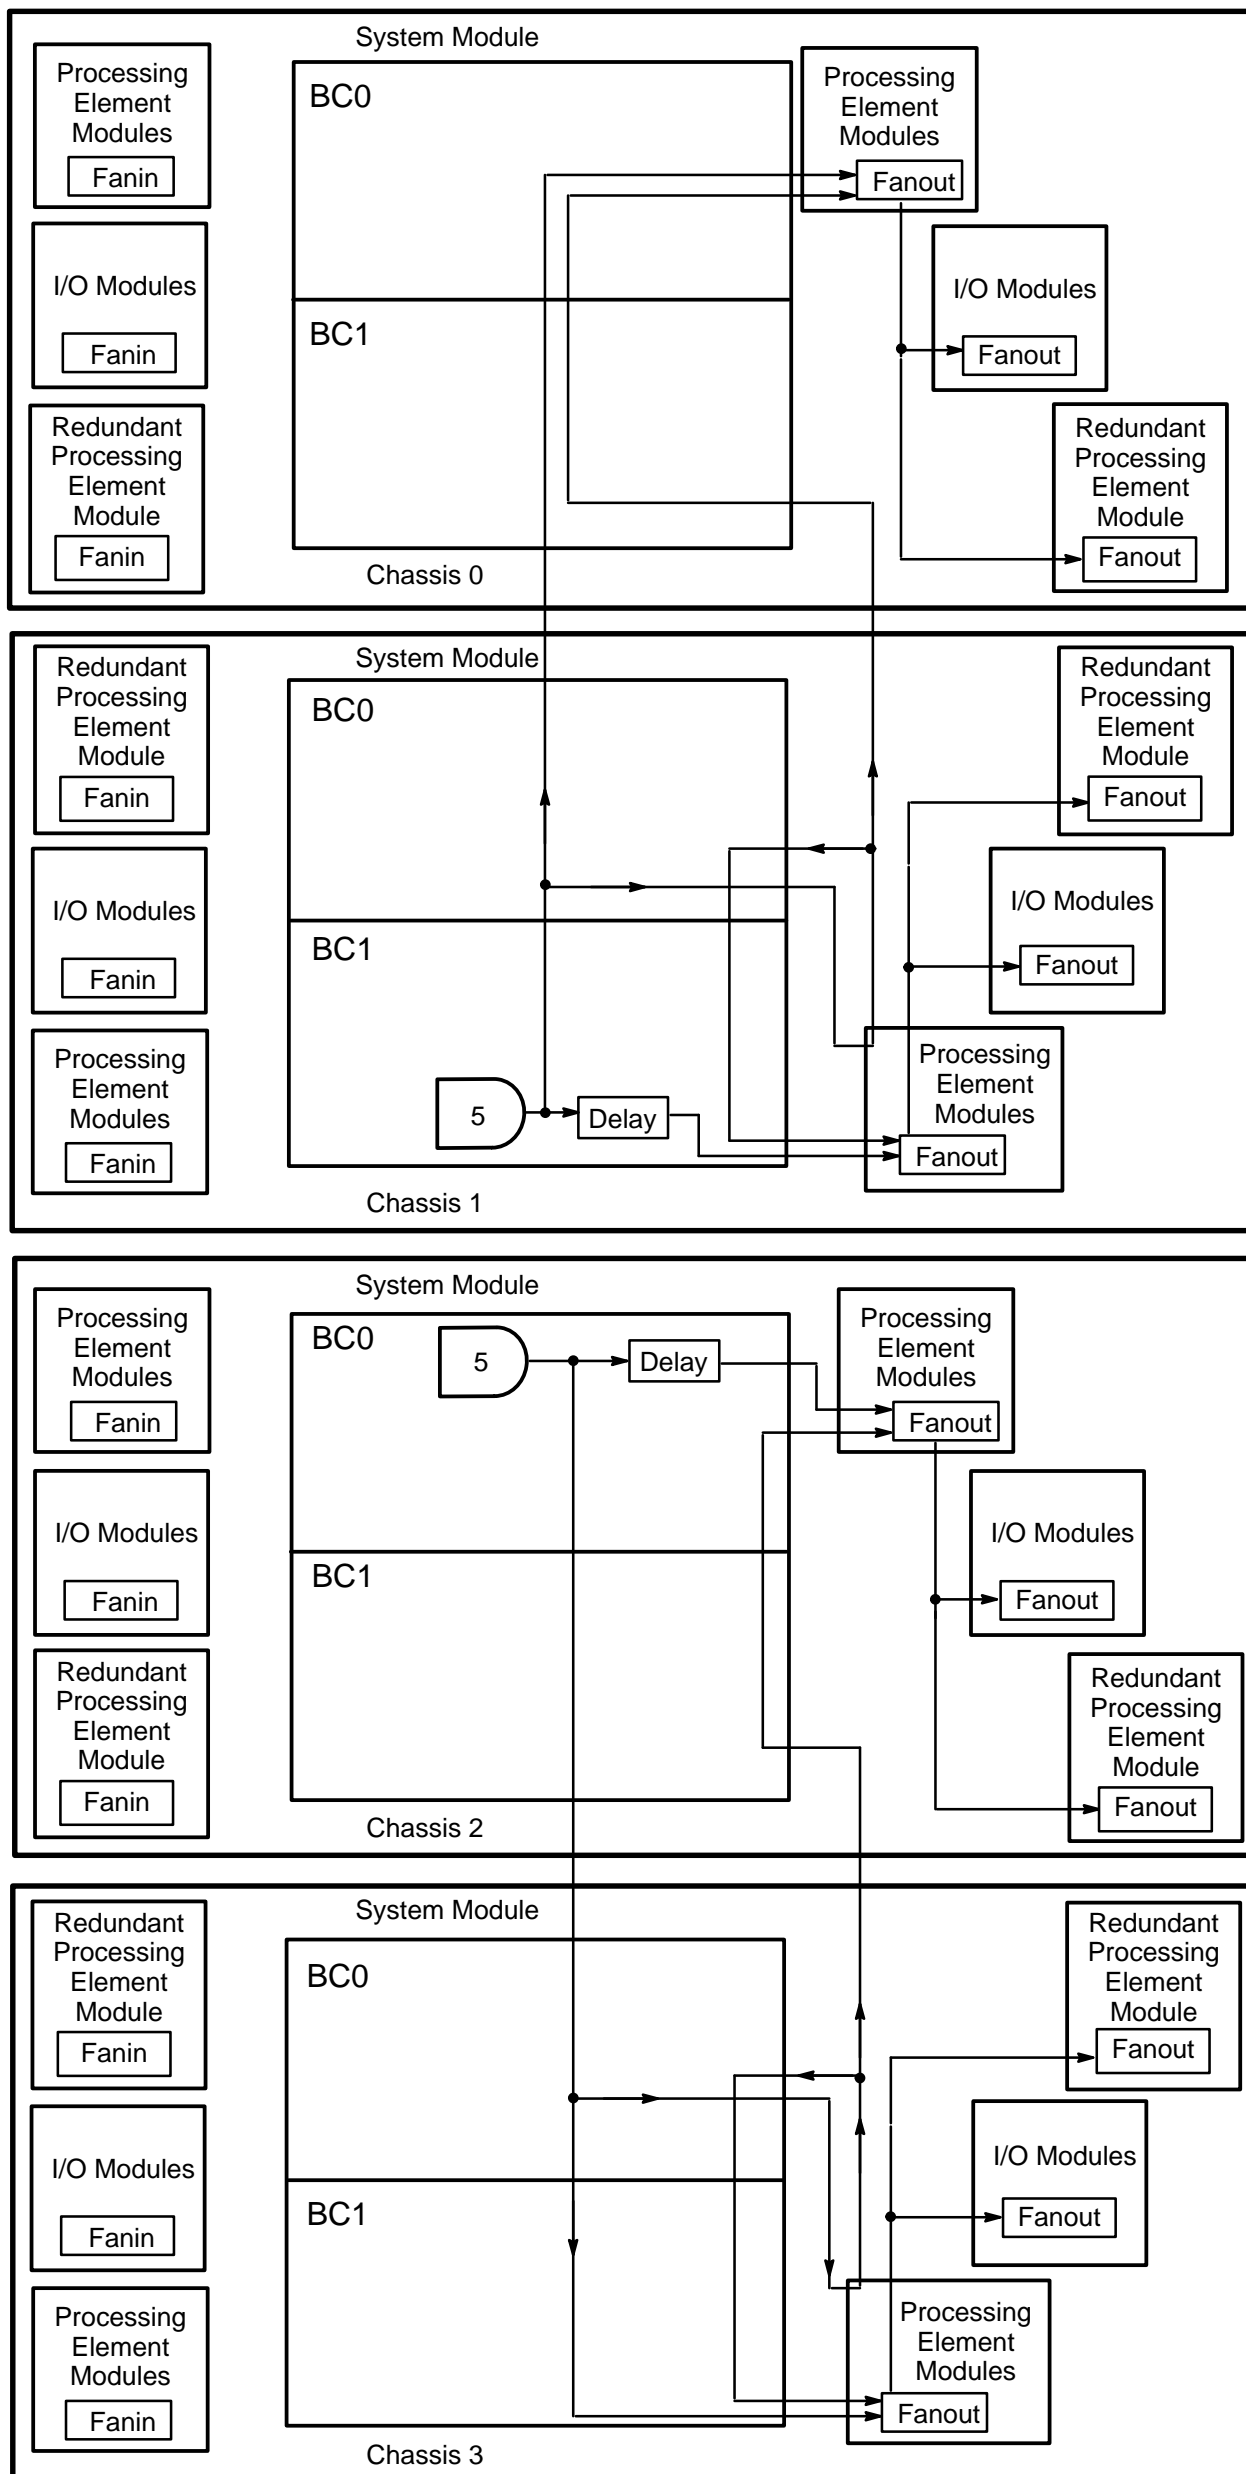
Figure 9-24. Fanout of Barrier Bits $n+0$ or $n+1$ in a Four Chassis System



1. AND barrier bits of I/O modules and redundant processing element module.
2. AND results of "AND 1" of chassis 0 and chassis 1 and of chassis 2 and chassis 3.
3. AND barrier bits of the processing element modules from chassis 0 and chassis 1 and from chassis 2 and chassis 3.
4. AND results of "AND 2" with the results of "AND 3."
5. AND results of "AND 4" from chassis 0 and chassis 1 with results of "AND 4" from chassis 2 and chassis 3.

A-11724

Figure 9-25. Fanin of Barrier Bits n+2 or n+3 in a Four Chassis System



A-11725

Figure 9-26. Fanout of Barrier Bits n+2 or n+3 in a Four Chassis System

Fanin Example of Barrier Bits 0 through 3 from PCB Bypass Point 000_{16}

The following text is an example of how the hardware of a 512 PE multiple cabinet (power of 2) system performs the fanin of barrier bits 0 through 3. While reading through this example, refer to Figure 9-27 and Figure 9-28.

Printed circuit board (PCB) bypass point 000_{16} is a PCB in level 1 of the AND tree. There are four PEs associated with this PCB: two from node 0 and two from node 1. When the microprocessors in the PE arrive at a barrier in a program, the microprocessors write a 1 to the appropriate bit of the BSR0 or BSR1 register in the AR option. The AR option sets the appropriate barrier bit and sends the barrier bits to the ED0 and ED10 options in four 4-bit transfers (for fanin).

The ED0 and ED10 options perform the 4-to-1 AND of the 4 bits from the four PEs. Keep in mind that the ED options within all of the nodes that are participating in the barrier are performing the same 4-to-1 AND. A forced 0 on the PC Board Select signal, which is presented to the ED0 option, directs the ED0 option to send barrier bits $n+0$ and $n+1$ to physical node (hereafter referred to as node) 100_{16} and node 102_{16} respectively. The PC Board Select signal on the ED10 option is a forced 1. This forced 1 directs the ED10 option to send barrier bits $n+2$ and $n+3$ to node 110_{16} and node 112_{16} respectively.

The ED options of node 100_{16} , node 102_{16} , node 110_{16} , and node 112_{16} receive the barrier bits from PCB bypass points 000_{16} , 002_{16} , 100_{16} , and 102_{16} . The ED0 option of node 100_{16} receives barrier bit $n+0$. The ED0 option of node 102_{16} receives barrier bit $n+1$. The ED10 option of node 110_{16} receives barrier bit $n+2$. And the ED10 option of node 112_{16} receives barrier bit $n+3$. Each of these ED options performs another 4-to-1 AND of the barrier bits and sends the results to the next level of the AND tree.

Level three of the AND tree involves the ED options of node 000_{16} , node 002_{16} , node 010_{16} , and node 012_{16} . Node 000_{16} receives barrier bit $n+0$ from node 100_{16} , node 104_{16} , node 300_{16} , and node 304_{16} . Node 002_{16} receives barrier bit $n+1$ from node 102_{16} , node 106_{16} , node 302_{16} , and node 306_{16} . Node 010_{16} receives barrier bit $n+2$ from node 110_{16} , node 114_{16} , node 310_{16} , and node 314_{16} . And node 012_{16} receives barrier bit $n+3$ from node 112_{16} , node 116_{16} , node 312_{16} , and node 316_{16} . Each of these ED options performs another 4-to-1 AND of the barrier bits and sends the results to the next level of the AND tree.

Level 4 of the AND tree involves the ED options of node 004_{16} , node 034_{16} , node 014_{16} , and node 024_{16} . Node 004_{16} receives barrier bit $n+0$ from node 000_{16} and node 030_{16} . Node 034_{16} receives barrier bit $n+1$ from node 002_{16} and node 032_{16} . Node 014_{16} receives barrier bit $n+2$ from node 010_{16} and node 020_{16} . Node 024_{16} receives barrier bit $n+3$ from node 012_{16} and node 022_{16} . Each of these ED options performs another 4-to-1 AND of the barrier bits and sends the results to the next level of the AND tree.

Level 5 of the AND tree involves the ED options of node 006_{16} , node 036_{16} , node 016_{16} , and node 026_{16} . Node 006_{16} receives barrier bit $n+0$ from node 004_{16} , node $00A_{16}$, node 404_{16} , and node $40C_{16}$. Node 036_{16} receives barrier bit $n+1$ from node 034_{16} , node $03A_{16}$, node 434_{16} , and node $43C_{16}$. Node 016_{16} receives barrier bit $n+2$ from node 014_{16} , node $01A_{16}$, node 414_{16} , and node $41C_{16}$. Node 026_{16} receives barrier bit $n+3$ from node 024_{16} , node $02A_{16}$, node 424_{16} , and node $42C_{16}$. Each of these ED options performs another 4-to-1 AND of the barrier bits and sends the results to the system PCB.

The BC options of the system PCB receive the final barrier bits from the processing element modules, the I/O modules, and the redundant processing element module. The BC options perform the final AND of these barrier bits and then fan out the result to the processing element modules, the I/O modules, and the redundant processing element module.

The barrier bits are fanned out to all of the PEs within the partition. This fanout is handled by the same nodes and options as the AND tree. The only difference is that instead of the ED options performing a 4-to-1 AND, they perform a 1-to-4 fanout.




Figure 9-27. Fanin of Barrier Bits $n+0$ through $n+3$ from PCB Bypass Point 000₁₆ (page 1 of 2)




Figure 9-28. Fanin of Barrier Bits $n+0$ through $n+3$ from PCB Bypass Point 000₁₆ (page 2 of 2)

10 Control and Status

Each PE in the CRAY T3D system contains several registers and signals used to control node functions and provide status on those functions. These registers and signals control functions of the microprocessor, the support circuitry, and the dynamic random access memory (DRAM) circuitry used in local memory.

Functional Description

The following subsections provide a detailed description of the signals and registers used to control the components of a node and to obtain the status of node functions.

Control Signals

The microprocessor and support circuitry in a PE use several signals to control PE functions. Figure 10-1 shows the main control signals between the microprocessor and support circuitry. These signals include the hardware interrupt pins, the instruction cache test mode pins, the cycle request pins, the cycle acknowledge pins, and the cache line invalidate pins.

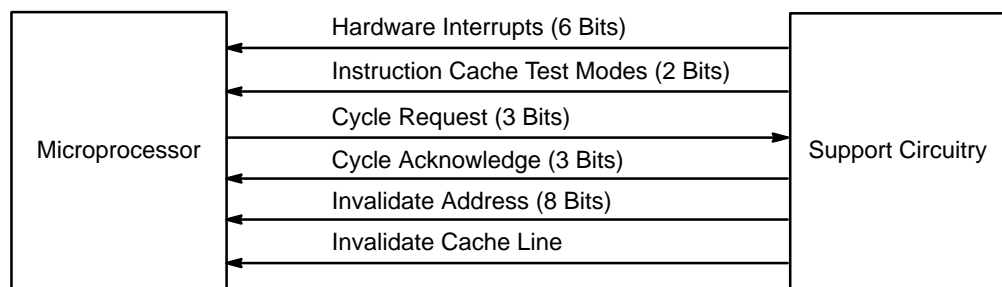


Figure 10-1. Microprocessor and Support Circuitry Control Signals

Hardware Interrupt Pins

Six hardware interrupts are generated by the support circuitry and sent to the microprocessor. Software must enable the hardware interrupts by setting the appropriate bits of the hardware interrupt enable register (HIER) in the microprocessor.

Software may read the value of the hardware interrupts by reading the value of the hardware interrupt request register (HIRR) in the microprocessor. Table 10-1 lists the hardware interrupts and shows the corresponding interrupt request pin, HIER register bit, HIRR register bit, and method of clearing the hardware interrupt. The following subsections describe each hardware interrupt.

Table 10-1. Hardware Interrupt Signals

Interrupt Request Pin	HIER Register Bit	HIRR Register Bit	Name	Cleared by
0	9	10	Heartbeat interrupt	Writing a 1 to bit 2 (Clear Heartbeat Interrupt bit) in the system control register (SCR)
1	10	11	Barrier interrupt	Reading the value of the barrier synchronization mask and interrupt (BSMI) register
2	11	12	Message interrupt	Writing to the message queue limit increment register (MQ_LIR)
3	12	5	BLT interrupt	Writing to the BLT control register (BLT_CR)
4	13	6	I/O interrupt	Refer to Section 7, "Input/output" for information on clearing the I/O interrupt bit
5	14	7	Error interrupt	Reading the value stored in the system status register (SSR)

Heartbeat Interrupt

The support circuitry sets the heartbeat interrupt each time it encounters a heartbeat event. A heartbeat event is triggered by system heartbeat signals. System heartbeat signals are generated by the deadstart I/O gateway and are distributed to each PE in the system at the same time. There are four types of heartbeat signals: a normal heartbeat, a deadman heartbeat, a selective reset heartbeat, and a global reset heartbeat.

The normal heartbeat signal is used to trigger a normal heartbeat event. A normal heartbeat event is used by the PEs to establish global time of day.

After receiving the normal heartbeat signal, the support circuitry sets the heartbeat interrupt. The microprocessor then clears the heartbeat interrupt by writing a 1 to bit 2 (clear heartbeat interrupt bit) in the SCR.

The deadman heartbeat signal triggers a deadman heartbeat event, which detects whether a PE is malfunctioning or in a hung condition.

After receiving the deadman heartbeat signal, the support circuitry sets the heartbeat interrupt and sets a deadman timer (DMT) latch. The microprocessor then clears the heartbeat interrupt and DMT latch by writing a 1 to bit 2 (clear heartbeat interrupt bit) in the SCR.

If the microprocessor does not clear the DMT latch before the support circuitry receives the next deadman heartbeat signal, the microprocessor may be malfunctioning or in a hung condition. Because of this possibility, the PE is reset.

During a reset caused by a deadman heartbeat time-out, the PE starts a reboot cycle and the support circuitry sets bit 5 (node enable vote bit) of the SCR to 0. After the microprocessor reboots, it reads and issues boot code instructions that were previously stored in the lower 8 Kbytes of local memory.

The selective reset heartbeat signal triggers a selective reset event. The deadstart I/O gateway performs the selective reset event to reset selected PEs in the system.

Before the selective reset heartbeat signal is distributed to each PE in the system, the deadstart I/O gateway informs each PE that will not be reset, to set bit 4 (selective reset mask bit) of the SCR to 1. After receiving the selective reset heartbeat signal, the support circuitry sets the heartbeat interrupt and, if previously set to 1, clears bit 4 of the SCR.

If bit 4 of the SCR was set to 0 when the PE received the selective reset heartbeat signal, the PE starts a reboot cycle and the support circuitry sets bit 5 (node enable vote bit) of the SCR to 0. After the microprocessor reboots, it reads and issues boot code instructions that were previously stored in the lower 8 Kbytes of local memory.

The global reset heartbeat signal triggers a global reset event. The deadstart I/O gateway performs a global reset to reset all of the PEs in the system.

After receiving the global reset heartbeat signal, the support circuitry sets the heartbeat interrupt. The PE then performs a reset.

During a global reset, the PE starts a reboot cycle, the support circuitry sets bit 5 (node enable vote bit) of the SCR to 0, and the network interface is master cleared. After the microprocessor reboots, it reads and issues boot code instructions that were previously loaded into the instruction cache memory from a serial read only memory (ROM).

Barrier Interrupt

The support circuitry sets the barrier interrupt when a barrier or eureka synchronization operation completes. Depending on the value of bit 3 (enable BSR0 interrupt bit) in the SCR, the barrier interrupt signals the microprocessor that one of two events has taken place. First, a barrier or eureka synchronization operation associated with the final barrier bits of barrier synchronization register 1 (BSR1) is completed. Second, a barrier or eureka synchronization operation associated with the final barrier bits in BSR0 and/or BSR1 is completed.

Message Interrupt

The support circuitry sets the message interrupt after the support circuitry receives an incoming message from the network interface and stores the message in the message queue. The support circuitry does not set the message interrupt when it rejects a message because the message queue is full. The message interrupt signals the microprocessor that one or more messages are in the message queue.

BLT Interrupt

The support circuitry sets the BLT interrupt when it receives an interrupt signal from the BLT. The BLT interrupt indicates that one of the following conditions occurred.

- The BLT is free and another transfer operation can be initiated.
- All outstanding BLT requests have completed.
- A BLT error occurred.

The microprocessor must read the values stored in the BLT status register (BLT_SR) to determine the cause of the BLT interrupt.

I/O Interrupt

The I/O interrupt is used only in the PEs that reside in the I/O gateways. For more information on the I/O interrupt, refer to Section 11, “LOSP and HISP.”

Error Interrupt

The support circuitry sets the error interrupt when any of the following error conditions occur:

- Network buffer parity error
- Network packet error
- Message queue oversubscribed
- Prefetch queue overrun error
- Prefetch queue underrun error
- Virtual PE number range error
- Illegal register access error
- DTB annex error
- Illegal write to protected low memory (PLM) – I/O gateway PEs only

More information on each of these errors is provided in “System Status Register” later in this section.

Instruction Cache Test Mode Pins

The instruction cache test mode pins are set by the support circuitry during a PE reset. These pins signal the microprocessor that code was loaded into the instruction cache and the instruction cache is valid or they inform the microprocessor that the instruction cache is invalid and instructions must be read from memory. Table 10-2 shows the value of bits 0 and 1 of the instruction cache test mode pins and lists the corresponding location for the boot code. Bit 1 is always set to 0.

Table 10-2. Instruction Cache Test Mode Pins

Bits 1 and 0	Description
00	Instructions were loaded from ROM and placed in the instruction cache. In this case, the instructions in the instruction cache are valid. These instructions are used during a global reset.
01	The instructions in the instruction cache are not valid and instructions must be read from memory. These instructions are used during a deadman time-out reset or a selective reset.

Cycle Request Pins

The microprocessor places information on the cycle request pins when the microprocessor requests that the support circuitry perform an operation. Table 10-3 shows the value of the cycle request pins and lists the corresponding cycle requests.

Table 10-3. Cycle Request Codes

Cycle Request Pins $2^2 - 2^0$	Cycle	Description
000	IDLE	The microprocessor has not requested that the support circuitry perform a function.
001	BARRIER	The microprocessor requested that the support circuitry perform a function associated with the memory barrier (MB) instruction.
010	FETCH	The microprocessor requested that the support circuitry read a word of data from memory and store the word in the prefetch queue. During this operation, the support circuitry prevents subsequent local memory operations from completing.
011	FETCHM	The microprocessor requested that the support circuitry read a word of data from memory and store the word in the prefetch queue. During this operation, the support circuitry does not prevent subsequent local memory operations from completing.
100	READ_BLOCK	The microprocessor requested that the support circuitry read data from memory or a memory-mapped register and transfer the data to the microprocessor.
101	WRITE_BLOCK	The microprocessor requested that the support circuitry transfer data from the microprocessor to memory or a memory-mapped register.
110	LDxL	The microprocessor requested that the support circuitry transfer information from a location in the DTB annex to the microprocessor.
111	STxC	The microprocessor requested that the support circuitry transfer information from the microprocessor to a location in the DTB annex.

Cycle Acknowledge Pins

The support circuitry places information on the cycle acknowledge pins to acknowledge a cycle request. Table 10-4 shows the value of the cycle acknowledge pins and lists the corresponding cycle acknowledges.

Table 10-4. Cycle Acknowledge Codes

Cycle Request Pins $2^2 - 2^0$	Cycle Acknowledge	Description
000	IDLE	The support circuitry is not performing a function for the microprocessor.
001	HARD_ERROR	The support circuitry encountered an error as defined in the following paragraph.
010	SOFT_ERROR	This is not used in the CRAY T3D system.
011	STxC_FAIL	The support circuitry completed a transfer of information from the microprocessor to a location in the DTB annex.
100	OK	The support circuitry successfully completed the requested function (except the STxC cycle).

The support circuitry sets the cycle acknowledge pins to indicate a **HARD_ERROR** when an address range error occurs. More information on errors is provided in “System Status Register” later in this section.

Cache Line Invalidate Pins

The support circuitry uses the cache line invalidate pins to signal the microprocessor when a line in the data cache memory should be discarded. There are two groups of cache line invalidate pins: the invalidate address pins and the invalidate cache line pins.

The support circuitry uses the invalidate address pins to identify which line in the data cache should be discarded. The support circuitry uses an invalidate cache line pin to signal the microprocessor when the cache line identified by the invalidate address pins should be discarded.

The cache line invalidate pins are used when the support circuitry performs a filtered or a nonfiltered cache line invalidate operation. More information on the data cache and the cache line invalidate operations is provided in Section 3, “Processing Element Node.”

Status Information

Each PE contains registers that indicate the status of PE functions. The status information includes error information and outstanding request information. Two registers provide this information: the system status register and the user control and status register. More information on these registers is provided in “Register Mapping” in this section.

Register Mapping

The following subsections describe the addressing and bit assignments for the memory-mapped registers (hereafter referred to as registers) used for control and status in the CRAY T3D system. Each subsection also provides a brief summary of the function of the register.

Table 10-5 is a summary of the control and status registers and their names. Table 10-5 also lists the partial physical address of each register as it appears on the address pins of the microprocessor.

Table 10-5. Control and Status Registers

Address	Register Name	Direction	Description
10410000 ₁₆	SCR	Write	System control register
10448000 ₁₆	SSR	Read	System status register
10144000 ₁₆	UCSR	Read or write	User control and status register
1044C000 ₁₆	DRAM_CR	Write	DRAM control register
10404000 ₁₆	BAR_TMGM	Write	Barrier timing and refresh register
106C0000 ₁₆	NET_ENA	Write	Network enable register
106D0000 ₁₆	NET_PFM	Write	Network performance monitor register
106E0000 ₁₆	NODE_CSR	Write	Network mode register
10140000 ₁₆	NOR	Read or write	No-operation register

NOTE: Because of multiplexed data paths, when one PE in a processing element node is modifying the contents of the NET_ENA, NET_PFM, or NODE_CSR register, the other PE in the node must not attempt to modify any of the shared registers at the same time. The shared registers include the X_WHOAMI, LPE_XLATE, ROUTE_LO, ROUTE_HI, NET_ENA, NET_PFM, NODE_CSR, and BLT registers.

Because the virtual address is defined by software, the addresses for each of the registers are given according to the partial physical address as it appears on the pins of the microprocessor.

System Control Register
Address 10410000_{16}

The system control register (SCR) is a 6-bit, write-only, system privileged register that controls the privileged mode resources of a PE.

Figure 10-2 shows the bit assignments for the SCR address as they appear on the address pins of the microprocessor.

HEX	1	0				4				1				0				0				0							
Binary	2^{28}	2^{27}	2^{26}	2^{25}	2^{24}	2^{23}	2^{22}	2^{21}	2^{20}	2^{19}	2^{18}	2^{17}	2^{16}	2^{15}	2^{14}	2^{13}	2^{12}	2^{11}	2^{10}	2^9	2^8	2^7	2^6	2^5	2^4	2^3	2^2	2^1	2^0
	1	x	x	x	x	x	1	0	0	0	0	1	0	0	x	x	x	x	x	x	x	x	x	x	x	x	x	x	x

NOTE: x = Don't Care

Figure 10-2. System Control Register Address Bit Format

Table 10-6 shows the bit format of the SCR, and the following subsections describe each bit of the register.

Table 10-6. SCR Format

Bits	Name
0	Low DTB annex write enable
1	High DTB annex write enable
2	Clear heartbeat interrupt
3	Enable barrier synchronization register 0 (BSR0) interrupts
4	Selective reset mask
5	Node enable vote
63 – 6	Not used

Low DTB Annex Write Enable
Bit 0

When set to 1, this bit enables entries 0 through 15 of the DTB annex so they may be written to by the user or the operating system. In this case, the support circuitry in a PE interprets entries 1 through 15 of the DTB annex as virtual PE numbers and function codes. The support circuitry also reads bit 11 (I/O node bit) of the PE number in the DTB annex entry as 0.

The support circuitry always interprets entry 0 of the DTB annex as a logical PE number and function code. In order for the support circuitry to read the correct value of bit 11 of DTB annex entry 0, bit 0 of the SCR must be set to 0. (This only affects PE numbers where bit 11 of the PE number may be set to 1.)

When set to 0, this bit disables entries 0 through 15 of the DTB annex so the user or the operating system cannot write to them. In this case, the support circuitry interprets entries 0 through 15 of the DTB annex as logical PE numbers and function codes. The support circuitry also reads bit 11 of the PE number in the DTB annex entry as the value it is set to in the DTB annex entry.

High DTB Annex Write Enable
Bit 1

When set to 1, this bit enables entries 16 through 31 of the DTB annex so the user or the operating system may write to them. In this case, the support circuitry in a PE interprets entries 16 through 31 of the DTB annex as virtual PE numbers and function codes.

When set to 0, this bit disables entries 16 through 31 of the DTB annex so they cannot be written to by the user or the operating system. In this case, the support circuitry interprets entries 16 through 31 of the DTB annex as logical PE numbers and function codes.

Clear Heartbeat Interrupt
Bit 2

This bit is used after the support circuitry has set the heartbeat interrupt to 1. When set to 1, the clear heartbeat interrupt bit signals the support circuitry to set the heartbeat interrupt for the microprocessor to 0. When set to 0, the clear heartbeat interrupt bit has no effect.

Enable BSR0 Interrupt
Bit 3

When set to 1, this bit signals the support circuitry to set the barrier interrupt to 1 if a barrier associated with BSR0 or BSR1 occurs. In this case, barrier interrupts from both BSR0 and BSR1 are enabled.

When set to 0, the enable BSR0 interrupt bit signals the support circuitry to set the barrier interrupt to 1 only if a barrier associated with BSR1 occurs. In this case, barrier interrupts from BSR0 are disabled.

NOTE: When bit 3 of the SCR is set to 0 and a barrier interrupt occurs, the barrier interrupt is not stored by the hardware and may be lost.

Selective Reset Mask
Bit 4

When set to 1, this bit signals the support circuitry in a PE not to reset when it receives a reset signal. In this case, the support circuitry does not perform a reset, but does clear the value of the selective reset mask bit to 0 after receiving the reset signal. When set to 0, the selective reset mask bit signals the support circuitry in a PE to reset the PE after receiving a selective reset heartbeat signal.

Only the deadstart I/O gateway can perform a selective reset. Before performing the reset, the deadstart I/O gateway informs all of the PEs that will not be reset to set their selective reset mask bit to 1. The deadstart I/O gateway then performs the selective reset. Only the PEs that have the selective reset mask bit cleared to 0 are reset. This enables the deadstart I/O gateway to reset only the PEs that cannot communicate with the deadstart I/O gateway or that need to be reset for other reasons.

Node Enable Vote
Bit 5

When set to 0, this bit indicates that the PE has disabled node communication with the interconnect network. A global reset, non-masked selective reset, or deadman reset also sets the value of the node enable vote bit to 0.

When both PEs in a PE node set their node enable vote bit to 1, the PE node can use the interconnect network. The node enable vote bit in both PEs must be set to 1 to enable the PE node to communicate using the interconnect network. Because there is only one PE in an input node or

output node of an I/O gateway, the input or output node can communicate using the interconnect network when the single PE sets the node enable vote bit to 1.

**System Status Register
Address 10448000₁₆**

The system status register (SSR) is a 10-bit, read-only, system privileged register that contains the status of PE error conditions related to the error interrupt and the HARD_ERROR cycle acknowledge. When the SSR is read, the error interrupt and the condition associated with each SSR bit are cleared.

While the interrupt handler software services a HARD_ERROR related error, an error associated with the error interrupt may occur. Reading the SSR while servicing a HARD_ERROR related error may clear status associated with the error interrupt. Care should be taken not to accidentally lose error interrupt information when servicing a HARD_ERROR related error.

Figure 10-3 shows the bit assignments for the SSR address as they appear on the pins of the microprocessor, and Table 10-7 shows the bit format of the SSR.

HEX	1	0				4				4				8				0				0				0			
Binary	2 ²⁸	2 ²⁷	2 ²⁶	2 ²⁵	2 ²⁴	2 ²³	2 ²²	2 ²¹	2 ²⁰	2 ¹⁹	2 ¹⁸	2 ¹⁷	2 ¹⁶	2 ¹⁵	2 ¹⁴	2 ¹³	2 ¹²	2 ¹¹	2 ¹⁰	2 ⁹	2 ⁸	2 ⁷	2 ⁶	2 ⁵	2 ⁴	2 ³	2 ²	2 ¹	2 ⁰
	1	x	x	x	x	x	1	0	0	0	1	0	0	1	0	x	x	x	x	x	x	x	x	x	x	x	x	x	x

NOTE: x = Don't Care

Figure 10-3. System Status Register Address Bit Format

Table 10-7. SSR Format

Bits	Name
0	Network buffer parity error
1	Network packet error
2	Message queue oversubscribed
3	Prefetch queue overrun
4	Prefetch queue underrun
5	Virtual PE number range error

Table 10-7. SSR Format (continued)

Bits	Name
6	8AM option Revision 7 – processing element node 8AI option Revision 8 – I/O gateway
7	Illegal register access
8	DTB annex error
9	Illegal write to protected low memory (PLM) – I/O gateway PEs only
63 – 10	These bits are not used

Network Buffer Parity Error
Bit 0

When set to 1, this bit indicates that a parity error has occurred in one of the input buffers in the network interface. Parity is checked on each byte of the packet phits (except the routing tag, destination, and data phits) as they are transferred from the network interface logic to the PE support circuitry. When a network buffer parity error occurs, the support circuitry sets the network buffer parity error bit to 1 and sets the error interrupt.

Network Packet Error
Bit 1

When set to 1, this bit indicates that a network packet error occurred. There are two types of network packet errors: a misrouted packet error and a packet parity error. If the network interface receives a packet and the value of the destination phit does not match the value stored in the network interface source (X_WHOAMI) register, a misrouted packet error occurs. When this occurs, the network interface sets the network packet error bit to 1, sets the error interrupt, converts the packet into an error message, and sends the error message to the appropriate PE in the node.

Likewise, the network interface detects that the bits of a packet header phit changed value after the source PE created them, a packet parity error occurs. When this occurs, the network interface sets the network packet error bit to 1, sets the error interrupt, converts the packet into an error message, and sends the error message to the appropriate PE in the node.

After receiving an error message, the support circuitry in a PE stores the error message in the message queue and sets the message interrupt. After receiving the first error message, the support circuitry discards all subsequent error messages until the SSR register is read to clear the existing error conditions.

Message Queue Oversubscribed
Bit 2

When set to 1, this bit indicates that the microprocessor started a message write when the value of the message queue limit counter was less than zero. When this occurs, the support circuitry sets the message queue oversubscribed bit to 1 and sets the error interrupt. Even though the message queue oversubscribed indication occurred, the support circuitry completes the message write operation.

If a message queue oversubscribed indication occurs, the operating system should immediately read the messages stored in the message queue. This action provides more space in the message queue.

Prefetch Queue Overrun
Bit 3

When set to 1, this bit indicates that a prefetch queue overrun error occurred. A prefetch queue overrun error occurs if the microprocessor has issued 16 prefetch requests without reading data from the prefetch queue and subsequently issues another prefetch request. When this occurs, the support circuitry sets the prefetch queue overrun bit to 1 and sets the error interrupt.

If a prefetch queue overrun indication occurs, the operating system should immediately read the data stored in the prefetch queue. This action provides more space in the prefetch queue.

Prefetch Queue Underrun
Bit 4

When set to 1, this bit indicates that a prefetch queue underrun error occurred. A prefetch queue underrun error occurs if the prefetch queue is empty and the microprocessor tries to read a word of data from the prefetch queue before issuing a prefetch request. When this occurs, the support circuitry sets the prefetch queue underrun bit to 1 and sets the error interrupt.

Virtual PE Number Range Error
Bit 5

When set to 1, this bit indicates that a virtual PE number range error occurred. A virtual PE number range error occurs if the virtual PE number read from the DTB annex is larger than the inverted value of the mask stored in the virtual PE range mask (VRT_RG) register. When this occurs, the support circuitry sets the virtual PE number range error bit to 1 and sets the error interrupt.

8AM Option Revision 7 – Processing Element Node
8AI Option Revision 8 – I/O Gateway
Bit 6

This bit indicates whether the 8AM option used in the support circuitry of a PE in a processing element node is pre-revision 7 or post-revision 7 (which includes revision 7). This bit is only valid immediately after the message queue tail pointer is reset (writing any value to the message queue tail pointer register (MQ_TP) resets the message queue tail pointer). When set to a 0, bit 6 of the SSR indicates the 8AM option is pre-revision 7. When set to 1, this bit indicates the 8AM option is revision 7 or later.

NOTE: When messaging is being performed, bit 6 of the SSR is undefined for pre-revision 7 8AM options.

This bit indicates whether the 8AI option used in an I/O gateway is pre-revision 8 or post-revision 8 (which includes revision 8). When set to 0, bit 6 of the SSR indicates the 8AI option is a pre-revision 8. When set to 1, this bit indicates the 8AI option is a revision 8 or later.

Illegal Register Access
Bit 7

When set to 1, this bit indicates that an illegal register access error occurred. An illegal register access error occurs if the microprocessor sets bit 28 of the partial physical address to 1 and sets the cycle request pins to indicate an operation other than a BARRIER, READ_BLOCK, or WRITE_BLOCK. When this occurs, the support circuitry sets the illegal register access bit to 1 and sets the error interrupt.

**DTB Annex Error
Bit 8**

When set to 1, this bit indicates that one or both types of DTB annex errors occurred. These errors are a DTB annex write error and an illegal memory function code error.

An illegal DTB annex write error occurs if the microprocessor attempts to write to an entry in the DTB annex that has been disabled by bit 0 or bit 1 of the SCR. When this occurs, the support circuitry sets the DTB annex error bit to 1 and sets the error interrupt.

An illegal memory function code error occurs if the function code read from the DTB annex is set to 3 and the microprocessor provides a READ_BLOCK, WRITE_BLOCK, or FETCH cycle request. When this occurs, the support circuitry sets the DTB annex error bit to 1 and sets the error interrupt. The support circuitry also aborts the code and, if READ_BLOCK was requested, returns undefined data to the microprocessor.

**Illegal Write to PLM
Bit 9**

When set to 1, this bit indicates that an illegal write to protected low memory (PLM) error occurred. An illegal write to PLM error occurs if the local PE, a remote PE, or any BLT tries to write to the lowest 8 Kbytes of local memory in an I/O gateway PE and if bit 10 of the DRAM control register indicates this is not valid. When this occurs, the support circuitry sets the illegal write to PLM bit to 1 and sets the error interrupt.

**User Status Register
Address 10144000₁₆**

The user control and status register (USR) is a 2-bit, general access, read only register. The USR provides the user with the ability to read status of outstanding remote references.

Figure 10-4 shows the bit assignments for the USR address as they appear on the address pins of the microprocessor.

HEX	1	0				1				4				4				0				0				0			
Binary	2 ²⁸	2 ²⁷	2 ²⁶	2 ²⁵	2 ²⁴	2 ²³	2 ²²	2 ²¹	2 ²⁰	2 ¹⁹	2 ¹⁸	2 ¹⁷	2 ¹⁶	2 ¹⁵	2 ¹⁴	2 ¹³	2 ¹²	2 ¹¹	2 ¹⁰	2 ⁹	2 ⁸	2 ⁷	2 ⁶	2 ⁵	2 ⁴	2 ³	2 ²	2 ¹	2 ⁰
	1	x	x	x	x	x	0	0	1	0	1	0	0	0	1	x	x	x	x	x	x	x	x	x	x	x	x	x	x

NOTE: x = Don't Care

Figure 10-4. User Control and Status Register Address Bit Format

Table 10-8 shows the bit format of the UCSR, and the following subsections describe each bit of the register.

Table 10-8. USR Format

Bits	Direction	Name
9 – 0	Not applicable	These bits are not used
10	Read	Remote writes are outstanding
11	Read	Prefetch requests are outstanding
63 – 12	Not applicable	These bits are not used

Remote Writes are Outstanding
Bit 10

When set to 1, this bit indicates that the microprocessor issued one or more remote write requests and the PE has not received all of the corresponding write response packets. When set to 0, this bit indicates that all remote write operations requested by the microprocessor have completed. Outstanding remote writes do not include outstanding message writes but do include fetch-and-increment register writes. Outstanding message writes are managed by the message queue mechanism.

Prefetch Requests Are Outstanding
Bit 11

When set to 1, this bit indicates that the PE issued one or more prefetch requests, however, the PE has not read the corresponding prefetch data from the prefetch queue. When set to 0, this bit indicates that the PE read all of the corresponding words from the prefetch queue.

Invalidate Control Register
Address 10544000₁₆

The invalidate control register (INV_CR) is a 10-bit, write only, system privileged register. The INV_CR provides the user with control of cache line invalidates from external writes.

NOTE: The INV_CR register is only used when the support circuitry contains a revision 7 or higher 8AM option. For more information on determining the revision of the 8AM option in the support circuitry, refer to “System Status Register” in this section.

Figure 10-5 shows the bit assignments for the INV_CR address as it appears on the address pins of the microprocessor.

HEX	1	0				1				4				4				0				0				0			
Binary	2 ²⁸	2 ²⁷	2 ²⁶	2 ²⁵	2 ²⁴	2 ²³	2 ²²	2 ²¹	2 ²⁰	2 ¹⁹	2 ¹⁸	2 ¹⁷	2 ¹⁶	2 ¹⁵	2 ¹⁴	2 ¹³	2 ¹²	2 ¹¹	2 ¹⁰	2 ⁹	2 ⁸	2 ⁷	2 ⁶	2 ⁵	2 ⁴	2 ³	2 ²	2 ¹	2 ⁰
	1	x	x	x	x	x	1	0	1	0	1	0	0	0	1	x	x	x	x	x	x	x	x	x	x	x	x	x	x

NOTE: x = Don't Care

Figure 10-5. Invalidate Control Register Address Bit Format

Table 10-9 shows the bit format of the INV_CR and the following subsections describe each bit of the register.

Table 10-9. INV_CR Format

Bits	Direction	Name
7 – 0	Write	Invalidate filter
8	Write	Enable invalidate filter
9	Write	Enable invalidates
63 – 10	Not applicable	These bits are not used

Invalidate Filter Bits 0 through 7

These bits contain the invalidate filter bits. The invalidate filter bits indicate which one of the 256 cache lines in the data cache memory will be invalidated if a corresponding location in local memory changes value. Bits 9 (enable invalidates bit) and 8 (enable invalidate filter bit) must be set to 1 to enable the invalidate filter.

Enable Invalidate Filter Bit 8

This bit selects whether all remote writes to local memory result in data cache invalidates or whether only a write associated with the invalidate filter address results in a data cache invalidate. When this bit and bit 9 (enable invalidates bit) are set to 1, bits 0 through 7 (invalidate filter bits) are enabled. When bit 8 is set to 0 and bit 9 is set to 1, the invalidate filter bits are not used. When bit 9 is set to 0, bit 8 is not used. Refer to Table 10-10.

Table 10-10. Enable Invalidation and Invalidation Filter Bits

Bits 9 and 8	Description
00	All invalidates disabled
01	All invalidates disabled
10	Invalidate filter bits are not used
11	Bits 0 through 7 (invalidate filter bits) are enabled

Enable Invalidates
Bit 9

When set to 1, this bit enables filtered or nonfiltered invalidate operations. Bit 8 (enable invalidate filter bit) indicates if the invalidate operation is filtered or nonfiltered. Refer again to Table 10-10. More information on the data cache and cache line invalidate operations is provided in Section 3, “Processing Element Node.”

DRAM Control Register
Address 1044C000₁₆

The DRAM control register (DRAM_CR) is a 13-bit, write-only, system privileged register that is used to configure local memory timing parameters and modes of operation. When configuring local memory, multiple values may have to be written sequentially to the DRAM_CR.

Figure 10-6 shows the bit assignments for the DRAM_CR address as they appear on the address pins of the microprocessor.

HEX	1	0				4				4				C				0				0				0			
Binary	2 ²⁸	2 ²⁷	2 ²⁶	2 ²⁵	2 ²⁴	2 ²³	2 ²²	2 ²¹	2 ²⁰	2 ¹⁹	2 ¹⁸	2 ¹⁷	2 ¹⁶	2 ¹⁵	2 ¹⁴	2 ¹³	2 ¹²	2 ¹¹	2 ¹⁰	2 ⁹	2 ⁸	2 ⁷	2 ⁶	2 ⁵	2 ⁴	2 ³	2 ²	2 ¹	2 ⁰
	1	x	x	x	x	x	1	0	0	0	1	0	0	1	1	x	x	x	x	x	x	x	x	x	x	x	x	x	x

NOTE: x = Don't Care

Figure 10-6. DRAM Control Register Address Bit Format

Table 10-11 shows the bit format of the DRAM_CR, and the following subsections describe each bit of the register.

Table 10-11. DRAM_CR Format

Bits	Name
4 – 0	Configuration parameter B
7 – 5	These bits are not used
11 – 8	Configuration parameter A
12	Configuration function select
63 – 13	Not used

Setting the overall timing configuration information for the DRAMs takes eight sequential transfers to the DRAM_CR. This operation is only done during system deadstart. The following subsections describe each transfer. The information written to the DRAM_CR during transfer 7 may be updated at any time following the initial configuration sequence.

NOTE: When writing a value into the DRAM_CR register, bit 0 of the system control register should be set to 0.

Also, a refresh reference may hang local memory if the reference issues after a system reset but before refresh references are disabled prior to configuring the DRAM. To avoid this condition, perform the following steps when configuring DRAM following a system reset:

- 1 . Disable refresh references by setting bits 8 through 11 of the barrier timing and refresh (BAR_TMG) register to 0.
- 2 . Wait at least 50 CPs before writing to the DRAM_CR. (This allows a possible refresh in progress to complete.)
- 3 . Configure the DRAM using the eight sequential transfers to the DRAM_CR, as described in the following subsections.
- 4 . Enable refresh references by setting bits 8 through 11 of the BAR_TMG register to the appropriate value. For more information on bits 8 through 11 of the BAR_TMG register, refer to “Barrier Timing and Refresh Register” later in this section.

Transfer 0

The first write to the DRAM_CR resets part of the logic that controls local memory refresh. Figure 10-7 shows the bit format of the DRAM_CR during the first transfer of a configuration sequence.

HEX	1	0				0				0				
Binary	²¹²	²¹¹	²¹⁰	²⁹	²⁸	²⁷	²⁶	²⁵	²⁴	²³	²²	²¹	²⁰	x = Don't Care
	1	0	0	0	0	x	x	x	0	0	0	0	0	

Figure 10-7. Transfer 0

Transfer 1

The second write to the DRAM_CR resets the remaining logic that controls local memory refresh. Figure 10-8 shows the bit format of the DRAM_CR during the second transfer of a configuration sequence.

HEX	1	1				0				0				
Binary	²¹²	²¹¹	²¹⁰	²⁹	²⁸	²⁷	²⁶	²⁵	²⁴	²³	²²	²¹	²⁰	x = Don't Care
	1	0	0	0	1	x	x	x	0	0	0	0	0	

Figure 10-8. Transfer 1

Transfer 2

The third write to the DRAM_CR indicates that this is the start of the configuration sequence. Figure 10-9 shows the bit format of the DRAM_CR during the third transfer of a configuration sequence.

HEX	1	0				0				0				
Binary	²¹²	²¹¹	²¹⁰	²⁹	²⁸	²⁷	²⁶	²⁵	²⁴	²³	²²	²¹	²⁰	x = Don't Care
	1	0	0	0	0	x	x	x	0	0	0	0	0	

Figure 10-9. Transfer 2

Transfer 3

The fourth write to the DRAM_CR loads the row address strobe precharge timing value (tRAS) in the DRAM timing parameters. Figure 10-10 shows the bit format of the DRAM_CR during the fourth transfer of a configuration sequence. Table 10-12 lists the configuration parameter b values and the corresponding tRAS values in clock periods (CPs). For example, if configuration parameter b is set to 01100 during the fourth transfer, the value of the tRAS is 8 CPs.

NOTE: The value of tRAS must always be smaller than the value of tNEND, which is set in transfer 5.

HEX	1	8				b				b				
Binary	2 ¹²	2 ¹¹	2 ¹⁰	2 ⁹	2 ⁸	2 ⁷	2 ⁶	2 ⁵	2 ⁴	2 ³	2 ²	2 ¹	2 ⁰	x = Don't Care b = Configuration Parameter B
	1	1	0	0	0	x	x	x	b	b	b	b	b	

Figure 10-10. Transfers 3 through 6

Transfer 4

The fifth write to the DRAM_CR loads the same-page end timing (tSEND) value in the DRAM timing parameters. Figure 10-10 also shows the bit format of the DRAM_CR during the fifth transfer of a configuration sequence. Table 10-12 lists the configuration parameter b values and the corresponding tSEND values in CPs. When set to 0, bit 4 of the tSEND configuration parameter b value indicates the timing is for a PEM. When set to 1, bit 4 of the tSEND configuration parameter b value indicates the timing is for an IOM.

In a PEM, the tSEND configuration parameter b value and the column address strobe pulse width timing (tCAS) configuration parameter b value must be identical. In an IOM, the configuration parameter b value of the tSEND should always be set to 10100 to select a 10-CP value (this actually sets the CAS cycle timing value in an IOM).

Table 10-12. Timing Information Transfers

Configuration Parameter B	tRAS Transfer 3	tSEND Transfer 4	tNEND Transfer 5		tCAS Transfer 6	
			Writes	Reads	Writes	Reads
00000	16	6	6	5	0	1
00001	1	7	7	6	1	2
00011	2	8	8	7	2	3
00010	3	9	9	8	3	4
00100	4	10	10	9	4	5
00101	5	11	11	10	5	6
00111	6	12	12	11	6	7
00110	7	13	13	12	7	8
01100	8	14	14	13	8	9
01101	9	15	15	14	9	10
01111	10	16	16	15	10	11
01110	11	17	17	16	11	12
01000	12	18	18	17	12	13
01001	13	19	19	18	13	14
01011	14	20	20	19	14	15
01010	15	21	21	20	15	16
10000	16	6	22	21	0	1
10001	1	7	23	22	1	2
10011	2	8	24	23	2	3
10010	3	9	25	24	3	4
10100	4	10	26	25	4	5
10101	5	11	27	26	5	6
10111	6	12	28	27	6	7
10110	7	13	29	28	7	8
11100	8	14	30	29	8	9
11101	9	15	31	30	9	10
11111	10	16	32	31	10	11
11110	11	17	33	32	11	12
11000	12	18	34	33	12	13

Table 10-12. Timing Information Transfers (continued)

Configuration Parameter B	tRAS Transfer 3	tSEND Transfer 4	tNEND Transfer 5		tCAS Transfer 6	
			Writes	Reads	Writes	Reads
11001	13	19	35	34	13	14
11011	14	20	36	35	14	15
11010	15	21	37	36	15	16

Transfer 5

The sixth write to the DRAM_CR loads the new-page end timing (tNEND) value in the DRAM timing parameters. Figure 10-10 also shows the bit format of the DRAM_CR during the sixth transfer of a configuration sequence. Table 10-12 lists the configuration parameter b values and the corresponding tNEND values in CPs. For example, if configuration parameter b is set to 10010 during the sixth transfer, the value of the tNEND is 25 CPs for writes and 24 CPs for reads.

Transfer 6

The seventh write to the DRAM_CR loads the column address strobe pulse width timing (tCAS) value in the DRAM timing parameters. Figure 10-10 also shows the bit format of the DRAM_CR during the seventh transfer of a configuration sequence. Table 10-12 lists the configuration parameter b values and the corresponding tCAS values in CPs. For example, if configuration parameter b value is set to 00100 during the seventh transfer, the value of the tCAS is 4 CPs for writes and 5 CPs for reads.

In a PEM, the configuration parameter b value of tCAS must be set to the same configuration parameter b value as tSEND. Bits 0 through 12 of the DRAM_CR must be set to the same value during transfer 4 and transfer 6.

In an IOM, the value of tCAS should always be set to a 4 clock pulse value for writes and 5 clock pulse value for reads (configuration parameter b set to 00100).

Transfer 7

This eighth transfer to the DRAM_CR loads the detailed timing and memory configuration values in the DRAM timing parameters. This information may be updated at any time following the initial configuration sequence. The information is updated by writing to the DRAM_CR once. Figure 10-11 shows the bit format of the DRAM_CR during the eighth transfer of a configuration sequence or when the DRAM_CR information is updated.

HEX	0	ra/p				be				r/c			
Binary	2 ¹²	2 ¹¹	2 ¹⁰	2 ⁹	2 ⁸	2 ⁷	2 ⁶	2 ⁵	2 ⁴	2 ³	2 ²	2 ¹	2 ⁰
	0	ra	p	x	x	x	x	x	be	r	r	c	c

NOTE: x = Don't Care
 c = Column Address stobe (CAS) Cycle Timing
 r = DRAM Readout Sample Timing
 be = DRAM Bus Enable Timing
 p = Disable Writes to Protected Low Memory — I/O Gateway PEs Only
 ra = Enable Read-ahead Memory Operation for Instruction Fetches

Figure 10-11. Transfer 7

In a PEM node, bits 0 and 1 of the DRAM_CR define the CAS cycle timing. (In an IOM node, the value of tSEND during transfer 4 sets the value of the CAS cycle timing.) Table 10-13 shows the values of the CAS cycle timing bits and lists the corresponding timing in clock pulses.

Table 10-13. CAS Cycle Timing Bit Values

Bits 1 and 0	CAS Cycle Timing
00	7 clock pulses
01	8 clock pulses
10	9 clock pulses
11	10 clock pulses

In a PEM node, bits 2 and 3 of the DRAM_CR define the DRAM readout sample timing. (In an IOM node, this value is fixed at five clock pulses.) Table 10-14 shows the values of the DRAM readout sample timing bits and lists the corresponding timing relative to the beginning of the column address strobe.

Bit 4 of the DRAM_CR sets the DRAM bus enable timing for the presentation of write data to the DRAM data bus. When set to 0, bit 4 sets the DRAM bus enable timing to 4 CPs. When set to 1, bit 4 sets the DRAM bus enable timing to 5 CPs.

Table 10-14. DRAM Readout Sample Timing Values

Bits 3 and 2	DRAM Readout Sample Timing
00	Clock pulse 3 after CAS
01	Clock pulse 4 after CAS
10	Clock pulse 5 after CAS
11	Clock pulse 6 after CAS

When set to 1 in an I/O gateway PE, bit 2¹⁰ of the DRAM_CR disables the ability to write to protected low memory (PLM), which is the lowest 8 Kbytes in physical memory. If a PE or BLT attempts to write to a disabled PLM, the support circuitry sets the illegal write to protected low memory bit in the SSR and sets the error interrupt. When set to 0, bit 10 of the DRAM_CR enables the ability to write to PLM.

When set to 1, bit 11 of the DRAM_CR signals the support circuitry in a PE to anticipate subsequent instruction fetch requests. When the support circuitry anticipates an instruction fetch, the support circuitry reads the next consecutive cache line location in memory and stores the information in the microprocessor read staging register of the support circuitry. When set to 0, bit 11 of the DRAM_CR signals the support circuitry in a PE not to anticipate subsequent instruction fetch requests. If this instruction fetch read-ahead operation and the data read-ahead operation (memory function code 6) are both enabled, system performance may be negatively affected.

Barrier Timing and Refresh Register

Address 10404000₁₆

The barrier timing and refresh (BAR_TMGM) register is a 12-bit, write-only, system privileged register. The BAR_TMGM register controls the timing of each physical barrier synchronization circuit and the timing of the DRAM refresh references.

Figure 10-12 shows the bit assignments for the BAR_TMGM register address as they appear on the address pins of the microprocessor.

HEX	1	0				4				0				4				0				0							
Binary	2 ²⁸	2 ²⁷	2 ²⁶	2 ²⁵	2 ²⁴	2 ²³	2 ²²	2 ²¹	2 ²⁰	2 ¹⁹	2 ¹⁸	2 ¹⁷	2 ¹⁶	2 ¹⁵	2 ¹⁴	2 ¹³	2 ¹²	2 ¹¹	2 ¹⁰	2 ⁹	2 ⁸	2 ⁷	2 ⁶	2 ⁵	2 ⁴	2 ³	2 ²	2 ¹	2 ⁰
	1	x	x	x	x	x	1	0	0	0	0	0	0	0	1	x	x	x	x	x	x	x	x	x	x	x	x	x	x

NOTE: x = Don't Care

Figure 10-12. Barrier Timing and Refresh Register Address Bit Format

Table 10-15 shows the bit format of the BAR_TMGM register, and the following paragraphs describe the function of bits 8 through 11 of the BAR_TMGM register. More information on bits 0 through 7 of the BAR_TMGM register is provided in Section 9, “Barrier Synchronization.”

Table 10-15. BAR_TMGM Bit Format

Bits	Name
1 – 0	These bits control the timing for physical barrier synchronization circuit 0.
3 – 2	These bits control the timing for physical barrier synchronization circuit 1.
5 – 4	These bits control the timing for physical barrier synchronization circuit 2.
7 – 6	These bits control the timing for physical barrier synchronization circuit 3.
11 – 8	These bits control the DRAM refresh timing.
63 – 12	These bits are not used.

Bits 8 through 11 define the number of CPs that occur between successive DRAM refresh references. The maximum number of CPs that can occur depends on the system clock period.

Table 10-16 lists the values for bits 8 through 11 of the BAR_TMG register and shows the corresponding timing interval for each value. To set the correct value for bits 8 through 11, select the system clock period in Table 10-16 that is closest to and greater than the actual system clock period. For example, if the system clock period is 6.6 ns, bits 8 through 11 should be set to 0011 (6.8-ns clock period).

Table 10-16. DRAM Refresh Timing Values

BAR_TMG Bits 11 – 8	Timing Interval in CPs	System Clock Period
0000	No refresh	Not applicable
0001	1088	7.2 ns
0011	1152	6.8 ns
0010	1216	6.5 ns
0100	1280	6.3 ns
0101	1344	5.9 ns
0111	1408	5.6 ns
0110	1472	5.4 ns
1100	1536	5.1 ns
1101	1600	4.9 ns
1111	1664	4.7 ns
1110	1728	4.6 ns
1000	1792	4.4 ns
1001	1856	4.3 ns
1011	1920	4.1 ns
1010	1984	4.0 ns

If bits 8 through 11 are set to 0, the timer is reset and refresh references are disabled. The timer can be disabled for testing and simulation.

Network Enable Register

Address **106C0000₁₆**

The network enable (NET_ENA) register is a 9-bit, write-only, system privileged register. The NET_ENA register enables or disables each of the inputs to the X-dimension switch, the Y-dimension switch, and the Z-dimension switch in the network interface.

Figure 10-13 shows the bit assignments for the NET_ENA address as they appear on the address pins of the microprocessor.

HEX	1	0				6				C				0				0				0							
Binary	2 ²⁸	2 ²⁷	2 ²⁶	2 ²⁵	2 ²⁴	2 ²³	2 ²²	2 ²¹	2 ²⁰	2 ¹⁹	2 ¹⁸	2 ¹⁷	2 ¹⁶	2 ¹⁵	2 ¹⁴	2 ¹³	2 ¹²	2 ¹¹	2 ¹⁰	2 ⁹	2 ⁸	2 ⁷	2 ⁶	2 ⁵	2 ⁴	2 ³	2 ²	2 ¹	2 ⁰
	1	x	x	x	x	x	1	1	0	1	1	0	0	x	x	x	x	x	x	x	x	x	x	x	x	x	x	x	x

NOTE: x = Don't Care

Figure 10-13. Network Enable Register Address Bit Format

Table 10-17 shows the bit format of the NET_ENA register. Bits 0 through 2 control the inputs to the X-dimension switch, bits 3 through 5 control the inputs to the Y-dimension switch, and bits 6 through 8 control the inputs to the Z-dimension switch. For example, if bit 2 of the NET_ENA register is set to 1, the input from the -X communication link to the X-dimension switch is enabled. If bit 2 of the NET_ENA register is set to 0, this input is disabled.

Table 10-17. NET_ENA Bit Format

Bits	Name
0	Enable PE and BLT input to X switch
1	Enable +X input to X switch
2	Enable -X input to X switch
3	Enable X-dimension input to Y switch
4	Enable +Y input to Y switch
5	Enable -Y input to Y switch
6	Enable Y-dimension input to Z switch
7	Enable +Z input to Z switch
8	Enable -Z input to Z switch
63 - 9	Not used

Network Performance Register**Address 106D0000₁₆**

The network performance (NET_PFM) register is an 8-bit, write-only, system privileged register. The NET_PFM controls which parameters of network performance are monitored.

Figure 10-14 shows the bit assignments for the NET_PFM address as they appear on the address pins of the microprocessor.

HEX	1	0				6				D				0				0				0							
Binary	2 ²⁸	2 ²⁷	2 ²⁶	2 ²⁵	2 ²⁴	2 ²³	2 ²²	2 ²¹	2 ²⁰	2 ¹⁹	2 ¹⁸	2 ¹⁷	2 ¹⁶	2 ¹⁵	2 ¹⁴	2 ¹³	2 ¹²	2 ¹¹	2 ¹⁰	2 ⁹	2 ⁸	2 ⁷	2 ⁶	2 ⁵	2 ⁴	2 ³	2 ²	2 ¹	2 ⁰
	1	x	x	x	x	x	1	1	0	1	1	0	1	x	x	x	x	x	x	x	x	x	x	x	x	x	x	x	x

NOTE: x = Don't Care

Figure 10-14. Network Performance Register Address Bit Format

Table 10-18 shows the bit format of the NET_PFM register, and the following subsections describe each bit of the register.

Table 10-18. NET_PFM Bit Format

Bits	Name
3 – 0	Parameter select
5 – 4	PE 0 switch or BLT select
7 – 6	PE 1 switch or BLT select
63 – 8	Not used

Parameter Select Bits 0 through 3

The Parameter Select bits select one of 15 network switch performance parameters. Table 10-19 shows the value of the parameter select bits and lists the corresponding performance parameters. For example, when the parameter select bits are set to 0001, the parameter selected is equal to the

number of references to virtual channel 1 (VC 1) from a PE or from the BLT in the X-dimension switch, or it is equal to the number of references to VC 1 from the previous dimension in the Y- or Z-dimension switch.

NOTE: The performance parameter is sent to the perf_cnt_h (1) pin of the microprocessor. The perf_cnt_h (0) pin of the microprocessor is not used.

Table 10-19. Network Switch Performance Parameters

Parameter Select Bits 3 – 0	Parameter
0000	Number of references to VC 0 from PE and BLT or previous dimension
0001	Number of references to VC 1 from PE and BLT or previous dimension
0010	Number of references to VC 2 from PE and BLT or previous dimension
0011	Number of references to VC 3 from PE and BLT or previous dimension
0100	Number of references to VC 0 from +dimension
0101	Number of references to VC 1 from +dimension
0110	Number of references to VC 2 from +dimension
0111	Number of references to VC 3 from +dimension
1000	Number of references to VC 0 from –dimension
1001	Number of references to VC 1 from –dimension
1010	Number of references to VC 2 from –dimension
1011	Number of references to VC 3 from –dimension
1100	Number of conflicts for the input from a PE and BLT or previous dimension
1101	Number of conflicts for the +dimension input
1110	Number of conflicts for the –dimension input
1111	This value is not used (forced to 1)

PE 0 Switch or BLT Select
Bits 4 and 5

The PE 0 Switch or BLT Select bits can be modified only by the microprocessor in PE 0. These bits select which network switch to monitor or select to monitor BLT conflicts (refer to Table 10-20).

Table 10-20. Switch or BLT Select Bits

Bits 5 and 4	Component to Monitor
00	X-dimension switch
01	Y-dimension switch
10	Z-dimension switch
11	BLT conflicts

PE 1 Switch or BLT Select
Bits 6 and 7

These bits function identically to bits 4 and 5, but can be modified only by the microprocessor in PE 1 (refer again to Table 10-20).

Network Mode Register
Address $106E0000_{16}$

The network mode (NODE_CSR) register is a 14-bit, write-only, system privileged register. The NODE_CSR enables or disables several types of error checking in the network interface and the BLT.

Figure 10-15 shows the bit assignments for the NODE_CSR address as they appear on the address pins of the microprocessor.

HEX	1	0				6				E				0				0				0				0			
Binary	2^{28}	2^{27}	2^{26}	2^{25}	2^{24}	2^{23}	2^{22}	2^{21}	2^{20}	2^{19}	2^{18}	2^{17}	2^{16}	2^{15}	2^{14}	2^{13}	2^{12}	2^{11}	2^{10}	2^9	2^8	2^7	2^6	2^5	2^4	2^3	2^2	2^1	2^0
	1	x	x	x	x	x	1	1	0	1	1	1	0	x	x	x	x	x	x	x	x	x	x	x	x	x	x	x	x

NOTE: x = Don't Care

Figure 10-15. Network Mode Register Address Bit Format

Table 10-21 shows the bit format of the NODE_CSR register, and the following subsections describe each bit of the register.

Table 10-21. NODE_CSR Bit Format

Bits	Description
0	Enable packet errors
1	Enable BLT remote index errors
3 – 2	Set incoming channel parity values
5 – 4	Set outgoing request channel parity values
6	Enable BLT address offset range errors
7	Enable BLT virtual PE range errors
10 – 8	Set outgoing response channel parity values
13 – 11	Barrier synchronization circuitry partition bits
63 – 14	These bits are not used

Enable Packet Errors Bit 0

When set to 1, this bit enables the detection of network packet errors. A network packet error occurs if the network interface receives a misrouted packet or if the network interface detects a parity error in the header of a packet. When set to 0, this bit disables the detection of network packet errors.

Enable BLT Remote Index Errors Bit 1

When set to 1, this bit enables error checking by the BLT of remote index values read from local memory during a BLT scatter or gather operation. When set to 0, this bit disables remote index error checking.

Set Incoming Channel Parity Values Bits 2 and 3

These bits control the value of the parity bits that are generated by the network interface when the network interface reads information out of the input buffers and sends the information to a PE. When set to 1, bit 3

indicates that both parity bits (one for bits 8 through 12 and one for bits 0 through 7) will be generated by the parity circuitry in the network interface. In this case, bit 2 is not used.

When set to 0, bit 3 indicates that both parity bits will be replaced by the value of bit 2. These bits provide the ability to force parity errors when information transfers from the input buffers in the network interface to a PE.

Set Outgoing Request Channel Parity Values Bits 4 and 5

These bits control the value of the parity bits that are generated by the network interface when the network interface creates outgoing request packets. When set to 1, bit 5 indicates that both parity bits (one for bits 8 through 12 and one for bits 0 through 7) will be generated by the parity circuitry in the network interface. In this case, bit 4 is not used.

When set to 0, bit 5 indicates that both parity bits will be replaced by the value of bit 4. These bits provide the ability to force parity errors when the network interface creates request packets.

Enable BLT Address Offset Range Errors Bit 6

When set to 1, this bit enables error checking by the BLT for address offset range errors. An address offset range error occurs if the address offset obtained from the remote index is larger than the value stored in the BLT remote limit register (BLT_RLR). When set to 0, this bit disables address offset range error checking by the BLT.

Enable BLT Virtual PE Range Errors Bit 7

When set to 1, this bit enables error checking by the BLT for virtual PE range errors. A virtual PE range error occurs if the virtual PE number obtained from the remote index is larger than the value allowed by the PE range mask stored in the BLT control register. When set to 0, this bit disables virtual PE range error checking by the BLT.

Set Outgoing Response Channel Parity Values
 Bits 8 through 10

These bits control the value of the parity bits that are generated by the network interface when the network interface creates outgoing response packets. When set to 1, bit 8 indicates that both parity bits will be generated by the parity circuitry in the network interface. In this case, bits 9 and 10 are not used.

When set to 0, bit 8 indicates that the parity bit for bits 0 through 7 will be replaced by the value of bit 9 and that the parity bit for bits 8 through 12 will be replaced by the value of bit 10. These bits provide the ability to force parity errors when the network interface creates outgoing response packets.

Barrier Synchronization Circuitry Partition Bits
 Bits 11 through 13

These bits are used to partition the barrier synchronization circuitry into smaller circuits. For more information on these bits, refer again to Section 9, “Barrier Synchronization.”

No-operation Register
 Address 10140000₁₆

The no-operation register (NOR) is a readable and writable general access register. The NOR may be used by software for debugging purposes. When data is written to or read from the NOR, there is no effect on the hardware.

Figure 10-16 shows the bit assignments for the NOR address as they appear on the address pins of the microprocessor.

HEX	1	0				1	4				0				0				0										
Binary	2 ²⁸ 1	2 ²⁷ x	2 ²⁶ x	2 ²⁵ x	2 ²⁴ x	2 ²³ x	2 ²² x	2 ²¹ 0	2 ²⁰ 1	2 ¹⁹ 0	2 ¹⁸ 1	2 ¹⁷ 0	2 ¹⁶ 0	2 ¹⁵ 0	2 ¹⁴ 0	2 ¹³ x	2 ¹² x	2 ¹¹ x	2 ¹⁰ x	2 ⁹ x	2 ⁸ x	2 ⁷ x	2 ⁶ x	2 ⁵ x	2 ⁴ x	2 ³ x	2 ² x	2 ¹ x	2 ⁰ x

NOTE: x = Don't Care

Figure 10-16. Invalidate Control Register Address Bit Format

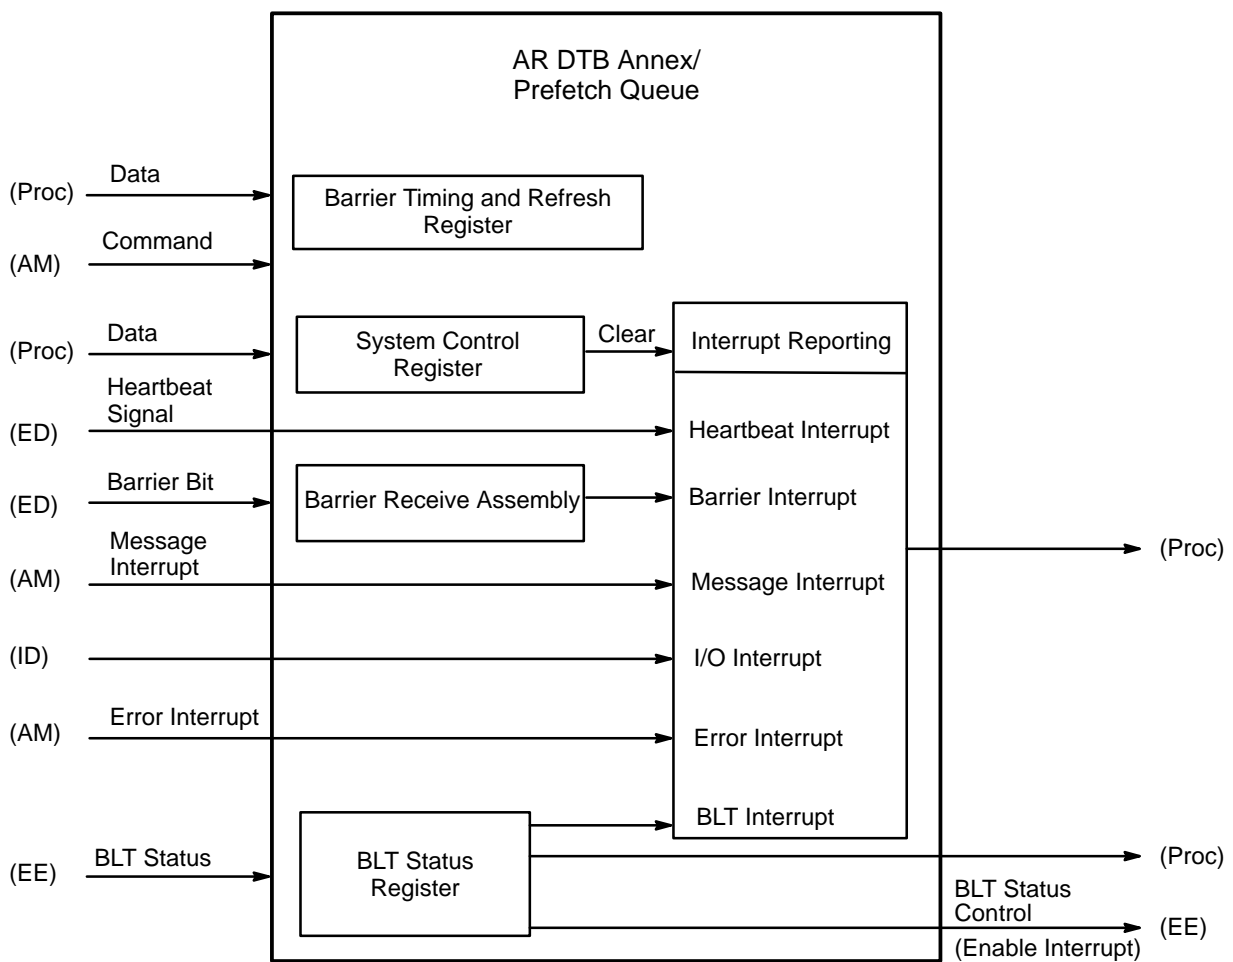
When data is read from the NOR, the data reflects the current state of the system status register (SSR); however, the state of the hardware is not affected when the NOR is read. More information on the SSR is provided in this section.

Hardware Description

The control and status circuitry is contained in the AR option, the AM option, and the EE option.

AR Option

The AR option contains the system control register and the barrier timing and refresh register. Refer to Figure 10-17. The AR option also sets all of the hardware interrupts.



A-11728

Figure 10-17. AR Option

System Control Register

The system control register controls the privileged mode resources of a PE. The following subsections describe the hardware associated with this register.

DTB Annex Write Enable Bits 0 and 1

The AR option uses bit 0 and bit 1 of the system control register as the write enable signal for the DTB annex. The AR option also sends these bits to the AM option.

The AM uses bits 0 and 1 to determine whether the PE number from the DTB annex is virtual. When the bit 0 is set to 1, the AM option interprets the PE number from locations 1 through 15 of the DTB annex as virtual. When the bit 1 is set to 1, the AM option interprets the PE number from locations 16 through 31 of the DTB annex as virtual.

The AM option also uses bits 0 and 1 of the system control register to disable writes to the DTB annex. When the microprocessor attempts to write to locations 1 through 15 of the DTB annex and bit 0 is set to 0 (write disabled), the AM option aborts the request and signals the AR option to set the error interrupt. The AM option also aborts the request and signals the AR option to set the error interrupt when the microprocessor attempts to write to locations 16 through 31 of the DTB annex and bit 1 is set to 0.

Clear Heartbeat Interrupt Bit 2

The AR option uses bit 2 of the system control register to clear the heartbeat interrupt (set it to 0).

Enable BSR0 Interrupt Bit 3

The AR option uses bit 3 of the system control register to enable the barrier interrupts from BSR0 and BSR1. When set to 1, interrupts from BSR0 and BSR1 are enabled. When set to 0, only the barrier interrupt from BSR1 is enabled.

Selective Reset Mask
Bit 4

The AR option uses bit 4 of the system control register to disable the PE reset. When set to 1, the AR option does not reset the PE support circuitry.

Node Enable Vote
Bit 5

The AR option sends bit 5 of the system control register to the EE option. The EE option ANDs the node enable vote bit from both PEs with bit 0 of the network enable register (Enable PE and BLT input to X switch). The result of this AND is sent to the X-dimension SR option. When the result of the AND is a 1, the PE and BLT input to the SR option is enabled. When the result is a 0, the PE and BLT input to the SR option is disabled.

Barrier Timing and Refresh Register

The barrier timing and refresh register controls the timing of each physical barrier synchronization circuit and the timing of the DRAM refresh references. The following subsections describe the hardware associated with this register.

Barrier Timing
Bits 0 through 7

The AR option uses bits 0 through 7 of the barrier timing and refresh register to adjust the transfer sequence of the barrier bits. The barrier timing bits are added to a 2-bit sync counter to generate an offset. This offset selects the appropriate barrier bits to be transferred. For example, when bits 0 and 1 of the barrier timing register are set to 00, barrier register bit 0 is transferred in CP 0 and bits 4, 8, and 12 are transferred in succession in the CPs that follow. Refer to Figure 10-18. When barrier timing register bits 0 and 1 are set to 10, barrier register bit 8 is transferred in CP 0 and bits 12, 0, and 4 are transferred in succession in the CPs that follow.

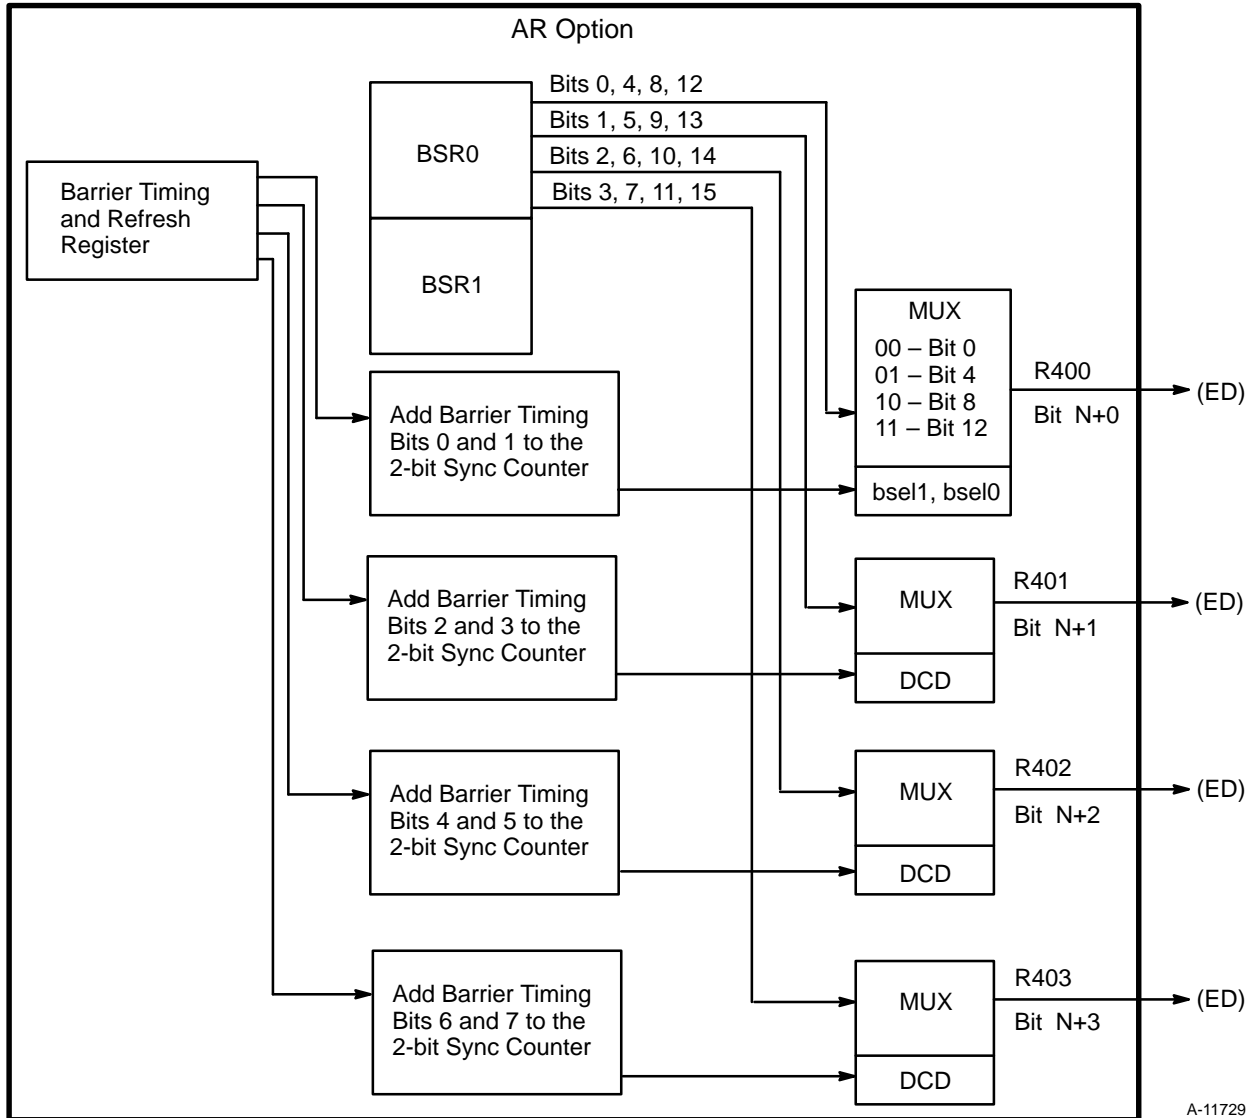


Figure 10-18. Barrier Timing Circuitry

Refresh Timing
Bits 8 through 11

The AR option uses the refresh timing bits 8 through 11 to generate the refresh at the proper time. A refresh reference is made at least once every 1,088 CPs. Using the refresh timing parameters, the AR option can change this interval from 1,088 to as many as 1,984 CPs. The interval is increased in 64-CP increments, as shown in Table 10-16.

Hardware Interrupts

The AR option also sets all of the hardware interrupts after receiving signals from various options within the PE support circuitry. These interrupts include the heartbeat interrupt, the barrier interrupt, the message interrupt, the BLT interrupt, the I/O interrupt, and the error interrupt.

Heartbeat Interrupt

The AR option sets the heartbeat interrupt after it receives the heartbeat signal from the ED option. The heartbeat interrupt remains set until the microprocessor sets the clear heartbeat signal in the system control register.

NOTE: The heartbeat signal is generated by the system circuitry and is fanned out to the modules within the chassis in the same manner as the barrier bits.

Barrier Interrupt

The AR option sets the barrier interrupt when it receives a 1 to 0 transition on a barrier result input and the barrier interrupt is enabled. The AR option clears the barrier interrupt after the microprocessor reads the contents of the barrier synchronization mask and interrupt register.

BLT Interrupt

The AR option sets the BLT interrupt when it receives a signal from the EE option that indicates that the block transfer is complete or that there is a BLT error. The microprocessor reads the BLT status register (also located in the AR option) to determine the cause of the BLT interrupt. The AR option clears the BLT interrupt after the microprocessor writes to the BLT control register.

I/O Interrupt

The AR options that reside on the I/O modules handle the I/O interrupt. For more information about the I/O interrupt, refer to Section 11, “LOSP and HISP.”

Error Interrupt

The AR option sets the error interrupt after receiving notification from the AM option that an error occurred. For a list of error conditions, refer to the “Hardware Interrupt Pins” subsection at the beginning of this section. The AR option clears the error interrupt after the microprocessor reads the contents of the system status register.

AM Option

The AM option contains the system status register, the user status register, the invalidate register, and the DRAM control register. The microprocessor loads the invalidate control register and the DRAM control register with parameter information using the AR option. Refer to Figure 10-19. The microprocessor also uses the AR option to read the information from the system status register and the user status register.

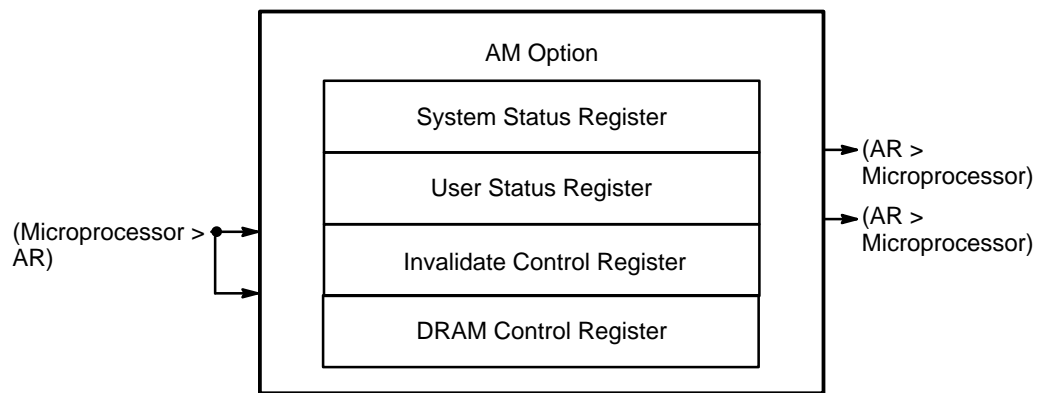


Figure 10-19. AM Option

System Status Register

The system status register contains the status of PE error conditions related to the error interrupt and the HARD_ERROR cycle acknowledge. When an error interrupt occurs, the microprocessor must read this register to determine the cause of the error. The following subsections describe the hardware associated with the system status register.

Network Buffer Parity Error
Bit 0

The AM option receives notification of a network buffer parity error from the ED option. Upon receiving this notification, the AM option sets bit 0 to a 1. When set to 1, bit 0 enables the AM option to generate control for writing to the message queue. The corrupted packet will be written into the message queue so that software can determine the cause of the error. The AM option also signals the AR option to set the message and error interrupts.

Network Packet Error
Bit 1

The AM option receives notification of a network packet error from the ED option. Upon receiving this notification, the AM option sets bit 1 to a 1. When set to 1, bit 1 enables the AM option to generate control for writing to the message queue. The corrupted packet will be written into the message queue so that software can determine the cause of the error. The AM option also signals the AR option to set the message and error interrupts.

Message Queue Oversubscribed
Bit 2

The AM option sets bit 2 of the system status register to a 1 when the message queue is full and it receives a message write request from the microprocessor. The oversubscribed condition exists because the AM option cannot reserve a location in the message queue for the outgoing message. The AM option needs to reserve this location in case the outgoing message is rejected by the destination PE. To notify the microprocessor of an error, the AM option also signals the AR option to set the error interrupt.

Prefetch Queue Overrun
Bit 3

The AM option sets bit 3 of the system status register to a 1 when the prefetch queue is full and the microprocessor issues another prefetch queue request. The overrun condition exists because there is no available space in the prefetch queue for the prefetch response data. To notify the microprocessor of an error, the AM option signals the AR option to set the error interrupt.

Prefetch Queue Underrun
Bit 4

The AM option sets bit 4 of the system status register to a 1 when the prefetch queue is empty and the microprocessor tries to read data from the prefetch queue. The underrun condition exists because there is no data in the prefetch queue to read. To notify the microprocessor of an error, the AM option also signals the AR option to set the error interrupt.

Virtual PE Number Range Error
Bit 5

The AM option sets bit 5 of the system status register to a 1 when the virtual PE number from the DTB annex does not fall within the range specified by the virtual PE range mask register. When the virtual PE number from the DTB annex does not fall within this range, the virtual PE number is not a legal PE number for the partition. To notify the microprocessor of an error, the AM option signals the AR option to set the error interrupt.

Illegal Register Access
Bit 7

The AM option sets bit 7 of the system status register to a 1 when the microprocessor issues a FETCH, FETCH_M, LDx_L, or a STx_C instruction and bit 28 of the partial physical address is set to a 1 (bit 28 set to 1 indicates a register request). The AM option notifies the microprocessor of the error by signaling the AR option to set the error interrupt.

DTB Annex Error
Bit 8

The AM option sets bit 8 of the system status register to a 1 when the microprocessor requests a write to a location in the DTB annex that is disabled or when the memory function code is set to 3 (011) and the microprocessor issues a READ_BLOCK, WRITE_BLOCK, or FETCH request. When a DTB annex error occurs, the AM option also signals the AR option to set the error interrupt.

**Illegal Write to PLM
Bit 9**

The AM option sets bit 9 of the system status register to a 1 when a local or remote memory request is for the PLM (low 8 Kbytes of local memory in an I/O gateway). The AM option notifies the microprocessor of an error by signaling the AR option to set the error interrupt.

User Status Register

The user status register provides the user with the ability to read the status of outstanding remote references. The following subsections describe the hardware associated with the user status register.

**Remote Writes are Outstanding
Bit 10**

The AM option sets bit 10 of the user control and status register when the outstanding write request counter is greater than 0. This indicates to the microprocessor that the remote write requests are not complete.

**Prefetch Requests Are Outstanding
Bit 11**

The AM option sets bit 11 of the user control and status register when the prefetch queue contains data. The following subsections describe the hardware associated with the invalidate control register.

Invalidate Control Register

The invalidate control register provides the user with control of cache line invalidates from external writes.

**Invalidate Filter
Bits 0 through 7**

The AM option compares bit 0 through 7 of the user control and status register to the address offset bits 5 through 12. When they are equal, the AM option sends the address to the AR option. The AR option then sends the address to the microprocessor using the iadr_h pins. This address indicates to the microprocessor which address to invalidate in the data cache.

Enable Invalidate Filter
Bit 8

When bit 8 of the user control and status register is set to a 1, the AM option is enabled to compare the invalidate filter bits and the address offset bits.

Enable Invalidates
Bit 9

When bit 9 of the user control and status register is set to a 1, the AM option is enabled to perform invalidation operations. The AM option signals the AR option that the invalidate operations are enabled. The AR option then signals the microprocessor of the invalidation using the `dinvreq_h` pin.

DRAM Control Register

The DRAM control register is used to configure local memory timing parameters and modes of operation. The following subsections describe the hardware associated with the DRAM control register.

Transfer 0 and Transfer 1

During transfer 0 and transfer 1, the DRAM control register is loaded with a bit pattern that instructs the AM option to reset the logic that controls local memory refresh. Refer again to Figure 10-7 and Figure 10-8 for the bit format of the DRAM_CR during the first and second transfer of a configuration sequence.

Transfer 2

During transfer 2, the DRAM control register is loaded with a bit pattern that instructs the AM option that this is the start of the configuration sequence. Refer again to Figure 10-9 for the bit format of the DRAM_CR during the third transfer of a configuration sequence.

Transfer 3
Row Address Strobe Precharge Timing Value
Bits 0 through 4, Bit 12

During transfer 3, bits 0 through 4 of the DRAM control register are loaded with the configuration parameter b value that specifies the row address strobe precharge timing (tRAS). Refer again to Table 10-12. The AM option sends this configuration parameter b value to the AE option. The AM option also sends bit 12 to the AE option. This bit indicates to the AE option that bits 0 through 4 contain configuration parameter b values.

The AE option uses the tRAS value to determine when to assert the RAS signal that is sent to the DRAM. The AE option asserts the RAS signal when the tRAS value equals the value of a bank cycle counter. The bank cycle counter begins incrementing by 1 each CP when the AM option initiates a new memory reference.

Transfer 4
Same-page End Timing Value
Bits 0 through 4, Bit 12

During transfer 4, bits 0 through 4 of the DRAM control register are loaded with the configuration parameter b value that specifies the same-page end timing (tSEND). Refer again to Table 10-12. The AM option sends this configuration parameter b value to the AE option. The AM option also sends bit 12 to the AE option. This bit indicates to the AE option that bits 0 through 4 contain configuration parameter b values.

The AE option uses the tSEND value to determine when to stop sending the RAS to the DRAM. The AE option stops sending the RAS signal when the bank cycle counter is cleared. The bank cycle counter is cleared when the bank cycle counter equals the tSEND value.

Transfer 5
New-page End Timing Value
Bits 0 through 4, Bit 12

During transfer 5, bits 0 through 4 of the DRAM control register are loaded with the configuration parameter b value that specifies the new-page end timing (tNEND). Refer again to Table 10-12. The AM option sends this configuration parameter b value to the AE option. The AM option also sends bit 12 to the AE option. This bit indicates to the AE option that bits 0 through 4 contain configuration parameter b values.

The AE option uses the tNEND value to determine when to stop sending the CAS to the DRAM. The AE option stops sending the CAS signal when the bank cycle counter is cleared. The bank cycle counter is cleared when the bank cycle counter equals the tNEND value.

Transfer 6
Column Address Strobe Pulse Width Timing Value
Bits 0 through 4, Bit 12

During transfer 6, bits 0 through 4 of the DRAM control register are loaded with the configuration parameter b value that specifies the column address strobe pulse width timing (tCAS). Refer again to Table 10-12. The AM option sends this configuration parameter b value to the AE option. The AM option also sends bit 12 to the AE option. This bit indicates to the AE option that bits 0 through 4 contain configuration parameter b values.

The AE option uses the tCAS value to determine when to assert the CAS signal that is sent to the DRAM. The AE asserts the CAS signal when the tCAS value equals the value of a bank cycle counter. The bank cycle counter begins incrementing by 1 each CP when the AM option initiates a new memory reference.

Transfer 7
Column Address Strobe Cycle Timing
Bits 0 and 1

During transfer 7, bits 0 and 1 of the DRAM control register specify the timing parameter for the interval between two CAS signals. Refer again to Table 10-12. The AM option uses these bits to determine when one memory reference is complete and when to initiate the next memory reference.

Transfer 7
DRAM Readout Sample Timing
Bits 2 and 3

During transfer 7, bits 2 and 3 of the DRAM control register specify the timing parameter for sampling read data. Refer again to Table 10-13. The AM option uses these bits to determine when it should instruct the AJ and AK options to sample the read data.

Transfer 7
DRAM Bus Enable
Bit 4

During transfer 7, bit 4 of the DRAM control register specifies the duration of the bank bus enable during write operations. When this bit is set to a 1, the AM option enables the bank bus for 5 CPs. When this bit is set to 0, the the AM option enables the bank bus for 4 CPs.

Transfer 7
DRAM Command
Bits 5 through 9

During transfer 7, bits 5 through 9 of the DRAM control register contain the command information that is sent to the AE option. The AM option sends the command information to the AE option over 2 CPs. Table 10-22 lists this configuration data.

Table 10-22. Configuration Data

1st Clock Period	2nd Clock Period
Bank select	CAS 0 mask
Chip select	CAS 1 mask
Write enable	CAS 2 mask
New page	CAS 3 mask
Go	Die select

Transfer 7
Disable Writes to Protected Low Memory
Bit 10

During Transfer 7, bit 10 of the DRAM control register specifies when the lower 8K of memory can be referenced. When the microprocessor attempts to write this lower memory and bit 10 is set to 1, the AM option aborts the write reference and signals the AE option to set the error interrupt. This portion of memory may be written to when bit 10 is set to 0.

Transfer 7
Instruction Fetch Read-ahead Mode
Bit 11

During transfer 7, bit 11 of the DRAM control register when set to a 1 specifies that the read-ahead operation is being performed. When this bit is set to a 1 and the data in the CPU read stage (CRS) in the AJ and AK options is valid (matches the address of the requested data), the AM option instructs the AJ and AK options to send the data in the CRS to the microprocessor. The AM option also instructs the AJ and AK options to read the data from the next successive address in local memory. This new data is stored in the CRS of the AJ and AK options.

When set to 0 and the data in the CRS is valid, the AM option also instructs the AJ and AK options to send the data in the CRS to the microprocessor, but no new data is read from local memory.

When bit 11 is set to 1 and the data is not valid in the CRS, the AM option instructs the AJ and AK options to read data from the specified address in local memory and send it to the microprocessor. Directly following this memory reference, the AM option instructs the AJ and AK options to read data from the next successive address and store it in the CRS.

When bit 11 is set to 0 and the data is not valid in the CRS, the AM option instructs the AJ and AK options to read data from the specified address in local memory and send it to the microprocessor (no other memory reference occurs).

Transfer 7
Configuration in Progress
Bit 12

During transfer 7, bit 12 of the DRAM control register is set to 0. This indicates to the AM option that the configuration is complete.

EE Option

The EE option contains the network enable register, the network performance register, and the network mode register. Refer to Figure 10-20.

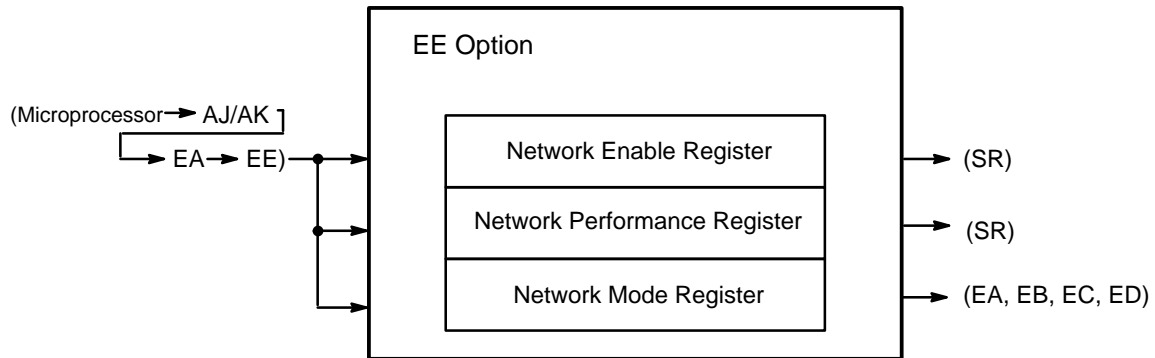


Figure 10-20. EE Option

Network Enable Register

The network enable register enables or disables each of the inputs to the X-dimension switch, the Y-dimension switch, and the Z-dimension switch in the network interface. Refer again to Table 10-17.

The EE option ANDs bit 0 of the network enable register with a PE Enable signal. The result of this AND is sent to the SR option in the X dimension. This SR option uses this signal to enable the input for the PE and BLT.

Bits 1 through 8 of the network enable register are also sent to the SR options. Bits 1 through 2 are sent to the X-dimension SR option, bits 3 through 5 are sent to the Y-dimension SR option, and bits 6 through 8 are sent to the Z-dimension SR option.

Network Performance Register

The network performance register controls which parameters of network performance are monitored. The following subsections describe the hardware associated with the network performance register.

Parameter Select
Bits 0 through 3

The EE option sends bits 0 through 3 of the network performance register to the SR option over 4 CPs. When the SR option has received all 4 bits, it uses the parameter select bits to select 1 of 15 network switch performance parameters. The SR option sends the network switch performance parameter to the ED option. Refer to Figure 10-21.




Figure 10-21. Performance Monitor Circuitry in the SR Option

PE 0 Switch or BLT Select Bits 4 and 5 and
PE 1 Switch or BLT Select Bits 6 and 7

The EE option sends the PE switch or BLT select bits to the ED option. The ED option uses the PE switch or BLT select bits to select one of four network switch performance parameters for each PE. The ED option sends the network switch parameters to the AR option in each PE. The AR options send the network switch parameter to the microprocessors. Refer to Figure 10-22.

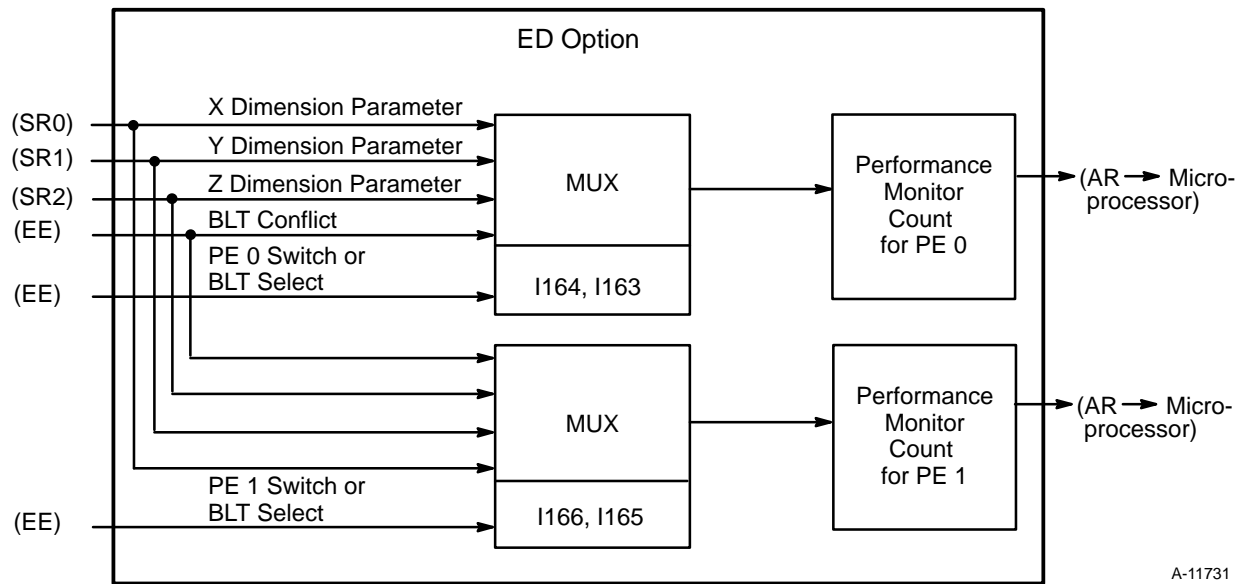


Figure 10-22. Performance Monitor Circuitry in the ED Option

Network Mode Register

The network mode register enables or disables several types of error checking in the network interface and the BLT. This register is located in the EE option. The following subsections describe how the EE option and other hardware use information from network mode register.

Enable Packet Errors
Bit 0

The EE option sends bit 0 of the network mode register to the EC option. When set to 1, the EC option checks for misrouted packets and parity errors in the header phits of the packet. When set to 0, this error checking circuitry is disabled.

Enable BLT Remote Index Errors
Bit 1

The EE option sends bit 1 of the network mode register to the EC option. When this bit is set to 1, the EC option is enabled to generate new check bits for the BLT remote index that was read from local memory. The EC option compares these new check bits to the old check bits for errors. When set to 0, the comparison between the new and old check bits is disabled.

Set Incoming Channel Parity Values
Bits 2 and 3

The EE option sends bits 2 and 3 of the network mode register to the ED option. When bit 3 is set to 1, the ED option selects the parity bits that were generated by the parity circuitry within the buffer to be compared to the parity bits that the ED option generates after the information is read out of the buffer.

When bit 3 is set to 0, the ED option selects the value of bit 2 for the parity bits to be compared with the parity bits that the ED option generates after the information is read out of the buffer. This provides the ability to force parity errors when information transfers from the input buffers in the network interface to a PE.

Set Outgoing Request Channel Parity Values
Bits 4 and 5

The EE option sends bits 4 and 5 of the network mode register to the EA option. When bit 5 is set to 1, the EA option sends the bits that were generated by the parity circuitry within the buffer to the destination PE. When bit 5 is set to 0, the EA option selects the value of bit 4 for the parity bits. This provides the ability to force parity errors when the EA option creates request packets.

Enable BLT Address Offset Range Errors
Bit 6

The EE option uses bit 6 of the network mode register to enable the reporting of address offset range errors. When this bit is set to 1, the error reporting is enabled. When set to 0, the error reporting is disabled.

Enable BLT Virtual PE Range Errors Bit 7

The EE option uses bit 7 of the network mode register to enable the error reporting of virtual PE range errors. When this bit is set to 1, the error reporting is enabled. When set to 0, the error reporting is disabled.

Set Outgoing Response Channel Parity Values Bits 8 through 10

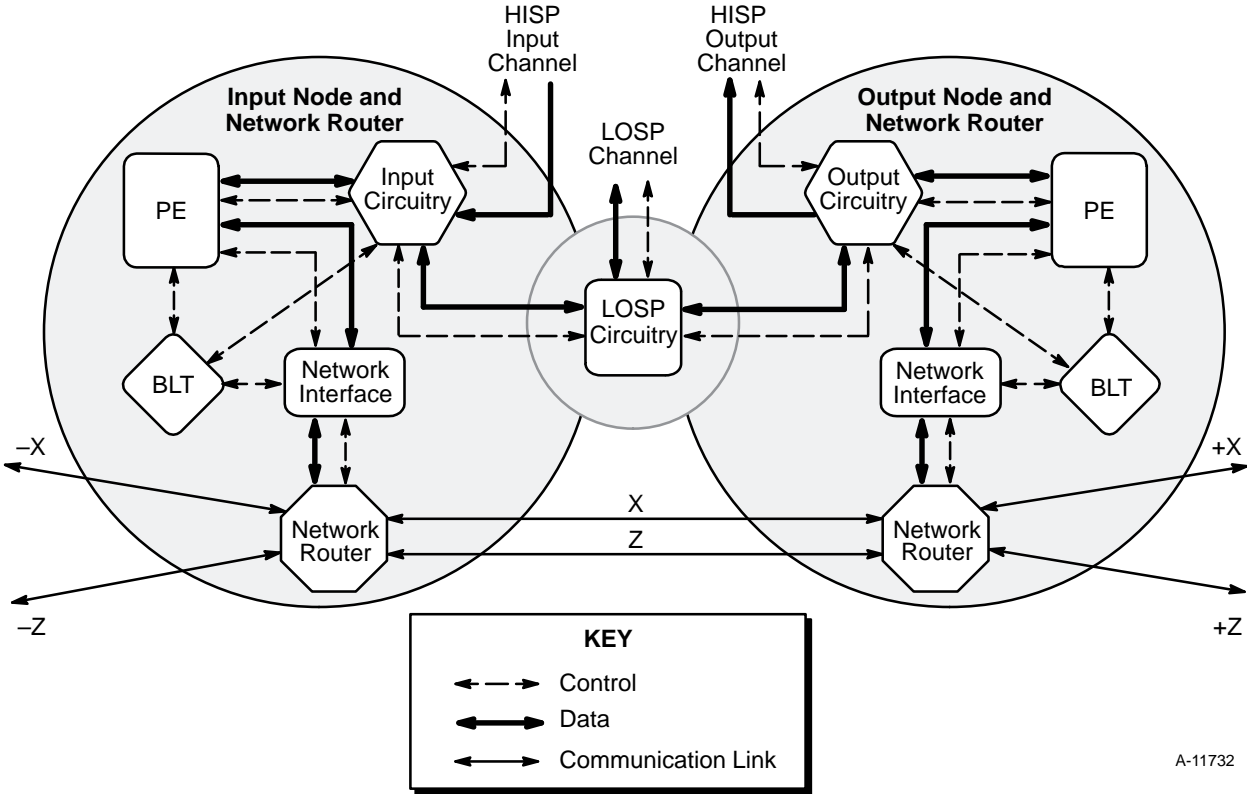
The EE option sends bits 8 through 10 of the network mode register to the EB option using the false outputs. When bit 8 is set to 1, the EB option selects the parity bits that were generated by the parity circuitry within the buffer to the destination PE. When bit 8 is set to 0, the EB option selects the value of bits 9 and 10 for the parity bits. This provides the ability to force parity errors when information transfers from the output buffers in the network interface to a destination PE.

Barrier Synchronization Circuitry Partition Bits Bits 11 through 13

The EE option sends bits 11 through 13 of the network mode register to the ED option. The ED option uses these bits to enable the bypass circuitry. Enabling the bypass circuitry sends the output of the fan-in circuitry to the input of the fan-out circuitry. For more information on the barrier synchronization fan-in and fan-out circuitry, refer to Section 9, “Barrier Synchronization.”

11 LOSP AND HISP CHANNELS

All input and output communication between the CRAY T3D system and the host system is performed through the I/O gateways. As was described in Section 1, "Architecture Overview," each I/O gateway contains an input node, an output node, and LOSP circuitry (refer to Figure 11-1).



A-11732

Figure 11-1. I/O Gateway

Each I/O gateway connects to one LOSP channel and one HISP channel. This section describes the format and use of the memory-mapped registers (hereafter referred to as registers) that are used in an I/O gateway. This section also includes descriptions of the master and slave I/O gateway registers.

LOSP Channel

LOSP channels transfer request and response information between the CRAY T3D system and a CPU or IOC in the host system. Either the CRAY T3D system or the host system can initiate a transfer over a LOSP channel.

Each LOSP channel is actually a pair of unidirectional channels (LOSP input and LOSP output). Figure 11-2 shows the signals used in these LOSP channels.

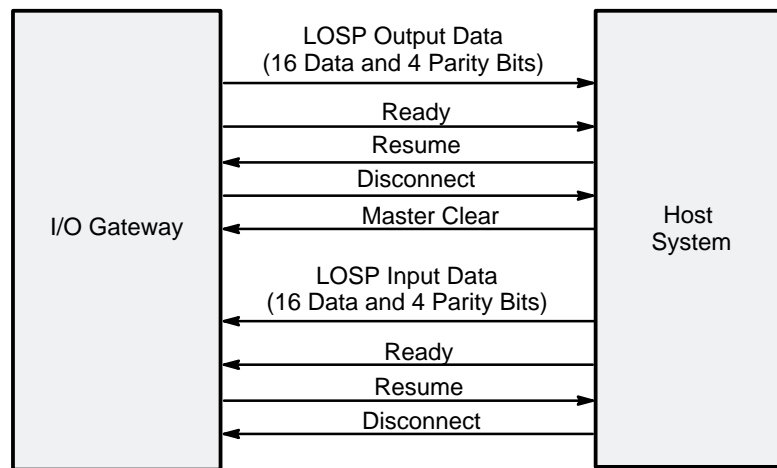


Figure 11-2. LOSP Channel Signals

LOSP data transfers over the LOSP channels in 16-bit parcels. Four parity bits that are used to check the data for errors are also transferred.

Both the input and output LOSP channels contain Ready, Resume, and Disconnect control signals. The Ready signal informs the receiver that 16 valid data bits are on the LOSP data signals. The Resume signal informs the sender that the data was received and new data may be transmitted. The Disconnect signal informs the receiver that the sender terminated the transfer.

Only the LOSP output channel contains a Master Clear signal. When the host system sets the Master Clear signal, the CRAY T3D system performs a global reset. The I/O gateway that receives the Master Clear signal is designated as the deadstart I/O gateway when the CRAY T3D system is initialized.

When information transfers over the LOSP input channel from the host system to the CRAY T3D system, the most significant bit of the first parcel transferred directs the information to the appropriate node. When this bit is 0, the information is for the output node. When this bit is 1, the information is for the input node.

When information transfers over the LOSP channel from the CRAY T3D system to the host system, the input node and output node share the LOSP channel. The first node to request a transfer over the LOSP channel controls the channel until that node sets the Disconnect LOSP Channel signal.

Input Channel

The host system transfers request and response information to an I/O gateway over the LOSP input channel. Table 11-1 shows the I/O registers used by the input node or output node in an I/O gateway when receiving LOSP data from the host system. Each node has a set of LOSP input registers; however, the LOSP_CFG register in the output node is not used for the LOSP input channel.

Table 11-1. LOSP Input Registers

Register Name	Definition	Description
LOSP_CFG	LOSP channel configuration	In the input node, this register sets parameters that control the Resume signal in the LOSP input channel.
IO_IFLAG	I/O interrupt flags	This register indicates what transfer conditions have occurred that may cause the I/O hardware interrupt to set.
IO_IMASK	I/O interrupt mask	This register indicates which transfer or error conditions will cause the I/O hardware interrupt to be set.
LOSP_DI	LOSP data in	This register is used to read the LOSP input data.
IO_ERR	Error register	This register indicates what error conditions occurred during the transfer.
IO_CLR	I/O clear	This register resets the bits in the IO_IFLAG and IO_ERR registers and clears the LOSP input circuitry.

The following subsections provide an overview of how these registers may be used to receive data from the host system over the LOSP input channel. Detailed descriptions of each register are provided in “Register Mapping” later in this section.

Initialization

Before performing the first LOSP input transfer, the microprocessor writes parameters into the LOSP_CFG register in the input node. These parameters set the pulse width of the Resume signal and set the amount of delay needed after the LOSP circuitry receives the leading edge of the Ready signal and before the LOSP circuitry samples the LOSP input data.

Interrupt Driven Versus Polled

The microprocessor in the input or output node of an I/O gateway uses one of two methods to monitor LOSP input transfers: interrupt driven or polled. When using the interrupt driven method, the microprocessor enables a hardware interrupt. When using the polled method, the microprocessor periodically checks the value of the IO_IFLAG register.

To enable the I/O hardware interrupt, bit 13 must be set to 1 in the hardware interrupt enable register (HIER) in the microprocessor. In addition to the HIER register, bits 7 and 8 of the IO_IMASK register should be set to 1 so the I/O hardware interrupt will be set when either of two LOSP input transfer conditions occur. These conditions are listed below. More information on each transfer condition is provided in the following subsection.

- LOSP input data available
- LOSP input disconnect

When using the polled method, the microprocessor may periodically check the value of its associated IO_IFLAG register. Selected bits in the IO_IFLAG register indicate when the previously listed LOSP input transfer conditions occur. The IO_IMASK register enables or disables the I/O hardware interrupt for certain transfer and error conditions; however, the IO_IFLAG register always indicates what conditions have occurred regardless of the value of the IO_IMASK register.

For simplicity, the following subsections describe a LOSP input transfer where the I/O hardware interrupt is enabled and will be set when a LOSP input data available or LOSP input disconnect transfer condition occurs (interrupt driven method). Also, the term “host system” in the following subsections refers to a CPU or IOC in the host system.

Receiving Data

When the host system initiates a LOSP data transfer to an I/O gateway, it places the first 16-bit parcel of data on the LOSP data signals and sets the Ready signal of the LOSP channel to 1. Bit 15 of the first parcel of LOSP data indicates which node in the I/O gateway will receive the LOSP data. When bit 15 is set to 1, the data is for the input node. When bit 15 is set to 0, the data is for the output node. Subsequent parcels of the transfer are for the node specified by bit 15 of the first parcel.

The host system may send several parcels of data to the I/O gateway. As the host system places each parcel on the LOSP data signals, it sets the Ready signal and waits for the I/O gateway to respond with the Resume signal (refer to Figure 11-3). After receiving the Resume signal, the host system sends the next parcel of LOSP data.

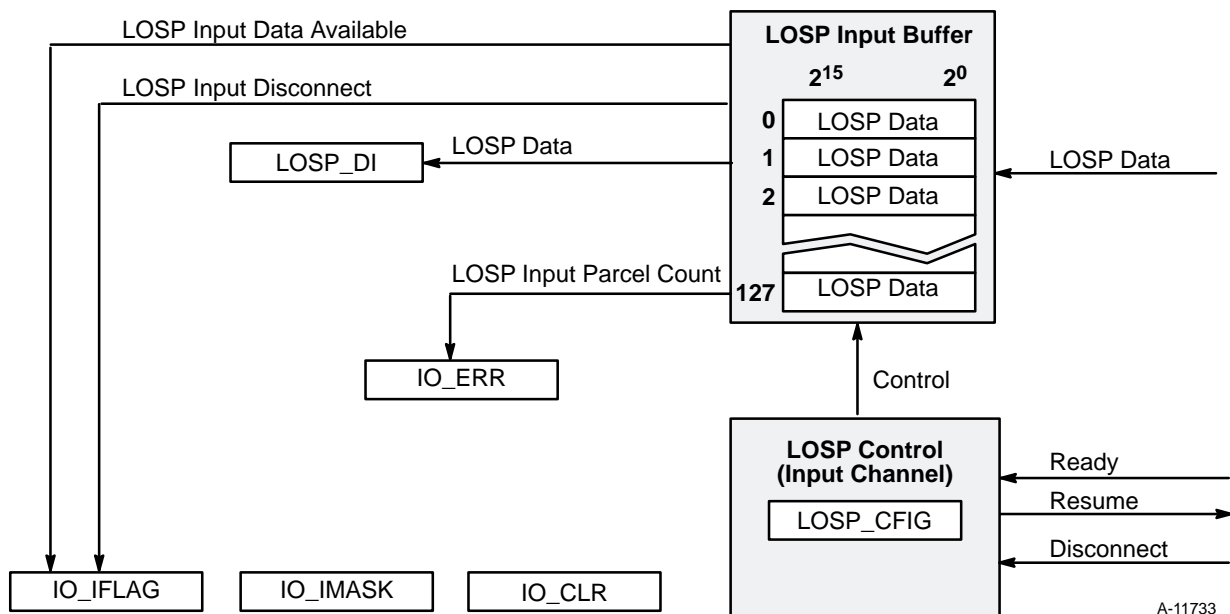


Figure 11-3. Receiving LOSP Input Data

The LOSP circuitry in the I/O gateway stores the first 4 parcels of data in the **LOSP_DI** register of the node specified by bit 15 of the first parcel transferred. The LOSP circuitry then stores subsequent parcels in the LOSP input buffer (refer again to Figure 11-3). The LOSP circuitry continues to store parcels in the LOSP input buffer until the host system sets the Disconnect signal instead of the Ready signal or until the LOSP input buffer fills with LOSP data.

Disconnect Received

When the host system sets the Disconnect signal, the LOSP circuitry sets the LOSP input data available bit (bit 7) in the IO_IFLAG register to 1. At this point, the LOSP input data available bit indicates that 0 to 132 parcels of data are stored in the LOSP_DI register and the LOSP input buffer. Because bit 7 of the IO_IMASK register is set to 1, the microprocessor in the input or output node receives the I/O hardware interrupt when the LOSP input data available bit is set to 1. Bit 15 of the first parcel of LOSP input data determines which microprocessor receives the hardware interrupt. (The following paragraphs assume bit 15 is set to 1.)

After receiving the I/O hardware interrupt, the microprocessor reads its associated IO_IFLAG register. When the LOSP input data available bit is set to 1, this bit signals the microprocessor that LOSP input data is available.

The microprocessor then reads the LOSP_DI register to receive the first word of LOSP input data. After the LOSP_DI is read, the LOSP circuitry transfers the next 4 parcels of LOSP input data from the LOSP input buffer to the LOSP_DI register. To reset the LOSP input data available bit to 0, the microprocessor must write a 1 to the clear LOSP input data available bit (bit 10) of the IO_CLR register.

When the microprocessor reads the LOSP_DI register again, it receives the second word of LOSP input data. The LOSP circuitry then transfers the next 4 parcels of data from the LOSP input buffer to the LOSP_DI register.

The microprocessor repeatedly reads the LOSP_DI register to read each word of LOSP input data. When the microprocessor reads the last word of the transfer from the LOSP_DI register, the support circuitry sets the LOSP input disconnect bit (bit 8) to 1 in the IO_IFLAG register. Because bit 8 of the IO_IMASK register is set to 1, this action also sets the I/O hardware interrupt. The microprocessor then reads the IO_IFLAG register to see what conditions have occurred.

When set to 1, the LOSP input disconnect bit indicates that the word read from the LOSP_DI register contained the last 0 to 4 parcels of the transfer. The microprocessor must read the valid parcel count (bits 5 through 7) from the IO_ERR register to determine the number of valid parcels that were read from the LOSP_DI register (refer again to Figure 11-3).

After reading the IO_ERR register and after performing error processing, the microprocessor resets the LOSP input disconnect bit to 0 in the IO_IFLAG register. It does this by writing a 1 to the clear LOSP input

disconnect bit (bit 9) of the IO_CLR register. When the LOSP input disconnect bit is reset to 0, it indicates that the input or output node is finished processing the LOSP input transfer and the node can receive another LOSP input transfer.

Full LOSP Input Buffer

When the buffer fills with LOSP input data before the host system sets the Disconnect signal, the LOSP circuitry sets the LOSP input data available bit of the IO_IFLAG register to 1. In this case, the LOSP input data available bit indicates that the input buffer and the LOSP_DI register contain a total of 132 parcels of data. The LOSP circuitry does not send the Resume signal to the host system. Because of this, the LOSP circuitry does not accept another parcel of LOSP data.

The microprocessor receives the I/O hardware interrupt when the LOSP input data available bit is set to 1. After receiving the I/O hardware interrupt, the microprocessor reads the IO_IFLAG register. Because the LOSP input data available bit is set to 1, this bit signals the microprocessor that LOSP input data is available.

The microprocessor then reads the first word of LOSP input data from the LOSP_DI register. After the LOSP_DI is read, the LOSP circuitry transfers the next 4 parcels of LOSP input data from the LOSP input buffer to the LOSP_DI register. To reset the LOSP input data available bit to 0, the microprocessor must write a 1 to the clear LOSP input data available bit (bit 10) of the IO_CLR register. The LOSP circuitry sends the Resume signal to the host system to accept another parcel of LOSP data.

When the microprocessor reads the LOSP_DI register again, it receives the second word of LOSP input data. The LOSP circuitry then transfers the next word of data from the LOSP input buffer to the LOSP_DI register.

The microprocessor continues to read from the LOSP_DI register. When the LOSP circuitry resets the LOSP Input Data Available bit to 1, the LOSP buffer has filled with data again, and the microprocessor must continue to read the LOSP_DI register.

When the LOSP circuitry sets the LOSP input disconnect bit to 1, the microprocessor has read the last 0 (none) to 4 parcels of the transfer from the LOSP_DI register. The microprocessor must read the valid parcel count from the IO_ERR register to find out how many parcels (0 to 4) of the last word read from the LOSP_DI register contain valid data.

When the microprocessor reads all of the LOSP input data before the host system sets the Disconnect signal, the microprocessor stalls, waiting for a value to return from the LOSP_DI register. The LOSP circuitry does not place data into the LOSP_DI register until it receives the LOSP data and a Disconnect signal from the host system.

Errors

Two types of errors may occur during a LOSP input transfer: a LOSP input sequence error or a LOSP input parity error. When either of these errors occur, the LOSP circuitry sets the appropriate bit in the IO_ERR register, which in turn sets the LOSP input error bit (bit 9) in the IO_IFLAG register (refer to Figure 11-4).




Figure 11-4. LOSP Input Errors

When enabled by the IO_IMASK register, the LOSP input error bit sets the I/O hardware interrupt; however, the LOSP input error interrupt is generally disabled to allow error processing to occur after the microprocessor reads the last word of a transfer from the LOSP_DI register.

LOSP Input Sequence Error

A LOSP input sequence error occurs when the host system places new LOSP data on the LOSP data signals and sets the Ready signal before the I/O gateway has read the previous LOSP data and set the Resume signal. When this occurs, the LOSP circuitry sets the LOSP input sequence error bit (bit 4) to 1 in the IO_ERR register, which in turn sets the LOSP input error bit to 1 in the IO_IFLAG register (refer again to Figure 11-4).

When the LOSP input error bit in the IO_IFLAG register is set to 1, an error occurred during the LOSP input transfer. After the transfer, the microprocessor reads the IO_ERR register to determine which error occurred. When the LOSP input sequence error bit is set to 1, a LOSP input sequence error occurred and the LOSP data the microprocessor received is not valid.

When the LOSP data is not valid, the node that received the data requests that the host system retransmit the data. The microprocessor in the node resets the LOSP Input Sequence Error bit to 0 by writing a 1 to the clear LOSP input error flags bit (bit 7) in the IO_CLR register. This action also resets both the LOSP input disconnect bit and the LOSP input error bit to 0 in the IO_IFLAG register.

LOSP Input Parity Error

A LOSP input parity error may occur due to one of two conditions. The first condition occurs when a parity error is detected on the LOSP data as it is transferred from the LOSP channel to the LOSP input buffer. The second condition occurs when a parity error is detected in the LOSP data as it is transferred from the LOSP input buffer to the LOSP_DI register.

In either condition, the LOSP circuitry sets the LOSP input parity error bit (bit 3) to 1 in the IO_ERR register only when the LOSP circuitry transfers the LOSP data from the LOSP buffer to the LOSP_DI register. When the LOSP input parity error bit is set to 1, the LOSP input error bit in the IO_IFLAG register sets to 1.

When the LOSP input error bit in the IO_IFLAG register is set to 1, an error occurred during the LOSP input transfer. After the transfer, the microprocessor reads the IO_ERR register to determine which error occurred. When the LOSP input parity error bit is set to 1, a LOSP input parity error occurred and the LOSP data received is not valid.

When the LOSP data is not valid, the node that received the data requests that the host system retransmit the data. The microprocessor in the node resets the LOSP input parity error bit to 0 by writing a 1 to the clear LOSP

input error flags bit (bit 7) in the IO_CLR register. This action also resets the LOSP input disconnect bit and the LOSP input error bit to 0 in the IO_IFLAG register.

Clearing the Channel

Different components of the LOSP input channel logic are cleared using selected bits in the IO_CLR register. These bits are clear LOSP input control (bit 0), clear LOSP input ready waiting bit (bit 5), clear LOSP buffer and priority (bit 6), clear LOSP input error flags (bit 7), and clear LOSP input disconnect (bit 9). When the IO_CLR register is in the input node, these bits may perform different functions than when it is in the output node. More information on the IO_CLR register is provided in “Register Mapping” later in this section.

Master Clear

When the host system sets the Master Clear signal, it signals the CRAY T3D system to perform a global reset. More information on the global reset is provided in Section 10, “Control and Status.” When an I/O gateway receives the Master Clear signal, it sets the master clear in bit (bit 12) of the IO_IFLAG register to 1 in both the input node and the output node (refer to Figure 11-5).

NOTE: In the hardware, the master clear in bit is also indicated by bit 12 of the IO_IMASK register. Because of this characteristic, the correct value of the master clear in bit is read by performing a logical OR between bit 12 of the IO_IFLAG register and bit 12 of the IO_IMASK register.

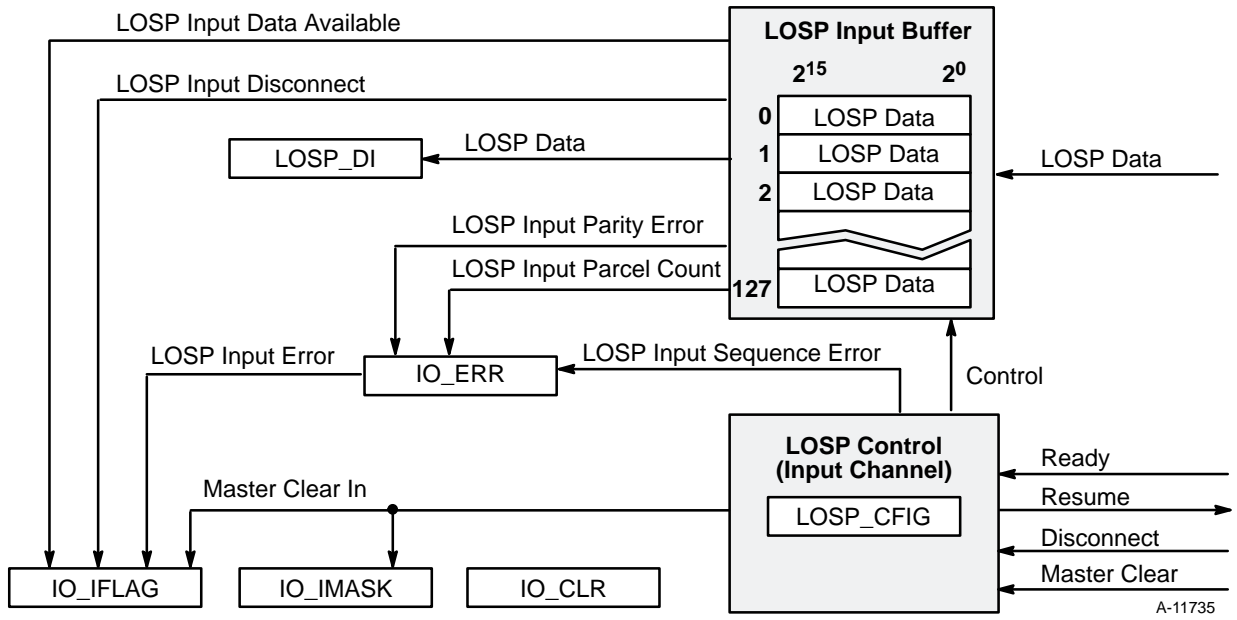


Figure 11-5. Master Clear

Output Channel

The I/O gateway transfers request and response information to the host system over the LOSP output channel. Table 11-2 shows the registers used by the input node or output node in an I/O gateway when sending LOSP data to the host system. Each node has a set of LOSP output registers; however, the LOSP_CFG register in the input node is not used for the output channel.

Table 11-2. LOSP Output Registers

Register Name	Definition	Description
LOSP_CFG	LOSP channel configuration	In the output node, this register sets parameters that control the Ready signal in the LOSP output channel.
IO_IFLAG	I/O interrupt flags	This register indicates what transfer conditions have occurred that may cause the I/O hardware interrupt to be set.
IO_IMASK	I/O interrupt mask	This register indicates which transfer or error conditions will cause the I/O hardware interrupt to be set.
LOSP_DO	LOSP data out	This register is used to send data to the LOSP output channel.
LOSP_DISC	LOSP output disconnect	This register sets the Disconnect signal on the LOSP output channel.
IO_ERR	Error register	This register indicates what error conditions occurred during the transfer.
IO_CLR	I/O clear	This register resets the bits in the IO_IFLAG and IO_ERR registers and clears the LOSP output circuitry.
IO_SYN	Syndrome register	This register contains the syndrome bits for a microprocessor single-bit or multiple-bit error.
DEC_SYN_CLR	Clear microprocessor syndrome register	This register clears microprocessor single-bit or multiple-bit errors.

The following subsections provide an overview of how these registers may be used to send LOSP data to the host system over the LOSP output channel. Detailed descriptions of each register are provided in “Register Mapping” later in this section.

Initialization

Before performing the first LOSP output transfer, the microprocessor in the output node writes parameters into the LOSP_CFG register. These parameters control the pulse width of the Ready signal.

Interrupt Driven Versus Polled

Like during LOSP input transfers, the microprocessor in the input or output node of an I/O gateway uses either the interrupt driven method or the polled method to monitor the LOSP output transfer.

The microprocessor must set bit 13 of the HIER register to 1 to receive the I/O hardware interrupt. In addition to the HIER register, bit 10 of the IO_IMASK register should be set to 1 so the I/O hardware interrupt will be set when a LOSP output complete transfer condition occurs. More information on this transfer condition is provided in the “Sending Data” subsection.

For simplicity, the following subsections describe a LOSP output transfer where the I/O hardware interrupt is enabled and will be set when a LOSP output complete transfer condition occurs. Also, the term host system in the following subsections refers to a CPU or IOC in the host system.

Sending Data

When the microprocessor in the input node or output node initiates a LOSP data transfer to the host system, it writes the first word of LOSP output data in the LOSP_DO register (refer to Figure 11-6). The microprocessor then waits until the LOSP circuitry indicates that the microprocessor has control of the LOSP output channel.

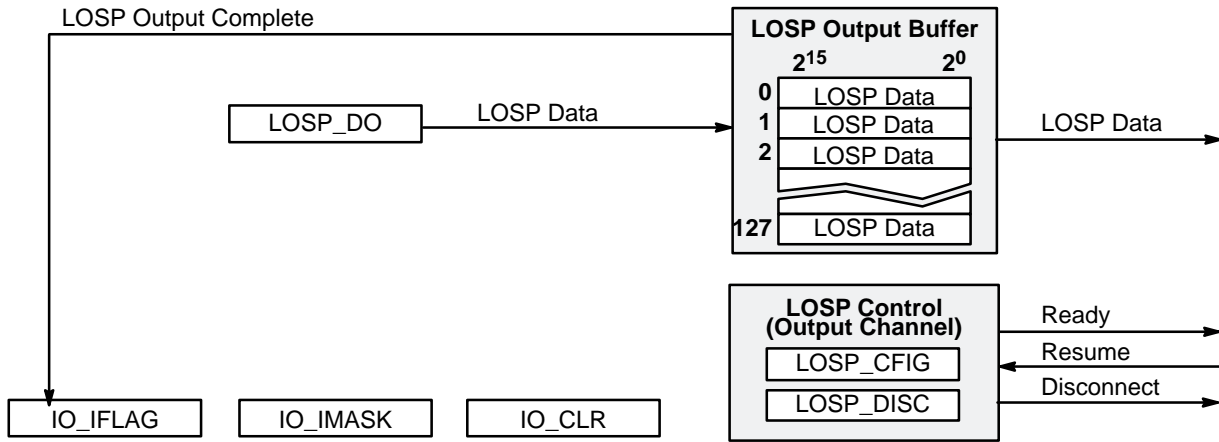


Figure 11-6. Sending LOSP Output Data

When the microprocessor in the input node and the microprocessor in the output node both write a word of data to their associated LOSP_DO registers at the same time, the LOSP circuitry selects which node has priority to use the LOSP output channel. The LOSP circuitry then transfers the 4 parcels of data from the LOSP_DO register in that node to the LOSP channel. After emptying the LOSP_DO register, the LOSP circuitry sets the LOSP output complete bit to 1 in the IO_IFLAG register of the node that has control of the LOSP output channel.

Because bit 10 of the IO_IMASK register is set to 1, when the LOSP output complete bit is set to 1, the I/O hardware interrupt is set. The microprocessor then reads the IO_IFLAG register to determine what condition set the I/O hardware interrupt. Because the LOSP output complete bit is set to 1, this bit indicates that the LOSP circuitry transferred the word of data from the LOSP_DO register to the LOSP channel and that the microprocessor has control of the LOSP output channel.

The microprocessor then writes the second word of LOSP data into the LOSP_DO register. This action causes the LOSP circuitry to reset the LOSP output complete bit to 0 in the IO_IFLAG register. The LOSP circuitry then transfers the 4 parcels of data from the LOSP_DO register into the LOSP output buffer and eventually to the LOSP channel.

After the microprocessor receives the I/O hardware interrupt that indicates it has control of the channel, the microprocessor can write up to 31 more words of data to the LOSP_DO register. When finished writing the last word of the transfer to the LOSP_DO register or when finished writing the 31st word of data to the LOSP_DO register, the microprocessor must wait for the I/O hardware interrupt to set.

When the microprocessor writes more than 31 words to the LOSP_DO register before the microprocessor receives the I/O hardware interrupt, the microprocessor may overwrite valid data that is stored in the LOSP output buffer. The I/O gateway hardware does not check for this condition, and all flow control must be maintained by software.

When the LOSP circuitry transfers the last word of LOSP data from the LOSP_DO register and LOSP output buffer to the LOSP output channel, the LOSP circuitry sets the LOSP output complete bit in the IO_IFLAG register to 1. When this bit is set to 1, the I/O hardware interrupt is set to the microprocessor. The microprocessor then reads the IO_IFLAG register to determine what condition set the I/O hardware interrupt.

When set to 1, the LOSP output complete bit indicates that all of the LOSP data that was written to the LOSP_DO register has been transferred to the LOSP output channel. When the microprocessor is finished with the transfer, the microprocessor writes any value to the LOSP_DISC register. This action causes the LOSP circuitry to set the Disconnect signal in the LOSP output channel, which terminates the transfer. The LOSP circuitry then determines which node has control of the LOSP output channel next and, when applicable, starts another LOSP output transfer.

When the microprocessor is not finished with the transfer, the microprocessor writes up to 31 more words of data to the LOSP_DO register. When finished writing the last word of the transfer or when finished writing the 31st word of data to the LOSP_DO register, the microprocessor must again wait for the I/O hardware interrupt to set due to a LOSP output complete condition. When this occurs, the microprocessor may terminate the LOSP output transfer by writing more data to the LOSP_DO register or by writing any value to the LOSP_DISC register.

Errors

Three types of errors may occur with a LOSP output transfer: LOSP output ROA parity errors, LOSP output sequence errors, and microprocessor SECEDED errors. When a LOSP output ROA parity error or a LOSP output sequence error occur, the LOSP circuitry sets the appropriate bit in the IO_ERR register, which in turn sets the LOSP output error bit (bit 11) in the IO_IFLAG register (refer to Figure 11-7). When microprocessor SECEDED errors occur, the LOSP circuitry sets the appropriate bits in the IO_IFLAG register.

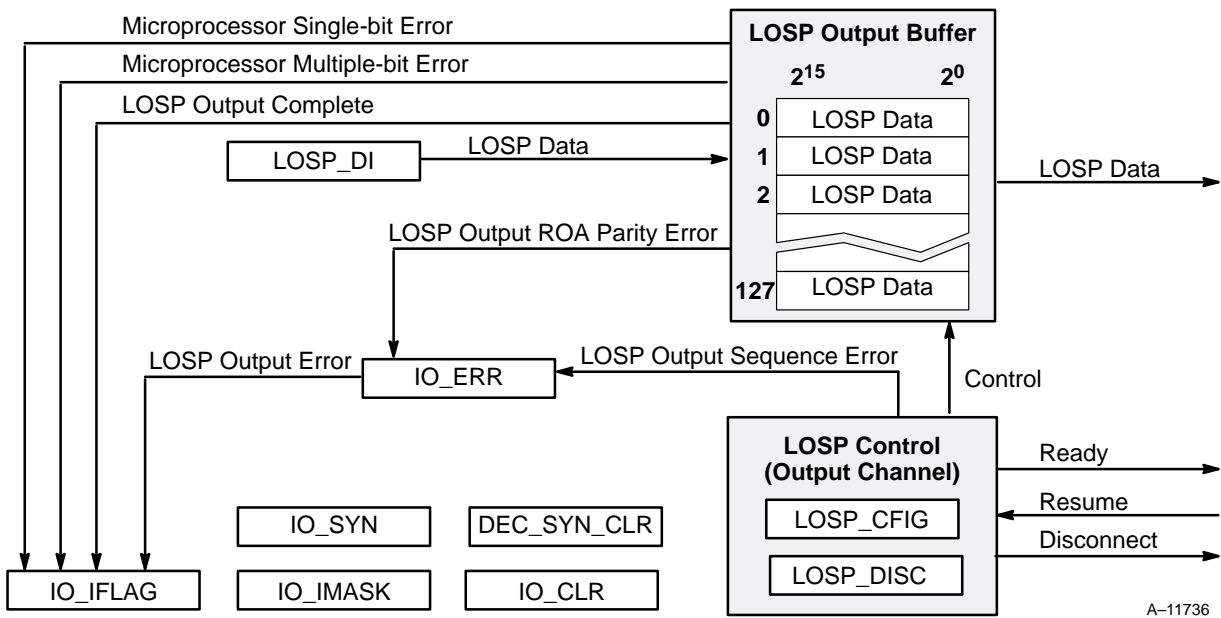


Figure 11-7. LOSP Output Errors

When enabled by the IO_IMASK register, the LOSP output error bit sets the I/O hardware interrupt; however, the LOSP output error interrupt is generally disabled to allow error processing to occur after the microprocessor writes the last word of the transfer into the LOSP_DO register. When enabled by the IO_IMASK register, the microprocessor SECEDED error bits of the IO_IFLAG register also set the I/O hardware interrupt. These interrupts are usually enabled during the transfer.

LOSP Output ROA Parity Error

A LOSP output ROA parity error occurs when the LOSP circuitry detects a parity error when reading a parcel of data out of the LOSP output buffer. This error is detected before the LOSP circuitry generates the parity bits for the LOSP output channel. When a LOSP output ROA parity error occurs, the LOSP circuitry sets the LOSP output channel parity bits so that the host system will also detect a parity error when it receives the LOSP data.

The LOSP circuitry also sets the LOSP output ROA parity error bit (bit 10) of the IO_ERR register to 1. This action causes the LOSP output error bit of the IO_IFLAG register to also set to 1 (refer again to Figure 11-7).

When the LOSP output error bit in the IO_IFLAG register is set to 1, an error occurred during the transfer. After the transfer, the microprocessor reads the IO_ERR register to determine which error occurred. When the LOSP output ROA parity error bit is set to 1, a parity error occurred and the LOSP data that was sent was not valid.

To clear the error, the microprocessor must write a 1 to the clear LOSP output error flags bit (bit 8) of the IO_CLR register. The microprocessor then writes any value to the LOSP_DISC register to terminate the transfer and set the LOSP output complete bit in the IO_IFLAG register to 0. The microprocessor can then attempt to send the LOSP output data to the host system again.

LOSP Output Sequence Error

A LOSP output sequence error occurs when the host system sets the Resume LOSP channel signal and the LOSP circuitry is not performing a LOSP output transfer. When this occurs, the LOSP circuitry sets the LOSP output sequence error bit (bit 11) to 1 in the IO_ERR register, which in turn sets the LOSP output error bit in the IO_IFLAG register to 1 (refer again to Figure 11-7). The microprocessor will detect this error only when the I/O hardware interrupt is enabled for error conditions, or when the microprocessor reads the value of the IO_IFLAG register.

To clear the error, the microprocessor must write a 0 to the clear LOSP output error flags bit (bit 8) of the IO_CLR register. The microprocessor then writes any value to the LOSP_DISC register to terminate the transfer requested by the host system.

Microprocessor SECDED Errors

A microprocessor SECDED error occurs when the I/O circuitry detects an error when performing SECDED on LOSP output data. The I/O circuitry performs SECDED on LOSP output data before writing the data into the LOSP_DO register. (A microprocessor SECDED error may also occur during a HISP output transfer. More information on the HISP output transfer is provided later in this section.)

When the I/O circuitry receives a 64-bit word of data and 14 check bits from a microprocessor (or BLT transfer during HISP output transfers), the circuitry generates a new set of check bits for the data. When the new check bits match the original check bits, none of the data or original check bits changed value after the microprocessor generated them. When the new check bits do not match the original check bits, one or more of the data or original check bits changed value after the microprocessor generated them.

When only one bit changed value, the I/O circuitry corrects the bit and sets the microprocessor single-bit error bit (bit 1) of the IO_IFLAG register to 1. The I/O circuitry also stores the syndrome bits for the single-bit error in the syndrome register (IO_SYN) of the node performing the transfer. The syndrome bits indicate which bit of a 32-bit halfword was corrected. More information on the microprocessor syndrome bits is provided in Section 3, "Processing Element Node."

After the I/O circuitry sets the microprocessor single-bit error bit to 1, and one or more single-bit errors occur, the microprocessor single-bit errors bit remains set to 1. The syndrome stored in the IO_SYN register corresponds to the last single-bit error that occurred. When one or more multiple-bit errors occur, the microprocessor single-bit error bit resets to 0. The syndrome stored in the IO_SYN register corresponds to the first multiple-bit error that occurred.

When more than one bit changed value, the I/O circuitry cannot correct the bits. In this case, the I/O circuitry sets the microprocessor multiple-bit error bit (bit 2) of the IO_IFLAG register to 1. The I/O circuitry also stores the syndrome bits for the multiple-bit error in the syndrome register (IO_SYN) of the node performing the transfer.

After the I/O circuitry sets the microprocessor multiple-bit error bit to 1, and one or more single-bit errors occur, the microprocessor multiple-bit errors bit remains set to 1. The syndrome stored in the IO_SYN register corresponds to the first multiple-bit error that occurred. When one or more multiple-bit errors occur, the microprocessor multiple-bit error bit remains set to 1. The syndrome stored in the IO_SYN register corresponds to the first multiple-bit error that occurred.

The microprocessor resets the microprocessor single-bit error and microprocessor multiple-bit error bits of the IO_IFLAG register by writing any value to the clear microprocessor syndrome register (DEC_SYN_CLR). Although the microprocessor SECDED bits of the IO_IFLAG register reset to 0, the syndromes in the IO_SYN register remain unchanged. The microprocessor syndromes in the IO_SYN register are valid only when one of the IO_IFLAG microprocessor SECDED bits is set to 1.

Clearing the Channel

Different components of the LOSP output channel are cleared using the following selected bits in the IO_CLR register.

- Clear LOSP output control (bit 1)
- Clear LOSP buffer and priority (bit 6)
- Clear LOSP output error flags (bit 8)

These bits may perform different functions when the IO_CLR register is in the input node compared to the output node. More information on each of these bits is provided in “Clear I/O Register” in “Register Mapping” later in this section.


HISP Channel

HISP channels transfer system data between the CRAY T3D system and the host system. A HISP channel connects two components: a master and a slave. The master controls the HISP channel by providing address information to the slave.

When a HISP channel connects a master I/O gateway to a CPU in the host system, the I/O gateway is the master of the HISP channel and the CPU is the slave. When a HISP channel connects a slave I/O gateway to an IOC, the IOC is the master of the HISP channel and the I/O gateway is the slave.

The data transfer rate of a HISP channel is 100 Mbytes/s or 200 Mbytes/s in each direction. The transfer rate can be changed by software.

Figure 11-8 shows the signals used in a HISP channel. Table 11-3 lists the HISP channel signals and describes each signal. In Figure 11-8, the master component is the I/O gateway and the slave component is the host system. The term input hereafter refers to a transfer from the host system to the I/O gateway, and the term output hereafter refers to a transfer from the I/O gateway to the host system.



A-11737

Figure 11-8. HISP Channel Signals

Table 11-3. HISP Channel Signals

Signal	Description
Data	Data is transferred between the master and slave in 64-bit words; 8 check bits that are used to perform SECDED are also transferred.
Address and Block Length	The address is the location in the slave where the first word that will be transferred resides. The block length is the total number of 64-bit words that will be transferred. Like the data signals, the transfer rate of address and block length information may be 100 Mbytes/s or 200 Mbytes/s.
Address Ready	The master sets the Address Ready signal to inform the slave that valid information is on the Address and Block Length signals.
Transmit Address	The slave sets the Transmit Address signal to inform the master that the slave is ready to receive the address and block length. The slave resets the Transmit Address signal after receiving the address and block length.
Address Error	The slave sets the Address Error signal when it detects a parity error on the address and block length information or when the master sets the Address Ready signal out of sequence.
Data Ready	The sender sets the Data Ready signal to inform the receiver that valid data is on the Data signals.
Transmit Data	The receiver sets the Transmit Data signal to inform the sender that the receiver is ready to accept data.
Last Word	The sender sets the Last Word signal to inform the receiver that the data on the Data signals is the last word of the transfer.
Unrecoverable Error	The slave sets the Unrecoverable Error signal to inform the master that the slave cannot recover from an error it has encountered.
Disable SECDED	The master sets the Disable SECDED signal to disable SECDED error correction on the HISP data.
Clear Channel	The master sets the Clear Channel signal to clear the HISP channel logic and inform the slave to reset the Unrecoverable Error and Address Error signals.

Master Input Channel

The master I/O gateway receives system data from the host system over the HISP input channel. Table 11-4 shows the I/O registers used by the input node in a master I/O gateway when receiving HISP data from the host system.

Table 11-4. HISP Input Registers

Register Name	Definition	Description
HISP_A_CFIG	HISP channel address configuration	This register sets several parameters for the HISP channel control signals.
HISP_D_CFIG	HISP data configuration	This register sets timing parameters for reading the HISP input data.
HISP_L_BL	HISP block length lower bits	This register contains bits 15 through 0, which specify the number of 64-bit words that will be transferred (block length).
HISP_U_BL	HISP block length upper bits	This register contains bits 17 and 16, which specify the number of 64-bit words that will be transferred (block length).
IO_IFLAG	I/O interrupt flags	This register indicates what transfer conditions have occurred that may cause the I/O hardware interrupt to be set.
IO_IMASK	I/O interrupt mask	This register indicates which transfer or error conditions will cause the I/O hardware interrupt to be set.
HISP_BUFFER (0 through 127)	HISP input buffer registers	These registers make up the HISP input buffer.
IO_ERR	Error register	This register indicates what error conditions occurred during the transfer.
IO_CLR	I/O clear	This register resets the bits in the IO_IFLAG and IO_ERR registers and clears the HISP input circuitry.
HISP_ADDR0	HISP address 0	This register is used to transfer address and block length information to the host system over the HISP channel Address and Block Length signals.
HISP_ADDR1	HISP address 1	This register is also used to transfer address and block length information to the host system over the HISP Address and Block Length signals.

Table 11-4. HISP Input Registers (continued)

Register Name	Definition	Description
IO_SYN	Syndrome register	This register contains the syndrome bits for a HISP single-bit or multiple-bit error.
HISP_SYN_CLR	Clear HISP syndrome register	This register clears HISP single-bit or multiple-bit errors.

Requesting a Transfer

To request data from the host system, a PE in a processing element node sends information to the PE in the input node of a master I/O gateway. The request information contains the following parameters.

- Number of 64-bit words requested
- Address where the first requested word will be stored
- Stride value and transfer type (constant stride or scatter)
- Address in the host system of the first requested word

The microprocessor in the input node PE then interprets this information and generates values for the input node memory-mapped registers.

Initialization

Before performing HISP input transfers, the microprocessor sets parameters in the HISP_A_CFIG and HISP_D_CFIG registers that control the HISP input channel signals. The HISP_A_CFIG register contains values that set the pulse width and period of the Address Ready signal and set the format of the address and block length information. The HISP_D_CFIG register contains values that set the time needed for the input node to sample data from the HISP input channel.

Interrupt Driven Versus Polling

The microprocessor in the input node of a master I/O gateway monitors the HISP input transfer by enabling the I/O and BLT hardware interrupts (interrupt driven method) or by periodically checking the value of its associated IO_IFLAG register and BLT_SR (polled method). The following subsections describe a HISP input transfer where the I/O and BLT hardware interrupts are enabled.

The microprocessor must set bit 13 of the HIER register to 1 in order to receive the I/O hardware interrupt and must also set bit 12 of the HIER register to 1 in order to receive the BLT hardware interrupt. In addition, the microprocessor sets bits 6, 4, and 3 of the IO_IMASK register to 1 so the microprocessor will receive the I/O hardware interrupt when any of the following transfer conditions occur. (More information on each transfer condition is provided in the following subsections.)

- HISP data transfer complete
- HISP single-bit error
- HISP multiple-bit error

BLT Parameters

The BLT in the input node is used to transfer HISP data from the HISP input buffer to the destination PEs. Before performing a HISP input transfer, the microprocessor writes the appropriate values to the BLT memory-mapped registers. For example, Table 11-5 shows sample BLT register values for a HISP input transfer. For more information on the BLT registers, refer again to Section 8, “Block Transfer Engine.”

Table 11-5. Sample BLT Input Parameters

Register Name	Value	Description
BLT_LAR	7000000 ₁₆	Set the local address to the address of the HISP_BUFFER0 register, which is the first entry of the HISP input buffer
BLT_LSR	1	Set the local stride to 1 word
BLT_RMR	FFF000000 ₁₆	Set the remote mask so bits 35 through 24 of the remote index are the PE number and bits 23 through 0 are the index offset
BLT_RIR	001000000 ₁₆	Set the remote index to a destination PE number of I/O = 0 Z = 0 Y = 0 X = 0 PE = 1 and an index offset of 0
BLT_RSR	1	Set the remote stride to 1 word
BLT_VLR	10000 ₁₆ (block length)	Set the vector length register so it matches the block length stored in the HISP_L_BL and HISP_U_BL registers †
BLT_RBR	88000 ₁₆	Set the remote base to an offset of 88000 ₁₆
BLT_RLR	97FFF ₁₆	Set the remote limit to an offset of 97FFF ₁₆
BLT_CR	18 ₁₆	Set the control register so that a constant stride write transfer is selected, the large packet size is disabled, logical PE numbers are enabled, the I/O mode is enabled, the 256 outstanding request count is enabled, and the PE range mask is set to 0
BLT_SR	5	Set the status register so that the microprocessor receives the BLT hardware interrupt when the BLT transfer is complete or when a BLT error occurs

† Because the vector length indicates the number of large packets (4-word body) the BLT generates, set the vector length to the block length divided by four. The block length must be a multiple of four in order to use large packet transfers. For more information on large packet transfer restrictions, refer to "Enable Large Packet Size" in Section 8, "Block Transfer Engine."

As soon as the microprocessor writes information into the BLT_CR, the BLT starts a transfer; however, because the BLT I/O mode is enabled, the BLT does not actually transfer HISP input data until it receives handshaking control signals from the HISP input circuitry.

I/O Registers

The microprocessor loads the HISP_L_BL and HISP_U_BL registers with the number of 64-bit words that will be transferred (block length). The HISP_L_BL register contains bits 0 through 15 of the block length, and the HISP_U_BL register contains bits 16 and 17 of the block length.

The HISP_ADDR0 and HISP_ADDR1 registers are the last registers to be written by the microprocessor. These registers contain the address and block length information that transfers to the host system over the Address and Block Length HISP channel signals (refer to Figure 11-9). The address is the address of the first word in the host system that will be transferred. The block length is the total number of 64-bit words that will be transferred and must be equal to the block length stored in the HISP_L_BL and HISP_U_BL registers.




Figure 11-9. Master HISP Input Transfer

Sending the Address and Block Length

As soon as the microprocessor writes address and block length information into the HISP_ADDR0 register, the input circuitry starts the HISP input transfer. The input circuitry first transfers the address and block length information from the HISP_ADDR0 and HISP_ADDR1 registers to the host system over the Address and Block Length HISP channel signals.

When the input node completes the transfer of address and block length information to the host system, the input node sets the HISP address transfer complete bit (bit 5) of the IO_IFLAG register to 1. This indicates that the address transfer is complete, but does not set the I/O hardware interrupt because bit 5 in the IO_IMASK register is not set to 1.

Receiving Data

The input circuitry receives HISP input data from the host system over the HISP data signals. When valid data is on the HISP data signals, the host system sets the Data Ready signal. After receiving the data, the input circuitry sets the Transmit Data signal to indicate that the host system may transmit new data. The input circuitry also decrements the block length stored in the HISP_L_BL and HISP_U_BL registers as it receives each word of HISP input data.

The microprocessor may periodically read the value of the block length to determine how many words are left to transfer. To obtain bits 0 through 15 of the block length, the microprocessor reads the HISP_L_BL register. To obtain bits 16 and 17 of the block length, the microprocessor reads bits 14 and 15 of the IO_ERR register (refer again to Figure 11-9).

The input circuitry stores each word of data in the HISP input buffer. The HISP input buffer is actually a set of 128 memory-mapped registers (HISP_BUFFER 0 through HISP_BUFFER 127) that operate as a circular buffer (refer again to Figure 11-9). Each HISP_BUFFER register has a unique address. More information on the HISP_BUFFER address is provided in “Register Mapping” later in this section.

The HISP input circuitry exchanges handshaking signals with the HISP input buffer on 16-word boundaries (eight 16-word blocks). As the input circuitry receives input data from the HISP channel, it stores the data in one block of the HISP input buffer. When the block is full, the input circuitry increments to the next block of the buffer. When all 8 blocks of the buffer are full, the input circuitry does not signal the host system that it

can receive more data. The input circuitry then waits until a block of the HISP input buffer empties before signaling the host system and receiving more data.

When the input circuitry stores a word into HISP_BUFFER 127 (last entry of the HISP input buffer) and the block length has not decremented to 0, the input circuitry resets the internal buffer address. This action causes the input circuitry to write the next word of HISP input data into HISP_BUFFER 0 (first entry of the HISP input buffer). When HISP_BUFFER 0 contains valid data, the input circuitry waits until the first HISP buffer block (HISP_BUFFER 0 through HISP_BUFFER 15) is empty before writing new data into HISP_BUFFER 0.

The BLT exchanges handshaking signals with the HISP input buffer on 32-word boundaries (four 32-word blocks). As soon as the input circuitry fills the first 32-word block in the HISP input buffer, the BLT reads data from the input buffer. The BLT reads 1 or 4 words from the buffer, creates a BLT write request packet, and sends the packet to the destination PE(s).

When the BLT reads all of the data from a 32-word block before the input circuitry fills the next 32-word block, the BLT waits until the input circuitry fills the next 32-word block or the input circuitry writes the last word of the transfer in the next 32-word block. When this occurs, the input circuitry signals the BLT that the BLT can read data from the next 32-word block.

When the BLT reads a word from the last entry of the HISP input buffer and the BLT vector length has not decremented to 0, the BLT continues to increment the local address by the local stride. For example, after reading a word from HISP_BUFFER 127 using the address $700007F_{16}$, the BLT increments the address by 1 (the local stride) to obtain a new address of 7000080_{16} ; however, the input circuitry does not use bits 7 through 11 of a HISP_BUFFER register address. Because of this characteristic, when the BLT provides an address of 7000080_{16} , the input circuitry interprets this address as the address for HISP_BUFFER 0 (7000000_{16}). More information on the HISP_BUFFER address is provided in “Register Mapping” later in this section.

The input circuitry continues to receive data from the HISP input channel and store the data in the HISP input buffer until the host system sets the Last Word HISP channel signal. Setting this signal indicates that the host system has placed the last word of the transfer on the data signals of the HISP channel. After storing the last word of data in the HISP input buffer, the input circuitry decrements the block length stored in the HISP_L_BL and the HISP_U_BL registers.

When the input circuitry receives the Last Word HISP channel signal, the block length stored in the HISP_L_BL and HISP_U_BL registers should decrement to 0. When the block length does not decrement to 0, a HISP block length error occurred. More information on HISP input errors is provided later in this section.

When the block length does decrement to 0, the input circuitry sets the HISP data transfer complete bit (bit 6) of the IO_IFLAG register to 1. Because bit 6 of the IO_IMASK register is set to 1, this action sets the I/O hardware interrupt. The I/O hardware interrupt signals the microprocessor that the input circuitry has received all of the HISP input data and has stored the data in the HISP input buffer.

When the BLT receives a response packet for all of the write request packets it generated during the HISP input transfer, the BLT sets the BLT hardware interrupt to 1. This action signals the microprocessor that the BLT transfer is complete and that the HISP data has been transferred from the HISP input buffer to the destination PE(s).

The BLT in an input node does not have to be used to transfer data from the HISP input buffer to the PEs in the system. Instead, the microprocessor in the input node may read data out of the HISP input buffer. The microprocessor uses different addresses to read data out of the HISP input buffer than the BLT uses. More information on the microprocessor addressing is provided in “HISP Buffer Registers” in “Register Mapping” later in this section.

Errors

Five types of errors may occur during a master HISP input transfer: HISP SECDDED errors, HISP ROA errors, HISP channel errors, HISP block length errors, and BLT errors. These errors are reported using the IO_ERR and IO_IFLAG registers (refer to Figure 11-10) and the BLT hardware interrupt. The following subsections describe each error.

Some of the HISP input errors indicate that the HISP input transfer should be terminated. To terminate the transfer, the microprocessor performs three steps.

First, the microprocessor aborts the BLT transfer. To do this, the microprocessor writes a status select code of two into the BLT_SR. For more information on the BLT abort code, refer to “BLT Abort” in Section 8, “Block Transfer Engine.”

Second, the microprocessor terminates the transfer with the host system. To do this, the microprocessor writes a 1 to the clear channel HISP signal bit (bit 12) of the HISP_A_CFG register. This bit must be set to 1 for 150ns or until the host system resets the Unrecoverable Error HISP channel signal.

Third, the microprocessor clears the HISP input circuitry. To do this, the microprocessor writes a 1 to the clear HISP address control and clear HISP data control bits (bits 2 and 3) of the IO_CLR register.

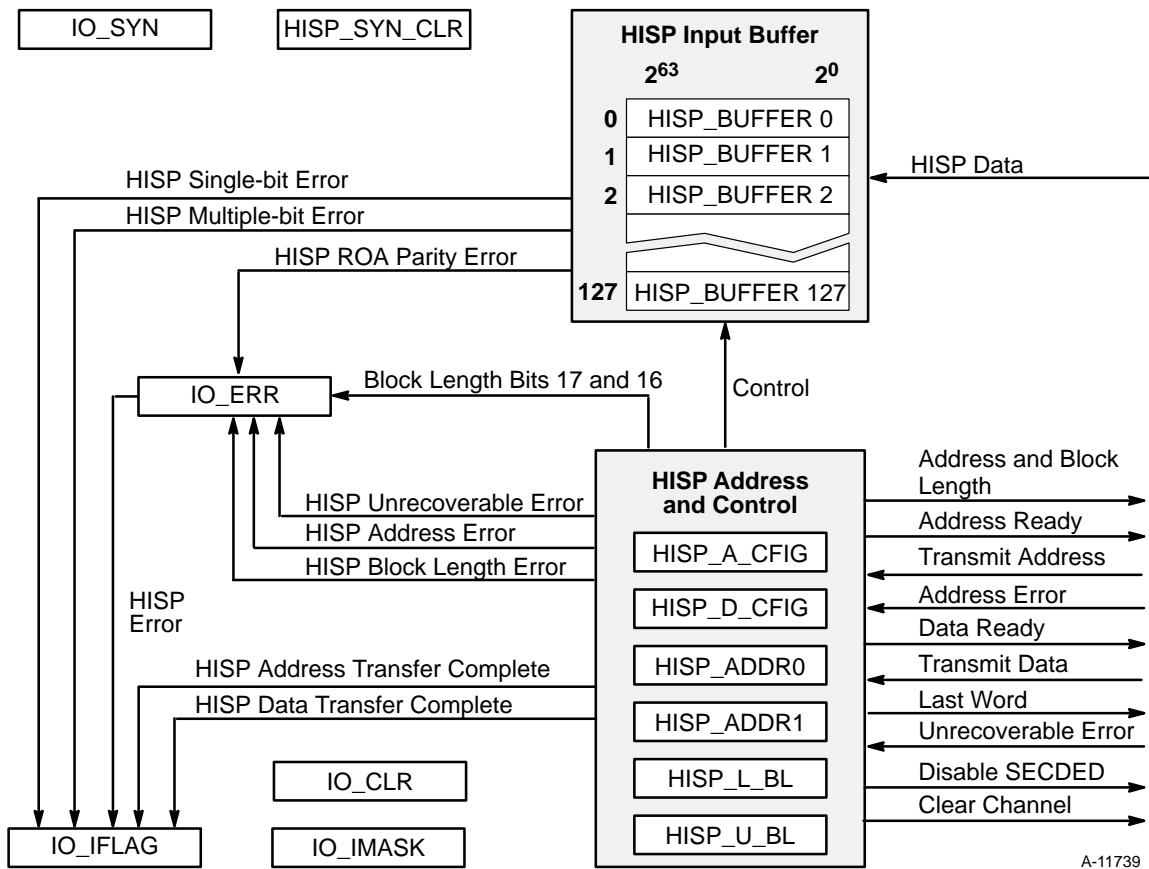


Figure 11-10. Master HISP Input Errors

HISP SECEDED Errors

The input circuitry performs SECEDED on the HISP input data before storing the data in the HISP input buffer. As the input circuitry receives each word of HISP data and check bits from the host system, the input circuitry generates a new set of check bits for the data. When the new set of check bits equals the original check bits, none of the data bits or

original check bits changed value during the transfer from the host system to the input circuitry. When the new set of check bits does not equal the original check bits, one or more of the data or check bits changed value during the transfer from the host system to the input circuitry.

When only one of the bits changed value, the input circuitry corrects the value of that bit. The input circuitry then sets the HISP single-bit error bit (bit 3) of the IO_IFLAG register to 1 and stores the syndrome for the single-bit error in the I/O syndrome (I/O_SYN) register. More information on the syndrome and I/O_SYN register is provided in “Syndrome Register” later in this section. The input circuitry then stores the data in the HISP input buffer and continues with the HISP input transfer.

After the I/O circuitry sets the HISP single-bit error bit to 1, and one or more single-bit errors occur, the HISP single-bit error bit remains set to 1. The syndrome stored in the IO_SYN register corresponds to the last single-bit error that occurred. When one or more multiple-bit errors occur, the HISP single-bit error bit resets to 0. The syndrome stored in the IO_SYN register corresponds to the first multiple-bit error that occurred.

When more than one of the data bits or original check bits changed value, the input circuitry cannot correct the value of the incorrect bits. In this case, the input circuitry sets the HISP multiple-bit error bit (bit 4) of the IO_IFLAG register to 1 and stores the syndrome for the multiple-bit error in the I/O syndrome (I/O_SYN) register. The input circuitry then stores the invalid word of data in the HISP input buffer and continues with the HISP input transfer.

After the I/O circuitry sets the HISP multiple-bit error bit to 1, and one or more single-bit errors occur, the HISP multiple-bit error bit remains set to 1. The syndrome stored in the IO_SYN register corresponds to the first multiple-bit error that occurred. When one or more multiple-bit errors occur, the HISP multiple-bit error bit remains set to 1. The syndrome stored in the IO_SYN register corresponds to the first multiple-bit error that occurred.

The microprocessor resets the HISP single-bit error and HISP multiple-bit error bits of the IO_IFLAG register to 0 by writing any value to the clear HISP syndrome register (HISP_SYN_CLR). Although the HISP SECEDED error bits of the IO_IFLAG register reset to 0, the syndrome in the IO_SYN register remains unchanged. The HISP syndrome in the IO_SYN register is valid only when one of the IO_IFLAG HISP SECEDED error bits is set to 1.

HISP ROA Parity Error

When the input circuitry stores a word of data in the HISP input buffer, it also generates parity for the data. When the BLT or the microprocessor reads a word of data out of the HISP input buffer, the input circuitry generates new parity for the data. When the new parity bits match the original parity bits, none of the HISP input data or parity bits changed value while being written into or read out of the HISP input buffer.

When the new parity bits do not match the original parity bits, the data read out of the HISP input buffer is not valid. When the data is not valid, the input circuitry sets the HISP ROA parity error bit (bit 8) of the IO_ERR register to 1. The input circuitry then continues the HISP input transfer. To reset the HISP ROA parity error bit to 0, the microprocessor must write a 1 to the clear HISP data control bit (bit 3) of the IO_CLR register.

HISP Channel Errors

The HISP channel errors are detected by the host system. There are two types of HISP channel errors: HISP address errors and HISP unrecoverable errors.

When the host system detects a parity error on the address and block length information it received, or when the host system detects an Address Ready signal that is out of synchronization, the host system sets the Address Error HISP channel signal. When the Address Error signal sets, the input circuitry sets the HISP address error bit (bit 0) of the IO_ERR register to 1, which in turn sets the HISP error bit in the IO_IFLAG register to 1 (refer again to Figure 11-10).

When the host system cannot recover from an error, the host system sets the Unrecoverable Error HISP channel signal. When the Unrecoverable Error signal sets, the input circuitry sets the HISP unrecoverable error bit (bit 1) of the IO_ERR register to 1, which in turn sets the HISP Error bit in the IO_IFLAG register to 1.

When the HISP address error bit or the unrecoverable error bit of the IO_ERR register is set to 1, the input circuitry must set the Clear Channel HISP channel signal to 1 (by writing a 1 to the clear channel HISP signal bit, bit 12, of the HISP_A_CFG register) for 150 ns or until both the Address Error and Unrecoverable Error HISP channel signals reset to 0. The HISP address error bit and the unrecoverable error bit of the IO_ERR register directly reflect the state of the Address Error and Unrecoverable Error HISP channel signals.

HISP Block Length Error

A HISP block length error occurs when the host system sets the Last Word HISP channel signal to indicate the end of a transfer and that the block length stored in the HISP_L_BL and HISP_U_BL registers has not decremented to 0. A HISP block length error also occurs when the block length stored in the HISP_L_BL and HISP_U_BL registers has decremented to 0, but the input circuitry has not received the Last Word HISP channel signal.

When a HISP block length error occurs, the input circuitry sets the HISP block length error bit (bit 2) of the IO_ERR register to 1, which also sets the HISP Error bit in the IO_IFLAG register to 1. The microprocessor resets the HISP block length error bit to 0 by writing a 1 to the clear HISP data control bit (bit 3) of the IO_CLR register.

Disable SECEDED

To disable the generation of check bits by the host system, the microprocessor can set the Disable SECEDED HISP Signal bit (bit 15) of the HISP_A_CFIG register to 1. This action sets the Disable SECEDED HISP channel signal to 1. The microprocessor should set this signal to 1 before a transfer, and it should reset the signal to 0 only after the transfer completes.

Clearing the Channel

Different components of the HISP input channel logic are cleared using selected bits in the IO_CLR register. These bits are clear HISP address control (bit 2), clear HISP data control (bit 3), and clear HISP ROA address (bit 4). More information on each of these bits is provided in “Clear I/O Register” in “Register Mapping” later in this section.

Slave Input Channel

The slave I/O gateway receives system data from an IOC over the HISP input channel. The input node in the slave I/O gateway contains the same registers as the input node in the master I/O gateway (refer again to Table 11-4); however, some of the HISP input registers in the slave I/O gateway are used differently than in the master I/O gateway.

The following subsections provide an overview of how the slave HISP input registers may be used to receive data from an IOC using the slave HISP input channel. Detailed descriptions of each register are provided in “Register Mapping” later in this section.

Initialization

Before performing HISP input transfers, the microprocessor sets parameters in the HISP_A_CFIG and HISP_D_CFIG registers that control the HISP input channel signals. The HISP_A_CFIG register contains values that set the time needed for the input circuitry to sample the address and block length from the HISP input channel. The HISP_D_CFIG register contains values that set the time needed for the input node to sample data from the HISP input channel.

Interrupt Driven Versus Polled

The microprocessor in a slave I/O gateway monitors a HISP input transfer in the same manner as the microprocessor in a master I/O gateway. After enabling the I/O and BLT hardware interrupts, the microprocessor sets bits 3 through 6 of the IO_IMASK register to 1 so the microprocessor will receive the I/O hardware interrupt when a HISP address transfer complete, HISP data transfer complete, HISP single-bit error, or HISP multiple-bit error transfer condition occurs.

Receiving the Address and Block Length

When the IOC starts the HISP input transfer, the IOC sends address and block length information to the slave I/O gateway over the Address and Block Length signals. The address points to the first location in the CRAY T3D system memory that will receive a word of data. The block length contains the total number of words that will be transferred.

As the input circuitry receives the address and block length information, it stores the information in the HISP_ADDR0 and HISP_ADDR1 registers (refer to Figure 11-11). The format of the HISP_ADDR0 and HISP_ADDR1 registers are different when the address and block length information transfers at 100 Mbytes/s or 200 Mbytes/s over the HISP channel. More information on the exact format of these registers is provided in “Register Mapping” later in this section.

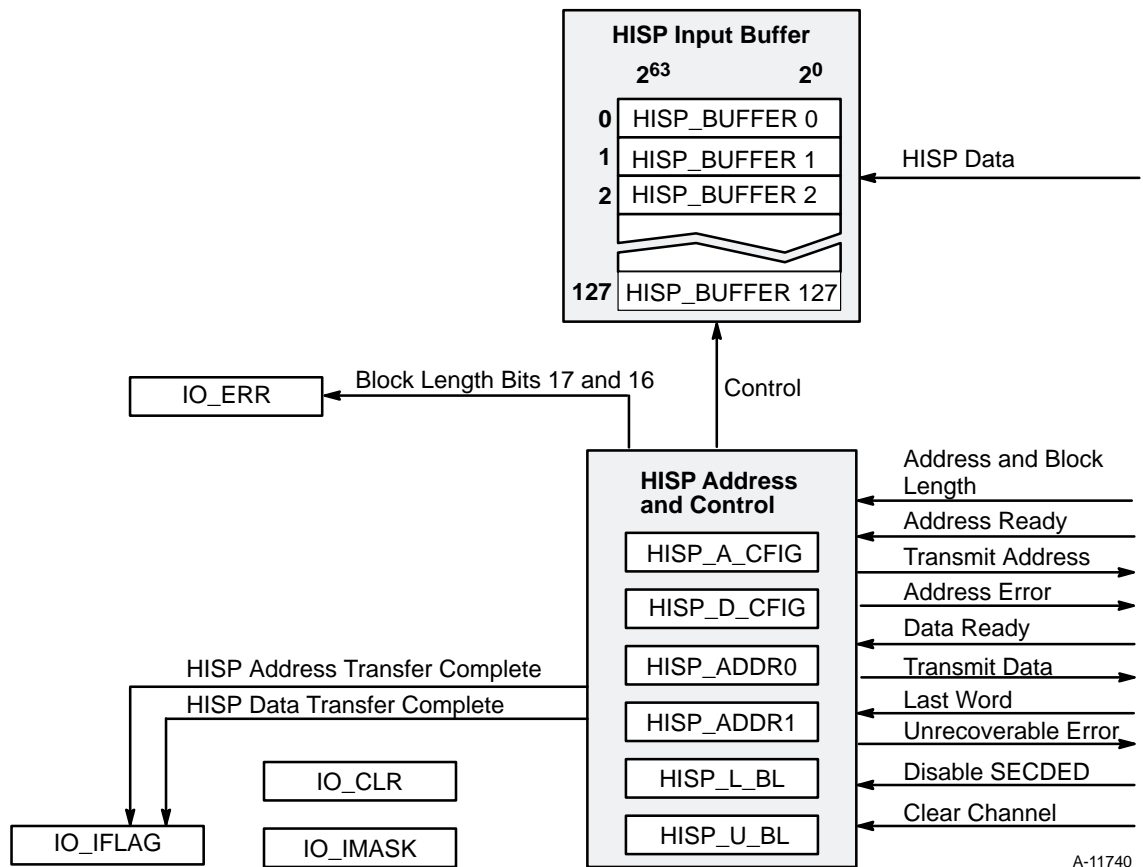


Figure 11-11. Slave HISP Input Transfer

When the address and block length transfer from the IOC to the input circuitry is complete, the input circuitry sets the HISP address transfer complete bit (bit 5) of the IO_IFLAG register. Because bit 5 of the IO_IMASK register is set to 1, the I/O hardware interrupt sets to 1 when the HISP address transfer complete bit sets to 1.

After receiving the I/O hardware interrupt, the microprocessor reads the IO_IFLAG register to determine what transfer conditions caused the interrupt to be set. When the HISP address transfer complete bit is set to 1, this bit indicates that valid address and block length information is stored in the HISP_ADDR0 and HISP_ADDR1 registers.

I/O Registers and BLT Parameters

After the microprocessor reads the HISP_ADDR0 and HISP_ADDR1 registers, it writes appropriate values to the BLT memory-mapped registers in the input node. These parameters program the BLT to transfer HISP input data from the HISP input buffer to the appropriate locations in the CRAY T3D system. Because the HISP channel transfers a fixed number of address bits with the address and block length information, the BLT is generally programmed to perform a constant stride write transfer with a stride of 1 only.

The microprocessor then writes the block length into the HISP_U_BL and HISP_L_BL registers. When the microprocessor writes data into the HISP_L_BL register, the Transmit Data signal is sent to the host system and the host system starts the transfer of data over the HISP channel.

Receiving Data

The slave I/O gateway receives HISP data in the same manner as the master I/O gateway. For more information on receiving HISP input data, refer again to “Receiving Data” under “Master Input Channel” earlier in this section.

Errors

Six types of errors may occur during a slave HISP input transfer: HISP SECEDED errors, HISP ROA parity errors, HISP block length errors, BLT errors, HISP address errors, and unrecoverable errors. Like in a master I/O gateway, these errors are reported using the IO_ERR and IO_IFLAG registers (refer to Figure 11-12) and the BLT hardware interrupt.

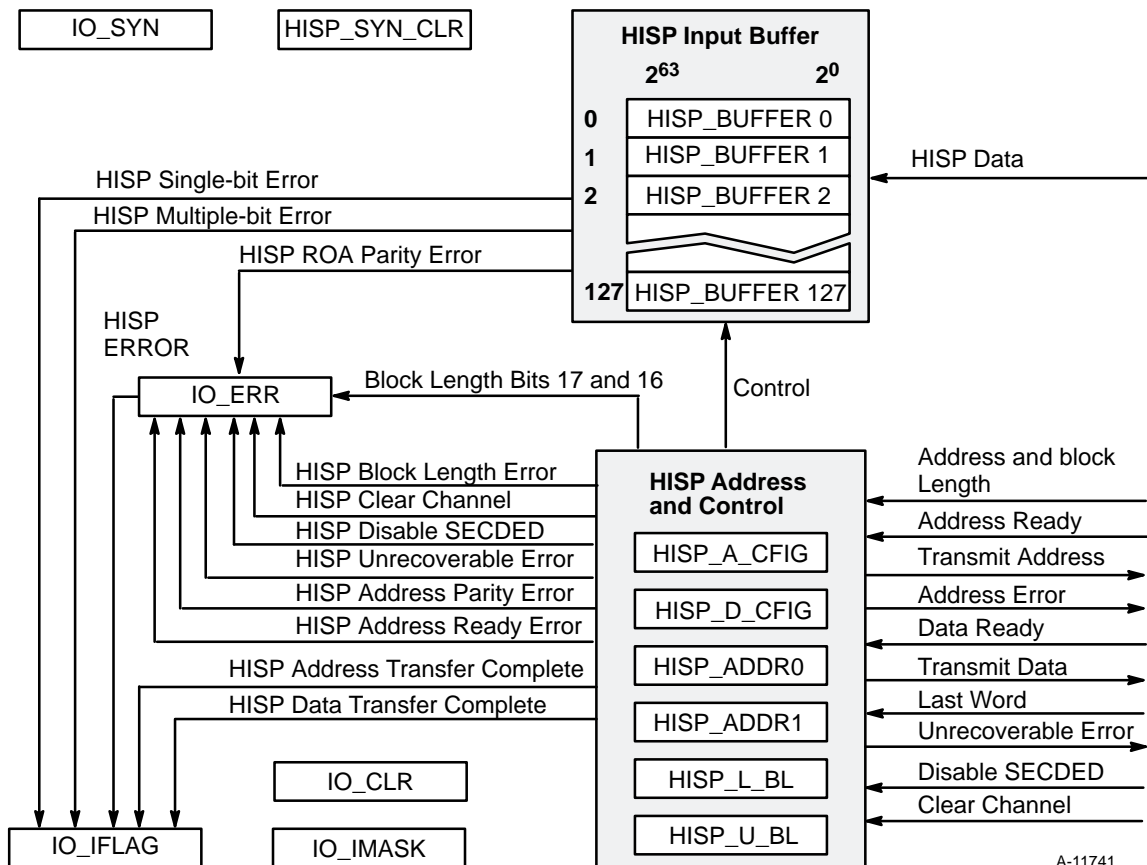


Figure 11-12. Slave HISP Input Errors

All of the errors, except HISP address errors and unrecoverable errors, are detected and reset in the same manner as in the master I/O gateway. For more information on these errors, refer to “Errors” under “Master Input Channel” earlier in this section. The following subsections describe the HISP address errors and unrecoverable errors.

Address Errors

Two types of address errors may occur when the input circuitry is receiving the address and block length from the host system: HISP address parity errors and HISP address ready errors.

As the input circuitry receives each transfer of address and block length, it generates new parity bits for the information. When the input circuitry detects a parity error, the input circuitry sets the HISP address parity error bit (bit 9) of the IO_ERR register to 1. This action also sets the HISP error bit (bit 0) of the IO_IFLAG register to 1.

When the input circuitry receives an Address Ready HISP channel signal when it was not expecting the signal, the input circuitry sets the HISP address ready error bit (bit 12) of the IO_ERR register to 1. The input circuitry also sets the HISP error bit of the IO_IFLAG register to 1.

When a HISP address parity error or a HISP address ready error occurs, the microprocessor should write a 1 to the address error HISP signal bit (bit 7) of the HISP_A_CFG register. This action sets the Address Error HISP channel signal and informs the host system that an address error occurred.

In reply to the Address Error signal, the host system sets the Clear Channel HISP channel signal to 1. When the input circuitry receives the Clear Channel signal, it sets the HISP clear channel bit (bit 0) of the IO_ERR register to 1 (refer again to Figure 11-12) and also sets the Unrecoverable Error HISP channel signal to 1.

The microprocessor clears a HISP address error by writing a 1 to the clear HISP address control bit (bit 2) of the IO_CLR register. This action resets the HISP address parity error and HISP address ready error bits of the IO_ERR register to 0.

The microprocessor then writes a 0 to the address error HISP signal bit (bit 7) and the unrecoverable error HISP signal bit (bit 8) of the HISP_A_CFG register. This action signals the host system that the address error has been reset.

Unrecoverable Errors

The input circuitry may encounter an error from which it cannot recover. When this occurs, the microprocessor writes a 1 to the unrecoverable error HISP signal bit (bit 8) of the HISP_A_CFG register. This action sets the Unrecoverable Error HISP channel signal to inform the host system of the error.

In reply to the Unrecoverable Error signal, the host system sets the Clear Channel HISP channel signal to 1. When the input circuitry receives the Clear Channel signal, it sets the HISP clear channel bit (bit 0) of the IO_ERR register to 1 (refer again to Figure 11-12).

The host system may set the Clear Channel signal even though the input circuitry has not set the Unrecoverable Error signal. When this occurs, the input circuitry, after receiving the Clear Channel signal, sets the HISP clear channel bit (bit 0) of the IO_ERR register to 1 and sets the Unrecoverable Error HISP channel signal to 1.

The Unrecoverable Error HISP channel signal remains set to 1 until the microprocessor writes a 0 to the unrecoverable error HISP signal bit (bit 8) of the HISP_A_CFG register. The microprocessor can check the value of the Unrecoverable Error HISP channel signal by reading the HISP unrecoverable error bit (bit 13) of the IO_ERR register.

The microprocessor clears an unrecoverable error by writing a 1 to the clear HISP address control and clear HISP data control bits (bits 2 and 3) of the IO_CLR register. This action resets the HISP address and data circuitry. The microprocessor then writes a 0 to the unrecoverable error HISP signal bit (bit 8) of the HISP_A_CFG register to signal the host system that the unrecoverable error has been reset.

Disable SECDED

The host system sets the Disable SECDED signal to 1 to indicate that the slave I/O gateway should not perform SECDED on HISP input data. When the host system sets the Disable SECDED signal, the input circuitry sets the HISP disable SECDED bit (bit 1) of the IO_ERR register to 1 (refer again to Figure 11-12). The HISP disable SECDED bit is reset to 0 when the host system resets the Disable SECDED HISP channel signal to 0.

Clearing the Channel

Like in the master I/O gateway, different components of the HISP input channel logic are cleared using selected bits in the IO_CLR register. These bits are clear HISP address control (bit 2), clear HISP data control (bit 3), and clear HISP ROA address (bit 4). More information on each of these bits is provided in “Clear I/O Register” in “Register Mapping” later in this section.

Master Output Channel

The master I/O gateway sends system data to the host system over the HISP output channel. Table 11-6 shows the I/O registers used by the output node in a master I/O gateway when sending HISP data to the host system.

Table 11-6. HISP Output Registers

Register Name	Definition	Description
HISP_A_CFG	HISP channel address configuration	This register sets several parameters for the HISP channel address control signals.
HISP_D_CFG	HISP data configuration	This register sets timing parameters for the HISP data control signals.
HISP_L_BL	HISP block length lower bits	This register contains bits 15 through 0, which specifies the number of 64-bit words that will be transferred (block length).
HISP_U_BL	HISP block length upper bits	This register contains bits 17 and 16, which specifies the number of 64-bit words that will be transferred (block length).
IO_IFLAG	I/O interrupt flags	This register indicates what transfer conditions have occurred that may cause the I/O hardware interrupt to be set.
IO_IMASK	I/O interrupt mask	This register indicates which transfer or error conditions will cause the I/O hardware interrupt to be set.
HISP_BUFFER 0 through 127	HISP output buffer registers	These registers make up the HISP output buffer.
IO_ERR	Error register	This register indicates what error conditions occurred during the transfer.
IO_CLR	I/O clear	This register resets the bits in the IO_IFLAG and IO_ERR registers and clears the HISP output circuitry.
HISP_ADDR0	HISP address 0	This register is used to transfer address and block length information to the host system over the HISP channel Address and Block Length signals.
HISP_ADDR1	HISP address 1	This register is also used to transfer address and block length information to the host system over the HISP Address and Block Length signals.

Table 11-6. HISP Output Registers (continued)

Register Name	Definition	Description
IO_SYN	Syndrom register	This register contains the syndrome bits for a microprocessor single-bit or multiple-bit error.
DEC_SYN_CLR	Clear microprocessor syndrom register	This register clears microprocessor single-bit or multiple-bit errors.

Requesting a Transfer

To request that data be sent to the host system, the PE in a processing element node sends information to the PE in the output node of a master I/O gateway. The request information contains the following parameters.

- Number of 64-bit words that will be sent
- Address in the host system where the first word will be stored
- Address in a PE where the first word of the transfer resides
- Stride value and transfer type (constant stride or gather)

The microprocessor in the output node PE then interprets this information and generates values for the output node memory-mapped registers.

Initialization

Before performing HISP output transfers, the microprocessor sets parameters in the HISP_A_CFG and HISP_D_CFG registers that control the HISP output channel signals. The HISP_A_CFG register contains values that set the pulse width and period of the Address Ready signal and set the format of the address and block length information. The HISP_D_CFG register contains values that set parameters for the Data Ready HISP channel signal and other data transfer parameters.

Interrupt Driven Versus Polled

The microprocessor in the output node of a master I/O gateway monitors the HISP output transfer by enabling the I/O and BLT hardware interrupts (interrupt driven) or by periodically checking the value of its associated IO_IFLAG register and BLT_SR (polled). The following subsections describe a HISP output transfer where the I/O and BLT hardware interrupts are enabled.

The microprocessor must set bit 13 of the HIER register to 1 to receive the I/O hardware interrupt and must set bit 12 of the HIER register to 1 to receive the BLT hardware interrupt. In addition, the microprocessor sets bits 6, 2, and 1 of the IO_IMASK register to 1 so the microprocessor will receive the I/O hardware interrupt when any of the following transfer conditions occur. (More information on each transfer condition is provided in the following subsections.)

- HISP data transfer complete
- Microprocessor single-bit error
- Microprocessor multiple-bit error

BLT Parameters

The BLT in the output node is used to transfer HISP data from the source PEs to the HISP output buffer. Before performing a HISP output transfer, the microprocessor writes the appropriate values to the BLT memory-mapped registers. For example, Table 11-7 shows sample BLT register values for a HISP output transfer. For more information on the BLT registers, refer again to Section 8, “Block Transfer Engine.”

Table 11-7. Sample BLT Output Parameters

Register Name	Value	Description
BLT_LAR	7000000 ₁₆	Set the local address to the address of the HISP_BUFFER 0 register, which is the first entry of the HISP output buffer
BLT_LSR	1	Set the local stride to 1 word
BLT_RMR	FFF000000 ₁₆	Set the remote mask so bits 35 through 24 of the remote index are the PE number and bits 23 through 0 are the index offset
BLT_RIR	001000000 ₁₆	Set the remote index to a destination PE number of I/O = 0, Z = 0, Y = 0, X = 0, PE = 1, and an index offset of 0
BLT_RSR	1	Set the remote stride to 1 word
BLT_VLR	10000 ₁₆ (block length)	Set the vector length register so it matches the block length stored in the HISP_L_BL and HISP_U_BL registers †
BLT_RBR	88000 ₁₆	Set the remote base to an offset of 88000 ₁₆
BLT_RLR	97FFF ₁₆	Set the remote limit to an offset of 97FFF ₁₆
BLT_CR	18 ₁₆	Set the control register so that a constant stride read transfer is selected, the large packet size is disabled, logical PE numbers are enabled, the I/O mode is enabled, the 256 outstanding request count is enabled, and the PE range mask is set to 0
BLT_SR	5	Set the status register so that the microprocessor receives the BLT hardware interrupt when the BLT transfer is complete or when a BLT error occurs

† Because the vector length indicates the number of large packets (4-word body) the BLT generates, set the vector length to the block length divided by four. The block length must be a multiple of four in order to use large packet transfers. For more information on large packet transfer restrictions, refer to "Enable Large Packet Size" in Section 8, "Block Transfer Engine."

As soon as the microprocessor writes to the BLT_CR, the BLT starts a transfer; however, because the BLT I/O mode is enabled, the BLT does not actually transfer data until it receives handshaking control signals from the HISP output circuitry.

I/O Registers

The microprocessor loads the HISP_L_BL and HISP_U_BL register with the number of 64-bit words that will be transferred (block length). The microprocessor then loads the address and block length information into the HISP_ADDR0 and HISP_ADDR1 registers (refer to Figure 11-13).

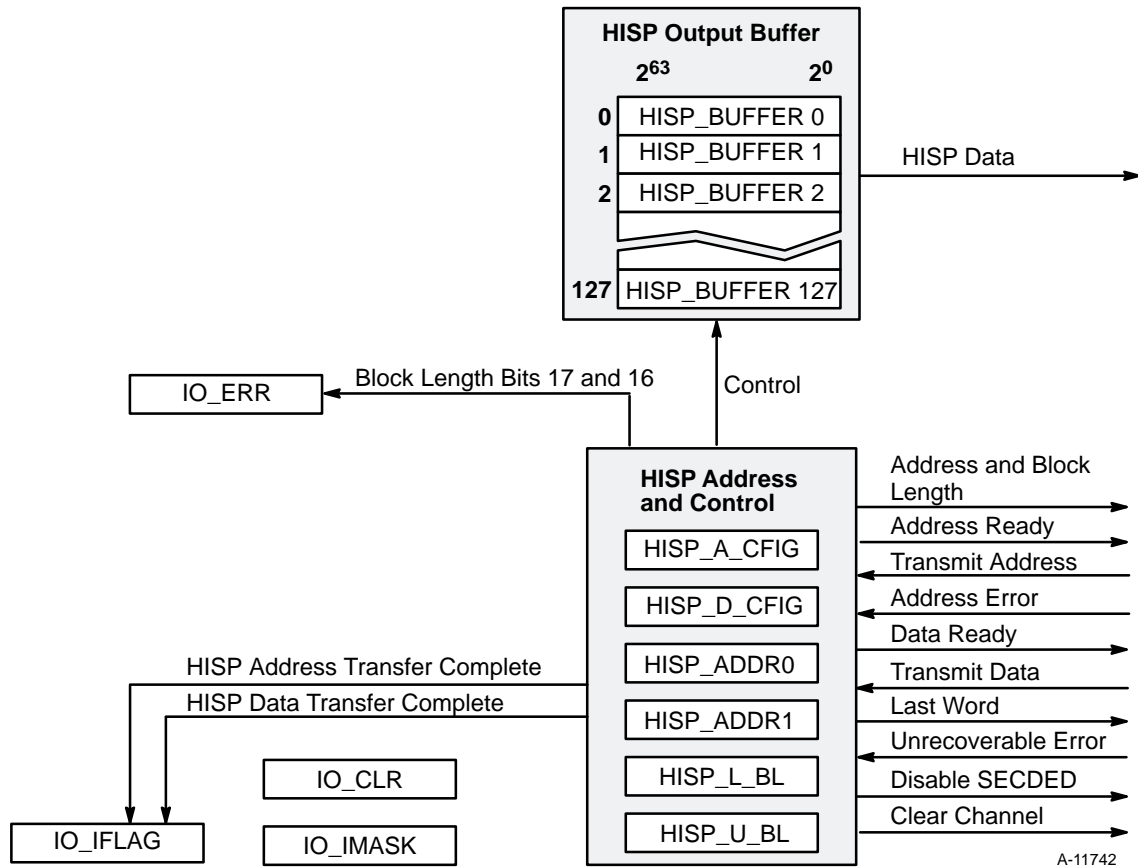


Figure 11-13. Master HISP Output Channel

Sending the Address and Block Length

As soon as the microprocessor writes address and block length information into the HISP_ADDR0 register, the output circuitry starts the HISP output transfer. The output circuitry first transfers the address and block length from the HISP_ADDR0 and HISP_ADDR1 registers to the host system over the Address and Block Length HISP channel signals.

When the output node completes the transfer of address and block length information to the host system, the output circuitry sets the HISP address transfer complete bit (bit 5) of the IO_IFLAG register to 1. This indicates that the address transfer is complete, but does not set the I/O hardware interrupt because bit 5 is not set to 1 in the IO_IMASK register.

Sending Data

The output circuitry sends HISP output data to the host system over the HISP data signals. When valid data is on the data signals, the output circuitry sets the Data Ready signal. After receiving the data, the host system sets the Transmit Data signal to indicate that the output circuitry may transmit new data. The output circuitry also decrements the block length stored in the HISP_L_BL and HISP_U_BL registers as it sends each word of HISP output data to the host system.

After receiving data from the BLT, the output circuitry temporarily stores each word of data in the HISP output buffer. Like the input buffer, the HISP output buffer is actually a set of 128 memory-mapped registers (HISP_BUFFER 0 through HISP_BUFFER 127) that operate as a circular buffer (refer again to Figure 11-13).

The BLT exchanges handshaking signals with the HISP output buffer on 32-word boundaries (four 32-word blocks). At the start of a transfer, the BLT writes data into the first 32-word block in the output buffer. The BLT creates read request packets, sends the packets to the source PE(s), receives 1- or 4-word response packets, and stores the data in the HISP output buffer.

When the BLT fills a 32-word block before the output circuitry reserves the next block of data, the BLT waits until the output circuitry reserves the next block of data. When this occurs, the output circuitry signals the BLT that the next block of data is reserved and the BLT can write data into that block.

When the BLT writes a word into the last entry of the HISP output buffer and the vector length has not decremented to 0, the BLT continues to increment the local address by the local stride. Because the output circuitry does not use bits 7 through 11 of the HISP_BUFFER address, this action causes the BLT to write data into the first entry of the HISP output buffer. When the first entry of the HISP output buffer is in a block that is not reserved, the BLT must wait until that block is reserved before writing data into the buffer.

The HISP output circuitry exchanges handshaking signals with the HISP output buffer on 16-word boundaries (eight 16-word blocks). After the BLT stores information into the first 16-word block, the output circuitry reads words of data from the output buffer and sends them to the HISP output channel. After reading all of the data from a block, the output circuitry increments to the next block.

When the next block of data is full or contains the last word of the transfer, the output circuitry reads data from the block and transfers the data to the HISP output channel. When the next block of data is not full, the output circuitry waits until the BLT fills the block or until the BLT writes the last word of the transfer into that block. When this occurs, the output circuitry reads data from the block and transfers the data to the HISP output channel.

When the output circuitry reads a word from HISP_BUFFER 127 (last entry of the HISP output buffer) and the block length has not decremented to 0, the output circuitry resets the internal buffer address. This action causes the output circuitry to read the next word of HISP output data from HISP_BUFFER 0 (first entry of the HISP output buffer). When HISP_BUFFER 0 does not contain valid data, the output circuitry waits until the first HISP buffer block (HISP_BUFFER 0 through HISP_BUFFER 15) is full or contains the last word of the transfer before reading data from HISP_BUFFER 0.

The BLT continues to create read request packets and send them to the source PE(s) until the vector length decrements to 0. After the BLT receives a response packet for all of the request packets it generated during the transfer, the BLT sets the BLT hardware interrupt. The BLT hardware interrupt indicates that the BLT has received all of the HISP output data from the PE(s) and has stored the data in the HISP output buffer.

The output circuitry continues to read data from the HISP output buffer and send the data to the HISP output channel until the block length decrements to 0. When the block length decrements to 0, it indicates that the output circuitry has read the last word of data for the transfer out of the HISP output buffer. When this occurs, the HISP output circuitry sets the Last Word HISP channel signal to indicate to the host system that the data signals contain the last word of the transfer.

The BLT in an output node does not have to be used to transfer data from the source PE(s) to the HISP output buffer. Instead, the microprocessor in the output node may write data into the HISP output buffer. The microprocessor uses different addresses to write data into the HISP output

buffer than the BLT uses. More information on the microprocessor addressing is provided in “HISP Buffer Registers” in “Register Mapping” later in this section.

Errors

Five types of errors may occur during a master HISP output transfer: microprocessor SECEDED errors, HISP ROA parity errors, HISP channel errors, HISP block errors, and BLT errors. These errors are reported using the IO_ERR and IO_IFLAG registers (refer to Figure 11-14) and the BLT hardware interrupt. The following subsections describe each error.




Figure 11-14. Master HISP Output Errors

Some of the HISP output errors indicate that the HISP output transfer should be terminated. To terminate the transfer, the microprocessor performs three steps.

First, the microprocessor aborts the BLT transfer. To do this, the microprocessor writes a status select code of two into the BLT_SR. For more information on the BLT abort code, refer to “BLT Abort” in Section 8, “Block Transfer Engine.”

Second, the microprocessor terminates the transfer with the host system. To do this, the microprocessor writes a 1 to the clear channel HISP signal bit (bit 12) of the HISP_A_CFIG register. This bit must be set to 1 for 150ns or until the host system resets the Unrecoverable Error HISP channel signal.

Third, the microprocessor clears the HISP output circuitry. To do this, the microprocessor writes a 1 to the clear HISP address control and clear HISP data control bits (bits 2 and 3) of the IO_CLR register.

Microprocessor SECDDED Errors

The output circuitry performs SECDDED on the HISP output data before writing the data into the HISP output buffer. Microprocessor SECDDED errors are detected, reported, and reset in the same manner described in “Microprocessor SECDDED Errors” under “LOSP Channel” earlier in this section.

HISP ROA Parity Error

When the BLT or microprocessor stores a word of data in the HISP output buffer, the output circuitry generates parity for the data. When the output circuitry reads a word of data out of the HISP output buffer, the output circuitry generates new parity for the data. When the new parity bits match the original parity bits, none of the HISP output data or parity bits changed value while being written into or read out of the HISP output buffer.

When the new parity bits do not match the original parity bits, the data read out of the HISP output buffer is not valid. When this occurs, the output circuitry sets the HISP ROA parity error bit (bit 8) of the IO_ERR register to 1. The output circuitry then continues the HISP output transfer. To reset the HISP ROA parity error bit to 0, the microprocessor must write a 1 to the clear HISP data control bit (bit 3) of the IO_CLR register.

HISP Channel Errors

The HISP channel errors are detected by the host system. These errors are detected, reported, and reset in the same manner described in “HISP Channel Errors” under “Master Input Channel” earlier in this section.

HISP Block Error

Two conditions will cause a HISP block error. One condition occurs when the microprocessor or BLT attempts to write a word into one of the four 32-word blocks in the HISP output buffer before the output circuitry has reserved that block for the write. The other condition occurs when the output circuitry attempts to read a word out of a 16-word block in the HISP output buffer before the block has been filled or contains the last word of the transfer.

When a HISP block error occurs, the output circuitry sets the HISP block error bit (bit 2) of the IO_ERR register to 1. This action also sets the HISP error bit (bit 0) of the IO_IFLAG register to 1. To reset the HISP block error bit to 0, the microprocessor writes a 1 to the clear HISP data control bit (bit 3) of the IO_CLR register.

BLT Error

When the BLT detects an error, it sets the BLT hardware interrupt to 1. For more information on the BLT errors, refer to “Status Register” in Section 8, “Block Transfer Engine.”

Disable SECDDED

To disable SECDDED in the host system, the microprocessor can set the Disable SECDDED HISP signal bit (bit 15) of the HISP_A_CFG register to 1. This action sets the Disable SECDDED HISP channel signal to 1. The microprocessor should set this signal to 1 before a transfer, and it should reset the signal to 0 only after the transfer completes.

Clearing the Channel

Different components of the HISP output channel logic are cleared using selected bits in the IO_CLR register. These bits are clear HISP address control (bit 2), clear HISP data control (bit 3), and clear HISP ROA address (bit 4). More information on each of these bits is provided in “Clear I/O Register” in “Register Mapping” later in this section.

Slave Output Transfers

The slave I/O gateway sends system data to an IOC over the HISP output channel. The output node in the slave I/O gateway contains the same registers as the output node in the master I/O gateway (refer again to Table 11-6); however, some of the HISP output registers in the slave I/O gateway are used differently than in the master I/O gateway.

The following subsections provide an overview of how the slave HISP output registers may be used to send data to an IOC using the slave HISP output channel. Detailed descriptions of each register are provided in “Register Mapping” later in this section.

Initialization

Before performing HISP output transfers, the microprocessor sets parameters in the HISP_A_CFIG and HISP_D_CFIG registers that control the HISP output channel signals. The HISP_A_CFIG register contains values that set the time needed for the output circuitry to sample the address and block length from the HISP channel. The HISP_D_CFIG register contains values that set parameters for the Data Ready HISP channel signal and other data transfer parameters

Interrupt Driven Versus Polled

The microprocessor in the output node of a slave I/O gateway monitors a HISP output transfer in the same manner as the microprocessor in a master I/O gateway. After enabling the I/O and BLT hardware interrupts, the microprocessor sets bits 6, 5, 2, and 1 of the IO_IMASK register to 1 so the microprocessor will receive the I/O hardware interrupt when any of the following transfer conditions occur. More information on each transfer condition is provided in the following subsections.

- HISP Address Transfer Complete
- HISP Data Transfer Complete
- Microprocessor Single-bit Error
- Microprocessor Multiple-bit Error

Receiving the Address and Block Length

When the IOC starts the HISP output transfer, the IOC sends address and block length information to the slave I/O gateway over the Address and Block Length signals. The address points to the location of the first word in the CRAY T3D system memory that will be transferred. The block length contains the total number of words that will be transferred.

As the output circuitry receives the address and block length information, it stores the information in the HISP_ADDR0 and HISP_ADDR1 registers (refer to Figure 11-15). The format of the HISP_ADDR0 and HISP_ADDR1 registers is different when the address and block length information transfers at 100 Mbytes/s than when the address and block length information transfers at 200 Mbytes/s over the HISP channel. More information on the exact format of these registers is provided in “Register Mapping” later in this section.

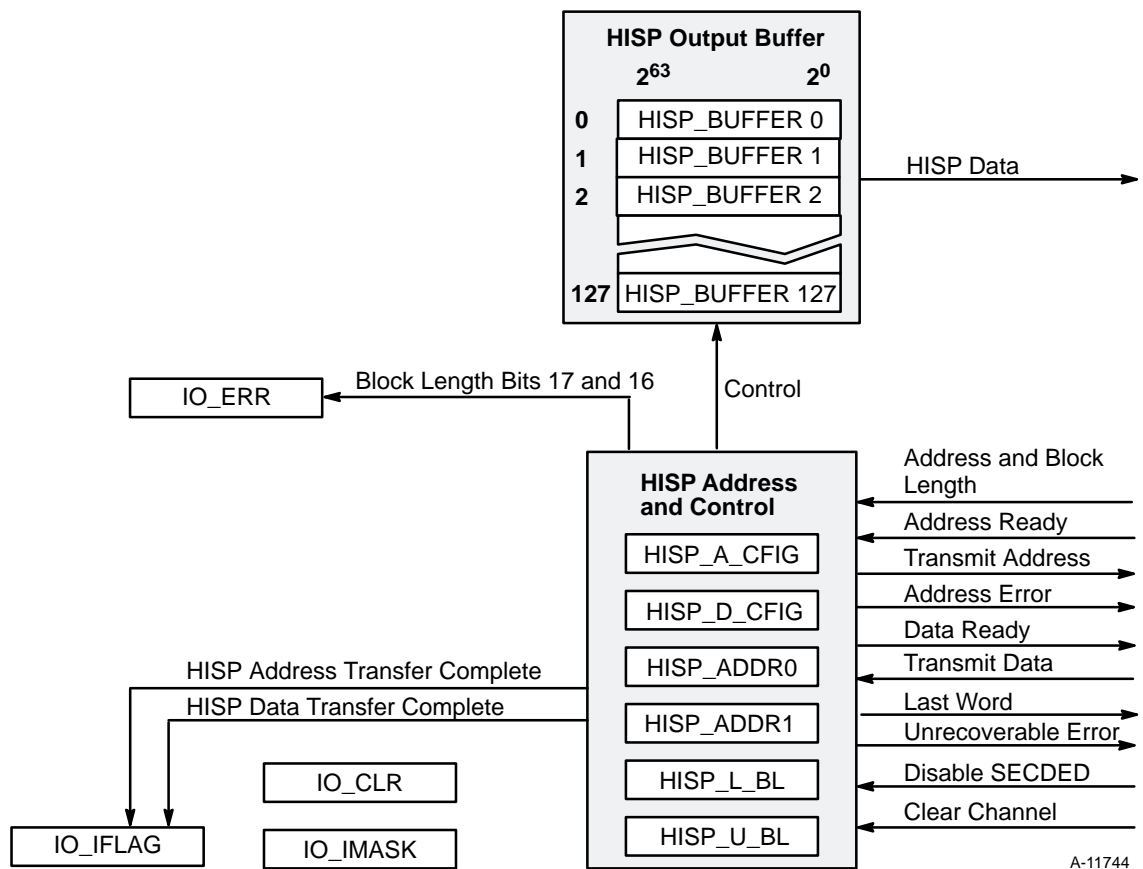


Figure 11-15. Slave HISP Output Transfer

When the address and block length transfer from the IOC to the output circuitry is complete, the output circuitry sets the HISP address transfer complete bit (bit 5) of the IO_IFLAG register. Because bit 5 of the IO_IMASK register is set to 1, the I/O hardware interrupt sets to 1 when the HISP address transfer complete bit sets to 1.

After receiving the I/O hardware interrupt, the microprocessor reads the IO_IFLAG register to determine what transfer conditions caused the interrupt to be set. When the HISP address transfer complete bit is set to 1, this bit indicates that valid address and block length information is stored in the HISP_ADDR0 and HISP_ADDR1 registers.

I/O Registers and BLT Parameters

After the microprocessor reads the HISP_ADDR0 and HISP_ADDR1 registers, it writes the block length into the HISP_L_BL and HISP_U_BL registers. The microprocessor then interprets the address information. When the microprocessor reads the HISP_ADDR0 register, the output circuitry reserves the first 32-word block in the HISP output buffer for the BLT.

After interpreting the address information, the microprocessor writes the appropriate values to the BLT memory-mapped registers in the output node. These parameters program the BLT to transfer data from the source PE(s) to the HISP output buffer. Because the HISP channel transfers a fixed number of address bits with the address and block length information, the BLT is generally programmed to perform a constant stride read transfer with a stride of 1 only.

Sending Data

The slave I/O gateway sends HISP data to the HISP output channel in the same manner as the master I/O gateway. For more information on sending HISP output data, refer again to “Sending Data” under “Master Output Channel” earlier in this section.

Errors

Seven types of errors may occur during a slave HISP output transfer: microprocessor SECEDED errors, HISP ROA parity errors, HISP channel errors, HISP block errors, BLT errors, HISP address errors, and unrecoverable errors. Like in a master I/O gateway, these errors are reported using the IO_ERR and IO_IFLAG registers (refer to Figure 11-16) and the BLT hardware interrupt.

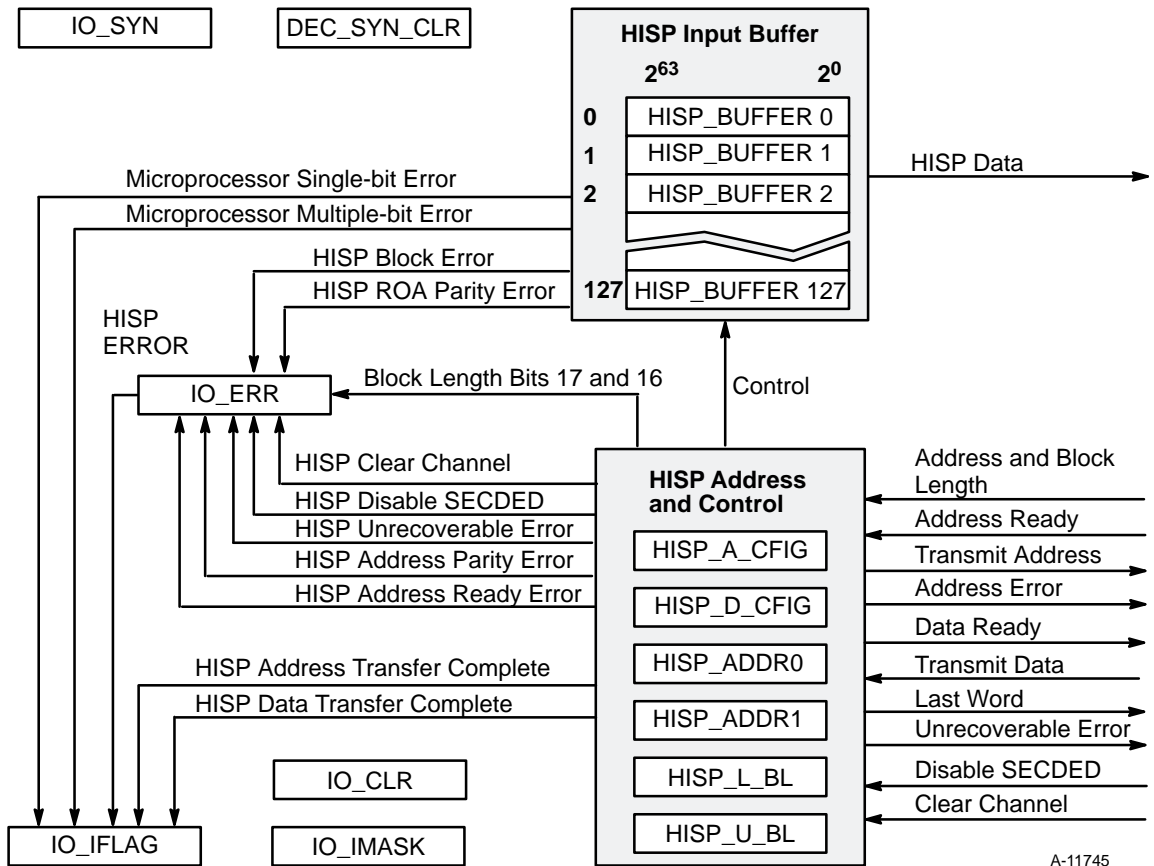


Figure 11-16. Slave HISP Output Errors

All of the errors, except HISP address errors and unrecoverable errors, are detected and reset in the same manner as in the master I/O gateway. For more information on these errors, refer to “Errors” under “Master Output Channel” earlier in this section.

The HISP address errors and unrecoverable errors are detected and reset in the same manner as in the input node of the slave I/O gateway. For more information on these errors, refer to “Address Errors” and “Unrecoverable Errors” under “Slave Input Channel” earlier in this section.

Disable SECDED

The host system sets the Disable SECDED signal to 1 in order to indicate that the slave I/O gateway should not generate check bits for HISP output data. When the output circuitry does not generate check bits, the output circuitry sets the value of the check bits to 0. When the host system sets the Disable SECDED signal, the output circuitry sets the HISP disable SECDED bit (bit 1) of the IO_ERR register to 1 (refer again to Figure 11-16). The HISP disable SECDED bit is reset to 0 when the host system resets the Disable SECDED HISP channel signal to 0.

Clearing the Channel

Like in the master I/O gateway, different components of the HISP output channel logic are cleared using selected bits in the IO_CLR register. These bits are clear HISP address control (bit 2), clear HISP data control (bit 3), and clear HISP ROA address (bit 4). More information on each of these bits is provided in “Clear I/O Register” in “Register Mapping” later in this section.

Register Mapping

The following subsections describe the addressing and bit assignments for the I/O memory-mapped registers (hereafter referred to as registers) used in the CRAY T3D system. Each subsection also provides a brief summary of the function of the register.

Table 11-8 is a summary of the input/output registers and their names. It also lists the partial physical address of each register as they appear on the address pins of the microprocessor.

Table 11-8. Register Assignments for I/O Node

Address	Register Name	Direction	Description
4000200 ₁₆	IO_IFLAG	Read	Interrupt flags
4000220 ₁₆	IO_IMASK	Read or write	Interrupt masks
4000C00 ₁₆	IO_ERR	Read	Error register
4001E00 ₁₆	IO_CLR	Write	I/O clear
4000000 ₁₆	LOSP_CFIG	Write	LOSP channel configuration registers
4000400 ₁₆	LOSP_DI	Read	LOSP data in
4000800 ₁₆	LOSP_DO	Write	LOSP data out
4000A00 ₁₆	LOSP_DISC	Write	LOSP out disconnect
4000020 ₁₆	HISP_A_CFIG	Write	HISP address configuration registers
4000040 ₁₆	HISP_D_CFIG	Write	HISP data configuration registers
4001A00 ₁₆	HISP_L_BL	Read or write	HISP lower block length (bits 15 – 0)
4001800 ₁₆	HISP_U_BL	Write	HISP upper block length (bits 17 and 16)
4001400 ₁₆	HISP_ADDR0	Read or write	HISP address 0
4001600 ₁₆	HISP_ADDR1	Read or write	HISP address 1
6000000 ₁₆	HISP_BUFFER	Read or write	HISP buffer base address
4001200 ₁₆	HISP_ROA	Write	HISP output ram-on-array address enter
4001000 ₁₆	IO_SYN	Read	DEC and HISP syndrome
4001C00 ₁₆	DEC_SYN_CLR	Write	Clear DEC syndrome
4001C20 ₁₆	HISP_SYN_CLR	Write	Clear HISP syndrome
4000C00 ₁₆	HB_GEN	Write	Heartbeat generation

Interrupt Flags Register (Address 4000200₁₆)

The interrupt flags (IO_IFLAG) register is a read-only register that indicates which transfer or error conditions have occurred. The values in the IO_IFLAG register always indicate the current I/O channel conditions and are not affected by the interrupt mask register.

Figure 11-17 shows the bit assignments for the IO_IFLAG register address as they appear on the address pins of the microprocessor, and Table 11-9 shows the bit format of the IO_IFLAG register.

HEX	0	4				0				0				0				2				0				0			
Binary	2 ²⁸	2 ²⁷	2 ²⁶	2 ²⁵	2 ²⁴	2 ²³	2 ²²	2 ²¹	2 ²⁰	2 ¹⁹	2 ¹⁸	2 ¹⁷	2 ¹⁶	2 ¹⁵	2 ¹⁴	2 ¹³	2 ¹²	2 ¹¹	2 ¹⁰	2 ⁹	2 ⁸	2 ⁷	2 ⁶	2 ⁵	2 ⁴	2 ³	2 ²	2 ¹	2 ⁰
	0	0	1	0	x	x	x	x	x	x	x	x	x	x	x	0	0	0	0	1	x	x	x	0	x	0	0	0	0

x = Don't Care

Figure 11-17. IO_IFLAG Register Address Bit Format

Table 11-9. IO_IFLAG Register Bit Format

Bits	Name	Nodes
0	HISP error	Input and output
1	Microprocessor single-bit error	Input and output
2	Microprocessor multiple-bit error	Input and output
3	HISP single-bit error	Input
4	HISP multiple-bit error	Input
5	HISP address transfer complete	Input and output
6	HISP data transfer complete	Input and output
7	LOSP input data available	Input and output
8	LOSP input disconnect	Input and output
9	LOSP input error	Input and output
10	LOSP output complete	Input and output
11	LOSP output error	Input and output
12	Master clear in	Input and output
14 – 13	Revision level	Input and output
15	Master or slave I/O node	Input and output
63 – 16	These bits are not used	Input and output

HISP Error (Bit 0)

This bit sets to 1 when specified error conditions occur during a HISP input or output data transfer. This bit is set for different error conditions in a master I/O gateway than it is in a slave I/O gateway. In a master I/O gateway, the HISP error bit sets to 1 when any of the following conditions occur.

- The HISP address error bit (bit 0) of the IO_ERR register sets to 1
- The HISP unrecoverable error bit (bit 1) of the IO_ERR register sets to 1
- The HISP block length error bit (input node, bit 2) or HISP block error bit (output node, bit 2) of the IO_ERR register sets to 1
- The HISP ROA parity error bit (bit 8) of the IO_ERR register sets to 1

The HISP Error bit in a master I/O gateway is reset to 0 only when the values of bits 0, 1, 2, and 8 of the IO_ERR register are 0. Bits 0 and 1 of the IO_ERR register are reset to 0 when the host system resets the Address Error and Unrecoverable Error HISP channel signals to 0. Bits 2 and 8 of the IO_ERR register are reset to 0 when the microprocessor writes a 1 to the clear HISP data control bit (bit 3) of the IO_CLR register.

In a slave I/O gateway, the HISP error bit sets to 1 when any of the following conditions occur.

- The HISP clear channel bit (bit 0) of the IO_ERR register changes state from 0 to 1 or from 1 to 0
- The HISP disable SECDED bit (bit 1) of the IO_ERR register changes state from 0 to 1 or from 1 to 0 (input node only)
- The HISP block length error bit (input node, bit 2) or HISP block error bit (output node, bit 2) of the IO_ERR register sets to 1
- The HISP ROA parity error bit (bit 8) of the IO_ERR register sets to 1
- The HISP address parity error bit (bit 9) of the IO_ERR register sets to 1
- The HISP address ready error bit (bit 12) of the IO_ERR register sets to 1

The HISP error bit in a slave I/O gateway is reset to 0 only when the values of bits 2, 8, 9, and 12 of the IO_ERR register are 0 and the microprocessor writes a 1 to the clear HISP data control bit (bit 3) of the IO_CLR register. Bits 2 and 8 of the IO_ERR register are reset to 0 when the microprocessor writes a 1 to the clear HISP data control bit (bit 3) of the IO_CLR register. Bits 9 and 12 of the IO_ERR register are reset to 0 when the microprocessor writes a 1 to the clear HISP address control bit (bit 2) of the IO_CLR register.

Microprocessor Single-bit Error (Bit 1)

This bit sets to 1 when the LOSP circuitry detects a single-bit error when a microprocessor writes data into the LOSP_DO register or when the HISP output node detects a single-bit error when the microprocessor or a BLT transfer writes data into the HISP output buffer. The microprocessor single-bit error bit is reset to 0 when the microprocessor writes any value to the DEC_SYN_CLR register.

When set to 1, and one or more single-bit errors occur, the microprocessor single-bit errors bit remains set to 1. The syndrome stored in the IO_SYN register corresponds to the last single-bit error that occurred.

When set to 1, and one or more multiple-bit errors occur, the microprocessor single-bit error bit resets to 0. The syndrome stored in the IO_SYN register corresponds to the first multiple-bit error that occurred.

Microprocessor Multiple-bit Error (Bit 2)

This bit sets to 1 when the LOSP circuitry detects a multiple-bit error when a microprocessor writes data into the LOSP_DO register or when the output node detects a multiple-bit error when the microprocessor or a BLT transfer writes data into the HISP output buffer. The microprocessor multiple-bit error bit is reset to 0 when the microprocessor writes any value to the DEC_SYN_CLR register.

When set to 1, and one or more single-bit errors occur, the microprocessor multiple-bit errors bit remains set to 1. The syndrome stored in the IO_SYN register corresponds to the first multiple-bit error that occurred.

When set to 1, and one or more multiple-bit errors occur, the microprocessor multiple-bit error bit remains set to 1. The syndrome stored in the IO_SYN register corresponds to the first multiple-bit error that occurred.

HISP Single-bit Error (Bit 3)

In the input node, this bit sets to 1 when the HISP input circuitry detects a single-bit error while performing SECDDED on HISP input data. The input circuitry performs SECDDED after receiving data from the HISP input channel and before storing the data in the HISP input buffer. The HISP single-bit error bit is reset to 0 when the microprocessor writes any value to the HISP_SYN_CLR register. In the output node, the HISP single-bit error bit is not used and is always set to 0.

When set to 1, and one or more single-bit errors occur, the HISP single-bit errors bit remains set to 1. The syndrome stored in the IO_SYN register corresponds to the last single-bit error that occurred.

When set to 1, and one or more multiple-bit errors occur, the HISP single-bit error bit resets to 0. The syndrome stored in the IO_SYN register corresponds to the first multiple-bit error that occurred.

HISP Multiple-bit Error (Bit 4)

In the input node, this bit sets to 1 when the HISP input circuitry detects a multiple-bit error while performing SECDDED on HISP input data. The input circuitry performs SECDDED after receiving data from the HISP input channel and before storing the data in the HISP input buffer. The HISP multiple-bit error bit is reset to 0 when the microprocessor writes any value to the HISP_SYN_CLR register. In the output node, the HISP multiple-bit error bit is not used and is always set to 0.

When set to 1, and one or more single-bit errors occur, the HISP multiple-bit errors bit remains set to 1. The syndrome stored in the IO_SYN register corresponds to the first multiple-bit error that occurred.

When set to 1, and one or more multiple-bit errors occur, the HISP multiple-bit error bit remains set to 1. The syndrome stored in the IO_SYN register corresponds to the first multiple-bit error that occurred.

HISP Address Transfer Complete (Bit 5)

In a master I/O gateway, this bit sets to 1 when the HISP circuitry completes the transfer of address and block length information to the host system. In a slave I/O gateway, this bit sets to 1 when the HISP circuitry receives the address and block length information from the host system and stores the information in the HISP_ADDR1 and HISP_ADDR0

registers. The HISP address transfer complete bit is reset to 0 when the microprocessor writes a 1 to the clear HISP address control bit (bit 2) of the IO_CLR register.

HISP Data Transfer Complete (Bit 6)

This bit sets to 1 when the last word HISP channel signal sets to 1 or when the block length stored in the HISP_U_BL and HISP_L_BL registers decrements to 0. The HISP data transfer complete bit resets to 0 when the microprocessor writes a 1 to the clear HISP data control bit (bit 3) of the IO_CLR register.

LOSP Input Data Available (Bit 7)

This bit sets to 1 when the I/O gateway receives the LOSP input channel Disconnect signal or when the LOSP circuitry completely fills the LOSP input buffer with data. LOSP input data may remain in the LOSP input buffer for an extended period of time; however, the LOSP channel bandwidth will be degraded. The LOSP input data available bit is reset to 0 by the following methods.

- Reading the LOSP_DI register
- Writing a 1 to the clear LOSP input control bit (bit 0) of the IO_CLR register

LOSP Input Disconnect (Bit 8)

This bit sets to 1 when the microprocessor reads the last word of LOSP input data from the LOSP_DI register. The microprocessor must read the valid parcel count (bits 5 through 7) from the IO_ERR register to determine the number of valid parcels (0 to 4) that were read from the LOSP_DI register. The LOSP input disconnect bit is reset to 0 by the following methods.

- Writing a 1 to the clear LOSP input control bit (bit 0) of the IO_CLR register
- Writing a 1 to the clear LOSP input disconnect bit (bit 9) of the IO_CLR register

LOSP Input Error (Bit 9)

This bit sets to 1 when the LOSP Input Parity Error bit (bit 3) or the LOSP Input Sequence Error bit (bit 4) of the IO_ERR register sets to 1. The LOSP Input Error bit is reset to 0 by the following methods.

- Writing a 1 to the Clear LOSP Input Control bit (bit 0) of the IO_CLR register
- Writing a 1 to the Clear LOSP Input Error Flags bit (bit 7) of the IO_CLR register.

LOSP Output Complete (Bit 10)

This bit sets to 1 when the LOSP circuitry transfers all of the LOSP output data from the LOSP_DO register and the LOSP output buffer to the LOSP output channel. When the word in the LOSP_DO register was the first word in the transfer, the LOSP output complete bit indicates that the microprocessor has control of the LOSP output channel. When the word in the LOSP_DO register was the last word of the transfer or the 33rd word of the transfer, the LOSP output complete bit indicates that the LOSP output buffer is empty and the microprocessor may terminate the transfer or write more words to the LOSP_DO register. The LOSP output complete bit is reset to 0 by the following methods.

- Writing to the LOSP_DO register
- Writing a 1 to the clear LOSP output control bit (bit 1) of the IO_CLR register

LOSP Output Error (Bit 11)

This bit sets to 1 when the LOSP output ROA parity error bit (bit 10) or the LOSP output sequence error bit (bit 11) of the IO_ERR register sets to 1. The LOSP output error bit is reset to 0 by the following methods.

- Writing a 1 to the clear LOSP output control bit (bit 1) of the IO_CLR register
- Writing a 1 to the clear LOSP output error flags bit (bit 8) of the IO_CLR register

Master Clear In (Bit 12)

This bit indicates the present value of the Master Clear LOSP input channel signal. When set to 1, the Master Clear signal instructs the CRAY T3D system to perform a global reset. During the reset, when the master clear in bit of the IO_IFLAG register in an I/O gateway is set to 1, this bit indicates that the I/O gateway is the deadstart I/O gateway.

When the Master Clear signal is set to 1, the master clear in bit of the IO_IFLAG register in both the input and output node of an I/O gateway is set to 1. The master clear in bit is reset to 0 only when the host system resets the Master Clear signal to 0.

NOTE: In the hardware, the master clear in bit is also indicated by bit 12 of the IO_IMASK register. Because of this characteristic, the correct value of the master clear in bit is read by performing a logical OR between bit 12 of the IO_IFLAG register and bit 12 of the IO_IMASK register.

Revision Level (Bits 13 and 14)

These bits indicate the revision level of the I/O module. For example, when set to 1 (01), these bits indicate the I/O node is on a revision 1 I/O module. Revision 1 I/O modules contain a buffered LOSP channel while the revision 0 I/O modules contain a nonbuffered LOSP channel.

Master or Slave I/O Node (Bit 15)

This bit indicates when the I/O node is part of a master I/O gateway or a slave I/O gateway. When set to 0, bit 15 indicates a master I/O gateway. When set to 1, bit 15 indicates a slave I/O gateway.

Interrupt Mask Register (Address 4000220₁₆)

The interrupt mask (IO_IMASK) register is a 13-bit, readable and writable register. The IO_IMASK register indicates which IO_IFLAG register bits, when set to 1, cause the I/O hardware interrupt to be set.

When a bit of the IO_IMASK register is set to 1, the corresponding bit of the IO_IFLAG register will set the I/O hardware interrupt. When a bit of the IO_IMASK register is set to 0, the corresponding bit of the IO_IFLAG register will not set the I/O hardware interrupt (refer to Figure 11-18).

Table 11-10 shows the bit format of the IO_IMASK register.

HEX	0	4				0				0				0				2		2		0							
Binary	2 ²⁸	2 ²⁷	2 ²⁶	2 ²⁵	2 ²⁴	2 ²³	2 ²²	2 ²¹	2 ²⁰	2 ¹⁹	2 ¹⁸	2 ¹⁷	2 ¹⁶	2 ¹⁵	2 ¹⁴	2 ¹³	2 ¹²	2 ¹¹	2 ¹⁰	2 ⁹	2 ⁸	2 ⁷	2 ⁶	2 ⁵	2 ⁴	2 ³	2 ²	2 ¹	2 ⁰
	0	0	1	0	x	x	x	x	x	x	x	x	x	x	x	x	0	0	0	1	x	x	x	1	x	0	0	0	0

NOTE: x = Don't Care

Figure 11-18. Interrupt Mask Register Address Bit Format

Table 11-10. IO_IMASK Register Bit Format

Bits	Name	Nodes
0	HISP error mask	Input and output
1	Microprocessor single-bit error mask	Input and output
2	Microprocessor multiple-bit error mask	Input and output
3	HISP single-bit error mask	Input
4	HISP multiple-bit error mask	Input
5	HISP address transfer complete mask	Input and output
6	HISP data transfer complete mask	Input and output
7	LOSP input data available mask	Input and output
8	LOSP input disconnect mask	Input and output
9	LOSP input error mask	Input and output
10	LOSP output complete mask	Input and output
11	LOSP output error mask	Input and output
12	Master clear in (Not a mask bit)	Input and output
63 – 13	These bits are not used	Input and output

I/O Channel Error Register (Address 4000C00₁₆)

The I/O channel error (IO_ERR) register is a 16-bit, read-only register that indicates when various error conditions exist for the HISP and LOSP channels. Figure 11-19 shows the bit assignments for the IO_ERR register address as they appear on the address pins of the microprocessor.

HEX	0	4				0				0				C				0				0							
Binary	2 ²⁸	2 ²⁷	2 ²⁶	2 ²⁵	2 ²⁴	2 ²³	2 ²²	2 ²¹	2 ²⁰	2 ¹⁹	2 ¹⁸	2 ¹⁷	2 ¹⁶	2 ¹⁵	2 ¹⁴	2 ¹³	2 ¹²	2 ¹¹	2 ¹⁰	2 ⁹	2 ⁸	2 ⁷	2 ⁶	2 ⁵	2 ⁴	2 ³	2 ²	2 ¹	2 ⁰
	0	0	1	0	x	x	x	x	x	x	x	x	x	x	x	x	0	1	1	0	x	x	x	x	x	0	0	0	0

NOTE: x = Don't Care

Figure 11-19. IO_ERR Register Address Bit Format

Master I/O Gateway

Table 11-11 shows the bit assignments for the IO_ERR register in a master I/O gateway, and the following subsections describe each bit.

Table 11-11. Master IO_ERR Register Bit Format

Bits	Name
0	HISP address error
1	HISP unrecoverable error
2	HISP block error in output node HISP block length error in input node
3	LOSP input parity error
4	LOSP input sequence error
7 – 5	LOSP input parcel count
8	HISP ROA parity error
9	Reserved
10	LOSP output ROA parity error
11	LOSP output sequence error
13 – 12	Reserved
15 – 14	Block length high bits 17 and 16
63 – 16	These bits are not used

HISP Address Error (Bit 0)

This bit indicates the present value of the Address Error HISP channel signal. When set to 1, the Address Error signal indicates that the host system detected a parity error on the address and block length information it received or indicates that the host system detected the Address Ready signal when it should not have been set. The HISP address error bit is reset to 0 when the host system resets the Address Error channel signal to 0.

To reset the Address Error channel signal, the microprocessor should write a 1 to the HISP clear channel bit (bit 12) of the HISP_A_CFG register. The clear channel bit should be set to 1 for a minimum of 150ns and should remain set when the host system has not reset the Address Error or Unrecoverable Error HISP channel signals.

HISP Unrecoverable Error (Bit 1)

This bit indicates the present value of the Unrecoverable Error HISP channel signal. When set to 1, the Unrecoverable Error signal indicates that the host system encountered an error from which it could not recover. The HISP unrecoverable error bit is reset to 0 when the host system resets the Unrecoverable Error Channel signal to 0.

To reset the Unrecoverable Error channel signal, the microprocessor should write a 1 to the HISP clear channel bit (bit 12) of the HISP_A_CFG register. The clear channel bit should be set to 1 for a minimum of 150ns and should remain set when the host system has not reset the Address Error or Unrecoverable Error HISP channel signals.

HISP Block Length Error or HISP Block Error (Bit 2)

In the input node, this bit sets to 1 when a block length error occurs. Two conditions will cause a block length error. One condition occurs when the host system sets the Last Word HISP channel signal and the block length in the HISP_U_BL and HISP_L_BL registers has not decremented to 0. The other condition occurs when the block length in the HISP_U_BL and HISP_L_BL registers has decremented to 0, but the host system has not set the Last Word HISP channel signal. The HISP block length error bit is reset to 0 when the microprocessor writes a 1 to the clear HISP data control bit (bit 3) of the IO_CLR register.

In the output node, the HISP block error bit sets to 1 when a block error occurs. Two conditions cause a block error. One condition occurs when the microprocessor or BLT attempts to write a word into one of the four

32-word blocks in the HISP output buffer before the output circuitry has enabled that block so it may be written to. The other condition occurs when the output circuitry attempts to read a word out of a 16-word block in the HISP output buffer before the block has been filled or before the block contains the last word of the transfer. The HISP block error bit is reset to 0 when the microprocessor writes a 1 to the clear HISP data control bit (bit 3) of the IO_CLR register.

LOSP Input Parity Error (Bit 3)

When set to 1, this bit indicates that a LOSP input parity error occurred. A LOSP input parity error may occur due to one of two conditions. The first condition occurs when a parity error is detected on the LOSP data when it is transferred from the LOSP channel to the LOSP input buffer. The second condition occurs when a parity error is detected on the LOSP data when it is transferred from the LOSP input buffer to the LOSP_DI register.

In either condition, the LOSP circuitry sets the LOSP input parity error bit to 1 only when the LOSP circuitry transfers the LOSP data from the LOSP buffer to the LOSP_DI register. When a parity error is detected on the LOSP data when it is transferred from the LOSP channel to the LOSP input buffer, the LOSP input parity error bit is not set to 1 until that word is read out of the LOSP input buffer. When the LOSP input parity error bit is set to 1, the LOSP input error bit (bit 9) of the IO_IFLAG register sets to 1. The LOSP input parity error bit is reset to 0 by the following methods.

- Writing a 1 to the clear LOSP input control bit (bit 0) of the IO_CLR register
- Writing a 1 to the clear LOSP input error flags bit (bit 7) of the IO_CLR register

LOSP Input Sequence Error (Bit 4)

When set to 1, this bit indicates that a LOSP input sequence error occurred. A LOSP input sequence error occurs when the host system places new LOSP data on the LOSP data signals and sets the Ready signal before the I/O gateway has read the previous LOSP data and set the Resume signal. When the LOSP input sequence error bit is set to 1, the LOSP input error bit (bit 9) of the IO_IFLAG register sets to 1. The LOSP input sequence error bit is reset to 0 by the following methods.

- Writing a 1 to the clear LOSP input control bit (bit 0) of the IO_CLR register
- Writing a 1 to the clear LOSP input error flags bit (bit 7) of the IO_CLR register

LOSP Input Parcel Count (Bits 5 – 7)

These bits indicate the number of valid parcels that were read from the LOSP_DI register. When the microprocessor reads the last word of LOSP input data from the LOSP_DI register, the LOSP circuitry sets the LOSP input disconnect bit (bit 8) of the IO_IFLAG register to 1. The microprocessor must then read the LOSP input parcel count bits to determine how many parcels (0 through 4) of the word read from the LOSP_DI register were valid.

HISP ROA Parity Error (Bit 8)

In the input node, this bit sets to 1 when the input circuitry detects a parity error on the HISP input data when the input circuitry reads the data out of the HISP input buffer. In the output node, this bit sets to 1 when the output circuitry detects a parity error on the HISP output data when the output circuitry reads the data out of the HISP output buffer. The HISP ROA parity error bit is reset to 0 when the microprocessor writes a 1 to the clear HISP data control bit (bit 3) of the IO_CLR register.

LOSP Output ROA Parity Error (Bit 10)

When set to 1, this bit indicates that a LOSP output ROA parity error occurred. A LOSP output ROA parity error occurs when the LOSP circuitry detects a parity error when reading a parcel of data out of the LOSP output buffer. This error is detected before the LOSP circuitry generates the parity bits for the LOSP output channel. When a LOSP

output ROA parity error occurs, the LOSP circuitry sets the LOSP output channel parity bits so that the host system will detect a parity error when it receives the LOSP data.

When the LOSP output ROA parity error bit is set to 1, the LOSP output error bit (bit 11) of the IO_IFLAG register sets to 1. The LOSP output ROA parity error bit is reset to 0 by the following methods.

- Writing a 1 to the clear LOSP output control bit (bit 1) of the IO_CLR register
- Writing a 1 to the clear LOSP output error flags bit (bit 8) of the IO_CLR register

LOSP Output Sequence Error (Bit 11)

When set to 1, this bit indicates that a LOSP output sequence error occurred. A LOSP output sequence error occurs when the host system sets the Resume LOSP channel signal and the LOSP circuitry is not performing a LOSP output transfer. When this bit sets to 1, the LOSP output error bit (bit 11) of the IO_IFLAG register sets to 1. The LOSP output sequence error bit is reset to 0 by the following methods.

- Writing a 1 to the clear LOSP output control bit (bit 1) of the IO_CLR register
- Writing a 1 to the clear LOSP output error flags bit (bit 8) of the IO_CLR register

Block Length High (Bits 14 – 15)

These bits contain the present value of bits 16 and 17 of the block length. Because the HISP_U_BL register is a write-only register, bits 16 and 17 of the block length must be read using bits 14 and 15 of the IO_ERR register.

Slave I/O Gateway

Table 11-12 shows the bit assignments for the IO_ERR register in a slave I/O gateway. Except for bits 0, 1, 9, 12, and 13, the IO_ERR register in a slave I/O gateway is identical to the IO_ERR register in a master I/O gateway. The following subsections describe only the bits that differ.

HISP Clear Channel (Bit 0)

This bit indicates the present value of the Clear Channel HISP channel signal. The host system sets the Clear Channel signal to 1 to reset the HISP circuitry in the slave I/O gateway. The HISP Clear Channel bit is reset to 0 when the host system resets the Clear Channel HISP channel signal to 0.

HISP Disable SECDED (Bit 1)

This bit indicates the present value of the Disable SECDED HISP channel signal. The host system sets the Disable SECDED signal to 1 to indicate that the slave I/O gateway should not perform SECDED on HISP input data or generate check bits for HISP output data. The HISP disable SECDED bit is reset to 0 when the host system resets the Disable SECDED HISP channel signal to 0.

Table 11-12. Slave IO_ERR Register Bit Format

Bits	Name
0	HISP clear channel
1	HISP disable SECDED
2	HISP block error in output node HISP block length error in input node
3	LOSP input parity error
4	LOSP input sequence error
7 – 5	LOSP input parcel count
8	HISP ROA parity error
9	HISP address parity error
10	LOSP output ROA parity error
11	LOSP output sequence error
12	HISP address ready error

Table 11-12. Slave IO_ERR Register Bit Format (continued)

Bits	Name
13	HISP unrecoverable error
15 – 14	Block length high bits 17 and 16
63 – 16	These bits are not used

HISP Address Parity Error (Bit 9)

This bit sets to 1 when the HISP circuitry detects a parity error on the address and block length information it receives from the host system over the Address and Block Length HISP channel signals. To reset the HISP address parity error bit to 0, the microprocessor writes a 1 to the clear HISP address control bit (bit 2) of the IO_CLR register.

HISP Address Ready Error (Bit 12)

This bit sets to 1 when the HISP circuitry detects an Address Ready HISP channel signal when it was not expected to be set. To reset the HISP address ready error bit to 0, the microprocessor writes a 1 to the clear HISP address control bit (bit 2) of the IO_CLR register.

HISP Unrecoverable Error (Bit 13)

This bit indicates the present value of the Unrecoverable Error HISP channel signal. When the host system sets the Clear Channel HISP channel signal, the I/O gateway automatically sets the Unrecoverable Error signal to 1. The Unrecoverable Error signal remains set to 1 until the microprocessor writes a 0 to the unrecoverable error HISP signal bit (bit 8 of the HISP_A_CFG register).

Clear I/O Register (Address 4001E00₁₆)

The clear I/O (IO_CLR) register is a 10-bit write-only register. The IO_CLR register clears various bits and hardware in the I/O gateway. Figure 11-20 shows the bit assignments for the IO_CLR register address as they appear on the address pins of the microprocessor.

HEX	0	4				0				0				1				E				0				0			
Binary	2 ²⁸	2 ²⁷	2 ²⁶	2 ²⁵	2 ²⁴	2 ²³	2 ²²	2 ²¹	2 ²⁰	2 ¹⁹	2 ¹⁸	2 ¹⁷	2 ¹⁶	2 ¹⁵	2 ¹⁴	2 ¹³	2 ¹²	2 ¹¹	2 ¹⁰	2 ⁹	2 ⁸	2 ⁷	2 ⁶	2 ⁵	2 ⁴	2 ³	2 ²	2 ¹	2 ⁰
	0	0	1	0	x	x	x	x	x	x	x	x	x	x	x	x	1	1	1	1	x	x	x	x	x	0	0	0	0

NOTE: x = Don't Care

Figure 11-20. Clear IO Register Address Bit Format

Table 11-13 shows the bit format of the IO_CLR register, and the following subsections describe each bit of the register.

Table 11-13. IO_CLR Register Bit Format

Bits	Name
0	Clear LOSP input control
1	Clear LOSP output control
2	Clear HISP address control
3	Clear HISP data control
4	Clear HISP ROA address
5	Clear LOSP input ready waiting bit (input node only)
6	Clear LOSP buffer and priority
7	Clear LOSP input error flags
8	Clear LOSP output error flags
9	Clear LOSP input disconnect
63 – 10	These bits are not used

Clear LOSP Input Control (Bit 0)

When the microprocessor in the input node writes a 1 to the clear LOSP input control bit of its associated IO_CLR register, following conditions occur.

- The LOSP input bits (bits 7 through 9) of the IO_IFLAG register reset to 0
- The LOSP input bits (bit 3 through 7) of the IO_ERR register reset to 0
- The LOSP input state machine logic resets

This action does not clear the resume waiting bit. More information on the resume waiting bit is provided in the following subsections. (Writing a 0 to the clear LOSP input control bit has no effect.)

When the microprocessor in the output node writes a 1 to the clear LOSP input control bit of its associated IO_CLR register, the following conditions occur.

- The LOSP input bits (bits 7 through 9) of the IO_IFLAG register reset to 0
- The LOSP input bits (bit 3 through 7) of the IO_ERR register reset to 0

This action does not clear the resume waiting bit and does not reset the LOSP input state machine logic. (Writing a 0 to the clear LOSP input control bit has no effect.)

Clear LOSP Output Control (Bit 1)

When the microprocessor in the output node writes a 1 to the clear LOSP output control bit of its associated IO_CLR register, the following conditions occur. (Writing a 0 to the clear LOSP output control bit has no effect.)

- The LOSP output bits (bits 10 and 11) of the IO_IFLAG register reset to 0
- The LOSP output bits (bit 10 and 11) of the IO_ERR register reset to 0
- The LOSP output state machine logic resets

When the microprocessor in the input node writes a 1 to the clear LOSP output control bit of its associated IO_CLR register, the following conditions occur. (Writing a 0 to the clear LOSP output control bit has no effect.)

- The LOSP output bits (bits 10 and 11) of the IO_IFLAG register reset to 0
- The LOSP output bits (bit 10 and 11) of the IO_ERR register reset to 0

Clear HISP Address Control (Bit 2)

When the microprocessor in the input or output node of a master I/O gateway writes a 1 to the clear HISP address control bit of its associated IO_CLR register, the hold paths of the Address Error and Unrecoverable Error HISP Channel signals are broken. This action causes a resampling of these signals on the channel. Bit 2 of the IO_CLR register should be set to 1 after the Clear Channel signal is asserted; setting the bit to 1 clears the Address Error and Unrecoverable Error signals in the I/O gateway circuitry.

When the microprocessor in the input or output node of a slave I/O gateway writes a 1 to the Clear HISP Address Control bit of its associated IO_CLR register, the following conditions occur. Writing a 0 to this bit has no effect.

- The HISP address parity error bit (bit 9) of the IO_ERR register is reset to 0
- The HISP address ready error bit (bit 12) of the IO_ERR register is reset to 0
- The HISP addressing circuitry is reset

Clear HISP Data Control (Bit 3)

When the microprocessor in an input or output node writes a 1 to the clear HISP data control bit of its associated IO_CLR register, the following conditions occur. (Writing a 0 to this bit has no effect.)

- The HISP clear channel bit (bit 0) of the IO_ERR register resets to 0 (slave I/O gateway only)
- The HISP disable SECDDED bit (bit 1) of the IO_ERR register resets to 0 (slave I/O gateway only)
- The HISP block error bit (bit 2 in the output node) or the HISP block length error bit (bit 2 in the input node) of the IO_ERR register resets to 0
- The HISP ROA parity error bit (bit 8) of the IO_ERR register resets to 0
- The HISP data circuitry is reset

In the master I/O gateway, when the microprocessor writes a 1 to the clear HISP data control bit of its associated IO_CLR register, the HISP addressing circuitry is also reset.

Clear HISP ROA Address (Bit 4)

When the microprocessor in an input node writes a 1 to the clear HISP ROA address bit of its associated IO_CLR register, the internal address for the HISP input buffer is reset to 0. When the microprocessor in an output node writes a 1 to the clear HISP ROA address bit of its associated IO_CLR register, the internal address for the HISP output buffer is reset to 0.

Clear LOSP In Ready Waiting Bit (Bit 5)

When the microprocessor in the input or output node writes a 1 to the clear LOSP input ready waiting bit of its associated IO_CLR register, the LOSP circuitry resets the ready waiting bit to 0. The ready waiting bit is located in the LOSP circuitry and indicates when the I/O gateway received a Ready signal but did not respond with the Resume signal. The ready waiting bit must be reset to 0 when an error occurs that prematurely terminates the LOSP input transfer and the host system sets the Ready signal after the transfer was terminated. (Writing a 0 to this bit has no effect.)

Clear LOSP Buffer and Priority (Bit 6)

When the microprocessor in the input node writes a 1 to the clear LOSP buffer and priority bit of its associated IO_CLR register, the LOSP input buffer is cleared. In addition, the priority of the LOSP input channel that was indicated by bit 15 of the first parcel of LOSP input data is reset. When this occurs, any LOSP input data that was in the LOSP input buffer is lost. (Writing a 0 to this bit has no effect.)

When the microprocessor in the output node writes a 1 to the clear LOSP buffer and priority bit of its associated IO_CLR register, the LOSP output buffer is cleared. When this occurs, any LOSP output data that was in the LOSP output buffer is lost. (Writing a 0 to this bit has no effect.)

Clear LOSP Input Error Flags (Bit 7)

When the microprocessor in the input or output node writes a 1 to the clear LOSP input error flags bit in its associated IO_CLR register, the following conditions occur. (Writing a 0 to this bit has no effect.)

- The LOSP input parity error bit (bit 3) and the LOSP input sequence error bit (bit 4) of the IO_ERR register reset to 0
- The LOSP input error bit (bit 9) of the IO_IFLAG register reset to 0

Clear LOSP Output Error Flags (Bit 8)

When the microprocessor in the input or output node writes a 1 to the clear LOSP output error flags bit of its associated IO_CLR register, the following conditions occur. (Writing a 0 to this bit has no effect.)

- The LOSP output ROA parity error bit (bit 10) and the LOSP output sequence error bit (bit 11) of the IO_ERR register reset to 0
- The LOSP output error bit (bit 11) of the IO_IFLAG register resets to 0

Clear LOSP Input Disconnect (Bit 9)

When the microprocessor in the input or output node writes a 1 to the clear LOSP input disconnect bit of its associated IO_CLR register, the LOSP input disconnect bit (bit 8) is reset to 0 in the IO_IFLAG register. Writing a 0 to the clear LOSP input disconnect bit has no effect.

Clear LOSP Input Data Available (Bit 10)

When the microprocessor in the input or output node writes a 1 to the clear LOSP input data available bit of its associated IO_CLR register, the LOSP input data available bit (bit 7) is reset to 0 in the IO_IFLAG register. Writing a 0 to the clear LOSP input data available bit has no effect.

LOSP Channel Configuration Register (Address 4000000₁₆)

The LOSP channel configuration (LOSP_CFG) register is a write-only register that controls timing parameters for the LOSP channel Ready or Resume control signals. Figure 11-21 shows the bit assignments for the LOSP_CFG register address as they appear on the address pins of the microprocessor.

HEX	0	4				0				0				0				0				0							
Binary	2 ²⁸	2 ²⁷	2 ²⁶	2 ²⁵	2 ²⁴	2 ²³	2 ²²	2 ²¹	2 ²⁰	2 ¹⁹	2 ¹⁸	2 ¹⁷	2 ¹⁶	2 ¹⁵	2 ¹⁴	2 ¹³	2 ¹²	2 ¹¹	2 ¹⁰	2 ⁹	2 ⁸	2 ⁷	2 ⁶	2 ⁵	2 ⁴	2 ³	2 ²	2 ¹	2 ⁰
	0	0	1	0	x	x	x	x	x	x	x	x	x	x	x	x	0	0	0	0	x	x	0	0	x	0	0	0	0

NOTE: x = Don't Care

Figure 11-21. LOSP Channel Configuration Register Address Bit Format

Do not change the value of the bits in the LOSP_CFG register unless the LOSP channel is idle. When the LOSP channel is not idle and the value of the bits in the LOSP_CFG register changes, the LOSP channel may hang. The LOSP_CFG register sets different parameters in the input node than it does in the output node.

Input Node

Table 11-14 shows the bit format of the LOSP_CFG register in an input node, and the following paragraphs describe each bit of the register.

Table 11-14. LOSP_CFG Register Format

Bits	Name
3 – 0	LOSP resume pulse width count
8 – 4	Delay count after leading edge of ready before sampling data
63 – 9	These bits are not used

Bits 0 through 3 of the LOSP_CFG register in the input node set the pulse width of the LOSP channel Resume signal. The value of bits 0 through 3 is calculated using the following equation:

$$\text{Resume Pulse Width Count} = \text{Pulse Width} / \text{Clock Period}$$

Currently in the CRAY T3D system, the clock period is 6.67 ns and the Resume signal pulse width should be 50 ns plus or minus 3 ns. Using these values, the resume pulse width count should be set to 7.

Bits 4 through 8 of the LOSP_CFIG register set the time delay after receiving the leading edge of the Ready signal before the LOSP circuitry samples the data. The value of bits 4 through 8 is calculated using the following equation:

$$\text{Data Sample Count} = (\text{Sample Time} / \text{Clock Period}) - 4$$

Currently in the CRAY T3D system, the clock period is 6.67 ns and the LOSP data sample time should be greater than 80 ns. Using these values, the data sample count should be set to 8.

Output Node

Table 11-15 shows the bit format of the LOSP_CFIG register in the output node, and the following paragraphs describe each bit of the register.

Table 11-15. LOSP_CFIG Register Bit Format in Output Node

Bits	Name
3 – 0	LOSP data ready pulse width count
63 – 4	These bits are not used

Bits 0 through 3 of the LOSP_CFIG register in the output node set the pulse width of the LOSP channel Ready signal. The value of bits 0 through 3 is calculated using the following equation:

$$\text{Ready Pulse Width Count} = \text{Pulse Width} / \text{Clock Period}$$

Currently in the CRAY T3D system, the clock period is 6.67 ns and the Ready signal pulse width should be 50 ns plus or minus 10 ns. Using these values, the resume pulse width count should be set to 7.

LOSP Channel Data In Register (Address 4000400₁₆)

The low speed channel data in (LOSP_DI) register is a 64-bit, read-only register. The LOSP_DI register contains a 64-bit word of LOSP input data that the LOSP circuitry assembles from four 16-bit parcels.

Figure 11-22 shows the bit assignments for the LOSP_DI register address as they appear on the address pins of the microprocessor.

HEX	0	4				0				0				0				4				0				0			
Binary	2 ²⁸	2 ²⁷	2 ²⁶	2 ²⁵	2 ²⁴	2 ²³	2 ²²	2 ²¹	2 ²⁰	2 ¹⁹	2 ¹⁸	2 ¹⁷	2 ¹⁶	2 ¹⁵	2 ¹⁴	2 ¹³	2 ¹²	2 ¹¹	2 ¹⁰	2 ⁹	2 ⁸	2 ⁷	2 ⁶	2 ⁵	2 ⁴	2 ³	2 ²	2 ¹	2 ⁰
	0	0	1	0	x	x	x	x	x	x	x	x	x	x	x	0	0	1	0	x	x	x	x	x	x	0	0	0	0

NOTE: x = Don't Care

Figure 11-22. LOSP_DI Register Address Bit Format

When the last word of data is read from the LOSP_DI register, the LOSP circuitry sets the LOSP Input Disconnect bit (bit 8) of the IO_IFLAG register to 1. The microprocessor then must read the valid parcel count (bits 5 through 7) of the IO_ERR register to determine how many parcels (0 through 4) of the word contain valid data.

When a LOSP input data word contains less than 4 valid parcels, the data will not be correctly aligned. When this occurs, the data must be shifted left by the number of parcels that were not valid. It is recommended that all transfers be full word transfers.

Table 11-16 shows the bit format of the LOSP_DI register.

Table 11-16. LOSP_DI Register Bit Format

Bits	Name
15 – 0	Bits 15 – 0 of a 64-bit word (parcel 0)
31 – 16	Bits 31 – 16 of a 64-bit word (parcel 1)
47 – 32	Bits 47 – 32 of a 64-bit word (parcel 2)
63 – 48	Bits 63 – 48 of a 64-bit word (parcel 3)

LOSP Channel Data Out Register (Address 4000800₁₆)

The LOSP channel data out (LOSP_DO) register is a 64-bit write-only register. The LOSP_DO register receives a 64-bit word of data from the microprocessor in the input or output node and sends the data through the LOSP output buffer to the LOSP output channel in 16-bit parcels.

Figure 11-23 shows the bit assignments for the LOSP_DO register address as they appear on the address pins of the microprocessor.

HEX	0	4				0				0				0				8				0				0				
Binary	2 ²⁸	2 ²⁷	2 ²⁶	2 ²⁵	2 ²⁴	2 ²³	2 ²²	2 ²¹	2 ²⁰	2 ¹⁹	2 ¹⁸	2 ¹⁷	2 ¹⁶	2 ¹⁵	2 ¹⁴	2 ¹³	2 ¹²	2 ¹¹	2 ¹⁰	2 ⁹	2 ⁸	2 ⁷	2 ⁶	2 ⁵	2 ⁴	2 ³	2 ²	2 ¹	2 ⁰	
	0	0	1	0	x	x	x	x	x	x	x	x	x	x	x	x	0	1	0	0	0	x	x	x	x	x	0	0	0	0

NOTE: x = Don't Care

Figure 11-23. LOSP_DO Register Address Bit Format

LOSP Output Channel Disconnect Register (Address 4000A00₁₆)

The LOSP output channel disconnect (LOSP_DISC) register is a 64-bit, write-only register. When the microprocessor in the input or output node writes any value to the LOSP_DISC register, the LOSP circuitry sets the Disconnect LOSP channel signal. The data written to the LOSP_DISC register is not used.

Figure 11-24 shows the bit assignments for the LOSP_DISC register address as they appear on the address pins of the microprocessor.

HEX	0	4				0				0				0				A				0				0			
Binary	2 ²⁸	2 ²⁷	2 ²⁶	2 ²⁵	2 ²⁴	2 ²³	2 ²²	2 ²¹	2 ²⁰	2 ¹⁹	2 ¹⁸	2 ¹⁷	2 ¹⁶	2 ¹⁵	2 ¹⁴	2 ¹³	2 ¹²	2 ¹¹	2 ¹⁰	2 ⁹	2 ⁸	2 ⁷	2 ⁶	2 ⁵	2 ⁴	2 ³	2 ²	2 ¹	2 ⁰
	0	0	1	0	x	x	x	x	x	x	x	x	x	x	x	x	0	1	0	1	x	x	x	x	x	0	0	0	0

NOTE: x = Don't Care

Figure 11-24. LOSP_DISC Register Address Bit Format

HISP Channel Address Configuration Register (Address 4000020₁₆)

The HISP channel address configuration (HISP_A_CFG) register is a write-only register that controls the HISP channel address timing. The HISP_A_CFG register also enables the various areas of SECEDED circuitry in the input node and output node.

Figure 11-25 shows the bit assignments for the HISP_A_CFG register address as they appear on the address pins of the microprocessor.

HEX	0				4				0				0				0				0				2		0		
Binary	2 ²⁸	2 ²⁷	2 ²⁶	2 ²⁵	2 ²⁴	2 ²³	2 ²²	2 ²¹	2 ²⁰	2 ¹⁹	2 ¹⁸	2 ¹⁷	2 ¹⁶	2 ¹⁵	2 ¹⁴	2 ¹³	2 ¹²	2 ¹¹	2 ¹⁰	2 ⁹	2 ⁸	2 ⁷	2 ⁶	2 ⁵	2 ⁴	2 ³	2 ²	2 ¹	2 ⁰
	0	0	1	0	x	x	x	x	x	x	x	x	x	x	x	x	0	0	0	0	x	x	0	1	x	0	0	0	0

NOTE: x = Don't Care

Figure 11-25. HISP Channel Address Configuration Register Address Bit Format

Do not change the value of the bits in the HISP_A_CFG register unless the HISP channel is idle. When the HISP channel is not idle and the value of the bits in the HISP_A_CFG register changes, the HISP channel may hang.

The HISP_A_CFG register has a different format in a master I/O gateway than it does in a slave I/O gateway.

Master I/O Gateway

Table 11-17 shows the bit format of the HISP_A_CFG register in a master I/O gateway, and the following subsections describe each bit of the register.

HISP Address Ready Pulse Width (Bits 0 – 3)

These bits set the pulse width of the HISP channel Address Ready signal. The value of bits 0 through 3 is calculated using the following equation:

$$\text{Address Ready Pulse Width Count} = (\text{Pulse Width} / \text{Clock Period}) - 1$$

Currently in the CRAY T3D system, the clock period is 6.67 ns and the Address Ready signal pulse width should be 25 ns when using 100 Mbytes/s mode and 18.75 ns when using 200 Mbytes/s mode. Using these values, the Address Ready pulse width count should be set to 3 in 100 Mbytes/s mode and set to 2 in 200 Mbytes/s mode.

Table 11-17. HISP_A_CFIG Register Bit Format

Bits	Name	Nodes
3 – 0	HISP address ready pulse width count	Input and output
8 – 4	HISP address ready period count	Input and output
9	Address and block length transfer count 0 = Three transfers 1 = Five transfers	Input and output
10	Input node – disable HISP SECEDED options Output node – disable HISP output SECEDED	Input and output
11	Enable HISP diagnostic mode	Output
12	Clear channel HISP signal	Input and output
13	Enable microprocessor ECC	Input and output
14	Enable loopback mode on HISP	Input and output
15	Disable SECEDED HISP signal	Output
63 – 16	These bits are not used	Input and output

HISP Address Ready Period (Bits 4 – 8)

These bits set the period of the HISP channel Address Ready signal. The value of bits 4 through 8 is calculated using the following equation:

$$\text{Address Ready Period Count} = (\text{Period} / \text{Clock Period}) - 1$$

Currently in the CRAY T3D system, the clock period is 6.67 ns and the Address Ready signal period should be 75 ns when using 100 Mbytes/s mode and 37.5 ns when using 200 Mbytes/s mode. Using these values, the Address Ready pulse width count should be set to 10 in 100 Mbytes/s mode and set to 5 in 200 Mbytes/s mode.

Address and Block Length Transfer Count (Bit 9)

This bit controls the number of transfers required to transmit the complete address and block length information over the HISP Address and Block Length channel signals. When set to 0, this bit enables 100 Mbyte/s protocol for the HISP Address and Block Length signals. The 100 Mbyte/s protocol uses three transfers to send the complete address and block length information.

When set to 1, this bit enables 200 Mbyte/s protocol for the HISP Address and Block Length signals. The 200 Mbyte/s protocol uses five transfers to send the complete address and block length information. More information on the format of the address and block length information is provided in “HISP Channel Address 0” and “HISP Channel Address 1” later in this section.

Disable Input HISP SECDED or Disable Output HISP SECDED (Bit 10)

This bit has a different meaning in the input node than it does in the output node. When set to 1 in the input node, this bit disables SECDED on data as it transfers from the HISP input channel to the HISP buffers in the input node. When set to 1 in the output node, this bit disables the generation of SECDED check bits that are sent out over the HISP output channel. When disabled, the output node sends 0's in place of the check bits. This bit should be set to 1 before a transfer and reset to 0 only when the transfer is complete.

Enable HISP Diagnostic Mode (Bit 11)

When set to 1 in the output node, this bit enables a diagnostic mode for the HISP output channel. When this bit is set to 1, bit 10 of the HISP_A_CFG must be set to 0. The enable HISP diagnostic ECC mode bit has no effect in the input node.

In diagnostic mode, software can introduce data errors on the channel and verify the operation of the error correction code (ECC) generation hardware located in the host system. The diagnostic mode uses the checkbits generated for the previous 64-bit word of data in place of the check bits generated for the current 64-bit word of data. After enabling the diagnostic mode, the first word transferred uses the checkbits from the previous word transferred.

Clear Channel HISP Signal (Bit 12)

This bit is used as the Clear Channel signal in the HISP channel. The clear channel bit should be set to 1 for a minimum of 150 ns and should remain set when the host system has not reset the Address Error or Unrecoverable Error HISP channel signals.

Enable Microprocessor ECC (Bit 13)

When set to 1, this bit enables the correction of microprocessor SECDED errors by the I/O circuitry. When set to 0, this bit disables the correction of microprocessor SECDED errors. Although the correction of microprocessor SECDED errors may be disabled, the errors are still detected and reported in the microprocessor single-bit error and microprocessor double-bit error bits (bits 1 and 2) of the IO_IFLAG register.

Enable Loopback Mode on HISP Channel (Bit 14)

When set to 1, this bit enables loopback testing of the HISP channel. During a loopback test, the input node of an I/O gateway is cabled to the output node. In loopback mode, the HISP channel transfers data between the nodes without first transferring the address and block length information.

Disable SECDED HISP Signal (Bit 15)

This bit is used as the Disable SECDED HISP channel signal and is only used in the output node. When used, this bit should be set to 1 before a transfer and reset to 0 only when the transfer is complete.

Slave I/O Gateway

Table 11-18 shows the bit format of the HISP_A_CFIG register in a slave I/O gateway. Bits 9, 11, 13, and 14 are identical to the bits in the HISP_A_CFIG register of a master I/O gateway. The following subsections describe only the bits that differ between the master and slave I/O gateways.

Table 11-18. Slave HISP_A_CFIG Register Bit Format

Bits	Name	Nodes
4 – 0	HISP address sample count	Input and output
6 – 5	Reserved	Input and output
7	Address error HISP signal	Output
8	Unrecoverable error HISP signal	Output
9	Address and block length transfer count 0 = Three transfers 1 = Five transfers	Input and output
10	Disable output SECCDED	Output
11	Enable HISP diagnostic ECC mode	Output
12	Reserved	Output
13	Enable microprocessor ECC	Input and output
14	Enable loopback mode on HISP	Input and output
15	Reserved	Input
63 – 16	These bits are not used	Input and output

HISP Address Sample Count (Bits 0 through 4)

These bits set the time needed from the leading edge of the Address Ready signal until the HISP circuitry samples the address and block length information. The value of bits 0 through 4 is calculated using the following equation:

$$\text{Address Sample Count} = (\text{Sample Time} / \text{Clock Period}) - 4$$

Currently in the CRAY T3D system, the clock period is 6.60 ns and the address sample time should be 62.5 ns when using 100 Mbytes/s mode and 37.5 ns when using 200 Mbytes/s mode. Using these values, the Address Ready pulse width count should be set to 5 in 100 Mbytes/s mode and set to 1 in 200 Mbytes/s mode.

Address Error HISP Signal (Bit 7)

This bit is used as the Address Error HISP channel signal. This bit should be set to 1 when the HISP address parity error bit (bit 9) or the HISP address ready error bit (bit 12) of the IO_ERR register is set to 1. The address error HISP signal bit should be reset to 0 when the host system sets the Clear Channel HISP Channel signal.

Unrecoverable Error HISP Signal (Bit 8)

This bit is used as the Unrecoverable Error HISP channel signal. The Unrecoverable Error HISP channel signal automatically sets to 1 when the host system sets the Clear Channel HISP channel signal. To reset the Unrecoverable Error HISP channel signal to 0, the microprocessor must write a 0 to the unrecoverable error HISP signal bit of the HISP_A_CFG register.

Disable Output SECEDED (Bit 10)

When set to 1 in the output node, this bit disables the generation of SECEDED check bits that are sent out the HISP output channel. When disabled, the output node sends 0's in place of the check bits. This bit should be set to 1 before a transfer and reset to 0 only when the transfer is complete. This bit has no effect in the input node of a slave I/O gateway.

HISP Data Configuration Register (Address 4000040₁₆)

The HISP data configuration (HISP_D_CFIG) register is a write only register used to configure the timing on data transfers over the HISP channel. The HISP_D_CFIG register is used differently in the input node than it is in the output node.

Figure 11-26 shows the bit assignments for the HISP_D_CFIG register address as they appear on the address pins of the microprocessor.

HEX	0				4				0				0				0				0				4				0			
Binary	2 ²⁸	2 ²⁷	2 ²⁶	2 ²⁵	2 ²⁴	2 ²³	2 ²²	2 ²¹	2 ²⁰	2 ¹⁹	2 ¹⁸	2 ¹⁷	2 ¹⁶	2 ¹⁵	2 ¹⁴	2 ¹³	2 ¹²	2 ¹¹	2 ¹⁰	2 ⁹	2 ⁸	2 ⁷	2 ⁶	2 ⁵	2 ⁴	2 ³	2 ²	2 ¹	2 ⁰			
	0	0	1	0	x	x	x	x	x	x	x	x	x	x	x	x	0	0	0	0	x	x	1	0	x	0	0	0	0			

NOTE: x = Don't Care

Figure 11-26. HISP Data Configuration Register Address Bit Format

Do not change the value of the bits in the HISP_D_CFIG register unless the HISP channel is idle. When the HISP channel is not idle and the value of the bits in the HISP_D_CFIG register change, the HISP channel may hang.

Input Node

Table 11-19 shows the bit format of the HISP_D_CFIG register in an input node, and the following paragraphs describe each bit of the register.

Table 11-19. HISP_D_CFIG Register Format in the Input Node

Bits	Name
3 – 0	HISP sample data count
63 – 4	These bits are not used

Bits 0 through 3 set the time delay after the HISP input circuitry receives the leading edge of the Data Ready signal and before the input circuitry samples the HISP data. The value of bits 0 through 3 is calculated using the following equation:

$$\text{HISP Sample Data Count} = (\text{Sample Time} / \text{Clock Period}) - 4$$

Currently in the CRAY T3D system, the clock period is 6.60 ns and the HISP data sample time should be 62.5 ns when using 100 Mbytes/s mode and 37.5 ns when using 200 Mbytes/s mode. Using these values, the HISP sample data count should be set to 5 in 100 Mbytes/s mode and set to 1 in 200 Mbytes/s mode.

Output Node

Table 11-20 shows the bit format of the HISP_D_CFIG register in the output node, and the following paragraphs describe each bit of the register.

Table 11-20. HISP_D_CFIG Register Bit Format in the Output Node

Bits	Name
3 – 0	Set HISP data ready count
7 – 4	Clear HISP data ready count
11 – 8	HISP data change count
15 – 12	HISP data transfer period count
63 – 16	These bits are not used

Bits 0 through 3 determine the time delay after the output circuitry receives the Transmit Data signal and before the output circuitry sets the Data Ready signal. Currently in the CRAY T3D system, the set HISP Data Ready Count should be set to the following value:

$$\text{Set HISP Data Ready Count} = 1$$

Bits 4 through 7 determine the time delay after the output circuitry sets Data Ready until the output circuitry clears Data Ready. The value of bits 4 through 7 is calculated using the following equation:

$$\text{Clear HISP Data Ready Count} = (\text{Data Ready Pulse Width} / \text{Clock Period}) + \text{Set Data Ready Count}$$

Currently in the CRAY T3D system, the clock period is 6.60 ns and the Data Ready pulse width should be 25 ns after the leading edge of Data Ready signal when using 100 Mbytes/s mode and 18.75 ns after the leading edge of Data Ready signal when using 200 Mbytes/s mode. Using these values, the clear HISP data ready count should be set to 5 in 100 Mbytes/s mode and set to 4 in 200 Mbytes/s mode.

Bits 8 through 11 determine the HISP data change count. The value of bits 8 through 11 is calculated using the following equation:

$$\text{HISP Data Change Count} = (\text{Data Change} / \text{Clock Period}) - 1 + \text{Set Data Ready Count}$$

Currently in the CRAY T3D system, the clock period is 6.60 ns and the data change value should be 25 ns after the leading edge of Data Ready signal when using 100 Mbytes/s mode and 18.75 ns after the leading edge of Data Ready signal when using 200 Mbytes/s mode. Using these values, the HISP data change count should be set to 4 in 100 Mbytes/s mode and set to 3 in 200 Mbytes/s mode.

Bits 12 through 15 determine the period of the HISP data transfer. The value of bits 12 through 15 is calculated using the following equation:

$$\text{HISP Data Transfer Period Count} = (\text{HISP Data Period} / \text{Clock Period}) - 1$$

Currently in the CRAY T3D system, the clock period is 6.60 ns and the HISP Data Transfer period should be 75 ns when using 100 Mbytes/s mode and 37.5 ns when using 200 Mbytes/s mode. Using these values, the HISP data transfer period count should be set to 10 in 100 Mbytes/s mode and set to 5 in 200 Mbytes/s mode.

HISP Channel Lower Block Length Register (Address 4001A00₁₆)

The HISP channel lower block length (HISP_L_BL) register is a readable and writable register that contains bits 0 through 15 of the block length for a HISP transfer. Figure 11-27 shows the bit assignments for the HISP_L_BL register address as they appear on the address pins of the microprocessor.

HEX	0				4				0				0				1				A				0				0			
Binary	2 ²⁸	2 ²⁷	2 ²⁶	2 ²⁵	2 ²⁴	2 ²³	2 ²²	2 ²¹	2 ²⁰	2 ¹⁹	2 ¹⁸	2 ¹⁷	2 ¹⁶	2 ¹⁵	2 ¹⁴	2 ¹³	2 ¹²	2 ¹¹	2 ¹⁰	2 ⁹	2 ⁸	2 ⁷	2 ⁶	2 ⁵	2 ⁴	2 ³	2 ²	2 ¹	2 ⁰			
	0	0	1	0	x	x	x	x	x	x	x	x	x	x	x	x	1	1	0	1	x	x	x	x	x	0	0	0	0			

NOTE: x = Don't Care

Figure 11-27. HISP Lower Block Length Register Address Bit Format

Table 11-21 shows the bit format of the HISP_L_BL register and describes each bit of the register.

Table 11-21. HISP_L_BL Register Bit Format

Bits	Name
15 – 0	Bits 15 through 0 of the block length
63 – 16	These bits are not used

During a HISP input transfer, the HISP circuitry decrements the block length stored in the HISP_L_BL and HISP_U_BL registers each time the HISP circuitry receives a word of HISP input data. When the block length decrements to 0 and the host system has set the Last Word HISP Channel signal, the transfer is complete. When the block length decrements to 0 and the host system has not set the Last Word HISP Channel signal, or when the host system has set the Last Word HISP Channel signal but the block length has not decremented to 0, a block length error occurs. For more information on the block length error, refer again to “I/O Channel Error Register” in this section.

During a HISP output transfer, the HISP circuitry decrements the block length stored in the HISP_L_BL and HISP_U_BL registers each time the HISP circuitry sends a word of HISP output data to the host system. When the block length decrements to 0, the HISP circuitry sets the Last Word HISP channel signal.

HISP Channel Upper Block Length Register (Address 4001800₁₆)

The HISP channel upper block length (HISP_U_BL) register is a write-only register that contains bits 16 and 17 of the block length for a HISP transfer. Figure 11-28 shows the bit assignments for the HISP_U_BL register address as they appear on the address pins of the microprocessor.

HEX	0				4				0				0				1				8				0				0			
Binary	2 ²⁸	2 ²⁷	2 ²⁶	2 ²⁵	2 ²⁴	2 ²³	2 ²²	2 ²¹	2 ²⁰	2 ¹⁹	2 ¹⁸	2 ¹⁷	2 ¹⁶	2 ¹⁵	2 ¹⁴	2 ¹³	2 ¹²	2 ¹¹	2 ¹⁰	2 ⁹	2 ⁸	2 ⁷	2 ⁶	2 ⁵	2 ⁴	2 ³	2 ²	2 ¹	2 ⁰			
	0	0	1	0	x	x	x	x	x	x	x	x	x	x	x	x	1	1	0	0	x	x	x	x	x	0	0	0	0			

NOTE: x = Don't Care

Figure 11-28. HISP Upper Block Length Register Bit Format

Table 11-22 shows the bit format of the HISP_U_BL register and describes each bit of the register. To read the value of the block length bits in the HISP_U_BL register, the microprocessor must read the block length high bits (bits 16 and 17) of the IO_ERR register.

Table 11-22. HISP_U_BL Register Bit Format

Bits	Name
1 – 0	Bits 17 and 16 of the block length
63 – 2	These bits are not used

The block length bits stored in the HISP_U_BL register are used as described in “HISP Channel Lower Block Length Register” earlier in this section.

HISP Channel Address 0 Register (Address 4001400₁₆)

The HISP channel address 0 (HISP_ADDR0) register is a 64-bit register that is used differently in a master I/O gateway than it is in a slave I/O gateway. In a master I/O gateway, the HISP_ADDR0 register is a write-only register used to transfer the address and block length information to the host system over the HISP channel. In a slave I/O gateway, the HISP_ADDR0 register is a read-only register that is used to receive the address and block length information from the host system over the HISP channel. Figure 11-29 shows the bit assignments for the HISP_ADDR0 address as they appear on the address pins of the microprocessor.

HEX	0				4				0				0				1				4				0				0			
Binary	2 ²⁸	2 ²⁷	2 ²⁶	2 ²⁵	2 ²⁴	2 ²³	2 ²²	2 ²¹	2 ²⁰	2 ¹⁹	2 ¹⁸	2 ¹⁷	2 ¹⁶	2 ¹⁵	2 ¹⁴	2 ¹³	2 ¹²	2 ¹¹	2 ¹⁰	2 ⁹	2 ⁸	2 ⁷	2 ⁶	2 ⁵	2 ⁴	2 ³	2 ²	2 ¹	2 ⁰			
	0	0	1	0	x	x	x	x	x	x	x	x	x	x	x	x	1	0	1	0	x	x	x	x	x	0	0	0	0			

NOTE: x = Don't Care

Figure 11-29. HISP Channel Address 0 Register Address Bit Format

In a master I/O gateway, the HISP circuitry transfers the address and block length information to the host system as soon as the microprocessor writes information into the HISP_ADDR0 register. In a slave I/O gateway, the HISP circuitry starts to receive HISP data as soon as the microprocessor reads information from the HISP_ADDR0 register.

The values used in the HISP_ADDR0 and HISP_ADDR1 registers are defined by software and are not directly used by the hardware in the I/O gateway. The block length stored in the HISP_ADDR0 and HISP_ADDR1 registers must match the block length stored in the HISP_L_BL and HISP_U_BL registers or an error will occur at the end of the transfer.

The HISP_ADDR0 register has a different format when the address and block length information transfers over the HISP channel in 100 MBytes/s mode than it does when the address and block length information transfers over the HISP channel in 200 MBytes/s mode. Also, the format of the HISP_ADDR0 register depends on the component that the I/O gateway connects to. The following subsections describe sample formats of the HISP_ADDR0 register.

100 Mbytes/s Mode

In 100 Mbytes/s mode, the HISP circuitry uses three transfers of 16 bits plus parity to send or receive the address and block length. Table 11-23 shows the bit format of the HISP_ADDR0 register when used in 100 Mbytes/s mode and connected to a CRAY Y-MP host system.

Table 11-23. HISP_ADDR0 Register Bit Format in 100 Mbytes/s Mode

Transfer	Bits	Name
1	14 – 0	Starting address bits 31 through 17
	15	Starting address bit 9
2	17 – 16	Block length bits 13 and 12
	26 – 18	Starting address bits 8 through 0
	27	Starting address bit 16
	31 – 28	Starting address bits 13 through 10
3	43 – 32	Block length bits 11 through 0
	45 – 44	Starting address bits 15 and 14
	47 – 46	These bits are not used
Not applicable	63 – 48	These bits are not used

200 Mbytes/s Mode

In 200 Mbytes/s mode, the HISP circuitry uses five transfers of 12 bits plus parity to send or receive the address and block length. Table 11-24 shows the bit format of the HISP_ADDR0 register when the HISP channel is used in 200 Mbytes/s mode and is connected to a CRAY C90 host system. The fifth transfer is stored in the HISP_ADDR1 register.

Table 11-24. HISP_ADDR0 Register Bit Format in 200 Mbytes/s Mode

Transfer	Bits	Name
1	3 – 0	Starting address bits 3 through 0
	7 – 4	Starting address bits 23 through 20
	11 – 8	Block length bits 3 through 0
	15 – 12	These bits are not used
2	19 – 16	Starting address bits 7 through 4
	23 – 20	Starting address bits 27 through 24
	27 – 24	Block length bits 7 through 4
	31 – 28	These bits are not used
3	35 – 32	Starting address bits 11 through 8
	39 – 36	Starting address bits 31 through 28
	43 – 40	Block length bits 11 through 8
	47 – 44	These bits are not used
4	51 – 48	Starting address bits 15 through 12
	55 – 52	These bits are not used
	59 – 56	Block length bits 15 through 12
	63 – 60	These bits are not used

HISP Channel Address 1 Register (Address 4001600₁₆)

The HISP channel address1 (HISP_ADDR1) register is a 16-bit register that contains the fifth transfer of address and block length information when the address and block length information transfers over the HISP channel in 200 Mbytes/s mode. When the address and block length information transfers over the HISP channel in 100 Mbytes/s mode, the HISP_ADDR1 register is not used.

Figure 11-30 shows the bit assignments for the HISP_ADDR1 register address as they appear on the address pins of the microprocessor.

HEX	0	4				0				0				1				6				0				0			
Binary	2 ²⁸	2 ²⁷	2 ²⁶	2 ²⁵	2 ²⁴	2 ²³	2 ²²	2 ²¹	2 ²⁰	2 ¹⁹	2 ¹⁸	2 ¹⁷	2 ¹⁶	2 ¹⁵	2 ¹⁴	2 ¹³	2 ¹²	2 ¹¹	2 ¹⁰	2 ⁹	2 ⁸	2 ⁷	2 ⁶	2 ⁵	2 ⁴	2 ³	2 ²	2 ¹	2 ⁰
	0	0	1	0	x	x	x	x	x	x	x	x	x	x	x	x	1	0	1	1	x	x	x	x	x	0	0	0	0

NOTE: x = Don't Care

Figure 11-30. HISP Channel Address 1 Register Address Bit Format

When the HISP_ADDR1 register is used, the microprocessor must read or write information from or to the HISP_ADDR1 register before reading or writing information from or to the HISP_ADDR0 register. Table 11-25 shows the bit format of the HISP_ADDR1 register.

Table 11-25. HISP_ADDR1 Register Bit Format

Transfer	Bits	Name
5	3 – 0	Starting address bits 19 through 16
	7 – 4	These bits are not used
	11 – 8	Block length bits 19 through 16
	15 – 12	These bits are not used
Not applicable	63 – 16	These bits are not used

HISP Buffer Registers (Address 7000000₁₆)

The HISP buffer registers (HISP_BUFFER) are used differently in the input node than they are in the output node. In the input node, the HISP_BUFFER registers make up the HISP input buffer. In the output node, the HISP_BUFFER registers make up the HISP output buffer. Figure 11-31 shows the bit assignments for the HISP_BUFFER register address as they appear on the address pins of the microprocessor.

HEX	0	7				0				0				0				0				0							
Binary	2 ²⁸	2 ²⁷	2 ²⁶	2 ²⁵	2 ²⁴	2 ²³	2 ²²	2 ²¹	2 ²⁰	2 ¹⁹	2 ¹⁸	2 ¹⁷	2 ¹⁶	2 ¹⁵	2 ¹⁴	2 ¹³	2 ¹²	2 ¹¹	2 ¹⁰	2 ⁹	2 ⁸	2 ⁷	2 ⁶	2 ⁵	2 ⁴	2 ³	2 ²	2 ¹	2 ⁰
	0	0	1	1	1	0	0	0	0	0	0	0	0	0	0	0	0	x	x	x	x	x	a	a	a	a	a	a	a

NOTE: x = Don't Care

Figure 11-31. HISP Buffer Register Address Bit Format

For the BLT, the HISP input or output buffer is made up of 128 HISP_BUFFER registers that are addressed sequentially using bits 0 through 7 of the HISP_BUFFER address (refer to Figure 11-32).

For the microprocessor, the HISP input or output buffer is made up of 128 HISP_BUFFER registers that are addressed by even numbers using bits 0 through 7 of the HISP_BUFFER address (refer again to Figure 11-32).




Figure 11-32. BLT and Microprocessor HISP_BUFFER Addresses

HISP Channel ROA Address Register (Address 4001200₁₆)

The HISP channel ROA address (HISP_ROA) register is a 15-bit, write-only register that is used to set the internal HISP input or output buffer addresses to a specific value. Figure 11-33 shows the bit assignments for the HISP_ROA register address as they appear on the address pins of the microprocessor.

HEX	0	4				0				0				1				2				0				0			
Binary	2 ²⁸	2 ²⁷	2 ²⁶	2 ²⁵	2 ²⁴	2 ²³	2 ²²	2 ²¹	2 ²⁰	2 ¹⁹	2 ¹⁸	2 ¹⁷	2 ¹⁶	2 ¹⁵	2 ¹⁴	2 ¹³	2 ¹²	2 ¹¹	2 ¹⁰	2 ⁹	2 ⁸	2 ⁷	2 ⁶	2 ⁵	2 ⁴	2 ³	2 ²	2 ¹	2 ⁰
	0	0	1	0	x	x	x	x	x	x	x	x	x	x	x	x	1	0	0	1	x	x	x	x	x	0	0	0	0

NOTE: x = Don't Care

Figure 11-33. HISP Channel ROA Register Address Bit Format

Table 11-26 shows the bit format of the HISP_ROA register. Seven bits of address reference each entry of the HISP input or output buffer. The 7 bits of address are replicated in the HISP_ROA register to allow for partitioning of the HISP buffer integrated circuits.

Table 11-26. HISP_ROA Register Bit Format

Bits	Name
6 – 0	Bits 6 through 0 of the buffer address
7	Reserved
14 – 8	Copy of bits 6 through 0 of the buffer address
63 – 15	These bits are not used

Before a HISP input or output transfer, the internal buffer address is reset to 0 when the HISP circuitry is cleared. During the transfer, the internal buffer address increments when the HISP circuitry stores HISP input data in the input buffer or when the HISP circuitry reads HISP output data from the output buffer. The HISP_ROA register is used only to set the internal buffer address for testing purposes.

Syndrome Register (Address 4001000₁₆)

The syndrome register (IO_SYN) is a 56-bit, read-only register. The IO_SYN register contains the SECDED syndrome bits for errors that were detected on HISP input data and errors that were detected on data received from a microprocessor or a BLT transfer. Figure 11-34 shows the bit assignments for the IO_SYN register address as they appear on the address pins of the microprocessor.

HEX	0	4				0				0				1				0				0				0			
Binary	2 ²⁸	2 ²⁷	2 ²⁶	2 ²⁵	2 ²⁴	2 ²³	2 ²²	2 ²¹	2 ²⁰	2 ¹⁹	2 ¹⁸	2 ¹⁷	2 ¹⁶	2 ¹⁵	2 ¹⁴	2 ¹³	2 ¹²	2 ¹¹	2 ¹⁰	2 ⁹	2 ⁸	2 ⁷	2 ⁶	2 ⁵	2 ⁴	2 ³	2 ²	2 ¹	2 ⁰
	0	0	1	0	x	x	x	x	x	x	x	x	x	x	x	x	1	0	0	0	x	x	x	x	x	0	0	0	0

NOTE: x = Don't Care

Figure 11-34. Syndrome Register Address Bit Format

Table 11-27 shows the bit format of the IO_SYN register, and the following subsections describe each bit of the register.

Table 11-27. IO_SYN Register Bit Format

Bits	Name	Node
6 – 0	Microprocessor syndrome for data bits 15 – 0	Input or output
7	Microprocessor multiple-bit error on data bits 31 – 0	Input or output
15 – 8	HISP syndrome on data bits 63 – 0	Input
22 – 16	Microprocessor syndrome on data bits 31 – 16	Input or output
23	Microprocessor multiple-bit error on data bits 31 – 0	Input or output
31 – 24	HISP syndrome on data bits 63 – 0 (Do not use)	Input
38 – 32	Microprocessor syndrome on data bits 47 – 32	Input or output
39	Microprocessor multiple-bit error on data bits 63 – 32	Input or output
47 – 40	Reserved	Input or output
54 – 48	Microprocessor syndrome on data bits 63 – 48	Input or output
55	Microprocessor multiple-bit error on data bits 63 – 32	Input or output
63 – 56	These bits are not used	Input or output

HISP Syndrome

The HISP syndrome bits, bits 8 through 15, are valid only when the HISP single-bit error bit (bit 3) or the HISP multiple-bit error bit (bit 4) of the IO_IFLAG register is set to 1. After the first single-bit error occurs, and one or more single-bit errors occur, the syndrome stored in the IO_SYN register corresponds to the last single-bit error that occurred. When one or more multiple-bit errors occur, the syndrome stored in the IO_SYN register corresponds to the first multiple-bit error that occurred.

After the first multiple-bit error occurs, and one or more single-bit errors occur, the syndrome stored in the IO_SYN register corresponds to the first multiple-bit error that occurred. When one or more multiple-bit errors occur, the syndrome stored in the IO_SYN register corresponds to the first multiple-bit error that occurred.

NOTE: Bits 24 through 31 of the IO_SYN register also contain the HISP syndrome; however, HISP syndrome bits 0 through 3 reside in bits 28 through 31 of the IO_SYN register and HISP syndrome bits 4 through 7 reside in bits 24 through 27 of the IO_SYN register.

Microprocessor Syndrome

The different groups of microprocessor syndrome bits in the IO_SYN register may contain different values. The values depend on the type of errors that occur and which data bits the error occurs in. The microprocessor syndrome bits of the IO_SYN register are valid only when the microprocessor single-bit error bit (bit 1) or the microprocessor multiple-bit error bit (bit 2) of the IO_IFLAG register is set to 1.

When a single-bit error occurs in data bits 0 through 15, bits 0 through 6 of the IO_SYN register contain the syndrome bits for this error. When a single-bit error occurs in data bits 16 through 31, bits 16 through 22 of the IO_SYN register contain the syndrome bits for this error. When a multiple-bit error occurs in data bits 0 through 31, bits 0 through 6 and bits 16 through 22 of the IO_SYN register both contain the syndrome bits for this error. In addition, bits 7 and 23 of the IO_SYN register set to 1.

Likewise, when a single-bit error occurs in data bits 32 through 47, bits 32 through 38 of the IO_SYN register contain the syndrome bits for this error. When a single-bit error occurs in data bits 48 through 63, bits 48 through 54 of the IO_SYN register contain the syndrome bits for this error. When a multiple-bit error occurs in data bits 32 through 63, bits 32

through 38 and bits 48 through 54 of the IO_SYN register contain the syndrome bits for this error. In addition, bit 55 and 39 of the IO_SYN register set to 1.

Like the HISP syndrome, the microprocessor syndrome bits contain the syndrome for the last single-bit error that occurred. However, if a multiple-bit occurs the microprocessor syndrome bits contain the syndrome for the first multiple-bit error that occurred.

Clear Microprocessor Syndrome Register (Address 4001C00₁₆)

The clear microprocessor syndrome register (DEC_SYN_CLR) is a write-only register. When the microprocessor writes any value to the DEC_SYN_CLR register, the microprocessor single-bit error and microprocessor multiple-bit error bits (bits 1 and 2) of the IO_IFLAG register reset to 0. When these bits are reset to 0, the microprocessor syndrome bits in the IO_SYN register are no longer valid.

Figure 11-35 shows the bit assignments for the DEC_SYN_CLR register address as they appear on the address pins of the microprocessor.

HEX	0				4				0				0				1				C				0				0			
Binary	2 ²⁸	2 ²⁷	2 ²⁶	2 ²⁵	2 ²⁴	2 ²³	2 ²²	2 ²¹	2 ²⁰	2 ¹⁹	2 ¹⁸	2 ¹⁷	2 ¹⁶	2 ¹⁵	2 ¹⁴	2 ¹³	2 ¹²	2 ¹¹	2 ¹⁰	2 ⁹	2 ⁸	2 ⁷	2 ⁶	2 ⁵	2 ⁴	2 ³	2 ²	2 ¹	2 ⁰			
	0	0	1	0	x	x	x	x	x	x	x	x	x	x	x	x	1	1	1	0	x	x	x	x	x	0	0	0	0			

NOTE: x = Don't Care

Figure 11-35. Clear Microprocessor Syndrome Register Address Bit Format

Clear HISP Syndrome Register (Address 4001C20₁₆)

The clear HISP syndrome register (HISP_SYN_CLR) is a write-only register. When the microprocessor writes any value to the HISP_SYN_CLR register, the HISP single-bit error and HISP multiple-bit error bits (bits 3 and 4) of the IO_IFLAG register reset to 0. When these bits are reset to 0, the HISP syndrome bits in the IO_SYN register are no longer valid.

Figure 11-36 shows the bit assignments for the HISP_SYN_CLR register address as they appear on the address pins of the microprocessor.

HEX	0				4				0				0				1				C				2				0			
Binary	2 ²⁸	2 ²⁷	2 ²⁶	2 ²⁵	2 ²⁴	2 ²³	2 ²²	2 ²¹	2 ²⁰	2 ¹⁹	2 ¹⁸	2 ¹⁷	2 ¹⁶	2 ¹⁵	2 ¹⁴	2 ¹³	2 ¹²	2 ¹¹	2 ¹⁰	2 ⁹	2 ⁸	2 ⁷	2 ⁶	2 ⁵	2 ⁴	2 ³	2 ²	2 ¹	2 ⁰			
	0	0	1	0	x	x	x	x	x	x	x	x	x	x	x	x	1	1	1	0	x	x	x	1	x	0	0	0	0			

NOTE: x = Don't Care

Figure 11-36. Clear HISP Syndrome Register Address Bit Format

Heartbeat Generation Register (Address 4000C00₁₆)

The heartbeat generation (HB_GEN) register is a 2-bit, write-only register that is used only in the input node. The HB_GEN register selects what type of system heartbeat signal will be sent to the PEs in the CRAY T3D system when the next heartbeat signal is distributed. Figure 11-37 shows the bit assignments for the HB_GEN register address as they appear on the address pins of the microprocessor.

HEX	0	4				0				0				0				C				0				0			
Binary	2 ²⁸	2 ²⁷	2 ²⁶	2 ²⁵	2 ²⁴	2 ²³	2 ²²	2 ²¹	2 ²⁰	2 ¹⁹	2 ¹⁸	2 ¹⁷	2 ¹⁶	2 ¹⁵	2 ¹⁴	2 ¹³	2 ¹²	2 ¹¹	2 ¹⁰	2 ⁹	2 ⁸	2 ⁷	2 ⁶	2 ⁵	2 ⁴	2 ³	2 ²	2 ¹	2 ⁰
	0	0	1	0	x	x	x	x	x	x	x	x	x	x	x	0	1	1	0	x	x	x	x	x	x	0	0	0	0

NOTE: x = Don't Care

Figure 11-37. Heartbeat Generation Register Address Bit Format

Table 11-28 shows the bit format of the HB_GEN register. The 2 least significant bits of the data written to this register determine how many clock periods the heartbeat will be set. For more information on the heartbeat signals, refer to the “Control and Status” Section.

Table 11-28. HB_GEN Register Bit Format

Bits	Name
1 – 0	00 – Normal heartbeat 01 – Deadman heartbeat 10 – Selective reset heartbeat 11 – Global reset heartbeat
63 – 2	These bits are not used

This register is written to on the I/O gateway which has been designated the deadstart I/O gateway. The deadstart I/O gateway is the I/O gateway that receives the Master Clear LOSP input signal.

Table 12-1. Page Reference for the Memory-mapped Registers

Register	Name	Page
Barrier mask and interrupt register	BSMI	9-21
Barrier synchronization function register	BSFR	9-20
Barrier synchronization register 0	BSR0	9-18
Barrier synchronization register 1	BSR1	9-19
Barrier timing and refresh register	BAR_TMKG	9-23, 10-29
Clear microprocessor syndrome register	DEC_SYN_CLR	11-100
Clear HISP syndrome register	HISP_SYN_CLR	11-101
Clear I/O register	I/O_CLR	11-71
Control register	BLT_CR	8-30
DRAM control register	DRAM_CR	10-21
Fetch-and-increment register	F & I	7-50
Heartbeat generation register	HB_GEN	11-102
HISP channel address 0 register	HISP_ADDR0	11-92
HISP channel address 1 register	HISP_ADDR1	11-95
HISP buffer register	HISP_BUFFER	11-96
HISP channel address configuration register	HISP_A_CFG	11-81
HISP data configuration register	HISP_D_CFG	11-87
HISP channel lower block length register	HISP_L_BL	11-90
HISP channel ram-on-array address enter register	HISP_ROA	11-97
HISP channel upper block length register	HISP_U_BL	11-91
Index vector register	BLT_IVR	8-23
Interrupt flag register	IO_IFLAG	11-56
Interrupt mask register	IO_IMASK	11-63
Invalidate control register	INV_CR	10-19
I/O channel error register	IO_ERR	11-64
Local address register	BLT_LAR	8-28
Local stride register	BLT_LSR	8-29
Logical PE number register	L_WHOAMI	6-42
LOSP channel configuration register	LOSP_CFG	11-77

Table 12-1. Page Reference for the Memory-mapped Registers (continued)

Register	Name	Page
LOSP channel data in register	LOSP_DI	11-79
LOSP channel data out register	LOSP_DO	11-80
LOSP output channel disconnect register	LOSP_DISC	11-80
Message queue limit increment register	MQ_LIR	7-52
Message queue tail pointer	MQ_TP	7-51
Network enable register	NET_ENA	10-31
Network interface PE adjust register	LPE_XLATE	6-46
Network interface source register	X_WHOAMI	6-45
Network mode register	NODE_CSR	9-25, 10-34
Network performance register	NET_PFM	10-32
No-operation register	NOR	10-37
Physical PE number register	P_WHOAMI	6-39
Prefetch queue head register	PRE_FETCH	7-53
Remote base register	BLT_RBR	8-26
Remote index register	BLT_RIR	8-18
Remote limit register	BLT_RLR	8-27
Remote mask register	BLT_RMR	8-24
Remote stride register	BLT_RSR	8-21
Routing tag look-up table low-order bits register	ROUTE_LO	6-48
Routing tag look-up table high-order bits register	ROUTE_HI	6-53
Status register	BLT_SR	8-36
Swaperand register	SWAP	7-49
Syndrome register	IO_SYN	11-98
System control register	SCR	10-10
System status register	SSR	10-13
User status register	USR	10-18
Vector length register	BLT_VLR	8-22
Virtual PE number register	V_WHOAMI	6-43
Virtual PE range mask register	VRT_RG	6-44

The following tables list the address, name, and location of all the CRAY T3D memory-mapped registers. They are grouped by option series.

Table 12-2. A-series Options

Hex Address	Name	Access	Register	Option(s)	Term(s)
1040000 ₁₆	P_WHOAMI	S	Physical PE number register	AR	ppe0–11
1040400 ₁₆	BAR_TMG	S	Barrier timing and refresh register	AR	btxtmg00–07, rtset6–9
1010800 ₁₆	BSR0	G	Barrier synchronization register 0	AR	breg00–07
1090C00 ₁₆	BSR1	G	Barrier synchronization register 1	AR	breg08–15
1041000 ₁₆	SCR	S	System control register	AR	anxen00–01, clrhb, enbir0, rmask, gonode
1041400 ₁₆	BLT_SR	S	BLT status register	AR	bltsts00–02, hldbtre, rdbltir, r603–605, r640
1041800 ₁₆	BSMI	S	Barrier mask and interrupt register	AR	bmsk00–15
1041C00 ₁₆	BSFR	S	Barrier mode register	AR	bfcn00–15
1014400 ₁₆	USR	G	User status register	AM	rwpondg, pqmt
1054400 ₁₆	INV_CR	S	Invalidate control register	AM	ifltr5–12, enifltr, eninv
1044800 ₁₆	SSR	S	System status register	AM	mer, exper, mqfer, pfoer, pfuer, vper, reger, anxer, lockerr
1044C00 ₁₆	DRAM_CR	S	DRAM control register	AM	rdtmg0–1, wttmg, ifram, lock8k, bktmg0–1, dcfgact, dcmd0–4
1045000 ₁₆	L_WHOAMI	S	Logical PE number register	AM	lpe00–11
1045400 ₁₆	V_WHOAMI	S	Virtual PE number register	AM	vpe1–10
1045800 ₁₆	VRT_RG	S	Virtual PE range mask register	AM	vrn01–10
1046000 ₁₆	MQ_LIR	S	Message queue increment limit	AM	mqlir
1046400 ₁₆	MQ_TP	S	Message queue tail pointer	AM	mqt06–17
1018000 ₁₆	PRE_FETCH	G	Prefetch queue head register	AR	pfq00–35
1030000 ₁₆	SWAP	G	Swap operand register	AJ, AK	swap00–19

S = system privileged, G = general access

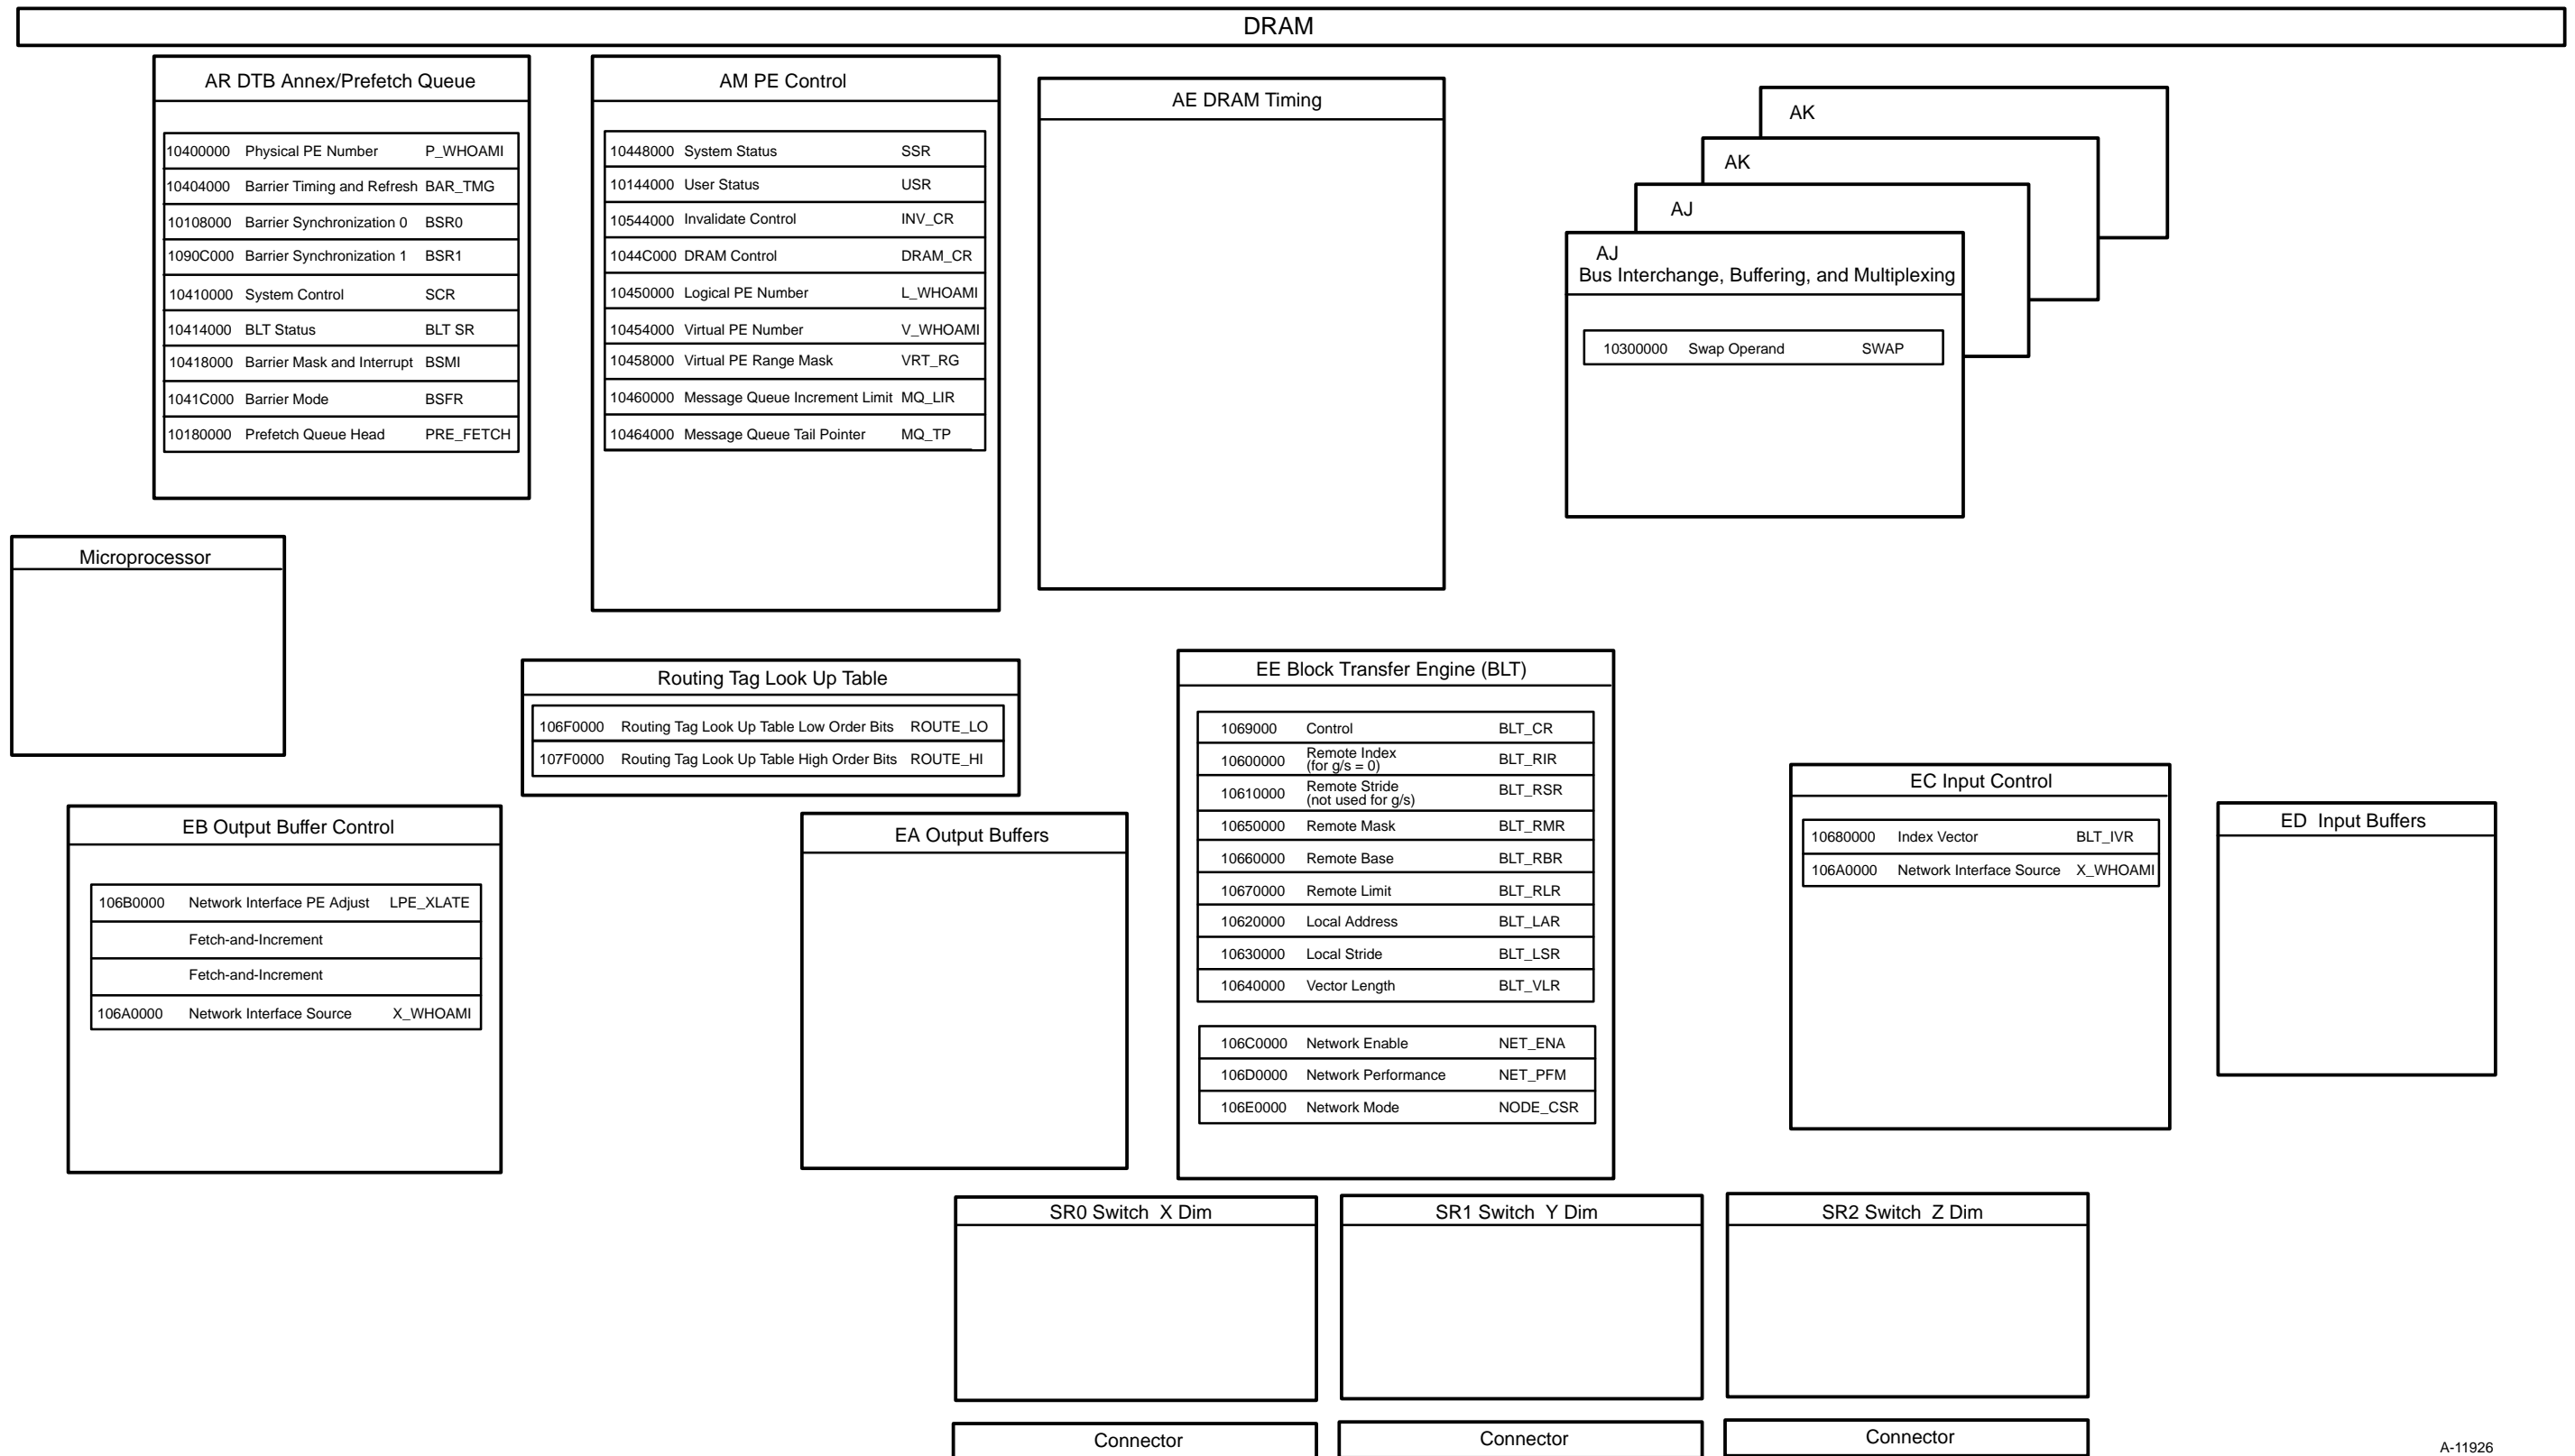
Table 12-3. E-series Options

Hex Address	Name	Access	Register	Option(s)	Term(s)
10600000 ₁₆	BLT_RIR	S	Remote index register	EE	b00–35
10610000 ₁₆	BLT_RSR	S	Remote stride register	EE	b40–75
10620000 ₁₆	BLT_LAR	S	Local address register	EE	l00–23
10630000 ₁₆	BLT_LSR	S	Local stride register	EE	l30–53
10640000 ₁₆	BLT_VLR	S	Vector length register	EE	x10–25
10650000 ₁₆	BLT_RMR	S	Remote mask register	EE	c00–35
10660000 ₁₆	BLT_RBR	S	Remote base register	EE	c40–63
10670000 ₁₆	BLT_RLR	S	Remote limit register	EE	c70–93
10680000 ₁₆	BLT_IVR	S	Index vector register	EC	j100–121
10690000 ₁₆	BLT_CR	S	Control register	EE	x00–4, x33–34, xs33–34, l100–110
106A0000 ₁₆	X_WHOAMI	S	Network interface source register	EB	a80–90
			Network interface source register	EC	j0–10
106B0000 ₁₆	LPE_XLATE	S	Network interface PE adjust register	EB	l100–130
106C0000 ₁₆	NET_ENA	S	Network enable register	EE	x60, r41–48
106D0000 ₁₆	NET_PFM	S	Network performance register	EE	r70–73, 74
106E0000 ₁₆	NODE_CSR	S	Network mode register	EE	r50–63
106F0000 ₁₆ 107F0000 ₁₆	ROUTE_LO ROUTE_HI	S	Routing tag look-up table	LT	q01–04
————	F & I	G	Fetch-and-increment register	EB	d00–31, d40–71

S = system privileged, G = general access

Table 12-4. I-series Options

Hex Address	Name	Register	Option(s)	Term(s)
04000000 ₁₆	LOSP_CFGI	LOSP channel configuration register	IC, ID	counta00–04, countb00–03
04000020 ₁₆	HISP_A_CFGI	HISP address configuration register	IC, ID	countc00–03, countd00–04, admode, r72–r73, r43, r44, r0–1, config0
04000040 ₁₆	HISP_D_CFGI	HISP data configuration register	IC, ID	rega00–03, regb00–03, regc00–03, regd00–03
04000200 ₁₆	IO_IFLAG	Interrupt flag register	IC, ID	intr00–11
04000220 ₁₆	IO_IMASK	Interrupt mask register	IC, ID	intrm00–11
04000400 ₁₆	LOSP_DI	LOSP data in register	IG, IH	l50–81, ld50–81
04000800 ₁₆	LOSP_DO	LOSP data out register	IG, IH	r40–47, l140–163
04000A00 ₁₆	LOSP_DISC	LOSP out disconnect register	IC, ID	r51
04000C00 ₁₆	HB_GEN	Heartbeat generation register	ID	data100–101
04000C00 ₁₆	IO_ERR	Error register	IC, ID	doutb00–15
04001000 ₁₆	IO_SYN	HISP syndrome	IE	r32–39, syn00–07
		DEC syndrome	IA	r50–53
04001200 ₁₆	HISP_ROA	HISP output ram-on-array address enter register	IG, IH	z23
04001400 ₁₆	HISP_ADDR0	HISP address 0 register	IG, IH	k40–71
04001600 ₁₆	HISP_ADDR1	HISP address 1 register	IG, IH	k72–79
04001800 ₁₆	HISP_U_BL	HISP upper block length register	IC, ID	htaa16–17
04001A00 ₁₆	HISP_L_BL	HISP lower block length register	IC, ID	htaa00–15
04001C00 ₁₆	DEC_SYN_CLR	Clear DEC syndrome register	IA	ecclr0
04001C20 ₁₆	HISP_SYN_CLR	Clear HISP syndrome register	IE	snc0
04001E00 ₁₆	IO_CLR	I/O clear register	IC, ID	liclrm, loclrm, clerrrm, hispc2, r27
07000000 ₁₆	HISP_BUFFER	HISP buffer base address register	IG, IH	ad20–26



A-11926

Figure 12-1. Processing Element Node Memory-mapped Registers

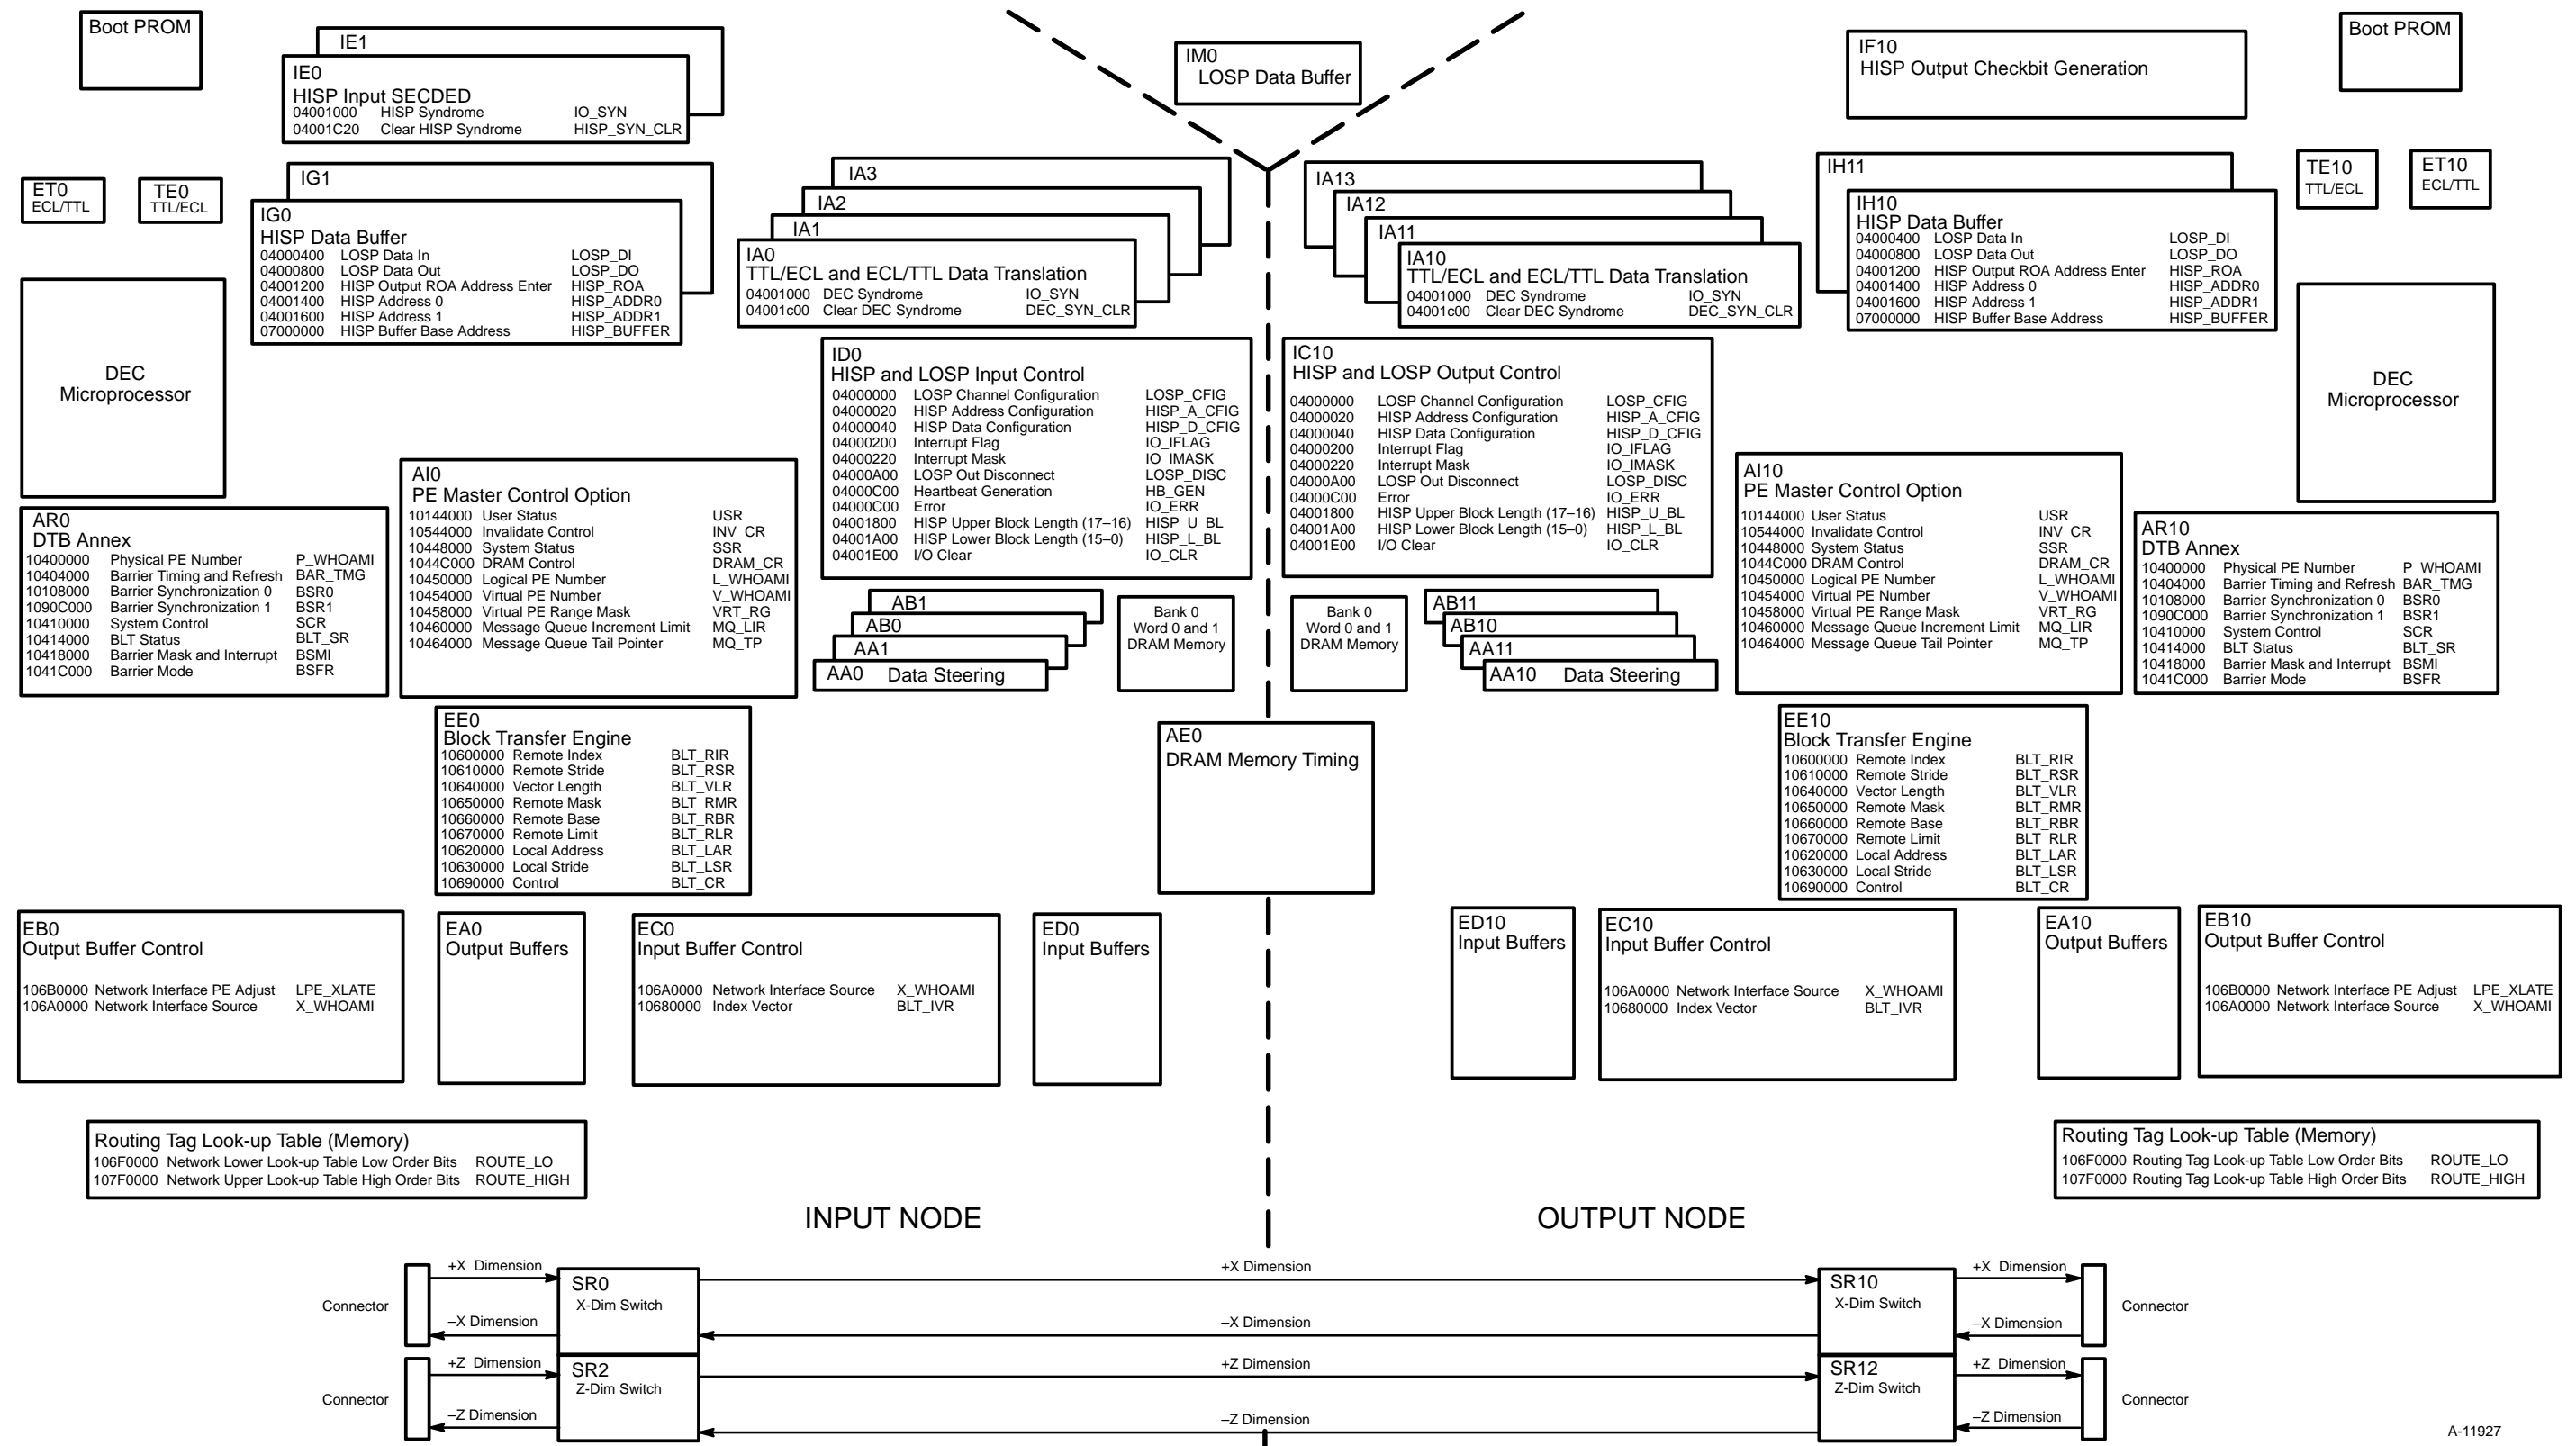
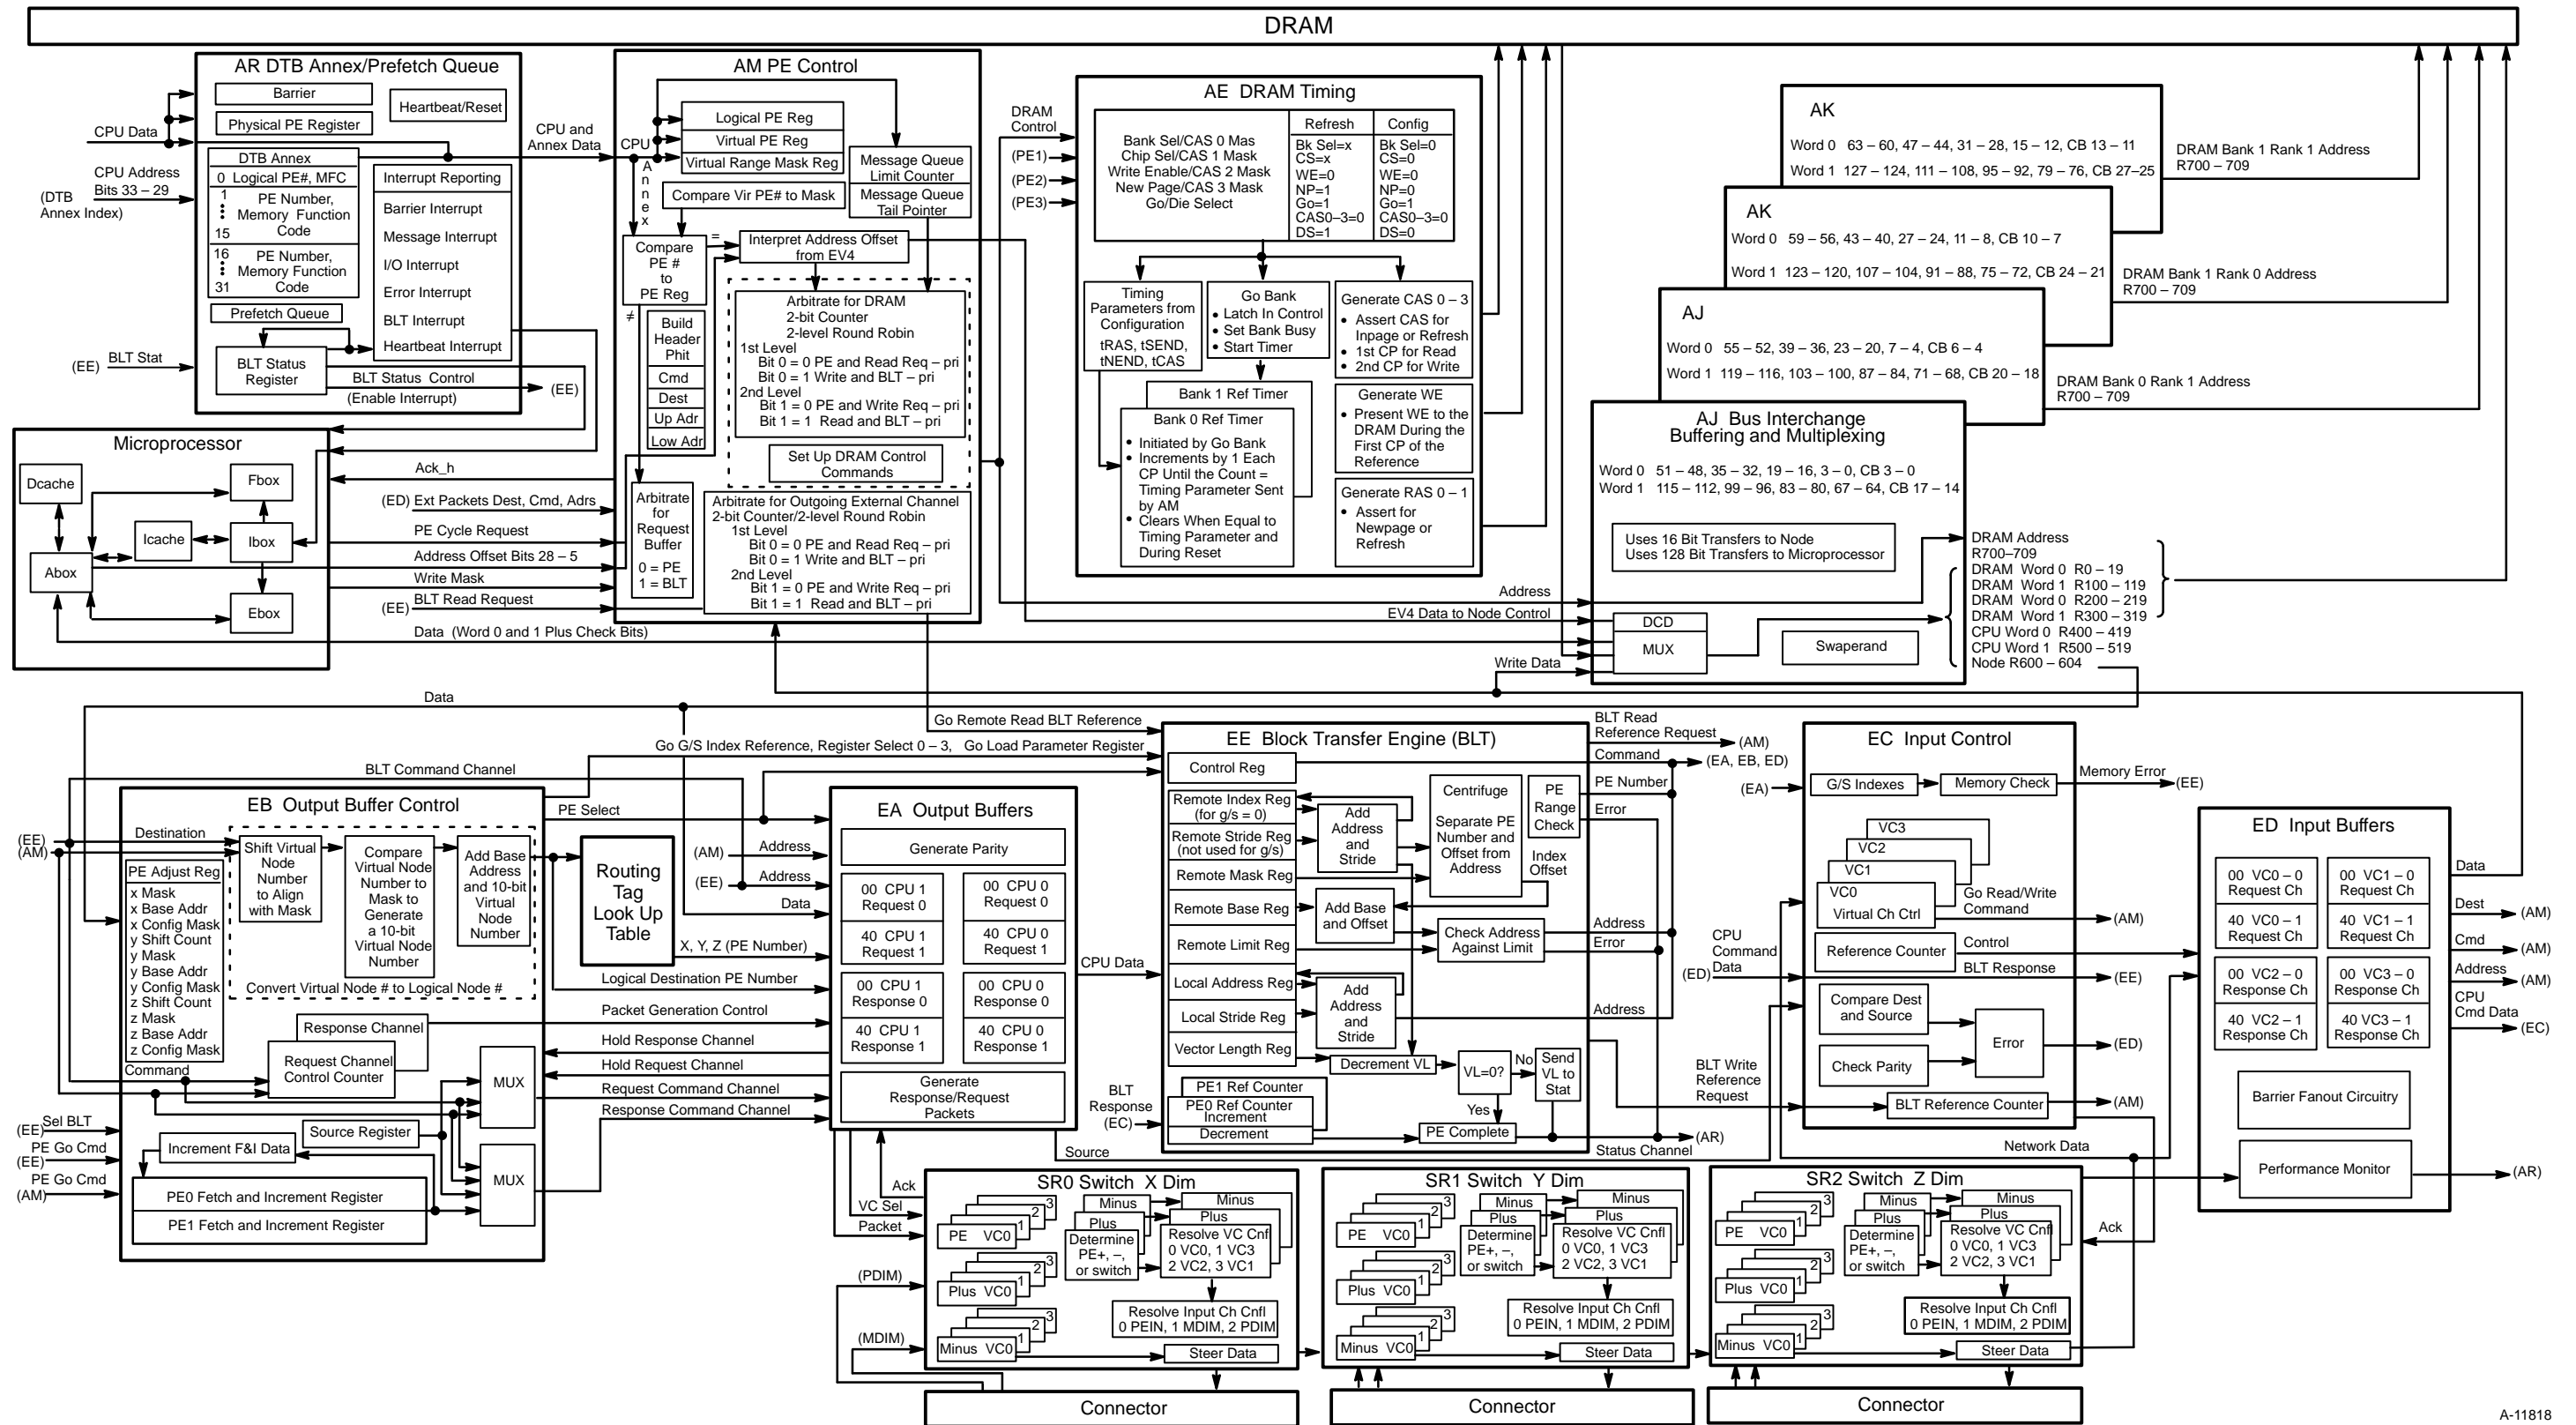


Figure 12-2. I/O Memory-mapped Registers

A-11927



A-11818

Figure 12-3. PE Module

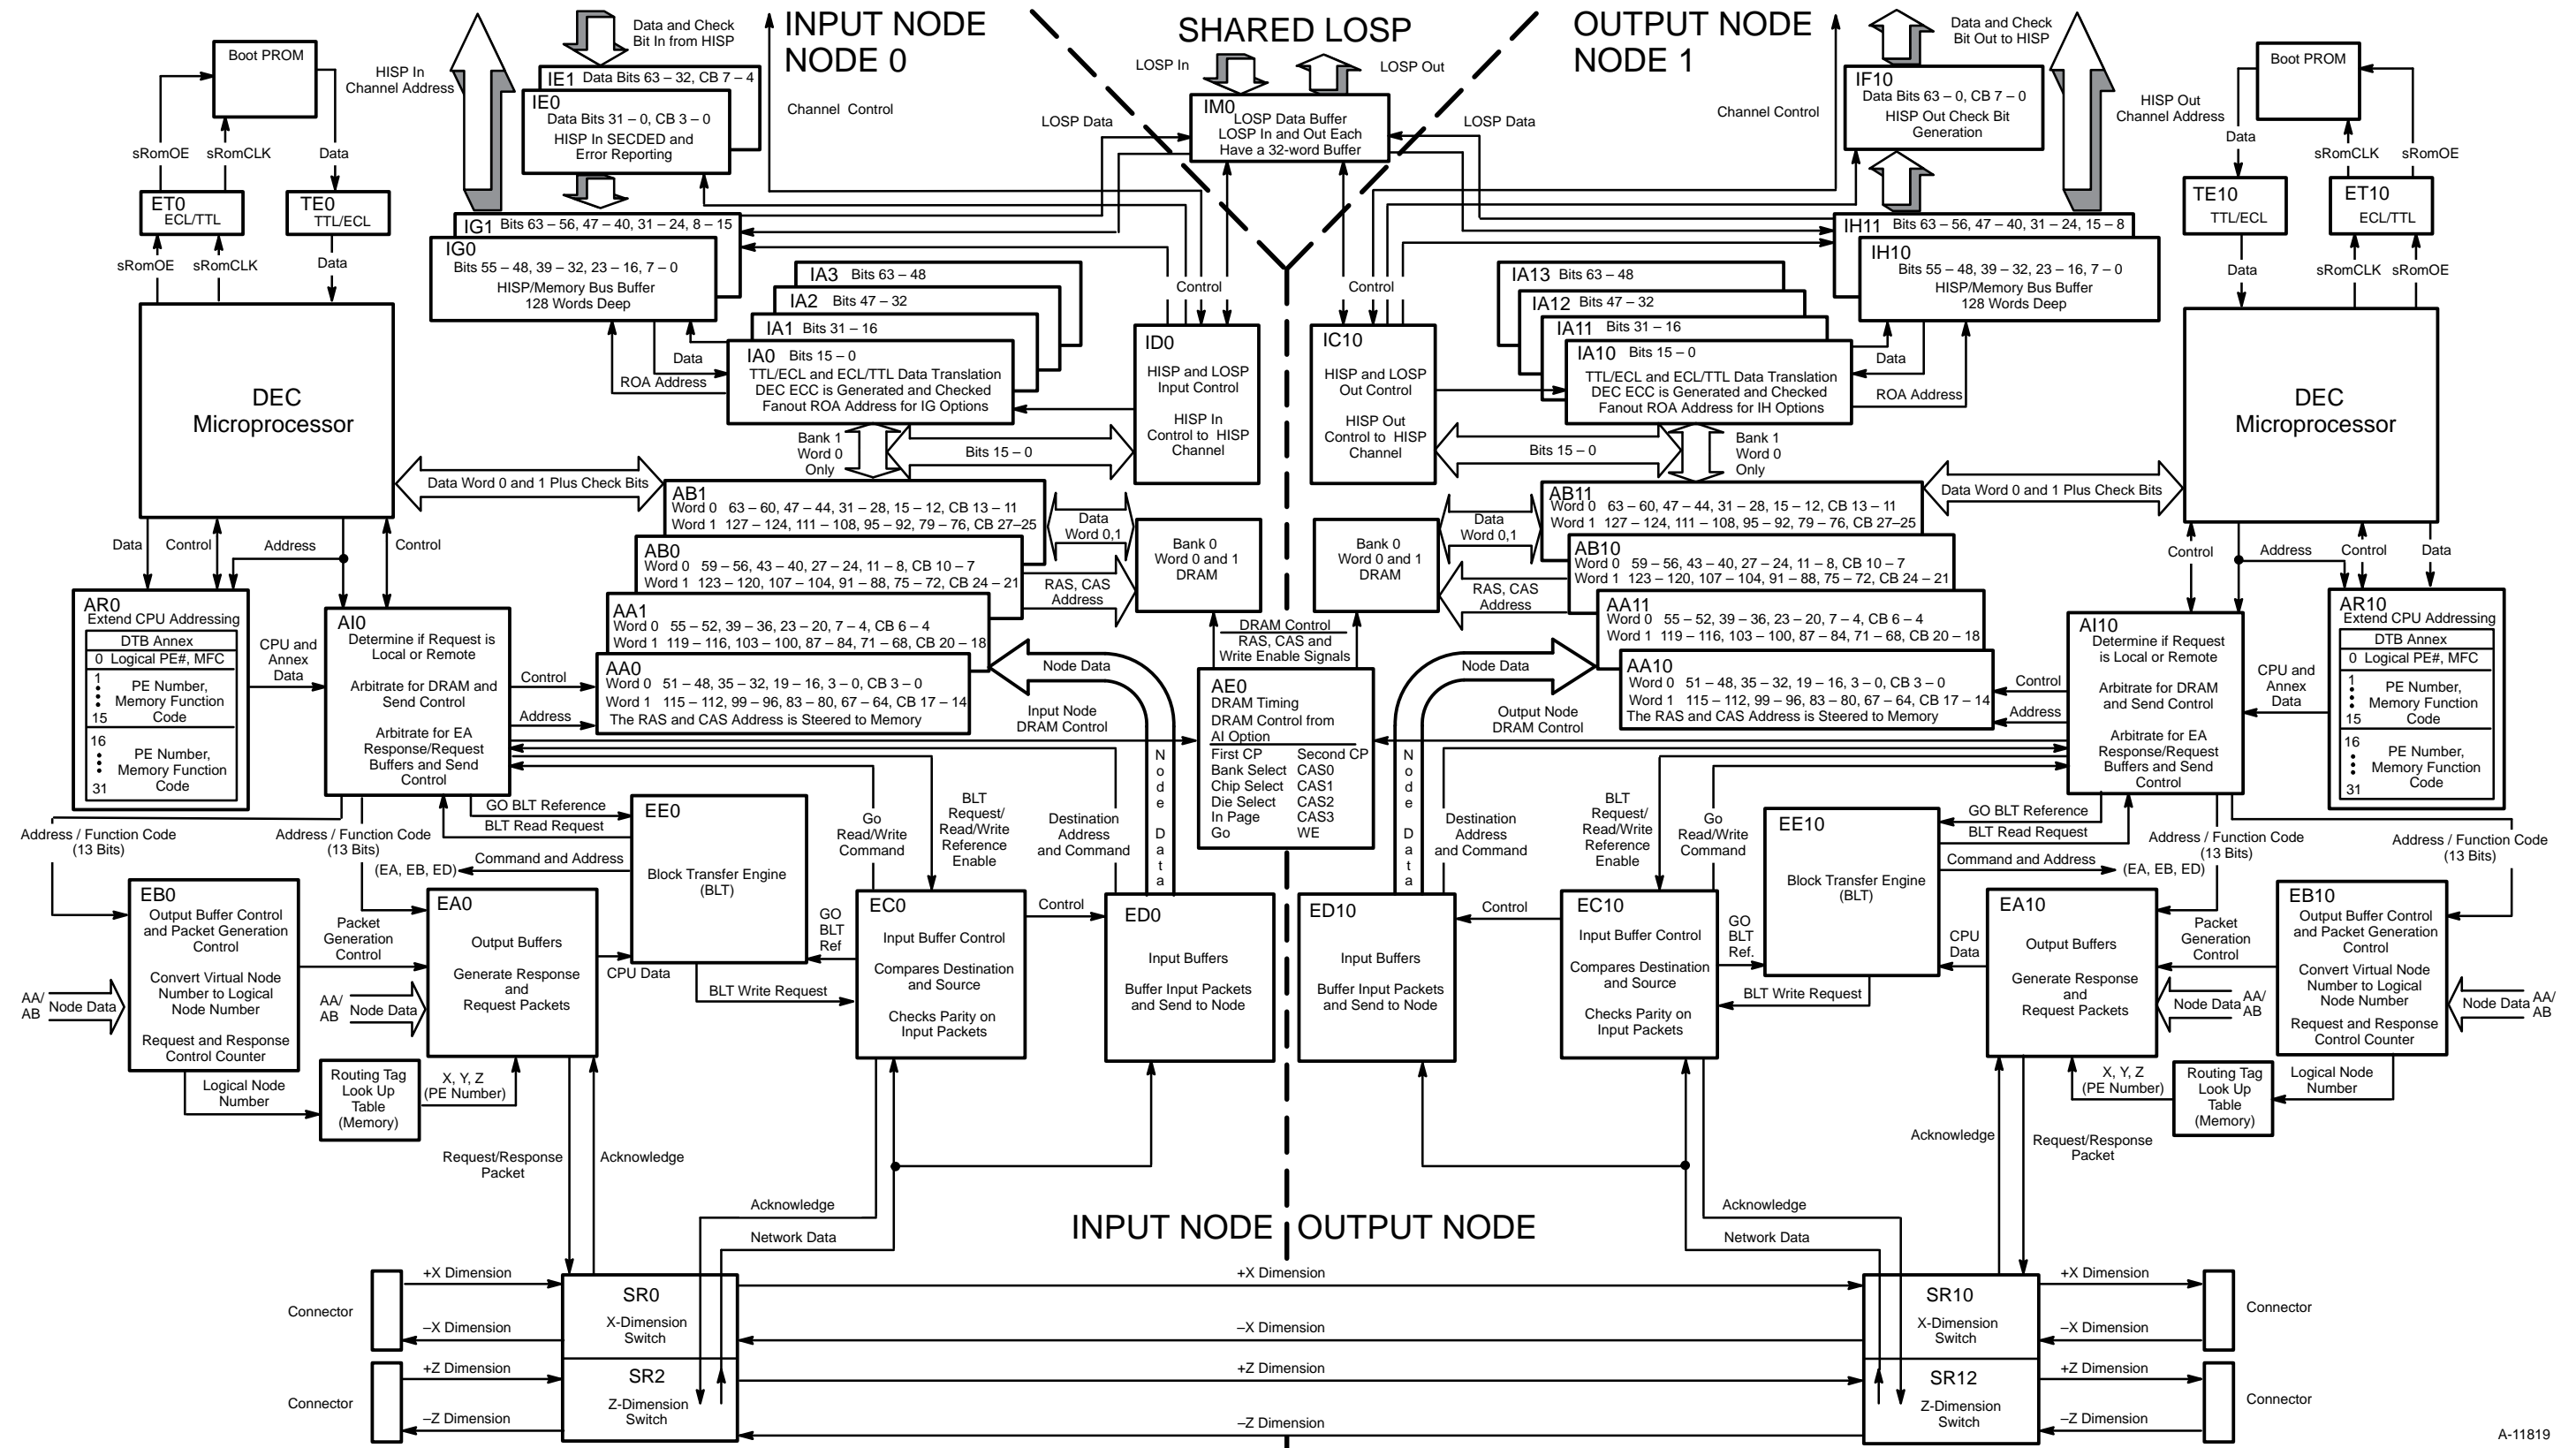


Figure 12-4. Master Input/Output Module

A-11819




Figure 12-6. Processing Element Module with Daughter Cards




Figure 12-8. I/O Module

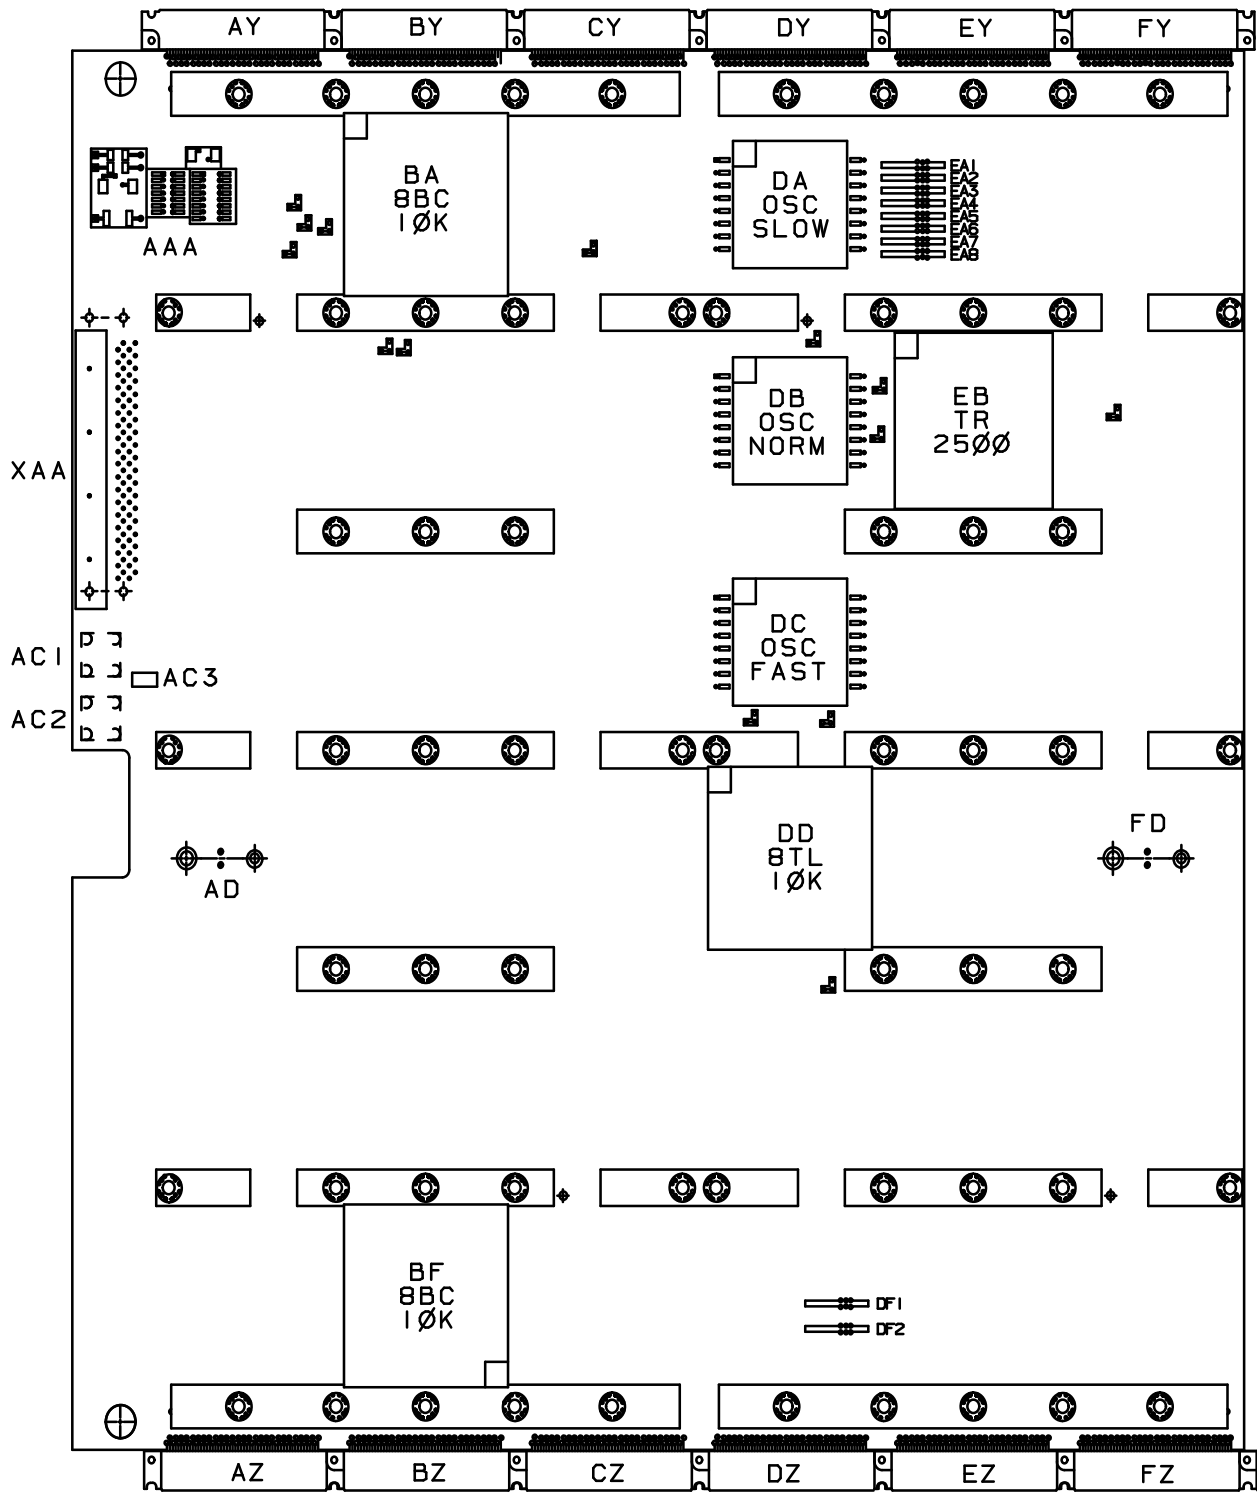


Figure 12-9. System Printed Circuit Board

13 CONTROL PANEL SWITCHES

This section describes the function of each of the switches on the CRAY T3D multiple-cabinet control panel and the CRAY T3D single-cabinet control panel.

Multiple-cabinet Control Panel Switches (6002 through 6015)

Figure 13-1 shows the multiple-cabinet control panel switches for the power of 2 (p2) systems. These switches are located on the upper half of the Y (left) side of a multiple-cabinet system chassis.

Not all of the chassis in a multiple-cabinet system have control panel switches. Table 13-1 lists the number of chassis in a multiple-cabinet system and shows the control panel locations for each system.

Table 13-1. Multiple-cabinet System Control Panel Locations (p2)

Number of Chassis	Control Panel Location
One Chassis	Chassis 0
Two Chassis	Chassis 0 only
Four Chassis	Chassis 0 and 3 only




Figure 13-1. Multiple-cabinet Control Panel Switches (p2)

LOSP ENABLE Switches

The LOSP ENABLE switches enable or disable the low-speed (LOSP) channel for each of the I/O gateways in the system (refer to Figure 13-2). When a LOSP channel is disabled, the external LOSP device may be powered up and down or may be disconnected without resetting the system.




Figure 13-2. Multiple-cabinet (p2) LOSP ENABLE Switches

Table 13-2 shows the switch labels and lists the corresponding I/O gateways that the switches control. In this table, each I/O gateway number is represented as a three-digit hexadecimal number. This number corresponds to the number read from the physical PE number (P_WHOAMI) register in the input node of the I/O gateway.

Table 13-2. Multiple-cabinet Switch Label and P_WHOAMI Values

Switch Label	P_WHOAMI Values			
	Chassis 0	Chassis 1	Chassis 2	Chassis 3
IOM 80	C20	C60	CA0	CE0
IOM 81	C02	C42	C82	CC2
IOM 82	C24	C64	CA4	CE4
IOM 83	C06	C46	C86	CC6
IOM 84	C08	C48	C88	CC8
IOM 85	C2A	C6A	CAA	CEA
IOM 86	C0C	C4C	C8C	CCC
IOM 87	C2E	C6E	CAE	CEE

CLOCK SELECT Switches

The CLOCK SELECT switches select one of three oscillators or an external generator for use as the system clock. Figure 13-3 shows the possible settings for the CLOCK SELECT switches. (The switch labeled N/A is not used.)




Figure 13-3. Multiple-cabinet (p2) CLOCK SELECT Switches

SYSTEM SIZE Switches

The SYSTEM SIZE switches indicate the number of chassis used in the system. The system module uses this value to configure the BC options for the appropriate system size. Figure 13-4 shows the possible settings for the SYSTEM SIZE switches. (The switch labeled N/A is not used.)




Figure 13-4. Multiple-cabinet (p2) SYSTEM SIZE Switches

CHASSIS SELECT Switches

The CHASSIS SELECT switches indicate the number of the chassis that the control panel resides on. This value is used to set bits 7 and 6 of the P_WHOAMI register for each PE in the chassis and is used to calculate the value for bits 7 and 6 of the P_WHOAMI register for each PE in the neighboring chassis.

Figure 13-5 shows the possible settings for the CHASSIS SELECT switches. Note that the chassis select 1 and chassis select 2 are not valid settings. This is because chassis 1 and chassis 2 in a multiple-cabinet system do not contain control panels. (The switch labeled N/A is not used.)




Figure 13-5. Multiple-cabinet (p2) CHASSIS SELECT Switches

HEARTBEAT INTERVAL Switches

The HEARTBEAT INTERVAL switches set the interval between system heartbeat signals, which are periodically sent to all of the PEs in the system. Figure 13-6 shows the possible settings for the HEARTBEAT INTERVAL switches and shows the corresponding heartbeat interval values.

NOTE: The recommended setting at this time is 10.3 ms but may be set to a different interval if desired by the site analyst.




Figure 13-6. Multiple-cabinet (p2) HEARTBEAT INTERVAL Values

HEARTBEAT ENABLE Switch

The HEARTBEAT ENABLE switch enables or disables the system heartbeat signal. The HEARTBEAT ENABLE switch must be enabled to run the operating system. When the HEARTBEAT ENABLE switch is disabled, software controls all events. (The system module does not generate any 1-CP heartbeat events.) Figure 13-7 shows the possible settings for the HEARTBEAT ENABLE switch.




Figure 13-7. Multiple-cabinet (p2) HEARTBEAT ENABLE Switch

HEARTBEAT SELECT Switch

The HEARTBEAT SELECT switch sets the interval at which heartbeat requests from the I/O gateways are serviced. Figure 13-8 shows the possible settings for the HEARTBEAT SELECT switch.




Figure 13-8. Multiple-cabinet (p2) HEARTBEAT SELECT Switch

REMOTE ENABLE Switch

The REMOTE ENABLE switch enables or disables remote access to the control panel settings via the RS-232 channel. When remote access is disabled, remote reads are enabled; however, remote writes are disabled. Also when remote access is disabled, the control panel switch settings override the remote settings.

Figure 13-9 shows the possible settings for the REMOTE ENABLE switch.




Figure 13-9. Multiple-cabinet (p2) REMOTE ENABLE Switch

Multiple-cabinet Control Panel Switches (6001, 6016 through 6099)

Figure 13-10 shows the multiple-cabinet control panel switches for the non-power of 2 (np2) systems. These switches are located on the upper half of the Y (left) side of a multiple-cabinet system chassis.

Not all of the chassis in a multiple-cabinet system have control panel switches. Table 13-3 lists the number of chassis in a multiple-cabinet system and shows the control panel locations for each system.

Table 13-3. Multiple-cabinet System Control Panel Locations (np2)

Number of Chassis	Control Panel Location
One Chassis	Chassis 0
Two Chassis	Chassis 0 only
Four Chassis	Chassis 0 and 3 only




Figure 13-10. Multiple-cabinet Control Panel Switches (np2)

LOSP ENABLE Switches

The LOSP ENABLE switches enable or disable the low-speed (LOSP) channel for each of the I/O gateways in the system (refer to Figure 13-11). When a LOSP channel is disabled, the external LOSP device may be powered up and down or may be disconnected without resetting the system.

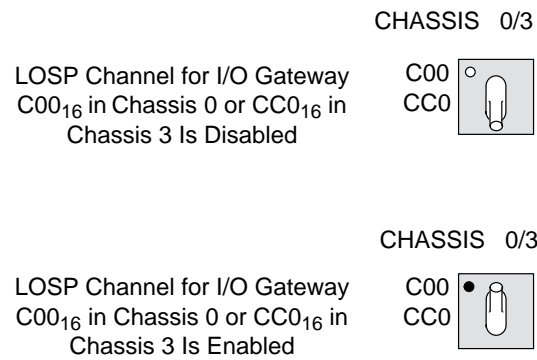


Figure 13-11. Multiple-cabinet (np2) LOSP ENABLE Switches

CLOCK SELECT Switches

The CLOCK SELECT switches select one of three oscillators or an external generator for use as the system clock. Figure 13-12 shows the possible settings for the CLOCK SELECT switches. (The switch labeled N/A is not used.)




Figure 13-12. Multiple-cabinet (np2) CLOCK SELECT Switches

SYSTEM SIZE Switches

The SYSTEM SIZE switches indicate the number of chassis used in the system. The system module uses this value to configure the BC options for the appropriate system size. Figure 13-13 shows the possible settings for the SYSTEM SIZE switches. (The switch labeled N/A is not used.)




Figure 13-13. Multiple-cabinet (np2) SYSTEM SIZE Switches

CHASSIS SELECT Switches

The CHASSIS SELECT switches indicate the number of the chassis that the control panel resides on. This value is used to set bits 7 and 6 of the P_WHOAMI register for each PE in the chassis and is used to calculate the value for bits 7 and 6 of the P_WHOAMI register for each PE in the neighboring chassis.

Figure 13-14 shows the possible settings for the CHASSIS SELECT switches. Note that the chassis select 1 and chassis select 2 are not valid settings. This is because chassis 1 and chassis 2 in a multiple-cabinet system do not contain control panels. (The switch labeled N/A is not used.)




Figure 13-14. Multiple-cabinet (np2) CHASSIS SELECT Switches

HEARTBEAT INTERVAL Switches

The HEARTBEAT INTERVAL switches set the interval between system heartbeat signals, which are periodically sent to all of the PEs in the system. Figure 13-15 shows the possible settings for the HEARTBEAT INTERVAL switches and shows the corresponding heartbeat interval values.

NOTE: The recommended setting at this time is 10.3 ms but may be set to a different interval if desired by the site analyst.




Figure 13-15. Multiple-cabinet (np2) HEARTBEAT INTERVAL Values

HEARTBEAT ENABLE Switch

The HEARTBEAT ENABLE switch enables or disables the system heartbeat signal. The HEARTBEAT ENABLE switch must be enabled to run the operating system. When the HEARTBEAT ENABLE switch is disabled, software controls all events. (The system module does not generate any 1-CP heartbeat events.) Figure 13-16 shows the possible settings for the HEARTBEAT ENABLE switch.




Figure 13-16. Multiple-cabinet (np2) HEARTBEAT ENABLE Switch

HEARTBEAT SELECT Switch

The HEARTBEAT SELECT switch sets the interval at which heartbeat requests from the I/O gateways are serviced. Figure 13-17 shows the possible settings for the HEARTBEAT SELECT switch.




Figure 13-17. Multiple-cabinet (np2) HEARTBEAT SELECT Switch

REMOTE ENABLE Switch

The REMOTE ENABLE switch enables or disables remote access to the control panel settings via the RS-232 channel. When remote access is disabled, remote reads are enabled; however, remote writes are disabled. Also when remote access is disabled, the control panel switch settings override the remote settings.

Figure 13-18 shows the possible settings for the REMOTE ENABLE switch.




Figure 13-18. Multiple-cabinet (np2) REMOTE ENABLE Switch

Single-cabinet Control Panel Switches

Figure 13-19 shows the single-cabinet control panel switches. These switches are located on the upper half of the Y (left) side of a single-cabinet system chassis. This control panel contains switches for both the CRAY T3D system and the CRAY Y-MP system.

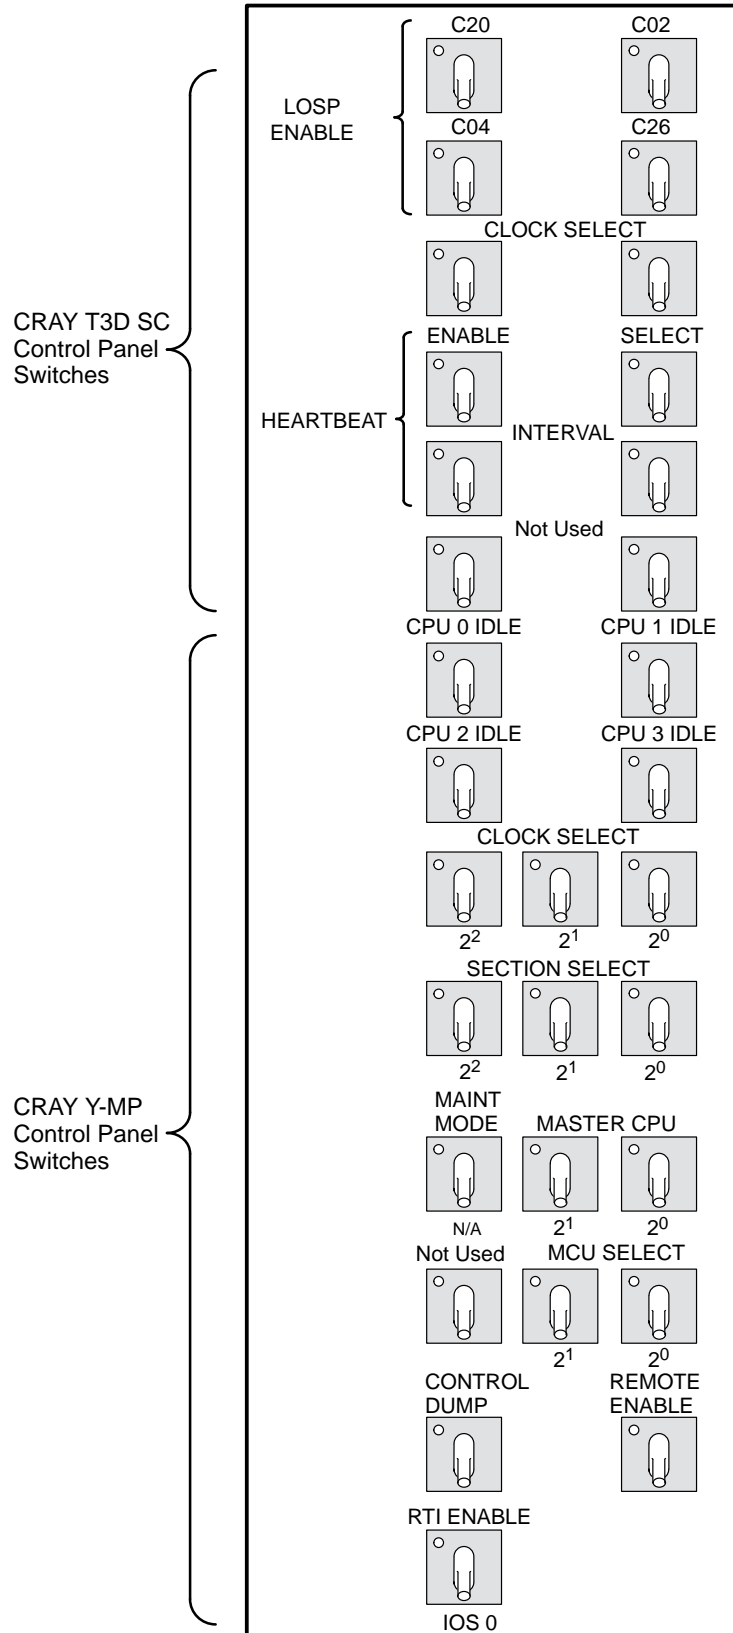


Figure 13-19. Single-cabinet Control Panel Switches

LOSP ENABLE Switches

The LOSP ENABLE switches enable or disable the LOSP channel for each of the I/O gateways in the system (refer to Figure 13-20). When a LOSP channel is disabled, the external LOSP device may be powered up and down or may be disconnected without resetting the system.




Figure 13-20. Single-cabinet LOSP ENABLE Switches

CLOCK SELECT Switches

The CLOCK SELECT switches select one of three oscillators or an external generator for use as the system clock. Figure 13-21 shows the possible settings for the CLOCK SELECT switches.




Figure 13-21. Single-cabinet CLOCK SELECT Switches

HEARTBEAT INTERVAL Switches

The HEARTBEAT INTERVAL switches set the interval between system heartbeat signals, which are periodically sent to all of the PEs in the system. Figure 13-22 shows the possible settings for the HEARTBEAT INTERVAL switches and shows the corresponding heartbeat interval values.

NOTE: The recommended setting at this time is 10.3 ms but may be set to a different interval if desired by the site analyst.




Figure 13-22. Single-cabinet HEARTBEAT INTERVAL Values

HEARTBEAT ENABLE Switch

The HEARTBEAT ENABLE switch enables or disables the system heartbeat signal. The HEARTBEAT ENABLE switch must be enabled to run the operating system. Figure 13-23 shows the possible settings for the HEARTBEAT ENABLE switch.




Figure 13-23. Single-cabinet HEARTBEAT ENABLE Switch

HEARTBEAT SELECT Switch

The HEARTBEAT SELECT switch sets the interval at which heartbeat requests from the I/O gateways are serviced. Figure 13-24 shows the possible settings for the HEARTBEAT SELECT switch.




Figure 13-24. Single-cabinet HEARTBEAT SELECT Switch

CPU IDLE Switches

The CPU IDLE switches inactivate selected CPUs. When a CPU is idled, the CPU Master Clear signal for the selected CPU is forced to a 1. Figure 13-25 shows some of the possible settings for the CPU IDLE switches.

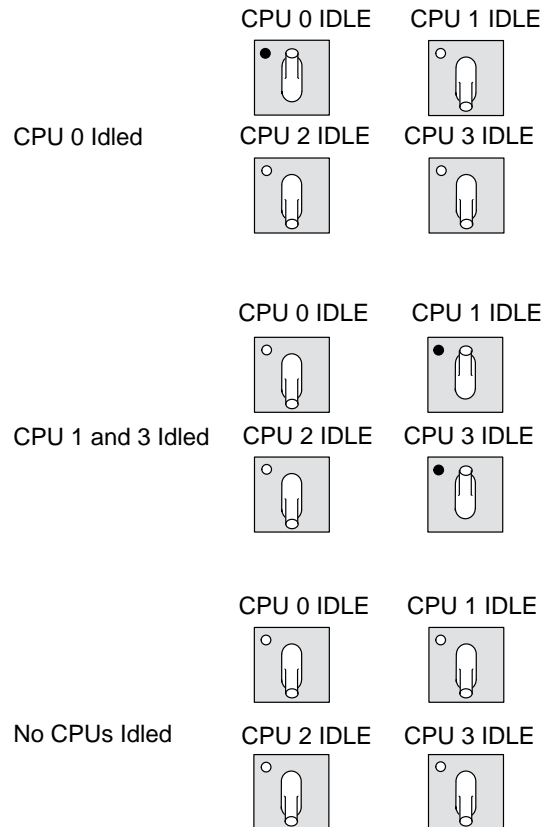


Figure 13-25. CRAY Y-MP CPU IDLE Switches

CLOCK SELECT Switches

The CLOCK SELECT switches select the clock speed for the CRAY Y-MP system. Figure 13-26 shows the possible settings for the CLOCK SELECT switches.




Figure 13-26. CRAY Y-MP CLOCK SELECT Switches

SECTION SELECT Switches

The SECTION SELECT switches select the number of memory sections used. Figure 13-27 shows the possible settings for the SECTION SELECT switches.




Figure 13-27. CRAY Y-MP SECTION SELECT Switches

MASTER CPU Switches

The MASTER CPU switches select one of the four CPUs as logical processor 0 or master CPU. Logical processor 0 is the CPU that is deadstarted first while running multi-CPU diagnostics, or it is the only CPU that executes the diagnostic code while running a single-CPU diagnostic.

Figure 13-28 shows the possible settings for the MASTER CPU switches.

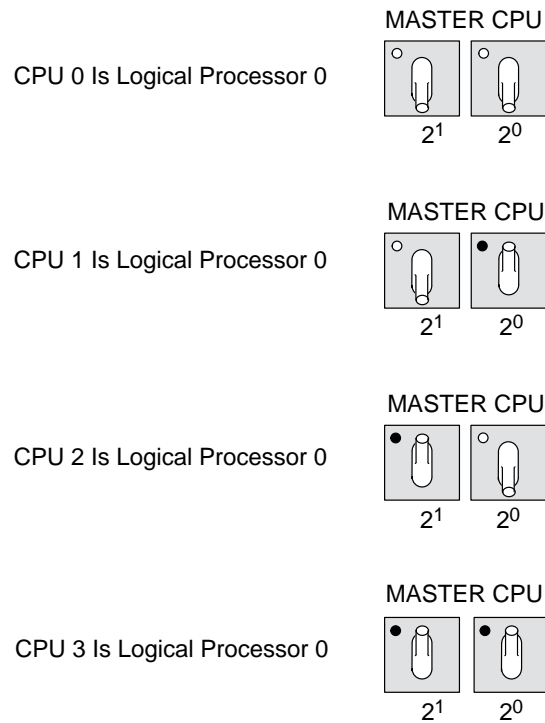


Figure 13-28. CRAY Y-MP MASTER CPU Switches

MAINTENANCE MODE Switch

The MAINTENANCE MODE switch enables and disables error correction on all CPU ports and special maintenance instructions. Maintenance instructions enable testing of SECCDED and register parity circuitry. Figure 13-29 shows the possible settings for the MAINTENANCE MODE switch and the corresponding maintenance mode values.




Figure 13-29. CRAY Y-MP MAINTENANCE MODE Switch

MCU SELECT Switches

The MCU SELECT switches select one of the four CPUs to handle the maintenance control unit (MCU) LOSP channel control signals. Figure 13-30 shows the possible settings for the MCU SELECT switches.

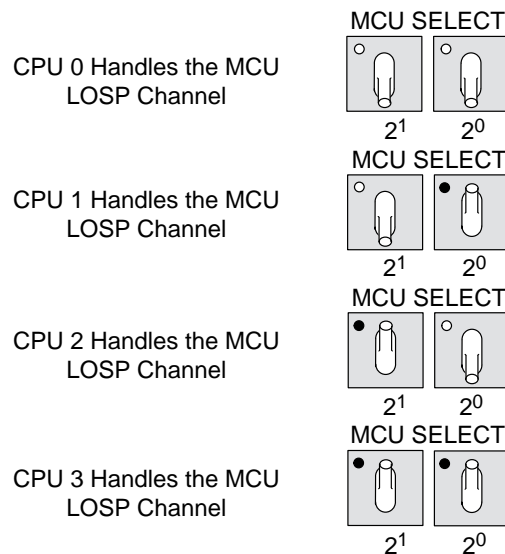


Figure 13-30. CRAY Y-MP MCU SELECT Switches

CONTROL DUMP Switch

The CONTROL DUMP switch enables or disables the control dump feature which allows the user to dump the contents of the instruction buffers to the maintenance workstation. Figure 13-31 shows the possible settings of the CONTROL DUMP switch.




Figure 13-31. CRAY Y-MP CONTROL DUMP Switch

REMOTE ENABLE Switch

The REMOTE ENABLE switch enables or disables remote access to the control panel settings via the RS-232 channel. When remote access is disabled, remote reads are enabled; however, remote writes are disabled. Also when remote access is disabled, the control panel switch settings override the remote settings.

Figure 13-32 shows the possible settings for the REMOTE ENABLE switch.

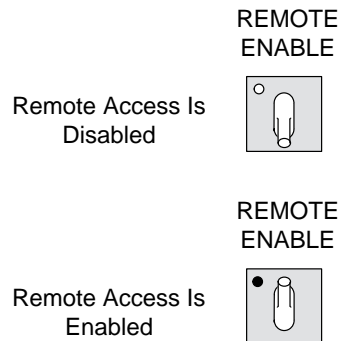


Figure 13-32. Single-cabinet REMOTE ENABLE Switch

RTI ENABLE Switch

The RTI ENABLE switch enables or disables the CRAY Y-MP CPU real-time interrupts from the PINT board on the IOS. Figure 13-33 shows the possible settings of the RTI enable switch.

NOTE: To run offline diagnostics, the RTI enable switch must be disabled.




Figure 13-33. CRAY Y-MP RTI ENABLE Switch

GLOSSARY

A

- Address Offset** The address offset is a byte oriented address that points to a byte in the local memory of a PE.
- Address Space Partition Bits** The address space partition bits divide the address space of the CRAY T3D system into local memory with SECDED protection, local memory without SECDED protection, and memory-mapped register space.
- AE Option** The AE option generates the CAS 0–3, RAS 0–1, and WE signals for the four PEs on a printed circuit board. Each node must initialize its portion of the AE option before it can make any memory references.
- AJ Option** The AJ option transfers data between local memory, the microprocessor, and the network interface logic. The AJ option also distributes the address offset to local memory and contains the swaperand memory-mapped register.
- AK Option** The AK option transfers data between local memory, the microprocessor, and the network interface logic. The AK option also distributes the address offset to local memory and contains the swaperand memory-mapped register.
- AM Option** The AM option provides control for local and remote memory references.
- AR Option** The AR option handles all hardware interrupts. The AR option also contains the DTB annex and the data prefetch queue.
- Atomic Swap Operation** An atomic swap operation retrieves a 64-bit word from a location in system memory and sends it to the requesting microprocessor. This location in system memory is then written with the contents of the swaperand register.

B

Barrier Synchronization	A barrier is used to synchronize the microprocessors within a partition. All microprocessors within the partition must reach the barrier before the microprocessors can continue with the program.
Barrier Synchronization Function Register (BSFR)	The barrier synchronization function register is a 16-bit, write-only, system privileged register. The BSFR contains a 16-bit mask that indicates which bits of BSR0 and BSR1 are used for barrier synchronization and which bits are used for eureka synchronization.
Barrier Synchronization Mask and Interrupt Register (BSMI)	The barrier synchronization mask and interrupt register is a 16-bit, readable and writable, system privileged register. When written to, the BSMI register indicates which bits in BSR0 and BSR1 are enabled for use by the PE. When read from, the BSMI register provides the current state of the barrier interrupts from BSR0 and BSR1 and then clears the interrupts.
Barrier Synchronization Register 0 (BSR0)	The barrier synchronization register 0 is an 8-bit, readable and writable, general access register. BSR0 contains the 8 least significant barrier bits for a PE. When read from, the value of BSR0 represents the value of bits 0 through 7 of BSR0. The remaining bits are not valid.
Barrier Synchronization Register 1 (BSR1)	The barrier synchronization register 1 is an 8-bit, readable and writable, general access register. BSR1 contains the 8 most significant barrier bits for a PE. When read from, the value of BSR1 represents the value of bits 8 through 15 of BSR1 and 0 through 7 of BSR0. The remaining bits are not valid.
Barrier Timing and Refresh Register (BAR_TMGM)	The barrier timing and refresh register is a 12-bit, write-only, system privileged register. The BAR_TMGM register controls the timing of each physical barrier synchronization circuit.
Block Transfer Engine (BLT)	Block Transfer Engine (BLT) – The block transfer engine is an asynchronous direct memory access device that redistributes system data. The block transfer engine operates independently from the PEs and contends with the PEs for access to memory.
BLT Centrifuge	The BLT centrifuge separates the remote index into a PE number and an index offset.
Body	The body of a packet contains memory addresses, system data, or both. The body has a variable length and may not be used in all packets.

Bypass Point A bypass point is a barrier synchronization circuit that consists of an AND gate and a fan-out block.

C

Cache Line Invalidate Pins The cache invalidate pins are used to signal the microprocessor when a cache line of data in the data cache should be discarded.

Cached Read Operation The cached read operation transfers data from system memory to the data cache in the microprocessor.

Clear HISP Syndrome Register (HISP_SYN_CLR) The clear HISP syndrome register is a write-only register. When the microprocessor writes any value to the HISP_SYN_CLR register, the HISP single-bit error and HISP multiple-bit error bits (bits 3 and 4) of the IO_IFLAG register reset to 0. When these bits are reset to 0, the HISP syndrome bits in the IO_SYN register are no longer valid.

Clear I/O Register (IO_CLR) The clear I/O register is a 10-bit write-only register. The IO_CLR register clears various bits and hardware in the I/O gateway.

Clear Microprocessor Syndrome Register (DEC_SYN_CLR) The clear microprocessor syndrome register is a write-only register. When the microprocessor writes any value to the DEC_SYN_CLR register, the microprocessor single-bit error and microprocessor multiple-bit error bits (bits 1 and 2) of the IO_IFLAG register reset to 0. When these bits are reset to 0, the microprocessor syndrome bits in the IO_SYN register are no longer valid.

Communication Links Communication links transfer data and control information between the nodes in the interconnect network. Each communication link connects two nodes in one dimension.

Control Register (BLT_CR) The control register is a 27-bit, write-only, system-privileged register. The BLT_CR contains bits that control the type of operation the BLT performs, define the range of legal PE numbers, and enable a transfer.

CRAY T3D System The CRAY T3D system is a multiprocessor multiple instruction, multiple data (MIMD) system with physically distributed, logically shared memory. It may also be used to emulate single instruction, multiple data (SIMD) or multicomputer MIMD systems.

Cycle Acknowledge Pins The cycle acknowledge pins signal the microprocessor that a cycle request completed successfully or that the PE support circuitry detected an error while performing the cycle request.

Cycle Request Code The cycle request code is a 3-bit code that the microprocessor supplies to the PE support circuitry to indicate it wants the PE support circuitry to perform an operation.

D

Data Cache Data cache memory is a small, high-speed random access memory that temporarily stores frequently or recently accessed data. The data cache memory is internal to the microprocessor and stores 256 32-byte lines (four 64-bit words in a line) of data.

Data Cache Miss A data cache miss occurs when the data that the microprocessor needs is not located in the data cache.

Data Prefetch Queue The data prefetch queue stores the data when the microprocessor issues a prefetch instruction. A prefetch instruction transfers one 64-bit word of data from the memory in a remote PE to the data prefetch queue. A programmer places the prefetch instruction in a program several instructions before the instruction that actually uses the prefetch data.

Data Translation Circuitry The data translation circuitry converts TTL signals from the PE into ECL signals. The data translation circuitry also converts ECL signals back to TTL signals.

Data Translation Look-aside Buffer (DTB) Annex The DTB Annex extends the addressing capabilities of the microprocessor. The DTB annex is a 32-entry RAM cell that is written by the microprocessor with virtual or logical PE numbers and memory function codes.

Daughter Card A daughter card is a printed circuit board that contains 2 million or 8 million words of dynamic random access memory.

Deadman Heartbeat Event The deadman heartbeat event detects whether a PE is malfunctioning or in a hung condition.

Destination PE Number The destination PE number indicates which node and PE will receive the packet.

Dimension Order Routing Dimension order routing is the process of information traveling through the network in the X dimension first, then through the Y dimension, and finally through the Z dimension. Once the information passes through the Z dimension, the information is at the destination node.

- DRAM Control Register (DRAM_CR)** The DRAM control register is a 13-bit, write-only, system privileged register that is used to configure local memory timing parameters and modes of operation. When configuring local memory, multiple values may have to be written sequentially to the DRAM_CR.
- DTB Annex Index** The DTB annex index is a 5-bit number that references one of the 32 entries in the physical DTB annex.
- Dynamic Random Access Memory (DRAM)** Local memory is composed of DRAM circuits that are mounted on daughter-card printed circuit boards. The DRAM daughter-card plugs into the processing element circuit board. The local memory size per PE is 2 Mwords or 8 Mwords. The total logically shared memory is local memory size multiplied by the total number of PEs in the system.

E

- EA Option** The EA option contains logic that buffers request and response information before it is sent outside the node. The EA option also generates parity for this information and assembles it into packet format.
- EB Option** The EB option contains the logic that controls packet generation. The EB option also contains logic that converts a virtual PE number to a logical PE number.
- EC Option** The EC option contains control logic used to steer the incoming packet to the correct PE within the node. The EC option also checks this incoming packet for errors.
- ED Option** The ED option contains logic that buffers the incoming packets. The ED option also checks the packet for parity errors when it is read out of the buffer.
- Eureka Synchronization** Eureka synchronization uses a point in program instructions where a microprocessor periodically checks to see when the first microprocessor associated with the eureka has finished its part of the program instructions.

F

- Fetch-and-increment Operation** The fetch-and-increment operation reads the contents of one of the fetch-and-increment registers, increments the contents by one, and stores the new value back into the fetch-and increment-register.

Fetch-and-increment Register	The fetch-and-increment register is a hardware register that increments each time information is read from this register. Each processing element node contains two fetch-and-increment registers, which function independently of the node.
Flit	A flit consists of 2 to 8 packet phits. When the flit contains fewer than 8 phits, the last phit of the flit is a tail phit (end of packet). A packet is broken into flits because of the flow control (handshaking protocol). For example, an SR option can send only 8 phits of a packet to the next SR option. The first SR option can not send another 8 phits of that packet until it receives an Acknowledge signal from the second SR option.
Folded Torus Network	A folded torus network consists of nodes that are physically placed so that the maximum wiring distance between the nodes is minimized. This type of network is also referred to as an interleaved network.

G

Global Reset Heartbeat Event	The global reset heartbeat event is used to reset all of the PEs in the system.
-------------------------------------	---

H

Header	The header of a packet contains a routing tag, packet type information, destination PE number, and source return address.
Heartbeat Generation Register (HB_GEN)	The heartbeat generation register is a 2-bit, write-only register that is used only in the input node. The HB_GEN register selects what type of system heartbeat signal (normal, deadman, selective reset, or global reset) will be sent to the PEs in the CRAY T3D system when the next heartbeat signal is distributed.
HISP Buffer Register (HISP_BUFFER)	The HISP buffer registers are used differently in the input node than they are in the output node. In the input node, the HISP_BUFFER registers make up the HISP input buffer. In the output node, the HISP_BUFFER registers make up the HISP output buffer.
HISP Channel	The HISP channel transfers data between the CRAY T3D system and the host system.
HISP Channel Address 0 Register (HISP_ADDR0)	The HISP channel address 0 register is a 64-bit register that is used differently in a master I/O gateway than it is in a slave I/O gateway. In a master I/O gateway, the HISP_ADDR0 register is a write-only register used to transfer the address and block length information to

the host system over the HISP channel. In a slave I/O gateway, the HISP_ADDR0 register is a read-only register that is used to receive the address and block length information from the host system over the HISP channel.

**HISP Channel
Address 1 Register
(HISP_ADDR1)**

The HISP channel address 1 register is a 16-bit register that contains the fifth transfer of address and block length information when the address and block length information transfers over the HISP channel in 200 Mbytes/s mode. When the address and block length information transfers over the HISP channel in 100 Mbytes/s mode, the HISP_ADDR1 register is not used.

**HISP Channel
Address
Configuration
Register
(HISP_A_CFG)**

The HISP channel address configuration register is a write-only register that controls the HISP channel address timing. The HISP_A_CFG register also enables the various areas of SECEDED circuitry in the input node and output node.

**HISP Channel
Lower Block Length
Register
(HISP_L_BL)**

The HISP channel lower block length register is a readable and writable register that contains bits 0 through 15 of the block length for a HISP transfer.

**HISP Channel ROA
Address Register
(HISP_ROA)**

The HISP channel ROA address register is a 15-bit, write-only register that is used to set the internal HISP input or output buffer addresses to a specific value.

**HISP Channel
Upper Block Length
Register
(HISP_U_BL)**

The HISP channel upper block length register is a write-only register that contains bits 16 and 17 of the block length for a HISP transfer.

**HISP Data
Configuration
Register
(HISP_D_CFG)**

The HISP data configuration register is a write only register used to configure the timing on data transfers over the HISP channel.

Hop

A hop is the movement of information from one communication link in the network to the next communication link.

Host System

The host system runs the operating system and software compilers for the CRAY T3D system. All programs written for the CRAY T3D system are compiled on the host system, but run in the CRAY T3D system. The host system may be any Cray Research Inc. computer system that has an input/output subsystem model E. Hosts systems

include the CRAY Y-MP E series computer systems, the CRAY Y-MP M90 series computer systems, and the CRAY C90 series computer systems.

I

I/O Channel Error Register (IO_ERR)	The I/O channel error register is a 16-bit, read-only register that indicates when various error conditions exist for the HISP and LOSP channels.
Index Vector Register (BLT_IVR)	The index vector register is a 24-bit, write-only, system-privileged register that contains the index vector. The index vector is a word-oriented address offset that points to a 64-bit word in local memory where a remote index is stored. The BLT uses the remote index when creating information for the header of a request packet during a gather or scatter operation.
Input Node	The input node receives data from the HISP channel, creates packets, and sends the packets to the PE nodes in the CRAY T3D system using the interconnect network. The input node contains one PE, a network interface, a block transfer engine, and HISP input circuitry.
Input/Output (I/O) Gateways	The I/O gateways transfer system data and control information between the host system and the CRAY T3D system. The CRAY T3D system contains two I/O gateways for every 64 processing element nodes; however, any I/O gateway can transfer data and control to any PE in the interconnect network. An I/O gateway contains an input node, an output node, and low-speed control circuitry.
Input/Output Module (IOM)	An I/O module consists of one coldplate and two printed circuit boards. Each printed circuit board contains one I/O gateway.
Instruction Cache Test Mode Pins	The instruction cache test mode pins are used to signal the microprocessor that code was loaded into the instruction cache and the instruction cache is valid or to inform the microprocessor that the instruction is invalid and instructions must be read from memory.
Interconnect Network	The interconnect network provides communication paths among the PE nodes and the input/output gateways in the CRAY T3D system. The interconnect network forms a three-dimensional matrix of paths that connect the nodes in X, Y, and Z dimensions.
Interleaved Network	An interleaved network consists of nodes that are physically placed so that the maximum wiring distance between the nodes is minimized. This type of network is also referred to as a folded torus network.

Interleaving	Interleaving is the physical placement of nodes so that the maximum wiring distance between the nodes is minimized.
Interrupt Flags Register (IO_IFLAG)	The interrupt flags register is a read-only register that indicates which transfer or error conditions have occurred. The values in the IO_IFLAG register always indicate the current I/O channel conditions and are not affected by the interrupt mask register.
Interrupt Mask Register (IO_IMASK)	The interrupt mask register is a 13-bit, readable and writable register. The IO_IMASK register indicates which IO_IFLAG register bits, when set to 1, cause the I/O hardware interrupt to be set.
Invalidate Control Register (INV_CR)	The invalidate control register is a 10-bit, write only, system privileged register. The INV_CR provides the user with control of cache line invalidates from external writes.

L

LDQ_L Instruction	The LDQ_L instruction is used to read an entry of the DTB annex.
Local Address Register (BLT_LAR)	The local address register is a 24-bit, write-only, system-privileged register that contains the local address. The local address is either the word-oriented address in local memory where data for a BLT write request pack resides or the address in local memory where data from a BLT read response packet will be placed.
Local Memory	Local memory consists of dynamic random access memory (DRAM) that stores system data. Each PE contains local memory.
Local Stride Register (BLT_LSR)	The local stride register is a 24-bit, write-only, system-privileged register that contains the local stride. The local stride is a positive value that the BLT repeatedly adds to the local address to generate successive local addresses during a transfer operation.
Logical Node Number	The logical node number indicates which logical node the PE resides in (processing element node, input node, or output node).
Logical PE Number	The logical PE number is a unique number that is assigned by the operating system to each physical PE used in the logical CRAY T3D system. The logical PE number contains the logical node number and the left or right PE information.
Logical PE Number Register (L_WHOAMI)	The logical PE number is a 12-bit, write-only, system privileged register that contains the logical PE number.

LOSP Channel	The LOSP (low-speed) channel transfers request and response information between the CRAY T3D system and the host system.
LOSP Channel Configuration Register (LOSP_CFIG)	The LOSP channel configuration register is a write-only register that controls timing parameters for the LOSP channel Ready or Resume control signals.
LOSP Channel Data In Register (LOSP_DI)	The LOSP channel data in register is a 64-bit, read-only register. The LOSP_DI register contains a 64-bit word of LOSP input data that the LOSP circuitry assembles from four 16-bit parcels.
LOSP Channel Data Out Register (LOSP_DO)	The LOSP channel data out register is a 64-bit write-only register. The LOSP_DO register receives a 64-bit word of data from the microprocessor in the input or output node and sends the data through the LOSP output buffer to the LOSP output channel in 16-bit parcels.

M

Massively Parallel Processor System (MPP)	A massively parallel processor system is a computer system that contains hundreds or thousands of microprocessors, each accompanied by a small local memory. The microprocessors in the MPP work simultaneously to solve computational problems.
Master I/O Gateway	A master I/O gateway controls the flow of information over the high-speed channel that connects the CRAY T3D system and the CPU in the host system.
Memory Barrier Operation (MB)	The memory barrier operation is used to guarantee that load and store instructions issued after the MB will not access memory until after any load or store instructions issued before the MB have accessed memory.
Memory-mapped Registers	Memory-mapped registers are registers that are physically located in the components of the processing element node; however, the microprocessor addresses the registers as if they were located in local memory.
Message	A message is a special packet that interrupts the microprocessor as soon as the PE receives the message. Messages can be used for data transfers and to indicate errors.
Message Queue	The message queue stores the arriving messages prior to sending them to the microprocessor for processing.

Message Queue Limit Increment Register (MQ_LIR)	The message queue limit increment register is a readable and writable, system privileged register. Software must use the MQ_LIR to manipulate the limit counter.
Message Queue Tail Pointer (MQ_TP)	The message queue tail pointer register is a 12-bit, readable and writable, system privileged register. The MQ_TP is a hardware counter that represents the slot in the message queue where the next message that arrives will be stored. The actual memory offset represented by the tail pointer is calculated by multiplying the tail pointer by 64 and adding the result to the message queue base address (40000_{16}).
Microprocessor	The microprocessor is a reduced instruction set computer (RISC) 64-bit microprocessor developed by Digital Equipment Corporation.
Multiple-cabinet Air-cooled System	In a multiple-cabinet air-cooled system, the CRAY T3D system is housed in one air-cooled cabinet that contains only CRAY T3D modules. The host system is housed in one or more separate cabinets.
Multiple Cabinet System	In a multiple cabinet system, the CRAY T3D system is housed in one or more cabinets that each contain only CRAY T3D modules. The host system is housed in one or more separate cabinets.

N

Network Enable Register (NET_ENA)	The network enable register is a 9-bit, write-only, system privileged register. The NET_ENA register enables or disables each of the inputs to the X-dimension switch, the Y-dimension switch, and the Z-dimension switch in the network interface.
Network Interface	The network interface formats data and control information before it is transmitted over the interconnect network. The network interface also receives and buffers incoming information from the interconnect network.
Network Interface PE Adjust Register (LPE_XLATE)	The network interface PE adjust register is a 31-bit, write-only, system privileged register that contains the shift, base, and configuration mask values used to translate a virtual node number into a logical node number.
Network Interface Source Register (X_WHOAMI)	The network interface source register is an 11-bit, write-only, system privileged register that contains the source logical node number.

Network Mode Register (NODE_CSR)	The network mode register is a 14-bit, write only, system privileged register. The NODE_CSR enables or disables several types of error checking in the network interface and the BLT. The NODE_CSR is also used to partition the physical barrier synchronization circuits.
Network Performance Register (NET_PFM)	The network performance register is an 8-bit, write-only, system privileged register. The NET_PFM controls which parameters of network performance are monitored.
Network Routers	The network routers receive and send packets over the communication links in the interconnect network.
Noncacheable Read Operation	The noncacheable read operation transfers data from system memory to a register in the microprocessor.
No-operation Register (NOR)	The no-operation register is a readable and writable general access register. The NOR may be used by software for debugging purposes. When data is written to or read from the NOR, there is no effect on the hardware.
Normal Heartbeat Event	The normal heartbeat event interrupts the PEs so that the PEs can establish global time of day.

O

Output Node	The output node receives the system data from the PE nodes through the interconnect network and sends the data to the host system over the HISP channel.
--------------------	--

P

Packet Type	The packet type information indicates whether the packet is a request or a response. The packet type information also indicates what type of transfer (read, write, etc.) the packet is requesting or responding to.
Packets	All information is transferred through the network in the form of a packet. A packet contains a header and a body.
Partial Physical Address	A partial physical address is an address that the PE support circuitry uses to determine where to route or retrieve data (local memory, memory-mapped register, or remote memory).
Partitions	A partition is a group of PE nodes that are assigned to one application.

PE	See processing element
PE Support Circuitry	The PE support circuitry extends the control and addressing functions of the microprocessor.
Phit	A phit is 16 bits of information that can be placed on a data channel in 1 clock period.
Physical PE Number	A physical PE number is a unique number assigned to each PE in the system.
Physical PE Number Register (P_WHOAMI)	The physical PE number register is a 12-bit, read-only, system privileged register that contains the physical PE number.
Physically Distributed, Logically Shared Memory	The CRAY T3D system memory is physically distributed because each PE contains local memory that provides a low latency, high bandwidth data path to the microprocessor in the PE. Memory is logically shared because any microprocessor can access the local memory of any PE through the interconnect network without interrupting the remote PE.
Prefetch Queue Head Register (PRE_FETCH)	The prefetch queue head register is a read only, general access register. When the microprocessor reads a word of data from the PRE_FETCH register, the support circuitry transfers 1 word from the head of the data prefetch queue to the microprocessor.
Processing Element (PE)	A processing element contains a microprocessor, local memory, and support circuitry that extends the control and address functions of the microprocessor.
Processing Element Module (PEM)	A processing element module consists of one coldplate and two printed circuit boards. Each printed circuit board contains two processing element nodes.
Processing Element Node	Each processing element node contains two processing elements, a network interface, and a block transfer engine. The processing element nodes perform all of the program instructions and store system data.

R

Read Ahead Operation	A read look ahead operation automatically reads and buffers the next consecutive cache-line size of data from local memory in preparation for the next read operation.
-----------------------------	--

Read Data Acknowledgement	The read data acknowledgement is a 3-bit code that signals the microprocessor whether or not it should place the read data it just received in the data cache and perform SECDED.
Remote Base Register (BLT_RBR)	The remote base register is a 24-bit, write-only, system-privileged register that contains the remote base. The remote base is the starting word-oriented address offset of a remote data structure that will be transferred by the BLT.
Remote Index Register (BLT_RIR)	The remote index register is a 36-bit, write-only, system-privileged register that contains the remote index. The remote index is a remote address that points to an aligned 64-bit word in memory. Byte address bits 0 through 2 are not represented in the remote index.
Remote Limit Register (BLT_RLR)	The remote limit register is a 24-bit, write-only, system-privileged register that contains the remote limit. The remote limit is the maximum remote word-oriented address offset that a user can access.
Remote Mask Register (BLT_RMR)	The remote mask register is a 36-bit, write-only, system-privileged register that contains the remote mask. The remote mask contains 12 bits that are set to 1 to indicate that the corresponding bits of the remote index are used for the PE number. The BLT_RMR also contains 24 bits set to 0 to indicate that the corresponding bits of the remote index are used for the index offset.
Remote Stride Register (BLT_RSR)	The remote stride register is a 36-bit, write-only, system-privileged register that contains the remote stride. The remote stride is a value that the BLT repeatedly adds to the remote index to generate successive remote indexes during a constant stride read or constant stride write operation.
Routing Tag	The routing tag contains three 5-bit fields. In each field, four of the bits contain the number of hops the packet will complete in one dimension. The remaining bit describes the direction the packet will travel in that dimension.
Routing Tag Look-up Table High-order Bits Register (ROUTE_HI)	The routing tag look-up table low-order bits registers are 6-bit, write-only, system privileged registers that each contain the 6 high-order bits of a routing tag.
Routing Tag Look-up Table Low-order Bits Register (ROUTE_LO)	The routing tag look-up table low-order bits registers are 12-bit, write-only, system privileged registers that each contain the 12 low-order bits of a routing tag.

S

Segment	A segment is a 4-bit number that divides the virtual address space into as many as 16 segments.
Selective Reset Heartbeat Event	The selective reset heartbeat event is used to reset selected PEs in the system.
Single Cabinet System	In a single cabinet system, the CRAY T3D system is housed in the same cabinet as the host system.
Slave I/O Gateway	A slave I/O gateway transfers information between the CRAY T3D system and an input/output cluster in the host system. The input/output cluster controls the flow of information over the high-speed channel that connects the CRAY T3D system and the host system.
Source Return Address	The source return address indicates which node originally sent the packet.
SR Option	Three SR options make up the network router in a processing element node. Two SR options make up the network router in an I/O gateway node.
Status Register (BLT_SR)	The status register is a 12-bit, read and write, system-privileged register that a PE may use to check the progress of a transfer and to initiate certain BLT control operations. Unlike the other BLT registers, there are two physical BLT_SR registers, one for each PE in a processing element node. The BLT_SR enables interrupts from the BLT to a PE and contains BLT status information.
STQ_C Instruction	The STQ_C instruction is used to store a new PE number and function code in an entry of the DTB annex.
Swaperand Register (SWAP)	The swaperand register is a 64-bit, write-only, general access register. The swaperand register contains 1 word of data that will be written to a memory location during an atomic swap operation.
Syndrome Register (IO_SYN)	The syndrome register is a 56-bit, read-only register. The IO_SYN register contains the SECDED syndrome bits for errors that were detected on HISP input data and errors that were detected on data received from a microprocessor or a BLT transfer.
System/Clock Module	The system/clock module consists of one coldplate and two printed circuit boards. One printed circuit board contains the clock circuitry and the other printed circuit board contains the system fanout circuitry.

System Control Register (SCR)	The system control register is a 6-bit, write-only, system privileged register that controls the privileged mode resources of a PE.
System Status Register (SSR)	The system status register is a 10-bit, read-only, system privileged register that contains the status of PE error conditions related to the error interrupt and the HARD_ERROR cycle acknowledge. When the SSR is read, the error interrupt and the condition associated with each SSR bit are cleared.

T

Torus Interconnect Topology	A torus interconnect topology contains communication links that connect the smallest numbered node in a dimension directly to the largest numbered node in the same dimension. This type of connection forms a ring where information can transfer from one node, through all of the nodes in the same dimension, and back to the original node.
------------------------------------	--

U

User Control and Status Register (USR)	The user control and status register is a 2-bit, general access, read only register. The USR provides the user with the ability to read status of outstanding remote references.
---	--

V

Vector Length Register (BLT_VLR)	The vector length register is a 16-bit, write-only, system-privileged register that contains the vector length. The vector length indicates how many request packets the BLT generates. When the value of the vector length register is set to 0, the BLT generates the maximum number of packets (65,536).
Virtual Address	The virtual address is a byte oriented address generated by a program compiler. This address may contain a segment, virtual index, and a virtual address offset.
Virtual Address Offset	A virtual address offset is a byte oriented address that points to a byte within a segment of the virtual address space.
Virtual Channels	The virtual channels are 24-bit channels that link the nodes together. There are four virtual channels labeled 0 through 3.
Virtual Index	A virtual index is a 5-bit number that references one of the entries in the DTB annex.

Virtual Node Number	The virtual node number refers to nodes within a partition.
Virtual PE Number	The virtual PE number is a number assigned by the operating system to a logical PE in a partition. The virtual PE number contains the virtual node number and the left or right PE information.
Virtual PE Range Mask	The virtual PE range mask contains a contiguous field of 0's in the legal bit positions for the size of the user partition. The remaining bits are set to 1.
Virtual PE Range Mask Register (VRT_RG)	The virtual PE range mask register is a 10-bit, write-only, system privileged register that contains the virtual PE range mask.
Virtual PE Number Register (V_WHOAMI)	The virtual PE number register is an 11-bit, write-only, system privileged register that contains the virtual PE number.

W

Write Buffer	Each microprocessor contains a write buffer that temporarily stores write data before it transfers to the support circuitry.
Write Operation	The write operation transfers data from the write buffer in the microprocessor to system memory.

INDEX

A

A-series options, differences between, 4-26
Abort (BLT), 8-40
ACK packet, 7-24, 7-25
Acknowledge
 delay, 2-24
 signal, 2-23–2-25
Address
 0 register (HISP_ADDR0), 11-92–11-94
 1 register (HISP_ADDR1), 11-95
 base node, 6-25
 and control, BLT generation of, 8-58
 error
 signal (HISP channel), 11-20, 11-21
 slave HISP input, 11-38
 fetch-and-increment register, 7-17, 7-50
 interpretation
 logical PE, 6-21
 microprocessor, 6-10
 network interface, 6-22–6-36, 6-56–6-64
 PE support circuitry, 6-19–6-21, 6-56
 virtual PE, 6-19
 local (BLT), 8-14
 microprocessor write buffer, 3-7
 offset, 6-18, 6-55
 decode of, 8-53
 generation, 8-11
 range error (BLT), 10-36
 virtual, 6-12
 partial physical, 6-12–6-18, 6-54–6-55
 phit layout, 3-43
 remote (BLT), 8-3–8-14
 space
 partial physical, 6-18
 partition, 6-16–6-18, 6-54–6-55
 tag, 3-3
 virtual, 6-10–6-12, 6-19
Address and block length (HISP)
 receiving, 11-34–11-36, 11-51–11-52
 sending, 11-27, 11-44–11-45

 signal, HISP channel, 11-20, 11-21
AE option, 3-30–3-33
AJ option, 3-34–3-38
AK option, 3-34–3-38
AM option. *See* Option: AM
AR option. *See* Option: AR
Arbitration, 3-41. *See also* Conflict(s); Local memory
ASP (address space partition), 6-16–6-18, 6-54–6-55
Atomic swap
 cached, 7-14–7-15, 7-79
 noncacheable, 7-9–7-12, 7-74–7-77

B

BAR_TMG register, 9-16, 9-23–9-25, 10-29–10-30
Barrier. *See also* Memory: barrier
 bits, fanin/fanout, 9-49–9-53.
 channel routing tag bit layout, 5-8
 function register, 9-35
 interrupt, 3-20, 10-2, 10-4, 10-12, 10-42
 enable bit, 10-39
 mask register, 9-35
Barrier synchronization, 9-1
 circuit(s)
 bit propagation timing, 9-16
 logical, 9-2–9-8
 partitioned, 9-11
 physical, 9-9–9-16, 9-34–9-47
 simplified, 9-3
 timing procedure values, 9-25
 mask and interrupt register, 9-7
 registers, 9-16–9-34
Barrier timing and refresh register (BAR_TMG), 9-17, 9-23–9-25, 10-29–10-30
Base node address, 6-25
BC option, 9-37–9-47

Block length, 11-27, 11-65, 11-68
 Block transfer engine. *See* BLT
 BLT. *See also* Option: AR; Option: EE
 hardware, 8-48–8-73
 input node, 4-6, 4-27
 interrupt, 3-20, 10-2, 10-4, 10-42
 complete, 8-46–8-47
 output node, 4-9, 4-38
 overview, 1-4, 1-6
 parameters (input)
 master, sample, 11-25
 slave, 11-36
 parameters (output)
 master, sample, 11-43
 slave, 11-52
 register write ordering, 8-47
 BLT_CR register, 8-30–8-35
 BLT_IVR register, 8-23–8-24
 BLT_LAR register, 8-28
 BLT_LSR register, 8-29
 BLT_RBR register, 8-26
 BLT_RIR register, 8-18–8-20
 BLT_RLR, register, 8-27
 BLT_RMR register, 8-24–8-25
 BLT_RSR register, 8-21
 BLT_SR register, 8-36–8-45
 BLT_VLR register, 8-22
 status, 8-41–8-42
 Body. *See* Packet
 Boot code segments, 5-2–5-7
 BSFR register, 9-6, 9-17, 9-19–9-20
 BSMI register, 9-7–9-8, 9-17, 9-21
 BSR0, 9-17, 9-18. *See also* Option: AR
 BSR1, 9-17, 9-19. *See also* Option: AR
 Buffer(s). *See also* Option: EA; Write: buffer
 base address register (HISP_BUFFER),
 11-96
 and input circuitry, 4-8
 and output circuitry, 4-10, 4-11
 Buffering, data. *See* Option: AJ; Option: AK
 Bypass point(s), 9-10, 9-14–9-15. *See also*
 Barrier synchronization: partition bits

C

Cabinet (illustrated), 1-13, 1-14
 Cable configuration, 4-2
 Cache coherence, 7-59
 Cache line. *See also* Data cache
 invalidate filter, 3-6
 invalidate pins, 10-8
 invalidation, 3-5–3-6
 Cache memory. *See* Data cache; Instruction
 cache
 Cached
 atomic swap, 7-14–7-15, 7-79
 read ahead operation, 7-13–7-14, 7-78–7-79
 read operation, 7-12–7-15, 7-77–7-79
 CAS, 10-24
 CAS cycle timing bit values, 10-27. *See also*
 Option(s): AE
 Centrifuge, 8-6–8-10
 Channel. *See also* Barrier; Communication
 link; Data channel; LOSP channel; HISP:
 channel; Register(s); Virtual channel
 command (AM to EA/EB option), 7-64
 control circuitry, 4-8, 4-10, 4-11
 data signals, 2-3
 phit type bits, 2-3, 2-4
 CHASSIS SELECT switches, 13-6, 13-14
 Chassis SYSTEM SIZE switches, 13-5, 13-13
 Check bits (0 through 31), 3-11
 Clear
 Channel signal (HISP channel), 4-20, 11-21
 heartbeat interrupt bit, 10-10, 10-11
 HISP syndrome register
 (HISP_SYN_CLR), 11-101
 I/O register (IO_CLR), 11-71–11-76
 microprocessor syndrome register
 (DEC_SYN_CLR), 11-100
 Clock
 central, 1-7
 module (system/clock), 1-12
 printed circuit board, 1-12
 CLOCK SELECT switches, 13-4, 13-12,
 13-20, 13-24

- Code(s)
 - cycle
 - acknowledge, 10-8
 - request, 7-1-7-5, 7-61, 10-7
 - deadstart, 5-3, 5-4
 - function, 6-15, 7-1-7-5, 7-61
 - nondeadstart, 5-3, 5-5-5-7
 - SADPAL, 5-1, 5-7-5-17
 - Command
 - from AM option, 7-90
 - channel (AM to EA/EB options), 7-64
 - field(s), 6-34
 - message, ACK and NACK, 7-24
 - packet, 7-38-7-39
 - phit, 3-42, 6-33, 6-34
 - Communication link(s). *See also* Torus
 - avoiding bad, 2-6, 2-11
 - channel signals, 2-3-2-4
 - dateline, 2-12-2-13, 6-48
 - illustrated, 2-2
 - interconnect network, 2-2-2-5
 - Complete interrupt, BLT, 8-46-8-47
 - Configuration
 - cabinet numbers (multiple), 6-41
 - cable, 4-2
 - mask values, 6-47
 - masks, 6-26, 6-27
 - message described, 5-6-5-7
 - multiple-cabinet, 1-14
 - setup by host system, 5-1
 - single-cabinet, 1-13
 - Conflict(s). *See also* Option: AM
 - arbitration, 3-41
 - monitoring of, 2-26
 - and packet transfer, 2-26-2-28
 - simultaneous
 - input, 2-26, 2-27
 - reference request, 2-27, 2-28
 - virtual channel
 - busy, 2-28
 - reserved, 2-28
 - Constant stride, 8-1
 - CONTROL DUMP switch, 13-28
 - Control panel
 - multiple cabinet
 - locations, 13-2, 13-10
 - remote access, 13-9, 13-17, 13-28
 - Control register (BLT_CR), 8-30-8-35
 - Control and status
 - information, 8-16
 - registers, 8-2
 - Conversion
 - logical PE number to routing tag, 6-30
 - virtual address to partial physical address, 6-10
 - virtual PE number to logical PE number, 6-23, 6-57-6-58
 - Counter
 - EB option, 3-56
 - message queue limit, 7-28
 - outstanding
 - request, 8-13-8-14
 - write request, 7-17
 - CPU
 - MCU SELECT switches, 13-27
 - real-time interrupts, 13-29
 - CPU data path control, 3-37
 - CPU IDLE switches, 13-23
 - Cycle
 - acknowledge code, 3-18, 10-8
 - function codes, 7-1-7-5, 7-61
 - request code, 3-18, 7-1-7-5, 7-61, 10-7
- ## D
- Data
 - packet, 6-37
 - buffering, 3-34-3-38
 - bus interface, 3-34-3-38
 - cache
 - coherence problems, 7-59
 - memory, 3-3-3-6
 - miss, 7-62
 - channel, AJ/AK to EA/EB options, 7-83
 - configuration register (HISP_D_CFGI), 11-87-11-89
 - muxing, 3-34-3-38
 - path control. *See* Path, data
 - phit, 6-37
 - steering, 7-73

- Dateline
 - communication link, 2-12–2-13
 - crossing, 6-48
 - Daughter card(s). *See also* Rank
 - address and control distribution, 7-66
 - data distribution, 3-25–3-27, 3-35–3-36
 - overview, 1-10
 - on PE module, 3-22
 - storage on, 3-23
 - Deadlock, 2-12
 - Deadman
 - heartbeat, 10-3
 - timer (DMT), 10-3
 - Deadstart. *See also* Initialization
 - entry point (SADPAL), 5-7–5-11
 - node, 5-2–5-4, 5-7–5-17
 - sequence, 5-1–5-7
 - DEC_SYN_CLR register, 11-100
 - Destination phit, 3-42, 6-34
 - Diagnostic code, 5-7, 5-10–5-12
 - Diagnostic monitor communication area (DMCA), memory locations, 5-9
 - Dimension order routing, 2-8–2-11
 - Dimensions. *See* Torus
 - Disable SECDDED signal, 11-20, 11-21
 - Disconnect signal, 11-6–11-7
 - DMCA memory locations, 5-9
 - DRAM. *See also* Daughter card; Local memory
 - control
 - command (from AM option), 3-31
 - register (DRAM_CR), 10-21–10-28, 10-47–10-51
 - reading from, 8-66
 - timing. *See* Option: AE
 - write data path control, 3-35
 - DRAM control and response packet header
 - phits, generation of, 8-64
 - DRAM_CR register, 10-21–10-28, 10-47–10-51
 - DT, 3-22. *See also* Daughter card
 - DTB annex. *See also* Option: AR
 - entry, 6-14
 - and partial physical address, 7-3
 - sample assignments, 6-16
 - error, 10-14, 10-17, 10-45
 - index, 6-14–6-16, 6-54
 - read/write, 7-5, 7-6, 7-61–7-62
- ## E
- E-series options, 12-4
 - EA option. *See* Option: EA
 - EB option. *See* Option: EB
 - EC option, 3-57–3-59
 - ECC generation, 4-15, 4-37
 - ECL-to-TTL translation. *See* Option: IA
 - ED option. *See* Option: ED
 - EE option. *See* Option: EE
 - Entry point. *See* Deadstart; Initialization
 - Error(s). *See also* Interrupt: flag register;
 - Packet: error
 - BLT, 8-16, 8-38–8-40, 11-49
 - checking. *See also* HISP: error; LOSP: errors; Network mode register; Option: EC; Option: IA
 - by BLT, 10-36
 - conditions, prefetch queue, 7-44–7-45
 - correction. *See also* Option: IE; SECDDED MAINTENANCE MODE switch, 13-27
 - DTB annex, 10-14, 10-17, 10-45
 - header phit, checking, 7-26
 - illegal register access, 10-16
 - illegal write to PLM, 10-17
 - interrupt, 3-20, 10-2, 10-5, 10-43. *See also* Option: AR
 - LOSP
 - input transfer, 11-8–11-10
 - output transfer, 11-16–11-19
 - messages (queue), 7-34–7-39, 10-15
 - range, 10-36, 10-45
 - ROA parity, 11-29, 11-32
 - SECDDED
 - HISP, 11-30–11-31
 - Microprocessor, 11-18–11-19
 - status (BLT), 8-39

Eureka synchronization, 9-1, 9-6. *See also*
 Barrier synchronization
 External outgoing path control, 3-37–3-38

F

Fanin/fanout. *See also* Option: BC
 of barrier bits, 9-49–9-53
 Fetch-and-increment
 operation, 7-17–7-21, 7-84–7-90
 register, 7-50
 FETCH instruction, 7-45–7-46, 7-93
 FETCH_M instruction, 7-45, 7-46–7-47, 7-93
 Field. *See* Command: field(s)
 FIFO. *See* Prefetch: queue
 Filtered cache line invalidation, 3-6
 Flit transfer, 2-23–2-25
 Function codes, 6-15, 7-61
 and cycle requests, 7-1–7-5, 10-7

G

Global Reset Heartbeat signal (deadstart), 5-1

H

Halfword, 7-16
 HARD_ERROR. *See* SSR
 Hardware interrupt. *See* Interrupt; Option:
 AR
 Hardware partition, 6-5–6-7
 HB_GEN register, 11-102
 Header, packet. *See* Packet: header
 Header phit. *See* Phit(s)
 Heartbeat. *See also* Option: AR; System:
 heartbeat signal
 generation register (HB_GEN), 11-102
 interrupt, 3-20, 10-2–10-4, 10-11, 10-42
 interval values, 13-15, 13-21
 request intervals, 13-7, 13-15, 13-21
 HEARTBEAT ENABLE switch, 13-8, 13-16,
 13-22

HEARTBEAT INTERVAL switch, 13-7,
 13-15, 13-21
 HEARTBEAT SELECT switch, 13-8, 13-16,
 13-22

HISP

channel. *See also* Initialization; Master:
 HISP; Register; Slave: HISP
 clearing, 11-33
 data control, 1-7
 master/slave components, 1-7
 signals, 4-3, 11-20, 11-21
 data
 check-bit generation, 4-10, 4-11
 receiving (slave I/O gateway), 11-36
 sending, 11-45–11-47, 11-52
 errors, 11-29–11-33, 11-37–11-39,
 11-47–11-49, 11-53–11-54,
 11-57–11-62
 input channel
 clearing, 11-39
 control, 4-16–4-20
 input circuitry, 1-7
 input data
 buffer, 4-21–4-22
 receiving, 11-27–11-29
 input registers, 11-22–11-23
 output
 channel control, 4-30–4-34
 circuitry, 1-7
 clearing, 11-49, 11-54
 data buffer, 4-35–4-36
 ECC generation, 4-37
 HISP_A_CFIG register, 11-81–11-86
 HISP_ADDR0 register, 11-92–11-94
 HISP_ADDR1 register, 11-95
 HISP_BUFFER registers, 11-96
 HISP_D_CFIG register, 11-87–11-89
 HISP_L_BL register, 11-90
 HISP_ROA address register, 11-97
 HISP_SYN_CLR register, 11-101
 HISP_U_BL register, 11-91
 Hop, defined, 2-8

- I**
- IA option, 4-15–4-16, 4-29–4-30
 - IC option, 4-30–4-34
 - ID option, 4-16–4-20
 - IE option, 4-23–4-24
 - IF option, 4-37
 - IG option, 4-21–4-22
 - IH option, 4-35–4-36
 - IJ option, 4-24
 - Illegal
 - register access, 10-14, 10-16, 10-45
 - write to PLM, 10-14, 10-17, 10-46
 - IM option, 4-39–4-42
 - Index. *See also* Centrifuge
 - data cache line, 3-4
 - DTB annex, 6-14–6-16, 6-54
 - memory error, 8-16
 - memory-mapped registers, 12-1–12-2
 - offset, conversion to address offset, 8-11
 - vector register (BLT_IVR), 8-23–8-24
 - virtual, 6-12
 - Initialization. *See also* Deadstart
 - entry point (SADPAL), 5-11–5-17
 - LOSP
 - input channel, 11-4
 - output channel, 11-13
 - master
 - input channel, 11-23
 - output channel, 11-41
 - slave
 - input channel, 11-34
 - output transfer, 11-50
 - system, 5-12
 - Inpage cycle, 3-31
 - Input. *See* LOSP input
 - Input channel. *See* LOSP; Master: HISP:
 - input channel
 - Input circuitry, 4-7–4-8, 4-14–4-24
 - Input node, 1-6, 4-4–4-8
 - hardware, 4-14–4-27
 - and local memory, 1-7
 - network router components, 2-18
 - processing element, 4-4–4-5
 - Instruction
 - FETCH, 7-45–7-46, 7-93
 - FETCH_M, 7-45, 7-46–7-47, 7-93
 - load, 3-4
 - MB, 7-6
 - store, 3-4
 - Instruction cache
 - memory, 10-6
 - test mode pins, 10-6
 - Instruction buffers CONTROL DUMP
 - switch, 13-28
 - Interconnect network. 2-1–2-28
 - components, 1-5, 2-1
 - in input node, 4-27
 - model of, 1-2
 - overview, 1-5
 - Interconnect topology (torus), 2-5–2-7
 - Interface. *See* Network: interface
 - Interface, data bus, 3-34–3-38
 - Interleaving of nodes, 2-7–2-8
 - Interrupt. *See also* Signal
 - barrier, 10-2, 10-4, 10-12
 - BLT, 10-2, 10-4
 - BLT complete, 8-46–8-47
 - CPU real-time, 13-29
 - enable BLT, 8-37
 - error, 10-2, 10-5
 - flag register (IO_IFLAG), 11-56–11-62
 - hardware, 3-20, 10-42–10-43
 - heartbeat, 10-2–10-4
 - I/O, 10-2, 10-5
 - mask register (IO_IMASK), 11-63
 - message, 10-2, 10-4
 - status, 8-38
 - Interrupt-driven monitoring
 - HISP
 - master input channel, 11-24
 - master output channel, 11-42
 - slave input channel, 11-34
 - output transfer, 11-50
 - LOSP input transfer, 11-4
 - INV_CR register, 10-19–10-21, 10-46–10-47
 - Invalidate
 - control register, 10-19–10-21, 10-46–10-47
 - filter, 3-6
 - pins, cache line, 3-5

I/O channel error register (IO_ERR),
 11-64–11-70

I/O gateway. *See also* Node: I/O gateway
 communication links, 2-1
 components, 1-6
 functional description, 4-1–4-13
 hardware description, 4-13–4-42
 I/O channel error, 11-64–11-70
 LOSEP ENABLE switch, 13-3, 13-11, 13-19
 master, 1-7, 4-14–4-24
 model of, 1-2
 network router switch, 2-17
 overview, 1-2, 1-6–1-7
 slave, 1-7, 11-69–11-70
 spare PE node locations, 6-52

I/O interrupt, 3-20, 10-2, 10-5, 10-42. *See also* Option: AR

I/O memory-mapped registers. *See also*
 Register
 illustrated, 12-9

I/O mode, enable, 8-34

I/O module, 1-11
 illustrated, 12-21

I/O registers
 master input channel, 11-26
 master output channel, 11-44
 slave input channel, 11-36
 slave output transfer, 11-52

IO_CLR register, 11-71–11-76

IO_ERR register, 11-64–11-70

IO_IFLAG register, 11-56–11-62

IO_IMASK register, 11-63

IO_SYN register, 11-98–11-100

IOM. *See* I/O module

IT option, 4-24

L

L_WHOAMI register, 6-21, 6-42, 6-65

Last Word signal, 11-20, 11-21

LDQ_L instruction, 7-6

Lines. *See* Data cache

Links, 2-2

Local address
 generation, 8-14–8-15
 register (BLT_LAR), 8-28

Local addressing register, 8-2

Local memory. *See also* Daughter card
 arbitration, 3-41
 data distribution in, 1-4
 distribution of, 1-3
 PE, 3-12
 hardware, 3-22–3-28
 input node (PE), 4-5, 4-6, 4-25
 modes, configuring, 10-47–10-51
 noncacheable atomic swap, read from, 7-74
 overview (PE), 1-3
 in PE mode, 1-7
 reference control. *See* Option: AM
 sharing of (PE), 3-12
 size of (PE), 3-12
 timing parameters, configuring,
 10-47–10-51

Local stride register (BLT_LSR), 8-29

Logic, PE support circuitry, 1-9

Logical barrier synchronization circuits,
 9-2–9-8

Logical partitions, 9-7–9-8

Logical PE
 address interpretation, 6-21, 6-56
 number
 BLT interpretation of, 8-34
 converting to routing tag, 6-29–6-30, 6-59
 converting from virtual PE number,
 6-57–6-58
 register, 6-21, 6-42, 6-65
 numbering, 6-2–6-4

Logical processor, selection of, 13-26

Logically shared memory, 1-3

Look-up tables. *See* Routing tag: look-up
 table

LOSP

- channel, 4-2, 11-2–11-19. *See also*
 - Initialization; Output circuitry
- configuration register (LOSP_CFG), 11-77–11-78
- data in register (LOSP_DI), 11-79
- data out register (LOSP_DO), 11-80
- output channel disconnect register (LOSP_DISC), 11-80
- overview, 11-2–11-3
- signals, 4-2, 11-2
- circuitry, 4-12–4-13
 - IM option, 4-39–4-42
 - transfer overview, 1-7
- data (sending), 11-13–11-15
- ENABLE switch, 13-3, 13-11, 13-19
- errors, 11-8–11-10, 11-16–11-19, 11-61
- input
 - channel, 11-3–11-11
 - data available, 11-60
 - disconnect, 11-60
 - output channel registers, 11-11–11-19
- LOSP_CFG register, 11-77–11-78
- LOSP_DI register, 11-79
- LOSP_DISC register, 11-80
- LOSP_DO register, 11-80
- LPE_XLATE register, 6-46–6-47, 6-65

M

MAINTENANCE MODE switch, 13-27

Mask(s)

- configuration, 6-26, 6-27
- microprocessor write buffer, 3-7
- values, configuration, 6-47
- X, 6-25
- Y, 6-25
- Z, 6-25

Master

clear, 5-1, 11-10

HISP

- input channel, 11-22–11-33
- output channel, 11-40–11-49
- output errors, 11-47–11-49

I/O

- gateway. *See* I/O gateway
- module, illustrated, 12-13
- printed circuit board, 1-11
- Master Clear signal (deadstart), 5-1
- MASTER CPU switches, 13-26
- MB instruction, 7-6
- MCU SELECT switches, 13-27
- Memory. *See also* Data cache: memory; Instruction cache; Local memory; Prefetch: queue; Remote: address generation
 - barrier
 - instruction, 7-61
 - operation, 7-5, 7-6
 - chips, data distribution to, 3-25–3-27
 - conflicts (AM option), 3-41
 - distribution overview, 1-3
 - function codes, 3-48, 6-15
 - logically shared, 1-3
 - physically distributed, 1-3
 - read, prefetch, 7-42–7-43
 - reading from/sending to microprocessor, 7-67
 - refresh, 3-28
 - remote, reading repeatedly, 7-60
 - SECTION SELECT switches, 13-25
- Memory-mapped registers, index of, 12-1–12-2. *See also* Register
- Message. *See also* Error: messages; Option:
 - AR
 - interrupt, 3-20, 10-2, 10-4
 - packet
 - command fields, 7-24
 - format (message queue), 7-26
 - queue, 7-26–7-33, 7-90
 - fully subscribed condition, 7-39
 - limit increment register (MQ_LIR), 7-52
 - oversubscribed, 10-13, 10-15, 10-44
 - pointers, 7-90
 - registers, 7-28
 - tail pointer register (MQ_TP), 7-51
 - reading, 7-33, 7-93
 - receiving (queue), 7-30–7-32
 - write address information, 7-22
 - writing, 7-22–7-25, 7-91–7-92

Messaging facility, 7-22–7-39, 7-90–7-93

Microprocessor

- address interpretation, 6-10–6-18
- boot code segments, 5-2–5-7
- components, 3-2
- functional description, 3-2–3-12
- hardware, 3-15–3-21
- input/output signals, 3-17–3-21
- overview, 1-3
- package dimension, 3-15
- packet header data paths, 7-57
- printed circuit board connection, 3-16
- read NACK message, 7-33
- SECEDED errors, 11-48

Mode, outstanding request counter, 8-14

Module

- I/O, 1-11
 - illustrated, 12-21
- master I/O (illustrated), 12-13
- processing element, 1-8–1-10
 - illustrated, 12-11, 12-19
 - spare, 1-10
- processing element, with daughter cards, 3-22
 - illustrated, 12-17
- slave I/O (illustrated), 12-15
- system/clock, 1-12
- types, 1-8

Monitoring

- HISP, 11-24, 11-34, 11-42, 11-50
- LOSP, 11-4

mppms, 5-1, 5-4, 5-7–5-11

MQ_LIR register, 7-52

MQ_TP register, 7-51

Multiple-bit errors, 4-23–4-24

Multiple-cabinet configuration, 1-14

Muxing, data, 3-34–3-38

N

NACK packet, 7-24, 7-25, 7-32

NEND, 10-26

NET_ENA register, 10-31, 10-52

NET_PFM register, 10-32–10-34,
10-53–10-54

Network. *See also* Interconnect network

- buffer parity error, 10-13–10-14
- enable register (NET_ENA), 10-31, 10-52
- interface, 1-4, 3-50–3-61. *See also* Option

EB; Option: EC; Option: ED

address interpretation, 6-22–6-37

connection, 6-22

input node, 4-5, 4-27

output node, 4-9, 4-38

packet creation, 6-31–6-37

packet header data paths, 7-57–7-58

PE adjust register (LPE_XLATE), 6-46

source register (X_WHOAMI), 6-45

switch inputs, enabling/disabling, 10-52

mode register (NODE_CSR), 9-25–9-33,

10-34–10-37, 10-54–10-56

packet error, 10-13, 10-14, 10-35

performance register (NET_PFM),

10-32–10-34, 10-53–10-54

Reset signal (deadstart), 5-1–5-2

router

input node, 4-5, 4-27

output node, 4-9, 4-38

overview, 1-5

processing element node, 2-19–2-28

Network Reset signal (deadstart), 5-1–5-2

Newpage cycle, 3-33

Node. *See also* Input node; Output node

deadstart, 5-2

failing, replacement, 6-2–6-3

interleaving, 2-7–2-8

I/O gateway, network router options, 2-20

nondeadstart, 5-11

number

logical, 6-2, 6-9

physical, 6-2

virtual, 6-5, 6-8, 6-9

processing element, 1-1, 1-3–1-4, 3-1–3-61

components, 1-4

connection to interconnect network, 1-5

model of, 1-2

and network router (illustrated), 3-1

network router options, 2-19

Node enable vote, 10-10, 10-12

NODE_CSR register, 9-25–9-33,

10-34–10-37, 10-54–10-56

- Noncacheable
 - atomic swap read, 7-9–7-12, 7-74–7-77
 - read operation, 7-8–7-12
 - normal, 7-62–7-73
 - Nondeadstart code, 5-3, 5-5–5-7
 - Nonfiltered cache line invalidation, 3-5–3-6
 - No-operation register (NOR), 10-37
 - Normal
 - cached read, 7-13, 7-77
 - noncacheable read operation, 7-8–7-9
 - 7-62–7-77
- O**
- Operating system partition, 6-5
 - Option(s). *See also* Signal
 - A-series, 4-26, 12-3
 - AE, 3-30–3-33
 - AJ, 3-34–3-38
 - AK, 3-34–3-38
 - AM
 - command from, 7-96
 - and control and status, 10-43–10-51
 - and DTB annex, 7-6
 - illustrated, 7-64
 - PE control, 3-39–3-44
 - AR
 - and barrier synchronization, 9-34–9-35
 - and control and status, 10-38–10-43
 - processing element node, 3-45–3-49
 - BC, 9-37–9-47
 - E-series, 12-4
 - EA
 - circuitry determining BLT memory-map reference, 8-52
 - network interface output buffer, 3-51–3-53
 - packet assembly, 3-51–3-53
 - parity generation, 3-51–3-53
 - remote read, 8-59
 - EB
 - BLT memory-map reference control, 8-55
 - network interface output buffer control, 3-54–3-56
 - and packet control, 3-53
 - remote read, 8-59
 - signals to EA, 3-52
 - EC, 3-57–3-59
 - ED
 - and barrier synchronization, 9-36–9-37
 - network interface input buffer, 3-60–3-61
 - performance monitor circuitry, 10-54
 - EE
 - BLT hardware, 8-48–8-49
 - and control and status, 10-52
 - initializing, 8-49–8-56
 - I-series, 12-5
 - IA
 - ECL-to-TTL translation and HISP ECC generation, 4-15–4-16
 - TTL-to-ECL translation and HISP error check, 4-29, 4-30
 - IC, 4-30–4-34
 - ID, 4-16–4-20
 - IE, 4-23–4-24
 - IF, 4-37
 - IG, 4-21–4-22
 - IH, 4-35–4-36
 - IJ, 4-24
 - IM, 4-39–4-42
 - IT, 4-24
 - network
 - interface (illustrated), 3-50
 - router, 2-19–2-28
 - output, master I/O gateway (illustrated), 4-28
 - PE support circuitry (illustrated), 3-29
 - SR, 2-19–2-28, 10-52, 10-53
 - Output channel disconnect register (LOSP_DISC), 11-80
 - Output circuitry, 4-10–4-11
 - differences, 4-38
 - master, 4-28–4-38
 - slave, 4-38
 - Output node, 1-7, 4-9–4-11, 4-28–4-38
 - Outstanding request counter, 8-13–8-14, 8-34

P

- P_WHOAMI register, 5-3, 5-5, 6-39–6-40, 6-65, 13-6
 - values, 13-3
- Package dimension, (microprocessor), 3-15
- Packet, 2-3, 2-14. *See also* Message;
 - Option(s): EA; Option(s): EC
 - assembly, 3-53, 3-55–3-56
 - body size control, 8-31–8-34
 - command fields, 7-38–7-39
 - creation, 6-31–6-37, 6-59–6-64
 - error, 10-44
 - message for (message queue), 7-34–7-39
 - misrouted, 7-34
 - parity, 10-14–10-15
 - format(s)
 - ACK, 7-24–7-25
 - data, 6-37
 - error, misrouted, 7-34, 10-14, 10-15
 - generic, 2-14
 - message (in queue), 7-26
 - messaging facility, 7-24
 - NACK, 7-24, 7-25, 7-32
 - response, 7-58
 - header, 6-33–6-36
 - data paths, 7-56–7-58
 - request, input, 7-70
 - transfer of, 2-15
 - types (per function), 6-31, 6-32
- Parameter. *See* BLT
- Parameter select register (NET_PFM), 10-32–10-34
- Parity
 - checking, 3-57–3-59
 - error(s). *See also* LOSP: circuitry
 - HISP ROA, 11-32, 11-48
 - LOSP output ROA, 11-9, 11-17
 - network buffer, 10-44
 - values, control of, 10-36
 - generation, 3-52, 3-55–3-56. *See also*
 - Option(s): EA
- Partial physical address, 6-12–6-18, 6-54–6-55, 7-3
- Partition(s), 6-5–6-7
 - A, B, C translation parameter, 6-26
 - address space, 6-54–6-55
 - bits, 6-16–6-18
 - barrier synchronization, 9-26–9-31, 9-32, 9-33
 - hardware, 6-5
 - logical, 9-7–9-8
 - operating system, 6-5
 - physical, 9-9–9-16
- Partitioned barrier synchronization circuit, 9-11
- Path, data
 - control, 3-35, 3-37–3-38
 - packet header, 7-56–7-58
 - through PE, 7-54–7-60
- PCB. *See* Printed circuit board
- PE, 1-3–1-4
 - adjust parameter bits, 6-57
 - circuit board (side view), 3-22
 - components, 1-3, 3-2
 - control. *See* Option(s): AM
 - hardware (nodes 0 and 1), 3-14
 - in input node, 1-6, 4-4–4-5, 4-25
 - circuitry, 4-25–4-27
 - in I/O gateway, 1-6–1-7
 - local, 6-19
 - module
 - with daughter card, 3-22, 12-17
 - illustrated, 12-11, 12-19
 - spare, 1-10
 - network interface options, 3-50
 - node. *See also* Node: processing element
 - bypass point, 9-14
 - circuitry, output node, 4-38
 - communication links, 2-1
 - mask, 8-13
 - registers (illustrated), 12-7
 - spare, 6-2, 6-52
 - numbering, 6-1–6-9
 - logical, 3-47, 6-2–6-4
 - physical, 6-1–6-2
 - virtual, 3-47, 6-5–6-9
 - in output node, 4-9
 - in PE node, 1-6
 - printed circuit board, 1-8, 1-10

- PE (continued)
 - range
 - check, 8-12–8-13
 - error, 8-12
 - mask, 8-35
 - remote, 6-20
 - Reset signal (deadstart), 5-2
 - spare, 6-2, 6-52
 - support circuitry, 1-9, 3-13, 3-29–3-49
 - address interpretation, 6-19–6-21, 6-56
 - PEM. *See* PE: module
 - Performance monitor circuitry
 - ED option, 10-54
 - SR option, 10-52, 10-53
 - Phit(s), 2-3–2-4
 - address, layout, 3-43
 - command, 3-42, 6-33, 6-34
 - control of, 2-3, 2-4
 - data, 6-37
 - destination, 3-42, 6-33, 6-34
 - header, error checking, 7-26
 - in packet header, 6-33–6-37
 - request address, 6-33, 6-35
 - response
 - address, 6-33, 6-36
 - header, generation of, 7-71
 - routing tag, 6-33
 - source, 6-33, 6-35
 - transfer of, 2-3
 - transfer in flits, 2-23–2-25
 - Physical
 - barrier synchronization circuitry, 9-9–9-16, 9-34–9-47
 - node number, 6-2
 - partitions, 9-9–9-16
 - PE number register (P_WHOAMI), 5-3, 5-5, 6-39–6-41, 6-65
 - PE numbering, 6-1–6-2
 - Physically distributed memory, 1-3
 - Pins
 - bi-directional data, 3-24
 - cache line invalidate, 3-5, 10-8
 - cycle acknowledge, 10-8
 - cycle request, 10-7
 - hardware interrupt, 10-2–10-5
 - instruction cache test mode, 10-6
 - PLM, 10-17
 - Pointer
 - message queue, 7-90
 - microprocessor write buffer, 3-7
 - prefetch queue, 7-41
 - Polled monitoring, 11-4
 - Polling
 - master input channel, 11-24
 - master output channel, 11-42
 - slave input channel, 11-34
 - slave output transfer, 11-50
 - Prefetch, 7-40–7-47
 - operation, 7-42–7-43, 7-94–7-96
 - queue, 3-48–3-49, 7-40–7-42, 7-94
 - head register (PRE_FETCH), 7-53
 - overflow, 10-13, 10-15, 10-44
 - reading from, 7-44, 7-97
 - underrun, 10-13, 10-15, 10-45
 - request, 7-42–7-43
 - outstanding, 10-18, 10-19
 - Printed circuit board
 - clock, 1-12
 - master I/O, 1-11
 - microprocessor connection to, 3-16
 - PE, 1-8, 1-10
 - system, 1-12
 - illustrated, 12-23
 - Priority. *See* Conflict(s)
 - Priority counter, disabling, 2-28
 - Processing element. *See* PE
 - Programming restrictions (for support circuitry operations), 7-60
 - Protected low memory (PLM), 10-17
- ## R
- RAM. *See* Prefetch: queue
 - RAM-on-array. *See* ROA buffer
 - Range
 - check, PE, 8-12–8-13
 - error(s)
 - address offset, 10-36
 - virtual PE, 10-36, 10-45
 - mask, PE, 8-35
 - Rank, 3-25–3-28
 - RAS, 3-30–3-33

- Read. *See* Cached: read operation;
 Noncacheable: read operation; Remote:
 read operation
 data. *See also* Path, data
 acknowledge bits, 3-19
 acknowledgment, 6-55
 microprocessor (from message queue), 7-33
 operations, 7-5, 7-7-7-15
- READ_BLOCK. *See* Read: operations
- Reading remote memory repeatedly, 7-60
- Real-time interrupts, 13-28
- Reduced instruction set computer. *See*
 Microprocessor
- Redundant (spare) PE, 6-2, 6-52
- Refresh, memory, 3-28
- Register(s). *See also* Index: memory-mapped
 registers
 access, illegal, 10-45
 addressing for, 6-37-6-54
 AM option, 3-44
 AR option, 3-49
 assignments, I/O node, 11-50
 BAR_TMG, 9-16, 9-23-9-25, 10-29-10-30
 barrier synchronization, 9-16-9-33
 BLT, 8-17-8-45. *See also* Option(s): EE
 functions, 8-3-8-16
 control and status, 8-2, 8-30-8-45,
 10-9-10-37
 BLT_CR, 8-30-8-35
 BLT_IVR, 8-23-8-24
 BLT_LAR, 8-28
 BLT_LSR, 8-29
 BLT_RBR, 8-26
 BLT_RIR, 8-18-8-20
 BLT_RLR, 8-27
 BLT_RMR, 8-24-8-25
 BLT_RSR, 8-21
 BLT_SR, 8-36-8-45
 BLT_VLR
 register, 8-22
 status, 8-41-8-42
 BSFR, 9-6, 9-17, 9-19-9-20
 BSMI, 9-7-9-8, 9-17, 9-21
 DEC_SYN_CLR, 11-100
 DRAM_CR, 10-21-10-28, 10-47-10-51
 EB option, 3-56
 EC option, 3-59
 F&I, 7-17, 7-50
 HB_GEN, 11-102
 HISP
 input, 11-22-11-33
 output, 11-40-11-49
 HISP_A_CFG, 11-81-11-86
 HISP_ADDR0, 11-92-11-94
 HISP_ADDR1, 11-95
 HISP_BUFFER, 11-96
 HISP_D_CFG, 11-87-11-89
 HISP_L_BL, 11-90
 HISP_ROA address, 11-97
 HISP_SYN_CLR, 11-101
 HISP_U_BL, 11-91
 I/O, 11-44
 illustrated, 12-9
 IO_CLR, 11-71-11-76
 IO_ERR, 11-64-11-70
 IO_IFLAG, 11-56-11-62
 IO_IMASK, 11-63
 IO_SYN, 11-98-11-100
 INV_CR, 10-19-10-21, 10-46-10-47
 L_WHOAMI, 6-21, 6-42, 6-65
 locations, address, 6-65
 LOSP
 input, 11-3-11-11
 output, 11-12-11-19
 LOSP_CFG, 11-77-11-78
 LOSP_DI, 11-3-11-9, 11-79
 LOSP_DO, 11-12-11-15, 11-80
 LOSP_DISC, 11-80
 LPE_XLATE, 6-46-6-47, 6-65
 message queue, 7-28
 MQ_TP, 7-51
 MQ_LIR, 7-52
 NET_ENA, 10-31
 NET_PFM, 10-32, 10-53-10-54
 NODE_CSR, 9-25-9-33, 10-34-10-37,
 10-54-10-56
 operational, 7-48-7-53
 PE node, illustrated, 12-7
 P_WHOAMI, 5-3, 5-5, 6-39-6-41, 6-65,
 13-6
 PRE_FETCH, 7-53

- Register(s) (continued)
- ROUTE_LO, 6-48–6-51, 6-65
 - ROUTE_HI, 6-53–6-54, 6-65
 - select bits, 8-55–8-56
 - SWAP, 7-49
 - V_WHOAMI, 6-43, 6-65
 - VRT_RG, 6-44, 6-65
 - X_WHOAMI, 6-45, 6-65
- Remote
- access to control panel, 13-9, 13-17, 13-28
 - address generation, 8-3–8-14
 - addressing registers, 8-18–8-27
 - base register (BLT_RBR), 8-26
 - index, 8-4–8-10
 - errors, checking, 10-35
 - indexing register (BLT_RIR), 8-18–8-20
 - limit register (BLT_RLR), 8-27
 - mask, 8-9
 - register (BLT_RMR), 8-24–8-25
 - memory, control of reference. *See*
 - Option(s): AM
 - memory location, reading repeatedly, 7-60
 - PE, 6-21
 - read operation (BLT), 8-57–8-68
 - request, generating, 7-68
 - stride, 8-4
 - register (BLT_RSR), 8-21
 - write
 - operation (BLT), 8-69–8-73
 - request, outstanding, 10-19
- REMOTE ENABLE switch, 13-9, 13-17, 13-28
- Request(s)
- address phit, 6-33, 6-35
 - buffer, 3-51
 - overview, 2-3
 - packet. *See also* Local address: generation
 - assembly, 6-62, 8-60
 - creating, 6-59–6-62
 - input, 7-70, 8-62
 - packet-header information, 7-56
 - PE, 8-45
 - remote, generating, 7-68
 - transfer, 11-23, 11-41
- Reset disable bit (AR option), 10-40
- Response(s)
- address phit, 6-33, 6-36
 - buffer, 3-51
 - header phits, generation of, 7-71
 - overview, 2-3
 - packet
 - assembly, 6-64, 8-67
 - creating, 6-62–6-64
 - header information, 7-58
 - header phits, 8-64
- RISC. *See* Microprocessor
- ROA buffer
- IG option, 4-21–4-22
 - output node, 4-35
- ROUTE_LO register, 6-48–6-51, 6-65
- ROUTE_HI register, 6-53–6-54, 6-65
- Router, 2-1, 2-15–2-18. *See also*
- Interconnect network
 - options, 2-19–2-28
- Routing, dimension order, 2-8–2-11
- Routing tag
- bit layout (barrier channel), 5-8
 - conversion from logical PE number, 6-29, 6-30
 - look-up table, 2-13, 6-48, 6-53, 8-60
 - high-order bits register (ROUTE_HI), 6-53
 - low-order bits register (ROUTE_LO), 6-48
 - and parity generation, 3-52
 - phit, 6-33
- RTI ENABLE switch, 13-29

S

- SADPAL code, 5-1, 5-7–5-17
- Scatter operation, 8-1
- SCR
 - and control and status, 10-10–10-13, 10-39–10-40
 - and the DTB annex index, 6-15–6-16

- SECEDED
 - and address space partition, 6-55
 - disable
 - HISP, 11-33
 - master HISP input, 11-49
 - slave HISP input, 11-39
 - slave HISP output, 11-54
 - errors
 - HISP, 11-30–11-31, 11-48
 - LOSP, 11-18
 - and input circuitry, 4-7–4-8
 - microprocessor, 3-9–3-12
- SECTION SELECT switches, 13-25
- Segment, virtual address, 6-10–6-11
- Selective reset mask (SCR), 10-10, 10-12
- SEND, 10-24
- Shared memory. *See* Local memory
- Signal(s). *See also* Interrupt; Option(s); Pins
 - Acknowledge, 2-23–2-24
 - Address and Block Length, 11-20, 11-21
 - Address Error, 11-20, 11-21, 11-86
 - Address Ready, 4-20, 11-20–11-21
 - assertion to memory chips, 3-33
 - barrier synchronization circuit, 9-16
 - Clear Channel, 4-20, 11-20–11-21
 - communication link channel, 2-2–2-3
 - control of, 10-1–10-8
 - data, 2-2, 2-3
 - decode, 4-16–4-19, 4-30–4-33
 - Disable SECEDED, 11-20, 11-21
 - Disconnect, 11-6–11-7
 - DRAM, generation of, 7-71
 - EB to EA option, 3-52
 - ED option packet control, 3-58, 3-61
 - Global Reset Heartbeat (deadstart), 5-1
 - hardware interrupt, 3-20, 10-2–10-5
 - heartbeat, 13-7–13-8, 13-15–13-16, 13-21–13-22
 - HISP channel, 4-3, 11-20, 11-21
 - IG option control, 4-21–4-22
 - Last Word, 11-20, 11-21
 - LOSP channel, 4-2, 11-2
 - Master Clear (deadstart), 5-1
 - microprocessor input/output, 3-17–3-21
 - Network Reset (deadstart), 5-1–5-2
 - packet assembly (EA option), 3-53
 - PE Reset (deadstart), 5-2
 - to ranks, 3-25–3-28
 - Transmit Address, 11-20, 11-21
 - Transmit Data, 4-20, 11-20, 11-21
 - virtual channel, 2-4–2-5
- Simplified barrier synchronization circuit, 9-3
- Single-bit errors. *See* Option(s): IE
- Single-cabinet configuration, 1-13
- Slave
 - HISP
 - input channel, 11-34–11-39
 - output transfer, 11-46
 - input circuitry, 4-24
 - I/O gateway. *See* I/O gateway
 - I/O module (illustrated), 12-15
 - output circuitry, 4-38
 - output transfers, 11-50–11-54
- Source phit, 6-35
- Spare PE, 6-2, 6-52
- Spare processing element module, 1-10
- SR option, 2-19–2-28, 10-52, 10-53
- SSR, 10-13–10-17, 10-43–10-46
- Stand-alone Diagnostic Privileged Architecture Library code (SADPAL), 5-1, 5-7–5-17
- Status
 - information, 10-9
 - register (BLT_SR), 8-36–8-45
- Status and control. *See also* Register(s):
 - control and status
 - information, 8-16
 - registers, 8-2
- STQ_C instruction, 7-6
- Swap. *See* Atomic swap; Read: operations
- Swaperand (SWAP) register, 7-49
- Switch(es)
 - BLT select bits, 10-34
 - CHASSIS SELECT, 13-6, 13-14
 - CLOCK SELECT, 13-4, 13-12, 13-20, 13-24
 - CONTROL DUMP, 13-28
 - control panel, illustrated, 13-2, 13-10, 13-18
 - CPU IDLE, 13-23
 - HEARTBEAT ENABLE, 13-8, 13-16, 13-22

Switch(es) (continued)
 HEARTBEAT INTERVAL, 13-7, 13-15, 13-21
 HEARTBEAT select, 13-8, 13-16, 13-22
 I/O gateway network router, 2-17
 LOSP ENABLE, 13-3, 13-11, 13-19
 MASTER CPU, 13-26
 MAINTENANCE MODE, 13-27
 MCU SELECT, 13-27
 REMOTE ENABLE, 13-9, 13-17, 13-28
 PE network routers, 2-15–2-18
 RTI ENABLE, 13-29
 SECTION SELECT, 13-25
 SYSTEM SIZE, 13-5, 13-13
 X-dimension, 2-15
 Y-dimension, 2-15
 Z-dimension, 2-15
 Synchronization. *See* Barrier
 synchronization; Eureka synchronization
 Syndrome register (IO_SYN), 11-98–11-100
 System
 clock. *See* Clock
 control register. *See* SCR
 heartbeat signal
 HEARTBEAT ENABLE switch, 13-8, 13-16, 13-22
 HEARTBEAT INTERVAL switch, 13-7, 13-15, 13-21
 status register. *See* SSR
 SYSTEM SIZE switches, 13-5, 13-13
 System/clock module, 1-12

T

Tables. *See* Routing tag
 Tail pointer, 7-30–7-31
 Timing values
 CAS cycle, 10-27
 DRAM refresh, 10-30
 Torus topology, 2-5–2-7
 Transfer. *See also* DRAM: control: register.
 timing information, 10-25–10-26

types (BLT), 8-31
 Translation, data, circuitry, 4-8, 4-10–4-11
 Transmit
 address signal, 11-20, 11-21
 data signal, 4-20, 11-20, 11-21
 TTL-to-ECL translation. *See* Option(s): IA

U

Unrecoverable error, 11-20, 11-21, 11-38–11-39
 User status register (USR), 10-18–10-19, 10-46

V

V_WHOAMI register, 6-43, 6-65
 Validation bit, 5-8
 VC. *See* Virtual channel
 Vector length register (BLT_VLR), 8-22
 Virtual address, 6-10–6-12
 converting to partial physical address, 6-13
 Virtual channel (VC), 2-3
 buffer, 2-5, 2-12–2-13
 conflicts, 2-27, 2-28
 select bits, 2-4–2-5
 signals, 2-4–2-5
 Virtual index, 6-12
 Virtual node, numbers, 6-5, 6-23, 6-25, 6-28
 Virtual PE
 address interpretation, 6-19–6-20, 6-56
 number, 6-5–6-9
 BLT interpretation of, 8-34
 converting to logical PE number, 6-23–6-28, 6-57–6-58
 range check, 8-12–8-13
 range error, 10-13, 10-16, 10-45
 range mask
 register (VRT_RG), 6-44, 6-65
 value, 6-20
 register (V_WHOAMI), 6-43, 6-65
 VRT_RG register, 6-44, 6-65

W

WE generation. *See* Option(s): AE
Wire length (SR options), 2-25
Write. *See also* Path, data
 buffer, 3-7-3-9
 operation, 7-15-7-17, 7-80-7-84. *See also*
 Remote: write: operation
 to PLM, illegal, 10-46
WRITE_BLOCK. *See* Write: operation

X

X-dimension switch, 2-15
X mask, 6-25

X_WHOAMI register, 6-45, 6-65

Y

Y-dimension switch, 2-15
Y mask, 6-25
Y shift, 6-23

Z

Z-dimension switch, 2-15
Z mask, 6-25
Z shift, 6-23