

KN210 CPU Module Set Technical Manual

Order Number EK-KN210-TM-001

**digital equipment corporation
maynard, massachusetts**

First Edition, June 1989

The information in this document is subject to change without notice and should not be construed as a commitment by Digital Equipment Corporation. Digital Equipment Corporation assumes no responsibility for any errors that may appear in this document.

The software described in this document is furnished under a license and may be used or copied only in accordance with the terms of such license.

No responsibility is assumed for the use or reliability of software on equipment that is not supplied by Digital Equipment Corporation or its affiliated companies.

Copyright © by Digital Equipment Corporation 1989

All Rights Reserved.

Printed in U.S.A.

The following are trademarks of Digital Equipment Corporation:

DEC	MicroVAX 3400	RT
DECmate	MicroVAX 3500	ThinWire
DECUS	MicroVAX 3600	UNIBUS
DECwriter	PDP	VAX
DIBOL	P/OS	VAXstation
LSI-11	Professional	VMS
MASSBUS	Q-bus	VT
MicroPDP-11	Q22-bus	VT100
MicroVAX	Rainbow	Work Processor
MicroVAX I	RSTS	
MicroVAX II	RSX	

digital[™]

FCC NOTICE: The equipment described in this manual generates, uses, and may emit radio frequency energy. The equipment has been type tested and found to comply with the limits for a Class A computing device pursuant to Subpart J of Part 15 of FCC Rules, which are designed to provide reasonable protection against such radio frequency interference when operated in a commercial environment. Operation of this equipment in a residential area may cause interference, in which case the user at his own expense may be required to take measures to correct the interference.

Contents

- About This Manual** xxi

- 1 Overview**
 - 1.1 Introduction 1-1
 - 1.2 R3000 RISC Processor 1-4
 - 1.3 Floating-Point Accelerator 1-5
 - 1.4 Cache Memory 1-5
 - 1.5 Memory Controller 1-5
 - 1.6 Diagnostic Processor 1-5
 - 1.7 MS650-BA Memory Modules 1-5
 - 1.8 MS650-AA Memory Modules 1-6
 - 1.9 DSSI Interface 1-6
 - 1.10 Ethernet Interface 1-7
 - 1.11 Q22-bus Interface 1-7
 - 1.12 System Support Functions 1-7
 - 1.13 Firmware 1-8
 - 1.14 Clock Functions 1-8

- 2 Installation and Configuration**
 - 2.1 Introduction 2-1
 - 2.2 Installing the KN210 CPU Module Set 2-1
 - 2.3 Configuring the KN210 2-3
 - 2.4 KN210 Connectors 2-4
 - 2.4.1 Console/Ethernet Connector 2-5
 - 2.4.2 DSSI/I/O Connector 2-7
 - 2.4.3 Memory Connector 2-9

2.5	H3602-SA CPU Cover Panel	2-11
-----	--------------------------------	------

3 Architecture

3.1	R3000 RISC Processor	3-1
3.1.1	Processor Features	3-2
3.1.1.1	General Purpose Registers	3-2
3.1.1.2	Instruction Set	3-3
3.1.2	Coprocessors	3-6
3.1.2.1	Coprocessor (0)	3-6
3.1.2.2	Coprocessor (1)	3-6
3.1.3	Memory Management	3-6
3.1.3.1	Operating Modes	3-9
3.1.3.2	EntryHi and EntryLo Registers	3-9
3.1.3.3	Index Register	3-11
3.1.3.4	Random Register	3-11
3.1.4	Exception Handling Registers	3-12
3.1.4.1	Cause Register	3-12
3.1.4.2	Exception Program Counter	3-15
3.1.4.3	Status Register	3-15
3.1.4.4	BadVaddr Register	3-16
3.1.4.5	Context Register	3-16
3.1.4.6	Processor Revision Identifier Register	3-17
3.1.4.7	Interrupt Status Register	3-17
3.1.4.8	Vector Read Registers	3-18
3.1.5	Exceptions	3-21
3.1.5.1	General Exception Vector	3-21
3.1.5.2	Reset Exception Vector	3-21
3.2	Floating-Point Accelerator	3-22
3.2.1	Floating-Point Accelerator Instructions	3-22
3.3	Cache Memory	3-23
3.3.1	Cache Organization	3-23
3.3.2	Cache Isolation	3-23
3.3.3	Cache Swapping	3-23
3.3.4	Cache Line Format	3-23
3.4	Main Memory System	3-24

3.4.1	Main Memory Organization	3-27
3.4.2	Main Memory Addressing	3-27
3.4.3	Main Memory Behavior on Writes	3-28
3.4.4	Main Memory Error Status Register	3-28
3.4.5	Main Memory Control and Diagnostic Status Register . .	3-32
3.4.6	Main Memory Error Detection and Correction	3-34
3.5	Console Serial Line	3-36
3.5.1	Console Registers	3-36
3.5.1.1	Console Receiver Control/Status Register	3-37
3.5.1.2	Console Receiver Data Buffer	3-37
3.5.1.3	Console Transmitter Control/Status Register	3-39
3.5.1.4	Console Transmitter Data Buffer	3-40
3.5.2	Break Response	3-41
3.5.3	Baud Rate	3-41
3.5.4	Console Interrupt Specifications	3-42
3.6	Time-of-Year Clock and Timers	3-42
3.6.1	R3000 Interval Timer Register	3-42
3.6.2	Time-of-Year Clock	3-43
3.6.3	Interval Timer	3-44
3.6.4	Programmable Timers	3-45
3.6.4.1	Timer Control Registers	3-45
3.6.4.2	Timer Interval Registers	3-47
3.6.4.3	Timer Next Interval Registers	3-48
3.6.4.4	Timer Interrupt Vector Registers	3-48
3.7	Boot and Diagnostic Facility	3-49
3.7.1	Boot and Diagnostic Register	3-49
3.7.2	Diagnostic LED Register	3-51
3.7.3	ROM Memory	3-52
3.7.3.1	ROM Socket	3-52
3.7.3.2	ROM Address Space	3-52
3.7.4	Battery Backed-Up RAM	3-53
3.7.5	KN210 Initialization	3-53
3.7.5.1	Power-Up Initialization	3-53
3.7.5.2	Hardware Reset	3-53
3.7.5.3	I/O Bus Initialization	3-53
3.7.5.4	Processor Initialization	3-54

3.7.6	Select Processor Register	3-54
3.7.7	Write Error Address Register	3-55
3.8	Q22-bus Interface	3-55
3.8.1	Q22-bus to Main Memory Address Translation	3-56
3.8.1.1	Q22-bus Map Registers	3-58
3.8.1.2	Accessing the Q22-bus Map Registers	3-59
3.8.1.3	Q22-bus Map Cache	3-60
3.8.2	CDAL Bus to Q22-bus Address Translation	3-62
3.8.3	Interprocessor Communication Register	3-62
3.8.4	Q22-bus Interrupt Handling	3-63
3.8.5	Configuring the Q22-bus Map	3-64
3.8.5.1	Q22-bus Map Base Address Register	3-64
3.8.6	System Configuration Register	3-65
3.8.7	DMA System Error Register	3-66
3.8.8	Q22-bus Error Address Register	3-69
3.8.9	DMA Error Address Register	3-69
3.8.10	Error Handling	3-70
3.9	Network Interface	3-72
3.9.1	Ethernet Overview	3-72
3.9.2	Network Interface Station Address ROM	3-74
3.9.3	LANCE Chip Overview	3-75
3.9.4	Network Interface Register Address Port	3-77
3.9.5	Network Interface Register Data Port	3-78
3.9.6	Network Interface Control and Status Register 0	3-78
3.9.7	Network Interface Control and Status Register 1	3-84
3.9.8	Network Interface Control and Status Register 2	3-84
3.9.9	Network Interface Control and Status Register 3	3-85
3.9.10	Network Interface Initialization Block	3-86
3.9.10.1	Network Interface Initialization Block Word 0	3-87
3.9.10.2	Network Interface Initialization Block Words 1-3	3-90
3.9.10.3	Network Interface Initialization Block Words 4-7	3-91
3.9.10.4	Network Interface Initialization Block Words 8,9	3-92
3.9.10.5	Network Interface Initialization Block Words 10,11	3-94
3.9.11	Buffer Management	3-95
3.9.12	Network Interface Receive Descriptor Ring	3-96
3.9.12.1	Receive Buffer Descriptors	3-97

3.9.13	Receive Buffers	3-102
3.9.14	Network Interface Transmit Descriptor Ring	3-103
3.9.14.1	Transmit Buffer Descriptors	3-104
3.9.15	Transmit Buffers	3-109
3.9.16	LANCE Operation	3-109
3.9.16.1	Switch Routine	3-110
3.9.16.2	Initialization Routine	3-110
3.9.16.3	Look-For-Work Routine	3-111
3.9.16.4	Receive Poll Routine	3-111
3.9.16.5	Receive Routine	3-111
3.9.16.6	Receive DMA Routine	3-112
3.9.16.7	Transmit Poll Routine	3-112
3.9.16.8	Transmit Routine	3-113
3.9.16.9	Transmit DMA Routine	3-113
3.9.16.10	Collision Detect Routine	3-114
3.9.17	LANCE Programming Notes	3-114
3.10	Mass Storage Interface	3-117
3.10.1	DSSI Bus Overview	3-117
3.10.2	Target Operation	3-119
3.10.3	Initiator Operation	3-121
3.10.3.1	Transmit Data Segment Links	3-121
3.10.4	Adding to a Buffer List	3-122
3.10.5	Mass Storage Interface Command Block (MSICB)	3-124
3.10.5.1	MSI Command Block Word 0	3-124
3.10.5.2	MSI Command Block Word 1	3-125
3.10.5.3	MSI Command Block Word 2	3-127
3.10.5.4	MSI Command Block Words 3-5	3-128
3.10.6	MSI Registers	3-128
3.10.6.1	MSI Control and Status Registers	3-128
3.10.6.2	List Pointer Registers	3-138
3.10.6.3	Diagnostic and Test Registers	3-139

4 KN210 Firmware

4.1	Firmware Capabilities	4-2
4.2	Power-Up	4-2
4.2.1	Initial Power-Up Test	4-2
4.2.2	Locating a Console Device	4-3
4.2.3	Operation and Function Switches	4-3
4.2.3.1	Operation Switch Set to Normal	4-4
4.2.3.2	Operation Switch Set to Maintenance	4-5
4.2.3.3	Operation Switch Set to Action	4-6
4.2.3.4	LED Codes	4-7
4.2.4	Interprocessor Interaction	4-8
4.2.4.1	Select Processor Register Operation	4-9
4.2.5	Power-Up Sequence	4-9
4.2.5.1	Normal Power-Up Operation	4-9
4.2.5.2	Maintenance Power-Up Operation	4-10
4.2.6	Processor Identification	4-10
4.2.6.1	Sys_Type Register Layout	4-10
4.3	Operating System Bootstrap	4-11
4.3.1	MDM Bootstrap	4-11
4.3.2	Operating System Bootstrap	4-12
4.3.3	Boot Process	4-13
4.3.4	Bootstrap Support Routines in the Console	4-14
4.3.5	Console Use of Memory Space	4-14
4.3.6	Boot Devices	4-15
4.3.6.1	Disk	4-15
4.3.6.2	Tape	4-15
4.3.6.3	Ethernet	4-15
4.3.7	Halts	4-16
4.4	KN210 Console Command Language	4-17
4.4.1	Maintenance Mode Console Command Language	4-17
4.4.2	Normal Mode Console Command Language	4-17
4.4.2.1	Control Characters	4-17
4.4.2.2	Lexical Conventions	4-18
4.4.2.3	Environment Variables	4-19
4.4.2.4	Commands	4-20
4.5	Diagnostics	4-23

4.6	PROM Entry Points	4-23
4.6.1	Argvize	4-23
4.6.2	Atob	4-23
4.6.3	Autoboot	4-24
4.6.4	Bevexcept	4-24
4.6.5	Bevutlb	4-24
4.6.6	Close	4-24
4.6.7	Dumpcmd	4-24
4.6.8	Exec	4-25
4.6.9	Getchar	4-25
4.6.10	Getenv	4-25
4.6.11	Gets	4-25
4.6.12	Halt	4-25
4.6.13	Help	4-25
4.6.14	Ioctl	4-26
4.6.15	Longjump	4-26
4.6.16	Lseek	4-26
4.6.17	Open	4-26
4.6.18	Parser	4-27
4.6.19	Printenvcmd	4-27
4.6.20	Printf	4-27
4.6.21	Puchar	4-27
4.6.22	Puts	4-28
4.6.23	Range	4-28
4.6.24	Read	4-28
4.6.25	Reboot	4-28
4.6.26	Reinit	4-29
4.6.27	Reset	4-29
4.6.28	Restart	4-29
4.6.29	Setenv	4-29
4.6.30	Setenvcmd	4-29
4.6.31	Setjmp	4-30
4.6.32	Showchar	4-30
4.6.33	Strcat	4-30
4.6.34	Strcmp	4-30
4.6.35	Strcpy	4-31

4.6.36	Strlen	4-31
4.6.37	Unsetenvcmd	4-31
4.6.38	Write	4-31
4.7	Supported Devices	4-32

5 Diagnostic Processor

5.1	Diagnostic Processor	5-1
5.1.1	Processor State	5-1
5.1.1.1	General Purpose Registers	5-1
5.1.1.2	Processor Status Longword	5-2
5.1.1.3	Internal Processor Registers	5-4
5.1.2	Data Types	5-8
5.1.3	Instruction Set	5-8
5.1.4	Memory Management	5-9
5.1.4.1	Translation Buffer	5-9
5.1.4.2	Memory Management Control Registers	5-10
5.1.5	Exceptions and Interrupts	5-11
5.1.5.1	Interrupts	5-11
5.1.5.2	Exceptions	5-13
5.1.5.3	Information Saved on a Machine Check Exception	5-16
5.1.5.4	System Control Block	5-21
5.1.5.5	Diagnostic Processor Hardware Detected Errors	5-24
5.1.5.6	Hardware Halt Procedure	5-25
5.1.6	System Identification	5-27
5.1.7	CVAX References	5-28
5.1.7.1	Instruction-Stream Read References	5-28
5.1.7.2	Data-Stream Read References	5-29
5.1.7.3	Write References	5-29

6 Maintenance Mode Firmware

6.1	Maintenance Mode Firmware Features	6-1
6.1.1	Halt Entry, Exit, and Dispatch	6-2
6.1.1.1	Halt Entry - Saving Processor State	6-2
6.1.1.2	Halt Exit - Restoring Processor State	6-3
6.1.1.3	Halt Dispatch	6-3
6.1.1.4	External Halts	6-4
6.1.2	Power-Up	6-5
6.1.2.1	Initial Power-Up Test	6-5
6.1.2.2	Locating a Console Device	6-5
6.1.3	Mode Switch Set to "Test"	6-6
6.1.4	Mode Switch Set to "Query"	6-7
6.1.5	Mode Switch Set to "Normal"	6-8
6.1.6	LED Codes	6-9
6.2	Console Service	6-9
6.2.1	Console Control Characters	6-10
6.2.2	Console Command Syntax	6-12
6.2.3	Console Command Keywords	6-12
6.2.4	Console Command Qualifiers	6-14
6.2.5	Command Address Specifiers	6-14
6.2.6	References to Processor Registers and Memory	6-17
6.2.7	Console Commands	6-18
6.2.7.1	BOOT	6-18
6.2.7.2	CONFIGURE	6-20
6.2.7.3	CONTINUE	6-22
6.2.7.4	DEPOSIT	6-22
6.2.7.5	EXAMINE	6-24
6.2.7.6	EXIT	6-27
6.2.7.7	FIND	6-28
6.2.7.8	HALT	6-29
6.2.7.9	HELP	6-29
6.2.7.10	INITIALIZE	6-31
6.2.7.11	MOVE	6-33
6.2.7.12	NEXT	6-35
6.2.7.13	REPEAT	6-36
6.2.7.14	SEARCH	6-37

xii Contents

6.2.7.15	SET	6-40
6.2.7.16	SHOW	6-44
6.2.7.17	START	6-48
6.2.7.18	TEST	6-48
6.2.7.19	UNJAM	6-52
6.2.7.20	X - Binary Load and Unload	6-52
6.2.7.21	! - Comment	6-54
6.2.8	Conventions for Tables 6-5 and 6-6	6-55
6.3	Bootstrapping	6-58
6.3.1	Boot Devices	6-58
6.3.2	Boot Flags	6-60
6.3.3	Preparing for the Bootstrap	6-62
6.3.4	Primary Bootstrap, VMB	6-63
6.3.5	Device Dependent Bootstrap Procedures	6-66
6.3.5.1	Disk and Tape Bootstrap Procedure	6-66
6.3.5.2	PROM Bootstrap Procedure	6-67
6.3.5.3	Network Bootstrap Procedure	6-68
6.3.5.4	Network Listening	6-69
6.4	Diagnostics	6-71
6.4.1	Error Reporting	6-73
6.4.2	Diagnostic Interdependencies	6-74
6.4.3	Areas Not Covered	6-75
6.5	Operating System Restart	6-75
6.5.1	Locating the RPB	6-76
6.6	Machine State on Power-Up	6-77
6.6.1	Main Memory Layout and State	6-77
6.6.1.1	Reserved Main Memory	6-78
6.6.1.2	PFN Bitmap	6-78
6.6.1.3	Scatter/Gather Map	6-79
6.6.1.4	Contents of Main Memory	6-79
6.6.2	CMCTL Registers	6-79
6.6.3	First Level Cache	6-80
6.6.4	Translation Buffer	6-80
6.6.5	Halt Protected Space	6-80
6.7	Public Data Structures and Entry Points	6-80
6.7.1	Maintenance Mode Firmware EPROM Layout	6-80

6.7.2	Call-Back Entry Points	6-82
6.7.2.1	CP\$GETCHAR_R4	6-82
6.7.2.2	CP\$MESSG_OUT_NOLF_R4	6-83
6.7.2.3	CP\$READ_WTH_PRMPRT_R4	6-84
6.7.3	SSC RAM Layout	6-85
6.7.4	Public Data Structures	6-85
6.7.4.1	Console Program Mailbox	6-85
6.7.4.2	Firmware Stack	6-87
6.7.4.3	Diagnostic State	6-87
6.7.4.4	User Area	6-87
6.8	Error Messages	6-87
6.8.1	Halt Code Messages	6-88
6.8.2	Console Error Messages	6-89
6.8.3	VMB error messages	6-91

A Specifications

A.1	Physical Specifications	A-1
A.2	Electrical Specifications	A-1
A.3	Environmental Specifications	A-2

B Address Assignments

B.1	R3000 Physical Address Space Map	B-1
B.2	CVAX Physical Address Space Map	B-4
B.2.1	External IPRs	B-6
B.3	Global Q22-bus Address Space Map	B-7

C Q22-bus Specification

C.1	Introduction	C-1
C.1.1	Master/Slave Relationship	C-2
C.2	Q22-bus Signal Assignments	C-3
C.3	Data Transfer Bus Cycles	C-6
C.3.1	Bus Cycle Protocol	C-7
C.3.2	Device Addressing	C-7
C.4	Direct Memory Access	C-17

xiv Contents

C.4.1	DMA Protocol	C-17
C.4.2	Block Mode DMA	C-18
C.4.2.1	DATBI Bus Cycle	C-23
C.4.2.2	DATBO Bus Cycle	C-24
C.4.3	DMA Guidelines	C-26
C.5	Interrupts	C-27
C.5.1	Device Priority	C-28
C.5.2	Interrupt Protocol	C-28
C.5.3	Q22-bus Four-Level Interrupt Configurations	C-32
C.6	Control Functions	C-34
C.6.1	Halt	C-34
C.6.2	Initialization	C-34
C.6.3	Power Status	C-34
C.7	Q22-bus Electrical Characteristics	C-34
C.7.1	Signal Level Specifications	C-35
C.7.2	Load Definition	C-35
C.7.3	120-Ohm Q22-bus	C-35
C.7.4	Bus Drivers	C-36
C.7.5	Bus Receivers	C-36
C.7.6	Bus Termination	C-37
C.7.7	Bus Interconnecting Wiring	C-38
C.7.7.1	Backplane Wiring	C-38
C.7.7.2	Intrabackplane Bus Wiring	C-39
C.7.7.3	Power and Ground	C-39
C.8	System Configurations	C-39
C.8.1	Power Supply Loading	C-43
C.9	Module Contact Finger Identification	C-43

D Acronyms

Index

Examples

4-1	Normal Diagnostic Countdown	4-4
4-2	Abnormal Diagnostic Countdown	4-5
4-3	Language Selection Menu	4-7
4-4	Normal Power-Up Sequence Algorithm	4-9
4-5	Maintenance Power-Up Sequence Algorithm	4-10
6-1	Language Selection Menu	6-7
6-2	Normal Diagnostic Countdown	6-8
6-3	Abnormal Diagnostic Countdown	6-8
6-4	Diagnostic Register Dump	6-73

Figures

1-1	KN210 Processor Module	1-2
1-2	KN210 I/O Module	1-3
1-3	KN210 Block Diagram	1-4
1-4	MS650 Memory Modules	1-6
2-1	Processor, I/O and Memory Module Placement	2-2
2-2	Cable Connections	2-3
2-3	KN210 Connector and LED Orientation	2-4
2-4	H3602-SA CPU Cover Panel	2-12
3-1	Instruction Set Types	3-3
3-2	R3000 Virtual to Physical Memory Map	3-7
3-3	R3000 to I/O Memory Map	3-8
3-4	TLB EntryHi Register	3-9
3-5	TLB EntryLo Register	3-10
3-6	Index Register	3-11
3-7	Random Register	3-11
3-8	Cause Register	3-12
3-9	Interrupt Acknowledge Daisy Chain	3-14
3-10	Status Register	3-15
3-11	Context Register	3-16
3-12	Processor Revision Identifier Register	3-17
3-13	Interrupt Status Register	3-18
3-14	Vector Read Register 0	3-19
3-15	Vector Read Register 1	3-20

3-16	Vector Read Register 2	3-20
3-17	Vector Read Register 3	3-20
3-18	Cache Line Format	3-23
3-19	Tag Field	3-24
3-20	Format for MEMCSR16	3-28
3-21	Format for MEMCSR17	3-32
3-22	Console Receiver Control/Status Register	3-37
3-23	Console Receiver Data Buffer	3-38
3-24	Console Transmitter Control/Status Register	3-39
3-25	Console Transmitter Data Buffer	3-40
3-26	Interval Timer Register	3-42
3-27	Time-of-Year Clock	3-44
3-28	Interval Timer	3-44
3-29	Timer Control Registers	3-46
3-30	Timer Interrupt Vector Register	3-48
3-31	Boot and Diagnostic Register	3-49
3-32	Diagnostic LED Register	3-51
3-33	Select Processor Register	3-54
3-34	Q22-bus to Main Memory Address Translation	3-57
3-35	Q22-bus Map Registers	3-58
3-36	Q22-bus Map Cache Entry	3-61
3-37	The Interprocessor Communication Register	3-62
3-38	Q22-bus Map Base Address Register	3-64
3-39	System Configuration Register	3-65
3-40	DMA System Error Register	3-67
3-41	Q22-bus Error Address Register	3-69
3-42	DMA Error Address Register	3-70
3-43	Ethernet Data Packet Format	3-73
3-44	Network Interface Station Address ROM Format	3-75
3-45	Network Interface Register Address Port	3-77
3-46	Network Interface Control and Status Register 0	3-79
3-47	Network Interface Control and Status Register 1	3-84
3-48	Network Interface Control and Status Register 2	3-85
3-49	Network Interface Control and Status Register 3	3-85
3-50	Network Interface Initialization Block	3-87
3-51	Network Interface Initialization Block Word 0	3-88

3-52	Network Interface Initialization Block Words 1-3	3-91
3-53	Network Interface Initialization Block Words 4-7	3-91
3-54	Network Interface Initialization Block Word 8	3-92
3-55	Network Interface Initialization Block Word 9	3-93
3-56	Network Interface Initialization Block Word 10	3-94
3-57	Network Interface Initialization Block Word 11	3-94
3-58	Network Interface Receive Descriptor Ring	3-97
3-59	Receive Buffer Descriptors	3-98
3-60	Receive Buffer Descriptor n Word 0	3-98
3-61	Receive Buffer Descriptor n Word 1	3-99
3-62	Receive Buffer Descriptor n Word 2	3-101
3-63	Receive Buffer Descriptor n Word 3	3-101
3-64	Receive Buffers	3-102
3-65	Network Interface Transmit Descriptor Ring	3-103
3-66	Transmit Buffer Descriptors	3-104
3-67	Transmit Buffer Descriptor n Word 0	3-104
3-68	Transmit Buffer Descriptor n Word 1	3-105
3-69	Transmit Buffer Descriptor n Word 2	3-106
3-70	Transmit Buffer Descriptor n Word 3	3-107
3-71	Transmit Buffers	3-109
3-72	DSSI Bus Sequences	3-118
3-73	Target Operation	3-120
3-74	Transmit Data Segment Links	3-121
3-75	Initiator Operation	3-123
3-76	MSI Command Block	3-124
3-77	MSI Command Block Word 0	3-124
3-78	MSI Command Block Word 1	3-125
3-79	MSI Command Block Word 2	3-127
3-80	MSI Control/Status Register	3-129
3-81	MSI DSSI Control Register	3-130
3-82	MSI DSSI Connection Register	3-132
3-83	MSI ID Register	3-136
3-84	MSI DSSI Timeout Register	3-137
3-85	MSI Target List Pointer Register	3-138
3-86	MSI Initiator List Pointer Register	3-139
3-87	MSI Diagnostic Control Register	3-140

3-88	MSI Diagnostic Register 0	3-141
3-89	MSI Diagnostic Register 1	3-142
3-90	MSI Diagnostic Register 2	3-144
3-91	MSI Clock Control Register	3-145
4-1	System Type Register	4-10
4-2	Bootblock Layout	4-13
5-1	General Purpose Register Bit Map	5-2
5-2	PSL Bit Map	5-3
5-3	Interrupt Registers	5-13
5-4	Information Saved on a Machine Check Exception	5-16
5-5	System Control Block Base Register	5-21
5-6	System Identification Register	5-27
5-7	System Type Register	5-28
6-1	VMB Boot Flags	6-61
6-2	Memory Layout at VMB Exit	6-65
6-3	Boot Block Format	6-67
6-4	RPB Signature Format	6-76
6-5	Memory Layout After Power-Up Diagnostics	6-77
6-6	KN210 Maintenance Mode EPROM Layout	6-81
6-7	KN210 SSC NVRAM Layout	6-85
6-8	NVR0	6-86
6-9	NVR1	6-86
6-10	NVR2	6-87
C-1	DATI Bus Cycle	C-9
C-2	DATI Bus Cycle Timing	C-11
C-3	DATO or DATOB Bus Cycle	C-12
C-4	DATO or DATOB Bus Cycle Timing	C-14
C-5	DATIO or DATIOB Bus Cycle	C-15
C-6	DATIO or DATIOB Bus Cycle Timing	C-16
C-7	DMA Protocol	C-19
C-8	DMA Request/Grant Timing	C-20
C-9	DATBI Bus Cycle Timing	C-21
C-10	DATBO Bus Cycle Timing	C-22
C-11	Interrupt Request/Acknowledge Sequence	C-29
C-12	Interrupt Protocol Timing	C-31
C-13	Position-Independent Configuration	C-33

C-14	Position-Dependent Configuration	C-33
C-15	Bus Line Terminations	C-37
C-16	Single-Backplane Configuration	C-40
C-17	Multiple Backplane Configuration	C-42
C-18	Typical Pin Identification System	C-43
C-19	Quad-Height Module Contact Finger Identification	C-44
C-20	Typical Q22-bus Module Dimensions	C-45

Tables

2-1	Console/Ethernet Connector (J1) Pinouts	2-5
2-2	DSSI Connector (I/O Module J2-Upper) Pinouts	2-8
2-3	I/O Connector (I/O Module J2-Lower) Pinouts	2-9
2-4	Memory Connector (Processor Module J2-Lower) Pinouts	2-10
2-5	H3602-SA Features and Controls	2-11
3-1	Instruction Summary	3-3
3-2	R3000 Interrupt Mapping	3-14
3-3	CPU Read Reference Timing	3-25
3-4	CPU Write Reference Timing	3-25
3-5	Q22-bus Interface Read Reference Timing	3-25
3-6	Q22-bus Interface Write Reference Timing	3-26
3-7	Error Syndromes	3-30
3-8	Console Registers	3-36
3-9	Baud Rate Select	3-42
3-10	Q22-bus Map	3-59
4-1	Operation and Function Switches	4-4
4-2	LED Codes	4-8
4-3	Console Memory Space	4-14
4-4	PROM_HALT Saved Registers	4-16
4-5	Control Characters	4-17
4-6	Environment Variables	4-19
4-7	KN210 Boot Devices	4-32
5-1	KN210 Internal Processor Registers	5-5
5-2	VAX Standard IPRs	5-7
5-3	KN210 Unique IPRs	5-8
5-4	Interrupts	5-11
5-5	Exceptions	5-15

5-6	System Control Block Format	5-22
5-7	Unmaskable Interrupts That can Cause a Halt	5-26
5-8	Exceptions That can Cause a Halt	5-26
6-1	Halt Action Summary	6-4
6-2	LED Codes	6-9
6-3	Command, Parameter, and Qualifier Keywords	6-13
6-4	Console Symbolic Addresses	6-14
6-5	Console Command Summary	6-55
6-6	Console Qualifier Summary	6-57
6-7	KN210 Supported Boot Devices	6-59
6-8	VMB Boot Flags	6-61
6-9	KN210 Network Maintenance Operations Summary	6-70
6-10	HALT Messages	6-88
6-11	Console Error Messages	6-89
6-12	VMB Error Messages	6-91
B-1	R3000 Memory Space	B-1
B-2	R3000 I/O Space	B-1
B-3	VAX Memory Space	B-4
B-4	CVAX I/O Space	B-4
B-5	External IPRs	B-6
B-6	Q22-bus Memory Space	B-7
B-7	Q22-bus I/O Space with BBS7 Asserted	B-7
C-1	Data and Address Signal Assignments	C-3
C-2	Control Signal Assignments	C-4
C-3	Power and Ground Signal Assignments	C-5
C-4	Spare Signal Assignments	C-5
C-5	Data Transfer Operations	C-6
C-6	Bus Signals for Data Transfers	C-6
C-7	Bus Pin Identifiers	C-45

About This Manual

This *KN210 CPU Module Set Technical Manual* documents the functional, physical, and environmental characteristics of the KN210 CPU module set and also includes some information on the MS650 memory expansion modules. The KN210 CPU module set consists of the KN210 processor module and the KN210 I/O module.

Intended Audience

This document is intended for a design engineer or applications programmer who is familiar with Digital's extended LSI-11 bus (Q22-bus) and the MIPS instruction set. This manual should be used along with the *VAX Architecture Reference Manual* as a programmer's reference to the module.

Organization

The manual is divided into six chapters and four appendices.

Chapter 1, **Overview**, introduces the KN210 processor module, the KN210 I/O module and MS650 memory modules, including module features and specifications.

Chapter 2, **Configuration and Installation**, describes the configuration and installation of the KN210 module set and MS650 modules in Q22-bus backplanes and system enclosures.

Chapter 3, **Architecture**, provides a description of KN210 registers, instruction set and memory.

Chapter 4, **KN210 Firmware**, describes the R3000 entry/dispatch code, boot diagnostics, device booting sequence, console program and console commands.

Chapter 5, **Diagnostic Processor**, describes the diagnostic processor registers.

Chapter 6, **Maintenance Mode Firmware**, describes the diagnostic processor entry/dispatch code, boot diagnostics, device booting sequence, console program and console commands.

Appendix A, **KN210 Specifications**, describes the physical, electrical, and environmental specifications for the KN210 CPU and I/O modules.

Appendix B, **Address Assignments**, provides a map of R3000 and diagnostic processor address space.

Appendix C, **Q22-bus Specification**, describes the low end member of Digital's bus family. All of Digital's microcomputers use the Q22-bus.

Appendix D, **Acronyms**, provides a list of the acronyms used in this manual.

Conventions

This manual uses the following conventions:

Convention	Meaning
Note	Provides general information you should be aware of.
Caution	Provides information to prevent damage to equipment.
<x:y>	Represents a bit field, a set of lines, or signals, ranging from x through y. For example, R0 <7:4> indicates bits 7 through 4 in general purpose register R0.
[x:y]	Represents a range of bytes, from y through x.
Return	Text within a box identifies a key, such as the Return key.
n	Boldface small n indicates variables.

Related Documents

You can order the following documents from Digital:

Document	Order Number
Microcomputer Interfaces Handbook	EB-20175-20
Microcomputers and Memories Handbook	EB-18451-20
VAX Architecture Handbook	EB-19580-20
VAX Architecture Reference Manual	EY-3459E-DP
BA213 Enclosure Maintenance	EK-189AA-MG

You can order these documents from:

Digital Equipment Corporation
Accessories and Supplies Group
P.O. Box CS2008
Nashua, NH 03061
Attention: Documentation Products

1

Overview

This chapter provides a brief overview of the KN210 CPU module set and MS650 memory modules.

1.1 Introduction

The KN210 CPU module set consists of the KN210 processor module, shown in Figure 1-1, and the KN210 I/O module, shown in Figure 1-2. Both are quad-height modules for the Q22-bus (extended LSI-11 bus). The KN210 processor module features the R3000 RISC processor. The KN210 I/O module features built-in DSSI and Ethernet controllers. The KN210 CPU module set is designed for use in high speed multiuser, multitasking environments.

The KN210 CPU module set is used in the DECsystem 5400. The DECsystem 5400 utilizes a BA213 enclosure. Refer to the *BA213 Enclosure Maintenance* for a detailed description of the enclosure.

The KN210 CPU module set and MS650 memory modules combine to form a RISC CPU/memory subsystem that uses the Q22-bus, DSSI bus, and Ethernet to communicate with mass storage and I/O devices. The KN210 CPU module set and MS650 modules mount in standard Q22-bus backplane slots that implement the Q22-bus in the AB rows and the CD interconnect in the CD rows. The KN210 CPU module set can support up to four MS650 modules, if enough Q22/CD slots are available.

The KN210 CPU module set communicates with the console device through the H3602-SA CPU cover panel, which also contains configuration switches, an LED display, and Ethernet connector.

1-2 Overview

Insert photo of KN210 CPU Module here.

Figure 1-1 KN210 Processor Module

Insert photo of KN210 I/O Module here.

Figure 1-2 KN210 I/O Module

The major functional blocks of the KN210 CPU module set are shown in Figure 1-3, and are described in the following paragraphs.

Figure 1-3 KN210 Block Diagram

1.2 R3000 RISC Processor

The central processor of the KN210 module set is the R3000 RISC processor chip. The R3000 chip implements two tightly coupled processors in a single VLSI chip. One processor is a 32-bit CPU, and the second is a system control processor (CP0).

The combined CPU/CP0 processors provide the following features:

- **32-bit operation**
- **A 5-stage pipeline**
- **On chip cache control**
- **On chip memory management**
- **Coprocessor interface**

1.3 Floating-Point Accelerator

The floating-point accelerator is implemented by the R3010 floating-point accelerator (FPA) chip. The R3010 FPA operates as a coprocessor for the R3000 processor and extends the R3000's instruction set to perform arithmetic operations on values in floating-point representations. The R3010 FPA interfaces with the R3000 processor to form a tightly-coupled unit with seamless integration of floating-point and fixed-point instruction sets.

1.4 Cache Memory

To maximize CPU performance, the KN210 module utilizes a system of cache memory. The cache memory is organized as two separate 64 Kbyte cache; one for instructions and the other for data.

1.5 Memory Controller

The main memory controller is implemented by a VLSI chip called the CMCTL. The CMCTL supports up to 64 Mbytes of ECC memory with a 500-550 ns cycle time for word transfers. This memory resides on up to four MS650 memory modules.

1.6 Diagnostic Processor

The diagnostic processor provides extensive diagnostic capabilities for the KN210 CPU module set.

1.7 MS650-BA Memory Modules

The MS650-BA memory modules are 16 Mbyte, 450 ns, 39-bit wide arrays (32-bit data and 7-bit ECC) implemented with 1 Mbyte dynamic RAMs in surface-mount packages. MS650-BA memory modules are single, quad-height, Q22-bus modules, as shown in Figure 1-4.

1.8 MS650-AA Memory Modules

The MS650-AA memory modules are 8 Mbyte, 450 ns, 39-bit wide arrays (32-bit data and 7-bit ECC) implemented with 256 Kbyte dynamic RAMs in zig-zag in-line packages (ZIPs). MS650-AA memory modules are single, quad-height, Q22-bus modules, as shown in Figure 1-4.

The MS650 modules communicate with the KN210 through the MS650 memory interconnect, which utilizes the CD rows of backplane slots 2 through 4, and a 50-pin ribbon cable. The KN210 memory subsystem supports a maximum of four memory modules.

Insert photo of both the MS650-AA and -BA memory modules here.

Figure 1-4 MS650 Memory Modules

1.9 DSSI Interface

The Digital small storage interconnect (DSSI) bus interface, located on the KN210 I/O module, is implemented by the SII chip and four 32K × 8 static RAMs. The DSSI interface allows the KN210 to transmit packets of data to, and receive packets of data from up to seven other DSSI devices. The KN210 system configurations contain one or more RF-series fixed disk devices, connected through the DSSI bus.

1.10 Ethernet Interface

The KN210 I/O module features an on-board network interface that is implemented through the LANCE chip and four $32K \times 8$ static RAMs. When used in conjunction with the H3602-SA CPU cover panel, this interface allows the KN210 CPU module set to be connected to either a ThinWire or standard Ethernet network.

The Ethernet interface includes registers for control and status reporting as well as a DMA controller, a 24 word transmit silo and a 24 word receive silo.

1.11 Q22-bus Interface

The Q22-bus interface is implemented by the CQBIC chip, located on the KN210 processor module. The CQBIC chip supports up to 16-word, block mode transfers between a Q22-bus DMA device and main memory, and up to 2-word, block mode transfers between the CPU and Q22-bus devices. The Q22-bus interface contains the following:

- A 16-entry map cache for the 8,192-entry, main memory-resident "scatter-gather" map, used for translating 22-bit Q22-bus addresses into 26-bit main memory addresses
- Interrupt arbitration logic that recognizes Q22-bus interrupt requests BR7-BR4
- Q22-bus termination (240Ω)

1.12 System Support Functions

System support functions are implemented by the system support chip (SSC), located on the KN210 processor module. The SSC provides console and boot code support functions, operating system support functions, timers, and many extra features, including the following:

- Word-wide ROM unpacking
- 1 Kbyte battery backed-up RAM
- Halt arbitration logic
- Console serial line
- Interval timer with 10 ms interrupts
- Time-of-year clock with support for battery back-up
- IORESET register

1-8 Overview

- Programmable CDAL bus timeout counter
- Two programmable timers similar in function to the VAX standard interval timer
- A register for controlling the diagnostic LEDs

1.13 Firmware

The firmware consists of 256 Kbytes of 16 bit-wide ROM, located on two 128K EPROMs (one for each processor). The firmware gains control when the processor halts, and contains programs that provide the following services:

- Board initialization
- Power-up self-testing of the KN210 processor and I/O modules, and the MS650 modules
- Emulation of a subset of the VAX standard console (automatic/manual bootstrap, automatic/manual restart, and a simple command language for examining/altering the state of the processor)
- Booting from supported Q22-bus devices, DSSI, or Ethernet
- Multilingual capability
- MOP support

1.14 Clock Functions

The clock chip provides the following functions:

- Generates two MOS clocks for the CPU and the main memory controller
- Generates three auxiliary clocks for other miscellaneous TTL logic
- Synchronizes reset signal for the CPU and the main memory controller
- Synchronizes data ready and data error signals for the CPU and the main memory controller

2

Installation and Configuration

2.1 Introduction

This chapter contains information required to install the KN210 CPU module set in a system. The following topics are discussed:

- Installing the KN210 CPU module set
- Configuring the KN210
- KN210 connectors
- H3602-SA CPU cover panel

2.2 Installing the KN210 CPU Module Set

The KN210 CPU module set can only be installed in BA213 backplanes.

The KN210 I/O and processor modules must be installed in slots 1 and 2, respectively, of the Q22/CD backplane (Figure 2–1). MS650 memory modules must be installed in slots immediately adjacent to the KN210 processor module. Up to four MS650 modules may be installed, occupying slots 3, 4, 5 and 6 respectively. A 50-pin ribbon cable is used to connect the KN210 processor and I/O modules to each other. Another 50-pin cable is used to connect the KN210 processor module to the MS650 memory module(s), as shown in Figure 2–2.

2-2 Installation and Configuration

Figure 2-1 Processor, I/O and Memory Module Placement

Figure 2-2 Cable Connections

2.3 Configuring the KN210

The following parameters must be configured on the KN210:

- Operation switch
- Function switch

2-4 Installation and Configuration

- Console serial line baud rate
- DSSI node ID
- Ethernet port connector selection

These parameters are configured using the H3602-SA CPU cover panel.

2.4 KN210 Connectors

The KN210 uses several connectors, and four rows of module fingers (A, B, C, and D) to communicate with the console device, main memory, the Q22-bus, the DSSI controller, and the Ethernet controller. The contact finger identification of the KN210 CPU and I/O modules is described in Appendix C.

The orientation of the connectors, and the LED indicators is shown in Figure 2-3.

Figure 2-3 KN210 Connector and LED Orientation

2.4.1 Console/Ethernet Connector

The Console/Ethernet connector (J1) is a 40-pin connector used to bus signals to/from the Ethernet interface and system console, and for configuration and display purposes. A multi-tapped 40-pin cable provides the physical connection between the KN210 processor and I/O boards and the H3602-SA CPU cover panel. Table 2-1 gives the pinouts for the console connector.

Table 2-1 Console/Ethernet Connector (J1) Pinouts

Pin	Mnemonic	Module	Meaning
01	XMIT- H	I/O	Transmit - output to the LAN interface
02	XMIT+ H	I/O	Transmit + output to the LAN interface
03	GND	I/O	Ground
04	GND	I/O	Ground
05	GND	I/O	Ground
06	RCV- H	I/O	Receive - input from the LAN interface
07	RCV+ H	I/O	Receive + input from the LAN interface
08	GND	I/O	Ground
09	GND	I/O	Ground
10	GND	I/O	Ground
11	COL- H	I/O	Collision
12	COL+ H	I/O	Collision
13	GND	I/O	Ground
14	GND	I/O	Ground
15	GND	I/O	Ground
16	GND	I/O	Ground
17	+12 V	I/O	Fused +12 Vdc
18	GND	I/O	Ground
19	DTR H	Processor	Data terminal ready
20	GND	Processor	Ground
21	TXD L	Processor	Transmit data
22	SPID0 L	—	Not used with H3602-SA
23	SPID1 L	—	Not used with H3602-SA

Table 2-1 (Cont.) Console/Ethernet Connector (J1) Pinouts

Pin	Mnemonic	Module	Meaning
24	RXD L	Processor	Receive data
25	RXD H	Processor	Receive data
26	SPID0 L	—	Not used with H3602-SA
27	+5 V	Processor	Fused +5 Vdc
28	CONBITRATE2 L	Processor	Console bit rate <02:00>. These three bits determine the console baud rate. They are configured using the select switch on the inside of the H3602-SA.
29	CONBITRATE1 L		
30	CONBITRATE0 L		
31	LEDCODE0 L	Processor	LED code register bits <03:00>.
32	LEDCODE1 L	Processor	When asserted each of these four output signals
33	LEDCODE2 L	Processor	Lights a corresponding LED on the module.
34	LEDCODE3 L	Processor	LEDCODE<03:00> are asserted (low) by power-up and by the negation of DCOK when the processor is halted. They are updated by boot and diagnostic programs from the BDR.
35	ENBHALT L	Processor	Halt enable. This input signal controls the response to an external halt condition. If ENBHALT is asserted (low), then the KN210 halts and enters the console program if any of the following occur: <ul style="list-style-type: none"> • The program executes a halt instruction in kernel mode. • The console detects a break character. • The Q22-bus halt line is asserted.

Table 2-1 (Cont.) Console/Ethernet Connector (J1) Pinouts

Pin	Mnemonic	Module	Meaning
			If ENBHALT is negated (high), then the halt line and break character are ignored and the ROM program responds to a halt instruction by restarting or rebooting the system. ENBHALT is read by software from the BDR.
36	BDCODE1 L	Processor	Boot and diagnostic code <01:00>. This 2-bit code indicates power-up mode, and is read by software from the BDR.
37	BDCODE0 L		
38	VBAT H	Processor	Battery backup voltage for the TODR clock.
39	CPUCODE1 L	Processor	CPU code <01:00>. This 2-bit code is read by software from the BDR. The configuration for the CPU code is as follows: 00 Normal operation 01 Reserved 10 Reserved 11 Reserved
40	CPUCODE0 L	Processor	

2.4.2 DSSI/I/O Connector

Connector J2 on the I/O module is a dual 50-pin connector (100 pins total) that provides two interfaces. The *lower* 50 pins are used as the interconnect between the I/O module and the processor module; the *upper* 50 pins are used for the DSSI controller.

2-8 Installation and Configuration

Table 2-3 lists the pinouts for the I/O interconnect portion of J2 (lower 50 pins). Table 2-2 lists the pinouts for the DSSI portion of J2 (upper 50 pins).

Table 2-2 DSSI Connector (I/O Module J2-Upper) Pinouts

Pin	Mnemonic	Pin	Mnemonic
51	DSSIDATA0 L	76	VTERM H
52	GND	77	VTERM H
53	DSSIDATA1 L	78	VTERM H
54	GND	79	GND
55	DSSIDATA2 L	80	Unused
56	GND	81	GND
57	DSSIDATA3 L	82	Unused
58	GND	83	GND
59	DSSIDATA4 L	84	Unused
60	GND	85	GND
61	DSSIDATA5 L	86	DSSIBSY L
62	GND	87	GND
63	DSSIDATA6 L	88	DSSIACK L
64	GND	89	GND
65	DSSIDATA7 L	90	DSSIRST L
66	GND	91	GND
67	DSSIPARITY L	92	Unused
68	GND	93	GND
69	Unused	94	DSSISEL L
70	GND	95	GND
71	Unused	96	DSSIC/D L
72	GND	97	GND
73	VTERM H	98	DSSIREQ L
74	VTERM H	99	GND
75	VTERM H	100	DSSII/O L

Table 2-3 I/O Connector (I/O Module J2-Lower) Pinouts

Pin	Signal	Pin	Signal
01	GND	26	BCDAL10 H
02	BCDAL9 H	27	GND
03	BCDAL8 H	28	BCDAL29 H
04	BCDAL7 H	29	BCDAL28 H
05	GND	30	BCDAL27 H
06	BCDAL6 H	31	GND
07	BCDAL5 H	32	BCDAL26 H
08	BCDAL4 H	33	BCDAL25 H
09	BCDAL3 H	34	BCDAL24 H
10	GND	35	BCDAL23 H
11	BCDAL2 H	36	GND
12	BCDAL1 H	37	BCDAL22 H
13	BCDAL0 H	38	BCDAL21 H
14	BCDAL19 H	39	BCDAL20 H
15	GND	40	BM3 L
16	BCDAL18 H	41	GND
17	BCDAL17 H	42	BM2 L
18	BCDAL16 H	43	BM1 L
19	BCDAL15 H	44	BM0 L
20	GND	45	Not Used
21	BCDAL14 H	46	GND
22	BCDAL13 H	47	Not Used
23	BCDAL12 H	48	IOPRESENT L
24	GND	49	BCDAL31 H
25	BCDAL11 H	50	BCDAL30 H

2.4.3 Memory Connector

Connector J2 on the processor module is a dual 50-pin connector (100 pins total) that provides two interfaces. The *lower* 50 pins are used as the memory interface between the KN210 module set and the MS650 memory modules; the *upper* 50 pins are used as the interconnect between the processor module and the I/O module.

2-10 Installation and Configuration

Table 2-4 lists the pinouts for the memory portion of J2 (lower 50 pins).
Table 2-3 lists the pinouts for the I/O interconnect portion of J2 (upper 50 pins).

Table 2-4 Memory Connector (Processor Module J2-Lower) Pinouts

Pin	Mnemonic	Pin	Mnemonic
01	GND	26	D MD10 H
02	D MD9 H	27	GND
03	D MD8 H	28	D MD29 H
04	D MD7 H	29	D MD28 H
05	GND	30	D MD27 H
06	D MD6 H	31	GND
07	D MD5 H	32	D MD26 H
08	D MD4 H	33	D MD25 H
09	D MD3 H	34	D MD24 H
10	GND	35	D MD23 H
11	D MD2 H	36	GND
12	D MD1 H	37	D MD22 H
13	D MD0 H	38	D MD21 H
14	D MD19 H	39	D MD20 H
15	GND	40	D MD38 H
16	D MD18 H	41	GND
17	D MD17 H	42	D MD37 H
18	D MD16 H	43	D MD36 H
19	D MD15 H	44	D MD35 H
20	GND	45	D MD34 H
21	D MD14 H	46	GND
22	D MD13 H	47	D MD33 H
23	D MD12 H	48	D MD32 H
24	GND	49	D MD31 H
25	D MD11 H	50	D MD30 H

2.5 H3602-SA CPU Cover Panel

The H3602-SA CPU cover panel is a special I/O panel that is used in BA213 enclosures. A multi-tapped 40-pin ribbon cable provides the physical connection between the H3602-SA CPU cover panel and the J1 connectors of both the KN210 processor and I/O modules. The H3602-SA fits over backplane slots 1 and 2, covering both the KN210 I/O and processor modules.

The H3602-SA CPU cover panel (Figure 2-4) includes the features and controls specified in Table 2-5.

Table 2-5 H3602-SA Features and Controls

Outside	Inside
Modified modular jack (MMJ) SLU connector	Baud rate rotary switch
Operation switch	Battery backup unit (BBU) for TODR clock
Hexadecimal LED display	40-pin cable connector
Function switch	List of baud rate switch settings
Standard/ThinWire Ethernet connectors	
Standard/ThinWire Ethernet selector	
Indicator LEDs	

2-12 Installation and Configuration

Figure 2-4 H3602-SA CPU Cover Panel

3

Architecture

This chapter describes KN210 registers, instruction set, and memory. The chapter covers the following KN210 topics:

- RISC processor
- Floating-point accelerator
- Cache memory
- Main memory system
- Console serial line
- Time-of-year clock and timers
- Boot and diagnostic facility
- Q22-bus interface
- Mass-storage interface
- Network interface

3.1 R3000 RISC Processor

The central processor of the KN210 is the R3000 RISC processor plus its associated R3010 floating-point unit.

3.1.1 Processor Features

The R3000 chip implements two tightly coupled processors in a single VLSI chip. One processor being a 32-bit CPU, and the second being a system control processor (CP0).

The combined CPU/CP0 processors provide these features.

- **32-bit operation** - The R3000 is a 32-bit machine with 32 32-bit registers, all addressing is 32-bits.
- **Pipelined** - The CPU contains a five stage pipeline capable of executing one instruction per 50 ns cycle.
- **On chip cache control** - The R3000 supports separate instruction and data caches of up to 256 Kbytes. (The KN210 supports 64 Kbytes each.) Both caches can be accessed in a single CPU cycle.
- **On chip memory management** - A four Gbyte virtual address space is mapped with a 64 entry fully associative translation lookaside buffer.
- **Coprocessor interface** - The R3000 provides a tightly coupled coprocessor interface for up to four coprocessors. The first CP0 is contained on the CPU chip and the second CP1 is the floating-point unit. CP2 and CP3 are unused in the KN210.
- **Write buffers** - The KN210 uses four R3020 4-word deep write buffers. All writes pass through these write buffers.

NOTE

Because all writes pass through the write buffers, software writing to I/O devices must wait for the write buffers to empty before reading from locations that could be affected by these I/O writes. This can be accomplished by conditionally branching on the coprocessor 3 condition. (BC3F branches if the WB is not empty and BC3T branches if the WB is empty.)

3.1.1.1 General Purpose Registers

The R3000 provides 32 32-bit general purpose registers, a 32-bit program counter, and 2 32-bit registers used in integer multiply and divide operations.

3.1.1.2 Instruction Set

R3000 instructions have three types, as shown in Figure 3-1. A summary of R3000 instructions is given in Table 3-1.

Figure 3-1 Instruction Set Types

Table 3-1 Instruction Summary

OP	Description	Group
LB	Load byte	Load/Store instructions
LBU	Load byte unsigned	
LH	Load halfword	
LHU	Load halfword unsigned	
LW	Load word	
LWL	Load word left	
LWR	Load word right	
SB	Store byte	
SH	Store halfword	
SW	Store word	
SWL	Store left	
SWR	Store right	
ADDI	Add immediate	Arithmetic instructions (ALU immediate)

Table 3-1 (Cont.) Instruction Summary

OP	Description	Group	
ADDIU	Add immediate unsigned	Arithmetic instructions (3-operand, register-type)	
SLTI	Set on less than immediate		
SLTIU	Set on less than immediate unsigned		
ANDI	AND immediate		
ORI	OR immediate		
XORI	Exclusive OR immediate		
LUI	Load upper immediate		
ADD	Add		
ADDU	Add unsigned		
SUB	Subtract		Shift instructions
SUBU	Subtract unsigned		
SLT	Set on less than		
SLTU	Set on less than unsigned		
AND	AND		
OR	OR		
XOR	Exclusive OR		
NOR	NOR		
SLL	Shift left logical		
SRL	Shift right logical		
SRA	Shift right arithmetic	Multiply/divide instructions	
SLLV	Shift left logical variable		
SRLV	Shift right logical variable		
SRAV	Shift right arithmetic variable		
MULT	Multiply		
MULTU	Multiply unsigned		
DIV	Divide		
DIVU	Divide unsigned		
MFHI	Move from HI		Jump and branch instructions
MTHI	Move to HI		
MFLO	Move from LO		
MTLO	Move to LO		
J	Jump		
JAL	Jump and link		
JR	Jump to register		

Table 3-1 (Cont.) Instruction Summary

OP	Description	Group
JALR	Jump and link register	
BEQ	Branch on equal	
BNE	Branch not equal	
BLEZ	Branch on less than or equal to zero	
BGTZ	Branch on greater than zero	
BLTZ	Branch on less than zero	
BGEZ	Branch on greater than or equal to zero	
BLTZAL	Branch on less than zero and link	
BGEZAL	Branch on greater than or equal to zero and link	
SYSCALL	System call	Special instructions
BREAK	Break	
LWCz	Load word from coprocessor	Coprocessor instructions
SWCz	Store word to coprocessor	
MTCz	Move to coprocessor	
MFCz	Move from coprocessor	
CTCZ	Move control to coprocessor	
CFCz	Move control from coprocessor	
COPz	Coprocessor operation	
BCzT	Branch on coprocessor z true	
BCzF	Branch on coprocessor z false	
MTCO	Move to CP0	System control coprocessor (CP0) instructions
MFCO	Move from CP0	
TLBR	Read indexed TLB entry	
TLBWI	Write indexed TLB entry	
TLBWR	Write random TLB entry	
TLBP	Probe TLB for matching entry	
RFE	Restore from exception	

3.1.2 Coprocessors

The R3000 can operate with up to four tightly coupled coprocessors, two of which are implemented in the KN210.

3.1.2.1 Coprocessor (0)

The system control coprocessor (CP0) is part of the R3000 chip itself. Its function is to support the virtual memory system and exception handling functions. It contains the translation lookaside buffer (TLB) plus the following registers:

- **EntryHi** - High half of a TLB entry.
- **EntryLo** - Low half of a TLB entry.
- **Index** - Programmable pointer into the TLB.
- **Random** - Pseudo-random pointer into the TLB.
- **Status** - Mode, interrupt enables, and diagnostic status information.
- **Cause** - Cause of last exception.
- **EPC** - Exception program counter.
- **Context** - Pointer into kernel's virtual page table entry array.
- **BadVA** - Most recent bad virtual address.
- **PRId** - Processor revision identification.

3.1.2.2 Coprocessor (1)

Coprocessor one (CP1) is the R3010 floating-point unit. The R3010 floating-point processor implements the IEEE arithmetic functions.

3.1.3 Memory Management

The R3000 provides for logical expansion of memory space by translating virtual address into physical addresses, 2 Gbytes for the kernel and 2 Gbytes for user, as shown in Figure 3-2 and Figure 3-3. Virtual addresses are mapped through the use of a translation lookaside buffer (TLB). This buffer contains 64 entries, each mapping a 4 Kbyte block. Controls are included for R/W accesses, cache or no-cache, and process identification.

Figure 3-2 R3000 Virtual to Physical Memory Map

3-8 Architecture

Figure 3-3 R3000 to I/O Memory Map

3.1.3.1 Operating Modes

The R3000 has two operating modes, kernel and user. All exceptions are handled in kernel mode. User mode is entered through the restore from exception (RFE) instruction.

User mode - 2 Gbytes of virtual address are available in this mode. Virtual addresses are extended with a 6-bit process identifier field allowing up to 64 user processes. All references in this mode are mapped through the TLB. Pages can be made cacheable or uncacheable through the use of bits within the TLB entries.

Kernel mode - contains four separate segments.

- **KUSEG** - references treated the same as user mode.
- **KSEG0** - references cached but not mapped.
- **KSEG1** - references not cached and not mapped.
- **KSEG2** - references mapped and cacheability determined through the TLB entries.

3.1.3.2 EntryHi and EntryLo Registers

These registers are used to read, write, or probe the TLB. During exceptions, these registers are loaded with information about the address that caused the exception. The format of these registers is the same as the format of a TLB entry. A TLB EntryHi register (EH) is shown in Figure 3-4. A TLB EntryLo register (EL) is shown in Figure 3-5.

Figure 3-4 TLB EntryHi Register

Data Bit	Definition
EntryHi<63:44>	(VPN) Virtual page number. Contains bits <31:12> of the virtual address.
EntryHi<43:38>	(PID) Process ID field. Allows multiple processes to share the TLB.
EntryHi<37:32>	EntryHi<5:0> Unused. Read as 0's.

Figure 3-5 TLB EntryLo Register

Data Bit	Definition
EntryLo<31:12>	(PFN) Page frame number. Contains bits <31:12> of the physical address.
EntryLo<11>	(N) Non-Cacheable. When set, the page mapped by this entry will not be cached.
EntryLo<10>	(D) Dirty. Setting this bit marks the page mapped by this entry as writeable.
EntryLo<9>	(V) Valid. This bit marks the page as valid, letting translation take place, otherwise a TLBL or TLBS miss will occur.
EntryLo<8>	(G) Global. When set, the R3000 will ignore the PID match requirement for translations through this entry.
EntryLo<7:0>	Unused. Read as 0's.

3.1.3.3 Index Register

The index register (IR) is used to index into the TLB when doing either TLB reads or TLB writes. The high-order bit returns status information for the probe instruction. The format for the index register is shown in Figure 3-6.

Figure 3-6 Index Register

Data Bit	Definition
Index<31>	(P) Probe failure. Set to (1) if the last probe instruction failed.
Index<30:14>	Unused. Read as 0's.
Index<13:8>	Indexes the TLB entry for TLB reads or writes.
Index<7:0>	Unused. Read as 0's.

3.1.3.4 Random Register

The random register (RR) is used when replacing TLB entries. The TLBWR instruction is used to write the TLB entry pointed to by this register. The first 8 TLB entries will never be written using the TLBWR instruction. The format for the random register is shown in Figure 3-7.

Figure 3-7 Random Register

Data Bit	Definition
Random<31:14>	Unused. Read as 0's.
Random<13:8>	Random index into the TLB ranging from 8 to 63.
Random<7:0>	Unused. Read as 0's.

3.1.4 Exception Handling Registers

Exceptions are handled through the use of six coprocessor zero (CP0) registers and one module specific interrupt status register. Software uses these registers during exception handling to determine the cause of the exception. These registers are described in the following sections.

3.1.4.1 Cause Register

The cause register (CR) is a 32-bit register used when servicing exceptions. It describes why the last exception was taken. All bits in this register are read only except for the SW bit. The format for the cause register is shown in Figure 3-8.

Figure 3-8 Cause Register

Data Bit	Definition
CR<31>	(BD) Branch delay. Set to 1 if last exception was taken while executing in a branch delay slot.
CR<30>	Unused. Read as zero.
CR<29:28>	CE<1:0> Coprocessor error. Indicates the unit number referenced when a coprocessor unusable exception is taken.
CR<27:16>	Unused. Read as 0's.

Data Bit	Definition																																													
CR<15:10>	IP<5:0> Interrupt pending. Indicates which of the external interrupts are pending. Table 3-2 shows the mapping of physical interrupt requests to the cause register IP<5:0> status bits. Figure 3-9 shows the flow of the interrupt acknowledge daisy chain.																																													
CR<9:8>	SW<1:0> Software interrupts. Indicates which software interrupts are pending. These bits are read/write.																																													
CR<7:6>	Unused. Read as 0's.																																													
CR<5:2>	(EC) Exception code. The exception code field is described in the following table.																																													
	<table border="1"> <thead> <tr> <th>Number</th> <th>Name</th> <th>Description</th> </tr> </thead> <tbody> <tr> <td>0 (10)</td> <td>INT</td> <td>External interrupt</td> </tr> <tr> <td>1</td> <td>MOD</td> <td>TLB modification exception</td> </tr> <tr> <td>2</td> <td>TLBL</td> <td>TLB miss exception (load or fetch)</td> </tr> <tr> <td>3</td> <td>TLBS</td> <td>TLB miss exception (store)</td> </tr> <tr> <td>4</td> <td>ADEL</td> <td>Address error exception (load or fetch)</td> </tr> <tr> <td>5</td> <td>ADES</td> <td>Address error exception (store)</td> </tr> <tr> <td>6</td> <td>IBE</td> <td>Bus error exception (fetch)</td> </tr> <tr> <td>7</td> <td>DBE</td> <td>Bus error exception (load or fetch)</td> </tr> <tr> <td>8</td> <td>SYS</td> <td>Syscall exception</td> </tr> <tr> <td>9</td> <td>BP</td> <td>Breakpoint exception</td> </tr> <tr> <td>10</td> <td>RI</td> <td>Reserved instruction exception</td> </tr> <tr> <td>11</td> <td>CPU</td> <td>Coprocessor unusable exception</td> </tr> <tr> <td>12</td> <td>OVF</td> <td>Arithmetic overflow exception</td> </tr> <tr> <td>13-15</td> <td>Reserved</td> <td></td> </tr> </tbody> </table>	Number	Name	Description	0 (10)	INT	External interrupt	1	MOD	TLB modification exception	2	TLBL	TLB miss exception (load or fetch)	3	TLBS	TLB miss exception (store)	4	ADEL	Address error exception (load or fetch)	5	ADES	Address error exception (store)	6	IBE	Bus error exception (fetch)	7	DBE	Bus error exception (load or fetch)	8	SYS	Syscall exception	9	BP	Breakpoint exception	10	RI	Reserved instruction exception	11	CPU	Coprocessor unusable exception	12	OVF	Arithmetic overflow exception	13-15	Reserved	
Number	Name	Description																																												
0 (10)	INT	External interrupt																																												
1	MOD	TLB modification exception																																												
2	TLBL	TLB miss exception (load or fetch)																																												
3	TLBS	TLB miss exception (store)																																												
4	ADEL	Address error exception (load or fetch)																																												
5	ADES	Address error exception (store)																																												
6	IBE	Bus error exception (fetch)																																												
7	DBE	Bus error exception (load or fetch)																																												
8	SYS	Syscall exception																																												
9	BP	Breakpoint exception																																												
10	RI	Reserved instruction exception																																												
11	CPU	Coprocessor unusable exception																																												
12	OVF	Arithmetic overflow exception																																												
13-15	Reserved																																													
CR<1:0>	Unused. Read as 0's.																																													

Table 3-2 R3000 Interrupt Mapping

Bit	Priority	Interrupt Request	Vector Register	Corresponding Diagnostic Processor IPL
IP<5>	2	FPU		
IP<4>	1	HALT		
IP<3>	3	PWRFL, MER1, MER0, WEAR		
IP<2>	4	100 Hz ->BIRQ7	VRR3	IPL17
IP<1>	5	DSSI ->NI ->BIRQ6	VRR2	IPL16
IP<0>	6	BIRQ5 ->Console ->Timers ->BIRQ4	VRR1 VRR0	IPL15 IPL14

Figure 3-9 Interrupt Acknowledge Daisy Chain

3.1.4.2 Exception Program Counter

The exception program counter (EPC) contains the virtual address of the instruction which caused the exception to be taken. If the instruction is in a branch delay slot, the EPC will contain the virtual address of the preceding branch or jump instruction.

3.1.4.3 Status Register

The status register (SR) is a 32-bit register containing various processor status. The format for the status register is shown in Figure 3–10.

Figure 3–10 Status Register

Data Bit	Definition
SR<31:28>	CU<3:0> Coprocessor usability. CU<3:0> controls the availability of the four possible coprocessors. Setting a CU bit to 1 enables the coprocessor.
SR<27:23>	Unused. Read as 0's.
SR<22>	(BEV) Bootstrap exception vector. When set to 1 causes the R3000 to use the alternate bootstrap vectors for UTLB miss and general exceptions.
SR<21>	(TS) Translation buffer shutdown. Set to 1 if the R3000 has disabled the translation buffer due to error.
SR<20>	(PE) Parity error. Set to 1 if a cache parity error occurs. Cleared by writing a 1.
SR<19>	(CM) Cache miss. Set to 1 if most recent D-cache load resulted in a miss when the D-cache is isolated.
SR<18>	(PZ) Parity zero. Setting to 1 forces parity bits to 0.
SR<17>	(SWC) Swap caches. Swaps I-cache and D-cache.

Data Bit	Definition
SR<16>	(ISC) Isolate cache. Isolates the D-cache from the memory system.
SR<15:10>	INTR<7:2> Interrupt mask. Setting these INTR bits to 1 enables their corresponding hardware interrupt.
SR<9:8>	INTR<1:0> Interrupt mask. Setting these INTR bits to 1 enables their corresponding software interrupt.
SR<7:6>	Unused. Read as 0's.
SR<5>	(KUO) Kernel/user mode, old. Set to 0 if kernel, 1 if user.
SR<4>	(IEO) Interrupt enable, old. Set to 1 to enable, 0 to disable.
SR<3>	(KUP) Kernel/user mode, previous. Set to 0 if kernel, 1 if user.
SR<2>	(IEP) Interrupt enable, previous. Set to 1 to enable, 0 to disable.
SR<1>	(KUC) Kernel/user mode, current. Set to 0 if kernel, 1 if user.
SR<0>	(IEC) Interrupt enable, current. Set to 1 to enable, 0 to disable.

3.1.4.4 BadVaddr Register

The BadVaddr register (BVA) saves the virtual address for any addressing exception.

3.1.4.5 Context Register

The context register (CR) is used by the UTLB miss handler. The CR saves some of the same information as the BadVaddr register. The format for the context register is shown in Figure 3-11.

Figure 3-11 Context Register

Data Bit	Definition
CR<31:21>	(PTEBase) Page table entry base. Holds the base for the page table entry.
CR<20:2>	(BadVPN) Bad virtual page number. Holds the failing virtual page number. Read only. Set by hardware.
CR<1:0>	Unused. Read as 0's.

3.1.4.6 Processor Revision Identifier Register

The processor revision identifier register (PRR) contains implementation and revision numbers for the R3000 chip. The format for the processor revision identification register is shown in Figure 3-12.

Figure 3-12 Processor Revision Identifier Register

Data Bit	Definition
PRR<31:16>	Unused. Read as 0's.
PRR<15:8>	IMP<7:0> Implementation identifier.
PRR<7:0>	REV<7:0> Revision identifier.

3.1.4.7 Interrupt Status Register

The interrupt status register, (R3000 address 1008 4000; diagnostic processor address 2008 4000) returns information concerning the four conditions which can cause an INTR3 interrupt. It can be accessed by either processor but only has meaning for the R3000. The diagnostic processor must make sure ISR<2:0> are 0 before transferring control to the R3000. The format for the interrupt status register is shown in Figure 3-13.

Figure 3-13 Interrupt Status Register

Data Bit	Definition
ISR<31:4>	Unused. Read as 1's.
ISR<3>	(HALT) Halt interrupt. Indicates a halt interrupt request is posted.
ISR<2>	(PRFL) Power fail interrupt. Indicates an impending power fail condition. Must be written zero (0) to enable subsequent power fail interrupts. Cleared on powerup.
ISR<1>	(MER1) Memory error one. Indicates a memory error interrupt has been generated either by the CQBIC or the CMCTL. Must be written as zero (0) to enable subsequent CQBIC or CMCTL memory error interrupts. Cleared on powerup.
ISR<0>	(MER0) Memory error zero. Indicates a write error has occurred and that the address contained in the WEAR register is valid. Must be written as 0 to enable subsequent MER0 interrupts. Cleared on powerup.

3.1.4.8 Vector Read Registers

The R3000 uses the vector read registers when servicing IRQ<2:0> interrupt requests. Reading any of these registers generates an interrupt acknowledge cycle on the peripheral bus (CP bus). Bits <15:0> of the data returned will be a unique vector from the device responding to the interrupt acknowledge. If no device responds, a **BUS ERROR EXCEPTION** will be taken by the R3000.

System software *must* guarantee interrupt requests from higher priority devices service before lower priority devices. This can be accomplished by reading the VRR register corresponding to the highest IRQ<2:0> pending in the cause/status register. The diagnostic processor has no access to these registers.

Register	IPL Generated
VRR0	14 (lowest priority)
VRR1	15
VRR2	16
VRR3	17 (highest priority)

NOTE

Because of the way interrupt requests are posted on the Q22-bus, interrupts at high priorities will also assert interrupt requests at lower levels. It is important that software read the VRR register corresponding to the highest IRQ pending in the CAUSE register. Lower level requests pending because of higher level interrupts will deassert at the same time the high level request deasserts.

Vector Read Register 0

Vector read register 0 (VRR0) cannot be accessed by the diagnostic processor. The format for VRR0 (R3000 address 1600 0050) is shown in Figure 3-14.

Figure 3-14 Vector Read Register 0

Vector Read Register 1

Vector read register 1 (VRR1) cannot be accessed by the diagnostic processor. The format for VRR1 (R3000 address 1600 0054) is shown in Figure 3-15.

Figure 3-15 Vector Read Register 1

Vector Read Register 2

Vector read register 2 (VRR2) cannot be accessed by the diagnostic processor. The format for VRR2 (R3000 address 1600 0058) is shown in Figure 3-16.

Figure 3-16 Vector Read Register 2

Vector Read Register 3

Vector read register 3 (VRR3) cannot be accessed by the diagnostic processor. The format for VRR3 (R3000 address 1600 005C) is shown in Figure 3-17.

Figure 3-17 Vector Read Register 3

3.1.5 Exceptions

The R3000 handles exceptions through the use of three different exception vectors. These exception vectors are the following:

- **BFC0 0000 virtual** - RESET exception vector.
- **8000 0000 virtual** - UTLB miss exception vector.
- **8000 0080 virtual** - General exception vector.

If the bootstrap exception vector (BEV) bit is set in the status register then the UTLB and general exception vector addresses will be changed, putting them in the ROM address space.

- **BFC0 0100 virtual** - UTLB miss exception vector.
- **BFC0 0180 virtual** - General exception vector.

3.1.5.1 General Exception Vector

The general exception vector handles the following exceptions:

- Address error
- Breakpoint
- Bus error
- Coprocessor unusable
- Interrupt
- Overflow
- Reserved instruction
- TLB miss
- TLB modified

3.1.5.2 Reset Exception Vector

The reset exception is taken on power up when the R3000 RESET signal is deasserted. The reset vector (BFC0 0000 virtual) resides in Kseg1 which is unmapped and uncached. It is also the first location in the ROM (physical 1FC0 0000).

When the reset exception occurs, the contents of all R3000 registers are undefined except for the following:

- The TS, SWc, KUc, and IEc bits of the status register are cleared.
- The BEV bit of the status register is set.
- The random register is cleared.

3.2 Floating-Point Accelerator

The KN210 floating-point accelerator (FPA) is implemented by a single VLSI chip called the R3010. The R3010 FPA operates as a coprocessor for the R3000 RISC processor and extends the R3000's instruction set to perform arithmetic operations on values in floating-point representations.

The R3010 FPA features 16 64-bit registers that can each be used to hold single-precision or double-precision values. The FPA also includes a 32-bit status/control register.

The 3010 FPA connects to the R3000 RISC processor to form a tightly-coupled unit with complete integration of floating-point and fixed-point instruction sets. Since each unit receives and executes instructions in parallel, some floating-point instructions can execute at the same single-cycle per instruction rate as fixed-point instructions.

3.2.1 Floating-Point Accelerator Instructions

Like the R3000 RISC processor, the R3010 FPA uses a load/store-oriented instruction set, with single-cycle loads and stores. Floating-point operations are started in a single cycle and their execution is overlapped with other fixed-point or floating-point operations.

The FPA monitors the R3000 processor instruction stream. If an instruction does not apply to the coprocessor, it is ignored. If an instruction does apply to the coprocessor, the FPA executes that instruction and transfers necessary result and exception data to the R3000 main processor.

3.3 Cache Memory

To maximize CPU performance, the KN210 utilizes two separate 64 Kbyte caches.

3.3.1 Cache Organization

The cache is organized as two separate caches; one for instructions and one for data. Each cache contains 16K entries (or 64 Kbyte). Both caches have the same organization; they are direct mapped, with a block size of one word (4-bytes). Fill size is either one word (4-bytes) or four words (16-bytes).

3.3.2 Cache Isolation

The data cache can be isolated from the memory system by setting the ISC bit in the status register. This allows software to flush a cache location in a single cycle without affecting the memory system.

3.3.3 Cache Swapping

To help in flushing the instruction cache, the caches can be swapped. This is done by setting the SWC bit in the status register. Precautions must be taken before setting this bit. The processor must be executing from an uncached region and must not have been executing loads or stores near to the time of doing the swap.

3.3.4 Cache Line Format

Each cache line is 60 bits long, and each cache can be accessed once per cycle. This is accomplished by alternately putting out instruction cache addresses and data cache addresses on each phase of the clock. These cache addresses are latched externally. The cache line format is shown in Figure 3–18. The tag field of the cache line format is shown in Figure 3–19.

Figure 3–18 Cache Line Format

Data Bit	Definition
<59:56>	(DataP) Holds the parity bits for the data field.
<55:53>	(TagP) Holds the parity bits for the tag field.
<52:32>	(TAG) Holds the valid bit and the page frame number.
<31:0>	(Data) Cached data.

Figure 3-19 Tag Field

Data Bit	Definition
<52>	(V) Holds the valid bit. The R3000 sets V whenever it writes a full word to the cache and clears V for any other writes. This allows software to flush a cache location by doing byte store operations.
<51:32>	(PFN) Page frame number.

3.4 Main Memory System

The KN210 includes a main memory controller implemented by a single VLSI chip called the CMCTL. The KN210 main memory controller communicates with the MS650 memory boards over the MS650 memory interconnect, which utilizes the CD interconnect for the address and control lines and a 50-pin, ribbon cable for the data lines. It supports up to four MS650 memory boards for a maximum of 64 Mbytes of ECC memory.

The memory controller can be configured for one or zero wait states under program control (using MEMCSR17<13>). For the KN210, MEMCSR17<13> must be set to one.

The controller supports synchronous word read references, and masked or unmasked synchronous write references generated by the CPU as well as synchronous twoword read references generated by cacheable CPU references that miss the cache. Table 3-3 gives CPU read reference timing. Table 3-4 gives CPU write reference timing.

Table 3-3 CPU Read Reference Timing

Data Type	CVAX Timing (ns)	R3000 Timing (ns)
Word	400	500-550
Twoword	600	No access
First word	400	No access
Second word	200	No access
Aborted reference	400	No access
Fourword	No access	1150-1200

Table 3-4 CPU Write Reference Timing

Data Type	CVAX Timing (ns)	R3000 Timing (ns)
Word	200	400-450
Word (masked)	500	700-750

The controller also supports asynchronous word and twoword DMA read references and masked and unmasked asynchronous word, twoword, threeword, and fourword DMA write references from the Q22-bus interface. Table 3-5 gives Q22-bus interface read reference timing. Table 3-6 gives Q22-bus interface write reference timing.

Table 3-5 Q22-bus Interface Read Reference Timing

Data Type	Timing (ns)
Word	500
Twoword	800
First word	500
Second word	300

Table 3-6 Q22-bus Interface Write Reference Timing

Data Type	Timing (ns)
Word	400
Word (masked)	600
Twoword	700
First word	400
Second word	300
Twoword (masked)	1100
First word	400
Second word	700
Threeword	1000
First word	400
Second word	300
Third word	300
Threeword (masked)	1400
First word	400
Second word	300
Third word	700
Fourword	1300
First word	400
Second word	300
Third word	300
Fourth word	300
Fourword (masked)	1700
First word	400
Second word	300
Third word	300
Fourth word	700

The timing in Table 3-6 assumes no exception conditions are encountered during the reference. Exception conditions will add the following amount of time if they are encountered during a reference:

Exception Condition	Time Added (ns)
Correctable error	100
Uncorrectable error	200 read
Uncorrectable error	100 write
CDAL parity error	100 write
Refresh collision	400

The main memory controller contains eighteen registers. Sixteen registers are used to configure each of the sixteen possible banks in main memory. One register is used to control the operating mode of all memory banks and one register captures state on main memory errors.

3.4.1 Main Memory Organization

Main memory is logically and physically divided into four boards that correspond to the four possible MS650 memory expansion modules that can be attached to a KN210. Each board can contain zero (no memory module present), 1 (as on the KN210), or 2 (MS650-AA present) memory bank(s). Each bank contains 1,048,576 (1M) words. Each word is divided into 4 data bytes and is stored with 7 ECC check bits, resulting in a memory array width of 39 bits.

3.4.2 Main Memory Addressing

The KN210 main memory controller is capable of controlling up to 16 banks of RAM, each bank containing 4 Mbytes of storage. Each bank of main memory has a programmable base address, determined by the state of bits <25:22> of the main memory configuration register associated with each bank.

A 4 Mbyte bank is accessed when bit <29> of the physical address is equal to 0, indicating a VAX memory space read/write reference, bits <28:26> of the physical address are equal to zero, indicating a reference within the range of the main memory controller, and the bank number of the bank matches bits <25:22> of the physical address. The remainder of the physical address (bits <21:2>) are used to determine the row and column of the desired word within the bank. The byte mask lines are ignored on read operations, but are used to select the proper byte(s) within a word during masked word write references.

3.4.3 Main Memory Behavior on Writes

On unmasked CPU write references, the main memory controller operates in dump and run mode, terminating the CDAL bus transaction after latching the data, but before checking CDAL bus parity, calculating the ECC check bits, and transferring the data to main memory.

On unmasked DMA write references by the Q22-bus interface: the data is latched; CDAL bus parity is NOT checked; the CDAL bus transaction is terminated; the ECC check bits are calculated; and the data is transferred to main memory.

On single masked CPU or DMA write references: the referenced word is read from main memory; the ECC code checked; the check bits recalculated to account for the new data byte(s); the CDAL transaction is terminated; and the word is rewritten.

On multiple transfer masked DMA writes, each word write is acknowledged, then the CDAL transaction is terminated.

3.4.4 Main Memory Error Status Register

The main memory status register (MEMCSR16) (R3000 address 1008 0140₁₆ diagnostic processor address 2008 0140₁₆) is used to capture main memory error data. The format for MEMCSR16 is shown in Figure 3-20.

MA-1112-87

Figure 3-20 Format for MEMCSR16

Data Bit	Definition
MEMCSR16<31>	RDS error. Read/Write to clear. When set, an uncorrectable ECC error occurred during a memory read or masked write reference. Cleared by writing a 1 to it. Writing a 0 has no effect. Undefined if MEMCSR16<7> (CDAL bus error) is set. Cleared on power-up and the negation of DCOK when the processor is halted.
MEMCSR16<30>	RDS high error rate. Read/Write to clear. When set, an uncorrectable ECC error occurred while the RDS error log request bit was set, indicating multiple uncorrectable memory errors. Cleared by writing a 1 to it. Writing a 0 has no effect. Undefined if MEMCSR16<7> (CDAL bus error) is set. Cleared on power-up and the negation of DCOK when the processor is halted.
MEMCSR16<29>	CRD error. Read/Write to clear. When set, a correctable (single bit) error occurred during a memory read or masked write reference. Cleared by writing a 1 to it. Writing a 0 has no effect. Undefined if MEMCSR16<7> (CDAL bus error) is set. Cleared by writing a 1, on power-up and the negation of DCOK when the processor is halted.
MEMCSR16<28:9>	Page address of error. Read only. This field identifies the page (512 byte block) containing the location that caused the memory error. In the event of multiple memory errors, the types of errors are prioritized and the page address of the error with the highest priority is captured. Errors with equal priority do not overwrite previous contents. Writes have no effect. Cleared on power-up and the negation of DCOK when the processor is halted. The types of error conditions follow in order of priority: <ol style="list-style-type: none"> <li data-bbox="594 1444 1255 1499">1. CDAL bus parity errors during a CPU write reference, as logged by the CDAL bus error bit. <li data-bbox="594 1520 1255 1604">2. Uncorrectable ECC errors during a CPU or DMA read or masked write reference, as logged by the RDS error log bit. <li data-bbox="594 1625 1255 1680">3. Correctable ECC errors during a CPU or DMA read or masked write reference, as logged by CRD error bit.

Data Bit	Definition
MEMCSR16<8>	DMA error. Read/Write to clear. When set, an error occurred during a DMA read or write reference. Cleared by writing a 1 to it. Writing a 0 has no effect. Cleared on power-up and the negation of DCOK when the processor is halted.
MEMCSR16<7>	CDAL bus error. Read/Write to clear. When set, a CDAL bus parity error occurred on a CPU write reference. Cleared by writing a 1 to it. Writing a 0 has no effect. Cleared on power-up and the negation of DCOK when the processor is halted.
MEMCSR16<6:0>	Error syndrome. Read only. This field stores the error syndrome. A non-zero syndrome indicates a detectable error has occurred. A unique syndrome is generated for each possible single bit (correctable) error. A list of these syndromes and their associated single bit errors is given in Table 3-7. Any non-zero syndrome that is not contained in Table 3-7 indicates a multiple bit (uncorrectable) error has occurred. This field handles multiple errors in the same manner as MEMCSR16<28:9>. Cleared on power-up and the negation of DCOK when the processor is halted.

Table 3-7 Error Syndromes

Syndrome <6:0>	Bit Position in Error
0000000	No error detected Data bits (0-32 decimal)
1011000	0
0011100	1
0011010	2
1011110	3
0011111	4
1011011	5
1011101	6
0011001	7
1101000	8
0101100	9

Table 3-7 (Cont.) Error Syndromes

Syndrome <6:0>	Bit Position in Error
0101010	10
1101110	11
0101111	12
1101011	13
1101101	14
0101001	15
1110000	16
0110100	17
0110010	18
1110110	19
0110111	20
1110011	21
1110101	22
0110001	23
0111000	24
1111100	25
1111010	26
0111110	27
1111111	28
0111011	29
0111101	30
1111001	31
	Check bits (32-38 decimal)
0000001	32
0000010	33
0000100	34
0001000	35
0010000	36
0100000	37
1000000	38
0000111	Result of incorrect check bits written on detection of a CDAL parity error.
All others	Multi-bit errors

3.4.5 Main Memory Control and Diagnostic Status Register

The main memory control and diagnostic status register (MEMCSR17) (R3000 address 1008 0144₁₆; diagnostic processor address 2008 0144₁₆) is used to control the operating mode of the main memory controller as well as to store diagnostic status information. This register is unique to CPU designs that use the CMCTL memory controller chip (Figure 3-21).

MA-1122-87

Figure 3-21 Format for MEMCSR17

Data Bit	Definition
MEMCSR17<31:13>	Unused. This field reads as zero and must be written as zero.
MEMCSR17<12>	CRD interrupt enable. Read/Write. When cleared, single-bit errors are corrected by the ECC logic, but no interrupt is generated. When set, single-bit errors are corrected by the ECC logic and they cause a diagnostic processor interrupt to be generated at IPL 1A with a vector of 54 ₁₆ . This bit has no effect on the capturing of error information in MEMCSR16, or on the reporting of uncorrectable errors. Cleared on power-up and the negation of DCOK when the processor is halted.

Data Bit	Definition
MEMCSR17 <11>	Force refresh request. Read/Write. When cleared, the refresh control logic operates in normal mode (refresh every 11.3 μ s). When set, one memory refresh operation occurs immediately after the MEMCSR write reference that set this bit. Setting this bit provides a mechanism for speeding up the testing of the refresh logic during manufacturing test of the controller chip. This bit is cleared by the memory controller upon completion of the refresh operation. Cleared on power-up and the negation of DCOK when the processor is halted.
MEMCSR17 <10>	Memory error detect disable. Read/Write. When set, error detection and correction (ECC) is disabled, so all memory errors go undetected. When cleared, error detection, correction, state capture and reporting (through MEMCSR16) is enabled. Cleared on power-up and the negation of DCOK when the processor is halted.
MEMCSR17 <9:8>	Unused. This field reads as zero and must be written as zero.
MEMCSR17 <7>	Diagnostic check mode. Read/Write. When set, the contents of MEMCSR17 <6:0> are written into the 7 ECC check bits of the location (even if a CDAL parity error is detected) during a memory write reference. When cleared, the 7 check bits calculated by the ECC generation logic are loaded into the 7 ECC check bits of the location during a write reference and a memory read reference will load the state of the 7 ECC check bits of the location that was read into MEMCSR17 <6:0>. Cleared on power up and the negation of DCOK when the processor is halted.

NOTE

Diagnostic check mode is restricted to unmasked memory write references. No masked write references are allowed when diagnostic check mode is enabled.

Data Bit	Definition
MEMCSR17 <6:0>	Check bits. Read/Write. When the diagnostic check mode bit is set, these bits are substituted for the check bits that are generated by the ECC generation logic during a write reference. When the diagnostic check mode bit is cleared, memory read references load the state of the 7 ECC check bits of the location that was read into MEMCSR16 <6:0>. Cleared on power-up and the negation of DCOK when the processor is halted.

3.4.6 Main Memory Error Detection and Correction

The KN210 main memory controller generates CDAL bus parity on CPU read references. The actions taken following the detection of a CDAL bus parity error depend on the type of write reference.

For unmasked diagnostic processor write references, incorrect check bits are written to main memory (potentially masking an as yet undetected memory error) along with the data. An interrupt is generated at IPL 1D through vector

60₁₆ on the next cycle and MCSR16 <7> is set. The incorrect check bits are determined by calculating the seven correct check bits, and complementing the three least significant bits.

For masked diagnostic processor write references: incorrect check bits are written to main memory (potentially masking an as yet undetected memory error) along with the data, unless an uncorrectable error is detected during the read portion; MEMCSR16 <7> is set; and a machine check abort is initiated. If an uncorrectable error is detected on the read portion, no write operation takes place. The incorrect check bits are determined by calculating the seven correct check bits, and complementing the three least significant bits.

The memory controller protects main memory by using a 32-bit modified Hamming code to encode the 32-bit data word with seven check bits. This allows the controller to detect and correct single-bit errors in the data field and detect single bit errors in the check bit field and double-bit errors in the data field. The most likely causes of these errors are failures in either the memory array or the 50-pin cable.

Upon detecting a correctable error on a read reference or the read portion of a masked write reference, the data is corrected (if it is in the data field). Before placing it on the CDAL bus, or back in main memory, a diagnostic processor interrupt is generated at IPL 1A through vector 54₁₆ (IRQ 3 interrupt for the R3000), bit <29> of MEMCSR16 is set, bits <28:9> of MEMCSR16 are loaded with the address of the page containing the location that caused the error, and bits <6:0> are loaded with the error syndrome which indicates which bit was in error. If the error was detected on a DMA reference, MEMCSR16 <8> is also set.

NOTE

The corrected data is not rewritten to main memory, so the single bit error will remain there until rewritten by software.

Upon detecting an uncorrectable error, the action depends on the type of reference being performed.

On a demand read reference: the affected row of the diagnostic processor cache is invalidated; bit <31> of MEMCSR16 is set; bits <28:9> of MEMCSR16 are loaded with the address of the page containing the location that caused the error; and bits <6:0> are loaded with the error syndrome which indicates that the error was uncorrectable and a machine check abort is initiated. If the read was a local-miss, global-hit read, or a read of the Q22-bus map, MEMCSR16 <8> and DSER <4> are also set, and DEAR <12:0> are loaded with the address of the page containing the location that caused the error.

On a request read reference: the diagnostic processor prefetch or fill cycle is aborted, but no machine check occurs; bit <31> of MEMCSR16 is set; bits <28:9> of MEMCSR16 are loaded with the address of the page containing the location that caused the error; and bits <6:0> are loaded with the error syndrome which indicates that the error was uncorrectable.

On the read portion of masked write reference: bit <31> of MEMCSR16 is set; bits <28:9> of MEMCSR16 are loaded with the address of the page containing the location that caused the error; and bits <6:0> are loaded with the error syndrome which indicates that the error was uncorrectable and a machine check abort is initiated on the diagnostic processor. An IRQ 3 interrupt is generated on the R3000 side.

On a DMA read reference: bit <31> and bit <8> of MEMCSR16 are set; bits <28:9> of MEMCSR16 are loaded with the address of the page containing the location that caused the error; bits <6:0> are loaded with the error syndrome which indicates that the error was uncorrectable; DSER <4> is set; DEAR <12:0> are loaded with the address of the page containing the location that caused the error; BDAL <17:16> are asserted on the Q22-bus along with the data to notify the receiving device (unless it was a map read by the Q22-bus interface during translation); and a diagnostic processor interrupt is generated at IPL 1D through vector 60₁₆. An IRQ 3 interrupt is generated on the R3000 side.

On a DMA masked write reference: bit <31> and bit <8> of MEMCSR16 are set; bits <28:9> of MEMCSR16 are loaded with the address of the page containing the location that caused the error; bits <6:0> are loaded with the error syndrome which indicates that the error was uncorrectable; DSER <4> is set, DEAR <12:0> are loaded with the address of the page containing the location that caused the error; IPCR <15> is set to notify the initiating device; and a diagnostic processor interrupt is generated at IPL 1D through vector 60₁₆. An IRQ 3 interrupt is generated on the R3000 side.

3.5 Console Serial Line

The console serial line provides the KN210 processor with a full duplex, RS-423 EIA, serial line interface, which is also RS-232C compatible. The only data format supported is 8-bit data with no parity and one stop bit. Four registers control the operation of the console serial line.

3.5.1 Console Registers

There are four registers associated with the console serial line unit. They are implemented in the SSC (Table 3-8).

Table 3-8 Console Registers

Register Name	Mnemonic
Console receiver control/status	RXCS
Console receiver data buffer	RXDB
Console transmit control/status	TXCS
Console transmit data buffer	TXDB

3.5.1.1 Console Receiver Control/Status Register

The console receiver control/status register (RXCS) is used to control and report the status of incoming data on the console serial line (Figure 3-22).

MA-1118-87

Figure 3-22 Console Receiver Control/Status Register

Data Bit	Definition
RXCS <31:8>	Unused. Read as zeros. Writes have no effect.
RXCS <7>	(RX DONE) Receiver done. Read only. Writes have no effect. This bit is set when an entire character has been received and is ready to be read from the RXDB register. This bit is automatically cleared when RXDB is read. It is also cleared on power-up and the negation of DCOK when the processor is halted.
RXCS <6>	(RX IE) Receiver interrupt enable. Read/Write. When set, this bit causes an interrupt to be requested at IPL 14 with an SCB offset of F8 if RX DONE is set. When cleared, interrupts from the console receiver are disabled. This bit is cleared on power-up and the negation of DCOK when the processor is halted.
RXCS <5:0>	Unused. Read as zeros. Writes have no effect.

3.5.1.2 Console Receiver Data Buffer

The console receiver data buffer (RXDB) is used to buffer incoming data on the serial line and capture error information (Figure 3-23).

MA-1119-87

Figure 3-23 Console Receiver Data Buffer

Data Bit	Definition
RXDB <31:16>	Unused. Always read as zero. Writes have no effect.
RXDB <15>	(ERR) Error. Read only. Writes have no effect. This bit is set if RBUF <14> or <13> is set. It is clear if these two bits are clear. This bit cannot generate a program interrupt. Cleared on power-up and the negation of DCOK when the processor is halted.
RXDB <14>	(OVR ERR) Overrun error. Read only. Writes have no effect. This bit is set if a previously received character was not read before being overwritten by the present character. Cleared by reading the RXDB, on power-up and the negation of DCOK when the processor is halted.
RXDB <13>	(FRM ERR) Framing error. Read only. Writes have no effect. This bit is set if the present character did not have a valid stop bit. Cleared by reading the RXDB, on power-up and the negation of DCOK when the processor is halted.
	NOTE Error conditions remain present until the next character is received, at which point, the error bits are updated.
RXDB <12>	Unused. This bit always reads as 0. Writes have no effect.

Data Bit	Definition
RXDB <11>	(RCV BRK) Received break. Read only. Writes have no effect. This bit is set at the end of a received character for which the serial data input remained in the space condition for 20 bit times. Cleared by reading the RXDB, on power-up and the negation of DCOK when the processor is halted.
RXDB <10:8>	Unused. These bits always read as 0. Writes have no effect.
RXDB <7:0>	Received data bits. Read only. Writes have no effect. These bits contain the last received character.

3.5.1.3 Console Transmitter Control/Status Register

The console transmitter control/status register (TXCS) is used to control and report the status of outgoing data on the console serial line (Figure 3–24).

MA-1120-87

Figure 3–24 Console Transmitter Control/Status Register

Data Bit	Definition
TXCS <31:8>	Unused. Read as zeros. Writes have no effect.
TXCS <7>	(TX RDY) Transmitter ready. Read only. Writes have no effect. This bit is cleared when TXDB is loaded and set when TXDB can receive another character. This bit is set on power-up and the negation of DCOK when the processor is halted.

Data Bit	Definition
TXCS <6>	(TX IE) Transmitter interrupt enable. Read/Write. When set, this bit causes an interrupt to be requested at IPL 14 with an SCB offset of FC if TX RDY is set. When cleared, interrupts from the console receiver are disabled. This bit is cleared on power-up and the negation of DCOK when the processor is halted.
TXCS <5:3>	Unused. Read as zeros. Writes have no effect.
TXCS <2>	(MAINT) Maintenance. Read/Write. This bit is used to facilitate a maintenance self-test. When MAINT is set, the external serial input is set to MARK and the serial output is used as the serial input. This bit is cleared on power-up and the negation of DCOK when the processor is halted.
TXCS<1>	Unused. Read as zero. Writes have no effect.
TXCS<0>	(XMIT BRK) Transmit break. Read/Write. When this bit is set, the serial output is forced to the space condition after the character in TXB<7:0> is sent. While XMIT BRK is set, the transmitter will operate normally, but the output line will remain low. Thus, software can transmit dummy characters to time the break. This bit is cleared on power-up and the negation of DCOK when the processor is halted.

3.5.1.4 Console Transmitter Data Buffer

The console transmitter data buffer (TXDB) is used to buffer outgoing data on the serial line (Figure 3-25).

MA-1123-87

Figure 3-25 Console Transmitter Data Buffer

Data Bit	Definition
TXDB<31:8>	Unused. Writes have no effect.
TXDB<7:0>	Transmitted data bits. Write only. These bits are used to load the character to be transmitted on the console serial line.

3.5.2 Break Response

The console serial line unit recognizes a break condition which consists of 20 consecutively received space bits. If the console detects a valid break condition, the RCV BRK bit is set in the RXDB register. If the break was the result of 20 consecutively received space bits, the FRM ERR bit is also set. If halts are enabled (ENBHALT asserted on the 40-pin connector), the KN210 will halt and transfer program control to diagnostic processor ROM location 2004 0000 when the RCV BRK bit is set if the diagnostic processor is running. If the R3000 is running an IRQ 4 interrupt is posted. RCV BRK is cleared by reading RXDB. Another mark followed by 20 consecutive space bits must be received to set RCV BRK again.

3.5.3 Baud Rate

The receive and transmit baud rates are always identical and are controlled by the SSC configuration register bits <14:12>.

The user selects the desired baud rate through pins <30:28> on the 40-pin system support connector (CONBITRATE <02:00>), which are configured using the select switch on the inside of the H3602-SA. The KN210 firmware reads this code from boot and diagnostic register bits <6:4> and loads it into SSC configuration register bits <14:12>. Operating systems will not cause the baud rate to be transferred. The baud rate is only set on power up.

Table 3-9 shows the baud rate select signal voltage levels (H or L), the corresponding INVERTED code as read in the boot and diagnostic register bits <6:4>, and the code that should be loaded into SSC configuration register bits <14:12>.

Table 3-9 Baud Rate Select

Baud Rate	CONBITRATE <2:0>	BDR <6:4>	SSC <14:12>
300	HHH	000	000
600	HHL	001	001
1200	HLH	010	010
2400	HLL	011	011
4800	LHH	100	100
9600	LHL	101	101
19200	LLH	110	110
38400	LLL	111	111

3.5.4 Console Interrupt Specifications

The console serial line receiver and transmitter both generate interrupts at IPL 14. The receiver interrupts with a vector of F8₁₆, while the transmitter interrupts with a vector of FC₁₆. See Section 3.1.4.

3.6 Time-of-Year Clock and Timers

The KN210 clocks include time-of-year clock (TODR), a subset interval clock (subset ICCS), and two additional programmable timers modeled after the VAX standard interval clock, plus a 100 Hz interval timer used as the R3000 interval clock.

3.6.1 R3000 Interval Timer Register

The interval timer register (ITR) provides a 100 Hz interval timer for the R3000. The format for the ITR (R3000 address 1008 4010; diagnostic processor address 2008 4010) is shown in Figure 3-26.

Figure 3-26 Interval Timer Register

Data Bit	Definition
ITR<31:8>	Unused. Read as ones, writing has no effect.
ITR<7>	(IS) Interrupt status. Read only. Status bit, shows when a 100 Hz interrupt is pending. Cleared by reading VRR2 for the R3000 and by an IPL16 interrupt acknowledge for the diagnostic processor.
ITR<6>	(IE) Interrupt enable. Read/Write. This bit enables and disables the ITR interval timer interrupts. When this bit is set, an interval timer interrupt is requested every 10 msec with an error of less than .01%. When this bit is clear, interval timer interrupts are disabled. This bit is cleared on power-up and during an I/O reset.
ITR<5>	Unused. Read as zeros, writing has no effect.
ITR<4>	(IOP) I/O present. Read only. Set to 1 if the I/O module is not present.
ITR<3:0>	Unused. Read as ones, writing has no effect.

Interval timer requests are posted at IPL 17 with a vector of C0: the interval timer is the highest priority device at this IPL for the R3000 (Section 3.1.4).

NOTE

The vector read back for the ITR interrupt when reading VRR2 comes from the R3000 ROM if the R3000 is running and from the diagnostic processor ROM if the diagnostic processor is running. ROM location 5C must be programmed with the vector 2400 00C0 (hex) for proper operation.

3.6.2 Time-of-Year Clock

The KN210 time-of-year clock (TODR) forms an unsigned 32-bit binary counter that is driven from a 100 Hz oscillator, so that the least significant bit of the clock represents a resolution of 10 ms, with less than .0025% error. The register counts only when it contains a non-zero value. This register is implemented in the SSC (Figure 3-27).

MA-1124-87

Figure 3-27 Time-of-Year Clock

The TODR (R3000 address 1014 006C; diagnostic processor address 2014 006C) is maintained during power failure by battery backup circuitry which interfaces, through the external connector, to a set of batteries which are mounted on the H3602-SA. The TODR will remain valid for greater than 162 hours when using the NiCad battery pack (three batteries in series).

The SSC configuration register contains a battery low (BLO) bit which, if set after initialization, the TODR is cleared, and will remain at zero until software writes a non-zero value into it.

NOTE

After writing a non-zero value into the TODR, software should clear the BLO bit by writing a 1 to it.

3.6.3 Interval Timer

The KN210 interval timer (ICCS), internal processor register 24, is implemented according to the *VAX Architecture Reference Manual* for subset processors. The interval clock control/status register (ICCS) is implemented as the standard subset of the standard VAX ICCS, while NICR and ICR are not implemented (Figure 3-28).

Figure 3-28 Interval Timer

Data Bit	Definition
ICCS<31:7>	Unused. Read as zeros, must be written as zeros.
ICCS<6>	(IE) Interrupt enable. Read/Write. This bit enables and disables the interval timer interrupts. When the bit is set, an interval timer interrupt is requested every 10 ms with an error of less than .01%. When the bit is clear, interval timer interrupts are disabled. This bit is cleared on power-up and the negation of DCOK when the processor is halted.
ICCS<5:0>	Unused. Read as zeros, must be written as zeros.

Interval timer requests are posted at IPL 16 with a vector of C0: the interval timer is the highest priority device at this IPL. See Section 3.1.4.

3.6.4 Programmable Timers

The KN210 features two programmable timers. They are accessed as I/O space registers and a control bit has been added which stops the timer upon overflow. If so enabled, the timers will interrupt at IPL 14 upon overflow. The interrupt vectors are programmable and are set to 78 and 7C by the firmware.

Each timer is composed of four registers: a timer **n** control register, a timer **n** interval register, a timer **n** next interval register, and a timer **n** interrupt vector register, where **n** represents the timer number (0 or 1).

3.6.4.1 Timer Control Registers

The KN210 has two timer control registers, one for controlling timer 0 (TCR0), and one for controlling timer 1 (TCR1). TCR0 is accessible at R3000 address

1014 0100₁₆ and diagnostic processor address 2014 0100₁₆. TCR1 is accessible at R3000 address 1014 0110₁₆ and diagnostic processor 2014 0110₁₆. These registers are implemented in the SSC (Figure 3–29).

MA-1125-87

Figure 3-29 Timer Control Registers

Data Bit	Definition
TCRn<31>	(ERR) Error. Read/Write to clear. This bit is set whenever the timer interval register overflows and INT is already set. Thus, the ERR indicates a missed overflow. Writing a 1 to this bit clears it. Cleared on power-up and the negation of DCOK when the processor is halted.
TCRn<30:8>	Unused. Read as zeros, must be written as zeros.
TCRn<7>	(INT) Read/Write to clear. This bit is set whenever the timer interval register overflows. If IE is set when INT is set, an interrupt is posted at IPL 14. Writing a 1 to this bit clears it. Cleared on power-up and the negation of DCOK when the processor is halted.
TCRn<6>	(IE) Read/Write. When this bit is set, the timer will interrupt at IPL 14 when the INT bit is set. Cleared on power-up and the negation of DCOK when the processor is halted.
TCRn<5>	(SGL) Read/Write. Setting this bit causes the timer interval register to be incremented by 1 if the RUN bit is cleared. If the RUN bit is set, then writes to the SGL bit are ignored. This bit always reads as 0. Cleared on power-up and the negation of DCOK when the processor is halted.

Data Bit	Definition
TCRn<4>	(XFR) Read/Write. Setting this bit causes the timer next interval register to be copied into the timer interval register. This bit is always read as 0. Cleared on power-up and the negation of DCOK when the processor is halted.
TCRn<3>	Unused. Read as zeros, must be written as zeros.
TCRn<2>	(STP) Read/Write. This bit determines whether the timer stops after an overflow when the RUN bit is set. If the STP bit is set at overflow, the RUN bit is cleared by the hardware at overflow and counting stops. Cleared on power-up and the negation of DCOK when the processor is halted.
TCRn<1>	Unused. Read as zeros, must be written as zeros.
TCRn<0>	(RUN) Read/Write. When set, the timer interval register is incremented once every microsecond. The INT bit is set when the timer overflows. If the STP bit is set at overflow, the RUN bit is cleared by the hardware at overflow and counting stops. When the RUN bit is clear, the timer interval register is not incremented automatically. Cleared on power-up and the negation of DCOK when the processor is halted.

3.6.4.2 Timer Interval Registers

The KN210 has two timer interval registers, one for timer 0 (TIR0), and one for timer 1 (TIR1). TIR0 is accessible at R3000 address 1014 0104₁₆; diagnostic processor address 2014 0104₁₆. TIR1 is accessible at diagnostic processor address 2014 0114₁₆; R3000 address 1014 0114₁₆.

The timer interval register is a read only register containing the interval count. When the run bit is 0, writing a 1 increments the register. When the RUN bit is 1, the register is incremented once every microsecond. When the counter overflows, the INT bit is set, and an interrupt is posted at IPL 14 if the IE bit is set. Then, if the RUN and STP bits are both set, the RUN bit is cleared and counting stops. Otherwise, the counter is reloaded. The maximum delay that can be specified is approximately 1.2 hours. This register is cleared on power-up and the negation of DCOK when the processor is halted.

3.6.4.3 Timer Next Interval Registers

The KN210 has two timer next interval registers, one for timer 0 (TNIR0), and one for timer 1 (TNIR1). TNIR0 is accessible at diagnostic processor address 2014 0108₁₆; R3000 address 1014 0108₁₆. TNIR1 is accessible at diagnostic processor address 2014 0118₁₆; R3000 address 1014 0118₁₆. These registers are implemented in the SSC.

This read/write register contains the value which is written into the timer interval register after overflow, or in response to a 1 written to the XFR bit. This register is cleared on power-up and the negation of DCOK when the processor is halted.

3.6.4.4 Timer Interrupt Vector Registers

The KN210 has two timer interrupt vector registers, one for timer 0 (TIVR0), and one for timer 1 (TIVR1). TIVR0 is accessible at R3000 address 1014 010C₁₆ and diagnostic processor address 2014 010C₁₆. TIVR1 is accessible at R3000 address 1014 011C₁₆ and diagnostic processor 2014 011C₁₆. These registers are implemented in the SSC and are set to 78 and 7C respectively by the resident firmware.

This read/write register contains the timer's interrupt vector. Bits <31:10> and <1:0> are read as 0 and must be written as 0. When TCRn<6> (IE) and TCRn<7> (INT) transition to 1, an interrupt is posted at IPL 14. When a timer's interrupt is acknowledged, the content of the interrupt vector register is passed to the CPU, and the INT bit is cleared. Interrupt requests can also be cleared by clearing either the IE or the INT bit. This register is cleared on power-up and the negation of DCOK when the processor is halted (Figure 3-30).

MA-1126-87

Figure 3-30 Timer Interrupt Vector Register

NOTE

Both timers interrupt at the same IPL (IPL 14) as the console serial line unit. When multiple interrupts are pending, the console serial line has priority over the timers, and timer 0 has priority over timer 1. Refer to Section 3.1.4.

3.7 Boot and Diagnostic Facility

The KN210 boot and diagnostic facility features four registers, two 40-pin ROM sockets with 128 Kbytes of read-only memory per processor, and 1 Kbyte of battery backed up RAM. The ROM and battery backed up RAM may be accessed by longword, word, halfword, or byte references.

The KN210 CPU module populates the ROM sockets with 256 Kbytes of ROM (or EPROM). This ROM contains the KN210 resident firmware. One 16-bit wide ROM is dedicated to each processor. If this ROM is replaced for special applications, the new ROM must initialize and configure the board, provide HALT and console emulation, as well as provide boot and diagnostic functionality.

3.7.1 Boot and Diagnostic Register

The boot and diagnostic register (BDR) is a byte-wide register located at the R3000 physical address $1008\ 4004_{16}$ (diagnostic processor physical address $2008\ 4004_{16}$). It can be accessed by either processor, but not by external Q22-bus devices. The BDR allows either processor to read various KN210 configuration bits. Only the low byte of the BDR should be accessed, bits $\langle 31:8 \rangle$ are undefined (Figure 3-31).

Figure 3-31 Boot and Diagnostic Register

Data Bit	Definition
BDR<31:8>	Undefined. Should not be read or written.
BDR<7>	(ENBHALT) Halt enable. Read only. Writes have no effect. This bit reflects the state of pin 35 (ENBHALT L) of the 40-pin connector. The assertion of this signal enables the halting of the CPU upon detection of a console break condition. On a power-up, the KN210 resident firmware reads the ENBHALT bit to decide whether to enter the console emulation program (ENBHALT set) or to boot the operating system (ENBHALT clear).
BDR<6:4>	(CONBITRATE) Console bit rate <02:00>. Read only. Writes have no effect. These three bits originate from pins <30:28> of the 40-pin connector. They reflect the setting of the baud rate select switch on the H3602-SA. These bits are read only on power up.

BDR<6:4>	Baud Rate
000	300
001	600
010	1200
011	2400
100	4800
101	9600
110	19200
111	38400

BDR<3:2>	(CPUCODE) CPU code <01:00>. Read only. Writes have no effect. These two bits originate from pins <40:39> of the 40-pin connector.
----------	---

CPUCODE <01:00>	Configuration
00	Power Up
01	Reserved
10	Reserved
11	Reserved

Data Bit	Definition
BDR<1:0>	(BDCODE) Boot and diagnostic code <01:00>. Read only. Writes have no effect. This 2-bit code reflects the status of the 40-pin connector (J1) pins <37:36>. The KN210 ROM programs use BDCODE <01:00> to determine the power up mode.

3.7.2 Diagnostic LED Register

The diagnostic LED register (DLEDR), R3000 address 1014 0030₁₆; diagnostic processor address 2014 0030₁₆, is implemented in the SSC and contains four read/write bits that control the external LED display. A 0 in a bit lights the corresponding LED; all four bits are cleared on power-up and the negation of DCOK when SCR<7> is clear to provide a power-up lamp test (Figure 3-32).

MA-X1447-87

Figure 3-32 Diagnostic LED Register

Data Bit	Definition
DLEDR<31:4>	Unused. Read as zeros, must be written as zeros.
DLEDR<3:0>	(DSPL) Display <3:0>. Read/Write. These four bits update an external LED display. Writing a 0 to a bit lights the corresponding LED. Writing a 1 to a bit turns its LED off. The display bits are cleared (all LEDs are lit) on power-up and the negation of DCOK when SCR<7> is clear.

3.7.3 ROM Memory

The KN210 supports 256 Kbytes of ROM memory for storing code (128 Kbytes per processor) for board initialization, board self-tests, and boot code. ROM memory may be accessed through byte, halfword, word and longword references. CDAL bus parity is neither checked nor generated on ROM references.

NOTE

The ROM size must be set in the SSC configuration register before attempting references outside the first 8 Kbyte block of the local ROM space.

3.7.3.1 ROM Socket

The KN210 provides two 40-pin ROM sockets to accept two 64K × 16-bit EPROMs or ROMs.

3.7.3.2 ROM Address Space

Two 64 Kbytes by 16-bit ROMs reside on the KN210 CPU module. One ROM is dedicated to each of the processors and cannot be accessed by the other.

The R3000 ROM address space goes from 1FC0 0000 to 1FC1 FFFF.

Writes to this address space by the R3000 will result in a Memerr Interrupt INT<3>.

The diagnostic processor ROM is located at 2004 0000 to 2005 FFFF in halt mode and 2006 0000 to 2007 FFFF in run mode. Note that the run mode ROM space reads exactly the same ROM code as the halt mode ROM space.

Writes to either of these address spaces by the diagnostic processor will result in a machine check.

Power-Up Modes

The boot and diagnostic ROM programs use boot and diagnostic code <1:0> to determine the power-up modes as described in Section 3.7.1.

3.7.4 Battery Backed-Up RAM

The KN210 contains 1 Kbyte of battery backed-up static RAM, for use as a console "scratchpad." The power for the RAM is provided through pins 38 (VBAT H) and 20, 18, 16-13, 10-8, 6-3 (GND) of the 40-pin connector.

The RAM is organized as a 256×32 -bit array. The array appears in a 1 Kbyte block of R3000 address space at 1014 0400-1014 07FF and in the diagnostic processor I/O page at addresses 2014 0400-2014 07FF.

This array is not protected by parity, and CDAL bus parity is neither checked nor generated on reads or writes to this RAM.

3.7.5 KN210 Initialization

There are three kinds of hardware initialization: power-up initialization, processor initialization, and I/O bus initialization.

3.7.5.1 Power-Up Initialization

Power-up initialization is the result of the restoration of power and includes a hardware reset, a processor initialization and an I/O bus initialization.

3.7.5.2 Hardware Reset

A KN210 hardware reset occurs on power-up, the negation of DCOK when SCR<7> is clear, on the assertion of BINIT on the Q22-bus. A hardware reset causes the hardware halt procedure to be initiated with a diagnostic processor halt code of 03. It also initializes some IPRs and most I/O page registers to a known state. The effect a hardware reset has on I/O space registers is documented in the description of the register.

3.7.5.3 I/O Bus Initialization

An I/O bus initialization occurs on power-up, the negation of DCOK, or as the result of an R3000 write to the IORESET register, or of a diagnostic processor MTPR to IPR 55 (IORESET), or a diagnostic processor write to the IORESET register (2014 00DC). An I/O bus initialization clears the IPCR and DSER, and causes the Q22-bus interface to acquire both the CDAL and Q22-bus, then assert the Q22-bus BINIT signal.

I/O Bus Reset Register

The I/O bus reset register (IORESET) is implemented in the SSC.

3.7.5.4 Processor Initialization

A processor initialization occurs on power-up, the negation of DCOK, the assertion of BINIT as the result of a console INITIALIZE command, and after a halt caused by a diagnostic processor error condition.

KN210 firmware must configure main memory, the local I/O page, and the Q22-bus map during a processor initialization.

3.7.6 Select Processor Register

The select processor register (SPR) determines which processor will control the memory and I/O subsystems. This register is write only, reading will produce unpredictable results. The format for the select processor register, (diagnostic processor address 2008 4008; R3000 address 1008 4008), is shown in Figure 3-33.

Figure 3-33 Select Processor Register

Data Bit	Definition
SPR<31>	(SP) Select processor. Write ONLY. When set to one (1) enables R3000 operation. When clear (0) enables diagnostic processor operation. The R3000 is disabled by stalling memory operations while the SP bit is clear. The diagnostic processor is disabled by requesting a DMG and holding the bus while the SP bit is set.
SPR<30:0>	Undefined. Reading will produce unpredictable results.

3.7.7 Write Error Address Register

The write error address register (WEAR) is a 32-bit register, located at 1700 0000, that latches the bus address for the first R3000 write cycle producing a bus error. The causes for this bus error on writes are accessing non-existent memory, doing a masked write to a memory location with incorrect ECC check bits, or any write to the CQBIC that results in an error. Any subsequent write which causes a bus error will not change the WEAR value until the MER0 bit has been cleared in the ISR register. The WEAR cannot be accessed by the diagnostic processor, and is undefined on powerup.

3.8 Q22-bus Interface

The KN210 includes a Q22-bus interface implemented through a single VLSI chip called the CQBIC. It contains a CDAL bus to Q22-bus interface that supports the following functions:

- A programmable mapping function (scatter-gather map) for translating 22-bit, Q22-bus addresses into 29-bit CDAL bus addresses that allows any page in the Q22-bus memory space to be mapped to any page in main memory.
- A direct mapping function for translating 29-bit CDAL addresses in the local Q22-bus address space and local Q22-bus I/O page into 22-bit, Q22-bus addresses.
- Masked and unmasked word reads and writes from the CPU to the Q22-bus memory and I/O space and the Q22-bus interface registers. Word reads and writes of the local Q22-bus memory space are buffered and translated into two-word, block mode, transfers on the Q22-bus. Word reads and writes of the local Q22-bus I/O space are buffered and translated into two, single-word transfers on the Q22-bus.
- Up to sixteen-word, block mode, writes from the Q22-bus to main memory. These words are buffered then transferred to main memory using two asynchronous DMA fourword transfers. For block mode writes of less than sixteen words, the words are buffered and transferred to main memory using the most efficient combination of fourword, twoword, and word asynchronous DMA transfers.

The maximum write bandwidth for block mode references is 3.3 Mbytes per second. Block mode reads of main memory from the Q22-bus cause the Q22-bus interface to perform an asynchronous DMA twoword read of main memory and buffer all four words, so that on block mode reads, the next three words of the block mode read can be delivered without any additional CDAL bus cycles. The maximum read bandwidth for Q22-bus block mode references is 2.4 Mbytes per second. Q22-bus burst mode DMA transfers result in single-word reads and writes of main memory.

- Transfers from the CPU to the local Q22-bus memory space, that result in the Q22-bus map translating the address back into main memory (local-miss, global-hit transactions).

The Q22-bus interface contains several registers for Q22-bus control and configuration, and error reporting.

The interface also contains Q22-bus interrupt arbitration logic that recognizes Q22-bus interrupt requests BR7-BR4 and translates them into CPU interrupts at levels 17-14.

The Q22-bus interface detects Q22-bus *no sack* timeouts, Q22-bus interrupt acknowledge timeouts, Q22-bus non-existent memory timeouts, main memory errors on DMA accesses from the Q22-bus and Q22-bus parity errors.

3.8.1 Q22-bus to Main Memory Address Translation

On DMA references to main memory, the 22-bit, Q22-bus address must be translated into a 29-bit main memory address. This translation process is performed by the Q22-bus interface by using the Q22-bus map. This map contains 8192 mapping registers, (one for each page in the Q22-bus memory space), each of which can map a page (512 bytes) of the Q22-bus memory address space into any of the 128K pages in main memory. Since local I/O space addresses cannot be mapped to Q22-bus pages, the local I/O page is inaccessible to devices on the Q22-bus.

Q22-bus addresses are translated to main memory addresses as shown in Figure 3-34.

MA-1145-87

Figure 3-34 Q22-bus to Main Memory Address Translation

At power up time, the Q22-bus map registers, including the valid bits, are undefined. External access to main memory is disabled as long as the interprocessor communication register LM EAE bit is cleared. The Q22-bus interface monitors each Q22-bus cycle and responds if the following three conditions are met:

1. The interprocessor communication register LM EAE bit is set.
2. The valid bit of the selected mapping register is set.
3. During read operations, the mapping register must map into existent main memory, or a Q22-bus timeout occurs. (During write operations, the Q22-bus interface returns Q22-bus BRPLY before checking for existent local memory; the response depends only on conditions 1 and 2 above.)

NOTE

In the case of local-miss, global-hit, the state of the LM EAE bit is ignored.

If the map cache does not contain the needed Q22-bus map register, then the Q22-bus interface will perform an asynchronous DMA read of the Q22-bus map register before proceeding with the Q22-bus DMA transfer.

3.8.1.1 Q22-bus Map Registers

The Q22-bus map contains 8192 registers (QMRs) that control the mapping of Q22-bus addresses into main memory. Each register maps a page of the Q22-bus memory space into a page of main memory (Table 3-10). These registers are implemented in a 32 Kbyte block of main memory, but are accessed through the CQBIC chip through a block of addresses in the I/O page.

The local I/O space address of each register was chosen so that register address bits <14:2> are identical to Q22-bus address bits <21:9> of the Q22-bus page which the register maps. The Q22-bus map registers (QMRs) have the format shown in Figure 3-35.

MA-X1450-87

Figure 3-35 Q22-bus Map Registers

Data Bit	Definition
QMR<31>	(V) Valid. Read/Write. When a Q22-bus map register is selected by bits <21:9> of the Q22-bus address, the valid bit determines whether mapping is enabled for that Q22-bus page. If the valid bit is set, the mapping is enabled, and Q22-bus addresses within the page controlled by the register are mapped into the main memory page determined by bits <28:9>. If the valid bit is clear, the mapping register is disabled, and the Q22-bus interface does not respond to addresses within that page. This bit is <i>undefined</i> on power-up and the negation of DCOK when the processor is halted.
QMR<30:20>	Unused. These bits always read as zero and must be written as zero.
QMR<19:0>	(A28-A9) Address bits <28:9>. Read/Write. When a Q22-bus map register is selected by a Q22-bus address, and if that register's valid bit is set, then these 20 bits are used as main memory address bits <28:9>. Q22-bus address bits <8:0> are used as main memory address bits <8:0>. These bits are <i>undefined</i> on power-up and the negation of DCOK when the processor is halted.

Table 3-10 Q22-bus Map

Register Address	Q22-bus Addresses Mapped (Hex)	Q22-bus Addresses Mapped (Octal)
1008 8000	00 0000 - 00 01FF	00 000 000 - 00 000 777
1008 8004	00 0200 - 00 03FF	00 001 000 - 00 001 777
1008 8008	00 0400 - 00 05FF	00 002 000 - 00 002 777
1008 800C	00 0600 - 00 07FF	00 003 000 - 00 003 777
1008 8010	00 0800 - 00 09FF	00 004 000 - 00 004 777
1008 8014	00 0A00 - 00 0BFF	00 005 000 - 00 005 777
1008 8018	00 0C00 - 00 0DFF	00 006 000 - 00 006 777
1008 801C	00 0E00 - 00 0FFF	00 007 000 - 00 007 777
.	.	.
.	.	.
.	.	.
1008 FFF0	3F F800 - 3F F9FF	17 774 000 - 17 774 777
1008 FFF4	3F FA00 - 3F FBFF	17 775 000 - 17 775 777
1008 FFF8	3F FC00 - 3F FDFD	17 776 000 - 17 776 777
1008 FFFC	3F FE00 - 3F FFFF	17 776 000 - 17 777 777

3.8.1.2 Accessing the Q22-bus Map Registers

Although the CPU accesses the Q22-bus map registers through aligned, masked word references, to the local I/O page (R3000 addresses 1008 8000₁₆ through 1008 FFFC₁₆; diagnostic processor addresses 2008 8000₁₆ through 2008 FFFC₁₆), the map actually resides in a 32 Kbyte block of main memory. The starting address of this block is controlled by the contents of the Q22-bus map base register. The Q22-bus interface also contains a 16-entry, fully associative, Q22-bus map cache to reduce the number of main memory accesses required for address translation.

NOTE

The system software must protect the pages of memory that contain the Q22-bus map from direct accesses that will corrupt the map or cause the entries in the Q22-bus map cache to become stale. Either of these conditions will result in the incorrect operation of the mapping function.

When the CPU accesses the Q22-bus map through the local I/O page addresses, the Q22-bus interface reads or writes the map in main memory. The Q22-bus interface does not have to gain Q22-bus mastership when accessing the Q22-bus map. Since these addresses are in the local I/O space, they are not accessible from the Q22-bus.

On a Q22-bus map read by the CPU, the Q22-bus interface decodes the local I/O space address (2008 8000 - 2008 FFFC for the diagnostic processor; 1008 8000 - 1008 FFFC for the R3000). If the register is in the Q22-bus map cache, the Q22-bus interface will internally resolve any conflicts between CPU and Q22-bus transactions (if both are attempting to access the Q22-bus map cache entries at the same time), then return the data. If the map register is not in the map cache, the Q22-bus interface will force the CPU to retry, acquire the CDAL bus, perform an asynchronous DMA read of the map register. On completion of the read, the CPU is provided with the data when its read operation is retried. A map read by the CPU does not cause the register that was read to be stored in the map cache.

On a Q22-bus map write by the CPU, the Q22-bus interface latches the data, then on the completion of the CPU write, acquires the CDAL bus and performs an asynchronous DMA write to the map register. If the map register is in the Q22-bus map cache, then the Cam Valid bit for that entry will be cleared to prevent the entry from becoming stale. A Q22-bus map write by the CPU does not update any cached copies of the Q22-bus map register.

3.8.1.3 Q22-bus Map Cache

To speed up the process of translating Q22-bus address to main memory addresses, the Q22-bus interface utilizes a fully associative, sixteen entry, Q22-bus map cache, which is implemented in the CQBIC chip.

If a DMA transfer ends on a page boundary, the Q22-bus interface will prefetch the mapping register required to translate the next page and load it into the cache, before starting a new DMA transfer. This allows Q22-bus block mode DMA transfers that cross page boundaries to proceed without delay. The replacement algorithm for updating the Q22-bus map cache is FIFO.

The cached copy of the Q22-bus map register is used for the address translation process. If the required map entry for a Q22-bus address (as determined by bits <21:9> of the Q22-bus address) is not in the map cache, then the Q22-bus interface uses the contents of the map base register to access main memory and retrieve the required entry. After obtaining the entry from main memory, the valid bit is checked. If it is set, the entry is stored in the cache and the Q22-bus cycle continues.

The format of a Q22-bus map cache entry is as shown in Figure 3-36.

MA-X1451-87

Figure 3-36 Q22-bus Map Cache Entry

Data Bit	Definition
CQMR<33>	(Cam Valid). When a mapping register is selected by a Q22-bus address, the Cam Valid bit determines whether the cached copy of the mapping register for that address is valid. If the Cam Valid bit is set, the mapping register is enabled, and addresses within that page can be mapped. If the Cam Valid bit is clear, the Q22-bus interface must read the map in local memory to determine if the mapping register is enabled. This bit is cleared on power-up, the negation of DCOK when the processor is halted, by setting the QMCIA (Q22-bus map cache invalidate all) bit in the interprocessor communication register, on writes to the IORESET register, by a write to the Q22-bus map base register, or by writing to the QMR that is being cached.
CQMR<32:20>	(QBUS ADR). These bits contain the Q22-bus address bits <21:9> of the page that this entry maps. This is the content addressable field of the 16 entry cache for determining if the map register for a particular Q22-bus address is in the map cache. These bits are <i>undefined</i> on power-up.
CQMR<19:0>	(Address bits A28-A9). When a mapping register is selected by a Q22-bus address, and if that register's Cam Valid bit is set, then these 20 bits are used as main memory address bits 28 through 9. Q22-bus address bits 8 through 0 are used as local memory address bits 8 through 0. These bits are <i>undefined</i> on power-up.

3.8.2 CDAL Bus to Q22-bus Address Translation

CDAL bus addresses within the local Q22-bus I/O space, addresses (2000 0000 - 2000 1FFF)₁₆ for the diagnostic processor; 1000 0000 - 1000 1FFF₁₆ for the R3000), are translated into Q22-bus I/O space addresses by using bits <12:0> of the CDAL address as bits <12:0> of the Q22-bus address and asserting BBS7. Q22-bus address bits <21:13> are driven as zeros.

CDAL bus addresses within the local Q22-bus memory space, addresses (3000 0000 - 303F FFFF)₁₆ for the diagnostic processor; 1400 0000 - 143F FFFF₁₆ for the R3000), are translated into Q22-bus memory space addresses by using bits <21:0> of the CDAL address as bits <21:0> of the Q22-bus address.

3.8.3 Interprocessor Communication Register

The interprocessor communication register (IPCR), diagnostic processor address 2000 1F40₁₆; R3000 address 1000 1F40₁₆, is a 16-bit register which resides in the Q22-bus I/O page address space and can be accessed by any device which can become Q22-bus master (including the KN210 itself). The IPCR, implemented in the CQBIC chip, is byte accessible, meaning that a write byte instruction can write to either the low or high byte without affecting the other byte.

The IPCR also appears at Q22-bus address 17 777 500 (Figure 3-37).

MA-X1452-87

Figure 3-37 The Interprocessor Communication Register

Data Bit	Definition
IPCR<15>	(DMA QME) DMA Q22-bus address space memory error. Read/Write to clear. This bit indicates that an error occurred when a Q22-bus device was attempting to read main memory. It is set if DMA system error register bit DSER<4> (main memory error) is set, or the CDAL bus timer expires. The main memory error bit indicates that an uncorrectable error occurred when an external device (or CPU) was accessing the KN210 local memory. The CDAL bus timer expiring indicates that the memory controller did not respond when the Q22-bus interface initiated a DMA transfer. This bit is cleared by writing a 1 to it, on power-up, by the negation of DCOK when the processor is halted, by writes to the IORESET register, and whenever DSER<4> is cleared.
IPCR<14>	(QM CIA) Q22-bus invalidate all. Write only. Writing a 1 to this bit clears the Cam Valid bits in the cached copy of the map. This bit always reads as zero. Writing a 0 has no effect.
IPCR<13:9>	(Unused) Read as zeros. Must be written as zeros.
IPCR<8>	Reserved for Digital use.
IPCR<7>	Unused. Read as zero. Must be written as zero.
IPCR<6>	Reserved for Digital use.
IPCR<5>	(LM EAE) Local memory external access enable. Read/Write when the KN210 is Q22-bus master. Read only when another device is Q22-bus master. When set, this bit enables external access to local memory (through the Q22-bus map). Cleared on power-up and by the negation of DCOK when the processor is halted.
IPCR<4:1>	Unused. Read as zeros. Must be written as zeros.
IPCR<0>	Reserved for Digital use.

3.8.4 Q22-bus Interrupt Handling

The KN210 responds to interrupt requests BR7-4 with the standard Q22-bus interrupt acknowledge protocol (DIN followed by IAK). The console serial line unit, the programmable timers, and the interprocessor doorbell request interrupts at IPL 14 and have priority over all Q22-bus BR4 interrupt requests.

3.8.5 Configuring the Q22-bus Map

The KN210 implements the Q22-bus map in an 8K word (32 Kbyte) block of main memory. This map must be configured by the KN210 firmware during a processor initialization by writing the base address of the uppermost 32 Kbyte block of good main memory into the Q22-bus map base register. The base of this map must be located on a 32 Kbyte boundary.

NOTE

This 32 Kbyte block of main memory must be protected by the system software. The only access to the map should be through local I/O page addresses (2008 8000 - 2008 FFFC₁₆ for the diagnostic processor; 1008 8000 - 1008 FFFC₁₆ for the R3000).

3.8.5.1 Q22-bus Map Base Address Register

The Q22-bus map base address register (QBMBR), diagnostic processor address 2008 0010₁₆; R3000 address 1008 0010₁₆, controls the main memory location of the 32 Kbyte block of Q22-bus map registers.

This read/write register is accessible by the CPU on a word boundary only. Bits <31:29,14:0> are unused and should be written as zero and will return zero when read.

A write to the map base register will flush the Q22-bus map cache by clearing the Cam Valid bits in all the entries.

The contents of this register are *undefined* on power up and the negation of DCOK when the processor is halted, and are not affected by BINIT being asserted on the Q22-bus (Figure 3-38).

MA-X1453-87

Figure 3-38 Q22-bus Map Base Address Register

3.8.6 System Configuration Register

The system configuration register (SCR), diagnostic processor address 2008 0000₁₆; R3000 address 1008 0000₁₆, contains a BHALT enable bit and a power OK flag.

The system configuration register (SCR) is word, halfword, and byte accessible. Programmable option fields are cleared on power-up and by the negation of DCOK when the processor is halted. The format of the SCR register is shown in Figure 3–39.

MA-X1454-87

Figure 3–39 System Configuration Register

Data Bit	Definition
SCR<31:16>	Unused. Read as zero. Must be written as zero.
SCR<15>	(POK) Power OK. Read only. Writes have no effect. This bit is set if the Q22-bus BPOK signal is asserted and clear if it is negated. This bit is cleared on power-up and by the negation of DCOK when the processor is halted.
SCR<14>	(BHALT EN) BHALT enable. Read/Write. This bit controls the effect the Q22-bus BHALT signal has on the CPU. When set, asserting the Q22-bus BHALT signal will halt the CPU and assert DSER<15>. When cleared, the Q22-bus BHALT signal will have no effect. This bit is cleared on power-up and by the negation of DCOK when the processor is halted.
SCR<13:11>	Unused. Read as zero. Must be written as zero.
SCR<10>	Reserved for Digital use.

Data Bit	Definition
SCR<9:8>	Unused. Read as zero. Must be written as zero.
SCR<7>	(ACTION ON DCOK NEGATION) Read/Write. When cleared, the Q22-bus interface asserts SYSRESET (causing a hardware reset of the board and control to be passed to the resident firmware through the hardware halt procedure with a halt code of 3) when DCOK is negated on the Q22-bus. When set, the Q22-bus interface asserts HALTIN (causing control to be passed to the resident firmware through the hardware halt procedure with a halt code of 2) when DCOK is negated on the Q22-bus. Cleared on power-up and the negation of DCOK when the processor is halted.
SCR<6:4>	Unused. Read as zero. Must be written as zero.
SCR<3:1>	Reserved for Digital use.
SCR<0>	Unused. Read as 0. Must be written as 0.

3.8.7 DMA System Error Register

The DMA system error register (DSER), diagnostic processor address 2008 0004₁₆; R3000 address 1008 0004₁₆, is one of three registers associated with Q22-bus interface error reporting.

The DMA system error register is implemented in the CQBIC chip, and logs main memory errors on DMA transfers, Q22-bus parity errors, Q22-bus non-existent memory errors, and Q22-bus no-grant errors. The Q22-bus error address register contains the address of the page in Q22-bus space which caused a parity error during an access by the local processor.

The DMA error address register contains the address of the page in local memory which caused a memory error during an access by an external device or the processor during a local miss global hit transaction. An access by the local processor which the Q22-bus interface maps into main memory will provide error status to the processor when the processor does a RETRY for a READ local miss-global hit, or by an interrupt in the case of a local-miss global-hit write.

The DSER is a word, halfword, or byte accessible read/write register available to the local processor. The bits in this register are cleared to 0 on power-up, by the negation of DCOK when the processor is halted, and by writes to the IORESET register. All bits are set to 1 to record the occurrence of an event. They are cleared by writing a 1, writing zeros has no effect (Figure 3-40).

MA-X1455-87

Figure 3-40 DMA System Error Register

Data Bit	Definition
DSER<31:16>	Unused. Read as 0. Must be written as 0.
DSER<15>	Q22-bus BHALT detected. Read/Write to clear. Set when the Q22-bus interface detects that the Q22-bus BHALT line was asserted and SCR<14> (BHALT ENABLE) is set. Cleared by writing a 1, writes to the IORESET register, on power-up and the negation of DCOK when the processor is halted.
DSER<14>	Q22-bus DCOK negation detected. Read/Write to clear. Set when the Q22-bus interface detects the negation of DCOK on the Q22-bus and SCR<7> (action on DCOK negation) is set. Cleared by writing a 1, writes to the IORESET register, on power-up and the negation of DCOK when the processor is halted.
DSER<13:8>	Unused. Read as zero. Must be written as zero.
DSER<7>	Master DMA NXM. Read/Write to clear. This bit is set when the CPU performs a demand Q22-bus read cycle or write cycle that does not reply after 10 μ s. During interrupt acknowledge cycles, or request read cycles, this bit is not set. It is cleared by writing a 1, on power-up, by the negation of DCOK when the processor is halted and by writes to the IORESET register.

Data Bit	Definition
DSER<6>	Unused. Read as zero. Must be written as zero.
DSER<5>	Q22-bus parity error. Read/Write to clear. This bit is set when the CPU performs a Q22-bus demand read cycle which returns a parity error. During interrupt acknowledge cycles or request read cycles, this bit is not set. It is cleared by writing a 1, on power-up, by the negation of DCOK when the processor is halted and by writes to the IORESET register.
DSER<4>	Main memory error. Read/Write to clear. This bit is set if an external Q22-bus device or local miss global hit receives a memory error while reading local memory. The IPCR<15> reports the memory error to the external Q22-bus device. It is cleared by writing a 1, on power-up, by the negation of DCOK when the processor is halted and by writes to the IORESET register.
DSER<3>	Lost error. Read/Write to clear. This bit indicates that an error address has been lost because of DSER<7,5,4,0> having been previously set and a subsequent error of either type occurs which would have normally captured an address and set either DSER<7,5,4,0> flag. It is cleared by writing a 1, on power-up, by the negation of DCOK when the processor is halted and by writes to the IORESET register.
DSER<2>	No grant timeout. Read/Write to clear. This bit is set if the Q22-bus does not return a bus grant within 10ms of the bus request from a CPU demand read cycle, or write cycle. During interrupt acknowledge or request read cycles this bit is not set. It is cleared by writing a 1, on power-up, by the negation of DCOK when the processor is halted and by writes to the IORESET register.
DSER<1>	Unused. Read as zero. Must be written as zero.
DSER<0>	DMA NXM. Read/Write to clear. This bit is set on a DMA transfer to a non-existent main memory location. This includes local-miss global-hit cycles and map accesses to non-existent memory. It is cleared by writing a 1, on power-up, by the negation of DCOK when the processor is halted and by writes to the IORESET register.

3.8.8 Q22-bus Error Address Register

The Q22-bus error address register (QBEAR), diagnostic processor address 2008 0008₁₆; R3000 address 1008 0008₁₆, is a read only, word accessible register which is implemented in the CQBIC chip. Its contents are valid only if DSER <5> (Q22-bus parity error) is set or if DSER<7> (Q22-bus timeout) is set.

Reading this register when DSER<5> and DSER<7> are clear will return *undefined* results. Additional Q22-bus parity errors that could have set DSER<5> or Q22-bus timeout errors that could have caused DSER<7> to set, will cause DSER<3> to set.

The QBEAR contains the address of the page in Q22-bus space which caused a parity error during an access by the on-board CPU which set DSER<5> or a master timeout which set DSER<7>.

Q22-bus address bits <21:9> are loaded into QBEAR bits <12:0>. QBEAR bits <31:13> always read as zeros (Figure 3-41).

MA-X1456-87

Figure 3-41 Q22-bus Error Address Register

NOTE

This is a read only register, if a write is attempted a machine check will be generated.

3.8.9 DMA Error Address Register

The DMA error address register (DEAR), diagnostic processor address 2008 000C₁₆; R3000 address 1008 000C₁₆, is a read only, word accessible register which is implemented in the CQBIC chip. It contains valid information only when DSER<4> (main memory error) is set or when DSER<0> (DMA NXM) is set. Reading this register when DSER<4> and DSER<0> are clear will return *undefined* data.

The DEAR contains the map translated address of the page in local memory which caused a memory error or non-existent memory error during an access by an external device or the Q22-bus interface for the CPU during a local-miss global-hit transaction or Q22-bus map access.

The contents of this register are latched when DSER<4> or DSER<0> is set. Additional main memory errors or non-existent memory errors have no effect on the DEAR until software clears DSER<4> and DSER<0>.

Mapped Q22-bus address bits <28:9> are loaded into DEAR bits <19:0>. DEAR bits <31:20> always read as zeros (Figure 3-42).

MA-X1457-87

Figure 3-42 DMA Error Address Register

NOTE

This is a read only register, if a write is attempted a machine check will be generated.

3.8.10 Error Handling

The Q22-bus interface does not generate or check CDAL bus parity.

The Q22-bus interface checks all CPU references to Q22-bus memory and I/O spaces to ensure that nothing but masked and unmasked word accesses are attempted. Any other type of reference will cause a diagnostic processor machine check abort to be initiated, an R3000 bus error exception on reads, and an IRQ3 interrupt on writes.

The Q22-bus interface maintains several timers to prevent incomplete accesses from hanging the system indefinitely. These include a 10 μ s non-existent memory timer for accesses to the Q22-bus memory and I/O spaces, a 10 μ s *no sack* timer for acknowledgement of Q22-bus DMA grants, and a 10 ms *no grant* timer for acquiring the Q22-bus.

If there is a non-existent memory (NXM) error (10 μ s timeout) while accessing the Q22-bus on a demand read reference, the associated row in the diagnostic processor cache is invalidated; DSER<7> is set; the address of the Q22-bus page being accessed is captured in QBEAR<12:0>; a diagnostic processor machine check abort is initiated; and an R3000 bus exception is generated.

If there is a NXM error on a prefetch read, or an interrupt acknowledge vector read, then the prefetch or interrupt acknowledge reference is aborted but no information is captured and no machine check occurs.

If there is a NXM error on a masked write reference, then DSER<7> is set, the address of the Q22-bus page being accessed is captured in QBEAR<12:0>, a diagnostic processor interrupt is generated at IPL 1D through vector 60₁₆, and an R3000 IRQ3 interrupt is generated.

If the Q22-bus interface does not receive an acknowledgement within 10 µs after it has granted the Q22-bus, then the grant is withdrawn, no errors are reported, and the Q22-bus interface waits 500 ns to clear the Q22-bus grant daisy chain before beginning arbitration again.

If the Q22-bus interface tries to obtain Q22-bus mastership on a CPU demand read reference and does not obtain it within 10 ms, then the associated row in the cache is invalidated, DSER<2> is set, and a machine check abort is initiated.

The Q22-bus interface also monitors Q22-bus signals BDAL<17:16> while reading information over the Q22-bus so that parity errors detected by the device being read from are recognized.

If a parity error is detected by another Q22-bus device on a diagnostic processor demand read reference to Q22-bus memory or I/O space, then the associated row in the diagnostic processor cache is invalidated, DSER<5> is set, the address of the Q22-bus page being accessed is captured in QBEAR<12:0>, and a machine check abort is initiated.

If a parity error is detected by another Q22-bus device on a prefetch request read by the diagnostic processor, the prefetch is aborted, the associated row in the diagnostic processor cache is invalidated, DSER<5> is set, the address of the Q22-bus page being accessed is captured in QBEAR<12:0>, but no machine check is generated.

The Q22-bus interface also monitors the backplane BPOK signal to detect power failures. If BPOK is negated on the Q22-bus, a power fail trap is generated, and the diagnostic processor traps through vector 0C₁₆. The R3000 gets an IRQ 3 interrupt request. The state of the Q22-bus BPOK signal can be read from SCR<15>. The Q22-bus interface continues to operate after generating the powerfail trap, until DCOK is negated.

3.9 Network Interface

The KN210 includes a network interface implemented by the LANCE chip, a 32×8 bit ROM and four $32K \times 8$ static RAMs, and located on the KN210 I/O module. When used in conjunction with the H3602-SA, this network interface allows the KN210 to be connected to either a thinwire or standard Ethernet network. It supports the Ethernet data link layer.

The network interface includes a word-wide 128 Kbyte NI buffer (the four 32×8 static RAMs) as well as four registers for control and status reporting, a DMA controller, a 24 word transmit silo and a 24 word receive silo (all resident in the LANCE chip). The DMA controller reads control information and writes status information from/to the 128 Kbyte NI buffer as well as transfers data between the NI buffer and either the transmit or receive silo. The DMA controller can perform bursts of up to eight halfword references.

Each reference (between the LANCE and the NI buffer) takes 600 ns and contains either a byte or word of data, resulting in a maximum burst duration of 4.8 μ s. The minimum time between bus requests is 8 μ s.

The CPU moves data between main memory and the NI buffer through programmed transfers.

3.9.1 Ethernet Overview

Ethernet is a serial bus that can support up to 1,024 nodes with a maximum separation of 2.8 km (1.7 miles). Data is passed over the Ethernet in Manchester encoded format at a rate of 10 million bits per second in variable-length packets. Each packet has the format shown in Figure 3-43.

The minimum size of a packet is 64 bytes, which implies a minimum data length of 46 bytes. Packets shorter than this are called *runt packets* and are treated as erroneous when received by the network controller.

Figure 3-43 Ethernet Data Packet Format

All nodes on the Ethernet have equal priority. The technique used to control access to the bus is carrier sense, multiple access, with collision detection (CSMA/CD). To access the bus, devices must first wait for the bus to clear (no carrier sensed). Once the bus is clear, all devices that want to access the bus have equal priority (multi-access), so they all attempt to transmit. After starting transmission, devices must monitor the bus for collisions (collision detection). If no collision is detected, the device may continue with transmission. If a collision is detected, then the device waits for a random amount of time and repeats the access sequence.

Ethernet allows point to point communication between two devices, as well as simultaneous communication between multiple devices. To support these two modes of communication, there are two types of network addresses, **physical** and **multicast**. These two types of addresses are both 48 bits (6 bytes) long and are described below.

A **Physical address** is the unique address associated with a particular station on the Ethernet, which should be distinct from the physical address of any other station on any other Ethernet.

A **Multicast address** is a multi-destination address associated with one or more stations on a given Ethernet (sometimes called a logical address).

Further, there are two kinds of multicast addresses, the *Multicast-group address* and the *Broadcast address*.

The *Multicast-group address* is an address associated by higher-level convention with a group of logically related stations.

The *Broadcast address* is a predefined multicast address which denotes the set of all the stations on the Ethernet.

Bit 0 (the least significant bit of the first byte) of an address denotes the type: it is 0 for physical addresses and 1 for multicast addresses. In either case the remaining 47 bits form the address value. A value of 48 ones is always treated as the broadcast address.

The hardware address of the KN210 module is determined at the time of manufacture and is stored in the network interface station address (NISA) ROM. Since every device that is intended to connect to an Ethernet network must have a unique physical address, the bit pattern blasted into the NISA ROM must be unique for each KN210. The multicast addresses to which the KN210 will respond are determined by the multicast address filter mask in the network interface initialization block.

3.9.2 Network Interface Station Address ROM

The network interface includes a byte-wide, 32-byte, socketed ROM called the network interface station address ROM (NISA ROM). One byte of this ROM appears in the low-order byte of each of 32 consecutive words in the address range 2008 4200 - 2008 427C₁₆. Bytes two and three of each word are *undefined*. The low-order byte of the first six words contain the 48-bit network physical address (NPA) of the KN210. The low-order byte in the remaining 26 words are unused. This address range is read only. Writes to this address range will result in a CDAL bus timeout and a machine check 83. The format for the NISA ROM is shown in Figure 3-44.

Figure 3-44 Network Interface Station Address ROM Format

3.9.3 LANCE Chip Overview

The LANCE chip is a microprogrammed controller which can conduct extensive operations independently of the central processor. There are four control and status registers (CSRs) within the LANCE chip which are programmed by the central processor (the MicroVAX CPU chip) to initialize the LANCE chip and start its independent operation. Once started, the LANCE uses its built-in DMA controller to directly access NI buffer RAM to get additional operating parameters and to manage the buffers it uses to transfer packets to and from the Ethernet. The LANCE uses three structures in the NI buffer:

1. **Network Interface Initialization Block**—24 bytes of contiguous memory starting on a word boundary. The initialization block is set up by the central processor and is read by the LANCE when the processor starts the LANCE's initialization process. The initialization block contains the system's network address and pointers to the receive and transmit descriptor rings.

2. **Descriptor Rings**—two logically circular rings of buffer descriptors, one ring used by the chip receiver for incoming data (the network interface receive descriptor ring) and one ring used by the chip transmitter for outgoing data (the network interface transmit descriptor ring). Each buffer descriptor in a ring is 8 bytes long and starts on an 8-byte boundary. It points to a data buffer elsewhere in memory, contains the size of that buffer, and holds various status information about the buffer's contents.
3. **Data Buffers**—contiguous portions of memory to buffer incoming packets (Receive Data Buffers) or outgoing packets (transmit data buffers). Data buffers must be at least 64 bytes long (100 bytes for the first buffer of a packet to be transmitted) and may begin on any byte boundary.

When the system is ready to begin network operation, the central processor sets up the network interface initialization block, the network interface receive descriptor ring, the network interface transmit descriptor ring, and their data buffers in the NI buffer RAM, and then starts the LANCE by writing to its CSRs. The LANCE performs its initialization process and then enters its polling loop. In this loop, it listens to the network for packets whose destination addresses are of interest and it scans the network interface transmit descriptor ring for descriptors which have been marked by the central processor to indicate that they contain outgoing data packets.

When the LANCE detects a network packet of interest, it receives and stores that packet in one or more receive buffers and marks their descriptors accordingly. When the LANCE finds a packet to be transmitted, it transmits the packet to the network and marks its descriptor when transmission is complete. Whenever the LANCE chip completes a reception or transmission (or encounters an error condition), it sets flags in NICSRO to signal the central processor (usually by an interrupt) that it has done something of interest.

3.9.4 Network Interface Register Address Port

The network interface register address port (NIRAP), address 1008 4404₁₆, is a 16-bit register that is implemented on all designs that use the LANCE chip. It is used to select which of the four CSRs is accessed through the network interface register data port. The format for the NIRAP is shown in Figure 3-45.

MA-X0047-88

Figure 3-45 Network Interface Register Address Port

Data Bit	Definition										
NIRAP <31:16>	Undefined. Should not be read or written.										
NIRAP <15:2>	Reserved. Ignored on write; read as zeros.										
NIRAP <1:0>	(CSRSEL) CSR select <1:0>. Read/Write. These bits select which of the four control and status registers (NICSR0-NICSR3) is accessible through the register data port. Cleared on power-up, by writing NICSR0 <2>, and by the negation of DCOK when SCR <7> is clear. Values are as follows:										
	<table border="1"> <thead> <tr> <th>Bits 1:0</th> <th>Register</th> </tr> </thead> <tbody> <tr> <td>0 0</td> <td>NICSR0</td> </tr> <tr> <td>0 1</td> <td>NICSR1</td> </tr> <tr> <td>1 0</td> <td>NICSR2</td> </tr> <tr> <td>1 1</td> <td>NICSR3</td> </tr> </tbody> </table>	Bits 1:0	Register	0 0	NICSR0	0 1	NICSR1	1 0	NICSR2	1 1	NICSR3
Bits 1:0	Register										
0 0	NICSR0										
0 1	NICSR1										
1 0	NICSR2										
1 1	NICSR3										

3.9.5 Network Interface Register Data Port

The network interface register data port (NIRDP), address 1008 4400₁₆, is a 16-bit register that is implemented on all designs that use the LANCE chip. It is used as a 16-bit window through which the CPU can read and write the control and status register (NICSR0-NICSR3) designated by the NIRAP.

Note that registers NICSR1, NICSR2, and NICSR3 are accessible only while NICSR0 <2>(STOP) is set. If NICSR0 <2> is clear (that is, the LANCE chip is active), attempts to read from those registers will return *undefined* data and attempts to write to them will be ignored. Accesses to a command and status register through the NIRDP do not alter the contents of the NIRAP. In normal operation only NICSR0 can be accessed, so the NIRAP should be configured so that NICSR0 is accessible through the NIRDP and left that way.

3.9.6 Network Interface Control and Status Register 0

The network interface control and status register 0 (NICSR0), address 1008 4400₁₆ when NIRAP <1:0> are set to 00, is a 16-bit register that is implemented on all designs that use the LANCE chip. This register is used to start and stop the operation of the LANCE chip and to monitor its status. All of its bits can be read at any time and none of its bits are affected by reading the register. The effects of a write operation are described individually for each bit. When power is applied to the system, all the bits in this register are cleared except the STOP bit which is set. The format for NICSR0 is shown in Figure 3-46.

Figure 3-46 Network Interface Control and Status Register 0

Data Bit	Definition
NICSR0 <31:16>	Undefined. Should not be read or written.
NICSR0 <15>	(ERR) Error summary. Read only. Writes have no effect. This bit is set whenever NICSR0 <14> (BABL), NICSR0 <13>(CERR), NICSR0 <12> (MISS), or NICSR0 <11> (MERR) are set. Cleared by clearing BABL, CERR, MISS and MERR, by setting NICSR0 <2>, on power-up and the assertion of IORESET.

Data Bit	Definition
NICSR0 <14>	(BABL) Transmitter timeout error. Read/Write to clear. This bit is set when the transmitter has been on the channel longer than the time required to send the maximum length packet. It will be set after 1519 data bytes have been transmitted (the LANCE will continue to transmit until the whole packet is transmitted or there is a failure). When this bit is set, NICSR0 <15> (ERR) and NICSR0 <7> (INTR) will also be set. Writing a 0 has no effect. Cleared by writing a 1, by setting NICSR0 <2>, on power-up and the assertion of IORESET.
NICSR0 <13>	(CERR) Collision error. Read/Write to clear. This bit is set when the collision input to the LANCE chip failed to activate within 2 microseconds after a LANCE initiated transmission is completed. This <i>collision after transmission</i> is a transceiver test feature. This function is also known as heartbeat or signal quality error test (SQE). When this bit is set, NICSR0 <15> is also set. Writing a 0 has no effect. Cleared by writing a 1, by setting NICSR0 <2>, on power-up and the assertion of IORESET.
NICSR0 <12>	(MISS) Missed packet. Read/Write to clear. This bit is set when the receiver loses a packet because it does not own a receive buffer. The MISS bit is not valid in internal loopback mode. When this bit is set, NICSR0 <15> and NICSR0 <7> bits are also set. Writing a 0 has no effect. Cleared by writing a 1, by setting NICSR0 <2>, on power-up and the assertion of IORESET.
NICSR0 <11>	(MERR) Memory error. Read/Write to clear. This bit is set when the LANCE chip attempts a DMA transfer and does not receive a ready response from the network interface buffer RAM within 25.6 microseconds after beginning the memory cycle. When MERR is set, the receiver and transmitter are turned off (NICSR0 <5:4> cleared). When this bit is set, NICSR0 <15> and NICSR0 <7> bits are also set. Writing a 0 has no effect. Cleared by writing a 1, by setting NICSR0 <2>, on power-up and the assertion of IORESET.

Data Bit	Definition
NICSR0 <10>	(RINT) Receive interrupt. Read/Write to clear. This bit is set when the LANCE chip updates an entry in the receive descriptor ring for the last buffer received or when reception is stopped due to a failure. When this bit is set, NICSR0 <7> is also set. Writing a 0 has no effect. Cleared by writing a 1, by setting NICSR0 <2>, on power-up and the assertion of IORESET.
NICSR0 <9>	(TINT) Transmitter interrupt. Read/Write to clear. This bit is set when the LANCE chip updates an entry in the transmit descriptor ring for the last buffer sent or when transmission is stopped due to a failure. When this bit is set, NICSR0 <7> is also set. Writing a 0 has no effect. Cleared by writing a 1, by setting NICSR0 <2>, on power-up and the assertion of IORESET.
NICSR0 <8>	(IDON) Initialization done. Read/Write to clear. This bit is set when the LANCE chip completes the initialization process which was started by setting NICSR0 <0> (INIT). When IDON is set, the LANCE chip has read the initialization block from memory and stored the new parameters. When this bit is set, NICSR0 <7> is also set. Writing a 0 has no effect. Cleared by writing a 1, by setting NICSR0 <2>, on power-up and the assertion of IORESET.
NICSR0 <7>	(INTR) Interrupt request. Read only. This bit is set whenever any of the bits NICSR0 <14> (BABL), NICSR0 <12> (MISS), NICSR0 <11> (MERR), NICSR0 <10> (RINT), NICSR0 <9> (TINT), or NICSR0 <8> (IDON) are set. When both this bit and NICSR0 <6> (INEA) are set, an interrupt request is posted on IRQ2 at IPL 16 with vector offset of D5 ₁₆ . Writing to this bit has no effect. Cleared by clearing BABL, MISS, MERR, RINT, TINT, and IDON, by setting NICSR0 <2>, on power-up and the assertion of IORESET.
NICSR0 <6>	(INEA) Interrupt enable. Read/Write. This bit controls whether the setting of the NICSR0 <7> (INTR) bit generates an interrupt request. When both this bit and NICSR0 <6> (INEA) are set, an interrupt request is posted at IPL 14 with vector offset of D4 ₁₆ . Cleared by setting NICSR0 <2>, on power-up and the assertion of IORESET.

Data Bit	Definition
NICSR0 <5>	(RXON) Receiver on. Read only. When set, this bit indicates that the receiver is enabled. This bit is set when initialization is completed (that is, when IDON is set, unless the DRX bit of the initialization block mode register was one) and then NICSR0 <1> (STRT) is set. Writing to this bit has no effect. Cleared by setting NICSR0 <2> or NICSR0 <11>, on power-up and the assertion of IORESET.
NICSR0 <4>	(TXON) Transmitter on. Read only. When set, this bit indicates that the transmitter is enabled. This bit is set when initialization is completed (that is, when IDON is set, unless the DTX bit of the initialization block mode register was one) and then NICSR0 <1> (STRT) is set. Writing to this bit has no effect. Cleared by setting NICSR0 <2>, NICSR0 <11>, NITMD2 <31> (UFLO), NITMD2 <30> (BUFF), or NITMD2 <26> (RTRY), on power-up and the assertion of IORESET.
NICSR0 <3>	(TDMD) Transmit demand. Read/Write. Setting this bit signals the LANCE chip to access the transmit descriptor ring without waiting for the polltime interval to elapse. This bit need not be set to transmit a packet; setting it merely hastens the chip's response to the insertion of a transmit descriptor ring entry by the host program. This bit is cleared by the LANCE chip when it recognizes the bit has been set (the bit may read as one for a short time after it is set, depending upon the level of activity in the LANCE chip). Writing a zero has no effect. This bit is also cleared by setting NICSR0 <2>, on power-up and the assertion of IORESET.
NICSR0 <2>	(STOP) Stop external activity. Read/Write. Setting this bit stops all external activity and clears the internal logic of the LANCE chip; this has the same effect on the LANCE chip as a hardware reset does. When set, the LANCE chip remains inactive until NICSR0 <1> (STRT) or NICSR0 <0> (INIT) are set. Setting STOP clears all the other bits in this register.

Data Bit	Definition
	After STOP has been set, the other three command and status registers (NICSR1, NICSR2, and NICSR3) must be reloaded before setting INIT or STRT (note that NICSR1, NICSR2, and NICSR3 may be accessed <i>only</i> while STOP is set). If the processor attempts to set STOP, INIT, and STRT at the same time, STOP takes precedence and neither STRT nor INIT is set. Writing zero has no effect. This bit is set on power-up and the assertion of IORESET. It is cleared by setting either INIT or STRT.
NICSR0 <1>	(STRT) Start operation. Read/Write. Setting this bit enables the LANCE chip to send and receive packets, perform DMA and manage its buffers. The STOP bit must be set prior to setting the STRT bit (setting STRT then clears STOP). Writing a 0 has no effect. Cleared by setting NICSR0 <2>, on power-up and the assertion of IORESET.
NICSR0 <0>	(INIT) Initialize. Read/Write. Setting this bit causes the LANCE chip to perform its initialization process, which reads the initialization block from the area in the network interface buffer RAM addressed by the contents of NICSR1 and NICSR2 through DMA. The STOP bit must be set prior to setting the INIT bit (setting INIT then clears STOP). Writing a zero has no effect. Cleared by setting NICSR0 <2>, on power-up and the assertion of IORESET.

NOTE

The INIT and STRT bits must not be set at the same time.

The proper initialization procedure is as follows:

- Set STOP in NICSR0
- Set up the initialization block in memory
- Load NICSR1 and NICSR2 with the starting address of the initialization block
- Set INIT in NICSR0
- Wait for IDON in NICSR0 to become set
- Set STRT in NICSR0 to begin operation

3.9.7 Network Interface Control and Status Register 1

The network interface control and status register 1 (NICSR1) (address 1008 4400₁₆ when NIRAP <1:0> are set to 01) is a 16-bit register that is implemented on all designs that use the LANCE chip. This register is used in conjunction with NICSR2 to supply the network interface buffer RAM address of the initialization block which the chip reads when it performs its initialization process. This register is accessible only if the STOP bit in NICSR0 is set. Bits <31:16> are undefined and bits <15:0> are read/write. On power-up, all the bits in this register are *undefined*. The format NICSR1 is shown in Figure 3-47.

MA-X0049-88

Figure 3-47 Network Interface Control and Status Register 1

Data Bit	Definition
NICSR1 <31:16>	Undefined. Should not be read or written.
NICSR1 <15:0>	(IADR15:0) Initialization block address bits <15:0>. Read/Write. These are the low-order sixteen bits of the network interface buffer RAM address of the first byte of the initialization block. Note that since the block must be aligned on a word boundary bit <0> must be zero.

3.9.8 Network Interface Control and Status Register 2

The network interface control and status register 2 (NICSR2) (address 1008 4400₁₆ when NIRAP <1:0> are set to 10) is a 16-bit register that is implemented on all designs that use the LANCE chip. This register is used in conjunction with NICSR1 to supply the network interface buffer RAM address of the initialization block which the chip reads when it performs its initialization process. This register is accessible only if the STOP bit in NICSR0 is set. Bits <31:16> are undefined bits <15:8> are reserved and bits <15:0> are read/write. On power-up, all the bits in this register are *undefined*. The format for NICSR2 is shown in Figure 3-48.

MA-X0050-88

Figure 3–48 Network Interface Control and Status Register 2

Data Bit	Definition
NICSR2 <31:16>	Undefined. Should not be read or written.
NICSR2 <15:8>	Reserved. Should not be read or written.
NICSR2 <7:0>	(IADR23:16) Initialization block address bits <23:16>. Read/Write. These are bits 23:16 of the network interface buffer RAM address of the first byte of the initialization block.

3.9.9 Network Interface Control and Status Register 3

The network interface control and status register 3 (NICSR3), address 1008 4400₁₆ when NIRAP <1:0> are set to 11, is a 16-bit register that is implemented on all designs that use the LANCE chip. This register controls certain aspects of the electrical interface between the LANCE chip and the system. It must be set by the on-board firmware as indicated for each bit. This register is accessible only if the STOP bit in NICSR0 is set. Bits <31:16> are undefined, bits <15:3> are reserved and bits <3:0> are read/write. The format for NICSR3 is shown in Figure 3–49.

MA-X0051-88

Figure 3–49 Network Interface Control and Status Register 3

Data Bit	Definition
NICSR3 <31:16>	Undefined. Should not be read or written.
NICSR3 <15:3>	Reserved. Read as zeros. Writes have no effect.
NICSR3 <2>	(BSPW) Byte swap. Read/Write. When this bit is set, the LANCE chip will swap the high and low bytes for DMA data transfers between the silo and the network interface buffer RAM in order to accommodate processors which consider address bits <15:08> to be the least significant byte of data. Cleared by setting NICSR0 <2>, on power-up and the assertion of IORESET. This bit must be set to 0 by the on-board firmware.
NICSR3 <1>	(ACON) ALE control. Read/Write. This bit controls the polarity of the signal emitted on the LANCE chip's ALE/AS pin during DMA operation. Cleared by setting NICSR0 <2>, on power-up and the assertion of IORESET. This bit must be set to 0 by the on-board firmware.
NICSR3 <0>	(BCON) Byte control. Read/Write. This bit controls the configuration of the byte mask and hold signals on the LANCE chip's pins during DMA operation. Cleared by setting NICSR0 <2>, on power-up and the assertion of IORESET. This bit must be set to 0 by the on-board firmware.

3.9.10 Network Interface Initialization Block

When the LANCE chip is initialized (by setting the INIT bit in NICSR0), it reads a 24-byte block of data called the network interface initialization block (NIIB) from the network interface buffer RAM using DMA accesses. The base address of the initialization block is formed by concatenating the contents of the NICSR1 and NICSR2. Since the NIIB must start on a word boundary, the low-order bit of the address must be zero. The initialization block is made up of twelve 16-bit words, (NIIBW0-NIIBW11), arranged as shown in Figure 3-50.

MA-X0053-88

Figure 3-50 Network Interface Initialization Block**3.9.10.1 Network Interface Initialization Block Word 0**

Word 0 of the network interface initialization block (NIIBW0), also referred to as the mode word, resides in the network interface buffer RAM at the base address of the initialization block. The mode word of the initialization block allows alteration of the LANCE chip's normal operation for testing and special applications. For normal operation the mode word is entirely zero. The format for NIIBW0 is shown in Figure 3-51.

MA-X0054-88

Figure 3-51 Network Interface Initialization Block Word 0

Data Bit	Definition
NIIBW0<15>	(PROM) Promiscuous mode. When set, all incoming packets are accepted regardless of their destination addresses. When cleared, only incoming packets with a destination address that matches the KN210's address are accepted (normal operating mode).
NIIBW0<14:7>	Reserved. Should be written with zeros.
NIIBW0<6>	(INTL) Internal loopback. This bit is used in conjunction with the NIIBW0<2> (LOOP) to control loopback operation. See the description of the LOOP bit, below.
NIIBW0<5>	(DRTY) Disable retry. When set, the LANCE chip will attempt only one transmission of a packet. If there is a collision on the first transmission attempt, a retry error (RTE) will be reported in the transmit buffer descriptor.
NIIBW0<4>	(COLL) Force collision. When set, the collision logic can be tested. The LANCE chip must be in internal loopback mode for COLL to be used. When COLL is set to one, a collision will be forced during the subsequent transmission attempt. This will result in 16 transmission attempts and a retry error (RTE) being reported in the transmit buffer descriptor.

Data Bit	Definition
NIIBW0<3>	<p>(DTCR) Disable transmit CRC. When cleared, the transmitter will generate and append a 4-byte CRC to each transmitted packet (normal operation). When DTCR is one the CRC logic is allocated instead to the receiver and no CRC is sent with a transmitted packet.</p> <p>During loopback, setting DTCR to zero will cause a CRC to be generated and sent with the transmitted packet, but no CRC check can be done by the receiver since the CRC logic is shared and cannot both generate and check a CRC at the same time. The CRC transmitted with the packet will be received and written into memory following the data where it can be checked by software.</p> <p>If DTCR is set to one during loopback, the driving software must compute and append a CRC value to the data to be transmitted. The receiver will check this CRC upon reception and report any error.</p>
NIIBW0 <2>	<p>(LOOP) Loopback control. This bit is used in conjunction with NIIBW0 <6> to perform internal loopback tests on the LANCE chip. During loopback the LANCE chip operates in full duplex mode. The maximum packet size is limited to 32 data bytes (in addition to which 4 CRC bytes may be appended). During loopback, the runt packet filter is disabled because the maximum packet is forced to be smaller than the minimum size Ethernet packet (64 bytes).</p> <p>Setting LOOP to one allows simultaneous transmission and reception for a packet constrained to fit within the silo. The chip waits until the entire packet is in the silo before beginning serial transmission. The incoming data stream fills the silo from behind as it is being emptied. Moving the received packet out of the silo into memory does not begin until reception has ceased.</p> <p>In loopback mode, transmit data chaining is not possible. Receive data chaining is allowed regardless of the receive buffer length. (In normal operation, the receive buffers must be 64 bytes long, to allow time for buffer lookahead.)</p>

Data Bit	Definition	
Valid loopback bit settings are:		
LOOP	INTL	Operation
0	x	Normal on-line operation
1	0	External loopback
1	1	Internal loopback
<p>Internal loopback allows the LANCE chip to receive its own transmitted packet without disturbing the network. The LANCE chip will not receive any packets from the network while it is in internal loopback mode.</p>		
<p>External loopback allows the LANCE chip to transmit a packet through the transceiver out to the network cable to check the operability of all circuits and connections between the LANCE chip and the network cable. Multicast addressing in external loopback is valid only when DTCR is one (user needs to append the 4 CRC bytes). In external loopback, the LANCE chip also receives packets from other nodes.</p>		
NIIBW0<1>	<p>(DTX) Disable transmitter. When set, the LANCE chip will not set the TXON bit in NICSRO at the completion of initialization. This will prevent the LANCE chip from attempting to access the transmit descriptor ring, hence no transmissions will be attempted.</p>	
NIIBW0<0>	<p>(DRX) Disable receiver. When set, the LANCE chip will not set the RXON bit in NICSRO at the completion of initialization. This will cause the LANCE chip to reject all incoming packets and to not attempt to access the receive descriptor ring.</p>	

3.9.10.2 Network Interface Initialization Block Words 1-3

Words 1-3 of the network interface initialization block (NIIBW1-3), reside in the network interface buffer RAM at the base address of the NIIB plus 2, 4, and 6 respectively. These three words contain the 48-bit NPA of the KN210 and are loaded by the resident firmware from the NISA ROM.

This address identifies the KN210 to the Ethernet network and must be unique within the domain of the network. The low-order bit (bit 0) of this address must be zero to indicate it is a physical rather than multicast address. The format for network interface initialization block words 1 through 3 is shown in Figure 3-52.

MA-X0055-88

Figure 3-52 Network Interface Initialization Block Words 1-3**3.9.10.3 Network Interface Initialization Block Words 4-7**

Words 4-7 of the network interface initialization block (NIIBW4-7), reside in network interface buffer RAM at the base address of the NIIB plus 8, 10, 12, and 14 respectively. These four words contain the 64-bit multicast address filter mask. The format for network interface initialization block words 4 through 7 is shown in Figure 3-53.

MA-X0056-88

Figure 3-53 Network Interface Initialization Block Words 4-7

Multicast Ethernet addresses are distinguished from physical network addresses by the presence of a one in bit 0 of the 48-bit address field. If an incoming packet contains a physical destination address (bit 0 is zero), then its entire 48 bits are compared with the network physical address and the packet is ignored if they are not equal. If the packet contains a multicast destination address which is all ones (the broadcast address), it is always accepted and stored regardless of the contents of the multicast address filter mask.

All other multicast addresses are processed through the multicast address filter to determine whether the incoming packet will be stored in a receive buffer. This filtering is performed by passing the multicast address field through the CRC generator. The high-order 6 bits of the resulting 32-bit CRC are used to select one of the 64 bits of the multicast address filter mask. (These high-order six bits represent in binary the number of the bit in the multicast address filter mask.)

If the bit selected from the multicast address filter mask is one, the packet is stored in a receive buffer; otherwise it is ignored. This mechanism effectively splits the entire domain of 2^{47} multicast addresses into 64 parts, and multicast addresses falling into each part will be accepted or ignored according to the value of the corresponding bit in the multicast address filter mask. The driver program must examine the addresses of the packets accepted by this partial filtering to complete the filtering task.

3.9.10.4 Network Interface Initialization Block Words 8,9

Words 8 and 9 of the network interface initialization block (NIIBW8,9), also referred to as the receive descriptor ring pointer, reside in network interface buffer RAM at the base address of the NIIB plus 16 and 18 respectively. These two words contain the starting address and the number of descriptors in the receive descriptor ring. The format for NIIBW8 is shown in Figure 3-54.

MA-X0057-88

Figure 3-54 Network Interface Initialization Block Word 8

Data Bit	Definition
NIIBW8 <15:0>	(RDRA <15:0>) Receive descriptor ring address <15:0>. This field contains bits <15:0> of the base address of the receive descriptor ring. Since the receive descriptor ring must start on an 8-byte boundary, bits <2:0> of this field must be zero.

The format for NIIBW9 is shown in Figure 3-55.

MA-X0058-88

Figure 3-55 Network Interface Initialization Block Word 9

Data Bit	Definition																		
NIIBW9<15:13>	(NRBD) Number of receive buffer descriptors. This field gives the number of receive buffer descriptors in the receive descriptor ring, expressed as a power of two:																		
	<table border="1"> <thead> <tr> <th>Value</th> <th>Number of Descriptors</th> </tr> </thead> <tbody> <tr> <td>000</td> <td>1</td> </tr> <tr> <td>001</td> <td>2</td> </tr> <tr> <td>010</td> <td>4</td> </tr> <tr> <td>011</td> <td>8</td> </tr> <tr> <td>100</td> <td>16</td> </tr> <tr> <td>101</td> <td>32</td> </tr> <tr> <td>110</td> <td>64</td> </tr> <tr> <td>111</td> <td>128</td> </tr> </tbody> </table>	Value	Number of Descriptors	000	1	001	2	010	4	011	8	100	16	101	32	110	64	111	128
Value	Number of Descriptors																		
000	1																		
001	2																		
010	4																		
011	8																		
100	16																		
101	32																		
110	64																		
111	128																		
NIIBW9<12:8>	Reserved; should be zeros.																		
NIIBW9<7:0>	(RDRA <23:16>) Receive descriptor ring address <23:16>. This field contains bits <23:16> of the base address of the receive descriptor ring.																		

3.9.10.5 Network Interface Initialization Block Words 10,11

Words 10 and 11 of the network interface initialization block (NIIBW10,11), also referred to as the transmit descriptor ring pointer, reside in network interface buffer RAM at the base address of the NIIB plus 20 and 22 respectively. These two words contain the starting address and the number of transmit buffer descriptors in the transmit descriptor ring. The format for NIIBW10 is shown in Figure 3-56.

MA-X0059-88

Figure 3-56 Network Interface Initialization Block Word 10

Data Bit	Definition
NIIBW10<15:0>	(TDRA <15:0>) Transmit descriptor ring address <15:0>. This field contains bits <15:0> of the base address of the transmit descriptor ring. Since the transmit descriptor ring must start on an 8-byte boundary, bits <2:0> of this field must be zero.

The format for NIIBW11 is shown in Figure 3-57.

MA-X0060-88

Figure 3-57 Network Interface Initialization Block Word 11

Data Bit	Definition																		
NIIBW11<15:13>	(NTBD) Number of transmit buffer descriptors. This field gives the number of transmit buffer descriptors in the transmit descriptor ring, expressed as a power of two:																		
	<table border="1"> <thead> <tr> <th>Value</th> <th>Number of Descriptors</th> </tr> </thead> <tbody> <tr> <td>000</td> <td>1</td> </tr> <tr> <td>001</td> <td>2</td> </tr> <tr> <td>010</td> <td>4</td> </tr> <tr> <td>011</td> <td>8</td> </tr> <tr> <td>100</td> <td>16</td> </tr> <tr> <td>101</td> <td>32</td> </tr> <tr> <td>110</td> <td>64</td> </tr> <tr> <td>111</td> <td>128</td> </tr> </tbody> </table>	Value	Number of Descriptors	000	1	001	2	010	4	011	8	100	16	101	32	110	64	111	128
Value	Number of Descriptors																		
000	1																		
001	2																		
010	4																		
011	8																		
100	16																		
101	32																		
110	64																		
111	128																		
NIIBW11 <12:8>	Reserved; should be zeros.																		
NIIBW11 <7:0>	(TDRA <23:16>) Transmit descriptor ring address <23:16>. This field contains bits <23:16> of the base address of the transmit descriptor ring.																		

3.9.11 Buffer Management

The LANCE chip manages its data buffers by using two rings of buffer descriptors which are stored in the network interface buffer RAM: the network interface receive descriptor ring and the network interface transmit descriptor ring. Each buffer descriptor points to a data buffer elsewhere in the network interface buffer RAM, contains the size of that buffer, and contains status information about that buffer's contents.

The starting location in the network interface buffer RAM of each ring and the number of descriptors in it are given to the LANCE chip through the NIIB during the chip initialization process. Each descriptor is 8 bytes long and must be aligned on an 8-byte boundary (the three low-order bits of its address must be zero). The descriptors in a ring are physically contiguous in the network interface buffer RAM and the number of descriptors must be a power of two. The LANCE keeps an internal index to its current position in each ring which it increments modulo the number of descriptors in the ring as it advances around each ring.

Once started, the LANCE chip polls each ring to find descriptors for buffers in which to receive incoming packets and from which to transmit outgoing packets, and revises the status information in buffer descriptors as it processes their associated buffers. When polling, the LANCE chip is limited to looking only one ahead of the descriptor with which it is currently working. The high speed of the data stream requires that each buffer be at least 64 bytes long to allow time to chain buffers for packets which are larger than one buffer. (The first buffer of a packet to be transmitted should be at least 100 bytes to avoid problems in case a late collision is detected.)

Each descriptor in a ring is "owned" either by the LANCE chip or by the host processor; this status is indicated by the OWN bit in each descriptor. Mutual exclusion is accomplished by the rule that each device can only relinquish ownership of a descriptor to the other device, it can never take ownership; and that each device cannot change any field in a descriptor or its associated buffer after it has relinquished ownership. When the host processor sets up the rings of descriptors before starting the LANCE chip, it sets the OWN bits such that the LANCE chip will own all the descriptors in the network interface receive descriptor ring (to be used by the LANCE to receive packets from the network) and the host will own all the descriptors in the network interface transmit descriptor ring (to be used by the host to set up packets to be transmitted to the network).

3.9.12 Network Interface Receive Descriptor Ring

The network interface receive descriptor ring (NIRDR) contains a receive buffer descriptor for each receive buffer (Figure 3-58). It is located in a contiguous block of the network interface buffer RAM whose base address is formed by concatenating the contents of NIIBW8 and NIIBW9 <7:0> (RDRA<23:0>). Since the NIRDR must start on an 8-byte boundary, bits <2:0> of this address must be zero. The size of the network interface receive descriptor ring can vary between 8 and 1024 bytes depending on the number of 8-byte descriptors it contains (Figure 3-58). The number of descriptors must be a power of two between one and 128 and is determined by NIIBW9 <15:13> (NRBD).

MA-X0061-88

Figure 3-58 Network Interface Receive Descriptor Ring

3.9.12.1 Receive Buffer Descriptors

Receive buffer descriptor **n** contains the base address and size of a receive buffer as well as status and error information. It is four words (eight bytes) in length and is located in the receive descriptor ring at base+8**n**.

A representation of a typical receive buffer descriptor (RBD) is shown in Figure 3-59.

MA-X0062-88

Figure 3-59 Receive Buffer Descriptors**Receive Buffer Descriptor n Word 0**

Word 0 of RBD **n** (RBD**n**W0) resides in the network interface buffer RAM at the base address of the NIIRD **R** + 8**n**. This word contains a portion of the base address of the associated receive buffer. The format for receive buffer descriptor **n** word 0 is shown in Figure 3-60.

MA-X0063-88

Figure 3-60 Receive Buffer Descriptor n Word 0

Data Bit	Definition
RBD n W0 <15:0>	(BADR) Buffer address. This field contains bits <15:0> of the 24-bit network address buffer RAM address of the start of the buffer associated with this descriptor. Written by the host; unchanged by the LANCE.

Receive Buffer Descriptor n Word 1

Word 1 of RBD **n** (RBD**n**W1) resides in the network interface buffer RAM at the base address of the NIIRDR +8**n**+2. This word contains a portion of the base address of the associated receive buffer as well as status and error information. The format for receive buffer descriptor **n** word 1 is shown in Figure 3-61.

MA-X0064-88

Figure 3-61 Receive Buffer Descriptor n Word 1

Data Bit	Definition
RBD n W1 <15>	(OWN) Owned flag. This bit indicates whether the descriptor is owned by the host (OWN = 0) or by the LANCE chip (OWN = 1). The LANCE clears OWN after filling the buffer associated with the descriptor with an incoming packet. The host sets OWN after emptying the buffer. In each case, this must be the last bit changed by the current owner, since changing OWN passes ownership to the other party and the relinquishing party must not thereafter alter anything in the descriptor or its buffer.
RBD n W1 <14>	(ERR) Error summary. This is the logical OR of the FRE, OFE, CHE and BUE bits in this word. Set by the LANCE chip and cleared by the host.
RBD n W1 <13>	(FRE) Framing error. This bit is set by the LANCE chip to indicate that the incoming packet stored in the buffer had both a non-integral multiple of 8 bits and a checksum error (CHE). It is cleared by the host.

Data Bit	Definition
RBD n W1 <12>	(OFE) Overflow error. This bit is set by the LANCE chip to indicate that the receiver has lost part or all of an incoming packet because it could not store it in the buffer before the chip's silo overflowed. Cleared by the host.
RBD n W1 <11>	(CHE) Checksum error. This bit is set by the LANCE chip to indicate that the received packet has an invalid CRC checksum. Cleared by the host.
RBD n W1 <10>	(BUE) Buffer error. This bit is set by the LANCE chip when it has used all its owned receive descriptors or when it could not get the next descriptor in time while attempting to chain to a new buffer in the midst of a packet. When a buffer error occurs, an overflow error (bit OFLO) also occurs because the LANCE continues to attempt to get the next buffer until its silo overflows. BUE is cleared by the host.
RBD n W1 <9>	(STP) Start of packet. This bit is set by the LANCE chip to indicate that this is the first buffer used for this packet. Cleared by the host.
RBD n W1 <8>	(ENP) End of packet. This bit is set by the LANCE chip to indicate that this is the last buffer used for this packet. When both STP and ENP are set in a descriptor, its buffer contains an entire packet; otherwise two or more buffers have been chained together to hold the packet. ENP is cleared by the host.
RBD n W1 <7:0>	(BADR <23:16>) Buffer address <23:16>. This field contains bits <23:16> of the 24-bit the NI buffer RAM address of the start of the buffer associated with this descriptor. Written by the host; unchanged by the LANCE.

Receive Buffer Descriptor n Word 2

Word 2 of RBD n (RBD n W2) resides in the network interface buffer RAM at the base address of the NIIRDR +8 n +4. This word contains the size of the associated receive buffer. The format for receive buffer descriptor n word 2 is shown in Figure 3-62.

MA-X0065-88

Figure 3-62 Receive Buffer Descriptor n Word 2

Data Bit	Definition
RBDnW2<15:12>	These bits must be set by the host to ones. Unchanged by the LANCE chip.
RBDnW2<11:0>	(BSZ <11:0>) Buffer size. This field contains the size (in bytes) of the associated receive buffer in two's complement form. Note that the minimum buffer size is 64 bytes (to allow enough time for chaining buffers) and the maximum buffer size is 1518 bytes (the largest legal Ethernet packet). Written by the host; unchanged by the LANCE chip.

Receive Buffer Descriptor n Word 3

Word 3 of RBD **n** (RBDnW3) resides in the network interface buffer RAM at the base address of the NIIRDR +8**n**+6. This word contains the size of the packet that was received. The format for receive buffer descriptor **n** word 3 is shown in Figure 3-63.

MA-X0066-88

Figure 3-63 Receive Buffer Descriptor n Word 3

Data Bit	Definition
RBDnW3<15:12>	These bits are reserved. They should be set to zeros by the host when it constructs the descriptor.
RBDnW3<11:0>	(PSZ <11:0>) Packet size. This field contains the size (in bytes) of the received packet. This field is valid only in a descriptor in which ENP is set (last buffer) and ERR is clear (no error). Set by the LANCE chip and cleared by the host.

3.9.13 Receive Buffers

Receive buffers are set up by the host by adding a receive buffer descriptor to the network interface receive buffer descriptor ring. These buffers are used for storing incoming Ethernet packets. An Ethernet packet may span multiple buffers, but a buffer cannot contain more than one Ethernet packet. The base address of a receive buffer is formed by concatenating the contents of RBDnW0 and RBDnW1 <7:0> (BADR). The size of a receive buffer is determined by RBDnW2 <11:0>.

Receive buffers are structured as shown in Figure 3-64.

MA-X0067-88

Figure 3-64 Receive Buffers

3.9.14 Network Interface Transmit Descriptor Ring

The network interface transmit descriptor ring (NITDR) contains a transmit buffer descriptor for each transmit buffer (Figure 3-65). It is located in a contiguous block of the network interface buffer RAM whose base address is formed by concatenating the contents of the NIIBW10 and NIIBW11 <7:0> (TDRA<23:0>). Since the NITDR must start on an 8-byte boundary, bits <2:0> of this address must be zero. The size of the network interface transmit descriptor ring can vary between 8 and 1024 bytes depending on the number of 8-byte descriptors it contains. The number of descriptors must be a power of 2 between 1 and 128 and is determined by NIIBW11 <15:13> (NTBD).

MA-X0068-88

Figure 3-65 Network Interface Transmit Descriptor Ring

3.9.14.1 Transmit Buffer Descriptors

Transmit buffer descriptor **n** contains the base address, size, of a transmit buffer as well as status and error information. It is four words (eight bytes) in length and is located in the transmit descriptor ring at base+8**n**.

A representation of a typical transmit buffer descriptor (TBD) is shown in Figure 3-66.

MA-X0069-88

Figure 3-66 Transmit Buffer Descriptors

Transmit Buffer Descriptor **n** Word 0

Word 0 of TBD **n** (TBD**n**W0) resides in the network interface buffer RAM at the base address of the NIIRDR +8**n**. This word contains a portion of the base address of the associated transmit buffer. The format for transmit buffer descriptor **n** word 0 is shown in Figure 3-67.

MA-X0070-88

Figure 3-67 Transmit Buffer Descriptor **n Word 0**

Data Bit	Definition
TBD n W0<15:0>	(BADR) Buffer address. This field contains bits <15:0> of the 24-bit network interface buffer RAM address of the start of the buffer associated with this descriptor. Written by the host; unchanged by the LANCE chip.

Transmit Buffer Descriptor n Word 1

Word 1 of TBD **n** (TBD**n**W1) resides in the network interface buffer RAM at the base address of the NIIRDR +8**n**+2. This word contains a portion of the base address of the associated transmit buffer as well as status and error information. The format for transmit buffer descriptor **n** word 1 is shown in Figure 3-68.

MA-X0071-88

Figure 3-68 Transmit Buffer Descriptor n Word 1

Data Bit	Definition
TBD n W1 <15>	(OWN) Owned flag. This bit indicates whether the descriptor is owned by the host (OWN = 0) or by the LANCE chip (OWN = 1). The LANCE clears OWN after filling the buffer associated with the descriptor with an incoming packet. The host sets OWN after emptying the buffer. In each case, this must be the last bit changed by the current owner, since changing OWN passes ownership to the other party and the relinquishing party must not thereafter alter anything in the descriptor or its buffer.
TBD n W1 <14>	(ERR) Error summary. This bit is the logical OR of the COE, CAE, UFE and RTE bits in this descriptor. Set by the LANCE chip and cleared by the host.
TBD n W1 <13>	(RSV) Reserved. The LANCE chip will write a zero in this bit.
TBD n W1 <12>	(MRE) More retries. The LANCE chip sets this bit when more than one retry was required to transmit the packet. Cleared by the host.

Data Bit	Definition
TBD n W1 <11>	(ORE) One retry. The LANCE chip sets this bit when exactly one retry was required to transmit the packet. Cleared by the host.
TBD n W1 <10>	(DEF) Deferred. The LANCE chip sets this bit when it had to defer while trying to transmit the packet. This occurs when the network is busy when the LANCE is ready to transmit. Cleared by the host.
TBD n W1 <9>	(STP) Start of packet. This bit is set by the host to indicate that this is the first buffer used for this packet. STP is not changed by the LANCE chip.
TBD n W1 <8>	(ENP) End of packet. This bit is set by the host to indicate that this is the last buffer used for this packet. When both STP and ENP are set in a descriptor, its buffer contains an entire packet; otherwise two or more buffers have been chained together to hold the packet. ENP is not changed by the LANCE chip.
TBD n W1 <7:0>	(BADR <23:16>) Buffer address <23:16>. This field contains bits <23:16> of the 24-bit network interface buffer RAM address of the start of the buffer associated with this descriptor. Written by the host; unchanged by the LANCE chip.

Transmit Buffer Descriptor n Word 2

Word 2 of TBD n (TBD n W2) resides in the network interface buffer RAM at the base address of the NIITDR +8 n +4. This word contains the size of the associated transmit buffer. The format for transmit buffer descriptor n word 2 is shown in Figure 3-69.

Figure 3-69 Transmit Buffer Descriptor n Word 2

Data Bit	Definition
TBD n W2<15:12>	These bits must be set by the host to ones. Unchanged by the LANCE chip.
TBD n W2 <11:0>	(BSZ <11:0>) Buffer size. This field contains the size (in bytes) of the associated transmit buffer in two's complement form. Note that the minimum buffer size is 64 bytes (to allow enough time for chaining buffers) and the maximum buffer size is 1518 bytes (the largest legal Ethernet packet). Written by the host; unchanged by the LANCE chip.

Transmit Buffer Descriptor n Word 3

Word 3 of TBD n (TBD n W3) resides in the network interface buffer RAM at the base address of the NIITDR +8 n +6. This word contains error information and a time domain reflectometer. The contents of this word are valid only when the ERR bit in TBD n W2 has been set by the LANCE chip.

The format for transmit buffer descriptor n word 3 is shown in Figure 3-70.

MA-X0073-88

Figure 3-70 Transmit Buffer Descriptor n Word 3

Data Bit	Definition
TBDnW3 <15>	(BUE) Buffer error. This bit is set by the LANCE chip during transmission when it does not find the ENP bit set in the current descriptor and it does not own the next descriptor. When BUE is set, the UFE bit (below) is also set because the LANCE chip continues to transmit until its silo becomes empty. BUE is cleared by the host.
TBDnW3 <14>	(UFE) Underflow error. This bit is set by the LANCE chip when it truncates a packet being transmitted because it has drained its silo before it was able to obtain additional data from a buffer in memory. UFE is cleared by the host.
TBDnW3 <13>	(RSV) Reserved. The LANCE chip will write a zero in this bit.
TBDnW3 <12>	(COE) Late collision error. This bit is set by the LANCE chip to indicate that a collision has occurred after the slot time of the network channel has elapsed. The LANCE chip does not retry after a late collision. COE is cleared by the host.
TBDnW3 <11>	(CAE) Loss of carrier error. This bit is set by the LANCE chip when the carrier present input to the chip becomes false during a transmission initiated by the LANCE. The LANCE chip does not retry after such a failure. CAE is cleared by the host.
TBDnW3 <10>	(RTE) Retries exhausted. This bit is set by the LANCE chip after 16 attempts to transmit a packet have failed due to repeated collisions on the network. (If the DRTY bit of network interface initialization block word 0 (mode word) is set, RTE will instead be set after only one failed transmission attempt.) RTE is cleared by the host.
TBDnW3 <9:0>	(TDR) Time domain reflectometer. These bits are the value of an internal counter which is set by the LANCE chip to count system clocks from the start of a transmission to the occurrence of a collision. This value is useful in determining the approximate distance to a cable fault; it is valid only when the RTE bit in this word is set.

3.9.15 Transmit Buffers

Transmit buffers are set up by the host by adding a transmit buffer descriptor to the network interface transmit buffer descriptor ring. These buffers are used for storing incoming Ethernet packets. An Ethernet packet may span multiple buffers, but a buffer cannot contain more than one Ethernet packet. The base address of a transmit buffer is formed by concatenating the contents of the $TBDnW0$ and $TBDnW1$ <7:0> (BADR). The size of a transmit buffer is determined by $TBDnW2$ <11:0>. Transmit buffers are structured as shown in Figure 3-71.

MA-X0074-88

Figure 3-71 Transmit Buffers

3.9.16 LANCE Operation

The LANCE chip operates independently of the host under control of its own internal microprogram. This section is a simplified description of the operation of the LANCE in terms of its principal microcode routines (these should not be confused with device driver programming in the host, which is not a part of this specification). These microcode routines make use of numerous temporary storage cells within the LANCE chip; most of these are not accessible from outside the chip but they are mentioned here when necessary to clarify the operation of the microcode.

Two such (conceptual) internal variables are of central importance: the pointers to the "current" entry in the receive descriptor ring and in the transmit descriptor ring, which are referred to below as TXP and RXP. Each of these designates the descriptor which the LANCE will use for the next operation of that type. If the descriptor designated by one of these pointers is not owned by the LANCE (the OWN bit is 0), then the LANCE can neither perform activity of that type nor advance the pointer.

For the transmit ring, the LANCE will do nothing until the host sets up a packet in the buffer and sets the OWN bit in the descriptor designated by the LANCE's TXP. (The host must keep track of the position of the TXP, since setting up a packet in some other descriptor will not be detected by the LANCE.) For the receive ring, if the LANCE does not own the descriptor designated by RXP, it cannot receive a packet. In both rings, when the LANCE finishes with a descriptor and relinquishes it to the host by clearing OWN, it then advances the ring pointer (modulo the number of entries in the ring).

When the LANCE begins activity using the current descriptor (that is, begins receiving or transmitting a packet), it may look ahead at the next descriptor and attempt to read its first three words in advance so it can chain to the next buffer in mid-packet without losing data. However, it does not actually advance its RXP or TXP until it has cleared the OWN bit in the current descriptor.

The LANCE is a very complex chip and this specification does not attempt to cover all the details of its operation. The chip purchase specification and the chip vendor's literature should also be consulted for more detailed information.

3.9.16.1 Switch Routine

Upon power on, the STOP bit is set and the INIT and STRT bits are cleared in NICSR0. The LANCE microprogram begins execution in the switch routine, which tests the INIT, STRT, and STOP bits. When the host sets either INIT or STRT, STOP is cleared. While STOP is set, if the host writes to NICSR1 and NICSR2, that data is stored for use by the initialization routine.

When the microprogram sees STOP cleared, it tests first the INIT bit and then the STRT bit. If INIT is set, it performs the initialization routine. Then if STRT is set, it begins active chip operation by jumping to the look-for-work routine. Control returns to the switch routine whenever the host again sets the STOP bit (which also clears the INIT and STRT bits). Note that the ring pointers RXP and TXP are not altered by the setting of either STOP or START; they are reset to the start of their rings only when INIT is set.

3.9.16.2 Initialization Routine

The initialization routine is called from the switch routine when the latter finds the INIT bit set. It reads the initialization block from the memory addressed by NICSR1 and NICSR2 and stores its data within the LANCE chip. This routine also sets the ring pointers RXP and TXP to the start of their rings (that is, to point to the descriptor at the lowest memory address in the ring).

3.9.16.3 Look-For-Work Routine

The look-for-work routine is executed while the LANCE is active and looking for work. It is entered from the switch routine when the STRT bit is set, and is returned to from the receive and transmit routines after they have received or transmitted a packet.

This routine begins by testing whether the receiver is enabled (bit RXON of NICSR0 is set). If so, it tries to have a receive buffer available for immediate use when a packet addressed to this system arrives. It tests its internal registers to see whether it has already found a receive descriptor owned by the LANCE and, if not, calls the receive poll routine to attempt to get a receive buffer.

Next the routine tests whether the transmitter is enabled (bit TXON of NICSR0 is set). If so, it calls the transmit poll routine to see whether there is a packet to be transmitted and to transmit it.

If there was no transmission and the TDMD bit of NICSR0 is not set, the microprogram delays 1.6 milliseconds and then goes to check the receive descriptor status again. If a packet was transmitted or the host has set TDMD, the delay is omitted so that multiple packets will be transmitted as quickly as possible.

If at any point in this routine the receiver detects an incoming packet whose destination address matches the station's physical address, is the broadcast address, or passes the multicast address filter (or if the PROM bit of NIIBW0 is set), the receive routine is called.

3.9.16.4 Receive Poll Routine

The receive poll routine is called whenever the receiver is enabled and the LANCE needs a free buffer from the receive descriptor ring. The routine reads the second word of the descriptor designated by RXP and, if the OWN bit in it is set, reads the first and third words also.

3.9.16.5 Receive Routine

The receive routine is called when the receiver is enabled and an incoming packet's destination address field matches one of the criteria described above. The routine has three sections: initialization, lookahead, and descriptor update.

In initialization, the routine checks whether a receive ring descriptor has already been acquired by the receive poll routine. If not, it makes one attempt to get the descriptor designated by RXP (if OWN is not set in it, MISS and ERR are set in NICSR0 and the packet is lost). The buffer thus acquired is used by the receive DMA routine to empty the silo.

In lookahead, the routine reads the second word of the next descriptor in the receive ring and, if the OWN bit is set, reads the rest of the descriptor and holds it in readiness for possible data chaining.

The descriptor update section is performed when either the current buffer is filled or the packet ends. If the packet ends but its total length is less than 64 bytes, it is an erroneous "runt packet" and is ignored: no status is posted in the descriptor, RXP is not moved, and the buffer will be reused for the next incoming packet (this is why a receive buffer must be at least 64 bytes long; otherwise the runt might be detected after advancing RXP).

If the packet ends (with or without error), the routine writes the packet length into MCNT, sets ENP and other appropriate status bits and clears OWN in the current descriptor, and sets RINT in NICSRO to signal the host that a complete packet has been received. Then it advances RXP and returns to the look-for-work routine.

If the buffer is full and the packet has not ended, chaining is required. The routine releases the current buffer by writing status bits into its descriptor (clearing OWN and ENP, in particular), makes current the next descriptor data acquired in the lookahead section, advances RXP, and goes to the lookahead section to prepare for possible additional chaining. Note that RINT is not set in NICSRO, although the host would find OWN cleared if it looked at the descriptor, and it could begin work on that section of the packet, since the mutual exclusion rule prevents the LANCE from going back and altering it.

3.9.16.6 Receive DMA Routine

The receive DMA routine is invoked asynchronously by the chip hardware during execution of the receive routine whenever the silo contains 16 or more bytes of incoming data or when the packet ends and the silo is not empty. It executes DMA cycles to drain data from the silo into the buffer designated by the current descriptor.

3.9.16.7 Transmit Poll Routine

The transmit poll routine is called by the look-for-work routine to see whether a packet is ready for transmission. It reads the second word of the descriptor designated by TXP and tests the OWN bit. If OWN is zero, the LANCE does not own the buffer and this routine returns to its caller. If OWN is set, the routine tests the STP bit, which should be set to indicate the start of a packet. If STP is clear, this is an invalid packet; the LANCE sets its OWN bit to return it to the host, sets TINT in NICSRO to notify the host, and advances TXP to the next transmit descriptor.

If both OWN and STP are set, this is the beginning of a packet, so the transmit poll routine reads the rest of the descriptor and then calls the transmit routine to transmit the packet. During this time the chip is still watching for incoming packets from the network and it will abort the transmit operation if one arrives.

3.9.16.8 Transmit Routine

The transmit routine is called from the transmit poll routine when the latter finds the start of a packet to be transmitted. This routine has three sections: initialization, lookahead, and descriptor update.

In initialization, the routine sets the chip's internal buffer address and byte count from the transmit descriptor, enables the transmit DMA engine, and starts transmission of the packet preamble. It then waits until the transmitter is actually sending the bit stream (including possible backoff-and-retry actions in case of collisions).

In lookahead, the transmit routine tests the current descriptor to see whether it is the last in the packet (the ENP bit is set). If so, no additional buffer is required so the routine waits until all the bytes from the current packet have been transmitted. If not, the routine attempts to get the next descriptor and hold it in readiness for data chaining, and then waits until all the bytes from the current buffer have been transmitted.

Descriptor update is entered when all the bytes from a buffer have been transmitted or an error has occurred. If there is no error and the buffer was not the last of the packet, the pre-fetched descriptor for the next buffer is made current for use by the transmit DMA routine. The routine writes the appropriate status bits and clears the OWN bits in the current descriptor and advances TXP. If this was the last buffer in the packet, the routine sets the TINT bit in NICSRO to notify the host and returns to the look-for-work routine. Otherwise it goes back to the lookahead section in this routine.

3.9.16.9 Transmit DMA Routine

The transmit DMA routine is invoked asynchronously by the chip hardware during execution of the transmit routine whenever the silo has 16 or more empty bytes. It executes DMA cycles to fill the silo with data from the buffer designated by the current descriptor.

3.9.16.10 Collision Detect Routine

This routine is invoked asynchronously by the chip hardware during execution of the transmit routine when a collision is detected on the network. It ensures that the jam sequence is transmitted, then backs up the chip's internal buffer address and byte count registers, waits for a pseudo-random backoff time, and then attempts the transmission again. If 15 retransmission attempts fail (a total of 16 attempts), it sends the microcode to the descriptor update routine to report an error in the current transmit descriptor (bits RTRY and ERR are set).

3.9.17 LANCE Programming Notes

The following are LANCE programming notes:

1. The interrupt signal is simply the OR of the interrupt-causing conditions. If another such condition occurs while the interrupt signal is already asserted, there will not be another active transition of the interrupt signal and another interrupt will not be generated. An interrupt service routine should use logic similar to the following to avoid losing interrupts:
 - a. Read NICSR0 and save the results in a register, say R2.
 - b. Clear the interrupt enable bit INEA in the saved data in R2.
 - c. Write NICSR0 with the saved data in R2. This will make the interrupt signal false because INEA is clear and will clear all the write-one-to-reset bits such as RINT, TINT and the error bits; it will not alter the STRT, INIT or STOP bits nor any interrupt-cause bits which came true after NICSR0 was read.
 - d. Write NICSR0 with only INEA to enable interrupts again.
 - e. Service all the interrupt and error conditions indicated by the flags in the data in R2.
 - f. Exit from the interrupt service routine.
2. An interrupt is signalled to the host only when the last buffer of a multibuffer (chained) packet is received or transmitted. However, the OWN bit in each descriptor is cleared as soon as the LANCE has finished with that portion of the packet, and the mutual exclusion rule makes it safe for the host to process such a descriptor and its buffer.

3. When a transmitter underflow occurs (UFE is set in a transmit descriptor because the silo is not filled fast enough), the LANCE will turn off its transmitter and the LANCE must be restarted to turn the transmitter back on again. This can be done by setting STOP in NICSR0 and then setting STRT in NICSR0 (DTX will still be clear in the chip's internal copy of NIIBW0). It is not necessary to set INIT to reread the initialization block.

Note that setting STOP will immediately terminate any reception which is in progress. If the status of a receive descriptor has been updated and its OWN bit is now clear, then the contents of its buffer are valid. If the incoming packet was chained into more than one buffer, however, the packet is only valid if its last buffer has been completed (the one with the ENP bit set).

4. The network controller hardware requires up to five seconds after power on to become stable. Self-test routines must delay at least this time before attempting to use the controller for either internal or external testing.
5. The CAE bit (loss of carrier) may be set in the transmit descriptor when a packet is sent in internal loopback mode. When the LANCE is operating in internal loopback mode and a transmission is attempted with a non-matching address, the LANCE will correctly reject that packet. If the next operation is an internal loopback transmission without first resetting the LANCE, the packet will not be sent and LCAR will be set in the transmit descriptor for that packet. The receive descriptor will still be owned by the LANCE. To avoid this problem, the LANCE should be reinitialized after each internal loopback packet.
6. The ONE flag is occasionally set in a transmit descriptor after a late collision. The LANCE does not attempt a retransmission even though ONE may be set. The host should disregard ONE if the COE flag is also set.
7. The chip's internal copy of NICSR1 may become invalid when the chip is stopped. The NICSR1 and NICSR2 registers should always be loaded prior to setting INIT to initialize the LANCE chip.
8. Attempting an external loopback test on a busy network can cause a silo pointer misalignment if a transmit abort occurs while the chip was preparing to transmit the loopback packet. The resulting retransmission may cause the transmitter enable circuit to hang, and the resulting illegal length transmission must be terminated by the jabber timer in the transceiver. It is unlikely that there will be a corrupted receive buffer because the reception that caused the transmit abort will usually not pass address recognition.

Since external loopback is a controlled situation it is possible to implement a software procedure to detect a silo pointer misalignment problem and prevent continuous transmissions. Since the test is being done in loopback the exact length and contents of the receive packet are known; thus the software can determine whether the data in the receive buffer has been corrupted.

On transmission, the diagnostic software should allow up to 32 retries before a hard error is flagged. This is not to say that 32 errors are allowed for each condition; the sum of all errors encountered in the test should not exceed 32. The diagnostic software should expect to get a transmit done interrupt with 1 millisecond of passing the transmit packet to the LANCE. If this does not occur, it should reset the LANCE and retry the test. This prevents a continuous transmission (babble) longer than the longest legal packet in case the LANCE has become hung.

9. When the chip is in internal loopback mode and a CRC error is forced, a framing error will also be indicated along with the CRC error. In external loopback, when a CRC error is forced only that error is indicated; a framing error is indicated only if the LANCE actually receives extra bits.
10. When transmit data chaining, a buffer error will be set in the current transmit descriptor if a late collision or retry error occurred while the LANCE was still transmitting data from the previous buffer. The BUE error in this case is an invalid error indication and should be ignored. BUE is valid only when UFE is also set.
11. When the host program sets up a packet for transmission in chained buffers, it should set the OWN bits in all the transmit buffers except the first one (that is, the one containing the STP bit), and then as its last act set the OWN bit in the first descriptor. Once that bit is set, the LANCE will start packet transmission and may encounter an underflow error if the subsequent descriptors for the packet are not available.
12. Do not set INIT and STRT in NICSRO at the same time. After stopping the chip, first set INIT and wait for IDON, then set STRT. If both are set at once, corrupt transmit or receive packets can be generated if RENA becomes true during the initialization process.

3.10 Mass Storage Interface

The KN210 contains a DSSI bus interface which is implemented by the SII chip and four 32K × 8 static RAMs, and is located on the KN210 I/O module. The interface allows the KN210 to transmit packets of data to, and receive packets of data from, up to seven other DSSI devices (typically RF type disk drives and TF type streaming tape drives). The KN210 also provides DSSI bus termination.

This interface contains 27 registers and 128 Kbytes of 32 bit wide RAM (MSI buffer RAM). The SII chip transfers data between the DSSI bus and the MSI buffer RAM, and the processor transfers data between the MSI buffer RAM and main memory.

3.10.1 DSSI Bus Overview

Some of the major characteristics of the DSSI bus are:

- Eight bit data path
- Eight devices supported
- Parity checking
- Distributed arbitration
- Synchronous operation
- Maximum bandwidth of 4M bytes/sec

Communication on the DSSI bus is limited to two devices at a time. Each device has a unique ID assigned to it.

When two devices communicate on the DSSI bus, one acts as the initiator, the other as the target. The initiator is the device that starts a DSSI bus transaction. The target device controls the remainder of the DSSI bus transaction. The direction of data flow is from the initiator to the target.

A DSSI bus transaction consists of six phases:

1. **WAIT**—During this phase, the initiator waits for the bus to become free.
2. **ARBITRATION**—During this phase, control of the bus is taken by the initiator with the highest ID.
3. **SELECTION**—During this phase, the initiator tries to make a logical connection with the target.

3-118 Architecture

4. **COMMAND OUT**—During this phase, the initiator sends the six bytes of command information specified in the command block to the target (Section 3.10.5).
5. **DATA OUT**—During this phase, the initiator sends from one to 4 Kbytes of data to the target.
6. **STATUS IN**—During this phase, the target sends one byte of status information on the transaction to the initiator. The initiator writes this byte to the status word in the command block.

A block diagram of DSSI bus sequences, showing these six phases, is given in Figure 3-72.

MA-X0075-88

Figure 3-72 DSSI Bus Sequences

The normal path follows vertically downward. Exception paths are listed below:

1. The initiator arbitrates and loses.
2. The target failed to respond or responded with an unexpected bus phase.
3. The operation was timed out or the target responded with unexpected phase.
4. The target detected a parity error or information mismatch in the command, or the target did not have any buffer space available.
5. The operation was timed out or the target responded with an unexpected phase.

3.10.2 Target Operation

When the KN210 is functioning as a target device, the SII chip expects receive buffers to be established in the 128 Kbyte MSI buffer RAM (addresses $1010\ 0000_{16}$ through $101F\ FFFF_{16}$). Receive buffers must be set up by the processor and start on 8-byte boundaries. These buffers consist of a command block (Section 3.10.5) and a receive data block. These buffers are linked together by the first word in the command block, and the MSI_TLP register is used to point to the first buffer in the list (Figure 3–73).

During target operation, the SII chip uses the MSI_TLP register to determine the address of the next free receive buffer to be used for this DSSI bus transaction. As the SII chip fills the buffer, it will reload the MSI_TLP for the next target transaction with the buffer's *thread word* (the first word in the command block). The target then places the DSSI bus in the status in phase, sends a status byte to the initiator and updates the status byte in its buffer's command block.

3-120 Architecture

Figure 3-73 Target Operation

3.10.3 Initiator Operation

When the KN210 is functioning as initiator device, the SII chip expects transmit buffers to be established in the 128 Kbyte MSI BUFFER RAM (addresses $1010\ 0000_{16}$ through $101F\ FFFF_{16}$). These buffers must be set up by the processor and start on 8-byte boundaries. These buffers consist of a command block (Section 3.10.5) and a transmit data block. These buffers are linked together by the first word in the command block, and the MSI_ILP register is used to point to the first buffer in the list.

3.10.3.1 Transmit Data Segment Links

The transmit data block is broken into one or more segments. These segments need not reside in contiguous locations in the MSI buffer RAM and are connected together by the link. Pictorially, the link appears as shown in Figure 3–74.

Figure 3–74 Transmit Data Segment Links

MSI Link Word 0

A definition of the bit fields of MSI link word 0 (MSILW0) is given below.

Data Bit	Definition
MSILW0 <15>	LNK. When set, this bit indicates that there is a data segment following the next one. When clear the next data segment is the last in this data block.
MSILW0 <14:0>	Length of next segment. This field contains the number of bytes in the next data segment.

MSI Link Word 1

A definition of the bit field of MSI link word 1 (MSILW1) is given below.

Data Bit	Definition
MSILW1 <15:0>	Address of next segment. This field contains bits <17:2> of the next 8-byte aligned data segment.

Each segment of data must be preceded by the above described link. The number of linked segments is only limited by the maximum size of the data block (4 Kbytes) (Figure 3-75).

When the KN210 is the initiator, the SII chip uses the MSI_IPL register to determine the address of the transmit buffer to be used for this DSSI transaction. As the SII chip is processing a transmit buffer, it loads an internal register with the second word of the link word as long as the LNK is enabled. This chaining continues until a LNK value of zero is encountered. The SII will then transfer the next segment and deposit the status of the entire transfer in the status area of the command block. The MSI_IPL register is then loaded with the buffer's thread word (the first word in the command block) for the next initiator operation. If an error of any kind occurs during the processing of a transmit buffer, the SII will stop the transmit operation by clearing the output enable bit MSI_DSCTRL <14>.

3.10.4 Adding to a Buffer List

The following enumerates the method required to dynamically add new buffers to the MSI_TLP and MSI_IPL lists:

1. Fill in the new buffer command block and ensure that the MSB of the status word is zero.
2. Make the thread word of the new buffer zero.
3. Replace the thread word of the last item on either the MSI_IPL or MSI_TLP list with the new thread word, pointing to the new buffer.
4. If the MSI_IPL or MSI_TLP is zero, load it with the address of the new buffer.

Figure 3-75 Initiator Operation

3.10.5 Mass Storage Interface Command Block (MSICB)

The MSI command block is a 12 byte data structure that the processor has to build at the start of all transmit and receive buffers in the MSI buffer RAM. The format for the MSI command block is shown in Figure 3-76.

Figure 3-76 MSI Command Block

3.10.5.1 MSI Command Block Word 0

Word 0 of the MSI command block (MSICBW0), also referred to as the thread word, resides in the MSI buffer RAM at the base address of the MSI command block. The thread word contains bits <17:2> of the base address of the next buffer. Bit 0 of this field must always be set to 0, since buffers must start on an 8-byte boundary. A thread word of zero indicates that there are no more buffers. The format for MSICBW0 is shown in Figure 3-77.

Figure 3-77 MSI Command Block Word 0

3.10.5.2 MSI Command Block Word 1

Word 1 of the MSI command block (MSICBW1), also referred to as the status word, resides in the MSI buffer RAM at the base+2 address of each MSI command block. This word indicates the status of the current DSSI transaction and is used by the processor to find out which buffers the SII chip has finished processing. The format for MSICBW1 is shown in Figure 3-78.

Figure 3-78 MSI Command Block Word 1

The bit fields in MSICBW1 represent the following:

Data Bit	Definition
MSICBW1<15>	(DNE) Done. When set, this indicates that the SII chip has used this buffer (either successfully or not). When clear the SII chip has not used this buffer. Note, if this bit is set when the SII chip begins processing a buffer, the buffer is not used.
MSICBW1<14:8>	Unused.
MSICBW1<7>	(RST). Reset. When set, a DSSI device reset the DSSI bus during this buffer's transaction.

NOTE

If a DSSI bus reset occurred before the SII chip reached status in phase, the SII chip will clear MSI_DSCTRL <7> (output enable bit) and interrupt the processor without writing any status.

Data Bit	Definition
MSICBW1<6>	(TMO) Timeout. When set, one of the MSI_DSTMO timers has expired. NOTE If the timeout occurred before the SII chip reached status in phase, the SII chip will clear MSI_DSCTRL<7> (output enable bit) and interrupt the processor without writing any status.
MSICBW1<5>	(XSM) Checksum. When set, the received checksum does not agree with that computed by the SII chip. Note the XSM bit is only valid when the KN210 is a target.
MSICBW1<4>	(BPH). Bad phase. When set, an illegal DSSI phase was entered by the target. Note the BPH bit is only valid when the KN210 is the initiator.
MSICBW1<3>	(STT) Status. When set, ACK was not returned by the target. Note the STT bit is only valid when the KN210 is the initiator.
MSICBW1<2>	(PHS) Phase. When set, the DSSI bus phase changed before the initiator expected. Note the PHS bit is only valid when the KN210 is the initiator.
MSICBW1<1>	(DSA). DSSI. When set, the target detected an error in the command bytes. Note the DSA bit is only valid when the KN210 is the target.
MSICBW1<0>	(PAR). Parity. When set, a parity error on the DSSI bus was detected.

Please note that the following cases will not cause status to be written in memory:

- DSSI bus reset occurs before status in phase is reached.
- Initiator selects a non-existent device (initiator timeout will cause a DSSI bus reset).
- Target disconnects from the DSSI bus before status in phase is reached.

3.10.5.3 MSI Command Block Word 2

Word 2 of the MSI command block (MSICBW2), also referred to as the command word, resides in the MSI buffer RAM at the base+4 address of each MSI command block. This word contains information regarding the transfer. The format for MSICBW2 is shown in Figure 3-79.

Figure 3-79 MSI Command Block Word 2

The bit fields in this memory word represent the following:

Data Bit	Definition
MSICBW2<15>	(IE). Interrupt enable. When set, the SII chip will interrupt the KN210 upon the completion (successful or not) of this transaction. When clear the SII chip will not generate an interrupt. Interrupts are posted on IRQ2 at IPL16 with a vector offset of C5 ₁₆ .
MSICBW2<14:3>	Unused.
MSICBW2<2:0>	(DEST ID). Destination ID. The ID of the target to be selected. This field is only used when the KN210 is the initiator.

3.10.5.4 MSI Command Block Words 3-5

Words 3-5 of the MSI command block (MSICBW3-5), also referred to as command bytes, reside in the MSI buffer RAM at the base+6 through base+10 address of each MSI command block. These 6 bytes are sent out during the command out phase by the initiator. Some of the information contained in these bytes are:

- The target and initiator IDs
- The number of data bytes which will be transferred by the initiator in the data out phase
- The DSSI opcode.

3.10.6 MSI Registers

The following is a description of the 16 registers needed to control the SII chip during DSSI bus operations.

NOTE

The other 11 registers are not used during DSSI operations and should not be accessed.

3.10.6.1 MSI Control and Status Registers

These five registers are used to configure, control and monitor the SII chip.

MSI Control/Status Register

The MSI control/status register (MSI_CSR) (address 1008 460C₁₆) contains control and status information about the general operation of the SII chip in regard to the DSSI bus, including various enable bits. The format of the mass storage control/status register is shown in Figure 3-80.

Figure 3-80 MSI Control/Status Register

Data Bit	Definition
MSI_CSR <31:5>	Unused. Reads return undefined results, writes have no effect.
MSI_CSR <4:3>	(MBZ) Must be zero. Read/Write. These bits must read as zero and be written as zero.
MSI_CSR <2>	(SLE) Selections. Read/Write. When set, the SII chip will respond to selections. When clear the SII chip will not respond to an initiator trying to select it. Cleared on power-up, or on the assertion of IORESET.
MSI_CSR <1>	(PCE) Parity check. Read/Write. When set, the SII chip reports parity errors. When clear the SII chip will continue to check parity but will not report any errors during the status in phase. Cleared on power-up, or on the assertion of IORESET.
MSI_CSR <0>	(IE) Interrupt enable. Read/Write. When set, interrupts are enabled. The SII chip posts interrupts when an error occurs or at the end of a transaction (successful or not). Interrupts are posted on IRQ2 at IPL16 with an offset of C5 ₁₆ . When clear interrupts are disabled. Cleared on power-up, or on the assertion of IORESET.

MSI DSSI Control Register

The MSI DSSI control register (MSI_DSCTRL) (address 1008 4644₁₆) contains information to control the SII chip. The format of the mass storage interface DSSI control register is shown in Figure 3-81.

Figure 3-81 MSI DSSI Control Register

Data Bit	Definition
MSI_DSCTRL <31:16>	Unused. Reads return undefined results. Writes have no effect.
MSI_DSCTRL <15>	(DSE) DSSI Enable. Read/Write. This bit must be set to one by the processor for the SII chip to work on a DSSI bus. This bit is cleared by the SII chip if: the SII chip selects or is selected by a non-DSSI device, the SII chip is selected with Attention. It is also cleared on power-up, or on the assertion of IORESET.
MSI_DSCTRL <14>	(OUT) Output enable. Read/Write. When set, the SII chip is enabled to send transmit buffers. This bit is cleared by the SII chip if: the MSI_IPL becomes zero, the initiator timer MSI_DSTMO <3:0> expires, or a transmit buffer is not terminated with ACK. It is also cleared on power-up, or on the assertion of IORESET.
MSI_DSCTRL <13:8>	Unused. Reads return undefined results, writes have no effect.

Data Bit	Definition
MSI_DSCTRL <7>	(CH7) Channel 7. Read/Write. This bit is used to determine if device 7 is a DSSI device. This bit must be set to one by the processor. Cleared on power-up, or on the assertion of IORESET.
MSI_DSCTRL <6>	(CH6) Channel 6. Read/Write. This bit is used to determine if device 6 is a DSSI device. This bit must be set to one by the processor. Cleared on power-up, or on the assertion of IORESET.
MSI_DSCTRL <5>	(CH5) Channel 5. Read/Write. This bit is used to determine if device 5 is a DSSI device. This bit must be set to one by the processor. Cleared on power-up, or on the assertion of IORESET.
MSI_DSCTRL <4>	(CH4) Channel 4. Read/Write. This bit is used to determine if device 4 is a DSSI device. This bit must be set to one by the processor. Cleared on power-up, or on the assertion of IORESET.
MSI_DSCTRL <3>	(CH3) Channel 3. Read/Write. This bit is used to determine if device 3 is a DSSI device. This bit must be set to one by the processor. Cleared on power-up, or on the assertion of IORESET.
MSI_DSCTRL <2>	(CH2) Channel 2. Read/Write. This bit is used to determine if device 2 is a DSSI device. This bit must be set to one by the processor. Cleared on power-up, or on the assertion of IORESET.
MSI_DSCTRL <1>	(CH1) Channel 1. Read/Write. This bit is used to determine if device 1 is a DSSI device. This bit must be set to one by the processor. Cleared on power-up, or on the assertion of IORESET.
MSI_DSCTRL <0>	(CH0) Channel 0. Read/Write. This bit is used to determine if device 0 is a DSSI device. This bit must be set to one by the processor. Cleared on power-up, or on the assertion of IORESET.

MSI DSSI Connection Register

The MSI DSSI connection register (MSI_CSTAT) (address 1008 4648₁₆) contains interrupt status related to SII chip connections. The format of the mass storage interface DSSI connection register is shown in Figure 3-82.

Figure 3-82 MSI DSSI Connection Register

Data Bit	Definition
MSI_CSTAT <31:16>	Unused. Reads return undefined results, writes have no effect.
MSI_CSTAT <15>	(CI) Composite interrupt. Read only. This bit is the composite error bit of the MSI_CSTAT register. It is the logical OR of bits MSI_CSTAT <13:11> and MSI_CSTAT <9:7>. When set, the processor will be interrupted on IRQ2 at IPL16 with a vector offset of C5 ₁₆ if interrupts are enabled. Cleared on power-up, or on the assertion of IORESET.
MSI_CSTAT <14>	(UNU) Unused. Reads return undefined results, writes have no effect.
MSI_CSTAT <13>	(RST) Reset asserted. Read/Write one to clear. When set, the DSSI bus was reset by one of the eight DSSI devices. The SII chip will automatically disconnect itself from the bus and interrupt the processor on IRQ2 at IPL16 with an offset of C5 ₁₆ . This bit is write one to clear and is also cleared on power-up, or on the assertion of IORESET.
MSI_CSTAT <12>	(BER) Bus error. Read/Write one to clear. This bit is set to one on any of the following conditions: <ul style="list-style-type: none"> <li data-bbox="594 1184 813 1211">• Buffer overflow <li data-bbox="594 1234 906 1262">• Req/Ack offset exceeded <li data-bbox="594 1285 870 1312">• Illegal phase change <p data-bbox="594 1329 1235 1409">While this bit is asserted, the SII chip will not receive or transmit data. This bit is write one to clear and is also cleared on power-up, or on the assertion of IORESET.</p>

Data Bit	Definition
MSI_CSTAT <11>	<p>(OBC) OUT_EN Bit cleared. Read/Write one to clear. This bit is set to one on any of the following conditions:</p> <ul style="list-style-type: none"> • The SII chip has received RSTIN. (The DSSI bus has been reset.) • The MSI_DSTMO (MSI_DSTMO <3:0> or MSI_DSTMO <7:4>) has expired. • As an initiator, the attached target disconnects unexpectedly. <p>This bit is write one to clear and is also cleared on power-up, or on the assertion of IORESET.</p>
MSI_CSTAT <10>	<p>(TZ) Target pointer zero. Read only. When set, the MSI_TLP register contains a value of zero. This bit is set on power-up, or on the assertion of IORESET.</p>
MSI_CSTAT <9>	<p>(BUF) Buffer service. Read/Write one to clear. When set, the SII chip has begun processing a transmit buffer destined for non-DSSI device. Note, this bit should always be zero since all devices must be DSSI. This bit is write one to clear and is also cleared on power-up, or on the assertion of IORESET.</p>
MSI_CSTAT <8>	<p>(LDN) List element done. Read/Write one to clear. When interrupts are enabled, this bit is set if the SII chip has completed a buffer, successfully or not. This bit is write one to clear and is also cleared on power-up, or on the assertion of IORESET.</p>
MSI_CSTAT <7>	<p>(SCH) State change. Read/Write one to clear. Set if MSI_DSCTRL <15> is cleared causing the SII chip to leave DSSI mode. This bit is write one to clear and is also cleared on power-up, or on the assertion of IORESET.</p>
MSI_CSTAT <6>	<p>(CON) Connected. Read only. When set, the SII is connected to another device on the DSSI bus. Clear while the SII chip is not connected to another device on the DSSI bus. This bit is cleared on power-up, or on the assertion of IORESET.</p>

Data Bit	Definition
MSI_CSTAT <5>	(DST) Destination. Read only. When set, the SII is the destination of the current transaction. In other words, this bit is set if the SII chip was selected by another device on the DSSI bus. This bit is cleared on power-up, or on the assertion of IORESET.
MSI_CSTAT <4>	(TGT) Target. Read only. When set, the SII chip is operating as a target during the current transaction. This bit is cleared on power-up, or on the assertion of IORESET.
MSI_CSTAT <3>	(SWA) Selected with attention. Read only. When set, the SII chip was selected with attention. This bit is write one to clear and is also cleared on power-up, or on the assertion of IORESET.
MSI_CSTAT <2>	(SIP) Selection in progress. Read only. When set, the SII chip is currently in a selection process. This is useful in determining if the desired target is unavailable. This bit is write one to clear and is also cleared on power-up, or on the assertion of IORESET.
MSI_CSTAT <1>	(LST) Lost. Read only. When set, the SII lost arbitration. It is cleared by the SII chip when it begins a selection process and on power-up or on the assertion of IORESET.
MSI_CSTAT <0>	(MBZ) Must be zero. Read Only. This bit will be read as zero.

MSI ID Register

The MSI ID register (MSI_ID) (address $1008\ 4610_{16}$) contains the three bit ID number of the KN210 on the DSSI bus. This value is placed on the DSSI bus during the selection phase so the target knows who selected it. The format of the mass storage interface ID register is shown in Figure 3–83.

Figure 3-83 MSI ID Register

Data Bit	Definition
MSI_ID <31:16>	Unused. Reads return undefined results, writes have no effect.
MSI_ID <15>	(I/O) Input/Output. Read/Write. When set, the KN210's ID is determined by MSI_ID <2:0>. When clear the KN210's ID is determined by on board jumpers and MSI_ID <2:0> will reflect the one's complement of the KN210's DSSI ID. Cleared on power-up, or on the assertion of IORESET. Note that if this bit is cleared, writing to this register has no effect.
MSI_ID <14:3>	Unused. Reads return undefined results, writes have no effect.
MSI_ID <2:0>	KN210 Bus ID. Read/Write. When MSI_ID <31> is clear (the normal operation configuration), this field contains the DSSI ID of the KN210, as determined by the on board jumpers. When MSI_ID <31> is set, any DSSI ID value may be input (used, for example, to temporarily override the on board jumpers for test or diagnostic purposes). Indeterminate on power-up, or on the assertion of IORESET.

MSI DSSI Timeout Register

The MSI DSSI timeout register (MSI_DSTMO) (address 1008 461C₁₆) contains the timeout values of the SII chip for both the initiator and target roles. Also contained in this register is a single enable bit that governs both timers. The format of the MSI DSSI timeout register is shown in Figure 3-84.

Figure 3-84 MSI DSSI Timeout Register

Data Bit	Definition
MSI_DSTMO <31:16>	Unused. Reads return undefined results, writes have no effect.
MSI_DSTMO <15>	(ENA) Enable. Read/Write. When set, both the DSSI target and DSSI initiator timers are enabled. When clear, both the DSSI target and DSSI initiator timer are disabled. Cleared on power up, or on the assertion of IORESET.
MSI_DSTMO <14:8>	Unused. Reads return undefined results, writes have no effect.
MSI_DSTMO <7:4>	Target timeout value. Read/Write. This field contains the number of 200 microsecond intervals which may elapse while the KN210 is the target. The timer starts from the point when the KN210 was selected ends at the next observed bus free phase. Cleared on power_up, or on the assertion of IORESET.
MSI_DSTMO <3:0>	Initiator timeout value. Read/Write. This field contains the number of 200 microsecond intervals which may elapse, from the last observed bus free phase, until the next observed bus free phase, while the KN210 is in the initiator role; or the number of 200 microsecond intervals which may elapse before the KN210, acting as a potential initiator, detects a bus free phase. Should the timer expire under either of these two conditions the SII chip will assert a DSSI bus reset. Cleared on power_up, or on the assertion of IORESET.

3.10.6.2 List Pointer Registers

These are the two registers used as address pointers for next incoming and outgoing data buffers.

MSI Target List Pointer Register

The MSI target list pointer register (MSI_TLP) (address 1008 463C₁₆) contains the address to which the SII chip will write the next free receive buffer. The SII chip will automatically reload the register with the receive buffer's thread word upon completion of the current transaction. Note this register must contain bits <17:2> of an 8-byte aligned address, therefore bit 0 will always be zero. The SII chip will interpret an address of 0000₁₆ as the end of a linked list. The format of the MSI target list pointer register is shown in Figure 3-85.

Figure 3-85 MSI Target List Pointer Register

Data Bit	Definition
MSI_TLP <31:16>	Unused. Reads return undefined results, writes have no effect.
MSI_TLP <15:1>	Address of next incoming buffer. Read/Write. This field contains bits 17:3 of the 8-byte aligned address to where the SII chip will find the next free receive buffer. Cleared on powerup, or on the assertion of IORESET.
MSI_TLP <0>	(MBZ) Must be zero. Read/Write. This bit is read as zero and must be written as zero.

NOTE

This register can only be written by the processor when the register is zero or MSI_DSCTRL <15> (DSE) is clear; all other attempts to write to this register have no effect.

MSI Initiator List Pointer Register

The MSI initiator list pointer register (MSI_IPL) (address 1008 4640₁₆) contains the address from which the SII chip will find the next transmit buffer. The SII chip will automatically reload this register with the transmit buffer's thread word upon completion of the current transaction. Note this register must contain bits <17:2> of an 8-byte aligned address, therefore bit 0 must always be zero. The SII chip will interpret an address of 0000₁₆ as the end of a linked list. The format of the MSI initiator list pointer register is shown in Figure 3–86.

Figure 3–86 MSI Initiator List Pointer Register

Data Bit	Definition
MSI_IPL <31:16>	Unused. Reads return undefined results, writes have no effect.
MSI_IPL <15:1>	Address of next outgoing buffer. Read/Write. This field contains bits 17:3 of the 8-byte aligned address of where the SII chip will find the next transmit buffer. Cleared on power_up, or on the assertion of IORESET.
MSI_TLP <0>	(MBZ) Must be zero. Read/Write. This bit will read as zero and must be written as zero.

NOTE

This register can only be written to by the processor when the register is zero or MSI_DSCTRL <15> (DSE) is clear; all other attempts to write to this register have no effect.

3.10.6.3 Diagnostic and Test Registers

This group of registers is used for test and diagnostic purposes only. They should never be used during normal operation.

MSI Diagnostic Control Register

The MSI diagnostic control register (MSI_DICTRL) at (address 1008 4654₁₆) allows the SII chip to be placed in one of three diagnostic test modes. The format of the MSI diagnostic control register is shown in Figure 3-87.

Figure 3-87 MSI Diagnostic Control Register

Data Bit	Definition
MSI_DICTRL<31:4>	Unused. Reads return undefined results. Writes have no effect.
MSI_DICTRL<3>	(ITM) Internal test mode. Read/Write. When set, the values written to MSI_DR0, MSI_DR1 and MSI_DR2 are to be looped back into the chip. This will enable the processor to insert test vectors into the chip during power-up diagnostics. Note that the MSI_DICTRL<1> (ELM) must be deasserted for this test to be meaningful. This bit is cleared on power-up, or on the assertion of IORESET.
MSI_DICTRL<2>	(PRE) Port enable. Read/Write. When set, the off-chip drivers to the DSSI port are enabled. After a reset, the KN210 will be disconnected from the bus (this bit will be zero). The primary purpose of this bit is to allow SII chip diagnostics to run without affecting the rest of the DSSI bus (PRE=0). This bit is cleared on power-up, or on the assertion of IORESET.

Data Bit	Definition
MSI_DICTRL<1>	(ELM) External loopback mode. Read/Write. When set, the SII chip is in external loopback mode. In this mode, MSI_DR0, MSI_DR1 and MSI_DR2 are used to directly control the DSSI data and control lines, as well as the external bus transceiver. Note an external loopback connector must be in place when using this test mode. This bit is cleared on power-up, or on the assertion of IORESET.
MSI_DICTRL<0>	(TST) Test mode. Read/Write. When set, to one (1), the SII chip is in test mode. This enables the user to replace the 20 MHz clock. The new clock is pulsed each time the MSI_CLOCK register is written. This bit is cleared on power-up, or on the assertion of IORESET.

MSI Diagnostic Register 0

The MSI diagnostic register 0 (MSI_DR0) (address 1008 4600₁₆) is used during internal and external loopback diagnostic tests. The fields in this register are used to emulate the data lines of the DSSI. The format of MSI diagnostic register 0 is shown in Figure 3–88.

Figure 3–88 MSI Diagnostic Register 0

Data Bit	Definition
MSI_DR0 <31:9>	Unused. Reads return undefined results, writes have no effect.
MSI_DR0 <8>	(PTY) Parity. Read/Write. This bit contains the parity bit for the data byte MSI_DR0 <7:0>. Indeterminate on power-up, or on the assertion of IORESET. Note, parity checking is only enabled if MSI_CSR <1> PCE is set to 1. The SII chip uses odd parity checking.

Data Bit	Definition
MSI_DR0 <7:0>	(DATA) Read/Write. This field contains the current byte on the data bus. Indeterminate on power-up, or on the assertion of IORESET.

NOTE

MSI_DR0 should NOT be used during normal operation.

MSI Diagnostic Register 1

The MSI diagnostic register 1 (MSI_DR1) (address 1008 4604₁₆) is used during internal and external loopback tests. In external loopback mode an external loopback connector in place allows values written into MSI_DR0 to be read back in MSI_DR1 and values written into MSI_DR1 to be read back in MSI_DR0. In internal loopback mode it acts as the DSSI bus emulating some of the DSSI control lines. Note that all the control lines are asserted high in internal loopback test mode. For more information on the SII chip modes see the description of the MSI diagnostic control register. The format of MSI diagnostic register 1 is shown in Figure 3-89.

Figure 3-89 MSI Diagnostic Register 1

Data Bit	Definition
MSI_DR1 <31:9>	Unused. Reads return undefined results, writes have no effect.
MSI_DR1 <8>	(BSY) Busy. Read/Write. In internal loopback test mode, MSI_DICTRL <3> set, this bit emulates the DSSI BSY bus signal. In external loopback mode this bit is linked to MSI_DR0 <8> (PTY) for driver testing. Indeterminate on power-up, or on the assertion of IORESET.
MSI_DR1 <7>	(SEL) Select. Read/Write. In internal loopback test mode, MSI_DICTRL <3> set, this bit emulates the DSSI SEL bus signal. In external loopback mode this bit is linked to MSI_DR0 <7> (DATA <7>) for driver testing. Indeterminate on power-up, the negation of DCOK when SCR <7>.
MSI_DR1 <6>	(RST) Reset. Read/Write. In internal loopback test mode, MSI_DICTRL <3> set, this bit emulates the DSSI RST bus signal. In external loopback mode this bit is linked to MSI_DR0 <6> (DATA <6>) for driver testing. Indeterminate on power-up, or on the assertion of IORESET.
MSI_DR1 <5>	(ACK) Acknowledge. Read/Write. In internal loopback test mode, MSI_DICTRL <3> set, this bit emulates the DSSI ACK bus signal. In external loopback mode this bit is linked to MSI_DR0 <5> (DATA <5>) for driver testing. Indeterminate on power-up, or on the assertion of IORESET.
MSI_DR1 <4>	(REQ) Request. Read/Write. In internal loopback test mode, MSI_DICTRL <3> set, this bit emulates the DSSI REQ bus signal. In external loopback mode this bit is linked to MSI_DR0 <4> (DATA <4>) for driver testing. Indeterminate on power-up, or on the assertion of IORESET.
MSI_DR1 <3>	(ATN) Attention. Read/Write. In internal loopback test mode, MSI_DICTRL <3> set, this bit emulates the DSSI ATN bus signal. In external loopback mode this bit is linked to MSI_DR0 <3> (DATA <3>) for driver testing. Indeterminate on power-up, or on the assertion of IORESET.
MSI_DR1 <2>	(MSG) Message. Read/Write. In internal loopback test mode, MSI_DICTRL <3> set, this bit emulates the DSSI MSG bus signal. In external loopback mode this bit is linked to MSI_DR0 <2> (DATA <2>) for driver testing. Indeterminate on power-up, or on the assertion of IORESET.

Data Bit	Definition
MSI_DR1 <1>	(C/D) Control/Data. Read/Write. In internal loopback test mode, MSI_DICTRL <3> set, this bit emulates the DSSI C/D bus signal. In external loopback mode this bit is linked to MSI_DR0 <1> (DATA <1>) for driver testing. Indeterminate on power-up, or on the assertion of IORESET.
MSI_DR1 <0>	(I/O) Input/Output. Read/Write. In internal loopback test mode, MSI_DICTRL <3> set, this bit emulates the DSSI I/O bus signal. In external loopback mode this bit is linked to MSI_DR0 <0> (DATA <0>) for driver testing. Indeterminate on power-up, or on the assertion of IORESET.

NOTE

The data written to this register in internal test mode may differ from that read back from it, since only certain bits are driven when configured as a target or initiator. See the register description of MSI_DR2 for more information on the internal test mode.

MSI Diagnostic Register 2

The MSI diagnostic register 2 (MSI_DR2) (address 1008 4608₁₆) is used by diagnostics to directly control the DC563 transceiver chip. The format of the MSI diagnostic register 2 is shown in Figure 3-90.

Figure 3-90 MSI Diagnostic Register 2

Data Bit	Definition
MSI_DR2 <31:4>	Unused. Reads return undefined results. Writes have no effect.
MSI_DR2 <3>	(IGS) Read/Write. This bit enables the DSSI bus drivers for ACK and ATN, placing the SII chip in the initiator role. Cleared on power-up, or on the assertion of IORESET.
MSI_DR2 <2>	(TGS) Read/Write. When set, this bit enables the DSSI bus drivers for I/O, C/D, MSG and ATN, placing the SII chip in the target role. Cleared on power-up, or on the assertion of IORESET.
MSI_DR2 <1>	(ARB) Arbitrate. Read/Write. This bit enables the decoding of ID0..ID2, putting the SII chip in the arbitration phase. Cleared on power-up, or on the assertion of IORESET.
MSI_DR2 <0>	(SBE) Read/Write. When set, the DC563 transceiver drives the DSSI data bus and parity lines. Cleared on power-up, or on the assertion of IORESET.

NOTE

Special care should be taken when writing to the MSI_DR2 register to avoid disturbing the DSSI bus during power-up diagnostics. This register should only be used when an external loopback connector is in place and not during normal operation.

MSI Clock Control Register

Writing to the MSI clock control register (MSI_CLOCK), address 1008 4658₁₆, generates a pulse which, in test mode (MSI_DICTRL <0> set to one), replaces the 20 MHz clock input. This can be used to allow the CPU to observe and sequence the various state machines inside the SII chip. The format of the MSI clock control register is shown in Figure 3–91.

Figure 3–91 MSI Clock Control Register

Data Bit	Definition
MSI_CLOCK <31:0>	Unused. Write only. Writing to this register generates a pulse which, in test mode (MSI_DICTRL <0> set to one), replaces the 20 MHz clock input.

MSI Internal State Registers (0-3)

These registers, at addresses 1008 465C₁₆, 1008 4660₁₆, 1008 4664₁₆, and 1008 4668₁₆, reflect the status of the SII chip's internal state machine when used in test mode (that is, MSI_DICTRL <0> set to one).

4

KN210 Firmware

This chapter describes the functional operation of the KN210 firmware. This firmware is used to provide an easy means to configure and bootstrap a KN210-based system and to detect and isolate system malfunctions.

The firmware described in this chapter resides on the KN210 processor module. The KN210 is an R3000-based Q22-bus CPU module set.

The KN210 provides a maximum of 256 Kbyte of EPROM for the firmware. The firmware resides on the KN210 processor module in two 128 Kbyte EPROMs (one per processor), which are arranged as 32-bit words and located at the R3000 restart location at physical address 1FC0 0000 and the diagnostic processor restart location at physical address 2004 0000.

The firmware uses the KN210 module LEDs and a console terminal to communicate diagnostic progress, display error conditions, and indicate the current mode of operation. Additionally, the console terminal is used as the primary operator interface when in console I/O mode.

NOTE

A console terminal is not required for operation, since the module set can be configured to bootstrap automatically. However, in most scenarios a console terminal is a recommended part of a standard configuration.

The KN210 firmware runs standalone, and does not require other software products for normal operation in *console I/O mode*. However, in most situations an operating system (or other image) is loaded in KN210 local memory and control is transferred to it. When the operating system is running, the processor is in *program I/O mode*. (The terms *console I/O mode* and *program I/O mode* refer to the context and usage of the console terminal.)

The KN210 supports ULTRIX. Additionally, the firmware supports bootstrap of MDM and other diagnostics images. Also, the console supports communication with an APT host in manufacturing environments through the console serial line.

KN210 firmware is an integral part of the module and does not require installation or support services.

4.1 Firmware Capabilities

The KN210 firmware provides the following services:

- Diagnostics that test all components on the board and verify the module is working correctly.
- Automatic/manual bootstrap of an operating system following processor halts.
- Automatic/manual restart of an operating system following processor halts.
- An interactive command language that allows the user to examine and alter the state of the processor.

The remainder of this chapter describes in detail the functions and external characteristics of the KN210 firmware.

4.2 Power-Up

On power-up, the KN210 firmware performs initial power-up tests, locates and identifies a console device, and performs the remaining diagnostics. Certain power-up actions are dependent on the state of the *Operation* and *Function* switches on the H3602-SA panel. The various functions determined by the *Operation* and *Function* switches are described in Section 4.2.3.

4.2.1 Initial Power-Up Test

At power-up, the diagnostic processor performs the Initial Power-Up Test (IPT). The purpose of the IPT is to verify that the console private NVRAM is valid and if invalid to test and initialize the NVRAM.

If the NVRAM contains invalid data, the IPT will initialize certain non-volatile data, such as the default boot device, to a known state. In any case, the IPT then initializes other data structures and performs a processor initialization. If the mode switch is set to *Test*, the IPT then tests the console serial line.

NOTE

All IPT failures are considered fatal, and the KN210 will hang with a value on the LEDs indicating the point of failure. Refer to Table 4-2 for the meaning of the LEDs.

4.2.2 Locating a Console Device

After the diagnostic processor IPT has completed successfully, the firmware attempts to locate a console device and find out what type of device it is. Normally, this is the device attached to the console serial line. In this case, the firmware will send out a device attributes escape sequence to the console serial line to determine the type of terminal attached and the functions it supports. Terminals that do not respond to the device attributes request correctly are assumed to be hardcopy devices.

NOTE

The KN210 is always an arbiter CPU and the console serial line is unconditionally treated as the system console.

Once a console device has been found, the firmware displays the KN210 banner message, similar to the following:

```
KN210-A V0.0
```

The banner message contains the processor name and the version of the firmware. The letter code in the firmware version indicates whether the firmware is pre-field test (X), field test (T) or an official release (V). The first digit indicates the major release number and the trailing digit indicates the minor release number.

4.2.3 Operation and Function Switches

The operation and function switches tell the KN210 to perform maintenance functions or to continue with the bootstrap process.

The operation switch has three settings: *Normal*, *Maintenance* and *Action*. When in *normal* mode, the function switch positions are *break enabled* and *break disabled*. If in *maintenance mode*, the function switch positions are *break enabled* and *break disabled*. If the position of the operation switch is set to *action*, the function switch positions are *test* and *query*.

4-4 KN210 Firmware

The operation and function switches are described in Table 4-1.

Table 4-1 Operation and Function Switches

Operation Switch Mode	Symbol	Function Switch Position	Symbol
Normal	→	<i>break enabled</i>	⊙
		<i>break disabled</i>	○
Maintenance	⊕	<i>break enabled</i>	⊙
		<i>break disabled</i>	○
Action	{°	<i>test</i>	⊙
		<i>query</i>	○

4.2.3.1 Operation Switch Set to Normal

When the operation switch is set to *normal* the diagnostic processor executes the power up diagnostics. In addition to message text, a countdown is displayed to indicate diagnostic test progress. A successful diagnostic countdown is shown in Example 4-1.

```

Performing normal system tests.
49..48..47..46..45..44..43..42..41..40..39..38..37..36..35..34..
33..32..31..30..29..28..27..26..25..24..23..22..21..20..19..18..
17..16..15..14..13..12..11..10..09..08..07..06..05..04..03..
Tests completed.
Memory Size: xxxxxxxx (0xxxxxxx) bytes
Ethernet Address: xx-xx-xx-xx-xx-xx
    
```

Example 4-1 Normal Diagnostic Countdown

After the diagnostics are executed, control is passed to the R3000 processor through BIT 31 of the select processor register (SPR) in the KN210 processor (Section 4.2.4.1), and the R3000 processor continues with the power-up sequence.

If the operation switch is set to *normal* and the environment variable **bootmode** is set to **a**, the R3000 processor attempts to find a booting device specified by the **setenv** command through the **bootpath** environment variable stored in non-volatile memory (Section 4.3.2).

When the operation switch is set to *normal* and the environment variable **bootmode** is not initialized or it is set to **d**, the R3000 enters console I/O mode and displays the R3000 console prompt (>>) to prompt the user for commands. Refer to Section 4.4 for console command descriptions.

If there are diagnostic failures, a diagnostic register dump is performed similarly to Example 4-2. The remaining diagnostics execute, the countdown continues, and the KN210 enters *maintenance mode* (Section 4.2.3.2). For a detailed description of the register dump refer to Section 4.5.

```
Performing normal system tests.
49..48..47..46..45..44..43..42..41..40..39..38..37..36..35..34..
?34 2 08 FF 00 0000

P1=00000000 P2=00000003 P3=00000031 P4=00000011 P5=00002000
P6=FFFFFFFF P7=00000000 P8=00000000 P9=00000000 P10=2005438F
r0=00114B98 r1=FFFFFFFF r2=2005D2F0 r3=55555555 r4=AAAAAAAA
r5=00000000 r6=AAAAAAAA r7=00000000 r8=00000000 ERF=80000180
33..32..31..30..29..28..27..26..25..24..23..22..21..20..19..18..
17..16..15..14..13..12..11..10..09..08..07..06..05..04..03..
Normal operation not possible.
>>>
```

Example 4-2 Abnormal Diagnostic Countdown

4.2.3.2 Operation Switch Set to Maintenance

When the operation switch is set to *maintenance*, the diagnostic processor executes the power-up diagnostics. In addition to message text, a countdown is displayed to indicate diagnostic test progress. A successful diagnostic countdown is shown in Example 4-1.

After the diagnostics are executed, the KN210 enters *maintenance mode* and displays the diagnostic processor console prompt (>>>) to prompt the user for commands (Section 6.2.7).

4.2.3.3 Operation Switch Set to Action

Two different operations are possible when the operation switch is set to *action*: *test* and *query*.

Function Switch Set to Test

If the function switch is set to *test*, the console serial line external loopback test is executed at the end of the IPT. The purpose of this test is to verify that the console serial line connections from the KN210 through the H3602-SA panel are intact.

NOTE

An external loopback connector should be inserted in the serial line connector on the H3602-SA panel before cycling power to invoke this test.

During this test the firmware toggles between the two states: active and passive, for a few seconds each with each displaying a different number on the LEDs.

During the active state (about 3 seconds long), the LEDs are set to 7. In this state, the firmware reads the baud rate and operation switch, then transmits and receives a character sequence. If the operation switch has been moved from the *action* position, the firmware exits the test and continues as if on a normal power-up.

During the passive state (about 7 seconds long), the LEDs are set to 3.

If at any time the firmware detects an error (parity, framing, overflow or no characters), the firmware hangs with a 7 displayed on the LEDs.

Function Switch Set to Query

If the KN210 designated console device supports DEC multinational character set (MCS) and either the NVRAM data is invalid or the function switch is set to *query* the firmware prompts for the console language. The firmware first displays the language selection menu shown in Example 4–3.

After the language query, the firmware invokes the ROM-based diagnostics, and the firmware enters *normal* mode if the diagnostics are passed successfully. However, if the diagnostics fail, the firmware enters *maintenance* mode.

- 1) Dansk
 - 2) Deutsch (Deutschland/Österreich)
 - 3) Deutsch (Schweiz)
 - 4) English (United Kingdom)
 - 5) English (United States/Canada)
 - 6) Español
 - 7) Français (Canada)
 - 8) Français (France/Belgique)
 - 9) Français (Suisse)
 - 10) Italiano
 - 11) Nederlands
 - 12) Norsk
 - 13) Português
 - 14) Suomi
 - 15) Svenska
- (1..15):

Example 4-3 Language Selection Menu

If no response is received within 30 seconds, the language defaults to English (United States/Canada).

NOTE

This action is only taken if the console device supports DEC MCS. Any console device that does not support DEC MCS, such as a VT100, defaults to English (United States/Canada).

After this inquiry, the firmware proceeds to *autoboot*.

4.2.3.4 LED Codes

In addition to the console diagnostic countdown a hexadecimal value is displayed on the diagnostic LEDs on the processor module and the H3602-SA panel. The purpose of the LED display is to improve fault isolation, when there is no console terminal or when the hardware is incapable of communicating with the console terminal. Table 4-2 lists all LED codes and the associated actions that are performed at power-up. The LED code is changed before the corresponding test or action is performed.

Table 4-2 LED Codes

LED Value	Actions
F	Initial state on power-up, no code has executed
E	Entered ROM, some instructions have executed
D	Waiting for power to stabilize (POK)
C	SSC and ROM tests
B	Diagnostic processor tests
A	R3000 tests
9	CMCTL and memory tests
8	CQBIC (Q22-bus) tests
7	Console loopback tests
6	DSSI subsystem tests
5	Ethernet subsystem tests
4	Diagnostic processor console I/O mode
3	R3000 console I/O mode
2	Diagnostic processor primary/secondary bootstrap
1	R3000 primary/secondary bootstrap
0	Operating system running

4.2.4 Interprocessor Interaction

The R3000 and the diagnostic processor can each select the other processor to start or continue execution. This is accomplished through the SPR register (Section 4.2.4.1).

4.2.4.1 Select Processor Register Operation

The select processor register (SPR) determines which processor is active. This is a write-only register.

If the diagnostic processor writes a 1 to the SPR (diagnostic processor executes the command EXIT), the diagnostic processor hangs on a DMA grant and the R3000 begins execution.

If the R3000 writes a 0 to the SPR (R3000 executes the command maint), the R3000 hangs on a RDBUSY or WRBUSY stall and the diagnostic processor begins execution.

4.2.5 Power-Up Sequence

The KN210 power-up sequence depends on whether the processor is in *normal* mode or in *maintenance* mode.

4.2.5.1 Normal Power-Up Operation

When the KN210 is set to operate in *normal* mode, the power-up sequence algorithm is as shown in Example 4-4.

1. Diagnostic processor powers-up (begins execution at a location pointed to by 2004 0000)
2. Diagnostic processor runs self test diagnostics
3. Diagnostic processor executes EXIT command (SPR set to 1)
4. Diagnostic processor hangs on a DMA grant
5. R3000 begins execution at an address stored in 1FC0 0000.
 - if the bootmode environment variable is set to 'a'
 - R3000 attempts to auto_boot
 - if bootpath environment variable is valid
 - the auto_boot succeeds
 - else
 - the R3000 waits for a command at its prompt (>>)
 - if the bootmode environment variable is not initialized (*) or if it is set to 'd'
 - the R3000 waits for a command at its prompt (>>)
 - when the R3000 waits for a command at its prompt (>>) console commands can be received
 - if the command is maint (SPR is set to 0) then
 - The R3000 hangs on a RDBUSY stall.
 - The diagnostic processor resumes execution.

Example 4-4 Normal Power-Up Sequence Algorithm

4.2.5.2 Maintenance Power-Up Operation

When the KN210 is set to operate in *maintenance* mode, the power-up sequence algorithm is as shown in Example 4-5.

1. Diagnostic processor powers-up (begins execution at a location pointed by 2004 0000)
2. Diagnostic processor runs self test diagnostics
3. Diagnostic processor enters console mode (diagnostic processor waits for a command at its prompt (>>>))
 - console commands can be received
 - if the command is EXIT (SPR is set to 1) then
 - The diagnostic processor hangs on a DMA grant
 - The R3000 begins execution

Example 4-5 Maintenance Power-Up Sequence Algorithm

4.2.6 Processor Identification

The KN210 provides the Sys_Type read only register at physical location 2004 0004 in the diagnostic processor firmware EPROM and through the R3000 environment variable **systype**, described in Section 4.4.2.3.

4.2.6.1 Sys_Type Register Layout

The layout of the Sys_Type register is shown in Figure 4-1. The Sys_Type register address is 2004 0004.

Figure 4-1 System Type Register

Field	Name	RW	Description
31:24	SYS_ TYPE	ro	This field identifies the type of system for a specific processor. (Reads as 06 for the KN210.)
23:16	VERSION	ro	This field identifies the resident version of the firmware EPROM encoded as two hexadecimal digits. For example, if the banner displays V5.0, then this field is 50 (hex).
15:8	SYS_ SUB_ TYPE	ro	This field identifies the particular system sub-type. (Reads as 01 for the KN210.)
7:0	HARDWARE	ro	This field identifies the hardware revision level.

4.3 Operating System Bootstrap

Bootstrapping is the process of loading and transferring control to an operating system. The KN210 supports bootstrap of ULTRIX-32. Additionally, the KN210 will boot MDM diagnostics and any user application image which conforms to the boot formats described in this manual.

On the KN210, a bootstrap occurs whenever a **boot** command is issued at the console in *maintenance* or *normal* mode or whenever the processor halts and the conditions specified in Section 4.3.7 for automatic bootstrap are satisfied.

4.3.1 MDM Bootstrap

MDM is booted by the diagnostic processor when in *maintenance* mode and when the **BOOT** command is executed at the diagnostic processor console prompt.

For a detailed description of the diagnostic processor boot process, refer to Chapter 6.

4.3.2 Operating System Bootstrap

The operating system is booted by the R3000 when in one of the following states:

- The operation switch is in the *normal* position and the environment variable **bootmode** is set to **a**.
- The operation switch is in the *normal* position and the environment variable **bootmode** is undefined or set to **d** and the **boot** command is typed.
- The operating system initiates a reboot operation.

One of three ports may be used for bootstrapping:

- KN210 I/O board Ethernet controller
- KN210 I/O board DSSI controller
- KN210 Q22-bus MSCP or TMSCP controller

Refer to Section 4.7 for a list of the supported devices.

The disk and tape bootstrap code has no knowledge of any file structure on disks or tapes. Because of this, the operating system must provide an image capable of being loaded through the boot block that in turn loads the remainder of the operating system.

The Ethernet bootstrap code is capable of booting using the MOP Ethernet bootstrap protocol. BOOTP is not supported.

The boot code allows the operator to set the default bootstrap device and file name by storing them in a non-volatile environment variable (REFERENCE>(ENVIRONMENT_VARIABLES)). Once set, the system normally uses this default for bootstraps. However, the operator can override the default manually.

The boot code interface also provides for command extension through the bootstrap mechanism. If after parsing a command line the boot code interface fails to recognize a keyword that begins a command, it looks for the environment variable **path** and if it finds the path, it prepends the value of **path** to the keyword and attempts to boot the resulting filename. Any remaining arguments following the initial keywords are passed as arguments to the program after it is loaded.

4.3.3 Boot Process

The boot process is as follows:

1. Determine the device to be booted from parsing the boot command or from the *bootpath* environment variable.
2. If the process is initiated by the boot command, set the *bootpath* environment variable to the boot device in non-volatile memory.
3. Open the boot device and read block zero (Figure 4-2).
4. Read the sequence of blocks described in the bootblock into main memory.
5. Transfer control to the loaded code at the start address found on the bootblock. Arguments supplied through the boot command are passed using the C-language (*argc,argv*) conventions. A pointer to the current environment table is also passed.
6. Any exception in this process forces the boot code to display the console prompt.

Figure 4-2 Bootblock Layout

4.3.4 Bootstrap Support Routines in the Console

The ULTRIX bootstrap loader contains no code for performing I/O or machine dependent operations. Instead, the loader calls a set of routines provided by the console. The address of these routines are contained in a transfer vector located in console ROM. All routines are called as normal C routines using the R3000 calling conventions.

The entry point routines can be grouped as follows:

- console - Invoke console program restarts, reboots, and so on
- saio - I/O support for standalone programs. Raw device I/O
- machine - Machine specific functions (cache flushes, interlock memory references, and so on)
- libc - Subset of the C library (strings, getenv, and so on)
- parser - Access to console command parser
- commands - Access to subset of console commands

Section 4.6 describes the console entry points.

4.3.5 Console Use of Memory Space

The KN210 console reserves the bottom 64 Kbytes of physical address space visible from kseg1. The 64 Kbytes are partitioned as follows:

Table 4-3 Console Memory Space

Address	O/S Interface	Console Program
Starting	0000 0000	0000 04FF
Ending	0000 0500	0000 FFFF

The operating system interface region is used to hold the boot block and the current exception handling code.

The console program region is used to hold the console program stack and static data structures.

4.3.6 Boot Devices

Three types of boot media are supported: disk, tape and Ethernet. When specifying a boot device, the following syntax is used:

dev(controller, unit, partition)

The **dev** field is a device mnemonic. The controller and unit fields specify the hardware controller and unit number of the device. The partition field specifies a software managed partition of the device. Partitions are only meaningful for disk devices. If not specified, controller, unit and partition default to zero. The parentheses are always required.

4.3.6.1 Disk

Only DSSI and MSCP disks are supported. A simple *bootblock* bootstrap is supported from disk. If the disk is partitioned, a bootblock per partition is possible.

rf() and ra()

4.3.6.2 Tape

Only TMSCP tapes are supported. A simple bootblock bootstrap is supported from tape.

tm()

4.3.6.3 Ethernet

Bootstrapping over the *Ethernet* is a very important bootstrap function. The MOP protocol is used for network bootstraps.

mop()

This is a standard Digital bootstrap protocol. With MOP, either the client or the server may specify the file to be downloaded. If the client supplies a file specification, then that file is provided by the server. If the client does not supply a file specification, then the server determines the file from a database of registered clients. In the latter case, the client must be known to the server and be in its data base. In the former case, the only restriction is that the file be known to the server.

4.3.7 Halts

The R3000 processor **does not** provide a halt input to unconditionally call the console program. The KN210 has a hardware halt interrupt at hardware interrupt request (IRQ) 5 (Section 3.1.4). The operating system interrupt handler must call the console at the PROM_HALT entry point. Interrupts must be left disabled. The console program will save the machine state and restore in response to the *continue* command.

Table 4-4 describes the R3000 registers that are saved when the PROM_HALT entry point is called. The *continue* command first restores the R3000 registers to the saved values and then restarts the operating system at the saved return address in register **ra**.

Table 4-4 PROM_HALT Saved Registers

Register	Name	Register	Name
\$2..\$3	v0..v1	\$4..\$7	a0..a3
\$16..\$23	s0..s7	\$26..\$27	k0-k1
\$28	gp	\$29	sp
\$30	fp	\$31	ra
	c0_sr		cause

The use of the *maint* command while the R3000 is in console I/O mode destroys the saved machine state. If the operating system is halted and *maintenance* mode is invoked, a subsequent attempt to restart the operating system will cause unpredictable results.

NOTE

Entering *maintenance* mode after an R3000 operating system destroys the saved machine state.

The *maintenance* mode halt process is described in Chapter 6.

4.4 KN210 Console Command Language

The following sections describe the console command language.

4.4.1 Maintenance Mode Console Command Language

When in *maintenance* mode, the KN210 console command language interface is the same as for the KA640 console language. One additional command is implemented, the **Exit** command. The command has no arguments and has the effect of switching the KN210 from *maintenance* mode to *normal* mode. Refer to Chapter 6 for a detailed description of the *maintenance* mode console commands.

4.4.2 Normal Mode Console Command Language

When the KN210 is in *normal* mode and in console I/O mode, it will read and interpret commands received on the console terminal. The commands are based on the DECstation 3100 (PMAx) command language.

4.4.2.1 Control Characters

Certain ASCII control characters have special meaning when typed on the console terminal. These characters are described in Table 4-5.

Table 4-5 Control Characters

Character	Function
Return	Ends a command line. Command characters are buffered until a carriage-return is received.
Delete	Deletes the previously typed character. If the console terminal is defined as <i>hardcopy</i> (environment variable term set to <i>hardcopy</i>) the deleted text is echoed surrounded by backslashes. If the console terminal is a CRT (environment variable term set to <i>crt</i>) each delete is echoed with the sequence "<BS><SP><BS>". Deletes received when there are no characters to be deleted are ignored.
Ctrl C	Causes the console to abort processing of a command.

Table 4-5 (Cont.) Control Characters

Character	Function
Ctrl Q	Causes console output to be discarded until the next Ctrl Q is entered, or until the next console prompt or error message is issued. Ctrl Q is also cancelled when Ctrl C is typed.
Ctrl Q	Resumes console output that was suspended by Ctrl S .
Ctrl R	Causes the current command line to be displayed, omitting any deleted characters.
Ctrl S	Suspends output on the console terminal until Ctrl Q is typed.
Ctrl U	Discards all characters accumulated for the current line.
Ctrl V	Supresses any special meaning associated with the next character.

4.4.2.2 Lexical Conventions

The console is case insensitive with respect to parsing commands but case is preserved when assigning values to environment variables.

All console commands are specified using US ASCII characters only. Values entered for environment variables however may contain any 8-bit character code.

Numeric values may be entered as decimal, hexadecimal, octal, and binary values. Decimal values are represented by a string of decimal digits with no leading zeros (123). Octal values are represented by a string of octal digits and a leading zero must be present (0177). Hexadecimal values are represented by hexadecimal digits preceded with a "0x" (0x3ff). Binary values are represented by binary digits preceded with a "0b" (0b1001).

4.4.2.3 Environment Variables

The KN210 console makes use of environment variables to pass information to the operating system. Some of the environment variables are maintained in non-volatile RAM so that their contents are not lost when power is removed.

Additional environment variables may be defined by the operator or are defined automatically by the console program but they are lost when power is lost. The environment variables shown in Table 4–6 are maintained in non-volatile memory.

Table 4–6 Environment Variables

Variable	Type	Function
baud ¹	init	Baud rate of the console terminal line. Read only. Possible values are 300, 600, 1200, 2400, 4800, 9600, 19200, and 38400.
bootpath	nvrAm	A string containing the complete boot path specification. The boot path has two fields; the boot device (Section 4.3.6), and the bootpath. An example of a bootpath definition is: ra(0,0,0)moon_lander.
bootmode	nvrAm	A one character code controlling what action the console is to take on power-up or following a reset. Two codes are defined: a for autoboot, and d to halt after performing power-up diagnostics. Allowable console devices are tty(0) for console serial line.
bitmap ¹	init	The hexadecimal address of good pages bitmap.
bitmaplen ¹	init	The length of the memory bitmap.
osconsole ¹	init	Always tty(0).
console	nvrAm	Always selects tty(0).

¹These environments variables are read only. Attempts to write from console I/O have no effect.

Table 4-6 (Cont.) Environment Variables

Variable	Type	Function
systype ¹	init	Read only. Contains information used to identify the processor. Bits 24:31 contain the CPU type. Bits 16:23 contain the system type (6 for the KN210). Bits 8:15 contain the firmware revision level and bits 0:7 contain the hardware version level.

¹These environments variables are read only. Attempts to write from console I/O have no effect.

4.4.2.4 Commands

The commands documented below are accepted by the console.

?

? [command-list]

Identical to the help command.

BOOT

boot [-f file] [-s] [-n] [args]

Boot loads the file following the flag -f. If the -f flag is not specified, then the file specified by the environment variable bootpath is loaded. If -s is specified, the operating system is booted in single user mode. If -n is specified, then the file is loaded but control is not passed to the program. If any arguments are present, then they are passed to the booted image using the standard argc/argv mechanism. If any argument begins with a (-) then it must be preceded with an additional (-) character.

CONTINUE

continue

The continue command causes the processor to begin execution at the address currently in the saved program counter. The processor state saved at the last console entry is restored before leaving console mode.

D

d [-bhw] address value

D deposits a single byte, halfword or word value at the indicated address.

DUMP

dump [-Bcdoux] [-bhw] range

Dump performs a formatted display of the contents of memory. Memory contents may be displayed (simultaneously) in hex (-x), unsigned decimal (-u), octal (-o), decimal (-d), as ASCII characters (-c), or as binary (-B). Memory contents may be dumped as bytes (-b), halfwords (-h), or words (-w). The range of memory to be dumped may be specified as base-address (a single value is dumped) base-address#count (a specified number of values are dumped) or base-address:limit-address (all values between the base address and the limit address are dumped).

E

e [-bhw] address

E displays the byte, halfword, or word at address.

FILL

fill [-bhw] [-vvalue] range

Fill sets the range of memory specified to the value specified. If no value is specified, then zero is used. Memory contents may be filled as bytes (-b), halfwords (-h), or words (-w). The range of memory to be filled is specified as in the dump command.

GO

go [entry]

Go transfers control to the indicated entry point address. If no entry address is supplied, then the entry point of the last program module loaded is used.

HELP

help [command-list]

Help displays a brief synopsis of the indicated command. If no command list is supplied then console displays a synopsis of each command.

INIT

init

Init performs a full initialization. The effect is identical to that performed at power-up or reset except that no diagnostics are executed.

MAINT

maint

This command causes the console to enter *maintenance* mode. Any saved program state is discarded.

PRINTENV

printenv [variable-list]

Printenv displays the indicated environment variable on the console terminal. If no variable is specified, all console environment variables are displayed.

SETENV

setenv variable value

Setenv assigns a value to the indicated environment variable.

UNSETENV

unsetenv variable

Unsetenv removes the indicated variable from the set of console environment variables. The environment variables stored in non-volatile memory are not affected by this command.

4.5 Diagnostics

The ROM based diagnostics constitute a major portion of the firmware on the KN210. These diagnostics run automatically on power-up when in *maintenance* mode, and can be executed interactively as a whole, or as individual tests. Refer to Chapter 6 for further information on KN210 diagnostics.

4.6 PROM Entry Points

The following entry points are defined in the KN210 PROM. This information is provided here only as a general guide as to what functions are available. All routines are called as normal C routines. The normal C conventions are assumed.

4.6.1 Argvize

Breaks a string into tokens and returns the number of tokens. The value returned by *Argvize* is *slp.strcnt (argc)* and *slp->strptrs* is the address of vector of strings (*argv*).

```

struct string_list{
char *strptrs[MAXSTRINGS];           /* Vector of string pointers */
char *strbuf[STRINGBYTES];          /* Strings themselves */
char *strp;                           /* Free ptr in strbuf. */
int strcnt;                            /* Number of strings in strptrs. */ }

int
argvize(str,slp)
char * str;
struct string_list *slp;

```

4.6.2 Atob

Converts ASCII to binary. Accepts all C numeric input formats. Returns pointer to any unconverted substring left after the numeric conversion.

```

char *
atob(str,intp)

```

4.6.3 Autoboot

Perform automatic bootstrap. Unlike **reboot**, the bootstrap is unconditional (although bootpath must be defined). If the autoboot fails, control passes to the console.

```
void
autoboot()
```

4.6.4 Bevexcept

Primitive exception handler for exceptions that occur while BEV bit is still set in system status register. Not a C routine.

```
bevexcept()
```

4.6.5 Bevutlb

Primitive exception handler for exceptions that occur while BEV bit is still set in system status register. Not a C routine.

```
bevutlb()
```

4.6.6 Close

Close a file.

```
int
close(fd)
int fd;
```

4.6.7 Dumpcmd

Invokes the console dump command. Returns 0 if successful, 1 if not.

```
int
dumpcmd(argc, argv)
int argc;
char **argv;
```

4.6.8 Exec

Loads a new program image.

```
void  
exec()
```

4.6.9 Getchar

Inputs a single character from the current console device.

```
int  
getchar()
```

4.6.10 Getenv

Returns the value of a console environment variable.

```
Char *  
getenv(name)  
char *name;
```

4.6.11 Gets

Gets a line of input from the console and places it in the caller's buffer. Normal line editing functions are performed on input. The address of the buffer is returned.

```
char *  
gets(buf)  
char *buf;
```

4.6.12 Halt

Saves the machine state and enters console I/O mode.

```
halt
```

4.6.13 Help

Invokes the console help command. Returns 0 if successful, 1 if not.

```
help
```

4.6.14 ioctl

Performs a device or file system specific operation on a device or file.

```
int
ioctl(fd, cmd, arg)
int fd;
int cmd;
int arg;
```

4.6.15 Longjump

Terminates execution in current context and continues from context saved in jmp_buf by a previous setjmp. Execution resumes with rval as the value returned from setjmp.

```
longjump(jmp_buf, rval)
struct jmp_buf *jmp_buf;
int rval;
```

4.6.16 Lseek

Positions a file into an arbitrary byte position and returns the new position as int.

```
int
lseek(fd, offset, direction);
int fd;
int offset;
int direction;
```

4.6.17 Open

Opens the indicated filename.

```
int
open(filename, flags)
char *filename;
int flags;
```

4.6.18 Parser

Inputs a command line from the console terminal, argvizes it and looks up the first token in the command table. If it is found, invokes the corresponding command with the standard argc/argv argument list. If the command is not found in the command table and search_path is not null, then attempts to perform the command by "exec'ing" the file specified by the concatenation of search_path and the first token.

```

struct cmd_table{
    char *name; /* Command name */
    int (*routine)(); /* Command routine */
    char *usage; /* Command usage string. */
}

int
parser(cmd_table, prompt, search_path)
struct cmd_table *cmd_table;
char *prompt;
char *search_path;

```

4.6.19 Printenvcmd

Invoke the console printenv command. Returns 0 if successful, 1 if not.

```

int
printenvcmd(argc, argv)
int argc;
char *argv;

```

4.6.20 Printf

Prints formatted values on the current console device.

```

void
printf(fmt, va_alist)
char *fmt;
va_dcl

```

4.6.21 Putchar

Outputs a single character to the current console device.

```

void
putchar(c)
char c;

```


4.6.22 Puts

Outputs a string to the current console device.

```
void  
puts(s)  
char *s;
```

4.6.23 Range

Parses a console range specification and returns the base address and the count. Returns 0 if the range is of the form address:address, 1 if the range is of the form address#cnt, and minus 1 if the range cannot be parsed.

```
int  
range(str, basep, cntp)  
char *str;  
unsigned *basep;  
unsigned *cntp;
```

4.6.24 Read

Read data from a file into a buffer. Return the actual number of bytes read to the caller

```
int  
read(fd, buf, cnt)  
int fd;  
char *buf;  
int cnt;
```

4.6.25 Reboot

Performs automatic test and reboots system if bootmode is **a**, otherwise enters console. The effect of this routine is to duplicate the normal console power-up/reset boot decision.

```
void  
reboot()
```

4.6.26 Reinit

Re-initializes the console monitor and enters command mode.

```
reinit()
```

4.6.27 Reset

Enter the prom at its entry point, the result is the same as if the reset button had been pressed or the power switch turned on.

```
void  
reset()
```

4.6.28 Restart

Re-enters the console monitor without re-initializing it.

```
void  
restart()
```

4.6.29 Setenv

Sets the value of an environment variable.

```
int  
setenv(name, value)  
char *name;  
char *value;
```

4.6.30 Setenvcmd

Invokes the console setenv command. Returns 0 if successful, 1 if not.

```
int  
setenvcmd(argc, argv)  
int argc;  
char **argv;
```

4.6.31 Setjmp

Saves the current register context in `jmp_buf` and then returns zero. Later, a call to `longjmp` causes the saved context to be restored and execution resumes again, this time with the return value specified by the **longjmp** routine. Used to regain control from exceptional conditions.

```
typedef jmp_buf[11]; /* Register context */
int
setjmp(jmp_buf)
jmp_buf jmp_buf;
```

4.6.32 Showchar

Outputs a single character to the current console device. Displays all printing characters normally. Displays blanks as `"\b"`, form feeds as `"\f"`, newlines as `"\n"`, returns as `"\r"`, tabs as `"\t"`. Display any other nonprinting character as `"\xxx"` where `xxx` is the octal code for the character.

```
void
showchar(c)
char c;
```

4.6.33 Strcat

Concatenates string 1 and string 2, returning a pointer to the result.

```
char *
strcat(str1, str2)
char *str1;
char *str2;
```

4.6.34 Strcmp

Compares string 1 and string 2. Returns 0 if they are the same, a negative value if string 1 is less than string 2, a positive value if string 1 is greater than string 2.

```
int
strcmp(str1, str2)
char *str1;
char *str2;
```

4.6.35 Strcpy

Copies string 2 to string 1 and returns a pointer to the first unmodified character in string 1. String 1 must be long enough to contain string 2.

```
char *
strcpy(str1, str2)
char *str1;
char *str2;
```

4.6.36 Strlen

Returns the number of bytes in a string.

```
int
strlen(str)
char *str;
```

4.6.37 Unsetenvcmd

Delete an environment variable.

```
int
unsetenvcmd(argc, argv)
int argc;
char *argv;
```

4.6.38 Write

Write data from a buffer into a file. Return the actual number of bytes written to the caller.

```
int
write(fd, buf, cnt)
int fd;
char *buf;
int cnt;
```

4.7 Supported Devices

The KN210 firmware supports operating system bootstrapping from the following devices:

Table 4-7 KN210 Boot Devices

Name	Protocol	Unit	Type
rf	DSSI	0-7	Disk
ra	MSCP	0-3	Disk
tm	TMSCP	0-3	Tape
mop	MOP		Ethernet

The DSSI and NI adaptors are in the KN210 I/O module. The MSCP devices are attached to the Q22-bus.

NOTE

Only Q22-bus MSCP or TMSCP devices are supported by the console bootstrap procedure.

5

Diagnostic Processor

5.1 Diagnostic Processor

The diagnostic processor of the KN210 is implemented by a VLSI chip called the CVAX.

5.1.1 Processor State

The processor state consists of that portion of the state of a process that is stored in processor registers rather than in memory. The processor state is composed of sixteen general purpose registers (GPRs), the processor status longword (PSL), and the internal processor registers (IPRs).

Non-privileged software can access the GPRs and the processor status word (bits <15:00> of the PSL). The IPRs and bits <31:16> of the PSL can only be accessed by privileged software. The IPRs are explicitly accessible only by the move to processor register (MTPR) and move from processor register (MFPR) instructions that can be executed only while running in kernel mode.

5.1.1.1 General Purpose Registers

The KN210 diagnostic processor implements 16 GPRs. These registers are used for temporary storage, as accumulators, and as base and index registers for addressing. These registers are denoted R0 - R15. The bits of a register are numbered from the right <0> through <31> (Figure 5-1).

MA-1100-87

Figure 5-1 General Purpose Register Bit Map

Certain of these registers have been assigned special meaning by the VAX-11 architecture:

- R15 is the program counter (PC). The PC contains the address of the next instruction byte of the program.
- R14 is the stack pointer (SP). The SP contains the address of the top of the processor defined stack.
- R13 is the frame pointer (FP). The VAX-11 procedure call convention builds a data structure on the stack called a *stack frame*. The FP contains the address of the base of this data structure.
- R12 is the argument pointer (AP). The VAX-11 procedure call convention uses a data structure called an *argument list*. The AP contains the address of the base of this data structure.

Consult the *VAX Architecture Reference Manual* for more information on the operation and use of these registers.

5.1.1.2 Processor Status Longword

The PSL is saved on the stack when an exception or interrupt occurs and is saved in the process control block (PCB) on a process context switch. Bits <15:00> may be accessed by non-privileged software, while bits <31:16> may only be accessed by privileged software. Processor initialization sets the PSL to 041F 0000₁₆. Figure 5-2 shows the processor status longword bit map.

Figure 5-2 PSL Bit Map

Data Bit	Definition
PSL <31>	(CM) Compatibility mode. This bit always reads as zero; loading a 1 into this bit has no effect.
PSL <30>	(TP) Trace pending
PSL <29:28>	Unused, must be written as zero.
PSL <27>	(FPD) First part done
PSL <26>	(IS) Interrupt stack
PSL <25:24>	(CUR) Current mode
PSL <23:22>	(PRV) Previous mode
PSL <21>	Unused, must be written as zero.
PSL <20:16>	(IPL) Interrupt priority level
PSL <15:8>	Unused, must be written as zero.

Data Bit	Definition
PSL <7>	(DV) Decimal overflow trap enable. This read/write bit has no effect on KN210 hardware; it can be used by macrocode which emulates VAX decimal instructions.
PSL <6>	(FU) Floating underflow fault enable
PSL <5>	(IV) Integer overflow trap enable
PSL <4>	(T) Trace trap enable
PSL <3>	(N) Negative condition code
PSL <2>	(Z) Zero condition code
PSL <1>	(V) Overflow condition code
PSL <0>	(C) Carry condition code

5.1.1.3 Internal Processor Registers

The KN210 IPRs can be accessed by using the MFPR and MTPR privileged instructions. Each IPR falls into one of the following seven categories:

1. Implemented by KN210 (in the CVAX chip)
2. Implemented by KN210 (in the SSC)
3. Implemented by KN210 (and all designs that use the CVAX chip) uniquely
4. Implemented by KN210 (and all designs that use the SSC) uniquely
5. Not implemented, timed out by the CDAL bus timer (in the SSC) after 4 μ s. Read as zero, NOP on write.
6. Access not allowed; accesses result in a reserved operand fault
7. Accessible, but not fully implemented; accesses yield *unpredictable* results

Refer to Table 5-1 for a listing of each of the KN210 IPRs, along with its mnemonic, its access type (read or write) and its category number.

Table 5-1 KN210 Internal Processor Registers

Decimal	Hex	Register Name	Mnemonic	Type	Category ¹
0	0	Kernel stack pointer	KSP	R/W	1
1	1	Executive stack pointer	ESP	R/W	1
2	2	Supervisor stack pointer	SSP	R/W	1
3	3	User stack pointer	USP	R/W	1
4	4	Interrupt stack pointer	ISP	R/W	1
7:5	7:5	Reserved			5
8	8	P0 base register	P0BR	R/W	1
9	9	P0 length register	P0LR	R/W	1
10	A	P1 base register	P1BR	R/W	1
11	B	P1 length register	P1LR	R/W	1
12	C	System base register	SBR	R/W	1
13	D	System length register	SLR	R/W	1
15:14	F:E	Reserved			5
16	10	Process control block base	PCBB	R/W	1
17	11	System control block base	SCBB	R/W	1
18	12	Interrupt priority level	IPL	R/W	1 I
19	13	AST level	ASTLVL	R/W	1 I
20	14	Software interrupt request	SIRR	W	1
21	15	Software interrupt summary	SISR	R/W	1 I
23:22	17:16	Reserved			5
24	18	Interval clock control/status	ICCS	R/W	3 I
25	19	Next interval count	NICR	W	5
26	1A	Interval count	ICR	R	5
27	1B	Time-of-year	TODR	R/W	2
28	1C	Console storage receiver status	CSRS	R/W	7 I
29	1D	Console storage receiver data	CSRD	R	7 I
30	1E	Console storage transmit status	CSTS	R/W	7 I

¹An I following the category number indicates that the register is initialized on power-up and by the negation of DCOK when the processor is halted.

Table 5-1 (Cont.) KN210 Internal Processor Registers

Decimal	Hex	Register Name	Mnemonic	Type	Category ¹
31	1F	Console storage transmit data	CSTD	W	7 I
32	20	Console receiver control/status	RXCS	R/W	4 I
33	21	Console receiver data buffer	RXDB	R	4 I
34	22	Console transmit control/status	TXCS	R/W	4 I
35	23	Console transmit data buffer	TXDB	W	4 I
36	24	Translation buffer disable	TBDR	R/W	5
37	25	Cache disable	CADR	R/W	3 I
38	26	Machine check error summary	MCESR	R/W	5
39	27	Memory system error	MSER	R/W	3 I
41:40	29:28	Reserved			5
42	2A	Console saved PC	SAVPC	R	3
43	2B	Console saved PSL	SAVPSL	R	3
47:44	2F:2C	Reserved			5
48	30	SBI fault/status	SBIFS	R/W	5
49	31	SBI silo	SBIS	R	5
50	32	SBI silo comparator	SBISC	R/W	5
51	33	SBI maintenance	SBIMT	R/W	5
52	34	SBI error register	SBIER	R/W	5
53	35	SBI timeout address register	SBITA	R	5
54	36	SBI quadword clear	SBIQC	W	5
55	37	I/O bus reset	IORESET	W	4
56	38	Memory management enable	MAPEN	R/W	1
57	39	TB invalidate all	TBIA	W	1
58	3A	TB invalidate single	TBIS	W	1
59	3B	TB data	TBDATA	R/W	5
60	3C	Microprogram break	MBRK	R/W	5

¹An I following the category number indicates that the register is initialized on power-up and by the negation of DCOK when the processor is halted.

Table 5-1 (Cont.) KN210 Internal Processor Registers

Decimal	Hex	Register Name	Mnemonic	Type	Category ¹
61	3D	Performance monitor enable	PMR	R/W	5
62	3E	System identification	SID	R	1
63	3F	Translation buffer check	TBCHK	W	1
64:127	40:7F	Reserved			6

¹An **I** following the category number indicates that the register is initialized on power-up and by the negation of DCOK when the processor is halted.

KN210 Diagnostic VAX Standard IPRs

VAX standard internal processor registers (IPRs) are classified as Category One IPRs. The *VAX Architecture Reference Manual* should be consulted for details on the operation and use of these registers.

The VAX standard IPRs listed in Table 5-2 are also referenced in other sections of this manual.

Table 5-2 VAX Standard IPRs

Number	Register Name	Mnemonic	Section
12	System base register	SBR	5.1.5.3
13	System length register	SLR	5.1.5.3
16	Process control block base	PCBB	5.1.5
17	System control block base	SCBB	5.1.5.4
18	Interrupt priority level	IPL	5.1.5.1
20	Software interrupt request	SIRR	5.1.5.1
21	Software interrupt summary	SISR	5.1.5.1
27	Time-of-year clock	TODR	3.6.2
56	Memory management enable	MAPEN	5.1.4.2
57	Translation buffer invalidate all	TBIA	5.1.4.2
58	Translation buffer invalidate single	TBIS	5.1.4.2
62	System identification	SID	5.1.6
63	Translation buffer check	TBCHK	5.1.4.2

KN210 Unique IPRs

Internal processor registers that are implemented uniquely on the KN210 are described in detail in this manual. Refer to the sections listed in Table 5-3 for a description of these registers.

Table 5-3 KN210 Unique IPRs

Number	Register Name	Mnemonic	Section
24	Interval clock control/status	ICCS	3.6.3
32	Console receiver control/status	RXCS	3.5.1.1
33	Console receiver data buffer	RXDB	3.5.1.2
34	Console transmit control/status	TXCS	3.5.1.3
35	Console transmit data buffer	TXDB	3.5.1.4
42	Console saved PC	SAVPC	5.1.5
43	Console saved PSL	SAVPSL	5.1.5

5.1.2 Data Types

The diagnostic processor supports the following subset of the VAX data types:

- Byte
- Word (16-bit)
- Longword (32-bit)
- Quadword (64-bit)
- Character string
- Variable length bit field

5.1.3 Instruction Set

The KN210 CPU implements the following subset of the VAX instruction set types in microcode:

- Integer arithmetic and logical
- Address
- Variable length bit field

- Control
- Procedure call
- Miscellaneous
- Queue¹
- Character string moves (MOV3, MOV5, CMPC3¹, CMPC5¹, LOCC¹, SCANC¹, SKPC¹, and SPANC¹)
- Operating system support

5.1.4 Memory Management

The KN210 implements full VAX memory management. System space addresses are virtually mapped through single-level page tables, and process space addresses are virtually mapped through two-level page tables.

5.1.4.1 Translation Buffer

To reduce the overhead associated with translating virtual addresses to physical addresses, the KN210 employs a 28-entry, fully associative, translation buffer for caching VAX PTEs in modified form. Each entry can store a modified PTE for translating virtual addresses in either the VAX process space, or VAX system space. The translation buffer is flushed whenever memory management is enabled or disabled (that is, by writes to IPR 56), any page table base or length registers are modified (that is, by writes to IPRs 8 - 13) and by writing to IPR 57 (TBIA) or IPR 58 (TBIS).

Each entry is divided into two parts: a 23-bit tag register and a 31-bit PTE register. The tag register is used to store the virtual page number (VPN) of the virtual page that the corresponding PTE register maps. The PTE register stores the 21-bit PFN field, the PTE.V bit, the PTE.M bit and an 8-bit partially decoded representation of the 4-bit VAX PTE PROT field, from the corresponding VAX PTE, as well as a translation buffer valid (TB.V) bit.

During virtual to physical address translation, the contents of the 28 tag registers are compared with the virtual page number field (bits <31:9> of the virtual address of the reference. If there is a match with one of the tag registers, then a translation buffer *hit* occurs, and the contents of the corresponding PTE register is used for the translation.

¹ These instructions were in the microcode assisted category on the KA630-AA (MicroVAX II) and therefore had to be emulated.

If there is no match, the translation buffer does not contain the necessary VAX PTE information to translate the address of the reference, and the PTE must be fetched from memory. Upon fetching the PTE, the translation buffer is updated by replacing the entry that is selected by the replacement pointer. Since this pointer is moved to the next sequential translation buffer entry whenever it is pointing to an entry that is accessed, the replacement algorithm is *not last used* (NLU).

5.1.4.2 Memory Management Control Registers

There are four IPRs that control the memory management unit (MMU): IPR 56 (MAPEN), IPR 57 (TBIA), IPR 58 (TBIS), and IPR 63 (TBCHK).

Memory management can be enabled/disabled through IPR 56 (MAPEN). Writing a 0 to this register with a MTPR instruction disables memory management, and writing a 1 to this register with a MTPR instruction enables memory management. Writes to this register flush the translation buffer. To determine whether or not memory management is enabled, IPR 56 is read using the MFPR instruction. Translation buffer entries that map a particular virtual address can be invalidated by writing the virtual address to IPR 58 (TBIS) using the MTPR instruction.

NOTE

Whenever software changes a valid PTE for the system or current process region, or a system PTE that maps any part of the current process page table, all process pages mapped by the PTE must be invalidated in the translation buffer.

The entire translation buffer can be invalidated by writing a 0 to IPR 57 (TBIA) using the MTPR instruction.

The translation buffer can be checked to see if it contains a valid translation for a particular virtual page by writing a virtual address within that page to IPR 63 (TBCHK) using the MTPR instruction. If the translation buffer contains a valid translation for the page, the condition code V bit (bit <1> of the PSL) is set.

NOTE

The TBIS, TBIA, and TBCHK IPRs are write only. The operation of a MFPR instruction from any of these registers is *undefined*.

5.1.5 Exceptions and Interrupts

Both exceptions and interrupts divert execution from the normal flow of control. An exception is caused by the execution of the current instruction and is typically handled by the current process (for example, an arithmetic overflow), while an interrupt is caused by some activity outside the current process and typically transfers control outside the process (for example, an interrupt from an external hardware device).

5.1.5.1 Interrupts

Interrupts can be divided into two classes: non-maskable, and maskable.

Non-maskable interrupts cause a halt through the hardware halt procedure which saves the PC, PSL, MAPEN<0> and a halt code in IPRs, raises the processor IPL to 1F and then passes control to the resident firmware. The firmware dispatches the interrupt to the appropriate service routine based on the halt code and hardware event indicators.

Non-maskable interrupts cannot be blocked by raising the processor IPL, but can be blocked by running out of the halt protected address space (except those non-maskable interrupts that generate a halt code of 3). Non-maskable interrupts with a halt code of 3 cannot be blocked since this halt code is generated after a hardware reset.

Maskable interrupts cause the PC and PSL to be saved, the processor IPL to be raised to the priority level of the interrupt (except for Q22-bus interrupts where the processor IPL is set to 17 independent of the level at which the interrupt was received) and the interrupt to be dispatched to the appropriate service routine through the system control block (SCB).

The various interrupt conditions for the KN210 are listed in Table 5-4 along with their associated priority levels and SCB offsets.

Table 5-4 Interrupts

Priority Level	Interrupt Condition	SCB Offset
Non-maskable	BDCOK and BPOK negated then asserted on Q22-bus (Power up) BDCOK negated then asserted while BPOK asserted on Q22-bus (SCR<7> clear)	¹

¹These conditions generate a hardware halt procedure with a halt code of 3 (hardware reset).

Table 5-4 (Cont.) Interrupts

Priority Level	Interrupt Condition	SCB Offset
	BDCOK negated then asserted while BPOK asserted on Q22-bus (SCR<7> set)	2
	BHALT asserted on Q22-bus	2
	BREAK generated by the console device	2
1F	Unused	
1E	BPOK negated on Q22-bus	0C
1D	CDAL bus parity error	60
	Q22-bus NXM on a write	60
	CDAL bus timeout during DMA	60
	Main memory NXM errors	60
	Uncorrectable main memory errors	60
1C - 1B	Unused	
1A	Correctable main memory errors	54
19 - 18	Unused	
17	BR7 L asserted	Q22-bus vector plus 200 ₁₆
16	Interval timer interrupt	C0
	Mass storage interface	C5
	Network interface	D5
	BR6 L asserted	Q22-bus vector plus 200 ₁₆
15	BR5 L asserted	Q22-bus vector plus 200 ₁₆
14	Console terminal	F8,F6 ₁₆
	Programmable timers	78,7C
	BR4 L asserted	Q22-bus vector plus 200 ₁₆
13 - 10	Unused	
0F - 01	Software interrupt requests	84-BC

²These conditions generate a hardware halt procedure with a halt code of 2 (external halt).

NOTE

Because the Q22-bus does not allow differentiation between the four bus grant levels (that is, a level 7 device could respond to a level 4 bus grant), the KN210 diagnostic processor raises the IPL to 17 after responding to interrupts generated by the assertion of either BR7 L, BR6 L, BR5 L, or BR4 L. The KN210 maintains the IPL at the priority of the interrupt for all other interrupts.

The interrupt system is controlled by three IPRs: IPR 18, the interrupt priority level register (IPL), IPR 20, the software interrupt request register (SIRR), and IPR 21, the software interrupt summary register (SISR).

The IPL is used for loading the processor priority field in the PSL (bits <20:16>). The SIRR is used for creating software interrupt requests. The SISR records pending software interrupt requests at levels 1 through 15. The format of these registers is shown in Figure 5–3.

Figure 5–3 Interrupt Registers

5.1.5.2 Exceptions

Exceptions can be divided into three types:

- Trap
- Fault
- Abort

A *trap* is an exception that occurs at the end of the instruction that caused the exception. After an instruction traps, the PC saved on the stack is the address of the next instruction that would have normally been executed and the instruction can be restarted.

A *fault* is an exception that occurs during an instruction, and that leaves the registers and memory in a consistent state such that the elimination of the fault condition and restarting the instruction will give correct results. After an instruction faults, the PC saved on the stack points to the instruction that faulted.

An *abort* is an exception that occurs during an instruction, leaving the value of the registers and memory *unpredictable*, such that the instruction cannot necessarily be correctly restarted, completed, simulated or undone. After an instruction aborts, the PC saved on the stack points to the instruction that was aborted (which may or may not be the instruction that caused the abort) and the instruction may or may not be restarted depending on the class of the exception and the contents of the parameters that were saved.

Exceptions are grouped into six classes:

- Arithmetic
- Memory management
- Operand reference
- Instruction execution
- Tracing
- System failure

A list of exceptions grouped by class is given in Table 5-5. Exceptions save the PC and PSL and in some cases, one or more parameters, on the stack. Most exceptions do not change the IPL of the processor (except the exceptions in serious system failures class, which set the processor IPL to 1F) and cause the exception to be dispatched to the appropriate service routine through the SCB (except for the interrupt stack not valid exception, and exceptions that occur while an interrupt or another exception are being serviced, which cause the exception to be dispatched to the appropriate service routine by the resident firmware).

The exceptions listed in Table 5-5 (except machine check) are described in greater detail in the *VAX Architecture Reference Manual*. The machine check exception is described in greater detail in Section 5.1.5.3. Exceptions that can occur while servicing an interrupt or another exception are listed in Table 5-8 in Section 5.1.5.6.

Table 5–5 Exceptions

Arithmetic Exceptions	Type	SCB Offset
Integer overflow	Trap	34
Integer divide-by-zero	Trap	34
Subscript range	Trap	34
<hr/>		
Memory Management Exceptions	Type	SCB Offset
Access control violation	Fault	20
Translation not valid	Fault	24
<hr/>		
Operand Reference Exceptions	Type	SCB Offset
Reserved addressing mode	Fault	1C
Reserved operand fault	Abort	18
<hr/>		
Instruction Execution Exceptions	Type	SCB Offset
Reserved/privileged instruction	Fault	10
Emulated instruction	Fault	C8, CC
Change mode	Trap	40-4C
Breakpoint	Fault	2C
<hr/>		
Tracing Exception	Type	SCB Offset
Trace	Fault	28
<hr/>		
System Failure Exceptions	Type	Offset
Interrupt stack not valid	Abort	¹
Kernel stack not valid	Abort	08

¹Dispatched by resident firmware rather than through the SCB

Table 5-5 (Cont.) Exceptions

System Failure Exceptions	Type	Offset
Machine check	Abort	04
CDAL bus parity errors		
Cache parity errors		
Q22-bus NXM errors		
Q22-bus device parity errors		
Q22-bus NO GRANT errors		
CDAL bus timeout errors		
Main memory NXM errors		
Main memory uncorrectable errors		

5.1.5.3 Information Saved on a Machine Check Exception

In response to a machine check exception the PSL, PC, four parameters, and a byte count are pushed onto the stack, as shown in Figure 5-4.

MA-1121-87

Figure 5-4 Information Saved on a Machine Check Exception

The meaning of this information and how it effects the recovery procedure is described in the following paragraphs.

Byte Count

<31:0> = 0000 0010₁₆, 16₁₀. This value indicates the number of bytes of information that follow on the stack (not including the PC and PSL).

Machine Check Code Parameter

Machine Check Code <31:0>—A code value that indicates the type of machine check that occurred. A list of the possible machine check codes (in hex) and their associated causes follows.

Memory Management Errors—These codes indicate that the microcode in the CVAX CPU chip detected an impossible situation while performing functions associated with memory management. The most likely cause of this type of a machine check is a problem internal to the CVAX chip. Machine checks due to memory management errors are **non-recoverable**. Depending on the current mode, either the current process or the operating system should be terminated. The state of the P0BR, P0LR, P1BR, P1LR, SBR, and SLR should be logged.

Hex Code	Error Description
5	The calculated virtual address for a process PTE was in the P0 space instead of the system space when the CPU attempted to access a process PTE after a translation buffer <i>miss</i> .
6	The calculated virtual address space for a process PTE was in the P1 space instead of the system space when the CPU attempted to access a process PTE after a translation buffer <i>miss</i> .
7	The calculated virtual address for a process PTE was in the P0 space instead of the system space when the CPU attempted to access a process PTE to change the PTE <M> bit before writing to a previously unmodified page.
8	The calculated virtual address for a process PTE was in the P1 space instead of the system space when the CPU attempted to access a process PTE to change the PTE <M> bit before writing to a previously unmodified page.

Interrupt Errors—This code indicates that the interrupt controller in the diagnostic processor requested a hardware interrupt at an unused hardware IPL. The most likely cause of this type of a machine check is a problem internal to the CVAX chip. Machine checks due to unused IPL errors are **non-recoverable**. A non-vectored interrupt generated by a serious error condition (memory error, power fail, or processor halt) has probably been lost. The operating system should be terminated.

Hex Code	Error Description
9	A hardware interrupt was requested at an unused Interrupt Priority Level (IPL).

Microcode Errors—This code indicates that an impossible situation was detected by the microcode during instruction execution. Note that most erroneous branches in the CVAX CPU microcode cause random microinstructions to be executed. The most likely cause of this type of machine check is a problem internal to the CVAX chip. Machine checks due to microcode errors are **non-recoverable**. Depending on the current mode, either the current process or the operating system should be terminated.

Hex Code	Error Description
A	An impossible state was detected during a MOVC3 or MOVC5 instruction (not move forward, move backward, or fill).

Read Errors—These codes indicate that an error was detected while the CVAX CPU was attempting to read from either the cache, main memory, or the Q22-bus. The most likely cause of this type of machine check must be determined from the state of the MSER, DSER, MEMCSR16, QBEAR, DEAR, and CBTCR. Machine checks due to read errors **may be recoverable**, depending on the state of the VAX CAN'T RESTART flag (captured in Internal State Information 2 <15>) and the FIRST PART DONE flag (captured in PSL <27>). If the FIRST PART DONE flag is set, the error is recoverable. If the FIRST PART DONE flag is cleared, then the VAX CAN'T RESTART flag must also be cleared for the error to be recoverable. Otherwise, the error is unrecoverable and depending on the current mode, either the current process or the operating system should be terminated. The information pushed onto the stack by this type of machine check is from the instruction that caused the machine check.

Hex Code	Error Description
80	An error occurred while reading an operand, a process PTE during address translation, or on any read generated as part of an interlocked instruction.

Hex Code	Error Description
81	An error occurred while reading a system page table entry (SPTE), during address translation, a process control block (PCB) entry during a context switch, or a system control block (SCB) entry while processing an interrupt.

Write Errors—These codes indicate that an error was detected while the CVAX CPU was attempting to write to either the cache, main memory, or the Q22-bus. The most likely cause of this type of machine check must be determined from the state of the MSER, DSER, MEMCSR16, QBEAR, DEAR, and CBTCR. Machine checks due to write errors are **non-recoverable** because the CPU is capable of performing many read operations out of the cache before a write operation completes. For this reason, the information that is pushed onto the stack by this type of machine check cannot be guaranteed to be from the instruction that caused the machine check.

Hex Code	Error Description
82	An error occurred while writing an operand, or a process page table entry (PPTE) to change the PTE <M> bit before writing a previously unmodified page.
83	An error occurred while writing a system page table entry (SPTE) to change the PTE <M> bit before writing a previously unmodified page, or a PCB entry during a context switch or during the execution of instructions that modify any stack pointers stored in the PCB.

Most Recent Virtual Address Parameter

Most Recent Virtual Address <31:0>—This field captures the contents of the virtual address pointer register at the time of the machine check. If a machine check, other than a machine check 81, occurs on a read operation, this field represents the virtual address of the location that is being read when the error occurred, plus four. If machine check 81 occurs, this field represents the physical address of the location that was being read when the error occurred, plus four.

If a machine check, other than a machine check 83, occurs on a write operation, this field represents the *virtual address* of a location that is being referenced either when the error occurred, or sometime after the error occurred, plus four. If a machine check 83 occurs, this field represents the *physical address* of the location that was being referenced either when the error occurred, or sometime after the error occurred, plus four. In other words, if the machine check occurs on a write operation, the contents of this field cannot be used for error recovery.

Internal State Information 1 Parameter

Internal State Information 1 is divided into four fields. The contents of these fields is described as follows:

<31:24>—This field captures the opcode of the instruction that was being read or executed at the time of the machine check.

<23:16>—This field captures the internal state of the CVAX CPU chip at the time of the machine check. The four most significant bits are equal to <1110> and the four least significant bits contain highest priority software interrupt <3:0>.

<15:8>—This field captures the state of CADR <7:0> at the time of the machine check.

<7:0>—This field captures the state of the MSER <7:0> at the time of the machine check.

Internal State Information 2 Parameter

Internal State Information 2 is divided into five fields. The contents of these fields is described as follows:

<31:24>—This field captures the internal state of the CVAX CPU chip at the time of the machine check. This field contains SC register <7:0>.

<23:16>—This field captures the internal state of the CVAX CPU chip at the time of the machine check. The two most significant bits are equal to 11 (binary) and the six least significant bits contain state flags <5:0>.

<15>—This field captures the state of the VAX CAN'T RESTART flag at the time of the machine check.

<14:8>—This field captures the internal state of the CVAX CPU chip at the time of the machine check. The three most significant bits are equal to <111> (binary) and the four least significant bits contain ALU condition codes.

<7:0>—This field captures the offset between the virtual address of the start of the instruction being executed at the time of the machine check (saved PC) and the virtual address of the location being accessed (PC) at the time of the machine check.

PC

PC <31:0>—This field captures the virtual address of the start of the instruction being executed at the time of the machine check.

PSL

PSL <31:0>—This field captures the contents of the PSL at the time of the machine check.

5.1.5.4 System Control Block

The system control block (SCB) consists of two pages in main memory that contain the vectors by which interrupts and exceptions are dispatched to the appropriate service routines. The SCB is pointed to by IPR 17, the system control block base register (SCBB), represented in Figure 5-5. The SCB format is presented in Table 5-6.

Figure 5-5 System Control Block Base Register

Table 5-6 System Control Block Format

SCB				
Offset	Interrupt/Exception	Type	Parameter Notes	
00	Unused			IRQ passive release on other VAXes
04	Machine check	Abort	4	Parameters depend on error type
08	Kernel stack not valid	Abort	0	Must be serviced on interrupt stack
0C	Power fail	Interrupt	0	IPL is raised to 1E
10	Reserved/privileged instruction	Fault	0	
14	Customer reserved instruction	Fault	0	XFC instruction
18	Reserved operand	Fault/Abort	0	Not always recoverable
1C	Reserved addressing mode	Fault	0	
20	Access control violation	Fault	2	Parameters are virtual address, status code
24	Translation not valid	Fault	2	Parameters are virtual address, status code
28	Trace pending (TP)	Fault	0	
2C	Breakpoint instruction	Fault	0	
30	Unused			Compatibility mode in other VAXes
34	Arithmetic	Trap/Fault	1	Parameter is type code
38:3C	Unused			
40	CHMK	Trap	1	Parameter is sign-extended operand word
44	CHME	Trap	1	Parameter is sign-extended operand word
48	CHMS	Trap	1	Parameter is sign-extended operand word

Table 5-6 (Cont.) System Control Block Format

SCB				
Offset	Interrupt/Exception	Type	Parameter	Notes
4C	CHMU	Trap	1	Parameter is sign-extended operand word
50	Unused			
54	Corrected read data	Interrupt	0	IPL is 1A (CRD L)
58:5C	Unused			
60	Memory error	Interrupt	0	IPL is 1D (MEMERR L)
64:6C	Unused			
78	Programmable timer 0	Interrupt	0	IPL is 14
7C	Programmable timer 1	Interrupt	0	IPL is 14
80	Unused			
84	Software level 1	Interrupt	0	
88	Software level 2	Interrupt	0	Ordinarily used for AST delivery
8C	Software level 3	Interrupt	0	Ordinarily used for process scheduling
90:BC	Software levels 4-15	Interrupt	0	
C0	Interval timer	Interrupt	0	IPL is 16 (INTTIM L)
C4	Mass storage interface	Interrupt	0	IPL is 14
C8	Emulation start	Fault	10	Same mode exception, FPD=0; parameters are opcode, PC, specifiers
CC	Emulation continue	Fault	0	Same mode exception, FPD=1: no parameters
D0	Unused			
D4	Network interface	Interrupt	0	IPL is 14
D8:DC	Unused			
E0:EC	Reserved for customer or CSS use			

Table 5-6 (Cont.) System Control Block Format

SCB				
Offset	Interrupt/Exception	Type	Parameter Notes	
F0:F4	Unused			Console storage registers on 11/750 and 11/730
F8	Console receiver	Interrupt	0	IPL is 14
FC	Console transmitter	Interrupt	0	IPL is 14
100:1FC	Adapter vectors	Interrupt	0	Not implemented by the KN210
200:3FC	Device vectors	Interrupt	0	Correspond to Q22-bus Vectors 000:1FC; KN210 appends the assertion of bit <9,0>
400:FFC	Unused	Interrupt	0	

5.1.5.5 Diagnostic Processor Hardware Detected Errors

The diagnostic processor is capable of detecting 11 types of error conditions during program execution.

1. CDAL bus parity errors indicated by MSER <6> (on a read) or MEMCSR16 <7> (on a write) being set.
2. Cache tag parity errors indicated by MSER <0> being set.
3. Cache data parity errors indicated by MSER <1> being set.
4. Q22-bus NXM errors indicated by DSER <7> being set.
5. Q22-bus NO SACK errors (no indicator).
6. Q22-bus NO GRANT errors indicated by DSER <2> being set.
7. Q22-bus device parity errors indicated by DSER <5> being set.
8. CDAL bus timeout errors indicated by DSER <4> (only on DMA) being set.
9. Main memory NXM errors indicated by DSER <0> (only on DMA) being set.
10. Main memory correctable errors indicated by MEMCSR16 <29> being set.

11. Main memory uncorrectable errors indicated by MEMCSR16 <31> and DSER <4> (only on DMA) being set.

These errors will cause either a machine check exception, a memory error interrupt, or a corrected read data interrupt, depending on the severity of the error and the reference type that caused the error.

5.1.5.6 Hardware Halt Procedure

The hardware halt procedure is the mechanism by which the hardware assists the firmware in emulating a processor halt. The hardware halt procedure saves the current value of the PC in IPR 42 (SAVPC), and the current value of the PSL, MAPEN<0>, and a halt code in IPR 43 (SAVPSL). The current stack pointer is saved in the appropriate internal register. The PSL is set to 041F 0000₁₆ (IPL=1F, kernel mode, using the interrupt stack) and the current stack pointer is loaded from the interrupt stack pointer. Control is then passed to the resident firmware at physical address 2004 0000₁₆ with the state of the CPU as follows:

Register	New Contents
SAVPC	Saved PC
SAVPSL<31:16, 7:0>	Saved PSL<31:16,7:0>
SAVPSL <15>	Saved MAPEN<0>
SAVPSL <14>	Valid PSL flag (unknown for halt code of 3)
SAVPSL <13:8>	Saved restart code
SP	Current interrupt stack
PSL	041F 0000 ₁₆
PC	2004 0000 ₁₆
MAPEN	0
ICCS	0 (for a halt code of 3)
MSER	0 (for a halt code of 3)
CADR	0 (for a halt code of 3, cache is also flushed)
SISR	0 (for a halt code of 3)
ASTLVL	0 (for a halt code of 3)
All else	Undefined

The firmware uses the halt code in combination with any hardware event indicators to dispatch the execution or interrupt that caused the halt to the appropriate firmware routine (either console emulation, power-up, reboot, or restart). Tables 5–7 and 5–8 list the interrupts and exceptions that can cause halts along with their corresponding halt codes and event indicators.

Table 5-7 Unmaskable Interrupts That can Cause a Halt

Halt Code	Interrupt Condition	Event Indicators
2	External Halt (CVAX HALTIN pin asserted) BHALT asserted on the Q22-bus. BDCOK negated and asserted on the Q22-bus while BPOK stays asserted (Q22-bus REBOOT/RESTART) and SCR <7> is set. BREAK generated by the console	DSER<15> DSER <14> RXDB <11>
3	Hardware Reset (CVAX RESET pin negated) BDCOK and BPOK negated then asserted on the Q22-bus (Power-up) BDCOK negated and asserted on the Q22-bus while BPOK stays asserted (Q22-bus REBOOT/RESTART) and SCR <7> is clear.	- -

Table 5-8 Exceptions That can Cause a Halt

Halt Code	Exception Condition
6	HALT instruction executed in kernel mode

Exceptions While Servicing an Interrupt or Exception

4	Interrupt stack not valid during exception
5	Machine check during normal exception
7	SCB vector bits <1:0> = 11
8	SCB vector bits <1:0> = 10
A	CHMx executed while on interrupt stack
B	CHMx executed to the interrupt stack
10	ACV or TNV during machine check exception
11	ACV or TNV during kernel stack not valid exception
12	Machine check during machine check exception

Table 5-8 (Cont.) Exceptions That can Cause a Halt

Halt Code	Exception Condition
13	Machine check during kernel stack not valid exception
19	PSL <26:24> = 101 during interrupt or exception
1A	PSL <26:24> = 110 during interrupt or exception
1B	PSL <26:24> = 111 during interrupt or exception
1D	PSL <26:24> = 101 during REI
1E	PSL <26:24> = 110 during REI
1F	PSL <26:24> = 111 during REI

5.1.6 System Identification

The system identification register (SID), IPR 62, is a read-only register implemented in the CVAX chip. This 32-bit, read-only register is used to identify the processor type and its microcode revision level (Figure 5-6).

MA-1101-87

Figure 5-6 System Identification Register

Field	Name	RW	Description
31:24	SYS_ TYPE	ro	This field identifies the type of system for a specific processor. (Reads as 06 for the KN210.)
23:16	VERSION	ro	This field identifies the resident version of the firmware EPROM encoded as two hexadecimal digits. For example, if the banner displays V5.0, then this field is 50 (hex).
15:8	SYS_ SUB_ TYPE	ro	This field identifies the particular system sub-type. (Reads as 01 for the KN210.)
7:0	HARDWARE	ro	This field identifies the hardware revision level.

In order to distinguish between different CPU implementations that use the same CPU chip, the KN210 implements a MicroVAX system type register (SYS_TYPE) at physical address 2004 0004₁₆. This 32-bit read-only register is implemented in the KN210 ROM. The format of this register is shown in Figure 5-7.

Figure 5-7 System Type Register

Data Bit	Definition
SYS_TYPE <31:24>	(SYS_TYPE) System type code. This field reads as 01 ₁₆ for all single-processor Q22-bus based systems.
SYS_TYPE <23:16>	(REV LEVEL) Revision level. This field reflects the revision level of the KN210 firmware.
SYS_TYPE <15:8>	(SYS_SUB_TYPE) System sub-type code. This field reads as 10 ₁₆ for the KN210.
SYS_TYPE <7:0>	Reserved for Digital use.

5.1.7 CVAX References

All references by the CPU can be classified into one of three groups:

- Request instruction-stream read references
- Demand data-stream read references
- Write references

5.1.7.1 Instruction-Stream Read References

The diagnostic processor has an instruction prefetcher with a 12-byte (3 longword) instruction prefetch queue (IPQ) for prefetching program instructions from either cache or main memory. Whenever there is an empty longword in the IPQ, and the prefetcher is not halted due to an error, the instruction prefetcher generates an aligned longword, request instruction-stream (I-stream) read reference.

5.1.7.2 Data-Stream Read References

Whenever data is immediately needed by the CPU to continue processing, a demand data-stream (D-stream) read reference is generated. More specifically, demand D-stream references are generated on operand, PTE, SCB, and PCB references.

When interlocked instructions, such as branch on bit set and set interlock (BBSSI) are executed, a demand D-stream read-lock reference is generated. Since the CPU does not impose any restrictions on data alignment (other than the aligned operands of the add aligned word interlocked (ADAWI) and interlocked queue instructions) and since memory can only be accessed one aligned longword at a time, all data read references are translated into an appropriate combination of masked and unmasked, aligned longword read references.

If the required data is a byte, a word within a longword, or an aligned longword, then a single, aligned longword, demand D-stream read reference is generated. If the required data is a word that crosses a longword boundary, or an unaligned longword, then two successive aligned longword demand D-stream read references are generated. Data larger than a longword is divided into a number of successive aligned longword demand D-stream reads, with no optimization.

5.1.7.3 Write References

Whenever data is stored or moved, a write reference is generated. Since the CPU does not impose any restrictions on data alignment (other than the aligned operands of the ADAWI and interlocked queue instructions) and since memory can only be accessed one aligned longword at a time, all data write references are translated into an appropriate combination of masked and unmasked aligned longword write references.

If the required data is a byte, a word within a longword, or an aligned longword, then a single, aligned longword, write reference is generated. If the required data is a word that crosses a longword boundary, or an unaligned longword, then two successive aligned longword write references are generated. Data larger than a longword is divided into a number of successive aligned longword writes.

6

Maintenance Mode Firmware

This chapter describes the KN210 maintenance mode functional firmware. The maintenance mode firmware is VAX-11 code which resides in EPROM on the KN210 processor module. The maintenance mode firmware gains control by one of the following methods:

- On the execution of the maintenance command from the KN210 operational (R3000) console
- On system power-up
- When the onboard CVAX CPU performs a processor restart operation

A processor restart operation is also called a processor halt. Halt, in this context, means only that the control is transferred to the firmware. It does not mean that the processor actually stops executing instructions.

6.1 Maintenance Mode Firmware Features

The maintenance mode firmware is located on one 128 Kbyte EPROM on the KN210 processor module, and provides the following features:

- Diagnostic tests that test all components on the board and verify that the module is working correctly
- Manual bootstrap of an operating system following processor halts
- Manual restart of an operating system following processor halts (Restart in this context is not the same as restarting or resetting the hardware.)
- An interactive command language that allows the user to examine and alter the state of the processor
- Multilanguage support for displaying critical system messages and handling LK201 country specific keyboards

6.1.1 Halt Entry, Exit, and Dispatch

The main purpose of this code is to save the state of the machine on halt entry, invoke the dispatcher, and restore the state of the machine on exit to program I/O mode.

6.1.1.1 Halt Entry - Saving Processor State

The entry code (residing at physical address 2004 0000) is executed whenever a halt occurs. The processor will halt for a variety of reasons. The reason for the halt is stored in PR\$_SAVPSL<13:8>(RESTART_CODE), IPR 43. A complete list of the halt reasons and the associated messages can be found in Table 6-10 in Section 6.8. PR\$_SAVPC, IPR 42, contains the value of the PC when the processor was halted. On a power-up, PR\$_SAVPC is undefined.

One of the first actions the firmware does after a halt is save the current LED code, then it writes an "E" to the diagnostic LEDs. This action occurs within several instructions upon entry into the firmware. The intent of this action is to let the user know that at least some instructions have been successfully executed.

The maintenance mode firmware unconditionally saves the following registers on any halt:

- R0 through R15, the general purpose registers
- PR\$_SAVPSL, the saved PSL register
- PR\$_SCBB, the system control block base register
- DLEDR, the diagnostic LED register
- SSCCR, the SSC configuration register
- ADxMCH & ADxMSK, the SSC address match and mask registers

NOTE

The SSC programmable timer registers are not saved. In some cases, such as bootstrap, the timers are used by the firmware and previous "time" context is lost.

Several registers are unconditionally set to pre-determined values by the firmware on any halt, processor init or bootstrap. This action ensures that the firmware itself can run and protects the board from physical damage.

Registers that fall into this category are:

- SSCCR, the SSC configuration register
- ADxMCH & ADxMSK, the SSC address match and mask registers

- CBTCR, the CDAL bus timeout control register
- TIVRx, the SSC timer interrupt vector registers

On every halt entry, the firmware sets the console serial line baud rate based on the value read from the BDR and extends the halt protection from 8 Kbytes to 128 Kbytes to include all of the EPROM.

6.1.1.2 Halt Exit - Restoring Processor State

When the firmware exits, it uses the currently defined saved context. This context is initially determined by what is saved on entry to the firmware, and may be modified by console commands, or automatic operations such as an automatic bootstrap on power-up.

When restoring the context, the firmware will flush both caches if enabled, and invalidate all translation buffer entries through the internal processor register PR\$_TBIA, IPR 57.

In restoring the context, the console pushes the user's PSL and PC onto the user's interrupt stack, then executes an REI from that stack. This implies that the user's ISP is valid before the firmware can exit. This is done automatically on a bootstrap. However, it is suggested that the SP is set to a valid memory location before issuing the START or CONTINUE command. Furthermore, the user should validate PR\$_SCBB prior to executing a NEXT command, since the firmware utilizes the trace trap vector for this function. At power-up, the user ISP is set to 200 (hex) and PR\$_SCBB is undefined.

6.1.1.3 Halt Dispatch

The action taken by the firmware on a halt is dependent primarily on the following information:

- The Function switch, BDR<7>(ENB_HALT)
- The halt action field, CPMBX<1:0>(HALT_ACTION)
- The halt code, PR\$_SAVPSL<13:8>(RESTART_CODE), in particular the power-up state

In general, the Function switch governs whether external halt conditions are recognized by the KN210. The halt action field in the console program mailbox is a 2-bit field used by operating systems to force the firmware to enter the console, restart, or reboot following a halt, regardless of the setting of the Function switch. The halt (or restart) code is automatically deposited in PR\$_SAVPSL on any processor restart operation. The action taken on a halt is summarized in Table 6-1.

Table 6-1 Halt Action Summary

Function Switch	Power-Up	Halt Action	Action
T	T	x	diagnostics, halt
T	F	0	halt
F	T	x	diagnostics, bootstrap, halt
F	F	0	restart, bootstrap, halt
x	F	1	restart, halt
x	F	2	bootstrap, halt
x	F	3	halt

"T" indicates that the condition is true.

"F" indicates that the condition is false.

"x" indicates that the condition is "don't care".

NOTE

This assumes that the operation switch is set to maintenance mode.

Multiple actions mean that the first action is taken and only if it fails is the next action taken. Diagnostics are an exception, if diagnostics fail the console is entered.

Because the KN210 does not support battery backed up main memory, an operating system restart operation is not attempted on a power-up.

6.1.1.4 External Halts

Several conditions can trigger an external halt (PR\$_{SAVPSL<13:8>(RESTART_CODE) = 2}), and different actions are taken depending on the condition.

An external halt can be caused by:

1. Pressing **Break** on the system console terminal, if halts are enabled (BDR<7>(ENB_HALT) = 1).
2. Assertion of the BHALT line on the Q22-bus, if the SCR<14>(BHALT_ENABLE) bit in the CQBIC is set.
3. Negation of DCOK, if the SCR<7>(DCOK_ACT) bit is set.

NOTE

The switch labeled RESTART on some BA213 system enclosures negates DCOK. The negation of DCOK may also be asserted by the DEQNA sanity timer, or any other Q22-bus module that chooses to implement the Q22-bus restart/reboot protocol.

6.1.2 Power-Up

On a power-up, the maintenance mode firmware performs actions that are unique to this condition. Among these actions are initial power-up tests, identifying a console device, language query, and the remaining diagnostics. Certain actions are dependent on the state of the mode switch on the H3602-SA panel which has three settings: "test," "query," and "normal." This section describes the sequence of events which occurs on power-up.

6.1.2.1 Initial Power-Up Test

The first action performed on power-up is the initial power-up test (IPT). The purpose of the IPT is to verify that the console private NVRAM is valid and if invalid to test and initialize the NVRAM. Prior to checking the NVRAM, the IPT waits for power to stabilize by monitoring SCR<5>(POK). Once power is stable, the IPT then tests to see if the backup batteries failed during the power failure by checking SSSCCR<31>(BLO). If the batteries failed, then the IPT will initialize certain non-volatile data, such as the default boot device, to a known state. In any case, the IPT then initializes other data structures and performs a processor initialization. If the mode switch is set to "test," the IPT then tests the console serial line as described in Section 6.1.3.

NOTE

All IPT failures are considered fatal, and the KN210 will appear to "hang" with a value on the LEDs indicating the point of failure. See Table 6-2 for the LED codes.

6.1.2.2 Locating a Console Device

After the IPT has completed successfully, the firmware attempts to locate a console device and find out what type of device it is. Normally, this is the device attached to the console serial line. In this case, the firmware will send out a device attributes escape sequence to the console serial line to determine the type of terminal attached and the functions it supports. Terminals that do not respond to the device attributes request correctly are assumed to be hardcopy devices.

6-6 Maintenance Mode Firmware

Once a console device has been found, the firmware displays the KN210 banner message, similar to the following:

```
KN210-A V0.0
```

The banner message contains the processor name and the version of the firmware. The letter code in the firmware version indicates whether the firmware is pre-field test ("X"), field test ("T") or an official release ("V"). The first digit indicates the major release number and the trailing digit indicates the minor release number.

Next, if the designated console device supports DEC multinational character set (MCS) and either the battery failed during power failure or the mode switch is set to "query," the firmware prompts for the console language. The firmware first displays the language selection menu shown in Example 6-1.

After the language query, the firmware invokes the ROM-based diagnostics, and eventually displays the console prompt.

6.1.3 Mode Switch Set to "Test"

If the mode switch is set to "test," the console serial line external loopback test is executed at the end of the IPT. The purpose of this test is to verify that the console serial line connections from the KN210 through the H3602-SA panel are intact.

NOTE

An external loopback connector should be inserted in the serial line connector on the H3602-SA panel prior to cycling power to invoke this test.

During this test the firmware toggles between two states, active and passive, each a few seconds long and each displaying a different number on the LEDs.

During the active state (about 3 seconds long), the LEDs are set to 7. In this state the firmware reads the baud rate and mode switch, then transmits and receives a character sequence. If the mode switch has been moved from the "test" position, the firmware exits the test and continues as if on a normal power-up.

During the passive state (about 7 seconds long), the LEDs are set to 3.

If at any time the firmware detects an error (parity, framing, overflow or no characters), the firmware hangs with a 7 on the LEDs.

6.1.4 Mode Switch Set to "Query"

If the mode switch is set to "query" (or the firmware detects that the battery failed during a power loss), the firmware queries the user for a language which is used for displaying critical system messages.

The language query menu is shown in Example 6-1.

- 1) Dansk
 - 2) Deutsch (Deutschland/Österreich)
 - 3) Deutsch (Schweiz)
 - 4) English (United Kingdom)
 - 5) English (United States/Canada)
 - 6) Español
 - 7) Français (Canada)
 - 8) Français (France/Belgique)
 - 9) Français (Suisse)
 - 10) Italiano
 - 11) Nederlands
 - 12) Norsk
 - 13) Português
 - 14) Suomi
 - 15) Svenska
- (1..15):

Example 6-1 Language Selection Menu

If no response is received within 30 seconds, the language defaults to English (United States/Canada).

NOTE

This action is only taken if the console device supports DEC MCS. Any console device that does not support DEC MCS, such as a VT100, defaults to English (United States/Canada).

After this inquiry, the firmware proceeds as if the mode switch were set to "normal," as described in Section 6.1.5.

6-8 Maintenance Mode Firmware

```
Performing normal system tests.  
49..48..47..46..45..44..43..42..41..40..39..38..37..36..35..34..  
33..32..31..30..29..28..27..26..25..24..23..22..21..20..19..18..  
17..16..15..14..13..12..11..10..09..08..07..06..05..04..03..  
Tests completed.
```

Example 6-2 Normal Diagnostic Countdown

6.1.5 Mode Switch Set to "Normal"

If the mode selected is "normal", then the next step in the power-up sequence is to execute the bulk of ROM-based diagnostics. In addition to message text, a *countdown* is displayed to indicate diagnostic test progress. A successful diagnostic countdown is shown in Example 6-2.

In the case of diagnostic failures, a diagnostic register dump is performed similar to that shown in Example 6-3. Depending on the failure, the remaining diagnostics may execute and the countdown continue. For a detailed description of the register dump refer to Section 6.4.

```
Performing normal system tests.  
49..48..47..46..45..44..43..42..41..40..39..38..37..36..35..34..  
  
?34 2 08 FF 00 0000  
  
P1=00000000 P2=00000003 P3=00000031 P4=00000011 P5=00002000  
P6=FFFFFFFF P7=00000000 P8=00000000 P9=00000000 P10=2005438F  
r0=00114B98 r1=FFFFFFFF r2=2005D2F0 r3=55555555 r4=AAAAAAAA  
r5=00000000 r6=AAAAAAAA r7=00000000 r8=00000000 ERF=80000180  
33..32..31..30..29..28..27..26..25..24..23..22..21..20..19..18..  
17..16..15..14..13..12..11..10..09..08..07..06..05..04..03..  
Normal operation not possible.
```

Example 6-3 Abnormal Diagnostic Countdown

If the diagnostics have successfully completed and halts are enabled, the firmware displays the console prompt (>>>) and enters console I/O mode.

6.1.6 LED Codes

In addition to the console diagnostic countdown, a hexadecimal value is displayed by the LEDs on the H3602-SA panel. The same value is displayed by the four red LEDs on the KN210 processor module. The purpose of the LED display is to improve fault isolation, when there is no console terminal or when the hardware is incapable of communicating with the console terminal. Table 6-2 lists all LED codes and the associated actions which are performed at power-up. The LED code is changed before the corresponding test or action is performed.

Table 6-2 LED Codes

LED Value	Actions
F	Initial state on power-up, no code has executed
E	Entered ROM, some instructions have executed
D	Waiting for power to stabilize (POK)
C	SSC and ROM tests
B	CVAX tests
A	R3000 tests
9	CMCTL and memory tests
8	CQBIC (Q22-bus) tests
7	Console loopback tests
6	DSSI subsystem tests
5	Ethernet subsystem tests
4	CVAX console I/O mode
3	R3000 console I/O mode
2	CVAX primary/secondary bootstrap
1	R3000 primary/secondary bootstrap
0	Operating system running

6.2 Console Service

The KN210 is by definition halted whenever the console program is running and the triple angle prompt (>>>) is displayed on the console terminal. When halted, the firmware provides most of the services of a standard VAX console through the device that is designated as the system console.

6.2.1 Console Control Characters

In console I/O mode, several characters have special meanings.

- **[Return]** — ends a command line. No action is taken on a command until after it is terminated by a carriage-return. A null line terminated by a **[Return]** is treated as a valid, null command. No action is taken, and the console re-prompts for input. **[Return]** is echoed as carriage-return, line feed.
- **[Rubout]** — deletes the character that the operator previously typed. What appears on the console terminal depends on whether the terminal is a video terminal or a hardcopy terminal. For hardcopy terminals, when a **[Rubout]** is pressed, the console echoes with a backslash (\), followed by the character being deleted. If the operator presses additional **[Rubout]**s, the additional characters deleted are echoed. When the operator types a non-rubout character, the console echoes another backslash, followed by the character typed. The result is to echo the characters deleted, surrounding them with backslashes.

For example:

```
The operator types: EXAMI;E<rubout><rubout>NE<CR>
The console echoes: EXAMI;E\E;\NE<CR>
The console sees the command line: EXAMINE<CR>
```

For video terminals, when **[Rubout]** is pressed the previous character is erased from the screen and the cursor is restored to its previous position.

The console does not delete characters past the beginning of a command line. If the operator presses more **[Rubout]**s than there are characters on the line, the extra rubouts are ignored. If **[Rubout]** is pressed on a blank line, it is ignored.

- **[Ctrl] C** — causes the console to echo ^C and to abort processing of a command. **[Ctrl] C** has no effect as part of a binary load data stream. **[Ctrl] C** clears **[Ctrl] S**, and reenables output stopped by **[Ctrl] O**.

- **Ctrl O** — causes the console to throw away transmissions to the console terminal until the next **Ctrl O** is entered. **Ctrl O** is echoed as ^O<CR> when it disables output, but is not echoed when it reenables output. Output is reenabled if the console prints an error message, or if it prompts for a command from the terminal. Displaying a REPEAT command does not reenables output. When output is reenabled for reading a command, the console prompt is displayed. Output is also enabled **Ctrl S**.
- **Ctrl Q** — resumes output to the console terminal. Additional **Ctrl Q**s are ignored. **Ctrl S** and **Ctrl Q** are not echoed.
- **Ctrl S** — stops output to the console terminal until **Ctrl Q** is typed. **Ctrl S** and **Ctrl Q** are not echoed.
- **Ctrl U** — causes the console to echo ^U<CR>, and deletes the entire line. If **Ctrl U** is typed on an empty line, it is echoed, and the console prompts for another command.
- **Ctrl R** — causes the console to echo <CR><LF> followed by the current command line. This function can be used to improve the readability of a command line that has been heavily edited. When **Ctrl C** is typed as part of a command line, the console deletes the line as it does with **Ctrl U**.
- **Break** - If the console is in console I/O mode, **Break** is equivalent to **Ctrl C**, but is echoed as "^C".

NOTE

If the local console is in program I/O mode and halts are disabled, **Break is ignored. If the console is in program I/O mode and halts are enabled, **Break** causes the processor to halt and enter console I/O mode.**

Control characters are typed by pressing the character key while holding down the **Ctrl** key.

If an unrecognized control character (ASCII code less than 32 decimal or between 128 and 159 decimal) is typed, it is echoed as up arrow followed by the character with ASCII code 64 greater. For example, BEL (ASCII code 7) is echoed as "^G", since capital G is ASCII code 7+64=71. When a control character is deleted with **Rubout**, it is echoed the same way. After echoing the control character, the console processes it like a normal character. Commands with control characters are invalid, unless they are part of a comment, and the console will respond with an error message.

Note that control codes from 128 to 159, the C1 control codes, cannot be entered by any present Digital terminal. The character with code 7 and the character with code 135 will both echo as "^G".

6.2.2 Console Command Syntax

The console accepts commands of lengths up to 80 characters. It responds to longer commands with an error message. The count does not include rubouts, rubbed out characters, or the terminating carriage-return.

Commands may be abbreviated. Abbreviations are formed by dropping characters from the end of a keyword, as long as the resulting keyword is still unique. Most commands can be uniquely expressed with their first character.

Multiple adjacent spaces and tabs are treated as a single space by the console. Leading and trailing spaces and tabs are ignored. Tabs are echoed as spaces.

Command qualifiers can appear after the command keyword, or after any symbol or number in the command. A qualifier is any contiguous set of non whitespace characters that is started with a slash (ASCII code 47 decimal).

All numbers (addresses, data, counts) are in hexadecimal. Note, though, that symbolic register-names number the registers in decimal. The console does not distinguish between upper and lower case either in numbers or in commands; both are accepted.

6.2.3 Console Command Keywords

The firmware implements a variant of the VAX SRM console command set. The only commands defined in the VAX SRM and not supported by the KN210 are MICROSTEP, LOAD, and @. The CONFIGURE, HELP, MOVE, SEARCH and SHOW commands have been added to the command set to facilitate system debugging and access to system parameters. In general, however, the KN210 console is similar to other VAX consoles.

Table 6-3 lists command and qualifier keywords.

Table 6-3 Command, Parameter, and Qualifier Keywords

Command Keywords		
Processor Control	Data Transfer	Console Control
B*OOT	D*EPOSIT	CONF*IGURE
C*ONTINUE	E*XAMINE	F*IND
H*ALT	M*OVE	R*EPEAT
I*NITIALIZE	SEA*RCH	SET
N*EXT	X	SH*OW
S*TART	EXIT	T*EST
U*NJAM		!
SET & SHOW Parameter Keywords		
BO*OT	BF*L(A)G	DE*VICE
DS*SI	ET*HERNET	H*OST
L*ANGUAGE	M*EMORY	Q*BUS
RL*V12	U*QSSP	VERS*ION
Qualifier Keywords		
Data Control	Address Space Control	Command Specific
/B	/G	/IN*STRUCTION
/W	/I	/NO*T
/L	/P	/R5: or /
/Q	/V	/RP*B or /ME*M
/N:	/M	/F*ULL
/S*TEP:	/U	/DU*P or
		/MA*INTENANCE
/WR*ONG		/DS*SI or /U*QSSP
		/DI*SK or /T*APE
		/SE*RVICE
<p>"*" indicates the minimal number of characters that are required to uniquely identify the keyword.</p>		

A complete summary of the console commands is provided in Table 6-5 following the command descriptions.

6.2.4 Console Command Qualifiers

All qualifiers in the console command syntax are global. That is, they may appear in any place on the command line after the command keyword.

All qualifiers have unique meanings throughout the console, regardless of the command. For example, the "/B" qualifier always means byte.

Table 6-6 is a summary of the qualifiers recognized by the KN210 console.

6.2.5 Command Address Specifiers

Several commands take an address or addresses as arguments. In the context of the console, an address has two components, the address space, and the offset into that space. The console supports 6 address spaces: physical memory (/P qualifier), virtual memory (/V qualifier), general purpose registers (/G qualifier), internal processor registers (/I qualifier), protected memory (/U qualifier), and the PSL (/M qualifier).

The address space that the console references is inherited from the previous console reference, unless explicitly specified. The initial address space reference is PHYSICAL.

The KN210 console supports symbolic references to addresses. A symbolic reference simultaneously defines the address space for a given symbol. Table 6-4 lists the symbolic addresses supported by the console grouped according to address space.

Table 6-4 Console Symbolic Addresses

Symbol	Address	Symbol	Address
/G—General Purpose Registers			
R0	00	R11	0B
R1	01	R12	0C
R2	02	R13	0D
R3	03	R14	0E
R4	04	R15	0F
R5	05	AP	0C
R6	06	FP	0D
R7	07	SP	0E
R8	08	PC	0F
R9	09	PSL	—

Note: All symbolic values in this table are in hexadecimal.

Table 6-4 (Cont.) Console Symbolic Addresses

Symbol	Address	Symbol	Address
/G—General Purpose Registers			
R10	0A		
/I - Internal Processor Registers			
PR\$_KSP	00	PR\$_SISR	15
PR\$_ESP	01	PR\$_ICCR	18
PR\$_SSP	02	PR\$_RXCS	20
PR\$_USP	03	PR\$_RXDB	21
PR\$_ISP	04	PR\$_TXCS	22
PR\$_POBR	08	PR\$_TXDB	23
PR\$_POLR	09	PR\$_TBDR	24
PR\$_P1BR	0A	PR\$_CADR	25
PR\$_P1LR	0B	PR\$_MCESR	26
PR\$_SBR	0C	PR\$_MSER	27
PR\$_SLR	0D	PR\$_SAVPC	2A
PR\$_PCBB	10	PR\$_SAVPSL	2B
PR\$_SCBB	11	PR\$_	37
		IORESET	
PR\$_IPL	12	PR\$_MAPEN	38
PR\$_ASTLV	13	PR\$_TBIA	39
PR\$_SIRR	14	PR\$_TBIS	3A
PR\$_NICR	19	PR\$_SID	3E
PR\$_ICR	1A	PR\$_TBCHK	3F
PR\$_TODR	1B		
/P - Physical (VAX I/O Space)			
QBIO	2000 0000	QBMEM	3000 0000
QBMBR	2008 0010		
ROM	2004 0000		
BDR	2008 4004		
DSCR	2008 0000	DSER	2008 0004
DMEAR	2008 0008	DSEAR	2008 000C
IPCRO	2000 1F40		
SSC_RAM	2014 0400	SSC_CR	2014 0010

Table 6-4 (Cont.) Console Symbolic Addresses

Symbol	Address	Symbol	Address
/P - Physical (VAX I/O Space)			
SSC_CDAL	2014 0020	SSC_DLEDR	2014 0030
SSC_	2014 0130	SSC_	2014 0134
AD0MAT		AD0MSK	
SSC_	2014 0140	SSC_	2014 0144
AD1MAT		AD1MSK	
SSC_TCR0	2014 0100	SSC_TIR0	2014 0104
SSC_TNIR0	2014 0108	SSC_TIVR0	2014 010C
SSC_TCR1	2014 0110	SSC_TIR1	2014 0114
SSC_TNIR1	2014 0118	SSC_TIVR1	2014 011C
MEMCSR0	2008 0100	MEMCSR1	2008 0104
MEMCSR2	2008 0108	MEMCSR3	2008 010C
MEMCSR4	2008 0110	MEMCSR5	2008 0114
MEMCSR6	2008 0118	MEMCSR7	2008 011C
MEMCSR8	2008 0120	MEMCSR9	2008 0124
MEMCSR10	2008 0128	MEMCSR11	2008 012C
MEMCSR12	2008 0130	MEMCSR13	2008 0134
MEMCSR14	2008 0138	MEMCSR15	2008 013C
MEMCSR16	2008 0140	MEMCSR17	2008 0144

Any Address Space

"@"	The last location successfully referenced in an EXAMINE or DEPOSIT command.
"+"	The location immediately following the last location successfully referenced in an EXAMINE or DEPOSIT command. For references to physical or virtual memory spaces, the location referenced is the last address, plus the size of the last reference (1 for byte, 2 for word, 4 for longword, 8 for quadword). For other address spaces, the address is the last address referenced plus one.

Table 6-4 (Cont.) Console Symbolic Addresses

Symbol	Address	Symbol	Address
Any Address Space			
"-"	The location immediately preceding the last location successfully referenced in an EXAMINE or DEPOSIT command. For references to physical or virtual memory spaces, the location referenced is the last address minus the size of this reference (1 for byte, 2 for word, 4 for longword, 8 for quadword). For other address spaces, the address is the last addressed referenced minus one.		
"@"	The location addressed by the last location successfully referenced in an EXAMINE or DEPOSIT command.		

6.2.6 References to Processor Registers and Memory

The KN210 console is implemented by macrocode executing from EPROM. Actual processor registers cannot be modified by the console command interpreter. When the console is entered, the console saves the processor registers in console memory and all command references to them are directed to the corresponding saved values, not to the registers themselves.

When the console reenters program I/O mode, the saved registers are restored and any changes become operative only then. References to processor memory are handled normally. The binary load and unload command cannot reference the console memory pages.

The following registers are saved by the console, and any direct reference to these registers will be intercepted by the console and the access will be to the saved copies:

- R0...R15 - the general purpose registers.
- PR\$_IPL - the interrupt priority level register.
- PR\$_SCBB - the system control block base register.
- PR\$_ISP - the interrupt stack pointer.
- PR\$MAPEN - the memory management enable register.

The following registers are also saved, yet may be accessed directly through console commands. Writing values to these registers may make the console inoperative.

- PR\$_SAVPC - the halt PC.
- PR\$_SAVPSL - the halt PSL.
- ADxMCH/ADxMSK - the SSC address decode and match registers.
- SSCCR - the SSC configuration register.
- DLEDR - the SSC diagnostic LED register.

6.2.7 Console Commands

The following sections define the commands accepted by the console, when it is in console I/O mode. The following conventions are used to describe command syntax:

- [] - denotes command elements that are optional.
- { } - denotes a command element.
- ... - denotes a list of command elements.

6.2.7.1 BOOT

Format :

BOOT [qualifier] [{boot_device}[:]]

Description :

The console initializes the processor and transfers execution to VMB. VMB attempts to boot the operating system from the specified device or the default boot device if none is specified. The console qualifies the bootstrap operation by passing a boot flags to VMB in R5. A more detailed description of the bootstrap process and how the default bootstrap device is determined is described in Section 6.3.

In the case where either the qualifiers or the device name is absent, then the corresponding default value is used. Explicitly stating the boot flags or the boot device overrides the current default value for the current boot request, but does not change the corresponding default value in NVRAM.

There are three mechanisms by which the default boot device and boot flags may be set.

1. The operating system may write a default boot device and flags into the appropriate locations in NVRAM (Section 6.7.3).

2. The user may explicitly set the default boot device and boot flags with the console SET BOOT and SET BFLAG commands respectively.
3. The console will prompt the user for the default boot device, if any of the following conditions are met:
 - The power-up mode switch is set to "query" mode.
 - The console detects that the battery failed, and therefore the contents of NVRAM are no longer valid.
 - The console detects that the default boot device has not been explicitly set by the user. Either a previous device query timed out and defaulted to ESA0 or neither (1) nor (2) has been performed. Simply stated, the console will prompt the user on each and every power-up for a default boot device, until such a request has been satisfied.

On power-up if no default boot device is specified in NVRAM, the console issues a list of potential bootable devices and then queries the user for a device name. If no device name is entered within 30 seconds, ESA0 is used. However, ESA0 does not become the "default" boot device.

Qualifiers :

- **/R5:{boot_flags}** Boot flags is a 32-bit hex value, that is passed to VMB in R5. No interpretation of this value is performed by the console. Refer to Figure 6-1 for the bit assignments of R5. A default boot flags longword may be specified using the SET BFLAG command and displayed with the SHOW BFLAG command.
- **/{boot_flags}** Equivalent to the form above.

Arguments :

- **[{boot_device}]** The boot device name may be any arbitrary character string, with a maximum length of 17 characters. Longer strings cause a "VAL TOO BIG" error message to be issued from the console. Otherwise the console makes no attempt at interpreting or validating the device name. The console converts the string to all upper case, and passes VMB a string descriptor to this device name in R0. A default boot device may be specified using the SET BOOT command and displayed with the SHOW BOOT command. The factory default device is the Ethernet port, ESA0.

6-20 Maintenance Mode Firmware

Examples :

```
>>>show boot
DUA0
>>>show bflag
0
>>>b                ! Boot using default boot flags and device.
(BOOT/R5:0 DUA0)

2..
-DUA0

>>>bo xqa0          ! Boot using default boot flags and specified device.
(BOOT/R5:0 XQA0)

2..
-XQA0

>>>boot/10         ! Boot using specified boot flags and default device.
(BOOT/R5:10 DUA0)

2..
-DUA0

>>>boot /r5:220 xqa0 ! Boot using specified boot flags and device.
(BOOT/R5:220 XQA0)

2..
-XQA0
```

6.2.7.2 CONFIGURE

Format :

CONFIGURE

Description :

CONFIGURE is similar to the VMS SYSGEN CONFIG utility. This feature simplifies system configuration by providing information that is typically available only with a running operating system.

The CONFIGURE command invokes an interactive mode that permits the user to enter Q22-bus device names, then generates a table of Q22-bus I/O page device CSR addresses and device vectors.

Qualifiers :

None

Arguments :

None

Examples :

```
>>>config
Enter device configuration, HELP, or EXIT
Device,Number? help
Devices:
  LPV11          KXJ11          DZQ11          DZV11          DFA01
                TSV05          RXV21          DRV11W         DRV11B         DPV11
  DMV11          DELQA          DEQNA          DESQA          RQDX3          KDA50
  RRD50          RQC25          KFQSA-DISK     TQK50          TQK70          TU81E
  RV20           KFQSA-TAPE     KMV11          IEQ11          DHQ11          DHV11
  CXA16          CXB16          CXY08          VCB01          QVSS           LNV11
  LNV21          QPSS           DSV11          ADV11C         AAV11C         AXV11C
  KWV11C         ADV11D         AAV11D         VCB02          IDV11B         DRV11J
  DRQ3B          VSV21          IBQ01          IDV11A         IDV11C         IDV11C
  IDV11D         IAV11A         IAV11B         MIRA           ADQ32          DTC04
  DESNA          IGQ11
Numbers:
  1 to 255, default is 1
Device,Number? rqdx3,2
Device,Number? dhv11
Device,Number? tqk50
Device,Number? tqk70
Device,Number? exit

Address/Vector Assignments
-772150/154 RQDX3
-760334/300 RQDX3
-774500/260 TQK50
-760444/304 TQK70
-760500/310 DHV11
>>>
```

6.2.7.3 CONTINUE

Format :

CONTINUE

Description :

The processor begins instruction execution at the address currently contained in the program counter. Processor initialization is not performed. The console enters program I/O mode. Internally, the continue command pushes the user's PC and PSL onto the user's ISP, and then executes an REI instruction. This implies that the user's ISP is pointing to some valid memory.

Qualifiers :

None

Arguments :

None

Examples :

```
>>>continue  
>>>
```

6.2.7.4 DEPOSIT

Format :

DEPOSIT [qualifier_list] {address} {data} [{data}...]

Description :

Deposits the data into the address specified. If no address space or data size qualifiers are specified, the defaults are the last address space and data size used in a DEPOSIT, EXAMINE, MOVE or SEARCH command. After processor initialization, the default address space is physical memory, the default data size is a longword and the default address is zero. If conflicting address space or data sizes are specified, the console ignores the command and issues an error message.

Qualifiers :

- **/B** — The data size is byte.
- **/W** — The data size is word.
- **/L** — The data size is longword.
- **/Q** — The data size is quadword.
- **/G** — The address space is the general purpose register set, R0 through R15. The data size is always long.
- **/I** — The address space is internal processor registers (IPRs). These are the registers only accessible by the MTPR and MFPR instructions. The data size is always long.
- **/M** — The address space is the processor status longword (PSL).
- **/P** — The address space is physical memory.
- **/V** — The address space is virtual memory. All access and protection checking occur. If the access would not be allowed to a program running with the current PSL, the console issues an error message. Virtual space DEPOSITs cause the PTE<M> bit to be set. If memory mapping is not enabled, virtual addresses are equal to physical addresses.
- **/U** — Access to console private memory is allowed. This qualifier also disables virtual address protection checks. On virtual address writes, the PTE<M> bit will not be set if the /U qualifier is present. This qualifier is not inherited, and must be respecified on each command.
- **/N:{count}** — The address is the first of a range. The console deposits to the first address, then to the specified number of succeeding addresses. Even if the address is the symbolic address (-), the succeeding addresses are at larger addresses. The symbolic address specifies only the starting address, not the direction of succession. For repeated references to preceding addresses, use REPEAT DEPOSIT - <DATA>.
- **/STEP:{size}** — The number to add to the current address. Normally this defaults to the data size, but is overridden by the presence of this qualifier. This qualifier is not inherited.
- **/WRONG** — The ECC bits for this data forced to a value of 3 (ECC bits of 3 will always generate a double bit error).

Arguments :

- **{address}** — A long word address that specifies the first location into which data is deposited. The address can be any legal address specifier as defined in Section 6.2.5 and Table 6-4.
- **{data}** — The data to be deposited. If the specified data is larger than the deposit data size, the console ignores the command and issues an error response. If the specified data is smaller than the deposit data size, it is extended on the left with zeros.
- **[{data}]** — Additional data to be deposited (up to a maximum of 6 values).

Examples :

```
>>>d/p/b/n:1FF 0 0      ! Clear first 512 bytes of physical memory.
>>>d/v/l/n:3 1234 5     ! Deposit 5 into four longwords starting at
                        ! virtual memory address 1234.
>>>d/n:8 R0 FFFFFFFF    ! Loads GPRs R0 through R8 with -1.
>>>d/n:200 - 0          ! Starting at previous address, clear 513 bytes.
>>>d/l/p/n:10/s:200 0 8 ! Deposit 8 in the first longword of
                        ! the first 17 pages in physical memory.
>>>
```

6.2.7.5 EXAMINE*Format :*

EXAMINE [qualifier_list] [{address}]

Description :

Examines the contents of the memory location or register specified by the address. If no address is specified, + is assumed. The display line consists of a single character address specifier, the hexadecimal physical address to be examined, and the examined data also in hexadecimal.

EXAMINE uses the same qualifiers as DEPOSIT. However, the /WRONG qualifier will cause examines to ignore ECC errors on reads from physical memory. Additionally, the examine command supports an /INSTRUCTION qualifier, which will disassemble the instructions at the current address.

Qualifiers :

- **/B** — The data size is byte.
- **/W** — The data size is word.
- **/L** — The data size is longword.
- **/Q** — The data size is quadword.
- **/G** — The address space is the general purpose register set, R0 through R15. The data size is always long.
- **/I** — The address space is internal processor registers (IPRs). These are the registers only accessible by the MTPR and MFPR instructions. The data size is always long.
- **/M** — The address space is the processor status longword (PSL).
- **/P** — The address space is physical memory. Note that when virtual memory is examined, the address space and address in the response are the translated physical address.
- **/V** — The address space is virtual memory. All access and protection checking occur. If the access would not be allowed to a program running with the current PSL, the console issues an error message. If memory mapping is not enabled, virtual addresses are equal to physical addresses.
- **/M** — The address space and display are the PSL. The data size is always long.
- **/U** — Access to console private memory is allowed. This qualifier also disables virtual address protection checks. This qualifier is not inherited, and must be respecified with each command.
- **/N:{count}** — The address is the first of a range. The console deposits to the first address, then to the specified number of succeeding addresses. Even if the address is the symbolic address (-), the succeeding addresses are at larger addresses. The symbolic address specifies only the starting address, not the direction of succession. For repeated references to preceding addresses, use REPEAT EXAMINE - <DATA>.
- **/STEP:{size}** — The number to add to the current address. Normally this defaults to the data size, but is overridden by the presence of this qualifier. This qualifier is not inherited.

6-26 Maintenance Mode Firmware

- **/WRONG** — ECC errors on this read access to main memory are ignored.
- **/INSTRUCTION** — Disassemble and display the VAX Macro-32 instruction at the specified address.

Arguments :

- **[{address}]** — A longword address that specifies the first location to be examined. The address can be any legal address specifier as defined in Section 6.2.5 and Table 6-4. If no address is specified, + is assumed.

Examples :

```
>>>e pc                                ! Examine the PC.
  G 0000000F FFFFFFFC
>>>e sp                                ! Examine the SP.
  G 0000000E 00000200
>>>e psl                               ! Examine the PSL.
  M 00000000 041F0000
>>>e/m                                 ! Examine PSL another way.
  M 00000000 041F0000
>>>e r4/n:5                             ! Examine R4 through R9.
  G 00000004 00000000
  G 00000005 00000000
  G 00000006 00000000
  G 00000007 00000000
  G 00000008 00000000
  G 00000009 801D9000
>>>e pr$_scbb                          ! Examine the SCBB, IPR 17.
  I 00000011 2004A000
>>>e/p 0                                ! Examine local memory 0.
  P 00000000 00000000
>>>e /ins 20040000                      ! Examine 1st byte of EPROM.
  P 20040000 11 BRB 20040019
>>>e /ins/n:5 20040019                 ! Disassemble from branch.
  P 20040019 D0 MOVL I^#20140000,@#20140000
  P 20040024 D2 MCOML @#20140030,@#20140502
  P 2004002F D2 MCOML S^#0E,@#20140030
  P 20040036 7D MOVQ R0,@#201404B2
  P 2004003D D0 MOVL I^#201404B2,R1
  P 20040044 DB MFPR S^#2A,B^44(R1)
>>>e/ins                                ! Look at next instruction.
  P 20040048 DB MFPR S^#2B,B^48(R1)
>>>
```

6.2.7.6 EXIT

Format :

EXIT

Description :

EXIT from maintenance mode and Start/Return to the operational (R3000) console. This command will idle the CVAX CPU. Executing the maintenance command on the operational (R3000) console will return control to CVAX and return the user to the maintenance mode prompt. The EXIT command implementation should comply with the following guidelines:

- Prior to execution of the command, all errors must be handled.
- Steps of the EXIT command are:
 1. Q22-bus map must be disabled.
 2. Turn off interrupts to the CVAX.
 3. SPR<31> set to 1.

Qualifiers :

None

Arguments :

None

Examples :

```
>>>exit           ! Return to operational console.
>>maint          ! Enter the maintenance console.
>>>
```


6.2.7.8 HALT

Format :

HALT

Description :

This command has no effect and is included for compatability with other consoles.

Qualifiers :

None

Arguments :

None

Examples :

```
>>>halt                ! Pretend to halt.  
>>>
```

6.2.7.9 HELP

Format :

HELP

Description :

This command has been included to help the console operator answer simple questions about command syntax and usage.

Qualifiers :

None

Arguments :

None

6-30 Maintenance Mode Firmware

Examples :

```
>>>help
```

Following is a brief summary of all the commands supported by the console:

```
UPPERCASE denotes a keyword that you must type in
|         denotes an OR condition
[]       denotes optional parameters
<>      denotes a field that must be filled in
         with a syntactically correct value
```

Valid qualifiers:

```
/B /W /L /Q /INSTRUCTION
/G /I /V /P /M
/STEP: /N: /NOT
/WRONG /U
```

Valid commands:

```
DEPOSIT [<qualifiers>] <address> [<datum> [<datum>]]
EXAMINE [<qualifiers>] [<address>]
MOVE [<qualifiers>] <address> <address>
SEARCH [<qualifiers>] <address> <pattern> [<mask>]
SET BFL(A)G <boot_flags>
SET BOOT <boot_device>[:]
SET HOST/DUP/DSSI <node_number> [<TASK>]
SET HOST/DUP/UQSSP </DISK | /TAPE> <controller_number> [<task>]
SET HOST/DUP/UQSSP <physical_CSR_address> [<task>]
SET HOST/MAINTENANCE/UQSSP/SERVICE <controller_number>
SET HOST/MAINTENANCE/UQSSP <physical_CSR_address>
SET LANGUAGE <language_number>
SHOW BFL(A)G
SHOW BOOT
SHOW DEVICE
SHOW DSSI
SHOW ETHERNET
SHOW LANGUAGE
SHOW MEMORY [ /FULL]
SHOW QBUS
SHOW UQSSP
SHOW VERSION
HALT
INITIALIZE
UNJAM
CONTINUE
EXIT
```



```

START <address>
REPEAT <command>
X <address> <count>
FIND [/MEMORY | /RPB]
TEST [<test_code> [<parameters>]]
BOOT [/R5:<boot_flags> | /<boot_flags>] [<boot_device>[:]]
NEXT [count]
CONFIGURE
HELP
>>>

```

6.2.7.10 INITIALIZE

Format :

INITIALIZE

Description :

A processor initialization is performed. The following registers are initialized.

Register	Initialized Value
PSL	041F 0000
IPL	1F
ASTLVL	4
SISR	0
ICCS	Bits <6> and <0> are clear, the rest are UNPREDICTABLE.
RXCS	0
TXCS	80
MAPEN	0
CPU cache	Disabled, all entries invalid
Instruction buffer	Unaffected
Console previous reference	Longword, physical, address 0
TODR	Unaffected

Register	Initialized Value
Main memory	Unaffected
General registers	Unaffected
Halt code	Unaffected
Bootstrap in progress flag	Unaffected
Internal restart in progress flag	Unaffected

The maintenance mode firmware performs the following additional initialization:

- The CDAL bus timer is initialized.
- The address decode and match registers are initialized.
- The programmable timer interrupt vectors are initialized.
- The BDR registers are read to determine the baud rate, and then the SSSCR is configured accordingly.
- All error status bits are cleared.

Qualifiers :

None

Arguments :

None

Examples :

```
>>>init
>>>
```

6.2.7.11 MOVE*Format :***MOVE [qualifier-list] {src_address} {dest_address}***Description :*

The console copies the block of memory starting at the source address to a block beginning at the destination address. Typically, this command is used with the /N: qualifier to transfer large blocks of data. The destination will correctly reflect the contents of the source, regardless of the overlap between the source and the data.

The MOVE command actually performs byte, word, longword, and quadword reads and writes as needed in the process of moving the data. Moves are only supported for the PHYSICAL and VIRTUAL address spaces.

Qualifiers :

- **/B** — The data size is byte.
- **/W** — The data size is word.
- **/L** — The data size is longword.
- **/Q** — The data size is quadword.
- **/P** — The address space is physical memory.
- **/V** — The address space is virtual memory. All access and protection checking occur. If the access would not be allowed to a program running with the current PSL, the console issues an error message. Virtual space MOVEs cause the destination PTE<M> bit to be set. If memory mapping is not enabled, virtual addresses are equal to physical addresses.
- **/U** — Access to console private memory is allowed. This qualifier also disables virtual address protection checks. On virtual address writes, the PTE<M> bit will not be set if the /U qualifier is present. This qualifier is not inherited, and must be respecified on each command.
- **/N:{count}** — The address is the first of a range. The console deposits to the first address, then to the specified number of succeeding addresses. Even if the address is the symbolic address (-), the succeeding addresses are at larger addresses. The symbolic address specifies only the starting address, not the direction of succession.

6-34 Maintenance Mode Firmware

- **/STEP:{size}** — The number to add to the current address. Normally this defaults to the data size, but is overridden by the presence of this qualifier. This qualifier is not inherited.
- **/WRONG** — On reads, ECC errors on the access of data in main memory are ignored. On writes, the ECC bits for this data are forced to a value of 3.

Arguments :

- **{src_address}** — A longword address that specifies the first location of the source data to be copied.
- **{dest_address}** — A longword address that specifies the destination of the first byte of data. These addresses may be any legal address specifier as defined in Section 6.2.5 and Table 6-4. If no address is specified, + is assumed.

Examples :

```
>>>e /n:4 0                                ! Observe destination.
P 00000000 00000000
P 00000004 00000000
P 00000008 00000000
P 0000000C 00000000
P 00000010 00000000
>>>e /n:4 200                              ! Observe source data.
P 00000200 58DD0520
P 00000204 585E04C1
P 00000208 00FF8FBB
P 0000020C 5208A8D0
P 00000210 540CA8DE
>>>move /n:4 200 0                         ! Move the data.
>>>e /n:4 0                                ! Observe the destination.
P 00000000 58DD0520
P 00000004 585E04C1
P 00000008 00FF8FBB
P 0000000C 5208A8D0
P 00000010 540CA8DE
>>>
```

6.2.7.12 NEXT

Format :

NEXT {count}

Description :

The NEXT command causes the processor to "step" the specified number of macro instructions. If no count is specified, "single-step" is assumed. The console does not however enter "Spacebar Step Mode" as described in the *VAX Architecture Reference Manual*, but rather returns to the console prompt.

The console uses the trace and trace pending bits in the PSL, and the SCB trace pending vector to implement the NEXT function. This creates the following restrictions on the usage of the NEXT command:

- If memory management is enabled, the NEXT command works if and only if the first page in SSC RAM is mapped somewhere in S0 (system) space.
- The NEXT command, due to the instructions executed in implementation, does not work where time critical code is being executed.
- The NEXT command elevates the IPL to 31 for long periods of time (milliseconds) while single stepping over several commands.
- Unpredictable results occur if the macro instruction being stepped over modifies the SCBB, or the trace trap entry. This means that the NEXT command cannot be used in conjunction with other debuggers. This also implies that the user should validate PR\$_SCCB before using the NEXT command.

Qualifiers :

None

Arguments :

- **{count}** — A value representing the number of macro instructions to execute.

6-36 Maintenance Mode Firmware

Examples :

```
>>>dep 1000 50D650D4          ! Create a simple program.
>>>dep 1004 125005D1
>>>dep 1008 00FE11F9
>>>e /instruction /n:5 1000   ! List it.
    P 00001000  D4 CLRL      R0
    P 00001002  D6 INCL      R0
    P 00001004  D1 CMPL      S^#05,R0
    P 00001007  12 BNEQ      00001002
    P 00001009  11 BRB       00001009
    P 0000100B  00 HALT
>>>dep pr$_scbb 200          ! Set up a user SCBB...
>>>dep pc 1000              ! ...and the PC.
>>>
>>>n                          ! Single...
    P 00001002  D6 INCL      R0
>>>n
    P 00001004  D1 CMPL      S^#05,R0
>>>n
    P 00001007  12 BNEQ      00001002
>>>n
    P 00001002  D6 INCL      R0
>>>n 5                          ! ...or multiple step the program.
    P 00001004  D1 CMPL      S^#05,R0
    P 00001007  12 BNEQ      00001002
    P 00001002  D6 INCL      R0
    P 00001004  D1 CMPL      S^#05,R0
    P 00001007  12 BNEQ      00001002
>>>n 7
    P 00001002  D6 INCL      R0
    P 00001004  D1 CMPL      S^#05,R0
    P 00001007  12 BNEQ      00001002
    P 00001002  D6 INCL      R0
    P 00001004  D1 CMPL      S^#05,R0
    P 00001007  12 BNEQ      00001002
    P 00001009  11 BRB       00001009
>>>n
    P 00001009  11 BRB       00001009
>>>
```

6.2.7.13 REPEAT

Format :

REPEAT {command}

Description :

The console repeatedly displays and executes the specified command. Press **[Ctrl] [C]** to stop the repeating. Any valid console command can be specified for the command with the exception of the REPEAT command.

Qualifiers :

None

Arguments :

- **{command}** — A valid console command other than REPEAT.

Examples :

```
>>>repeat e pr$_todr                                ! Watch the clock.
I 0000001B 5AFE78CE
I 0000001B 5AFE78D1
I 0000001B 5AFE78FD
I 0000001B 5AFE7900
I 0000001B 5AFE7903
I 0000001B 5AFE7907
I 0000001B 5AFE790A
I 0000001B 5AFE790D
I 0000001B 5AFE7910
I 0000001B 5AFE793C
I 0000001B 5AFE793F
I 0000001B 5AFE7942
I 0000001B 5AFE7946
I 0000001B 5AFE7949
I 0000001B 5AFE794C
I 0000001B 5AFE794F
I 0000001B 5^C
>>>
```

6.2.7.14 SEARCH*Format :*

SEARCH [qualifier_list] {address} {pattern} [{mask}]

Description :

The search command finds all occurrences of a pattern, and reports the addresses where the pattern was found. If the /NOT qualifier is present, all addresses where the pattern did not match are reported.

The command accepts an optional mask that indicates don't care bits. For example, to ignore bit 0 in the comparison, specify a mask of 1. The mask, if not present, defaults to 0.

Conceptually, a match condition occurs if the following condition is true:

```
(pattern AND NOT mask) EQUALS (data AND NOT mask)
```

where:

pattern -- is the target data.

mask -- is the optional don't care bitmask (which defaults to 0).

data -- is the data (byte, word, long, quad) at the current address.

The command reports the address if the match condition is true, and there is no /NOT qualifier, or if the match condition is false and there is a /NOT qualifier. Stating this in a tabular form:

/NOT Qualifier	Match Condition	Action
-----	-----	-----
absent	true	report address
absent	false	no report
present	true	no report
present	false	report address

The address is advanced by the size of the pattern (byte, word, long or quad), unless overridden by the /STEP qualifier.

Qualifiers :

- **/B** — The data size is byte.
- **/W** — The data size is word.
- **/L** — The data size is longword.
- **/Q** — The data size is quadword.
- **/P** — The address space is physical memory. Note that when virtual memory is examined, the address space and address in the response are the translated physical address.

- **/V** — The address space is virtual memory. All access and protection checking occur. If the access would not be allowed to a program running with the current PSL, the console issues an error message. If memory mapping is not enabled, virtual addresses are equal to physical addresses.
- **/U** — Access to console private memory is allowed. This qualifier also disables virtual address protection checks. This qualifier is not inherited, and must be respecified with each command.
- **/N:{count}** — The address is the first of a range. The first access is to the address specified, then subsequent accesses are made to succeeding addresses. Even if the address is the symbolic address (-), the succeeding addresses are at larger addresses. The symbolic address specifies only the starting address, not the direction of succession.
- **/STEP:{size}** — The number to add to the current address. Normally this defaults to the data size, but is overridden by the presence of this qualifier. This qualifier is not inherited.
- **/WRONG** — ECC errors on read accesses to main memory are ignored.
- **/NOT** — Inverts the sense of the match.

Arguments :

- **{start_address}** — A longword address that specifies the first location subject to the search. This address can be any legal address specifier as defined in Section 6.2.5 and Table 6-4. If no address is specified, + is assumed.
- **{pattern}** — The target data.
- **[{mask}]** — A longword containing the bits in the target which are to be masked out.

Examples :

```

>>>dep /p/l/n:1000 0 0           ! Clear some memory.
>>>
>>>dep 300 12345678             ! Deposit some "search" data.
>>>dep 401 12345678
>>>dep 502 87654321
>>>
>>>search /n:1000 /st:1 0 12345678 ! Search for all occurrences...
P 00000300 12345678             ! ...of 12345678 on any byte...
P 00000401 12345678             ! ...boundary.
>>>search /n:1000 0 12345678     ! Then try on longword...
P 00000300 12345678             ! ...boundaries.
>>>search /n:1000 /not 0 0       ! Search for all non-zero...
P 00000300 12345678             ! ...longwords.
P 00000400 34567800
P 00000404 00000012
P 00000500 43210000
P 00000504 00008765
>>>search /n:1000 /st:1 0 1 FFFFFFFE ! Search for "odd" longwords...
P 00000502 87654321             ! ...on any boundary.
P 00000503 00876543
P 00000504 00008765
P 00000505 00000087
>>>search /n:1000 /b 0 12        ! Search for all occurrences...
P 00000303 12                   ! ...of the byte 12.
P 00000404 12
>>>search /n:1000 /st:1 /w 0 FE11 ! Search for all words which...
>>>                               ! ...could be interpreted as...
>>>                               ! ...a "spin" (10$: brb 10$).
>>>                               ! Note, none found.

```

6.2.7.15 SET*Format :*

SET {parameter} {value}

Description :

Sets the indicated console parameter to the indicated value. The following are console parameters and their acceptable values:

Parameters :

- **BFL(A)G** — Set the default R5 boot flags. The value must be a hexadecimal number of up to 8 hex digits.
- **BOOT** — Set the default boot device. The value must be a valid device name as specified in Section 6.2.7.1 on the BOOT command.

- **HOST** — Invoke the DUP or MAINTENANCE driver on the selected node. Only SET HOST /DUP accepts a value parameter. The hierarchy of the SET HOST qualifiers listed below suggests the appropriate usage. Each qualifier only supports the additional qualifiers at levels below it.

/DUP — Use the DUP protocol to examine/modify parameters of a device on either the DSSI bus or the Q22-bus. The optional value for SET HOST /DUP is a task name for the selected DUP driver to execute.

NOTE

The KN210 DUP driver only supports SEND DATA IMMEDIATE messages, and hence those devices which also support them.

/DSSI — Select the DSSI node, where node is a number from 0 to 7.

/UQSSP — Select the Q22-bus device using one of the following three methods.

/DISK n — Specify the disk controller number, where **n** is from 0 to 255. (The resulting fixed address for **n=0** is 20001468 and the floating rank for **n>0** is 26.)

/TAPE n — Specify the tape controller number, where **n** is from 0 to 255. (The resulting fixed address for **n=0** is 20001940 and the floating rank for **n>0** is 30.)

csr_address — Specify the Q22-bus I/O page CSR address for the device.

/MAINTENANCE — Use the MAINTENANCE protocol to examine/modify KFQSA EEPROM configuration parameters. Note that SET HOST /MAINTENANCE does not accept a task value.

/UQSSP —

/SERVICE n — Specify the KFQSA controller number **n** of a KFQSA in service mode, where **n** is from 0 to 3. (The resulting fixed address of a KFQSA in service mode is 20001910+4***n**.)

csr_address — Specify the Q22-bus I/O page CSR address for the KFQSA.

- **LANGUAGE** — Set console language and keyboard type. If the current console terminal does not support the Digital multinational character set (MCS), then this command has no effect and the console remains in English message mode. Acceptable values are 1 through 15 and have the following meaning:

- 1) Dansk
- 2) Deutsch (Deutschland/Österreich)
- 3) Deutsch (Schweiz)
- 4) English (United Kingdom)
- 5) English (United States/Canada)
- 6) Español
- 7) Français (Canada)
- 8) Français (France/Belgique)
- 9) Français (Suisse)
- 10) Italiano
- 11) Nederlands
- 12) Norsk
- 13) Português
- 14) Suomi
- 15) Svenska

Qualifiers :

On a per parameter basis

Arguments :

None

Examples :

```
>>>
>>>set bflag 220
>>>
>>>set boot dua0
>>>
>>>set host / dup / dssi 0
Starting DUP server...
```

```

DSSI Node 0 (SUSAN)
Copyright 1988 Digital Equipment Corporation
DRVEXR V1.0 D 5-Jul-1988 15:33:06
DRVST V1.0 D 5-Jul-1988 15:33:06
HISTRY V1.0 D 5-Jul-1988 15:33:06
ERASE V1.0 D 5-Jul-1988 15:33:06
PARAMS V1.0 D 5-Jul-1988 15:33:06
DIRECT V1.0 D 5-Jul-1988 15:33:06
End of directory

```

```

Task Name? params
Copyright 1988 Digital Equipment Corporation

```

```
PARAMS> stat path
```

ID	Path	Block	Remote Node	DGS_S	DGS_R	MSG_S_S	MSG_S_R
0	PB	FF811ECC	Internal Path	0	0	0	0
6	PB	FF811FD0	KFQSA KFX V1.0	0	0	0	0
1	PB	FF8120D4	KAREN RFX V101	0	0	0	0
4	PB	FF8121D8	WILMA RFX V101	0	0	0	0
5	PB	FF8122DC	BETTY RFX V101	0	0	0	0
2	PB	FF8123E0	DSSI1 VMS V5.0	0	0	14328	14328
3	PB	FF8124E4	3 VMB BOOT	0	0	61	61

```
PARAMS> exit
Exiting...
```

```
Task Name?
```

```

Stopping DUP server...
>>>
>>>set host /dup/dssi 0 params
Starting DUP server...

```

```

DSSI Node 0 (SUSAN)
Copyright 1988 Digital Equipment Corporation

```

```
PARAMS> show node
```

Parameter	Current	Default	Type	Radix
NODENAME	SUSAN	RF30	String	Ascii B

```
PARAMS> show allclass
```

Parameter	Current	Default	Type	Radix
ALLCLASS	1	0	Byte	Dec B

```
PARAMS> exit
Exiting...
```

6-44 Maintenance Mode Firmware

```
Stopping DUP server...
>>>
>>>set host /maint/ugssp 20001468
UQSSP Controller (772150)

Enter SET, CLEAR, SHOW, HELP, EXIT, or QUIT
Node  CSR Address  Model
0      772150      21
1      760334      21
4      760340      21
5      760344      21
7      ----- KFQSA -----
? help
Commands:
    SET <node> /KFQSA           set KFQSA DSSI node number
    SET <NODE> <CSR_address> <model> enable a DSSI device
    CLEAR <NODE>                disable a DSSI device
    SHOW                        show current configuration
    HELP                        print this text
    EXIT                        program the KFQSA
    QUIT                        don't program the KFQSA

Parameters:
    <NODE>                       0 to 7
    <CSR_address>                 760010 to 777774
    <model>                       21 (disk) or 22 (tape)
? set 6 /kfqsa
? show
Node  CSR Address  Model
0      772150      21
1      760334      21
4      760340      21
5      760344      21
6      ----- KFQSA -----
? exit
Programming the KFQSA...
>>>
>>>set language 5
>>>
```

6.2.7.16 SHOW

Format :

SHOW {parameter}

Description :

Displays the console parameter indicated.

Parameters :

- **BFL(A)G** — Show the default R5 boot flags.
- **BOOT** — Show the default boot device.
- **DEVICE** — Show a list of all devices in the system.
- **ETHERNET** — Show the hardware Ethernet address for all Ethernet adapters that can be found. Displays as blank, if no Ethernet adapter is present.
- **DSSI** — Show the status of all nodes that can be found on the DSSI bus. For each node on the DSSI bus, the console displays the node number, the node name, and the boot name and type of the device, if available. The command does not include the bootability of the device. The node that issues the command reports a node name of "*". The device information is obtained from the media type field of the MSCP command GET UNIT STATUS. In the case where the node is not running, or is not capable of running an MSCP server, then no device information is displayed.
- **LANGUAGE** — Show the console language and keyboard type. Refer to the corresponding SET LANGUAGE command for the meaning.
- **MEMORY** — Show main memory configuration on a board by board basis. Also report the addresses of bad pages, as defined by the bitmap.
 - **/FULL** Additionally show the normally inaccessible areas of memory, such as, the PFN bitmap pages, the console scratch memory pages, and the Q22-bus scatter/gather map pages.
- **QBUS** — Show all Q22-bus I/O addresses that respond to an aligned word read. For each address, the console displays the address in the VAX I/O space in hex, the address as it would appear in the Q22-bus I/O space in octal, and the word data that was read in hex. This command may take several minutes to complete, so the user may want to issue a **Ctrl C** to terminate the command. The command disables the scatter/gather map for the duration of the command.
- **UQSSP** — Show the status of all disks and tapes that can be found on the Q22-bus which support the UQSSP protocol. For each such disk or tape on the Q22-bus, the console displays the controller number, the controller CSR address, and the boot name and type of each device connected to the controller. The command does not indicate the bootability of the device.

The device information is obtained from the media type field of the MSCP command GET UNIT STATUS. In the case where the node is not running, or is not capable of running, an MSCP server, then no device information is displayed.

- **VERSION** — Show the current version of the firmware.

Qualifiers :

On a per parameter basis

Arguments :

None

Examples :

```
>>>
>>> show bflag
00000220
>>> show boot
DUA0
>>>
>>>show device
DSSI Node 0 (SUSAN)
-DIA0 -rf(0,0,*) (RF71)

DSSI Node 1 (KAREN)
-DIA1 -rf(1,1,*) (RF71)

DSSI Node 7 (*)

UQSSP Tape Controller 0 (772150)
-MUA0 -tm(0,0) (TK70)

Ethernet Adapter
-ESA0 -se -mop() (08-00-2B-0C-C4-75)
>>>
>>> show dssi
DSSI Node 0 (SUSAN)
-DIA0 -rf(0,0,*) (RF71)

DSSI Node 1 (KAREN)
-DIA1 -rf(1,1,*) (RF71)
```



```
DSSI Node 7 (*)
>>>
>>>show ethernet
Ethernet Adapter
-ESAO -se -mop() (08-00-2B-0C-C4-75)
>>>
>>>show language
English (United States/Canada)
>>>
>>>show memory
Memory 0: 00000000 to 003FFFFFF, 4MB, 0 bad pages

Total of 4MB, 0 bad pages, 98 reserved pages
>>>
>>>show memory/full
Memory 0: 00000000 to 003FFFFFF, 4MB, 0 bad pages

Total of 4MB, 0 bad pages, 98 reserved pages

Memory Bitmap
-003F8000 to 003FFFFFF, 2 pages

Console Scratch Area
-003F4000 to 003F7FFF, 32 pages

Qbus Map
-003F8000 to 003FFFFFF, 64 pages

Scan of Bad Pages
>>>
>>>show qbus
Scan of Qbus I/O Space
-20001468 (772150) = 4000 (154) RQDX3/KDA50/RRD50/RQC25/X-DISK
-2000146A (772152) = 0B40
-20001940 (774500) = 0000 (260) TQK50/TQK70/TU81E/RV20/X-TAPE
-20001942 (774502) = 0BC0
-20001F40 (777500) = 0020 (004) IPCR

Scan of Qbus Memory Space
>>>
>>>show uqssp
UQSSP Tape Controller 0 (772150)
-MUA0 -tm(0,0) (TK70)
>>>
>>>show version
KN210--A Vn.n
>>>
```

6.2.7.17 START

Format :

START [{address}]

Description :

The console starts instruction execution at the specified address. If no address is given, the current PC is used. If memory mapping is enabled, macro instructions are executed from virtual memory, and the address is treated as a virtual address. The START command is equivalent to a DEPOSIT to PC, followed by a CONTINUE. No INITIALIZE is performed.

Qualifiers :

None

Arguments :

- **[{address}]** — The address at which to begin execution. This is loaded in the user's PC.

Examples :

```
>>>start 1000
```

6.2.7.18 TEST

Format :

TEST [{test_number} [{test_arguments}]]

Description :

The console invokes a diagnostic test program specified by the test number. If a test number of 0 is specified, the power-up script is executed. The console accepts an optional list of up to five additional hexadecimal arguments.

A more detailed explanation of the diagnostics may be found in Section 6.4.

Qualifiers :

None

Arguments :

- **{test_number}** — A two digit hexadecimal number specifying the test to be executed.
- **{test_arguments}** — Up to five additional test arguments. These arguments are accepted but no meaning is attached to them by the console. For the interpretation of these arguments, consult the test specification for each individual test.

Examples :

```
>>>                ! Execute the power-up diagnostic script
>>>                ! Warning...this has the same affect as a power-up!
>>>t 0

49..48..47..46..45..44..43..42..41..40..39..38..37..36..35..34..
33..32..31..30..29..28..27..26..25..24..23..22..21..20..19..18..
17..16..15..14..13..12..11..10..09..08..07..06..05..04..03..
```

```
>>>
>>>                ! List all of the diagnostic tests.
>>>
>>>t 9e
```

Test #	Address	Name	Parameters
	2004C000	De_SCB	
	2004B200	CP_SCB	
C6	2004D6DC	SSC_powerup	*****
C7	2004D79E	CBTCR timeout	***
34	2004D858	ROM logic test	*
33	2004D920	CMCTL_powerup	*
32	2004D968	CMCTL regs	MEMCSR0_addr *****
91	2004DA8C	CQBIC_powerup	**
90	2004DB1C	CQBIC regs	*
80	2004DB9E	CQBIC_memory	*****
60	2004E1C5	Console serial	start_baud end_baud *****
52	2004E512	Prog timer	which_timer wait_time_us ***
53	2004E7D8	TOY clock	repeat_test_250ms_ea Tolerance ***
55	2004E983	Interval timer	*
5A	2004E9F8	VAX CMCTL CDAL	dont_report_memory_bad repeat_count *
45	2004EB08	cache_mem_cqbic	start_addr end_addr addr_incr ****
46	2004EDF0	Cache1_diag_md	addr_incr *****
9E	2004F41A	List diags	*
81	2004F440	MSCP-QBUS test	IP_csr *****
82	2004F602	DELQA	device_num_addr ****
C1	2004F7DD	SSC RAM	*
C2	2004F9A4	SSC RAM ALL	*
C5	2004FB20	SSC regs	*

6-50 Maintenance Mode Firmware

```

56 2004FC14 SII_ext_loopbck ***
5C 2004FF19 SII_initiator run_test *****
5D 20050C0C SII_target run_test *****
58 200523D2 DSSI_reset port_no time_secs *
57 200527B8 SII_memory incr test_pattern run_test *****
5E 200527CD NI_memory incr test_pattern run_test *****
5B 20052BF6 SII_registers run_test ****
5F 20052DC4 NI_Test do_extl where run_test *****
54 20053967 Virtual_mode *****
41 20053C3C Board_Reset ***
42 20053DDD Check_for_intrs ***
31 20053E20 MEM_Setup_CSRS *****
30 2005434F MEM_Bitmap *** mark_Hard_SBEs *****
4F 2005478A MEM_Data start_add end_add add_incr cont_on_err *****
4E 20054948 MEM_Byte start_add end_add add_incr cont_on_err *****
4D 20054A5D MEM_Address start_add end_add add_incr cont_on_err *****
4C 20054BEF MEM_ECC_Error start_add end_add add_incr cont_on_err *****
4B 200550B0 MEM_Maskd_Errs start_add end_add add_incr cont_on_err *****
4A 2005528A MEM_Correction start_add end_add add_incr cont_on_err *****
49 2005549D MEM_FDM_Logic *** cont_on_err *****
48 20055A6C MEM_Addr_shrts start_add end_add * cont_on_err pat2 pat3 **
47 200560C3 MEM_Refresh start end incr cont_on_err time_seconds ****
40 2005624E MEM_Count_Errs First_board Last_board Soft_errs_allowed ***
44 20056568 Cache_memory addr_incr *****
9D 200568B4 Utilities Expnd_err_msg get_mode init_LEDs clr_ps_cnt
9C 200569B8 List_CPU_regs *
9F 200571CD Create_script *****
70 200578F0 R3000_cache *
71 20057907 R3000_fpu *
72 2005791C R3000_tlb *
73 20057931 R3000_reg_intrf loop run_test ***
74 2005794B R3000_buf_intrf incr pattern start_add end_add run_test **
75 20057965 R3000_mem_intrf incr pattern start_add end_add **
76 2005797F R3000_interrupt run_test
77 20057999 R3000_mov_inver incr * start_add end_add **

>>>
>>> ! Show the diagnostic state.
>>>
>>>t fe

```

Maintenance Mode Firmware 6-51

```
bitmap=00BF3400, length=0C00, checksum=007E
busmap=00BF8000
return_stack=201406A4
subtest_pc=2004EBB0
timeout=00000001, error=00, de_error=00
de_error_vector=00, severity_code=02, total_error_count=0000
previous_error=00000000, 00000000, 00000000, 00000000, 00000000
last_exception_pc=2004EBDA
flags=21FFFD7F, test_flags=20
highest_severity=00
led_display=06
console_display=00
save_mchk_code=80, save_err_flags=000000
parameter_1=00000000 2=00000000 3=00000000 4=00000000 5=00000000
parameter_6=00000001 7=00000000 8=2004EBE0 9=00000000 10=20056056
```

>>>

>>> ! Display the CPU registers.

>>>

>>>t 9c

```
TOY =76BA1D75 ICCS =00000000
TCR0 =00000000 TIR0 =00000000 TNIR0=00000000 TIVR0=00000078
TCR1 =00000001 TIR1 =02BD7971 TNIR1=0000000F TIVR1=0000007C
RXCS =00000000 RXDB =0000000D TXCS =00000000 TXDB =00000030
MSER =00000000 CADR =0000000C
BDR =FFFFFFD0 DLEDR=0000000C SCCR=00D45033 CBTCR=C0000004
SCR =0000C000 DSER =00000080 QBEAR=0000000F DEAR =00000000
QBMBR=00000000 IPCRn=0020

MEM_FRU 1 MEMCSR_0=80000015 1=00000015 2=00000015 3=00000015
MEM_FRU 2 MEMCSR_4=80400016 5=80800016 6=00000016 7=00000016
MEM_FRU 3 MEMCSR_8=00000000 9=00000000 10=00000000 11=00000000
MEM_FRU 4 MEMCSR12=00000000 13=00000000 14=00000000 15=00000000
MEMCSR16=00000044 17=0000203C
```

Ethernet SA = 08-00-2B-0B-25-65 NICSR0=0004

```
SII MSIDR0 =01FF MSIDR1 =0002 MSIDR2 =0000 MSICSR =0010
MSIIDR =8007 MSITR =0000 MSITLP =0000 MSIILP =0000
MSIDSCR=80FF MSIDSSR=8500 MSIDCR =0008
```

>>>

6.2.7.19 UNJAM

Format :

UNJAM

Description :

An I/O bus reset is performed. This is implemented by writing 1 to IPR 55.

Qualifiers :

None

Arguments :

None

Examples :

```
>>>unjam  
>>>
```

6.2.7.20 X - Binary Load and Unload

Format :

**X {address} {count} <CR> {line_checksum}
{data} {data_checksum}**

Description :

The X command is for use by automatic systems communicating with the console. It is not intended for use by operators.

The console loads or unloads (that is, writes to memory, or reads from memory) the specified number of data bytes, starting at the specified address through the console serial line.

If bit 31 of the count is clear, data is to be received by the console, and deposited into memory. If bit 31 of the count is set, data is to be read from memory and sent by the console. The remaining bits in the count are a positive number indicating the number of bytes to load or unload.

The console accepts the command upon receiving the carriage return. The next byte the console receives is the command checksum, which is not echoed. The command checksum is verified by adding all command characters, including the checksum and separating whitespace, (but not including the terminating carriage return or rubouts or characters deleted by rubout), into an 8 bit register initially set to zero. If no errors occur, the result is zero. If the command checksum is correct, the console responds with the input prompt and either sends data to the requester or prepares to receive data. If the command checksum is in error, the console responds with an error message. The intent is to prevent inadvertent operator entry into a mode where the console is accepting characters from the keyboard as data, with no escape mechanism possible.

If the command is a load (bit 31 of the count is clear), the console responds with the input prompt, then accepts the specified number of bytes of data for depositing to memory, and an additional byte of received data checksum. The data is verified by adding all data characters and the checksum character into an 8 bit register initially set to zero. If the final contents of the register is non-zero, the data or checksum are in error, and the console responds with an error message.

If the command is a binary unload (bit 31 of the count is set), the console responds with the input prompt, followed by the specified number of bytes of binary data. As each byte is sent, it is added to a checksum register initially set to zero. At the end of the transmission, the 2's complement of the low byte of the register is sent.

If the data checksum is incorrect on a load, or if memory errors or line errors occur during the transmission of data, the entire transmission is completed, and then the console issues an error message. If an error occurs during loading, the contents of the memory being loaded are UNPREDICTABLE.

Echo is suppressed during the receiving of the data string and checksums.

To avoid treating flow control characters from the terminal as valid command line checksums, all flow control is terminated at the reception of the carriage return terminating the command line.

It is possible to control the console serial line through the use of the control characters (`Ctrl C`, `Ctrl S`, `Ctrl O` and so on) during a binary unload. It is not possible during a binary load, as all received characters are valid binary data.

Data being loaded with a binary load command must be received by the console at a rate of at least one byte every 60 seconds. The command checksum that precedes the data must be received by the console within 60 seconds of the carriage return that terminates the command line. The data checksum must be received within 60 seconds of the last data byte. If any of these timing requirements are not met, the console aborts the transmission by issuing an error message and prompting for input.

The entire command, including the checksum, can be sent to the console as a single burst of characters at the console serial line's specified character rate. The console is able to receive at least 4 Kbytes of data in a single X command.

Qualifiers :

None

Arguments :

None

Examples :

None

6.2.7.21 ! - Comment

Format :

!

Description :

The comment command is used to document command sequences. The comment character can appear anywhere on the command line. All characters following the comment character are ignored.

Qualifiers :

None

Arguments :

None

Examples :

```
>>>! The console ignores this line.  
>>>
```


6.2.8 Conventions for Tables 6-5 and 6-6

Table 6-5 lists a complete summary of the console commands. Table 6-6 is a summary of the qualifiers recognized by the diagnostic processor console.

- UPPERCASE denotes the command or qualifier keyword.
- {} denotes a mandatory item which must be syntactically correct.
- [] denotes an optional item.
- ! denotes an "or" condition.
- *boot_flags*, *count*, *size*, *address*, & *parameters* denote hex longword values.
- *boot_device* denotes a legal boot device name.
- *csr_address* denotes a Q22-bus I/O page CSR address.
- *controller_number* denotes a controller number from 0 to 255.
- *language_type* denotes the language value, from 1 to 15.
- *command* denotes a console command other than REPEAT.
- *data*, *pattern*, & *mask* denote hex values of the current size.
- *test_number* denotes hex byte test number.

Table 6-5 Console Command Summary

Command	Qualifiers	Argument	Other(s)
BOOT	/R5:{boot_flags} /{boot_flags}	[{boot_device}]	—
CONTINUE	—	—	—
DEPOSIT	/B /W /L /Q — /G /I /V /P /M /U /N:{count} /STEP:{size} /WRONG	{address}	{data} [{data}]
EXAMINE	/B /W /L /Q — /G /I /V /P /M /U /N:{count} /STEP:{size} /WRONG /INSTRUCTION	[{address}]	—
EXIT	—	—	—
FIND	/MEM /RPB	—	—

Table 6-5 (Cont.) Console Command Summary

Command	Qualifiers	Argument	Other(s)
HALT	—	—	—
HELP	—	—	—
INITIALIZE	—	—	—
MOVE	/B /W /L /Q — /V /P /U /N:{count} /STEP:{size} /WRONG	{src_address}	{dest_address}
NEXT	—	[[count]]	—
REPEAT	—	{command}	—
SEARCH	/B /W /L /Q — /V /P /U /N:{count} /STEP:{size} /WRONG /NOT	{start_address}	{pattern} [[mask]]
SET	—	{bitmap}	—
BFL(A)G	—	{device_string}	—
SET BOOT	—	{node_number}	[[task]]
SET HOST	/DUP /DSSI	{node_number}	[[task]]
SET HOST	/DUP /UQSSP {DISK ! /TAPE }	{controller_number}	[[task]]
SET HOST	/DUP /UQSSP /MAINTENANCE /UQSSP /SERVICE /MAINTENANCE /UQSSP	{csr_address} {controller_number} {csr_address}	—
SET LANGUAGE	—	{language_type}	—
SHOW	—	—	—
BFL(A)G	—	—	—
SHOW BOOT	—	—	—
SHOW DSSI	—	—	—
SHOW ETHERNET	—	—	—
SHOW LANGUAGE	—	—	—
SHOW MEMORY	/FULL	—	—
SHOW QBUS	—	—	—

Table 6-5 (Cont.) Console Command Summary

Command	Qualifiers	Argument	Other(s)
SHOW	—	—	—
UQSSP	—	—	—
SHOW VERSION	—	—	—
START	—	{address}	—
TEST	—	{test_number}	[{parameters}]
UNJAM	—	—	—
X	—	{address}	{count}

Table 6-6 Console Qualifier Summary**Data Control**

/B	Byte, legal for memory references only.
/W	Word, legal for memory references only.
/L	Longword, the default for GPR and IPR references.
/Q	Quadword, legal for memory references only.
/N:{count}	Specify number of additional operations.
/STEP:{size}	Override the default step incrementing size with the value specified for the current reference.
/WRONG	Ignore ECC bits on reads. Use an ECC value of 3 on writes.

Address Space Control

/G	General Purpose Registers
/I	Internal Processor Registers
/V	Virtual memory
/P	Physical memory, both VAX memory and I/O spaces
/U	Protected memory (ROMs, SSC RAM, PFN bitmap)
/M	Machine state (PSL)

Table 6-6 (Cont.) Console Qualifier Summary

Command Specific	
/INSTRUCTION	EXAMINE command only. Disassemble the instruction at address specified.
/NOT	SEARCH command only. Invert the sense of the match.
/R5:{boot_flags}, /{boot_flags}	BOOT command only. Specify a function bitmap to pass to VMB through R5. Refer to Figure 6-1 for a bit description of R5. Either form of the command is acceptable.
/RPB, /MEM	FIND command only. Search for valid RPB or good block of memory.
/DUP, /DSSI, /UQSSP, /DISK, /TAPE, /MAINTENANCE,	SET HOST command only. Refer to command description for usage.
/SERVICE	

6.3 Bootstrapping

Bootstrapping is the process of loading and transferring control to an operating system. The KN210 maintenance mode supports bootstrap MDM diagnostics and any user application image which conforms to the boot formats described herein.

On the KN210, a bootstrap occurs whenever a BOOT command is issued at the console or whenever the processor halts and the conditions specified in the Table 6-1 for automatic bootstrap are satisfied.

6.3.1 Boot Devices

The maintenance mode firmware passes the address of a descriptor of the boot device name to VMB through R0. The device name used for the bootstrap operation is one of the following:

- ESA0, if no default boot device has been specified
- The default boot device specified at initial power-up or through a SET BOOT command
- The boot device name explicitly specified in a BOOT command line

The device name may be any arbitrary character string, with a maximum length of 17 characters. Longer strings cause an error message to be issued to the console. Otherwise the console makes no attempt at interpreting or validating the device name. The console converts the string to all upper case, and passes VMB the address of a string descriptor for the device name in R0.

Table 6–7 correlates the boot device names expected in a BOOT command with the corresponding supported devices.

Table 6–7 KN210 Supported Boot Devices

Boot Name ¹	Controller Type	Device Type(s)
Disk:		
[node\$]DIAn	On-board DSSI	RF30, RF71
DUcn	RQDX3	RD52, RD53, RD54, RX33, RX50
	MSCP	
	KDA50 MSCP	RA70, RA80, RA81, RA82, RA90
	KFQSA	RF30, RF71
	MSCP	
	KLESI	RC25
Tape:		
[node\$]MIAn	On-board DSSI	TF70
MUcn	TQK50 MSCP	TK50
	TQK70 MSCP	TK70
	KFQSA	TF _{xx}
	MSCP	
	KLESI	TU81E
Network:		
ESA0	On-board Enet	—

¹ Boot device names consist of minimally a two letter device code, followed by a single character controller letter (A...Z), and terminating in a device unit number (0...65535). DSSI device names may optionally include a node prefix, consisting of either a node number (0...7) or a node name (a string of up to 8 characters), terminating in a "\$".

Table 6-7 (Cont.) KN210 Supported Boot Devices

Boot Name¹	Controller Type	Device Type(s)
Network:		
XQcn	DEQNA	—
	DELQA	—
	DESQA	—
PROM:		
PRA0	MRV11	—

¹ Boot device names consist of minimally a two letter device code, followed by a single character controller letter (A...Z), and terminating in a device unit number (0...65535). DSSI device names may optionally include a node prefix, consisting of either a node number (0...7) or a node name (a string of up to 8 characters), terminating in a "\$".

NOTE

Table 6-7 presents a definitive list of boot devices which the KN210 supports. However, the KN210 will likely boot other devices which adhere to the MSCP standards.

6.3.2 Boot Flags

The action of VMB is qualified by the value passed to it in R5. R5 contains boot flags that specify conditions of the bootstrap. The firmware passes to VMB either the R5 value specified in the BOOT command or the default boot flag value specified with a SET BFLAG command.

Figure 6-1 shows the location of the boot flags used by VMB in the boot flag longword and Table 6-8 describes each flag's function.

Figure 6-1 VMB Boot Flags**Table 6-8 VMB Boot Flags**

Field	Name	Description
0	RPB\$V_CONV	Conversational bootstrap.
3	RPB\$V_BBLOCK	Secondary bootstrap from bootblock. When this bit is set, VMB reads logical block number 0 of the boot device and tests it for conformance with the bootblock format. If in conformance, the block is executed to continue the bootstrap. No attempt to perform a Files-11 bootstrap is made.
4	RPB\$V_DIAG	Diagnostic bootstrap. When this bit is set, the load image requested over the network is [SYS0.SYSMAINT]DIAGBOOT.EXE.
5	RPB\$V_BOOBPT	Bootstrap breakpoint. If this flag is set, a breakpoint instruction is executed in VMB and control is transferred to XDELTA prior to boot.
6	RPB\$V_HEADER	Image header. If this bit is set, VMB transfers control to the address specified by the file's image header. If this bit is not set, VMB transfers control to the first location of the load image.
8	RPB\$V_SOLICT	File name solicit. When this bit is set, VMB prompts the operator for the name of the application image file. A maximum of a 39 character file specification is permitted.

Table 6-8 (Cont.) VMB Boot Flags

Field	Name	Description
9	RPB\$V_HALT	Halt before transfer. When this bit is set, VMB halts before transferring control to the application image.
31:28	RPB\$V_TOPSYS	This field can be any value from 0 through F. This flag changes the top level directory name for the system disks with multiple operating systems. For example, if TOPSYS is 1, the top level directory name is [SYS1...].

6.3.3 Preparing for the Bootstrap

Prior to dispatching to the primary bootstrap (VMB), the firmware initializes the system to a known state. The initialization sequence is the following:

1. Check CPMBX<2>(BIP). If it is set, bootstrap fails.
2. If this is an automatic bootstrap, print the message "Loading system software" on the console terminal.
3. Validate the boot device name. If none exists, supply a list of available devices and prompt user for a device. If no device is entered within 30 seconds, use ESA0.
4. Write a form of this BOOT request including the active boot flags and boot device on the console, for example "(BOOT/R5:0 DUA0)".
5. Set CPMBX<2>(BIP).
6. Initialize the Q22-bus scatter/gather map.
 - a. Clear IPCR<5>(LMEAE).
 - b. Perform an UNJAM.
 - c. Map all vacant Q22-bus pages to the corresponding page in local memory and validate each entry if that page is good.
 - d. Perform an INIT.
 - e. Set IPCR<5>(LMEAE).
7. Validate the PFN bitmap. If invalid, rebuild it.
8. Search for a 128 Kbyte contiguous block of good memory as defined by the PFN bitmap. If 128 Kbytes cannot be found, the bootstrap fails.

9. Initialize the general purpose registers.

R0 = address of descriptor of the boot device name or 0 if none specified
R2 = length of PFN bitmap in bytes
R3 = address of PFN bitmap
R4 = time of day from PR\$_TODR at power-up
R5 = boot flags
R10 = halt PC value
R11 = halt PSL value (without halt code and map enable)
AP = halt code
SP = base of 128 Kbyte good memory block + 512
PC = base of 128 Kbyte good memory block + 512
R1, R6, R7, R8, R9, FP = 0

10. Copy the VMB image from EPROM to local memory beginning at the base of the 128 Kbyte good memory block + 512.
11. Exit from the firmware to memory resident VMB.

On entry to VMB the processor is running at IPL 31 on the interrupt stack with memory management disabled.

6.3.4 Primary Bootstrap, VMB

Virtual memory boot (VMB) is the primary bootstrap for booting VAX processors. On the KN210, VMB is resident in the firmware and is copied into main memory before control is transferred to it. VMB then loads the secondary bootstrap image and transfers control to it.

NOTE

In certain cases, VMB actually loads the operating system directly. However, for the purpose of this discussion "secondary bootstrap" refers to any VMB loadable image.

VMB inherits a well defined environment and is responsible for further initialization. The following summarizes the operation of VMB:

1. Initialize a two page SCB on the first page boundary above VMB.
2. Allocate a three page stack above the SCB.
3. Initialize the restart parameter block (RPB).
4. Initialize the secondary bootstrap argument list.
5. If not a PROM boot, locate a minimum of 3 consecutive valid QMRs.

6. Write "2" to the diagnostic LEDs and display "2.." on the console to indicate that VMB is searching for the device.
7. Optionally, solicit from the console a "Bootfile: " name.
8. Write the name of the boot device from which VMB will attempt to boot on the console, "-DUA0".
9. Copy the secondary bootstrap from the boot device into local memory above the stack. If this fails, the bootstrap fails.
10. Clear CPMBX<2>(BIP) and CPMBX<3>(RIP).
11. Write "0" to the diagnostic LEDs and display "0.." on the console to indicate that VMB is now transferring control to the loaded image.
12. Transfer control to the loaded image with the following register usage:
 - R5** = transfer address in secondary bootstrap image
 - R10** = base address of secondary bootstrap memory
 - R11** = base address of RPB
 - AP** = base address of secondary boot parameter block
 - SP** = current stack pointer

If the bootstrap operation fails, VMB relinquishes control to the console by halting with a HALT instruction.

NOTE

VMB makes no assumptions about the location of Q22-bus memory. However, VMB searches through the Q22-bus map registers (QMRs) for the first QMR marked *valid*. VMB requires minimally 3 and maximally 129 contiguous *valid* maps to complete a bootstrap operation. If the search exhausts all map registers or there are fewer than the required number of *valid* maps, a bootstrap cannot be performed. It is recommended that a suitable block of Q22-bus memory address space be available (unmapped to other devices) for proper operation.

The following is a sample console display of a successful automatic bootstrap:

```
Loading system software.  
(BOOT/R5:0 DUA0)
```

```
2..  
-DUA0  
1..0..
```

After a successful bootstrap operation, control is passed to the secondary bootstrap image with the memory layout as shown in Figure 6-2.

Figure 6-2 Memory Layout at VMB Exit

In the event that an operating system has an extraordinarily large secondary bootstrap which overflows the 128 Kbytes of good memory, VMB loads the remainder of the image in memory above the good block. However, if there are not enough contiguous good pages above the block to load the remainder of the image, the bootstrap fails.

6.3.5 Device Dependent Bootstrap Procedures

As mentioned earlier, the KN210 supports bootstrapping from a variety of boot devices. The following sections describe the various device dependent boot procedures.

6.3.5.1 Disk and Tape Bootstrap Procedure

The disk and tape bootstrap supports Files-11 lookup (supporting only the ODS level 2 file structure) or the boot block mechanism (used in PROM boot also). Of the standard DEC operating systems VMS and ELN use the Files-11 bootstrap procedure and Ultrix-32 uses the boot block mechanism.

VMB first attempts a Files-11 lookup, unless the RPB\$V_BBLOCK boot flag is set. If VMB determines that the designated boot disk is a Files-11 volume, it searches the volume for the designated boot program, usually [SYS0.SYSEXEXE]SYSBOOT.EXE. However, VMB can request a diagnostic image or prompt the user for an alternate file specification (Section 6.3.2). If the boot image cannot be found, VMB fails.

If the volume is not a Files-11 volume or the RPB\$V_BBLOCK boot flag was set, the boot block mechanism proceeds as follows:

1. Read logical block 0 of the selected boot device (this is the boot block).
2. Validate that the contents of the boot block conform to the boot block format (Figure 6-3).
3. Use the boot block to find and read in the secondary bootstrap.
4. Transfer control to the secondary bootstrap image, just as for a Files-11 boot.

The format of the boot block must conform to that shown in Figure 6-3.

Figure 6-3 Boot Block Format

6.3.5.2 PROM Bootstrap Procedure

The PROM bootstrap uses a variant of the boot block mechanism. VMB searches through Q22-bus memory on 16 Kbyte boundaries for a valid PROM *signature block*, the second segment of the boot block defined in Figure 6-3.

At each boundary, VMB :

1. Validates the readability of that Q22-bus memory page.
2. If readable, check to see if it contains a valid PROM signature block.

If verification passes, the PROM image will be copied into main memory and VMB will transfer control to that image at the offset specified in the PROM bootblock. If not, the next page will be tested.

NOTE

It is not necessary that the boot image actually reside in PROM. Any boot image in Q22-bus memory space with a valid signature block on a 16 Kbyte boundary is a candidate.

The PROM image is copied into main memory in 127 page "chunks" until the entire PROM is moved. All destination pages beyond the primary 128 Kbyte block are verified to make sure they are marked good in the PFN bitmap. The PROM must be copied contiguously and if all required pages cannot fit into the memory immediately following the VMB image, the boot fails.

6.3.5.3 Network Bootstrap Procedure

Whenever a network bootstrap is selected on a KN210, VMB makes continuous attempts to boot from the network. VMB uses the DNA maintenance operations protocol (MOP) as the transport protocol for network bootstraps and other network operations. Once a network boot has been invoked, VMB turns on the designated network link and repeats load attempts, until either a successful boot occurs or it is halted from the operator console.

The KN210 supports the load of a standard operating system, a diagnostic image, or a user designated program through network bootstraps. The default image is a standard operating system. However, a user may select an alternate image by setting either the RPB\$V_DIAG bit or the RPB\$V_SOLICIT bit in the boot flag longword R5. Note, that the RPB\$V_SOLICIT bit has precedence over the RPB\$V_DIAG bit. Hence, if both bits are set, then the solicited file is requested. (Refer to Figure 6-1 for the usage of these bits.)

NOTE

VMB accepts a maximum of a 39 character file specification for solicited boots. If the network server is running VMS, the following defaults apply to the file specification: the directory MOM\$LOAD:, and an extension .SYS. Therefore, the 39 character file specification need only consist of the filename if the default directory and extension attributes are used.

The KN210 VMB uses the MOP program load sequence for bootstrapping the module and the MOP "dump/load" protocol type for load related message exchanges. The MOP message types used in the exchange are listed in Table 6-9.

VMB, the requester, starts by sending a REQ_PROGRAM message in the appropriate envelope to the MOP "dump/load" multicast address. It then waits for a response in the form of a VOLUNTEER message from another node on the network, the MOP load server. If a response is received, then the destination address is changed from the multicast address to the node address of the server. The same REQ_PROGRAM message is retransmitted to the server as an acknowledge which initiates the load.

Next, VMB begins sending REQ_MEM_LOAD messages in response to either:

- MEM_LOAD message, while there is still more to load
- MEM_LOAD_w_XFER, if it is the end of the image
- PARAM_LOAD_w_XFER, if it is the end of the image and operating system parameters are required

The "load number" field in the load messages is used to synchronize the load sequence. At the beginning of the exchange, both the requester and server initialize the load number. The requester only increments the load number if a load packet has been successfully received and loaded. This forms the acknowledge to each exchange. The server will resend a packet with a specific load number, until it sees a request with the load number incremented. The final acknowledge is sent by the requester and has a load number equivalent to the load number of the appropriate LOAD_w_XFER message + 1.

During the boot sequence if no response is made to the REQ_PROGRAM message in the current time-out limit, the time-out limit is increased linearly (by 4 seconds) up to a maximum of about 4 minutes. The initial time-out limit is 8 seconds.

6.3.5.4 Network Listening

While the KN210 is waiting for a load volunteer during bootstrap, it listens on the network for other maintenance messages directed to the node and periodically identifies itself at 8 to 12 minute intervals. In particular, this listener supplements the MOP functions of the VMB load requester typically found in bootstrap firmware and supports the following:

- A remote console server that generates unsolicited SYSTEM_ID messages every 8 to 12 minutes and solicited SYSTEM_ID messages in response to REQUEST_ID messages, as well as, COUNTERS messages in response to REQ_COUNTERS messages.

6-70 Maintenance Mode Firmware

- A loopback server that responds to Ethernet LOOPBACK messages by echoing the message to the requester.
- An IEEE 802.2 responder which replies to both XID and TEST messages.

During network operation the firmware listens only to MOP "Load/Dump," MOP "Remote Console," and Ethernet "Loopback Assistance" messages protocols directed to the Ethernet physical address of the node. All other Ethernet protocols are filtered by the network device driver. Additionally, IEEE 802.3 messages are also processed by the network listener.

The MOP functions and message types which are supported by the KN210 are summarized in the following tables.

Table 6-9 KN210 Network Maintenance Operations Summary

Function	Role	Transmit		Receive
MOP Ethernet and IEEE 802.3 Messages ¹				
Dump	Requester	—		—
	Server	—		—
Load	Requester	REQ_	to solicit	VOLUNTEER
		PROGRAM ²		
		REQ_MEM_	to solicit &	MEM_LOAD
	LOAD	ACK	or	MEM_LOAD_w_
			or	PARAM_LOAD_
				w_XFER
	Server	—		—
Console	Requester	—		—
	Server	SYSTEM_ID ³	in response to	REQUEST_ID

¹ All unsolicited messages are sent in Ethernet (MOP V3) and IEEE 802.2 (MOP V4), until the MOP version of the server is known. All solicited messages are sent in the format used for the request.

²The initial REQ_PROGRAM message is sent to the dumpload multicast address. If an assistance VOLUNTEER message is received, then the responder's address is used as the destination to repeat the REQ_PROGRAM message and for all subsequent REQ_MEM_LOAD messages.

³SYSTEM_ID messages are sent out every 8 to 12 minutes to the remote console multicast address and on receipt of a REQUEST_ID message they are sent to the initiator.

Table 6-9 (Cont.) KN210 Network Maintenance Operations Summary

Function	Role	Transmit		Receive
MOP Ethernet and IEEE 802.3 Messages ¹				
		COUNTERS	in response to	REQ_ COUNTERS BOOT
Loopback	Requester Server	— LOOPED_ DATA ⁴	in response to	— LOOP_DATA
IEEE 802.2 Messages⁵				
Exchange ID	Requester Server	— XID_RSP	in response to	— XID_CMD
Test	Requester Server	— TEST_RSP	in response to	— TEST_CMD

¹ All unsolicited messages are sent in Ethernet (MOP V3) and IEEE 802.2 (MOP V4), until the MOP version of the server is known. All solicited messages are sent in the format used for the request.

⁴LOOPED_DATA messages are sent out in response to LOOP_DATA messages. These messages are actually in Ethernet LOOP TEST format, not in MOP format, and when sent in Ethernet frames omit the additional length field (padding is disabled).

⁵IEEE 802.2 support of XID and TEST is limited to Class 1 operations.

6.4 Diagnostics

The ROM based diagnostics constitute a major portion of the firmware on the KN210 and consist of an initial power-up test and a series of functional component diagnostics. These diagnostics run automatically on power-up and when in *maintenance* mode. They can be executed interactively as a whole, or as individual tests using the TEST command (Section 6.2.7.18). This section summarizes their operation.

The purpose of the ROM based diagnostics is multifaceted:

1. During powerup, they determine if enough of the KN210 is working to allow the console to run.
2. During the manufacturing process, they verify that the board was correctly built.
3. In the field, they verify that the board is operational, and to report all detected errors.
4. They allow sophisticated users and Field Service technicians to run individual diagnostics interactively, with the intent of isolating errors to the FRU.

To accommodate these requirements, the diagnostics have been designed as a collection of individual parameterized tests. A data structure, called a *script*, and a program, called the *diagnostic executive*, orchestrate the running of these tests in the right order with the right parameters.

A script is a data structure that points to various tests. There are several scripts, one for the field, and several for manufacturing, depending on where on the manufacturing line the board is. Sophisticated users may also create their own scripts interactively. Additionally, the script contains other information:

- What parameters need to be passed to the test.
- What is to be displayed, if anything, on the console.
- What is to be displayed, if anything, on the LED.
- What to do on errors (halt, loop, or continue).
- Where the tests may be run from. For example, there are certain tests that can only be run from the EEROM. Other tests are position independent code (PIC), and may be run from EEROM, or main memory in the interests of execution speed.

The diagnostic executive interprets scripts to determine what tests are to be run. There are several built-in scripts on the KN210 that are used for manufacturing, power-up, and Field Service personnel. The diagnostic executive automatically invokes the correct script based on the current environment of the KN210. Any script can be explicitly run with the TEST command from the console terminal.

The diagnostic executive is also responsible for controlling the tests so that when errors occur, they can be caught and reported to the user. The executive also ensures that when the tests are run, the machine is left in a consistent and well defined state.

6.4.1 Error Reporting

Before a console is established, the only error reporting is through the KN210's diagnostic LEDs (and any LEDs on other boards). Once a console has been established, all errors detected by the diagnostics are also reported by the console. When possible, the diagnostics issue an error summary on the console. Example 6-4 shows a typical error display.

```
?34 2 08 FF 00 0000 (1)
P1=00000000 P2=00000003 P3=00000031 P4=00000011 P5=00002000 (2)
P6=FFFFFFFF P7=00000000 P8=00000000 P9=00000000 P10=2005438F (3)
r0=00114B98 r1=FFFFFFFF r2=2005D2F0 r3=55555555 r4=AAAAAAAA (4)
r5=00000000 r6=AAAAAAAA r7=00000000 r8=00000000 ERF=80000180 (5)
```

Example 6-4 Diagnostic Register Dump

The numbers in parenthesis on the right side of Example 6-4 refer to lines of the display and are not a part of the diagnostic dump. The information on these lines is summarized as follows:

1. Test summary containing six hexadecimal fields.
 - a. ?34, **test** identifies the diagnostic test.
 - b. 2, **severity** is the level of a test failure, as dictated by the script. Failure of a severity level 2 test causes the display of this five-line error printout, and halts an autoboot to console I/O mode. An error of severity level 1 displays the first line of the error printout, but does not interrupt an autoboot. Most tests have a severity level of 2.
 - c. 08, **error** is a number, that in conjunction with listing files, isolates to within a few instructions where the diagnostic detected the error. This field is also called the subtestlog.

- d. FF, **de_error** is a code with which the diagnostic executive signals the diagnostic's state and any illegal behavior. This field indicates a condition that the diagnostic expects on detecting a failure. The possible codes are:
 - FF - Normal error exit from diagnostic
 - FE - Unanticipated interrupt
 - FD - Interrupt in cleanup routine
 - FC - Interrupt in interrupt handler
 - FB - Script requirements not met
 - FA - No such diagnostic
 - EF - Unanticipated exception in executive
 - e. 00, **vector** is the SCB vector (if non-zero) through which an unexpected exception or interrupt trapped, when the **de_error** field indicates an unexpected exception or interrupt (FE or EF).
 - f. 0000, **count** is the number of previous errors that have occurred.
2. P1...P5 are the first five longwords of the diagnostic state. This is internal information that is used by repair personnel.
 3. P6...P10 are the last five longwords of the diagnostic state.
 4. R0...R4 are the first five GPRs at the moment the error was detected.
 5. R5...R8 are additional GPRs and ERF is a diagnostic summary longword.

NOTE

This information only applies to CVAX-based tests. R3000 tests (70-76) will display different registers. This information is only intended to be an example of an error display.

6.4.2 Diagnostic Interdependencies

When running tests interactively on an individual basis, users should be aware that certain tests may be dependent on some state set up from a previous test. In general, tests should not be run out of order.

6.4.3 Areas Not Covered

The goal has been to achieve the highest possible coverage on the KN210 and the memory boards. However, the testing of the KN210 while running with memory management turned on is minimal. Also, due to the way the firmware is implemented (a polled environment running at IPL 31), the testing of interrupts is not extensive.

These diagnostics are not intended to be used as system level tests. There are no tests that completely verify that access to the Q22-bus will work. Thus, a disk, a controller, the backplane, or portions of the CQBIC may be faulty, and the diagnostics may not detect the fault. Such a fault may result later on as an inability to boot.

6.5 Operating System Restart

An operating system restart is the process of bringing up the operating system from a known initialization state following a processor halt. This procedure is often called *restart* or *warmstart*, and should not be confused with a processor restart which results in firmware entry.

On the KN210 a restart occurs if the conditions specified in Table 6-1 are satisfied.

To restart a halted operating system, the firmware searches system memory for the restart parameter block (RPB), a data structure constructed for this purpose by VMB. If a valid RPB is found, the firmware passes control to the operating system at an address specified in the RPB.

The firmware keeps a *restart in progress* (RIP) flag in CPMBX which it uses to avoid repeated attempts to restart a failing operating system. An additional RIP flag is maintained by the operating system in the RPB.

The firmware uses the following algorithm to restart the operating system:

1. Check CPMBX<3>(RIP). If it is set, restart fails.
2. Print the message "Restarting system software" on the console terminal.
3. Set CPMBX<3>(RIP).
4. Search for a valid RPB. If none is found, restart fails.
5. Check the operating system RPB\$L_RSTRTFLG<0>(RIP) flag. If it is set, restart fails.
6. Write "0" on the diagnostic LEDs.

7. Dispatch to the restart address, RPB\$L_RESTART, with :

SP = the physical address of the RPB plus 512

AP = the halt code

PSL = 041F 0000

PR\$_MAPEN = 0.

If the restart is successful, the operating system must clear CPMBX<3>(RIP).

If restart fails, the firmware prints "Restart failure" on the system console.

6.5.1 Locating the RPB

The RPB is a page aligned control block which can be identified by the first three longwords. The format of the RPB *signature* is shown in Figure 6-4.

Figure 6-4 RPB Signature Format

The firmware uses the following algorithm to find a valid RPB:

1. Search for a page of memory that contains its address in the first longword. If none is found, the search for a valid RPB has failed.
2. Read the second longword in the page (the physical address of the restart routine). If it is not a valid physical address, or if it is zero, return to step 1. The check for zero is necessary to ensure that a page of zeros does not pass the test for a valid RPB.
3. Calculate the 32 bit twos-complement sum (ignoring overflows) of the first 31 longwords of the restart routine. If the sum does not match the third longword of the RPB, return to step 1.
4. A valid RPB has been found.

6.6 Machine State on Power-Up

This section describes the state of the KN210 after a power-up halt.

The descriptions in this section assume a machine with no errors, that the machine has just been turned on and that only the power-up diagnostics have been run. The state of the machine is not defined if individual diagnostics are run or during any other halts other than a power-up halt (SAVPSL<13:8>(RESTART_CODE) = 3).

The following sections describe data structures that are guaranteed to be constant over future versions of the maintenance mode firmware. Placement and/or existence of any other structure(s) is not implied.

6.6.1 Main Memory Layout and State

Main memory is tested and initialized by the firmware on power-up. Figure 6-5 is a diagram of how main memory is partitioned after diagnostics.

Figure 6-5 Memory Layout After Power-Up Diagnostics

6.6.1.1 Reserved Main Memory

In order to build the scatter/gather map and the bitmap, the firmware attempts to find a physically contiguous page aligned 64 Kbyte block of memory at the highest possible address that has no multiple bit errors. Single bit errors are tolerated in this section.

Of the 64 Kbytes, the upper 32 Kbytes are dedicated to the Q22-bus scatter/gather map, as shown in Figure 6-5. Of the lower 32 Kbytes, up to 16 Kbytes at the bottom of the block is allocated to the page frame number (PFN) bitmap. The size of the PFN bitmap is dependent on the extent of physical memory, each bit in the bitmap maps one page (512 bytes) of memory. The remainder of the block between the bitmap and scatter/gather map (minimally 16 Kbytes) is allocated for the firmware.

6.6.1.2 PFN Bitmap

The PFN bitmap is a data structure that indicates which pages in memory are deemed useable by operating systems. The bitmap is built by the diagnostics as a side effect of the memory tests on power-up. The bitmap always starts on a page boundary. The bitmap requires 1 Kbyte for every 4 Mbytes of main memory. Therefore, an 8 Mbyte system requires 2 Kbytes; 16 Mbyte requires 4 Kbytes; 32 Mbyte requires 8 Kbytes; and a 64 Mbyte requires 16 Kbytes. The bitmap does not map itself or anything above it. There may be memory above the bitmap which has both good and bad pages.

Each bit in the PFN bitmap corresponds to a page in main memory. There is a one to one correspondence between a PFN (origin 0) and a bit index in the bitmap. A one in the bitmap indicates that the page is good and can be used. A zero indicates that the page is bad and should not be used. By default, a page is flagged bad, if a multiple bit error occurs when referencing the page. Single bit errors, regardless of frequency, will not cause a page to be flagged bad.

The PFN bitmap is protected by a checksum stored in the NVRAM. The checksum is a simple byte wide, two's complement checksum. The sum of all bytes in the bitmap and the bitmap checksum should result in zero. Operating systems that modify the bitmap are encouraged to update this checksum to facilitate diagnosis by service personnel.

6.6.1.3 Scatter/Gather Map

On power-up, the scatter/gather map is initialized by the firmware to map to the first 4 Mbytes of main memory. Main memory pages will not be mapped, if there is a corresponding page in Q22-bus memory, or if the page is marked bad by the PFN bitmap.

On a processor halt other than power-up, the contents of the scatter/gather map is undefined, and is dependent on operating system usage.

Operating systems should not move the location of the scatter/gather map, and should access the map only on aligned longwords through the local I/O space of 2008 8000 to 2008 FFFC, inclusive. The Q22-bus map base register, (QMBR) is set up by the firmware to point to this area, and should not be changed by software.

6.6.1.4 Contents of Main Memory

The contents of main memory are undefined after the diagnostics have run. Typically, non-zero test patterns will be left in memory.

The diagnostics scrub all of main memory, so that no power-up induced errors remain in the memory system. On the KN210 memory subsystem, the state of the ECC bits and the data bits are undefined on initial power-up. This can result in single and multiple bit errors if the locations are read before written because the ECC bits are not in agreement with their corresponding data bits. An aligned longword write to every location (done by diagnostics) eliminates all power-up induced errors.

6.6.2 CMCTL Registers

The maintenance mode firmware assigns bank numbers to CMCTL registers in ascending order, without attempting to disable physical banks that contain errors. High order unused banks are set to zero. Error loggers should capture the following bits from each MEMCSR register:

MEMCSR<31> (bank enable bit). As the firmware always assigns banks in ascending order, knowing which banks are enabled is sufficient information to derive the bank numbers. MEMCSR<1:0> (bank usage). This field determines the size of the banks on the particular memory board.

Additional information should be captured from the MCSR16 and MCSR17 as needed.

6.6.3 First Level Cache

The first level cache is tested during the power-up diagnostics, flushed and then turned off. The cache is again turned on by the BOOT and the INIT command. Otherwise, the state of the first level cache is disabled.

6.6.4 Translation Buffer

The CPU translation buffer is tested by diagnostics on power-up, but not used by the firmware since it runs in physical mode. The translation buffer can be invalidated by using PR\$_TBIA, IPR 57.

6.6.5 Halt Protected Space

Halt protected space is from 2004 0000 to 2005 FFFF, inclusive, for the 128 Kbytes of maintenance mode firmware on the KN210. The halt unprotected space is from 2006 0000 to 2008 FFFF.

The diagnostic processor firmware always runs in halt protected space. When passing control to the bootstrap, the firmware exits the halt protected space, so if halts are enabled, and the halt line is asserted, the processor will then halt before booting.

The SSC decodes both spaces (256 Kbyte). That is, the ROMs appear twice in the address space. However, the halt protected space is set to 128 Kbyte, the size of the EPROMs.

6.7 Public Data Structures and Entry Points

This section describes public data structures and subroutine entry points that are public and are guaranteed to be compatible over future versions of the maintenance mode firmware.

6.7.1 Maintenance Mode Firmware EPROM Layout

The KN210 maintenance mode uses one 128 Kbyte EPROM. Approximately 120 Kbytes of this is used for code, with the remaining reserved for future expansion and customer usage. There are two copies of the firmware, one in halt protected space, and one in halt unprotected space. Both copies are identical.

Figure 6-6 KN210 Maintenance Mode EPROM Layout

The first instruction executed on halts is a branch around the system ID extension (SIE) and the callback entry points. This allows these public data structures to reside in fixed locations in the EPROM.

The callback area entry points provide a simple interface to the currently defined console for VMB and secondary bootstraps. This is documented further in the next section.

The EPROM checksum is a longword checksum from 2004 0000 to the checksum inclusive. The diagnostics use this to determine that the EPROMs can correctly be read.

The memory between the checksum and the 4 page user area at the end of the EPROMs is reserved for DIGITAL for future expansion of the maintenance mode firmware. The contents of this area is set to FF.

The four pages reserved for customer use are at the top of the EPROMs, and start at address 2005 F800 (halt protected space) or 2007 F800 (halt unprotected space). These areas are not "burned" and may be returned by OEMs or end users. The area is not tested by the maintenance mode firmware and is not included in the checksum.

6.7.2 Call-Back Entry Points

The maintenance mode firmware provides several entry points that facilitate I/O to the designated console device. Users of these entry points do not need to be aware of the console device type, be it a video terminal or workstation.

The primary intent of these routines is to provide a simple console device to VMB and secondary bootstraps, before operating systems load their own terminal drivers.

These are JSB (subroutine as opposed to procedure) entry points located in fixed locations in the firmware. These locations branch to code that in turn calls the appropriate routines.

All of the entry points are designed to run at IPL 31 on the interrupt stack in physical mode. Virtual mode is not supported. Due to internal firmware architectural restrictions, users are encouraged to only call into the halt protected entry points. These entry points are listed below.

CP\$GET_CHAR_ R4	2004 0008
CP\$MESSG_OUT_ NOLF_R4	2004 000C
CP\$READ_WTH_ PRMPT_R4	2004 0010

6.7.2.1 CP\$GETCHAR_R4

This routine returns the next character entered by the operator in R0. A timeout interval can be specified. If the timeout interval is zero, no timeout is generated. If a timeout is specified and if timeout occurs, a value of 18 (CAN) is returned instead of normal input.

Registers R0, R1, R2, R3 and R4 are modified by this routine, all others are preserved.

```

;-----
; Usage with timeout:
movl    #timeout_in_tenths_of_second,r0 ; Specify timeout.
jsb     @CP$GET_CHAR_R4                 ; Call routine.
cmpb    r0,#^x18                        ; Check for timeout.
beql    timeout_handler                 ; Branch if timeout.
; Input is in R0.

;-----
; Usage without timeout:
clrl    r0                               ; Specify no timeout.
jsb     @CP$GET_CHAR_R4                 ; Call routine.
; Input is in R0.

;-----

```

6.7.2.2 CP\$MSG_OUT_NOLF_R4

This routine outputs a message to the console. The message is specified either by a message code or a string descriptor. The routine distinguishes between message codes and descriptors by requiring that any descriptor be located outside of the first page of memory. Hence, message codes are restricted to values between 0 and 511.

Registers R0, R1, R2, R3 and R4 are modified by this routine, all others are preserved.

```

;-----
; Usage with message code:
movzbl  #console_message_code,r0        ; Specify message code.
jsb     @CP$MSG_OUT_NOLF_R4             ; Call routine.

;-----
; Usage with a message descriptor (position dependent).
movaq   5$,r0                           ; Specify address of desc.
jsb     @CP$MSG_OUT_NOLF_R4             ; Call routine.
.
.
5$:     .ascid /This is a message/       ; Message with descriptor.

;-----
; Usage with a message descriptor (position independent).

```

6-84 Maintenance Mode Firmware

```
pushab 5$ ; Generate message desc.
pushl #10$-5$ ; on stack.
movl sp,r0 ; Pass desc. addr. in R0.
jsb @CP$MSG_OUT_NOLF_R4 ; Call routine.
clrq (sp)+ ; Purge desc. from stack.
.
.
5$: .ascii /This is a message/ ; Message.
10$: ;
;-----
```

6.7.2.3 CP\$READ_WTH_PRMP_T_R4

This routine outputs a prompt message and then inputs a character string from the console. When the input is accepted, **Delete**, **Ctrl U** and **Ctrl R** functions are supported.

As with CP\$MSG_OUT_NOLF_R4, either a message code or the address of a string descriptor is passed in R0 to specify the prompt string. A value of zero results in no prompt.

A descriptor of the input string is returned in R0 and R1. R0 contains the length of the string and R1 contains the address. This routine inputs the string into the console program string buffer and therefore the caller need not provide an input buffer. Successive calls however destroy the previous contents of the input buffer.

Registers R0, R1, R2, R3 and R4 are modified by this routine, all others are preserved.

```
;-----
; Usage with a message descriptor (position independent).
pushab 10$ ; Generate prompt desc.
pushl #10$-5$ ; on stack.
movl sp,r0 ; Pass desc. addr. in R0.
jsb @CP$READ_WTH_PRMP_T_R4 ; Call routine.
clrq (sp)+ ; Purge prompt desc.
. ; Input desc in R0 and R1.
.
5$: .ascii /Prompt> / ; Prompt string.
10$: ;
;-----
```

6.7.3 SSC RAM Layout

The maintenance mode firmware uses the 1 Kbyte of battery backed up (BBU) RAM in the SSC for storage of firmware specific data structures and other information that must be preserved across power cycles. This nonvolatile RAM (NVRAM) resides in the SSC chip starting at address 2014 0400. The NVRAM should not be used by the operating systems except as documented below. This NVRAM is not reflected in the bitmap built by the firmware.

Figure 6–7 KN210 SSC NVRAM Layout

6.7.4 Public Data Structures

The following is a list of the public data structures in NVRAM used by the console.

Fields that are designated as reserved and/or internal use should not be written, since there is no protection against such corruption.

6.7.4.1 Console Program Mailbox

The console program mailbox (CPMBX) is a software data structure located at the beginning of NVRAM (2014 0400). The CPMBX is used to pass information between the maintenance mode firmware and diagnostics, VMB, or an operating system. It consists of three bytes referred to here as NVR0, NVR1, and NVR2 shown in Figures 6–8 through 6–10.

Figure 6-8 NVR0

Field	Name	Description
7:4	LANGUAGE	This field specifies the current selected language for displaying halt and error messages on terminals which support MCS.
3	RIP	If set, a restart attempt is in progress. This flag must be cleared by the operating system, if the restart succeeds.
2	BIP	If set, a bootstrap attempt is in progress. This flag must be cleared by the operating system if the bootstrap succeeds.
1:0	HLT_ ACT	Processor halt action - this field, in conjunction with the conditions specified in Table 6-1, is used to control the automatic restart/bootstrap procedure. HLT_ACT is normally written by the operating system. 0 : Restart; if that fails, reboot; if that fails, halt. 1 : Restart; if that fails, halt. 2 : Reboot; if that fails, halt. 3 : Halt.

Figure 6-9 NVR1

Field	Name	Description
2	MCS	If set, indicates that the attached terminal supports multinational character set (MCS). If clear, MCS is not supported.
1	CRT	If set, indicates that the attached terminal is a CRT. If clear, indicates that the terminal is hardcopy.

Figure 6-10 NVR2

Field	Name	Description
7:0	KEYBOARD	This field indicates the national keyboard variant in use.

6.7.4.2 Firmware Stack

This section contains the stack that is used by all of the firmware, with the exception of VMB, which has its own built in stack.

6.7.4.3 Diagnostic State

This area is used by the firmware resident diagnostics. This section is not documented here.

6.7.4.4 User Area

The KN210 console reserves the last longword (address 2014 07FC) of the NVRAM for customer use. This location is not tested by the console firmware. Its value is undefined.

6.8 Error Messages

The error messages issued by the maintenance mode firmware fall into three categories, halt code messages, VMB error messages, and console messages.

6.8.1 Halt Code Messages

Except on power-up, which is not treated as an error condition, the following messages are issued by the firmware whenever the processor halts.

For example:

```
?06 HLT INST
PC = 800050D3
```

The number preceding the halt message is the *halt code*. This number is obtained from SAVPSL<13:8>(RESTART_CODE), IPR 43, which is written on any CVAX processor restart operation.

Table 6-10 HALT Messages

Code	Message	Description
?02	EXT HLT	External halt, caused by either console BREAK condition or Q22-bus BHALT_L.
_03	—	Power-up, no halt message is displayed; _03 is not displayed. However, the presence of the firmware banner and diagnostic countdown indicates this halt reason.
?04	ISP ERR	In attempting to push state onto the interrupt stack during an interrupt or exception, the processor discovered that the interrupt stack was mapped NO ACCESS or NOT VALID.
?05	DBL ERR	The processor attempted to report a machine check to the operating system, and a second machine check occurred.
?06	HLT INST	The processor executed a HALT instruction in kernel mode.
?07	SCB ERR3	The SCB vector had bits <1:0> equal to 3.
?08	SCB ERR2	The SCB vector had bits <1:0> equal to 2.
?0A	CHM FR ISTK	A change mode instruction was executed when PSL<IS> was set.
?0B	CHM TO ISTK	The SCB vector for a change mode had bit <0> set.
?0C	SCB RD ERR	A hard memory error occurred while the processor was trying to read an exception or interrupt vector.

Table 6-10 (Cont.) HALT Messages

Code	Message	Description
?10	MCHK AV	An access violation or an invalid translation occurred during machine check exception processing.
?11	KSP AV	An access violation or translation not valid occurred during processing of a kernel stack not valid exception.
?12	DBL ERR2	Double machine check error. A machine check occurred while trying to service a machine check.
?13	DBL ERR3	Double machine check error. A machine check occurred while trying to service a kernel stack not valid exception.
?19	PSL EXC5 ¹	PSL<26:24> = 5 on interrupt or exception.
?1A	PSL EXC6 ¹	PSL<26:24> = 6 on interrupt or exception.
?1B	PSL EXC7 ¹	PSL<26:24> = 7 on interrupt or exception.
?1D	PSL REI5 ¹	PSL<26:24> = 5 on an REI instruction
?1E	PSL REI6 ¹	PSL<26:24> = 6 on an REI instruction.
?1F	PSL REI7 ¹	PSL<26:24> = 7 on an REI instruction.

¹For the last six cases, the VAX architecture does not allow execution on the interrupt stack while in a mode other than kernel. In the first three cases, an interrupt is attempting to run on the interrupt stack while not in kernel mode. In the last three cases, an REI instruction is attempting to return to a mode other than kernel and still run on the interrupt stack.

6.8.2 Console Error Messages

These error messages are issued in response to a console command that has error(s).

Table 6-11 Console Error Messages

Code	Message	Description
?20	CORRPTN	The console program database has been corrupted.
?21	ILL REF	Illegal reference. The requested reference would violate virtual memory protection, the address is not mapped, the reference is invalid in the specified address space, or the value is invalid in the specified destination.
?22	ILL CMD	The command string cannot be parsed.

Table 6-11 (Cont.) Console Error Messages

Code	Message	Description
?23	INV DGT	A number has an invalid digit.
?24	LTL	The command was too large for the console to buffer. The message is issued only after receipt of the terminating carriage return.
?25	ILL ADR	The address specified falls outside the limits of the address space.
?26	VAL TOO LRG	The value specified does not fit in the destination.
?27	SW CONF	Switch conflict, for example, two different data sizes are specified for an EXAMINE command.
?28	UNK SW	The switch is unrecognized.
?29	UNK SYM	The symbolic address in an EXAMINE or DEPOSIT command is unrecognized.
?2A	CHKSM	The command or data checksum of an X command is incorrect. If the data checksum is incorrect, this message is issued, and is not abbreviated to <i>Illegal command</i> .
?2B	HLTED	The operator entered a HALT command.
?2C	FND ERR	A FIND command failed either to find the RPB or 128 Kbytes of good memory.
?2D	TMOUT	During an X command, data failed to arrive in the time expected (60 seconds).
?2E	MEM ERR	A machine check occurred with a code of 80 (hex) or 81 (hex), indicating a read or write memory error.
?2F	UNXINT	Unexpected interrupt or exception.
?30	UNIMPLEMENTED	Unimplemented function.
?31	QUAL NOVAL	Qualifier does not take a value.
?32	QUAL AMBG	Ambiguous qualifier.
?33	QUAL REQ VAL	Qualifier requires a value.
?34	QUAL OVERF	Too many qualifiers.
?35	ARG OVERF	Too many arguments.
?36	AMBG CMD	Ambiguous command.
?37	INSUF ARG	Insufficient arguments.

6.8.3 VMB error messages

Table 6-12 lists errors issued by VMB.

Table 6-12 VMB Error Messages

Code	Message	Description
?40	NOSUCHDEV	No bootable devices found.
?41	DEVASSIGN	Device is not present.
?42	NOSUCHFILE	Program image not found.
?43	FILESTRUCT	Invalid boot device file structure.
?44	BADCHKSUM	Bad checksum on header file.
?45	BADFILEHDR	Bad file header.
?46	BADIRECTORY	Bad directory file.
?47	FILNOTCNTG	Invalid program image format.
?48	ENDOFFILE	Premature end of file encountered.
?49	BADFILENAME	Bad file name given.
?4A	BUFFEROVF	Program image does not fit in available memory.
?4B	CTRLERR	Boot device I/O error.
?4C	DEVINACT	Failed to initialize boot device.
?4D	DEVOFFLINE	Device is offline.
?4E	MEMERR	Memory initialization error.
?4F	SCBINT	Unexpected SCB exception or machine check.
?50	SCB2NDINT	Unexpected exception after starting program image.
?51	NOROM	No valid ROM image found.
?52	NOSUCHNODE	No response from load server.
?53	INSFMAPREG	Invalid memory configuration.
?54	RETRY	No devices bootable, retrying.

A

Specifications

This appendix contains the physical, electrical and environmental specifications for the KN210 CPU module and the KN210 I/O module.

A.1 Physical Specifications

The KN210 CPU module, KN210 I/O module and MS650-AA are quad-height modules with the following dimensions:

Dimension	Measurement
Height	10.457 (+0.015/-0.020) inches
Length	8.430 (+0.010/-0.010) inches
Width	0.375 inches maximum (nonconductive) 0.343 inches maximum (conductive)

NOTE

Width, as defined for Digital modules, is the height of components above the surface of the module.

A.2 Electrical Specifications

The power requirements for the KN210 CPU module are as follows:

+5 V $\pm 5\%$	+12V $\pm 5\%$
4.5 A maximum	0.13 A maximum

Typical currents are 10 percent less than the specified maximum.

A-2 Specifications

The bus loads for the KN210 CPU module are as follows:

- 3.5 ac loads
- 1.0 dc loads

The power requirements for the KN210 I/O module are as follows:

+5 V $\pm 5\%$	+12V $\pm 5\%$
5.0 A maximum	0.23 A maximum

Typical currents are 10 percent less than the specified maximum.

The KN210 I/O module does not present any loads to the Q22-bus.

A.3 Environmental Specifications

The environmental specifications for the KN210 CPU module and the KN210 I/O module are as follows:

Operating Conditions

Temperature	5°C (41°F) to 60°C (140°F) with a rate of change no greater than $20 \pm 2^\circ\text{C}$ ($36 \pm 4^\circ\text{F}$) per hour at sea level. Derate maximum temperature by 1.8°C for each 1000 meters (1°F for each 1000 ft) of altitude above sea level.
Humidity	0% to 95% noncondensing, with a maximum wet bulb temperature of 32°C (90°F) and a minimum dew point temperature of 2°C (36°F).
Altitude	Up to 2,400 meters (8,000 ft) with a rate of change no greater than 300 meters per min (1000 ft per min).

Nonoperating Conditions Less Than 60 Days

Temperature	-40°C to $+66^\circ\text{C}$ (-40°F to $+151^\circ\text{F}$) with a rate of change no greater than $11 \pm 2^\circ\text{C}$ ($20 \pm 4^\circ\text{F}$) per hour at sea level. Derate the maximum temperature by 1.8°C for each 1000 meters (1°F for each 1000 ft) of altitude above sea level.
Humidity	Up to 95% noncondensing.
Altitude	Up to 4,900 meters (16,000 ft) with a rate of change no greater than 600 meters per min (2000 feet per min).

Nonoperating Conditions Greater Than 60 Days

Temperature	+5°C to +60°C (+40°F to +140°F) with a rate of change no greater than $20 \pm 2^\circ\text{C}$ ($36 \pm 4^\circ\text{F}$) per hour at sea level. Derate the maximum temperature by 1.8°C for each 1000 meters (1°F for each 1000 ft) of altitude above sea level.
Humidity	10% to 95% noncondensing, with a maximum wet bulb temperature of 32°C (90°F) and a minimum dew point temperature of 2°C (36°F).
Altitude	Up to 2,400 meters (8,000 ft) with a rate of change no greater than 300 meters per min (1000 ft per min).

B

Address Assignments

B.1 R3000 Physical Address Space Map

Table B-1 lists the R3000 memory space.

Table B-1 R3000 Memory Space

Contents	Address Range
Local memory space (Up to 64 Mbytes)	0000 0000 - 03FF FFFF
Reserved memory space	0400 0000 - 1FFF FFFF

Table B-2 lists the R3000 I/O memory space.

Table B-2 R3000 I/O Space

Contents	Address Range
Reserved Q-22 bus I/O space	1000 0000 - 1000 0007
Q-22 bus floating address space	1000 0008 - 1000 07FF
User reserved Q-22 bus I/O space	1000 0800 - 1000 0FFF
Reserved Q-22 bus I/O space	1000 1000 - 1000 1F3F
Interprocessor communication register	1000 1F40
Reserved Q-22 bus I/O space	1000 1F48 - 1000 1FFF
Reserved I/O module address space	1000 2000 - 1003 FFFF
Two copies of local ROM	1004 0000 - 1007 FFFF
DMA system configuration register	1008 0000
DMA system error register	1008 0004
DMA master error address register	1008 0008
DMA slave error address register	1008 000C

B-2 Address Assignments

Table B-2 (Cont.) R3000 I/O Space

Contents	Address Range
Q-22 bus map base register	1008 0010
Reserved	1008 0014 - 1008 00FF
Main memory configuration register 00	1008 0100
Main memory configuration register 01	1008 0104
Main memory configuration register 02	1008 0108
Main memory configuration register 03	1008 010C
Main memory configuration register 04	1008 0110
Main memory configuration register 05	1008 0114
Main memory configuration register 06	1008 0118
Main memory configuration register 07	1008 011C
Main memory configuration register 08	1008 0120
Main memory configuration register 09	1008 0124
Main memory configuration register 10	1008 0128
Main memory configuration register 11	1008 012C
Main memory configuration register 12	1008 0130
Main memory configuration register 13	1008 0134
Main memory configuration register 14	1008 0138
Main memory configuration register 15	1008 013C
Main memory error status register	1008 0140
Main memory control/diagnostic status register	1008 0144
Reserved I/O module address space	1008 0148 - 1008 3FFF
Interrupt status	1008 4000
Boot and diagnostic register	1008 4004
Select processor register	1008 4008
100 Hz clock	1008 4010
Reserved I/O module address space	1008 4014 - 1008 7FFF
Q-22 bus map registers	1008 8000 - 1008 FFFF
Reserved I/O module address space	1009 0000 - 1013 FFFF
SSC base address register	1014 0000
SSC configuration register	1014 0010
CDAL bus timeout control register	1014 0020
Diagnostic LED register	1014 0030
Reserved I/O module address space	1014 0034 - 1014 0068
Time-of-year register	1014 006C
Reserved	1014 0070 - 1014 007C
Console receiver control/status	1014 0080

Table B-2 (Cont.) R3000 I/O Space

Contents	Address Range
Console receiver data buffer	1014 0084
Console transmitter control/status	1014 0088
Console transmitter data buffer	1014 008C
Reserved	1014 0090 - 1014 00DB
I/O system reset register	1014 00DC
Reserved	1014 00E0
ROM data register	1014 00F0
Bus timeout counter	1014 00F4
Interval timer	1014 00F8
Reserved	1014 00FC - 1014 00FF
Timer 0 control register	1014 0100
Timer 0 interval register	1014 0104
Timer 0 next interval register	1014 0108
Timer 0 interrupt vector	1014 010C
Timer 1 control register	1014 0110
Timer 1 interval register	1014 0114
Timer 1 next interval register	1014 0118
Timer 1 interrupt vector	1014 011C
Address decode match register	1014 0130
Address decode mask register	1014 0134
Reserved	1014 0138 - 1014 033F
Battery backed-up RAM	1014 0400 - 1014 07FF
Local Q-22 bus memory space	1400 0000 - 143F FFFF
Reserved (4 copies local Q-22 memory)	1440 0000 - 14FF FFFF
Reserved	1500 0000 - 1600 004F
Vector read register 0	1600 0050 ¹
Vector read register 1	1600 0054 ¹
Vector read register 2	1600 0058 ¹
Vector read register 3	1600 005C ¹
Reserved	1600 0060 - 16FF FFFF
Write error address register	1700 0000 ¹
Reserved	1700 0004 - 17BF FFFF
ROM	1FC0 0000 - 1FC1 FFFF ¹
Reserved	1FC2 0000 - FFFF FFFF

¹Only accessible from R3000 processor.

B.2 CVAX Physical Address Space Map

Table B-3 lists the CVAX memory space.

Table B-3 VAX Memory Space

Contents	Address Range
Local memory space (Up to 64 Mbytes)	0000 0000 - 03FF FFFF
Reserved memory space	0400 0000 - 1FFF FFFF

Table B-4 lists the VAX input/output memory space.

Table B-4 CVAX I/O Space

Contents	Address Range
Reserved Q22-bus I/O space	2000 0000 - 2000 0007
Q22-bus floating address space	2000 0008 - 2000 07FF
User reserved Q22-bus I/O space	2000 0800 - 2000 0FFF
Reserved Q22-bus I/O space	2000 1000 - 2000 1F3F
Interprocessor communication register	2000 1F40
Reserved Q22-bus I/O space	2000 1F48 - 2000 1FFF
Reserved I/O module address space	2000 2000 - 2003 FFFF
MicroVAX system type register (in ROM)	2004 0004
Local ROM - halt mode	2004 0000 - 2005 FFFF
Local ROM - run mode	2006 0000 - 2007 FFFF
DMA system configuration register	2008 0000
DMA system error register	2008 0004
DMA master error address register	2008 0008
DMA slave error address register	2008 000C
Q22-bus map base register	2008 0010
Reserved	2008 0014 - 2008 00FF
Main memory configuration register 00	2008 0100
Main memory configuration register 01	2008 0104
Main memory configuration register 02	2008 0108
Main memory configuration register 03	2008 010C
Main memory configuration register 04	2008 0110
Main memory configuration register 05	2008 0114
Main memory configuration register 06	2008 0118
Main memory configuration register 07	2008 011C
Main memory configuration register 08	2008 0120

Table B-4 (Cont.) CVAX I/O Space

Contents	Address Range
Main memory configuration register 09	2008 0124
Main memory configuration register 10	2008 0128
Main memory configuration register 11	2008 012C
Main memory configuration register 12	2008 0130
Main memory configuration register 13	2008 0134
Main memory configuration register 14	2008 0138
Main memory configuration register 15	2008 013C
Main memory error status register	2008 0140
Main memory control/diagnostic status register	2008 0144
Reserved I/O module address space	2008 0148 - 2008 3FFF
Interrupt status	2008 4000
Boot and diagnostic register	2008 4004
Select processor register	2008 4008
100 Hz clock	2008 4010
Reserved I/O module address space	2008 4014 - 2008 7FFF
Q-22 bus map registers	2008 8000 - 2008 FFFF
Reserved I/O module address space	2009 0000 - 2013 FFFF
SSC base address register	2014 0000
SSC configuration register	2014 0010
CDAL bus timeout control register	2014 0020
Diagnostic LED register	2014 0030
Reserved I/O module address space	2014 0034 - 2014 0068
Time-of-year register	2014 006C
Reserved	2014 0070 - 2014 007C
Console receiver control/status	2014 0080
Console receiver data buffer	2014 0084
Console transmitter control/status	2014 0088
Console transmitter data buffer	2014 008C
Reserved	2014 0090 - 2014 00DB
I/O system reset register	2014 00DC
Reserved	2014 00E0
ROM data register	2014 00F0
Bus timeout counter	2014 00F4
Interval timer	2014 00F8
Reserved	2014 00FC - 2014 00FF
Timer 0 control register	2014 0100

B-6 Address Assignments

Table B-4 (Cont.) CVAX I/O Space

Contents	Address Range
Timer 0 interval register	2014 0104
Timer 0 next interval register	2014 0108
Timer 0 interrupt vector	2014 010C
Timer 1 control register	2014 0110
Timer 1 interval register	2014 0114
Timer 1 next interval register	2014 0118
Timer 1 interrupt vector	2014 011C
Address decode match register	2014 0130
Address decode mask register	2014 0134
Reserved	2014 0138 - 2014 033F
Battery backed-up RAM	2014 0400 - 2014 07FF
Reserved I/O module address space	2014 0800 - 2FFF FFFF
Local Q-22 bus memory space	3000 0000 - 303F FFFF
Reserved I/O module address space	3040 0000 - 3FFF FFFF

B.2.1 External IPRs

Several of the internal processor registers (IPRs) on the KN210 CPU module are implemented in the SSC rather than in the CVAX chip. These registers are referred to as external IPRs, and are listed in Table B-5.

Table B-5 External IPRs

IPR Number	Register Name	Abbreviation
27	Time-of-year register	TODR
28-31	Reserved	
32	Console receiver control/status	RXCS
33	Console receiver data buffer	RXDB
34	Console transmitter control/status	TXCS
35	Console transmitter data buffer	TXDB
55	I/O system reset register	IORESET

B.3 Global Q22-bus Address Space Map

The addresses and memory contents of the Q22-bus memory space are listed in Table B-6.

Table B-6 Q22-bus Memory Space

Contents	Address Range
Q22-bus memory space (octal)	0000 0000 - 1777 7777

The contents and addresses of the Q22-bus I/O space with BBS7 asserted are listed Table B-7.

Table B-7 Q22-bus I/O Space with BBS7 Asserted

Contents	Address Range
Q22-bus I/O space (octal)	1776 0000 - 1777 7777
Reserved Q22-bus I/O space	1776 0000
Q22-bus floating address space	1776 0010 - 1776 3777
User reserved Q22-bus I/O space	1776 4000 - 1776 7777
Reserved Q22-bus I/O space	1777 0000 - 1777 7477
Interprocessor communication register	1777 7500
Reserved Q22-bus I/O space	1777 7502 - 1777 7777

C

Q22-bus Specification

C.1 Introduction

The Q22-bus, also known as the extended LSI-11 bus, is the low-end member of Digital's bus family. All of Digital's microcomputers, such as the MicroVAX I, MicroVAX II, MicroVAX 3500, MicroVAX 3600, and MicroPDP-11 use the Q22-bus.

The Q22-bus consists of 42 bidirectional and 2 unidirectional signal lines. These form the lines along which the processor, memory, and I/O devices communicate with each other.

Addresses, data, and control information are sent along these signal lines, some of which contain time-multiplexed information. The lines are divided as follows:

- Sixteen multiplexed data/address lines — BDAL<15:00>
- Two multiplexed address/parity lines — BDAL<17:16>
- Four extended address lines — BDAL<21:18>
- Six data transfer control lines — BBS7, BDIN, BDOUT, BRPLY, BSYNC, BWTBT
- Six system control lines — BHALT, BREF, BEVNT, BINIT, BDCOK, BPOK
- Ten interrupt control and direct memory access control lines — BIAKO, BIAKI, BIRQ4, BIRQ5, BIRQ6, BIRQ7, BDMGO, BDMR, BSACK, BDMGI

In addition, a number of power, ground, and space lines are defined for the bus. Refer to Table C-1 for a detailed description of these lines.

C-2 Q22-bus Specification

The discussion in this appendix applies to the general 22-bit physical address capability. All modules used with the KN210 CPU module must use 22-bit addressing.

Most Q22-bus signals are bidirectional and use terminations for a negated (high) signal level. Devices connect to these lines by way of high-impedance bus receivers and open collector drivers. The asserted state is produced when a bus driver asserts the line low.

Although bidirectional lines are electrically bidirectional (any point along the line can be driven or received), certain lines are functionally unidirectional. These lines communicate to or from a bus master (or signal source), but not both. Interrupt acknowledge (BIAK) and direct memory access grant (BDMG) signals are physically unidirectional in a daisy-chain fashion. These signals originate at the processor output signal pins. Each is received on device input pins (BIAKI or BDMGI) and is conditionally retransmitted through device output pins (BIAKO or BDMGO). These signals are received from higher priority devices and are retransmitted to lower priority devices along the bus, establishing the position-dependent priority scheme.

C.1.1 Master/Slave Relationship

Communication between devices on the bus is asynchronous. A master/slave relationship exists throughout each bus transaction. Only one device has control of the bus at any one time. This controlling device is termed the bus master, or arbiter. The master device controls the bus when communicating with another device on the bus, termed the slave.

The bus master (typically the processor or a DMA device) initiates a bus transaction. The slave device responds by acknowledging the transaction in progress and by receiving data from, or transmitting data to, the bus master. Q22-bus control signals transmitted or received by the bus master or bus slave device must complete the sequence according to bus protocol.

The processor controls bus arbitration, that is, which device becomes bus master at any given time. A typical example of this relationship is a disk drive, as master, transferring data to memory as slave. Communication on the Q22-bus is interlocked so that, for certain control signals issued by the master device, there must be a response from the slave in order to complete the transfer. It is the master/slave signal protocol that makes the Q22-bus asynchronous. The asynchronous operation precludes the need for synchronizing with, and waiting for, clock pulses.

Since bus cycle completion by the bus master requires response from the slave device, each bus master must include a timeout error circuit that aborts the bus cycle if the slave does not respond to the bus transaction within 10 μ s. The actual time before a timeout error occurs must be longer than the reply time of the slowest peripheral or memory device on the bus.

C.2 Q22-bus Signal Assignments

Table C-1 lists the data and address signal assignments. Table C-2 lists the control signal assignments. Table C-3 lists the power and ground signal assignments. Table C-4 lists the spare signal assignments.

Table C-1 Data and Address Signal Assignments

Data and Address Signal	Pin Assignment
BDAL0	AU2
BDAL1	AV2
BDAL2	BE2
BDAL3	BF2
BDAL4	BH2
BDAL5	BJ2
BDAL6	BK2
BDAL7	BL2
BDAL8	BM2
BDAL9	BN2
BDAL10	BP2
BDAL11	BR2
BDAL12	BS2
BDAL13	BT2
BDAL14	BU2
BDAL15	BV2
BDAL16	AC1
BDAL17	AD1
BDAL18	BC1
BDAL19	BD1
BDAL20	BE1
BDAL21	BF1

Table C-2 Control Signal Assignments

Control Signal	Pin Assignment
Data Control	
BDOUT	AE2
BRPLY	AF2
BDIN	AH2
BSYNC	AJ2
BWTBT	AK2
BBS7	AP2
Interrupt Control	
BIRQ7	BP1
BIRQ6	AB1
BIRQ5	AA1
BIRQ4	AL2
BIAKO	AN2
BIAKI	AM2
DMA Control	
BDMR	AN1
BSACK	BN1
BDMGO	AS2
BDMGI	AR2
System Control	
BHALT	AP1
BREF	AR1
BEVNT	BR1
BINIT	AT2
BDCOK	BA1
BPOK	BB1

Table C-3 Power and Ground Signal Assignments

Power and Ground	Pin Assignment
+5 B (battery) or +12 B (battery)	AS1
+12 B	BS1
+5 B	AV1
+5	AA2
+5	BA2
+5	BV1
+12	AD2
+12	BD2
+12	AB2
-12	AB2
-12	BB2
GND	AC2
GND	AJ1
GND	AM1
GND	AT1
GND	BC2
GND	BJ1
GND	BM1
GND	BT1

Table C-4 Spare Signal Assignments

Spare	Pin Assignment
SSpare1	AE1
SSpare3	AH1
SSpare8	BH1
SSpare2	AF1
MSpareA	AK1
MSpareB	AL1
MSpareB	BK1
MSpareB	BL1
PSpare1	AU1
ASpare2	BU1

C.3 Data Transfer Bus Cycles

Data transfer bus cycles, executed by bus master devices, transfer 32-bit words or 8-bit bytes to or from slave devices. In block mode, multiple words can be transferred to sequential word addresses, starting from a single bus address. Data transfer bus cycles are listed and defined in Table C-5.

Table C-5 Data Transfer Operations

Bus Cycle	Definition	Function (with respect to the bus master)
DATI	Data word input	Read
DATO	Data word output	Write
DATOB	Data byte output	Write-byte
DATIO	Data word input/output	Read-modify-write
DATIOB	Data word input/byte output	Read-modify-write byte
DATBI	Data block input	Read block
DATBO	Data block output	Write block

The bus signals listed in Table C-6 are used in the data transfer operations described in Table C-5.

Table C-6 Bus Signals for Data Transfers

Signal	Definition	Function
BDAL<21:00> L	22 data/address lines	BDAL<15:00> L are used for word and byte transfers. BDAL<17:16> L are used for extended addressing, memory parity error (16), and memory parity error enable (17) functions. BDAL<21:18> L are used for extended addressing beyond 256 Kbytes.
BSYNC L	Bus cycle control	Indicates bus transaction in progress.
BDIN L	Data input indicator	Strobe signals

Table C-6 (Cont.) Bus Signals for Data Transfers

Signal	Definition	Function
BDOUT L	Data output indicator	Strobe signals
BRPLY L	Slave's acknowledge of bus cycle	Strobe signals
BWTBT L	Write/byte control	Control signals
BBS7	I/O device select	Indicates address is in the I/O page.

Data transfer bus cycles can be reduced to five basic types: DATI, DATO(B), DATIO(B), DATBI, and DATBO. These transactions occur between the bus master and one slave device selected during the addressing part of the bus cycle.

C.3.1 Bus Cycle Protocol

Before initiating a bus cycle, the previous bus transaction must have been completed (BSYNC L negated) and the device must become bus master. The bus cycle can be divided into two parts: addressing and data transfer. During addressing, the bus master outputs the address for the desired slave device, memory location, or device register. The selected slave device responds by latching the address bits and holding this condition for the duration of the bus cycle until BSYNC L becomes negated. During data transfer the actual data transfer occurs.

C.3.2 Device Addressing

Device addressing of a data transfer bus cycle comprises an address setup and deskew time, and an address hold and deskew time. During address setup and deskew time, the bus master does the following operations:

- Asserts BDAL<21:00> L with the desired slave device address bits.
- Asserts BBS7 L if a device in the I/O page is being addressed.
- Asserts BWTBT L if the cycle is a DATO(B) or DATBO bus cycle.

C-8 Q22-bus Specification

During this time, the address, BBS7 L, and BWTBT L signals are asserted at the slave bus receiver for at least 75 ns before BSYNC goes active. Devices in the I/O page ignore the nine high-order address bits BDAL<21:13>, and instead, decode BBS7 L along with the 13 low-order address bits. An active BWTBT L signal during address setup time indicates that a DATO(B) or DATBO operation follows, while an inactive BWTBT L indicates a DATI, DATBI, or DATIO(B) operation.

The address hold and deskew time begins after BSYNC L is asserted.

The slave device uses the active BSYNC L bus received output to clock BDAL address bits, BBS7 L, and BWTBT L into its internal logic. BDAL<21:00> L, BBS7 L, and BWTBT L remain active for 25 ns minimum after the BSYNC L bus receiver goes active. BSYNC L remains active for the duration of the bus cycle.

Memory and peripheral devices are addressed similarly, except for the way the slave device responds to BBS7 L. Addressed peripheral devices must not decode address bits on BDAL<21:13> L. Addressed peripheral device can respond to a bus cycle when BBS7 L is asserted (low) during the addressing of the cycle. When asserted, BBS7 L indicates that the device address resides in the I/O page (the upper 4K address space). Memory devices generally do not respond to addresses in the I/O page; however, some system applications may permit memory to reside in the I/O page for use as DMA buffers, read-only memory bootstraps, and diagnostics.

DATI

The DATI bus cycle, shown in Figure C-1, is a read operation. During DATI, data is input to the bus master. Data consists of 16-bit word transfers over the bus. During data transfer of the DATI bus cycle, the bus master asserts BDIN L 100 ns minimum after BSYNC L is asserted. The slave device responds to BDIN L active as follows:

- Asserts BRPLY L 0 ns minimum (8 ns maximum to avoid bus timeout) after receiving BDIN L, and 125 ns maximum before BDAL bus driver data bits are valid.
- Asserts BDAL<21:00> L with the addressed data and error information 0 ns (minimum) after receiving BDIN, and 125 ns (maximum) after assertion of BRPLY.

MA-1074-87

Figure C-1 DATI Bus Cycle

When the bus master receives BRPLY L, it does the following:

- Waits at least 200 ns deskew time and then accepts input data at BDAL<17:00> L bus receivers. BDAL <17:16> L are used for transmitting parity errors to the master.
- Negates BDIN L 200 ns minimum to 2 μ s maximum after BRPLY L goes active.

The slave device responds to BDIN L negation by negating BRPLY L and removing read data from BDAL bus drivers. BRPLY L must be negated 100 ns maximum prior to removal of read data. The bus master responds to the negated BRPLY L by negating BSYNC L.

Conditions for the next BSYNC L assertion are as follows:

- BSYNC L must remain negated for 200 ns minimum.
- BSYNC L must not become asserted within 300 ns of previous BRPLY L negation.

Figure C-2 shows DATI bus cycle timing.

NOTE

Continuous assertion of BSYNC L retains control of the bus by the bus master, and the previously addressed slave device remains selected. This is done for DATIO(B) bus cycles where DATO or DATOB follows a DATI without BSYNC L negation and a second device addressing operation. Also, a slow slave device can hold off data transfers to itself by keeping BRPLY L asserted, which causes the master to keep BSYNC L asserted.

DATOB

DATOB, shown in Figure C-3, is a write operation. Data is transferred in 32-bit words (DATO) or 8-bit bytes (DATOB) from the bus master to the slave device. The data transfer output can occur after the addressing part of a bus cycle when BWTBT L has been asserted by the bus master, or immediately following an input transfer part of a DATIOB bus cycle.

MA-1084-87

Figure C-2 DATI Bus Cycle Timing

C-12 Q22-bus Specification

MA-1081-87

Figure C-3 DATO or DATOB Bus Cycle

The data transfer part of a DATOB bus cycle comprises a data setup and deskew time and a data hold and deskew time.

During the data setup and deskew time, the bus master outputs the data on BDAL<15:00> L at least 100 ns after BSYNC L assertion. BWTBT L remains negated for the length of the bus cycle. If the transfer is a byte transfer, BWTBT L remains asserted. If it is the output of a DATIOB, BWTBT L becomes asserted and lasts the duration of the bus cycle.

During a byte transfer, BDAL<00> L selects the high or low byte. This occurs in the addressing part of the cycle. If asserted, the high byte (BDAL<15:08> L) is selected; otherwise, the low byte (BDAL<07:00> L) is selected. An asserted BDAL 16 L at this time forces a parity error to be written into memory if the memory is a parity-type memory. BDAL 17 L is not used for write operations. The bus master asserts BDOUT L at least 100 ns after BDAL and BDWTBT L bus drivers are stable. The slave device responds by asserting BRPLY L within 10 μ s to avoid bus timeout. This completes the data setup and deskew time.

During the data hold and deskew time, the bus master receives BRPLY L and negates BDOUT L, which must remain asserted for at least 150 ns from the receipt of BRPLY L before being negated by the bus master. BDAL<17:00> L bus drivers remain asserted for at least 100 ns after BDOUT L negation. The bus master then negates BDAL inputs.

During this time, the slave device senses BDOUT L negation. The data is accepted and the slave device negates BRPLY L. The bus master responds by negating BSYNC L. However, the processor does not negate BSYNC L for at least 175 ns after negating BDOUT L. This completes the DATOB bus cycle. Before the next cycle, BSYNC L must remain unasserted for at least 200 ns. Figure C-4 shows DATOB bus cycle timing.

DATIOB

The protocol for a DATIOB bus cycle is identical to the addressing and data transfer part of the DATI and DATOB bus cycles, and is shown in Figure C-5. After addressing the device, a DATI cycle is performed as explained earlier; however, BSYNC L is not negated. BSYNC L remains active for an output word or byte transfer (DATOB). The bus master maintains at least 200 ns between BRPLY L negation during the DATI cycle and BDOUT L assertion. The cycle is terminated when the bus master negates BSYNC L, as described for DATOB. Figure C-6 illustrates DATIOB bus cycle timing.

C-14 Q22-bus Specification

MA-1080-87

Figure C-4 DATO or DATOB Bus Cycle Timing

MA-1082-87

Figure C-5 DATIO or DATIOB Bus Cycle

C-16 Q22-bus Specification

MA-1060-87

Figure C-6 DATIO or DATIOB Bus Cycle Timing

C.4 Direct Memory Access

The direct memory access (DMA) capability allows direct data transfer between I/O devices and memory. This is useful when using mass storage devices (for example, disks) that move large blocks of data to and from memory. A DMA device needs to be supplied with only the starting address in memory, the starting address in mass storage, the length of the transfer, and whether the operation is read or write. When this information is available, the DMA device can transfer data directly to or from memory. Since most DMA devices must perform data transfers in rapid succession or lose data, DMA devices are given the highest priority.

DMA is accomplished after the processor (normally bus master) has passed bus mastership to the highest priority DMA device that is requesting the bus. The processor arbitrates all requests and grants the bus to the DMA device electrically closest to it. A DMA device remains bus master until it relinquishes its mastership. The following control signals are used during bus arbitration:

- BDMGI L DMA grant input
- BDMGO L DMA grant output
- BDMR L DMA request line
- BSACK L bus grant acknowledge

C.4.1 DMA Protocol

A DMA transaction can be divided into the following three phases:

- Bus mastership acquisition phase
- Data transfer phase
- Bus mastership relinquishment phase

During the bus mastership acquisition phase, a DMA device requests the bus by asserting BDMR L. The processor arbitrates the request and initiates the transfer of bus mastership by asserting BDMGO L.

The maximum time between BDMR L assertion and BDMGO L assertion is DMA latency. This time is processor-dependent. BDMGO L/BDMGI L is one signal that is daisy-chained through each module in the backplane.

It is driven out of the processor on the BDMGO L pin, enters each module on the BDMGI L pin, and exits on the BDMGO L pin. This signal passes through the modules in descending order of priority until it is stopped by the requesting device. The requesting device blocks the output of BMDGO L and asserts BSACK L. If BDMR L is continuously asserted, the bus hangs.

During the data transfer phase, the DMA device continues asserting BSACK L. The actual data transfer is performed as described earlier.

The DMA device can assert BSYNC L for a data transfer 250 ns minimum after it received BDMGI L and its BSYNC L bus receiver is negated.

During the bus mastership relinquishment phase, the DMA device gives up the bus by negating BSACK L. This occurs after completing (or aborting) the last data transfer cycle (BRPLY L negated). BSACK L can be negated up to a maximum of 300 ns before negating BSYNC L.

NOTE

If multiple data transfers are performed during this phase, consideration must be given to the use of the bus for other system functions, such as memory refresh (if required).

Figure C-7 shows the DMA protocol, and Figure C-8 shows DMA request/grant timing.

C.4.2 Block Mode DMA

For increased throughput, block mode DMA can be implemented on a device for use with memories that support this type of transfer. In a block mode transaction, the starting memory address is asserted, followed by data for that address, and data for consecutive addresses.

By eliminating the assertion of the address for each data word, the transfer rate is almost doubled.

There are two types of block mode transfers, DATBI (input) and DATBO (output). The DATBI bus cycle is described in Section C.4.2.1 and illustrated in Figure C-9.

The DATBO bus cycle is described in Section C.4.2.2 and illustrated in Figure C-10.

Figure C-7 DMA Protocol

C-20 Q22-bus Specification

MA-1078-87

Figure C-8 DMA Request/Grant Timing

MA-1088-87

Figure C-9 DATBI Bus Cycle Timing

C-22 Q22-bus Specification

MA-1087-87

Figure C-10 DATBO Bus Cycle Timing

C.4.2.1 DATBI Bus Cycle

Before a DATBI block mode transfer can occur, the DMA bus master device must request control of the bus. This occurs under conventional Q22-bus protocol.

A block mode DATBI transfer is executed as follows:

- **Address device memory**—the address is asserted by the bus master on TADDR<21:00> along with the negation of TWTBT. The bus master asserts TSYNC 150 ns minimum after gating the address onto the bus.
- **Decode address**—the appropriate memory device recognizes that it must respond to the address on the bus.
- **Request data**—the address is removed by the bus master from TADDR<21:00> 100 ns minimum after the assertion of TSYNC. The bus master asserts the first TDIN 100 ns minimum after asserting TSYNC. The bus master asserts TBS7 50 ns maximum after asserting TDIN for the first time. TBS7 remains asserted until 50 ns maximum after the assertion of TDIN for the last time. In each case, TBS7 can be asserted or negated as soon as the conditions for asserting TDIN are met. The assertion of TBS7 indicates the bus master is requesting another read cycle after the current read cycle.
- **Send data**—the bus slave asserts TRPLY 0 ns minimum (8000 ns maximum to avoid a bus timeout) after receiving RDIN. The bus slave asserts TREF concurrent with TRPLY if, and only if, it is a block mode device which can support another RDIN after the current RDIN. The bus slave gates TDATA<15:00> onto the bus 0 ns minimum after receiving RDIN and 125 ns maximum after the assertion of TRPLY.

NOTE

Block mode transfers must not cross 16-word boundaries.

- **Terminate input transfer**—the bus master receives stable RDATA<15:00> from 200 ns maximum after receiving RRPLY until 20 ns minimum after the negation of RDIN. (The 20 ns minimum represents total minimum receiver delays for RDIN at the slave and RDATA<15:00> at the master.) The bus master negates TDIN 200 ns minimum after receiving RRPLY.

- **Operation completed**—the bus slave negates TRPLY 0 ns minimum after receiving the negation of RDIN. If RBS7 and TREF are both asserted when TRPLY negates, the bus slave prepares for another DIN cycle. RBS7 is stable from 125 ns after RDIN is received until 150 ns after TRPLY negates. If TBS7 and RREF were both asserted when TDIN negated, the bus master asserts TDIN 150 ns minimum after receiving the negation of RRPLY and continues with the timing relationship in send data above. RREF is stable from 75 ns after RRPLY asserts until 20 ns minimum after TDIN negates. (The 0 ns minimum represents total minimum receiver delays for RDIN at the slave and RREF at the master.)

NOTE

The bus master must limit itself to not more than eight transfers unless it monitors RDMR. If it monitors RDMR, it may perform up to 16 transfers as long as RDMR is not asserted at the end of the seventh transfer.

- **Terminate bus cycle**—if RBS7 and TREF were not both asserted when TRPLY negated, the bus slave removes TDATA<15:00> from the bus 0 ns minimum and 100 ns maximum after negating TRPLY. If TBS7 and RREF were not both asserted when TDIN negated, the bus master negates TSYNC 250 ns minimum after receiving the last assertion of RRPLY and 0 ns minimum after the negation of that RRPLY.
- **Release the bus**—the DMA bus master negates TSACK 0 ns after negation of the last RRPLY. The DMA bus master negates TSYNC 300 ns maximum after it negates TSACK. The DMA bus master must remove RDATA<15:00>, TBS7, and TWTBT from the bus 100 ns maximum after clearing TSYNC.

At this point the block mode transfer is complete, and the bus arbitration logic in the CPU enables processor-generated TSYNC or issues another bus grant (TDMGO) if RDMR is asserted.

C.4.2.2 DATBO Bus Cycle

Before a block mode transfer can occur, the DMA bus master device must request control of the bus. This occurs under conventional Q22-bus protocol.

A Block mode DATBO transfer is executed as follows:

- **Address device memory**—the address is asserted by the bus master on TADDR<21:00> along with the assertion of TWTBT. The bus master asserts TSYNC 150 ns minimum after gating the address onto the bus.

- **Decode address**—the appropriate memory device recognizes that it must respond to the address on the bus.
- **Send data**—the bus master gates TDATA<15:00> along with TWTBT 100 ns minimum after the assertion of TSYNC. TWTBT is negated. The bus master asserts the first TDOUT 100 ns minimum after gating TDATA<15:00>.

NOTE

During DATBO cycles, TBS7 is undefined.

- **Receive data**—the bus slave receives stable data on RDATA<15:00> from 25 ns minimum before receiving RDOUT until 25 ns minimum after receiving the negation of RDOUT. The bus slave asserts TRPLY 0 ns minimum after receiving RDOUT. The bus slave asserts TREF concurrent with TRPLY if, and only if, it is a block mode device which can support another RDOUT after the current RDOUT.

NOTE

Block mode transfers must not cross 16-word boundaries.

- **Terminate output transfer**—the bus master negates TDOUT 150 ns minimum after receiving RRPLY.
- **Operation completed**—the bus slave negates TRPLY 0 ns minimum after receiving the negation of RDOUT. If RREF was asserted when TDOUT negated and if the bus master wants to transfer another word, the bus master gates the new data on TDATA<15:00> 100 ns minimum after negating TDOUT. RREF is stable from 75 ns maximum after RRPLY asserts until 20 ns minimum after RDOUT negates. (The 20 ns minimum represents minimum receiver delays for RDOUT at the slave and RREF at the master). The bus master asserts TDOUT 100 ns minimum after gating new data on TDATA<15:00> and 150 ns minimum after receiving the negation of RRPLY. The cycle continues with the timing relationship in receive data above.

NOTE

The bus master must limit itself to not more than 8 transfers unless it monitors RDMR. If it monitors RDMR, it may perform up to 16 transfers as long as RDMR is not asserted at the end of the seventh transfer.

- **Terminate bus cycle**—if RREF was not asserted when RRPLY negated or if the bus master has no additional data to transfer, the bus master removes data on TDATA<15:00> from the bus 100 ns minimum after negating TDOUT. If RREF was not asserted when TDOUT negated, the bus master negates TSYNC 275 ns minimum after receiving the last RRPLY and 0 ns minimum after the negation of the last RRPLY.
- **Release the bus**—the DMA bus master negates TSACK 0 ns after negation of the last RRPLY. The DMA bus master negates TSYNC 300 ns maximum after it negates TSACK. The DMA bus master must remove TDATA, TBS7, and TWTBT from the bus 100 ns maximum after clearing TSYNC.

At this point the block mode transfer is complete, and the bus arbitration logic in the CPU enables processor-generated TSYNC or issues another bus grant (TDMGO) if RDMR is asserted.

C.4.3 DMA Guidelines

The following is a list of DMA guidelines:

- Systems with memory refresh over the bus must not include devices that perform more than one transfer per acquisition.
- Bus masters that do not use block mode are limited to four DATI, four DATO, or two DATIO transfers per acquisition.
- Block mode bus masters that do not monitor BDMR are limited to eight transfers per acquisition.
- If BDMR is not asserted after the seventh transfer, block mode bus masters that do monitor BDMR may continue making transfers until the bus slave fails to assert BREF, or until they reach the total maximum of 16 transfers. Otherwise, they stop after eight transfers.

C.5 Interrupts

The interrupt capability of the Q22-bus allows an I/O device to temporarily suspend (interrupt) current program execution and divert processor operation to service the requesting device. The processor inputs a vector from the device to start the service routine (handler). Like the device register address, hardware fixes the device vector at locations within a designated range below location 001000. The vector indicates the first of a pair of addresses. The processor reads the contents of the first address, the starting address of the interrupt handler. The contents of the second address is a new processor status word (PS).

The new PS can raise the interrupt priority level, thereby preventing lower-level interrupts from breaking into the current interrupt service routine. Control is returned to the interrupted program when the interrupt handler is ended. The original interrupted program's address (PC) and its associated PS are stored on a stack. The original PC and PS are restored by a return from interrupt (RTI or RTT) instruction at the end of the handler. The use of the stack and the Q22-bus interrupt scheme can allow interrupts to occur within interrupts (nested interrupts), depending on the PS.

Interrupts can be caused by Q22-bus options or the MicroVAX CPU. Those interrupts that originate from within the processor are called traps. Traps are caused by programming errors, hardware errors, special instructions, and maintenance features.

The following Q22-bus signals are used in interrupt transactions:

Signal	Definition
BIRQ4 L	Interrupt request priority level 4
BIRQ5 L	Interrupt request priority level 5
BIRQ6 L	Interrupt request priority level 6
BIRQ7 L	Interrupt request priority level 7
BIAKI L	Interrupt acknowledge input
BIAKO L	Interrupt acknowledge output
BDAL<21:00>	Data/address lines
BDIN L	Data input strobe
BRPLY L	Reply

C.5.1 Device Priority

The Q22-bus supports the following two methods of device priority:

- Distributed arbitration — priority levels are implemented on the hardware. When devices of equal priority level request an interrupt, priority is given to the device electrically closest to the processor.
- Position-defined arbitration — priority is determined solely by electrical position on the bus. The closer a device is to the processor, the higher its priority is.

C.5.2 Interrupt Protocol

Interrupt protocol on the Q22-bus has three phases:

- Interrupt request
- Interrupt acknowledge and priority arbitration
- Interrupt vector transfer phase

The interrupt request phase begins when a device meets its specific conditions for interrupt requests. For example, the device is ready, done, or an error occurred. The interrupt enable bit in a device status register must be set. The device then initiates the interrupt by asserting the interrupt request line(s). BIRQ4 L is the lowest hardware priority level and is asserted for all interrupt requests for compatibility with previous Q22-bus processors. The level at which a device is configured must also be asserted. A special case exists for level 7 devices that must also assert level 6. The following list gives the interrupt levels and the corresponding Q22-bus interrupt request lines. For an explanation, refer to Section C.5.3.

Interrupt Level	Lines Asserted by Device
4	BIRQ4 L
5	BIRQ4 L, BIRQ5 L
6	BIRQ4 L, BIRQ6 L
7	BIRQ4 L, BIRQ6 L, BIRQ7 L

Figure C-11 shows the interrupt request/acknowledge sequence.

MA-1065-87

Figure C-11 Interrupt Request/Acknowledge Sequence

The interrupt request line remains asserted until the request is acknowledged.

During the interrupt acknowledge and priority arbitration phase, the processor acknowledges interrupts under the following conditions:

- The device interrupt priority is higher than the current PS<7:5>.
- The processor has completed instruction execution and no additional bus cycles are pending.

The processor acknowledges the interrupt request by asserting BDIN L, and 150 ns minimum later asserting BIAKO L. The device electrically closest to the processor receives the acknowledge on its BIAKI L bus receiver.

At this point, the two types of arbitration must be discussed separately. If the device that receives the acknowledge uses the four-level interrupt scheme, it reacts as follows:

- If not requesting an interrupt, the device asserts BIAKO L and the acknowledge propagates to the next device on the bus.
- If the device is requesting an interrupt, it must check that no higher-level device is currently requesting an interrupt. This is done by monitoring higher-level request lines. The following table lists the lines that need to be monitored by devices at each priority level:

Device Priority Level	Line(s) Monitored
4	BIRQ5, BIRQ6
5	BIRQ6
6	BIRQ7
7	–

In addition to asserting levels 7 and 4, level 7 devices must drive level 6. This is done to simplify the monitoring and arbitration by level 4 and 5 devices. In this protocol, level 4 and 5 devices need not monitor level 7 because level 7 devices assert level 6. Level 4 and 5 devices become aware of a level 7 request because they monitor the level 6 request. This protocol has been optimized for level 4, 5, and 6 devices, since level 7 devices are very seldom necessary.

- If no higher-level device is requesting an interrupt, the acknowledge is blocked by the device. (BIAKO L is not asserted.) Arbitration logic within the device uses the leading edge of BDIN L to clock a flip-flop that blocks BIAKO L. Arbitration is won and the interrupt vector transfer phase begins.
- If a higher-level request line is active, the device disqualifies itself and asserts BIAKO L to propagate the acknowledge to the next device along the bus.

Signal timing must be considered carefully when implementing four-level interrupts (Figure C-12).

MA-1076-87

Figure C-12 Interrupt Protocol Timing

If a single-level interrupt device receives the acknowledge, it reacts as follows:

- If not requesting an interrupt, the device asserts BIAKO L and the acknowledge propagates to the next device on the bus.
- If the device was requesting an interrupt, the acknowledge is blocked using the leading edge of BDIN L, and arbitration is won. The interrupt vector transfer phase begins.

The interrupt vector transfer phase is enabled by BDIN L and BIAKI L. The device responds by asserting BRPLY L and its BDAL<15:00> L bus driver inputs with the vector address bits. The BDAL bus driver inputs must be stable within 125 ns maximum after BRPLY L is asserted. The processor then inputs the vector address and negates BDIN L and BIAKO L. The device then negates BRPLY L and 100 ns maximum later removes the vector address bits. The processor then enters the device's service routine.

NOTE

Propagation delay from BIAKI L to BIAKO L must not be greater than 500 ns per Q22-bus slot. The device must assert BRPLY L within 10 µs maximum after the processor asserts BIAKI L.

C.5.3 Q22-bus Four-Level Interrupt Configurations

If you have high-speed peripherals and desire better software performance, you can use the four-level interrupt scheme. Both position-independent and position-dependent configurations can be used with the four-level interrupt scheme.

Figure C-13 shows the position-independent configuration. This allows peripheral devices that use the four-level interrupt scheme to be placed in the backplane in any order. These devices must send out interrupt requests and monitor higher-level request lines as described. The level 4 request is always asserted from a requesting device regardless of priority. If two or more devices of equally high priority request an interrupt, the device physically closest to the processor wins arbitration. Devices that use the single-level interrupt scheme must be modified, or placed at the end of the bus, for arbitration to function properly.

Figure C-13 Position-Independent Configuration

Figure C-14 shows the position-dependent configuration. This configuration is simpler to implement. A constraint is that peripheral devices must be inserted with the highest priority device located closest to the processor, and the remaining devices placed in the backplane in decreasing order of priority (with the lowest priority devices farthest from the processor). With this configuration, each device has to assert only its own level and level 4. Monitoring higher-level request lines is unnecessary. Arbitration is achieved through the physical positioning of each device on the bus. Single-level interrupt devices on level 4 should be positioned last on the bus.

Figure C-14 Position-Dependent Configuration

C.6 Control Functions

The following Q22-bus signals provide control functions:

Signal	Definition
BREF L	Memory refresh (also block mode DMA)
BHALT L	Processor halt
BINIT L	Initialize
BPOK H	Power OK
BDCOK H	DC power OK

C.6.1 Halt

Assertion of BHALT L for at least 25 ns interrupts the processor, which stops program execution and forces the processor unconditionally into console I/O mode.

C.6.2 Initialization

Devices along the bus are initialized when BINIT L is asserted. The processor can assert BINIT L as a result of executing a reset instruction as part of a power-up or power-down sequence. BINIT L is asserted for approximately 10 μ s when reset is executed.

C.6.3 Power Status

Power status protocol is controlled by two signals, BPOK H and BDCOK H. These signals are driven by an external device (usually the power supply).

C.7 Q22-bus Electrical Characteristics

The input and output logic levels for Q22-bus signals are given in Section C.7.1.

C.7.1 Signal Level Specifications

The signal level specifications for the Q22-bus are as follows:

Input Logic Level

TTL logical low	0.8 Vdc maximum
TTL logical high	2.0 Vdc minimum

Output Logic Level

TTL logical low	
TTL logical high	0.4 Vdc maximum
	2.4 Vdc minimum

C.7.2 Load Definition

AC loads make up the maximum capacitance allowed per signal line to ground. A unit load is defined as 9.35 pF of capacitance. DC loads are defined as maximum current allowed with a signal line driver asserted or unasserted. A unit load is defined as 210 μ A in the unasserted state.

C.7.3 120-Ohm Q22-bus

The electrical conductors interconnecting the bus device slots are treated as transmission lines. A uniform transmission line, terminated in its characteristic impedance, propagates an electrical signal without reflections. Since bus drivers, receivers, and wiring connected to the bus have finite resistance and nonzero reactance, the transmission line impedance is not uniform, and introduces distortions into pulses propagated along it. Passive components of the Q22-bus (such as wiring, cabling, and etched signal conductors) are designed to have a nominal characteristic impedance of 120 ohms.

The maximum length of interconnecting cable, excluding wiring within the backplane, is limited to 4.88 m (16 ft).

C.7.4 Bus Drivers

Devices driving the 120-ohm Q22-bus must have open collector outputs and meet the following specifications:

DC Specifications

- Output low voltage when sinking 70 mA of current is 0.7 V maximum.
- Output high leakage current when connected to 3.8 Vdc is 25 μ A (even if no power is applied, except for BDCOK H and BPOK H).
- These conditions must be met at worst-case supply temperature, and input signal levels.

AC Specifications

- Bus driver output pin capacitance load should not exceed 10 pF.
- Propagation delay should not exceed 35 ns.
- Skew (difference in propagation time between slowest and fastest gate) should not exceed 25 ns.
- Transition time (from 10% to 90% for positive transition—rise time, from 90% to 10% for negative transition—fall time) must be no faster than 10 ns.

C.7.5 Bus Receivers

Devices that receive signals from the 120-ohm Q22-bus must meet the following requirements:

DC Specifications

- Input low voltage maximum is 1.3 V.
- Input high voltage minimum is 1.7 V.
- Maximum input current when connected to 3.8 Vdc is 80 μ A (even if no power is applied).

These specifications must be met at worst-case supply voltage, temperature, and output signal conditions.

AC Specifications

- Bus receiver input pin capacitance load should not exceed 10 pF.
- Propagation delay should not exceed 35 ns.
- Skew (difference in propagation time between slowest and fastest gate) should not exceed 25 ns.

C.7.6 Bus Termination

The 120-ohm Q22-bus must be terminated at each end by an appropriate terminator, as shown in Figure C-15. This is to be done as a voltage divider with its Thevenin equivalent equal to 120 ohms and 3.4 V (nominal). This type of termination is provided by an REV11-A refresh/boot/terminator, BDV11-AA, KPV11-B, TEV11, or by certain backplanes and expansion cards.

MA-1071-87

Figure C-15 Bus Line Terminations

Each of the several Q22-bus lines (all signals whose mnemonics start with the letter B) must see an equivalent network with the following characteristics at each end of the bus:

Bus Termination Characteristic	Value
Input impedance (with respect to ground)	120 ohm +5%, -15%
Open circuit voltage	3.4 Vdc +5%
Capacitance load	Not to exceed 30 pF

NOTE

The resistive termination can be provided by the combination of two modules. (The processor module supplies 220 ohms to ground. This, in parallel with another 220-ohm card, provides 120 ohms.) Both terminators must reside physically within the same backplane.

C.7.7 Bus Interconnecting Wiring

The following sections give specific information about bus interconnecting wiring.

C.7.7.1 Backplane Wiring

The wiring that connects all device interface slots on the Q22-bus must meet the following specifications:

- The conductors must be arranged so that each line exhibits a characteristic impedance of 120 ohms (measured with respect to the bus common return).
- Crosstalk between any two lines must be no greater than 5%. Note that worst-case crosstalk is manifested by simultaneously driving all but one signal line and measuring the effect on the undriven line.
- DC resistance of the signal path, as measured between the near-end terminator and the far-end terminator module (including all intervening connectors, cables, backplane wiring, and connector-module etch) must not exceed 20 ohms.
- DC resistance of the common return path, as measured between the near-end terminator and the far-end terminator module (including all intervening connectors, cables, backplane wiring and connector-module etch) must not exceed an equivalent of 2 ohms per signal path. Thus, the composite signal return path dc resistance must not exceed 2 ohms divided by 40 bus lines, or 50 milliohms. Note that although this common return path is nominally at ground potential, the conductance must be part of the bus wiring. The specified low impedance return path must be provided by the bus wiring as distinguished from the common system or power ground path.

C.7.7.2 Intrabackplane Bus Wiring

The wiring that connects the bus connector slots within one contiguous backplane is part of the overall bus transmission line. Owing to implementation constraints, the nominal characteristic impedance of 120 ohms may not be achievable. Distributed wiring capacitance in excess of the amount required to achieve the nominal 120-ohm impedance may not exceed 60 pF per signal line per backplane.

C.7.7.3 Power and Ground

Each bus interface slot has connector pins assigned for the following dc voltages. The maximum allowable current per pin is 1.5 A. +5 Vdc must be regulated to 5% with a maximum ripple of 100 mV pp. +12 Vdc must be regulated to 3% with a maximum ripple of 200 mV pp.

- +5 Vdc — three pins (4.5 A maximum per bus device slot)
- +12 Vdc — two pins (3.0 A maximum per bus device slot)
- Ground — eight pins (shared by power return and signal return)

NOTE

Power is not bused between backplanes on any interconnecting bus cables.

C.8 System Configurations

Q22-bus systems can be divided into two types:

- Systems containing one backplane
- Systems containing multiple backplanes

Before configuring any system, three characteristics for each module in the system must be identified.

- Power consumption — +5 Vdc and +12 Vdc are the current requirements.
- AC bus loading — the amount of capacitance a module presents to a bus signal line. AC loading is expressed in terms of ac loads, where one ac load equals 9.35 pF of capacitance.
- DC bus loading—the amount of dc leakage current a module presents to a bus signal when the line is high (undriven). DC loading is expressed in terms of dc loads, where one dc load equals 210 μ A (nominal).

C-40 Q22-bus Specification

Power consumption, ac loading, and dc loading specifications for each module are included in the *Microcomputer Interfaces Handbook*.

NOTE

The ac and dc loads and the power consumption of the processor module, terminator module, and backplane must be included in determining the total loading of a backplane.

Rules for configuring single-backplane systems are as follows:

- When using a processor with 220-ohm termination, the bus can accommodate modules that have up to 20 ac loads before additional termination is required (Figure C-16). If more than 20 ac loads are included, the other end of the bus must be terminated with 120 ohms, and then up to 35 ac loads may be present.
- With 120-ohm processor termination, up to 35 ac loads can be used without additional termination. If 120-ohm bus termination is added, up to 45 ac loads can be configured in the backplane.
- The bus can accommodate modules up to 20 dc loads (total).
- The bus signal lines on the backplane can be up to 35.6 cm (14 in.) long.

MA-1072-87

Figure C-16 Single-Backplane Configuration

Rules for configuring multiple backplane systems are as follows:

- Figure C-17 shows that up to three backplanes can make up the system.
- The signal lines on each backplane can be up to 25.4 cm (10 in.) long.
- Each backplane can accommodate modules that have up to 22 ac loads. Unused ac loads from one backplane may not be added to another backplane if the second backplane loading exceeds 22 ac loads. It is desirable to load backplanes equally, or with the highest ac loads in the first and second backplanes.
- DC loading of all modules in all backplanes cannot exceed 20 loads.
- Both ends of the bus must be terminated with 120 ohms. This means the first and last backplanes must have an impedance of 120 ohms. To achieve this, each backplane can be lumped together as a single point. The resistive termination can be provided by a combination of two modules in the backplane – the processor providing 220 ohms to ground in parallel with an expansion paddle card providing 250 ohms to give the needed 120-ohm termination.

Alternately, a processor with 120-ohm termination would need no additional termination on the paddle card to attain 120 ohms in the first box. The 120-ohm termination in the last box can be provided in two ways: the termination resistors may reside either on the expansion paddle card, or on a bus termination card (such as the BDV11).

- The cable(s) connecting the first two backplanes is 61 cm (2 ft) or more in length.
- The cable(s) connecting the second backplane to the third backplane is 122 cm (4 ft) longer or shorter than the cable(s) connecting the first and second backplanes.
- The combined length of both cables cannot exceed 4.88 m (16 ft).
- The cables used must have a characteristic impedance of 120 ohms.

C-42 Q22-bus Specification

MA-1073-87

Figure C-17 Multiple Backplane Configuration

C.8.1 Power Supply Loading

Total power requirements for each backplane can be determined by obtaining the total power requirements for each module in the backplane. Obtain separate totals for +5 V and +12 V power. Power requirements for each module are specified in the *Microcomputer Interfaces Handbook*.

When distributing power in multiple backplane systems, do not attempt to distribute power through the Q22-bus cables. Provide separate, appropriate power wiring from each power supply to each backplane. Each power supply should be capable of asserting BPOK H and BDCOK H signals according to bus protocol; this is required if automatic power-fail/restart programs are implemented, or if specific peripherals require an orderly power-down halt sequence. The proper use of BPOK H and BDCOK H signals is strongly recommended.

C.9 Module Contact Finger Identification

Digital's plug-in modules all use the same contact finger (pin) identification system. A typical pin is shown in Figure C-18.

MA-1054-87

Figure C-18 Typical Pin Identification System

The Q22-bus is based on the use of quad-height modules that plug into a 2-slot bus connector. Each slot contains 36 lines (18 lines on both the component side and the solder side of the circuit board).

Slots, row A, and row B include a numeric identifier for the side of the module. The component side is designated side 1, the solder side is designated side 2, as shown in Figure C-19.

MA-1079-87

Figure C-19 Quad-Height Module Contact Finger Identification

Letters ranging from A through V (excluding G, I, O, and Q) identify a particular pin on a side of a slot. Table C-7 lists and identifies the bus pins of the quad-height module. A bus pin identifier ending with a 1 is found on the component side of the board, while a bus pin identifier ending with a 2 is found on the solder side of the board.

The positioning notch between the two rows of pins mates with a protrusion on the connector block for correct module positioning.

The dimensions for a typical Q22-bus module are represented in Figure C-20.

MA-1091-87

Figure C-20 Typical Q22-bus Module Dimensions

Table C-7 Bus Pin Identifiers

Bus Pin	Signal	Definition
AA1	BIRQ5 L	Interrupt request priority level 5.
AB1	BIRQ6 L	Interrupt request priority level 6.
AC1	BDAL16 L	Extended address bit during addressing protocol; memory error data line during data transfer protocol.
AD1	BDAL17 L	Extended address bit during addressing protocol; memory error logic enable during data transfer protocol.

Table C-7 (Cont.) Bus Pin Identifiers

Bus Pin	Signal	Definition
AE1	SSPARE1 (alternate +5B)	Special spare — not assigned or bused in Digital's cable or backplane assemblies. Available for user connection. Optionally, this pin can be used for +5 V battery (+5 B) back-up power to keep critical circuits alive during power failures. A jumper is required on Q22-bus options to open (disconnect) the +5 B circuit in systems that use this line as SSPARE1.
AF1	SSPARE2	Special spare — not assigned or bused in Digital's cable or backplane assemblies. Available for user interconnection. In the highest priority device slot, the processor can use this pin for a signal to indicate its run state.
AH1	SSPARE3 SRUN	Special spare — not assigned or bused simultaneously in Digital's cable or backplane assemblies; available for user interconnection. An alternate SRUN signal can be connected in the highest priority set.
AJ1	GND	Ground — system signal ground and dc return.
AK1	MSPAREA	Maintenance spare — normally connected together on the backplane at each option location (not a bused connection).
AL1	MSPAREB	Maintenance spare — normally connected together on the backplane at each option location (not a bused connection).
AM1	GND	Ground — system signal ground and dc return.

Table C-7 (Cont.) Bus Pin Identifiers

Bus Pin	Signal	Definition
AN1	BDMR L	DMA request — a device asserts this signal to request bus mastership. The processor arbitrates bus mastership between itself and all DMA devices on the bus. If the processor is not bus master (it has completed a bus cycle and BSYNC L is not being asserted by the processor), it grants bus mastership to the requesting device by asserting BDMGO L. The device responds by negating BDMR L and asserting BSACK L.
AP1	BHALT L	Processor halt — when BHALT L is asserted for at least 25 μ s, the processor services the halt interrupt and responds by halting normal program execution. External interrupts are ignored but memory refresh interrupts in Q22-bus operations are enabled if W4 on the M7264 and M7264-YA processor modules is removed and DMA request/grant sequences are enabled. The processor executes the ODT microcode, and the console device operation is invoked.
AR1	BREF L	Memory refresh — asserted by a DMA device. This signal forces all dynamic MOS memory units requiring bus refresh signals to be activated for each BSYNC L/BDIN L bus transaction. It is also used as a control signal for block mode DMA.

CAUTION

The user must avoid multiple DMA data transfers (burst or hot mode) that could delay refresh operation if using DMA refresh. Complete refresh cycles must occur once every 1.6 ms if required.

Table C-7 (Cont.) Bus Pin Identifiers

Bus Pin	Signal	Definition
AS1	+12 B or +5 B	+12 Vdc or +5 V battery back-up power to keep critical circuits alive during power failures. This signal is not bused to BS1 in all of Digital's backplanes. A jumper is required on all Q22-bus options to open (disconnect) the backup circuit from the bus in systems that use this line at the alternate voltage.
AT1	GND	Ground — system signal ground and dc return.
AU1	PSPARE 1	Spare — not assigned. Customer usage not recommended. Prevents damage when modules are inserted upside down.
AV1	+5 B	+5 V battery power — secondary +5 V power connection. Battery power can be used with certain devices.
BA1	BDCOK H	DC power OK — a power supply generated signal that is asserted when the available dc voltage is sufficient to sustain reliable system operation.
BB1	BPOK H	Power OK — asserted by the power supply 70 ms after BDCOK is negated when ac power drops below the value required to sustain power (approximately 75% of nominal). When negated during processor operation, a power-fail trap sequence is initiated.
BC1	SSPARE4 BDAL18 L (22-bit only)	Special spare in the Q22-bus — not assigned. Bused in 22-bit cable and backplane assemblies. Available for user interconnection.
BD1	SSPARE5 BDAL19 L (22-bit only)	CAUTION These pins may be used by manufacturing as test points in some options.
BE1	SSPARE6 BDAL20 L	In the Q22-bus, these bused address lines are address lines <21:18>. Currently not used during data time.

Table C-7 (Cont.) Bus Pin Identifiers

Bus Pin	Signal	Definition
BF1	SSPARE7 BDAL21 L	In the Q22-bus, these bused address lines are address lines <21:18>. Currently not used during data time.
BH1	SSPARE8	Special spare — not assigned or bused in Digital's cable and backplane assemblies. Available for user interconnection.
BJ1	GND	Ground — system signal ground and dc return.
BK1 BL1	MSPAREB MSPAREB	Maintenance spare — normally connected together on the backplane at each option location (not a bused connection).
BM1	GND	Ground — system signal ground and dc return.
BN1	BSACK L	This signal is asserted by a DMA device in response to the processor's BDMGO L signal, indicating that the DMA device is bus master.
BP1	BIRQ7 L	Interrupt request priority level 7.
BR1	BEVNT L	External event interrupt request — when asserted, the processor responds by entering a service routine through vector address 1008. A typical use of this signal is as a line time clock (LTC) interrupt.
BS1	+12 B	+12 Vdc battery back-up power (not bused to AS1 in all of Digital's backplanes).
BT1	GND	Ground — system signal ground and dc return.
BU1	PSPARE2	Power spare 2 — not assigned a function and not recommended for use. If a module is using -12 V (on pin AB2), and, if the module is accidentally inserted upside down in the backplane, -12 Vdc appears on pin BU1.
BV1	+5	+5 V power — normal +5 Vdc system power.
AA2	+5	+5 V power — normal +5 Vdc system power.

Table C-7 (Cont.) Bus Pin Identifiers

Bus Pin	Signal	Definition
AB2	-12	-12 V power — -12 Vdc power for (optional) devices requiring this voltage. Each Q22-bus module that requires negative voltages contains an inverter circuit that generates the required voltage(s). Therefore, -12 V power is not required with Digital's options.
AC2	GND	Ground — system signal ground and dc return.
AD2	+12	+12 V power — +12 Vdc system power.
AE2	BDOUT L	Data output — when asserted, BDOUT implies that valid data is available on BDAL<0:15> L and that an output transfer, with respect to the bus master device, is taking place. BDOUT L is deskewed with respect to data on the bus. The slave device responding to the BDOUT L signal must assert BRPLY L to complete the transfer.
AF2	BRPLY L	Reply — BRPLY L is asserted in response to BDIN L or BDOUT L and during IAK transactions. It is generated by a slave device to indicate that it has placed its data on the BDAL bus or that it has accepted output data from the bus.
AH2	BDIN L	Data input — BDIN L is used for two types of bus operations. <ul style="list-style-type: none"> • When asserted during BSYNC L time, BDIN L implies an input transfer with respect to the current bus master, and requires a response (BRPLY L). BDIN L is asserted when the master device is ready to accept data from the slave device. • When asserted without BSYNC L, it indicates that an interrupt operation is occurring. The master device must deskew input data from BRPLY L.

Table C-7 (Cont.) Bus Pin Identifiers

Bus Pin	Signal	Definition
AJ2	BSYNC L	Synchronize — BSYNC L is asserted by the bus master device to indicate that it has placed an address on BDAL<0:17> L. The transfer is in process until BSYNC L is negated.
AK2	BWTBT L	Write/byte — BWTBT L is used in two ways to control a bus cycle. <ul style="list-style-type: none"> • It is asserted at the leading edge of BSYNC L to indicate that an output sequence (DATO or DATOB), rather than an input sequence, is to follow. • It is asserted during BDOUT L, in a DATOB bus cycle, for byte addressing.
AL2	BIRQ4 L	Interrupt request priority level 4 — a level 4 device asserts this signal when its interrupt enable and interrupt request flip-flops are set. If the PS word bit 7 is 0, the processor responds by acknowledging the request by asserting BDIN L and BIAKO L.
AM2 AN2	BIAKI L BIAKO L	Interrupt acknowledge — in accordance with interrupt protocol, the processor asserts BIAKO L to acknowledge receipt of an interrupt. The bus transmits this to BIAKI L of the device electrically closest to the processor. This device accepts the interrupt acknowledge under two conditions. <ul style="list-style-type: none"> • The device requested the bus by asserting BIRQn L (where n= 4, 5, 6 or 7) • The device has the highest priority interrupt request on the bus at that time.

Table C-7 (Cont.) Bus Pin Identifiers

Bus Pin	Signal	Definition
		If these conditions are not met, the device asserts BIAKO L to the next device on the bus. This process continues in a daisy chain fashion until the device with the highest interrupt priority receives the interrupt acknowledge signal.
AP2	BBS7 L	Bank 7 select — the bus master asserts this signal to reference the I/O page (including that part of the page reserved for nonexistent memory). The address in BDAL<0:12> L when BBS7 L is asserted is the address within the I/O page.
AR2 AS2	BDMGI L BDMGO L	<p>Direct memory access grant — the bus arbitrator asserts this signal to grant bus mastership to a requesting device, according to bus mastership protocol. The signal is passed in a daisy-chain from the arbitrator (as BDMGO L) through the bus to BDMGI L of the next priority device (the device electrically closest on the bus).</p> <p>This device accepts the grant only if it requested to be the bus master (by a BDMR L). If not, the device passes the grant (asserts BDMGO L) to the next device on the bus. This process continues until the requesting device acknowledged the grant.</p> <p>CAUTION DMA device transfers must not interfere with the memory refresh cycle.</p>
AT2	BINIT L	Initialize — this signal is used for system reset. All devices on the bus are to return to a known, initial state; that is, registers are reset to zero, and logic is reset to state 0. Exceptions should be completely documented in programming and engineering specifications for the device.

Table C-7 (Cont.) Bus Pin Identifiers

Bus Pin	Signal	Definition
AU2 AV2	BDAL0 L BDAL1 L	Data/address lines — these two lines are part of the 16-line data/address bus over which address and data information are communicated. Address information is first placed on the bus by the bus master device. The same device then either receives input data from, or outputs data to, the addressed slave device or memory over the same bus lines.
BA2	+5	+5 V power — normal +5 Vdc system power.
BB2	-12	-12 V power (voltage not supplied) — -12 Vdc power for (optional) devices requiring this voltage.
BC2	GND	Ground — system signal ground and dc return.
BD2	+12	+12 V power — +12 V system power.
BE2 BF2 BH2 BJ2 BK2 BL2 BM2 BN2 BP2 BR2 BS2 BT2 BU2 BV2	BDAL2 L BDAL3 L BDAL4 L BDAL5 L BDAL6 L BDAL7 L BDAL8 L BDAL9 L BDAL10 L BDAL11 L BDAL12 L BDAL13 L BDAL14 L BDAL15 L	Data/address lines — these 14 lines are part of the 16-line data/address bus.

D

Acronyms

This appendix lists and defines the acronyms that are most frequently used in this manual.

ACRONYM	DEFINITION
ACV	Access control violation
AIE	Alarm interrupt enable
ANSI	American National Standards Institute
AP	Argument pointer
ASTLVL	Asynchronous system trap level
BBU	Battery back-up unit
BCD	Binary coded decimal
BDR	Boot and diagnostic register
BM	Byte mask
BRS	Baud rate select signals
CMCTL	CVAX memory controller chip
CPMBX	Console program mailbox
CQBIC	CVAX Q22-bus interface chip
CRC	Cyclic redundancy check
CSR	Control and status register
CSTD	Console storage transmit data
CSTS	Console storage transmit status
DEAR	DMA error address register
DIP	Dual in-line package
DM	Data mode
DMA	Direct memory access

D-2 Acronyms

ACRONYM	DEFINITION
DSE	Daylight saving enable
DSER	DMA system error register
DSSI	Digital small storage interconnect
EDITPC	EDIT packed to character string
EIA	Electronic Industries Association
EPC	Exception program counter
EPROM	Erasable programmable read-only memory
ERR	Error signal
ESP	Executive stack pointer
FP	Frame pointer
FPA	Floating-point accelerator
FPU	Floating point unit
GPR	General purpose register
ICCS	Interval clock control and status register
ICR	Interval count register
IORESET	I/O bus reset register
IPCR	Interprocessor communication register
IPL	Interrupt priority level
IPR	Internal processor register
IR	Index register
ISP	Interrupt stack pointer
ISR	Interrupt status register
KSP	Kernel stack pointer
LANCE	Local area network controller chip
LSI	Large scale integration
MAPEN	Memory management mapping enable register
MBRK	Microprogram break register
MBZ	Must be zero
MCESR	Machine check error summary register
MCS	Multinational character set
MFPR	Move from process register
MMU	Memory management unit
MOP	Maintenance operation protocol
MOS	Metal oxide semiconductor
MSER	Memory system error register
MSI	Mass storage interface
MTPR	Move to process register
NI	Network interface

ACRONYM	DEFINITION
NICR	Next interval count register
NIRAP	Network interface register address port
NIRDP	Network interface register data port
NPA	Network physical address
NXM	Nonexistent memory
P0BR	P0 base register
P1BR	P1 base register
PC	Program counter
PCB	Process control block
PCBB	Process control block base
PIE	Periodic interrupt enable
P0LR	P0 length register
P1LR	P1 length register
PMR	Performance monitor enable register
P0PT	P0 page table
P1PT	P1 page table
PROM	Programmable read only memory
PRR	Processor revision identifier register
PSL	Processor status longword
PSW	Processor status word
PTE	Page table entry
QBEAR	Q22-bus error address register
RAM	Random-access memory
RBD	Receive Buffer Descriptor
RPB	Restart parameter block
RR	Random register
RXCS	Console receiver control/status register
RXDB	Console receiver data buffer
SAVPC	Console saved PC register
SAVPSL	Console saved PSL register
SBR	System base register
SCA	System communications architecture
SCB	System control block
SCBB	System control block base
SCR	System configuration register
SID	System identification register
SIE	System identification extension
SIRR	Software interrupt request register

D-4 Acronyms

ACRONYM	DEFINITION
SISR	Software interrupt summary register
SLR	System length register
SLU	Serial line unit
SP	Stack pointer
SPT	System page table
SQWE	Square-wave enable
SR	Status register
SSC	System support chip
SSP	Supervisor stack pointer
TBCHK	Translation buffer check register
TBD	Transmit buffer descriptor
TBDATA	Translation buffer data
TBDR	Translation buffer disable register
TBIA	Translation buffer invalidate all
TBIS	Translation buffer invalidate single
TLB	Translation lookaside buffer
TNV	Translation not valid
TODR	Time-of-year register
TXCS	Console transmit control/status register
TXDB	Console transmit data buffer
UIE	Update interrupt enable
UIP	Update in progress bit
USP	User stack pointer
VLSI	Very large scale integration
VPN	Virtual page number
VRR	Vector read register
VRT	Valid RAM and time bit
VMB	Virtual memory bootstrap
XFC	Extended Function Call
ZIP	Zig-zag in-line package

Index

A

- Abort, 5–14
- Accessing the Q22-bus map registers, 3–59
- Adding to a buffer list, 3–122

B

- Backplane wiring, C–38
- BadVaddr register, 3–16
- Baud rate, 3–41
- Block mode DMA, C–18
- BOOT, 6–18
- Boot and diagnostic facility, 3–49
- Boot and diagnostic register, 3–49
- Boot devices, 4–15, 6–58
- Boot flags, 6–60
- Bootstrap
 - conditions, 4–11
- Bootstrapping, 6–58
- Break response, 3–41
- Buffer management, 3–95
- Bus cycle protocol, C–7
- Bus drivers, C–36
- Bus interconnecting wiring, C–38
- Bus receivers, C–36
- Bus termination, C–37

C

- Cache isolation, 3–23
- Cache line format, 3–23
- Cache memory, 1–5, 3–23

- Cache organization, 3–23
- Cache swapping, 3–23
- Call-back entry points, 6–82
- Cause register, 3–12
- CDAL bus to Q22-bus address translation, 3–62
- Clock functions, 1–8
- CMCTL registers, 6–79
- Collision detect routine, 3–114
- Command address specifiers, 6–14
- ! - Comment, 6–54
- CONFIGURE, 6–20
- Configuring the KN210, 2–3
- Configuring the Q22-bus map, 3–64
- Console command keywords, 6–12
- Console command qualifiers, 6–14
- Console commands, 6–18
- Console command syntax, 6–12
- Console control characters, 6–10
- Console error messages, 6–89
- Console interrupt specifications, 3–42
- Console program mailbox, 6–85
- Console receiver control/status register, 3–37
- Console receiver data buffer, 3–37
- Console registers, 3–36
- Console serial line, 3–36
- Console service, 6–9
- Console transmitter control/status register, 3–39
- Console transmitter data buffer, 3–40
- Contents of main memory, 6–79

2 Index

- Context register, 3–16
- Continue, 4–20
- CONTINUE, 6–22
- Control functions, C–34
- Coprocessor 0, 3–6
- Coprocessor 1, 3–6
- Coprocessors, 3–6
- CVAX initial power-up test, 4–2
- CVAX Physical address space map, B–4
- CVAX references, 5–28

D

- D, 4–21
- Data-stream read references, 5–29
- Data transfer bus cycles, C–6
- Data types, 5–8
- DATBI bus cycle, C–23
- DATBO bus cycle, C–24
- DEPOSIT, 6–22
- Device addressing, C–7
- Device dependent bootstrap procedures, 6–66
- Device priority, C–28
- Diagnostic and test registers, 3–139
- Diagnostic interdependencies, 6–74
- Diagnostic LED register, 3–51
- Diagnostic processor, 1–5, 5–1
- Diagnostic processor hardware detected errors, 5–24
- Diagnostics, 4–23, 6–71
- Direct memory access, C–17
- Disk and tape bootstrap procedure, 6–66
- DMA error address register, 3–69
- DMA guidelines, C–26
- DMA protocol, C–17
- DMA system error register, 3–66
- DSSI bus, 3–117
- Dump, 4–21

E

- E, 4–21
- Electrical specifications, A–1
- EntryHi and EntryLo registers, 3–9
- Environmental specifications, A–2
- Environment variables, 4–19
- Error handling, 3–70
- Error messages, 6–87
- Ethernet, 3–72
- EXAMINE, 6–24
- Exception program counter, 3–15
- Exceptions, 5–13
- Exceptions and interrupts, 5–11
- EXIT, 6–27
- External halts, 6–4
- External IPRs, B–6

F

- Fault, 5–14
- Fill, 4–21
- FIND, 6–28
- Firmware, 1–8
- Floating-point accelerator, 3–22
- Floating-Point accelerator, 1–5
- FPA instructions, 3–22
- Function switch, 4–3

G

- General exception vector, 3–21
- General purpose registers, 5–1
- Global Q22-bus address space map, B–7
- Go, 4–21

H

- H3602-SA CPU cover panel, 2–11
- Halt, C–34
- HALT, 6–29
- Halt code messages, 6–88
- Halt dispatch, 6–3
- Halt entry, 6–2

Halt entry, exit and dispatch, 6–2
 Halt exit, 6–3
 Halt protected space, 6–80
 Hardware halt procedure, 5–25
 Help, 4–22
 HELP, 6–29

I

Index register, 3–11
 Information saved on a machine
 check, 5–16
 Init, 4–22
 Initialization, C–34
 Initialization routine, 3–110
 INITIALIZE, 6–31
 Initial power-up test, 6–5
 Initiator operation, 3–121
 Instruction set, 5–8
 Instruction-stream read references,
 5–28
 Internal processor registers, 5–4
 Interprocessor communication
 register, 3–62
 Interprocessor interaction, 4–8
 Interrupt errors, 5–17
 Interrupt protocol, C–28
 Interrupts, 5–11, C–27
 Interrupt status register, 3–17
 Interval timer, 3–44
 Intrabackplane bus wiring, C–39

K

Kernel mode, 3–9
 KN210 Connectors, 2–4
 KN210 Firmware, 4–1

L

LANCE chip, 3–75
 LANCE operation, 3–109
 LANCE programming notes, 3–114
 LED codes, 6–9

Lexical conventions, 4–18
 List pointer registers, 3–138
 Load definition, C–35
 Locating a console device, 6–5
 Locating the RPB, 6–76
 Look-for-work routine, 3–111

M

Machine state on power-up, 6–77
 Main memory addressing, 3–27
 Main memory behavior on writes,
 3–28
 Main memory control and diagnostic
 status register, 3–32
 Main memory error detection and
 correction, 3–34
 Main memory error status register,
 3–28
 Main memory layout and state,
 6–77
 Main memory organization, 3–27
 Main memory system, 3–24
 Maint, 4–22
 Maintenance mode EPROM layout,
 6–80
 Maintenance mode firmware, 6–1
 Maintenance power-up operation,
 4–10
 Mass storage interface, 3–117
 Mass storage interface command
 block, 3–124
 Memory controller, 1–5
 Memory management, 5–9
 Memory management control
 registers, 5–10
 Memory management errors, 5–17
 Microcode errors, 5–18
 Mode switch set to "Normal", 6–8
 Mode switch set to "Query", 6–7
 Mode switch set to "Test", 6–6
 Module contact finger identification,
 C–43
 MOVE, 6–33

4 Index

- MS650-AA Memory modules, 1–6
- MS650-BA Memory modules, 1–5
- MSI clock control register, 3–145
- MSI command block word 0, 3–124
- MSI command block word 1, 3–125
- MSI command block word 2, 3–127
- MSI command block words 3-5,
3–128
- MSI control/status register, 3–128
- MSI diagnostic control register,
3–140
- MSI diagnostic register 0, 3–141
- MSI diagnostic register 1, 3–142
- MSI diagnostic register 2, 3–144
- MSI DSSI connection register,
3–132
- MSI DSSI control register, 3–130
- MSI DSSI timeout register, 3–136
- MSI ID register, 3–135
- MSI initiator list pointer register,
3–139
- MSI internal state registers, 3–146
- MSI link word 0, 3–121
- MSI link word 1, 3–122
- MSI target list pointer register,
3–138

N

- Network bootstrap procedure, 6–68
- Network interface, 3–72
- Network interface control and status
register, 3–78
- Network interface control and status
register 1, 3–84
- Network interface control and status
register 2, 3–84
- Network interface control and status
register 3, 3–85
- Network interface initialization
block word 0, 3–87
- Network interface initialization
block words 10,11, 3–94

- Network interface initialization
block words 1-3, 3–90
- Network interface initialization
block words 4-7, 3–91
- Network interface receive descriptor
ring, 3–96
- Network interface register address
port, 3–77
- Network interface register data port,
3–78
- Network interface transmit
descriptor ring, 3–103
- Network listening, 6–69
- NEXT, 6–35
- NI initialization block, 3–86
- NI initialization block words 8,9,
3–92
- NISA ROM, 3–74
- Nonoperating conditions greater
than 60 days, A–3
- Nonoperating conditions less than
60 days, A–2

O

- 120-Ohm Q22-bus, C–35
- Operating conditions, A–2
- Operating system restart, 6–75
- Operating systems, 4–2
- Operation switch, 4–3

P

- PFN bitmap, 6–78
- Physical specifications, A–1
- Power status, C–34
- Power supply loading, C–43
- Power-up, 6–5
- Preparing for the bootstrap, 6–62
- Primary bootstrap, VMB, 6–63
- Printenv, 4–22
- Processor revision identifier register,
3–17
- Processor state, 5–1

Processor status longword, 5-2
 Programmable timers, 3-45
 PROM bootstrap procedure, 6-67
 Public data structures, 6-85
 Public data structures and entry points, 6-80

Q

Q22-bus electrical characteristics, C-34
 Q22-bus error address register, 3-69
 Q22-bus four-level interrupt configurations, C-32
 Q22-bus interface, 1-7, 3-55
 Q22-bus interrupt handling, 3-63
 Q22-bus map base address register, 3-64
 Q22-bus map cache, 3-60
 Q22-bus map registers, 3-58
 Q22-bus signal assignments, C-3
 Q22-bus to main memory address translation, 3-56

R

R3000 console command language, 4-17
 R3000 general purpose registers, 3-2
 R3000 instruction set types, 3-3
 R3000 interval timer register, 3-42
 R3000 memory management, 3-6
 R3000 physical address space map, B-1
 R3000 RISC processor, 1-4
 R3000 sys_type environment variable, 4-10
 Random register, 3-11
 Read errors, 5-18
 Receive buffer descriptor **n** word 0, 3-98

Receive buffer descriptor **n** word 1, 3-99
 Receive buffer descriptor **n** word 2, 3-100
 Receive buffer descriptor **n** word 3, 3-101
 Receive buffer descriptors, 3-97
 Receive buffers, 3-102
 Receive DMA routine, 3-112
 Receive poll routine, 3-111
 Receive routine, 3-111
 References to processor registers and memory, 6-17
 REPEAT, 6-36
 Reserved main memory, 6-78
 Reset exception vector, 3-21
 Restoring processor state, 6-3
 RISC processor, 3-1
 ROM memory, 3-52

S

Saving processor state, 6-2
 Scatter/gather map, 6-79
 SEARCH, 6-37
 Select processor register operation, 4-9
 SET, 6-40
 Setenv, 4-22
 SHOW, 6-44
 SIE, 4-10
 Signal level specifications, C-35
 START, 6-48
 Status register, 3-15
 Switch routine, 3-110
 System configuration register, 3-65
 System configurations, C-39
 System control block, 5-21
 System identification, 5-27
 System support functions, 1-7

T

Target operation, 3-119

6 Index

TEST, 6–48
Time-of-year clock, 3–43
Time-of-year clock and timers, 3–42
Timer control registers, 3–45
Timer interrupt vector registers,
3–48
Timer interval registers, 3–47
Timer next interval registers, 3–48
Translation buffer, 5–9
Transmit buffer descriptor **n** word 0,
3–104
Transmit buffer descriptor **n** word 1,
3–105
Transmit buffer descriptor **n** word 2,
3–106
Transmit buffer descriptor **n** word 3,
3–107
Transmit buffer descriptors, 3–104
Transmit buffers, 3–109
Transmit data segment links, 3–121
Transmit DMA routine, 3–113
Transmit poll routine, 3–112

Transmit routine, 3–113
Trap, 5–13

U

UNJAM, 6–52
Unsetenv, 4–22
User mode, 3–9

V

Vector read registers, 3–18
VMB error messages, 6–91

W

Write errors, 5–19
Write references, 5–29

X

X - binary load and unload, 6–52