

PERKIN-ELMER

MODEL 3205 SYSTEM INSTRUCTION SET

Reference Manual

50-022 R00

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

The hardware description in this document is intended solely for use in operation, installation, maintenance, or repair of Perkin-Elmer equipment. Use of this document for all other purposes, without prior written approval from Perkin-Elmer is prohibited.

Any approved copy of this manual must include the Perkin-Elmer copyright notice.

Warning: This equipment generates, uses, and can radiate radio frequency energy and if not installed and used in accordance with the instructions manual, may cause interference to radio communications. It has been 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 interference when operated in a commercial environment. Operation of this equipment in a residential area is likely to cause interference in which case the user at his own expense will be required to take whatever measures may be required to correct the interference.

The Perkin-Elmer Corporation, Data Systems Group, 2 Crescent Place, Oceanport, New Jersey 07757

© 1984 by The Perkin-Elmer Corporation

Printed in the United States of America

TABLE OF CONTENTS

PREFACE

xiii

CHAPTERS

1 SYSTEM DESCRIPTION

1.1	INTRODUCTION	1-1
1.2	PROCESSOR	1-4
1.2.1	Program Status Word (PSW)	1-4
1.2.1.1	Register Set Select (R)	1-5
1.2.1.2	Condition Code (C, V, G, L)	1-6
1.2.1.3	Location Counter (LOC)	1-7
1.2.2	General Registers	1-7
1.2.3	Floating Point Registers	1-7
1.3	PROCESSOR INTERRUPTS	1-7
1.4	RESERVED MEMORY LOCATIONS	1-8
1.5	DATA FORMATS	1-9
1.5.1	Fixed Point Data	1-9
1.5.2	Floating Point Data	1-10
1.5.3	Logical Data	1-10
1.5.4	Decimal String Data	1-10
1.5.5	Alphanumeric String Data	1-10
1.6	DATA ALIGNMENT	1-11
1.7	INSTRUCTION ALIGNMENT	1-11
1.8	INSTRUCTION FORMATS	1-11
1.8.1	Branch Instruction Formats	1-14
1.8.2	Programming Examples	1-14
1.8.3	Register to Register (RR) Format	1-16
1.8.4	Short Form (SF) Format	1-17
1.8.5	Register and Indexed Storage One (RX1) Format	1-18
1.8.6	Register and Indexed Storage Two (RX2) Format	1-20
1.8.7	Register and Indexed Storage Three (RX3) Format	1-23
1.8.8	Register and Immediate Storage One (RI1) Format	1-25

CHAPTERS (Continued)

1.8.9	Register and Immediate Storage Two (RI2) Format	1-27
1.8.10	Register and Indexed Storage/Register and Indexed Storage (RXX) Format	1-29
2	SYSTEM CONTROL	
2.1	INTRODUCTION	2-1
2.2	CONFIGURATION	2-1
2.3	CONSOLE SWITCHES AND INDICATORS	2-4
2.3.1	Key-Operated Security Lock	2-4
2.3.2	Control Switches	2-4
2.4	OPERATING INSTRUCTIONS	2-5
2.4.1	Power Up	2-5
2.4.2	Entering Console Service	2-6
2.4.3	Initial Program Load (IPL)	2-6
2.5	SYSTEM TERMINAL COMMANDS	2-6
2.5.1	Select an Address and Examine (@)	2-6
2.5.2	Increment and Examine Next Location (+)	2-7
2.5.3	Decrement and Examine Prior Location (-)	2-7
2.5.4	Modify Current Location (=)	2-7
2.5.5	Examine General Register (R)	2-7
2.5.6	Modify General Register (=)	2-8
2.5.7	Examine Single Precision Floating Point Register (F)	2-8
2.5.8	Modify Single Precision Floating Point Register (=)	2-8
2.5.9	Examine Double Precision Floating Point Register (D)	2-8
2.5.10	Modify Double Precision Floating Point Register (=)	2-9
2.5.11	Examine Program Status Word (PSW) (P)	2-9
2.5.12	Modify Program Status Word (PSW) (=)	2-9
2.5.13	Execute Single Instruction (>)	2-9
2.5.14	Enter Run Mode (<)	2-9
2.6	MEMORY INITIALIZATION	2-9
2.7	SYSTEM TERMINAL PROGRAMMING INSTRUCTIONS	2-11
3	LOGICAL OPERATIONS	
3.1	INTRODUCTION	3-1
3.2	LOGICAL DATA FORMATS	3-1

CHAPTERS (Continued)

3.3	OPERATIONS	3-2
3.3.1	Boolean Operations	3-2
3.3.2	Translation	3-2
3.3.3	List Processing	3-3
3.4	LOGICAL INSTRUCTION FORMATS	3-4
3.5	LOGICAL INSTRUCTIONS	3-4
3.5.1	Load (L, LR, LI)	3-7
3.5.2	Load Immediate Short (LIS)	3-8
3.5.3	Load Complement Short (LCS)	3-9
3.5.4	Load Halfword (LH, LHI)	3-10
3.5.5	Load Address (LA)	3-11
3.5.6	Load Real Address (LRA)	3-12
3.5.7	Load Halfword Logical (LHL)	3-16
3.5.8	Load Multiple (LM)	3-17
3.5.9	Load Byte (LB, LBR)	3-18
3.5.10	Exchange Halfword Register (EXHR)	3-19
3.5.11	Exchange Byte Register (EXBR)	3-20
3.5.12	Store (ST)	3-21
3.5.13	Store Halfword (STH)	3-22
3.5.14	Store Multiple (STM)	3-23
3.5.15	Store Byte (STB, STBR)	3-24
3.5.16	Compare Logical (CL, CLR, CLI)	3-25
3.5.17	Compare Logical Halfword (CLH, CLHI)	3-27
3.5.18	Compare Logical Byte (CLB)	3-29
3.5.19	AND (N, NR, NI)	3-30
3.5.20	AND Halfword (NH, NHI)	3-31
3.5.21	OR (O, OR, OI)	3-32
3.5.22	OR Halfword (OH, OHI)	3-33
3.5.23	Exclusive-OR (X, XR, XI)	3-34
3.5.24	Exclusive-OR Halfword (XH, XHI)	3-35
3.5.25	Test Immediate (TI)	3-36
3.5.26	Test Halfword Immediate (THI)	3-37
3.5.27	Shift Left Logical (SLL, SLLS)	3-39
3.5.28	Shift Right Logical (SRL, SRLS)	3-40
3.5.29	Shift Left Halfword Logical (SLHL, SLHLS)	3-41
3.5.30	Shift Right Halfword Logical (SRHL, SRHLS)	3-42
3.5.31	Rotate Left Logical (RLL)	3-43
3.5.32	Rotate Right Logical (RRL)	3-45
3.5.33	Test and Set (TS)	3-47
3.5.34	Test Bit (TBT)	3-48
3.5.35	Set Bit (SBT)	3-49
3.5.36	Reset Bit (RBT)	3-50
3.5.37	Complement Bit (CBT)	3-51
3.5.38	Cyclic Redundancy Check (CRC12, CRC16)	3-52
3.5.39	Translate (TLATE)	3-54
3.5.40	Add To List (ATL, ABL)	3-58
3.5.41	Remove From List (RTL, RBL)	3-60

CHAPTERS (Continued)

4 BRANCHING

4.1	INTRODUCTION	4-1
4.2	OPERATIONS	4-1
4.2.1	Decision Making	4-1
4.2.2	Subroutine Linkage	4-2
4.3	BRANCH INSTRUCTION FORMATS	4-2
4.4	BRANCH INSTRUCTIONS	4-2
4.4.1	Branch on True (BTC, BTCR, BTBS, BTFS)	4-3
4.4.2	Branch on False (BFC, BFCR, BFBS, BFFS)	4-5
4.4.3	Branch and Link (BAL, BALR)	4-7
4.4.4	Branch on Index Low or Equal (BXLE)	4-9
4.4.5	Branch on Index High (BXH)	4-11
4.5	EXTENDED BRANCH MNEMONICS	4-13
4.5.1	Branch on Carry (BC, BCR, BCS)	4-15
4.5.2	Branch on No Carry (BNC, BNCR, BNCS)	4-16
4.5.3	Branch on Equal (BE, BER, BES)	4-17
4.5.4	Branch on Not Equal (BNE, BNER, BNES)	4-18
4.5.5	Branch on Low (BL, BLR, BLS)	4-19
4.5.6	Branch on Not Low (BNL, BNLR, BNLS)	4-20
4.5.7	Branch on Minus (BM, BMR, BMS)	4-21
4.5.8	Branch on Not Minus (BNM, BNMR, BNMS)	4-22
4.5.9	Branch on Plus (BP, BPR, BPS)	4-23
4.5.10	Branch on Not Plus (BNP, BNPR, BNPS)	4-24
4.5.11	Branch on Overflow (BO, BOR, BOS)	4-25
4.5.12	Branch on No Overflow (BNO, BNOR, BNOS)	4-26
4.5.13	Branch on Zero (BZ, BZR, BZS)	4-27
4.5.14	Branch on Not Zero (BNZ, BNZR, BNZS)	4-28
4.5.15	Branch (Unconditional) (B, BR, BS)	4-29
4.5.16	No Operation (NOP, NOPR)	4-30

5 FIXED POINT ARITHMETIC

5.1	INTRODUCTION	5-1
5.2	FIXED POINT DATA FORMATS	5-1
5.3	FIXED POINT NUMBER RANGE	5-2
5.4	OPERATIONS	5-2
5.5	CONDITION CODE	5-3
5.6	FIXED POINT INSTRUCTION FORMATS	5-4
5.7	FIXED POINT INSTRUCTIONS	5-4
5.7.1	Add (A, AR, AI, AIS)	5-6
5.7.2	Add Halfword (AH, AHI)	5-8

CHAPTERS (Continued)

5.7.3	Add to Memory (AM)	5-10
5.7.4	Add Halfword to Memory (AHM)	5-12
5.7.5	Subtract (S, SR, SI, SIS)	5-14
5.7.6	Subtract Halfword (SH, SHI)	5-16
5.7.7	Compare (C, CR, CI)	5-18
5.7.8	Compare Halfword (CH, CHI)	5-20
5.7.9	Multiply (M, MR)	5-22
5.7.10	Multiply Halfword (MH, MHR)	5-24
5.7.11	Divide (D, DR)	5-26
5.7.12	Divide Halfword (DH, DHR)	5-30
5.7.13	Shift Left Arithmetic (SLA)	5-33
5.7.14	Shift Left Halfword Arithmetic (SLHA)	5-35
5.7.15	Shift Right Arithmetic (SRA)	5-36
5.7.16	Shift Right Halfword Arithmetic (SRHA)	5-38
5.7.17	Convert to Halfword Value Register (CHVR)	5-39
6	FLOATING POINT ARITHMETIC	
6.1	INTRODUCTION	6-1
6.2	FLOATING POINT DATA FORMATS	6-2
6.3	FLOATING POINT NUMBER	6-3
6.3.1	Floating Point Number Range	6-4
6.3.2	Normalization	6-5
6.3.3	Equalization	6-6
6.3.4	True Zero	6-7
6.3.5	Exponent Overflow	6-8
6.3.6	Exponent Underflow	6-8
6.3.7	Guard Digits and R*Rounding	6-9
6.3.8	Conversion from Decimal	6-10
6.4	CONDITION CODE	6-11
6.5	FLOATING POINT INSTRUCTIONS	6-11
6.5.1	Load Unnormalized Floating Point (LU, LUR)	6-14
6.5.2	Load Floating Point (LE, LER, LEGR)	6-15
6.5.3	Load Positive Floating Point Register (LPER)	6-17
6.5.4	Load Complement Floating Point Register (LCER)	6-19
6.5.5	Load Multiple Floating Point (LME)	6-20
6.5.6	Load General Register from Floating Point Register (LGER)	6-21
6.5.7	Store Floating Point (STE)	6-22
6.5.8	Store Multiple Floating Point (STME)	6-23
6.5.9	Add Floating Point (AE, AER)	6-24
6.5.10	Subtract Floating Point (SE, SER)	6-26
6.5.11	Compare Floating Point (CE, CER)	6-28
6.5.12	Multiply Floating Point (ME, MER)	6-29
6.5.13	Divide Floating Point (DE, DER)	6-31
6.5.14	Fix Register (FXR)	6-33
6.5.15	Float Register (FLR)	6-35

CHAPTERS (Continued)

6.5.16	Load Unnormalized Double Precision Floating Point (LW, LWR)	6-37
6.5.17	Load Double Precision Floating Point (LD, LDR, LDGR)	6-38
6.5.18	Load Positive Double Precision Register (LPDR)	6-39
6.5.19	Load Complement Double Precision Register (LCDR)	6-40
6.5.20	Load Multiple Double Precision Floating Point (LMD)	6-41
6.5.21	Load General Registers from Double Precision Floating Point Register (LGDR)	6-42
6.5.22	Store Double Precision Floating Point (STD)	6-43
6.5.23	Store Multiple Double Precision Floating Point (STMD)	6-44
6.5.24	Add Double Precision Floating Point (AD, ADR)	6-45
6.5.25	Subtract Double Precision Floating Point (SD, SDR)	6-47
6.5.26	Compare Double Precision Floating Point (CD, CDR)	6-49
6.5.27	Multiply Double Precision Floating Point (MD, MDR)	6-50
6.5.28	Divide Double Precision Floating Point (DD, DDR)	6-52
6.5.29	Fix Register Double Precision (FXDR)	6-54
6.5.30	Float Register Double Precision (FLDR)	6-55
6.5.31	Load Single Precision Floating Point Register from Double (LED, LEDR)	6-56
6.5.32	Load Double Precision Floating Point Register from Single (LDE, LDER)	6-57
6.5.33	Store Double Precision Floating Point Register in Single Precision Memory (STDE)	6-58
7	STRING OPERATIONS	
7.1	INTRODUCTION	7-1
7.2	DECIMAL DATA FORMAT DEFINITIONS	7-1
7.2.1	Packed Decimal	7-1
7.2.2	Unpacked (Zoned) Decimal	7-2
7.3	DECIMAL AND ALPHANUMERIC STRING INSTRUCTION FORMATS	7-3
7.4	STRING INSTRUCTIONS	7-3
7.4.1	Load Packed Decimal String as Binary (LPB)	7-4
7.4.2	Store Binary as Packed Decimal String (STBP)	7-5
7.4.3	Move Translated Until (MVTU)	7-6
7.4.4	Move (MOVE, MOVEP)	7-8
7.4.5	Compare (CPAN, CPANP)	7-10
7.4.6	Pack and Move (PMV, PMVA)	7-12
7.4.7	Unpack and Move (UMV, UMVA)	7-14

CHAPTERS (Continued)

8	HIGH-SPEED DATA HANDLING INSTRUCTIONS	
8.1	INTRODUCTION	8-1
8.2	DATA HANDLING INSTRUCTION FORMATS	8-1
8.3	DATA HANDLING INSTRUCTIONS	8-1
8.3.1	Process Byte (PB)	8-2
8.3.2	Process Byte Register (PBR)	8-4
9	INPUT/OUTPUT (I/O) OPERATIONS	
9.1	INTRODUCTION AND CONFIGURATION OF INPUT/OUTPUT (I/O) SYSTEM	9-1
9.2	DEVICE CONTROLLERS	9-1
9.2.1	Device Addressing	9-1
9.2.2	Processor/Controller Communication	9-2
9.2.3	Interrupt Queuing	9-2
9.3	INTERRUPT SERVICE POINTER TABLE (ISPT)	9-2
9.4	CONTROL OF INPUT/OUTPUT (I/O) OPERATIONS	9-3
9.5	STATUS MONITORING INPUT/OUTPUT (I/O)	9-4
9.6	INTERRUPT DRIVEN INPUT/OUTPUT (I/O)	9-4
9.7	SELECTOR CHANNEL (SELCH) INPUT/OUTPUT (I/O)	9-6
9.7.1	Selector Channel (SELCH) Devices	9-6
9.7.2	Selector Channel (SELCH) Operation	9-6
9.7.3	Selector Channel (SELCH) Programming	9-7
9.8	INPUT/OUTPUT (I/O) INSTRUCTION FORMATS	9-8
9.9	INPUT/OUTPUT (I/O) INSTRUCTIONS	9-8
9.9.1	Output Command (OC, OCR)	9-9
9.9.2	Sense Status (SS, SSR)	9-10
9.9.3	Read Data (RD, RDR)	9-11
9.9.4	Read Halfword (RH, RHR)	9-12
9.9.5	Write Data (WD, WDR)	9-13
9.9.6	Write Halfword (WH, WHR)	9-14
9.9.7	Autoload (AL)	9-15
9.9.8	Simulate Channel Program (SCP)	9-17
9.10	AUTO DRIVER CHANNEL	9-18
9.11	CHANNEL COMMAND BLOCK (CCB)	9-18
9.11.1	Subroutine Address	9-19
9.11.2	Buffers	9-20
9.11.3	Translation	9-20
9.11.4	Check Word	9-21

CHAPTERS (Continued)

9.11.5	Channel Command Word (CCW)	9-21
9.11.6	Valid Channel Command Codes	9-23
9.11.7	General Auto Driver Channel Programming Procedure	9-25
10	STATUS SWITCHING AND INTERRUPTS	
10.1	INTRODUCTION	10-1
10.2	PROGRAM STATUS WORD (PSW) AND RESERVED MEMORY LOCATIONS	10-1
10.2.1	Program Status Word (PSW)	10-3
10.2.1.1	Catastrophic System Failure (CSF)	10-3
10.2.1.2	Memory Access Level Field (LVL)	10-3
10.2.1.3	Floating Point Masked Mode (FLM)	10-3
10.2.1.4	Interruptible Instruction in Progress (IIP)	10-4
10.2.1.5	Wait State (W)	10-4
10.2.1.6	Input/Output (I/O) Interrupt Mask (I)	10-5
10.2.1.7	Machine Malfunction Interrupt Enable (M)	10-5
10.2.1.8	Floating Point Underflow Interrupt Enable (FLU)	10-6
10.2.1.9	Relocation/Protection Enable (R/P)	10-6
10.2.1.10	System Queue Service (SQS) Interrupt Enable (Q)	10-6
10.2.1.11	Protect Mode Enable (P)	10-7
10.2.1.12	Register Set Select Field (R)	10-7
10.2.1.13	Condition Code (C, V, G, L)	10-8
10.2.2	Program Status Word (PSW) Location Counter (LOC)	10-8
10.2.3	Reserved Memory Locations	10-9
10.3	INTERRUPT TIMING AND PRIORITY	10-10
10.3.1	Maskable and Nonmaskable Interrupts	10-10
10.3.2	Interrupt Timing	10-12
10.3.3	Interrupt Precedence	10-12
10.3.4	Interruptible Instructions	10-13
10.4	PROCESSOR MODES	10-14
10.4.1	Console Mode	10-14
10.4.2	Run Mode	10-15
10.4.3	Single Step Mode	10-16
10.5	STATUS SWITCHING	10-17
10.5.1	Illegal Instruction Interrupt	10-18
10.5.2	Data Format Fault Interrupt	10-18
10.5.2.1	Alignment Faults	10-19
10.5.2.2	Invalid Digit Faults	10-19
10.5.3	Relocation/Protection (MAT) Fault Interrupt	10-20
10.5.4	Machine Malfunction Interrupt	10-20
10.5.4.1	Early Power Fail (EPF) Detect and Automatic Shutdown	10-22

CHAPTERS (Continued)

10.5.4.2	Power Restore	10-24
10.5.4.2.1	If the Loader Storage Unit (LSU) Is Disabled	10-24
10.5.4.2.2	If the Loader Storage Unit (LSU) Is Enabled	10-25
10.5.4.3	Noncorrectable Memory Error	10-25
10.5.4.4	Nonconfigured Memory Address	10-26
10.5.5	Input/Output (I/O) Device Interrupts	10-27
10.5.5.1	Priority Levels	10-27
10.5.5.2	Immediate Interrupt - Auto Driver Channel Operation	10-28
10.5.6	Simulated Interrupt (SINT)	10-29
10.5.7	System Queue Service (SQS) Interrupt	10-29
10.5.8	Supervisor Call (SVC) Interrupt	10-31
10.5.9	System Breakpoint Interrupt	10-31
10.5.10	Arithmetic Fault Interrupt	10-32
10.6	STATUS SWITCHING INSTRUCTIONS	10-33
10.6.1	Load Program Status Word (LPSW)	10-34
10.6.2	Load Program Status Word Register (LPSWR)	10-35
10.6.3	Exchange Program Status Register (EPSR)	10-36
10.6.4	Simulate Interrupt (SINT)	10-37
10.6.5	Supervisor Call (SVC)	10-38
10.6.6	System Breakpoint (BRK)	10-39
10.6.7	Privileged System Function (PSF)	10-40
10.6.7.1	Read Error Logger (REL)	10-41
10.6.7.2	Load Process Segment Table Descriptor (LPSTD)	10-43
10.6.7.3	Load Shared Segment Table Descriptor (LSSTD)	10-44
10.6.7.4	Store Process State (STPS)	10-45
10.6.7.5	Load Process State (LDPS)	10-46
10.6.7.6	Save Interruptible State (ISSV)	10-48
10.6.7.7	Restore Interruptible State (ISRST)	10-49
10.6.7.8	Store Byte, No Error Correction Code (ECC) (XSTB)	10-50
10.6.7.9	Reset Memory Voltage Failure (RMVF)	10-51
11	MEMORY MANAGEMENT	
11.1	INTRODUCTION	11-1
11.2	TRANSLATION FROM VIRTUAL TO REAL ADDRESS	11-2
11.3	ADDRESS SPACE	11-6
11.3.1	Virtual Address (VA)	11-6
11.3.1.1	Segment Field	11-6
11.3.1.2	Offset and Page Field	11-7
11.3.2	Selection of Virtual or Physical Addressing	11-7
11.4	SHARED AND PRIVATE SEGMENTS	11-7
11.4.1	Segment Table Descriptors (STDs) and Their Use	11-8
11.4.1.1	Format of a Segment Table Descriptor (STD)	11-8

CHAPTERS (Continued)

11.4.2	Setting the Virtual Address Space Size	11-8
11.5	SEGMENT TABLE ENTRIES (STEs)	11-9
11.5.1	Segment Table Entry (STE) Size	11-10
11.5.2	Segment Tables	11-10
11.5.3	Hardware Segment Table Entry (HSTE)	11-10
11.5.4	Software Segment Table Entry (SWSTE)	11-13
11.6	MEMORY ADDRESS TRANSLATOR (MAT) FAULTS	11-17
11.6.1	Conditions that Cause Memory Address Translator (MAT) Faults	11-17
11.6.1.1	Process Segment Table (PST) or Shared Segment Table (SST) Size Exceeded Fault	11-17
11.6.1.2	Nonpresence Fault	11-18
11.6.1.3	Access Level Fault	11-18
11.6.1.4	Access Mode Faults	11-18
11.6.1.5	Segment Limit Fault	11-19
11.6.2	Fault Precedence	11-19
11.6.3	Memory Address Translator (MAT) Fault Handling Routine	11-19
11.6.4	Reexecution of Faulting Instructions	11-20
11.6.5	Effect of System Initialization on the Memory Address Translator (MAT)	11-21
11.7	MEMORY MANAGEMENT INSTRUCTIONS	11-21
11.7.1	Load Process Segment Table Descriptor (LPSTD)	11-22
11.7.2	Load Shared Segment Table Descriptor (LSSTD)	11-23

APPENDIXES

A	OPCODE MAP	A-1
B	INSTRUCTION SUMMARY - ALPHABETICAL BY MNEMONIC	B-1
C	INSTRUCTION SUMMARY - NUMERICAL BY OPCODE	C-1
D	ARITHMETIC REFERENCES	D-1
E	INPUT/OUTPUT (I/O) REFERENCES	E-1
F	CONSOLE SERVICE ROUTINE FLOWCHART	F-1

FIGURES

1-1	Model 3205 Processor Block Diagram	1-3
1-2	Program Status Word	1-4
1-3	Register Set Numbering	1-6

FIGURES (Continued)

1-4	Instruction Formats	1-13
1-5	Sample Program	1-15
1-6	RXRX Formats	1-31
2-1	Consolette	2-1
2-2	Model 6100 Keyboard Layout	2-3
3-1	Logical Data	3-1
3-2	Translation Table Entry	3-2
3-3	Circular List Definition	3-3
3-4	Circular List	3-4
3-5	LRA Example	3-15
3-6	Flowchart for CRC Generation	3-53
3-7	List Processing Instructions	3-61
5-1	Fixed Point Data Formats	5-1
6-1	Exponent Overflow	6-8
6-2	Exponent Underflow	6-8
7-1	Packed Decimal Format	7-1
7-2	Unpacked Decimal Format	7-2
9-1	Channel Command Block	9-19
9-2	Channel Command Word	9-22
9-3	Auto Driver Channel Flowchart	9-26
10-1	Program Status Word	10-2
10-2	Schematic Diagram of Interrupt System Architecture	10-11
10-3	Machine Malfunction Status Word	10-22
11-1	Flowchart of MAT Process	11-2
11-2	PSTD and SSTD Registers	11-3
11-3	MAT Translation, Private Segment	11-4
11-4	MAT Translation, Shared Segment	11-5
11-5	Virtual Address	11-6
11-6	STE and SWSTE	11-9
11-7	Bit Representation of HSTE	11-10

TABLES

1-1	PSW BITS	1-5
1-2	RESERVED MEMORY LOCATIONS	1-8
1-3	OPERAND ABBREVIATIONS	1-12
2-1	SYSTEM TERMINAL SUPPORT COMMAND SUMMARY	2-2
4-1	DECISION TABLE	4-1

TABLE (Continued)

5-1	FIXED POINT FORMAT RELATIONS	5-2
9-1	VALID REDUNDANCY CHECKS	9-23
9-2	CHANNEL COMMAND WORD	9-24
10-1	PSW BITS	10-2
10-2	RESERVED MEMORY LOCATIONS	10-9
11-1	SEGMENT ACCESS FIELD SETTINGS	11-12
D-1	POWERS OF TWO	D-1
D-2	POWERS OF SIXTEEN	D-2
D-3	HEXADECIMAL ADDITION AND SUBTRACTION	D-3
D-4	HEXADECIMAL MULTIPLICATION AND DIVISION	D-4
D-5	MATHEMATICAL CONSTANTS	D-5
D-6	FRACTION CONVERSION	D-6
D-7	INTEGER CONVERSION	D-7
E-1	ASCII/HEXADECIMAL CONVERSION	E-1
E-2	ASCII/CARD CODE CONVERSION	E-2
E-3	STANDARD-PREFERRED ADDRESS TABLE	E-3
INDEX		Ind-1

PREFACE

This manual provides programming and operating information for the Perkin-Elmer Model 3205 System. The programmer is provided with information on the 32-bit system architecture and the unique memory management scheme, as well as a description of each instruction in the repertoire. The instruction descriptions include valuable system-related information presented in the form of programming notes and instruction examples.

Chapter 1 is a general description of the Model 3205 System, processor interrupts, registers, instruction formats and reserved memory locations. System control, including system commands, operator and programming instructions, and memory initialization, is discussed in Chapter 2. Chapter 3 is comprised primarily of the logical instruction set with a brief description of logical data formats and operations. Each instruction is outlined by its assembler notation, opcode and format, accompanied by a discussion of its operation, the status of its condition code and an example. Chapter 4 details branching operation and instructions. Chapters 5 and 6 list fixed and floating point instructions, and Chapter 7 discusses string operations and instructions. In Chapter 8, data handling instructions including the process byte register are discussed. Chapter 9 deals with the input/output (I/O) operations including status monitoring, instruction formats and the channel command block (CCB). Chapter 10 discusses the program status word (PSW), reserved memory locations, interrupt timing and priority, processor modes and status switching. The bits and function of the memory address translator (MAT) are described in Chapter 11.

Information pertaining to the system control terminal is given mainly to show how to access memory, modify locations and single step the system for troubleshooting.

For information on the contents of all Perkin-Elmer 32-bit manuals, see the 32-Bit Systems User Documentation Summary.

CHAPTER 1 SYSTEM DESCRIPTION

1.1 INTRODUCTION

The Model 3205 Processor is designed to meet the needs for low cost and reliability in a 32-bit minicomputer. The architecture has improved error recovery capabilities for those applications where fault tolerance is a necessity, and allows direct addressing of up to 4Mb of memory implemented in the MOS

Through the use of 32-bit general registers and a comprehensive instruction set, the Model 3205 Processor provides fullword data processing and direct memory addressing up to a limit of 4Mb. See Figure 1-1 for a block diagram of the system. The instruction set includes:

- Load/store halfwords, fullwords and multiple words
- Fixed point arithmetic on halfwords and fullwords
- Logical operations (AND, OR, Exclusive-OR, compare and test) on halfwords and fullwords
- Logical and arithmetic shifts and rotation on halfwords and fullwords
- Bit manipulation
- Floating point arithmetic on single (32-bit) and double (64-bit) precision operands
- List operations
- Data handling operations
- Input/output (I/O)
- Byte manipulations
- Privileged system functions
- Storage-to-storage functions
- Decimal conversion

With this enriched repertoire and direct memory addressing, coding and debugging time is reduced to a minimum.

Eight sets of sixteen 32-bit general registers are provided. Register set selection is controlled by bits in the program status word (PSW). Register to Register (RR) instructions permit operations between any of the 16 registers in the current set, eliminating redundant loads and stores. The multiple register set organization reduces the overhead that would otherwise be incurred in saving and restoring registers when responding to interrupts.

The memory address translator (MAT) provides automatic program segmentation, relocation and protection. The protect mode enables detection of privileged instructions. These two features are invaluable in process control, data communication and time-sharing operations because they prevent a running program from interfering with system integrity.

The Model 3205 System supports up to 4Mb of directly addressable MOS memory. Error correction is standard and is performed across every 16-bit halfword in memory using a 6-bit modified hamming-code. All single bit errors are detected and corrected; all double bit errors and most multiple bit errors are detected. The memory error logger indicates the location of the latest faulty memory chip.

The Model 3205 System microcode implements an autodriver channel that automatically acknowledges all I/O interrupts and performs much of the required overhead before activating an interrupt service routine (ISR) if enabled. The autodriver channel can perform data transfers with character translation, longitudinal or cyclic redundancy checking (CRC), and data buffer chaining transparent to the user.

8059




Figure 1-1 Model 3205 Processor Block Diagram

TABLE 1-1 PSW BITS

BIT	MNEMONIC	MEANING
0	CSF	Must be zero; IF SET, CATASTROPHIC SYSTEM FAILURE
1:9		Unused; must be zero
10:11	LVL	Memory access level
12		Reserved; must be zero
13	FLM	Floating point arithmetic masked mode
14	IIP	Interruptible instruction in progress
15		Reserved; must be zero
16	W	Wait state
17	I	I/O interrupt mask
18	M	Machine malfunction interrupt mask
19	FLU	Floating point arithmetic underflow mask
20		Reserved; must be zero
21	R/P	Relocation/protection interrupt mask
22	Q	System queue interrupt mask
23	P	Protect mode
24:27	R	Register set select bits
28:31	C,V,G,L	Condition code
32:39		Reserved; must be zero
40:63		Program address (LOC)

1.2.1.1 Register Set Select (R)

Bits 24:27 of the PSW are used to designate the current register set. Register sets are numbered 0 through 15. The processor has eight sets of general registers (see Figure 1-3).

REGISTER SET NUMBER	DESIGNATION
0	RESERVED FOR INTERRUPTS
1 2 3 4 5 6	MAY BE ALLOCATED BY THE OS FOR GENERAL PURPOSE USE
7 8 9 10 11 12 13 14	UNIMPLEMENTED SETS
15	GENERAL PURPOSE

Figure 1-3 Register Set Numbering

1.2.1.2 Condition Code (C, V, G, L)

Bits 28:31 of the PSW contain the condition code. As part of the execution of certain instructions, the state of the condition code can be changed to indicate the nature of the result. Not all instructions affect the condition code. The state of the condition code can be tested with conditional branch instructions. Each bit in the condition code is set if the corresponding condition occurred as a result of the last instruction that affected the condition code. The normal interpretation of these bits is:

C	V	G	L
1	0	0	0
0	1	0	0
0	0	1	0
0	0	0	1

Arithmetic carry, borrow or shifted carry
 Arithmetic overflow
 Greater than zero
 Less than zero

1.2.1.3 Location Counter (LOC)

The LOC contains the address of the instruction currently being executed by the processor and points to that instruction until it has successfully completed execution. Once this execution is completed, the LOC is incremented by 2, 4, 6, 8, 10 or 12 (depending upon the instruction executed), and the next instruction is fetched. In the case of a branch instruction, the LOC is loaded with the address to which control is being transferred, and the next instruction is fetched from that address.

If an instruction is not successfully completed due to a fault or other interrupting condition, the LOC contains the address of the faulting or interrupted instruction. When a program interruption is due to an incorrect branch address, the LOC contains the branch address and not the location of the branch instruction.

1.2.2 General Registers

The processor has eight register sets, numbered 0 through 6 and 7 through 15 (see Figure 1-3). There are 16 registers in each set and each register is 32 bits wide. Register set selection is determined by the states of bits 24:27 of the current PSW. Registers 1 through 15 of any set can be used as index registers.

When an interrupt occurs, the processor loads pertinent information into preselected registers of the register set selected by the new PSW. See Chapter 10 for details of this operation.

1.2.3 Floating Point Registers

There are eight single precision floating point registers, each 32 bits wide. These registers are identified by the even numbers 0 through 14.

There are eight double precision floating point registers, each 64 bits wide. These registers are also identified by the even numbers 0 through 14 and are separate from the single precision floating point registers. Floating point operations must always specify the registers with even numbers.

1.3 PROCESSOR INTERRUPTS

The PSW that is loaded in the processor at any point in time is called the current PSW. If either the status word or both the LOC and status word are changed, a status switch is said to have occurred. This status switch can be caused explicitly by executing special instructions, or it can be forced to occur by an interrupt or fault. At the time of a status switch, the current PSW that is saved is called the old PSW. The PSW that replaces the current PSW is the new PSW.

Interrupt conditions cause the entire PSW to be replaced by a new PSW, thus breaking the usual sequential flow of instruction execution. When an interrupt condition occurs, the processor saves its current PSW either in memory or in a pair of general registers belonging to the register set selected by the new PSW. The processor loads information related to the interrupt condition in other registers of this same set. A new PSW is loaded from a memory location reserved for the specific interrupt condition. The immediate interrupt is an exception to the rule. In this case, the status portion of the new PSW, bits 0:31, is forced to a preset value, and the LOC is loaded from a memory location reserved for that interrupting device. See Chapter 10 for details of interrupt processing.

1.4 RESERVED MEMORY LOCATIONS

Physical memory locations X'0' through X'2CF' are called reserved memory locations. These locations contain the various new PSWs and other information needed to handle interrupts, as shown in Table 1-2.

TABLE 1-2 RESERVED MEMORY LOCATIONS

LOCATION	MEANING
X'000000'-X'00001F'	Reserved; must be zero
X'000020'-X'000027'	Machine malfunction interrupt old PSW
X'000028'-X'00002B'	Used by console service microcode
X'00002C'-X'00002F'	Machine malfunction load memory (LM) block start address
X'000030'-X'000037'	Illegal instruction interrupt new PSW
X'000038'-X'00003F'	Machine malfunction interrupt new PSW
X'000040'-X'000043'	Machine malfunction status word
X'000044'-X'000047'	Machine malfunction virtual (program) address
X'000048'-X'00004F'	Arithmetic fault interrupt new PSW
X'000050'-X'00007F'	Bootstrap loader and device definition table
X'000080'-X'000083'	System queue pointer
X'000084'-X'000087'	Power fail save area pointer

TABLE 1-2 RESERVED MEMORY LOCATIONS (Continued)

LOCATION	MEANING
X'000088'-X'00008F'	System queue service interrupt new PSW
X'000090'-X'000097'	Relocation/protection interrupt new PSW
X'000098'-X'00009B'	Supervisor call (SVC) new PSW <u>status</u>
X'00009C'-X'0000BB'	SVC new PSW <u>LOC</u> values (<u>16</u> halfwords)
X'0000BC'-X'0000BF'	Reserved; must be zero
X'0000C0'-X'0000C7'	Reserved; must be zero
X'0000C8'-X'0000CF'	Data format fault new PSW
X'0000D0'-X'0002CF'	Interrupt service pointer table
X'0002D0'-X'0004CF'	Expanded interrupt service pointer table
X'0004D0'-X'0008CF'	Expanded interrupt service pointer table

These reserved locations play an important role in both interrupt and I/O processing (see Chapters 9 and 10).

All LOC values are subject to MAT relocation if the new PSW enables the MAT (bit 21=1). All other pointers contain absolute addresses not subject to MAT relocation.

1.5 DATA FORMATS

The processor performs logical operations on single bits, bytes, halfwords, fullwords and doublewords. This data can represent a fixed point number, a floating point number, logical information, a bit or byte array, or a decimal or alphanumeric byte string.

1.5.1 Fixed Point Data

Fixed point arithmetic operands can be either halfwords or fullwords. In fullword multiply and divide operations, doubleword operands are manipulated. Fixed point data is treated as 15-bit signed integers in the halfword format. Positive numbers are expressed in true binary form with a sign bit of 0. Negative numbers are represented in two's complement form with a sign bit of 1. The numerical value of zero is represented with all bits 0. See Chapter 5 for details of fixed point data representation.

In fixed point arithmetic and logical operations between a fullword register and a halfword operand, the halfword operand is expanded to a fullword by propagating the most significant bit into the high-order bits before the operation is started. This permits the use of halfword-to-fullword operations with consistent results and provides space economy, since small values do not require fullword locations.

Arithmetic operations on fixed point halfword quantities can produce results not entirely consistent with those obtained in a 16-bit processor. If this problem exists, the Convert to Halfword Value Register (CHVR) instruction can be used to adjust the result and the condition code, making them consistent with the same operations in a 16-bit processor.

1.5.2 Floating Point Data

A floating point number consists of a 7-bit exponent in excess-64 notation and a signed fraction. The quantity expressed by this number is the product of the fraction and the number 16 raised to the power represented by the exponent. Each floating point value requires a fullword or a doubleword, of which eight bits are used for the sign and exponent. The remaining bits are used for the fraction. See Chapter 6 for details of floating point data representation.

Floating point operations take place between the contents of two floating point registers, a floating point operand contained in a fullword or doubleword in memory, or a general register or pair of general registers.

1.5.3 Logical Data

Logical operations manipulate bytes, halfwords and fullwords. In addition, it is possible to perform logical operations on single bits located in bit arrays. See Chapter 3 for details of logical data representation.

1.5.4 Decimal String Data

Decimal strings are strings of consecutive bytes in memory that begin and end on byte boundaries. Information contained in a decimal string can represent packed or unpacked decimal data. See Chapter 7 for details of decimal data formats and operations.

1.5.5 Alphanumeric String Data

Alphanumeric strings are strings of consecutive bytes in memory that begin and end on byte boundaries. Information contained in an alphanumeric string can represent any character stream including decimal string data. See Chapter 7 for details of alphanumeric string data format and operations.

1.6 DATA ALIGNMENT

The following discussion is unique to the Model 3205 System implementation and is presented for information only. Any program that misuses a processor feature by taking advantage of a peculiarity of one implementation cannot work on a different implementation.

Locations in main memory are numbered consecutively, beginning at address '000000'.

Bytes of information are addressed by their specific hexadecimal address. Two bytes form a halfword. Halfwords have an even address, which is the address of the lower addressed byte in the pair. Two halfwords comprise a fullword. A fullword address is a multiple of four (four bytes) and is the address of the lower addressed halfword in the pair. A data format fault is generated if a fullword access is directed to an address that has bit 30 or 31 set, or if a halfword access is directed to an address that has bit 31 set.

The common assembly language (CAL) assembler generates an error flag if it sees halfword operations directed to an odd byte address or if it sees fullword operations directed to an address other than a fullword address.

1.7 INSTRUCTION ALIGNMENT

User level instructions are always aligned on halfword boundaries. Any halfword address is valid regardless of the length of the instruction word. The CAL assembler generates boundary errors if the assembled LOC for an instruction becomes odd. At the machine level, an attempt to make the instruction LOC odd by branching or causing a status switch results in a data format fault.

1.8 INSTRUCTION FORMATS

Instruction formats provide a concise method of representing required operations for easy interpretation by the processor. Figure 1-4 shows the eight basic formats. Table 1-3 is a list of abbreviations and their meanings as used in Figure 1-4.


TABLE 1-3 OPERAND ABBREVIATIONS

ABBREVIATION	MEANING
OP	Operation code
R1	First operand register
R2	Second operand register
N	A 4-bit immediate value
X2	Second operand, single index register
D2	Second operand displacement
FX2	Second operand, first index register
SX2	Second operand, second index register
A2	Second operand, direct address
I2	Second operand, immediate value
L1	Specifies the length of the first operand
L2	Specifies the length of the second operand
OPMOD	Specifies a particular instruction within the class specified by OP
ADD1	The effective first operand address
ADD2	The effective second operand address


Many instructions can be expressed in two or more formats. This feature provides flexibility in data organization and instruction sequencing. When working with the CAL assembler, it is unnecessary to specify the instruction format. The assembler selects the most economical format and supplies the required bits in the machine code. When double indexing is implied, the assembler always chooses the RX3 format.

557-1


REGISTER-TO-REGISTER (RR)




SHORT FORMAT (SF)




REGISTER AND INDEXED STORAGE (RX1)




REGISTER AND INDEXED STORAGE 2 (RX2)




REGISTER AND INDEXED STORAGE 3 (RX3)



REGISTER AND IMMEDIATE STORAGE 1 (RI1)



REGISTER AND IMMEDIATE STORAGE 2 (RI2)



REGISTER AND INDEXED STORAGE, REGISTER AND INDEXED STORAGE (RXXR)




Figure 1-4 Instruction Formats

1.8.1 Branch Instruction Formats

Branch instructions use the Register-to-Register (RR), the Short Form (SF) and all variations of the Register and Indexed Storage (RX) formats. In the conditional branch instructions, however, the R1 field does not specify a register; instead, it contains a mask value (labeled M1 in the instruction descriptions). This mask value is tested with the condition code. The CAL assembler provides a series of extended branch mnemonics, making it possible to identify a conditional branch without specifying the mask value explicitly.

1.8.2 Programming Examples

Each of the following examples refers to the sample assembly language program shown in Figure 1-5. Note the use of symbolic equivalents for general registers. The machine code generated and the result of each instruction are dependent upon the physical and logical placement of the instructions, respectively.


PROG= S3200 ASSEMBLED BY CAL 03-066R05-01 (32-BIT)

		1	S3200	PROG	SERIES 3200 INSTRUCTION FORMAT EXAMPLES	
		2		CROSS		
		3		NORX3		
	0000 0005	5	R5	EQU	5	GENERAL REGISTER 5
	0000 0006	6	R6	EQU	6	GENERAL REGISTER 6
	0000 0007	7	R7	EQU	7	GENERAL REGISTER 7
	0000 0008	8	R8	EQU	8	GENERAL REGISTER 8
	0000 0009	9	R9	EQU	9	GENERAL REGISTER 9
	0000 000A	10	R10	EQU	10	GENERAL REGISTER 10
	0000 000B	11	R11	EQU	11	GENERAL REGISTER 11
000000I	245E	13	SF	LIS	R5,R14	(R5) = '0000000E'
000002I	0865	15	RR	LR	R6,R5	(R6) = '0000000E'
000004I	4050 1000	17	RX1.EX1	STH	R5,X*1000'	(X*1000') = X*000E'
000008I	4056 0FF2	19	RX1.EX2	STH	R5,X*0FF2'(R6)	(X*1000') = X*000E'
00000CI	4050 8004 =000014I	21	RX2.EX1	STH	R5,LOC1	(LOC1) = X*000E'
000010I	4300 8004 =000018I	22		B	R11.EX1	
000014I	0000 0000	23	LOC1	DC	F*0'	TWO HALFWORDS OF STORAGE
000018I	C890 8000	25	RI1.EX1	LHI	R9,X*8000'	(R9) = Y*FFFF8000'
00001CI	C895 8000	27	RI1.EX2	LHI	R9,X*8000'(R5)	(R9) = Y*FFFF800F'
000020I	F8A0 0000 8000	29	RI2.EX1	LI	R10,X*8000'	(R10) = Y*00008000'
000026I	F8BA 0001 7FFE	31	RI2.EX2	LI	R11,Y*17FFE'(R10)	(R11) = Y*00017FFE'
00002CI	4050 FFE4 =000014I	33	RX2.EX2	STH	R5,LOC1	(LOC1) = X*000E'
000030I	4056 FFD2 =000006I	35	RX2.EX3	STH	R5,LOC1-14(R6)	(LOC1) = X*000E'
000034I	5870 4001 0000	37	RX3.EX1	L	R7,Y*10000'	(R7) = (Y*010000')
00003AI	5885 4601 FFE4	39	RX3.EX2	L	R8,Y*20300'-29(R5,R6)	(R8) = (Y*020000')
000040I	4300 FFBC =000000I	40		B	SF	
000044I		42		END		

LOCATION COUNTER OBJECT INFORMATION STATEMENT NUMBER LABEL OP-CODE OPERANDS COMMENTS

Figure 1-5 Sample Program

1.8.3 Register to Register (RR) Format



In this 16-bit format, bits 0:7 contain the operation code; bits 8:11 contain the R1 field; and bits 12:15 contain the R2 field. In most RR instructions, the register specified by R1 contains the first operand, and the register specified by R2 contains the second operand.


Example:

<u>Machine Code</u>	<u>Label</u>	<u>Assembler Notation</u>
0865	RR	LR R6,R5

Diagram illustrating the bit fields for the example instruction:

- Second operand (bits 8-11)
- First operand (bits 12-15)
- Load Register (LR) instruction opcode (bits 0-7)

1.8.4 Short Form (SF) Format




This 16-bit format provides space economy when working with small values. Bits 0:7 contain the operation code; bits 8:11 contain the R1 field; and bits 12:15 contain the N field. In arithmetic and logical operations, the register specified by R1 contains the first operand. The N field contains a 4-bit immediate value used as the second operand.

Example:

Machine Code	Label	Assembler Notation
<u>245E</u>	SF	LIS R5,14

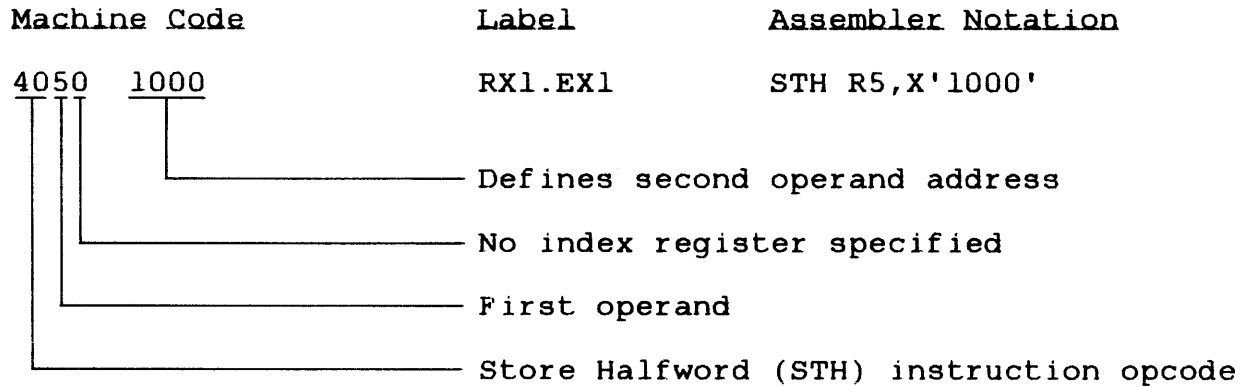
1.8.5 Register and Indexed Storage One (RX1) Format

563




This is a 32-bit format in which bits 0:7 contain the operation code; bits 8:11 contain the R1 field; bits 12:15 contain the X2 field; bits 16 and 17 must be zero; and bits 18:31 contain the D2 field. In general, the register specified by R1 contains the first operand. The second operand is located in memory at the address obtained by adding the contents of the second operand index register (specified by X2) and the 14-bit absolute address contained in the D2 field.

Example:



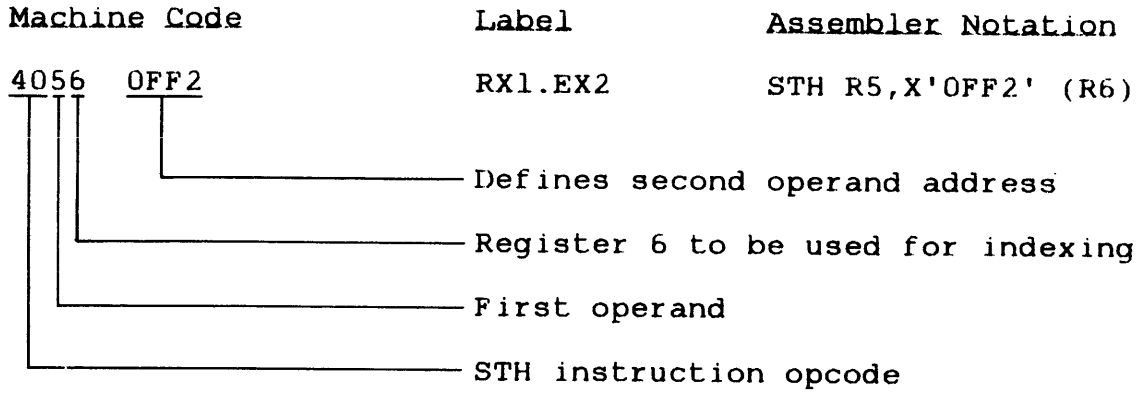
The second operand address is calculated as follows:

564-2




No indexing is specified; therefore, the second operand address is X'1000'.

Example:



The second operand address is calculated as follows:


565-1



Second operand address = contents of D2 field + contents of index register 6 (see Figure 1-5)
 = X'OFF2' + Y'0000000E'
 = Y'00001000'

1.8.6 Register and Indexed Storage Two (RX2) Format

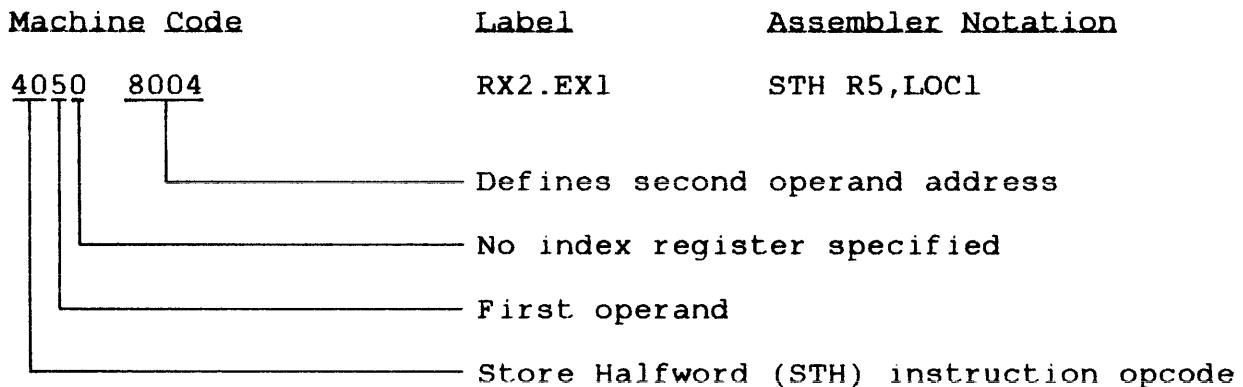
566



This format provides relative addressing capability in a 32-bit instruction word. Bits 0:7 contain the operand code; bits 8:11 contain the R1 specification; bits 12:15 contain the X2 specification; bit 16 must always be one; and bits 17:31 contain the relative displacement, D2.


In the RX2 format, the register specified by R1 contains the first operand. The address of the second operand, in memory, is calculated by adding the value contained in the incremented LOC (the address of the next sequential instruction) and the sum of (1) the 32-bit representation of the 15-bit signed number contained in the D2 field, and (2) the contents of the index register specified by X2. Negative numbers in the D2 field are expressed in two's complement notation.

Example 1:



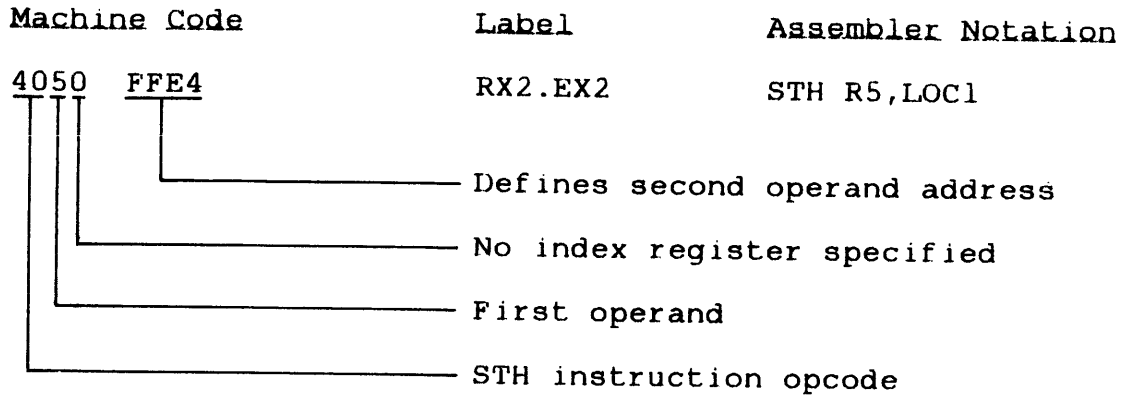
The second operand address is calculated as follows:

567-1




Second operand address = 32-bit expansion of contents of D2 field
 + contents of incremented LOC (see Figure 1-5)
 = Y'00000004' + Y'00000010'
 = Y'00000014'

Example 2:



The second operand address is calculated as follows:




Second operand address = 32-bit expansion of contents of D2 field
 + contents of incremented LOC (see
 Figure 1-5)
 = Y'FFFFFFE4' + Y'00000030'
 = Y'00000014'

Example 3:

Machine Code	Label	Assembler Notation
4056 FFD2	RX2.EX3	STH R5,LOC1-14 (R6)

The second operand address is calculated as follows:


569-1



Second operand address = 32-bit expansion of D2 field + contents of incremented LOC + contents of index register 6 (see Figure 1-5)
 = Y'FFFFFFD2' + Y'00000034' + Y'0000000E'
 = Y'00000014'

1.8.7 Register and Indexed Storage Three (RX3) Format

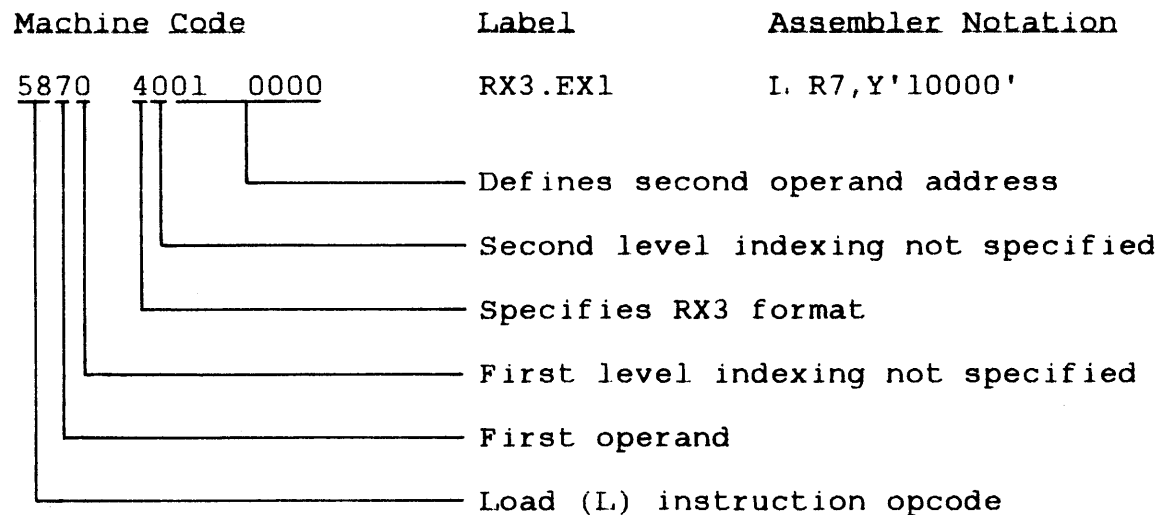
570



This is a 48-bit format in which double indexing is permitted. Bits 0:7 contain the operation code; bits 8:11 contain the R1 specification; bits 12:15 contain the first index specification, FX2; bit 16 must be zero; bit 17 must be one; bits 18 and 19 must be zero; bits 20:23 contain the second index specification, SX2; and bits 24:47 contain a 24-bit address, A2. Second level indexing is allowed even if first level indexing is not specified.


In general, the first operand is contained in the register specified by R1. The second operand is located in memory. Its memory address is obtained by adding the contents of the first index register and the contents of the second index register, then adding to this result the contents of the A2 field.

Example 1:



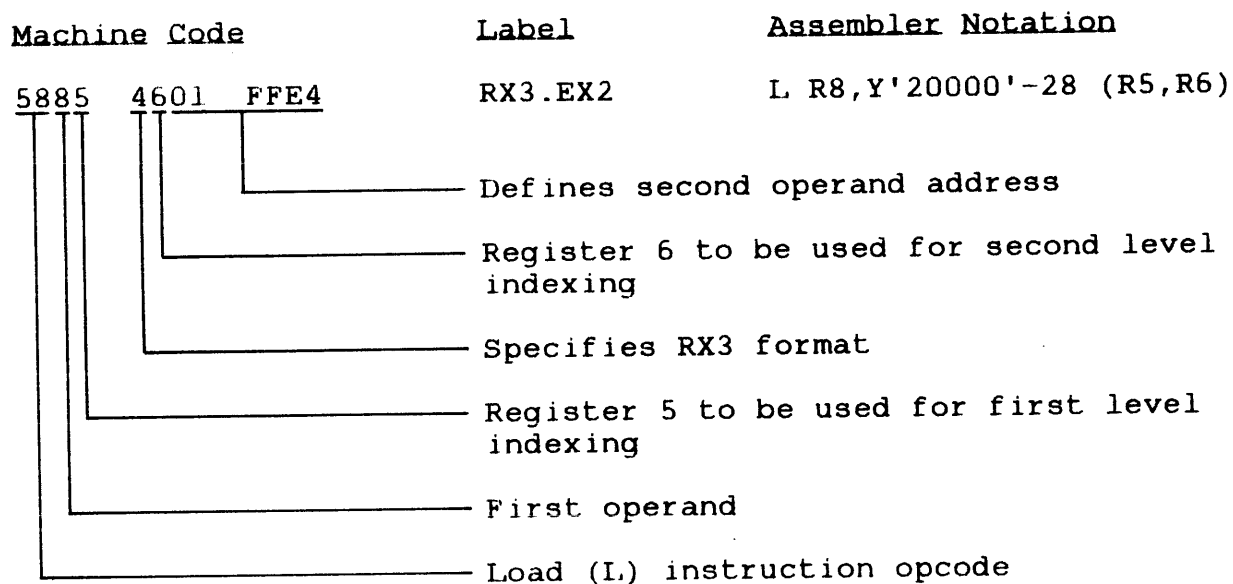
The second operand address is calculated as follows:

571-1




Second operand address = contents of A2 field
 = Y'00010000'

Example 2:



The second operand address is calculated as follows:


572-1



Second operand address = contents of A2 field + contents of index register 6 + contents of index register (see Figure 1-5)
 = Y'0001FFE4' + Y'0000000E' + Y'0000000E'
 = Y'00020000'

1.8.8 Register and Immediate Storage One (RI1) Format

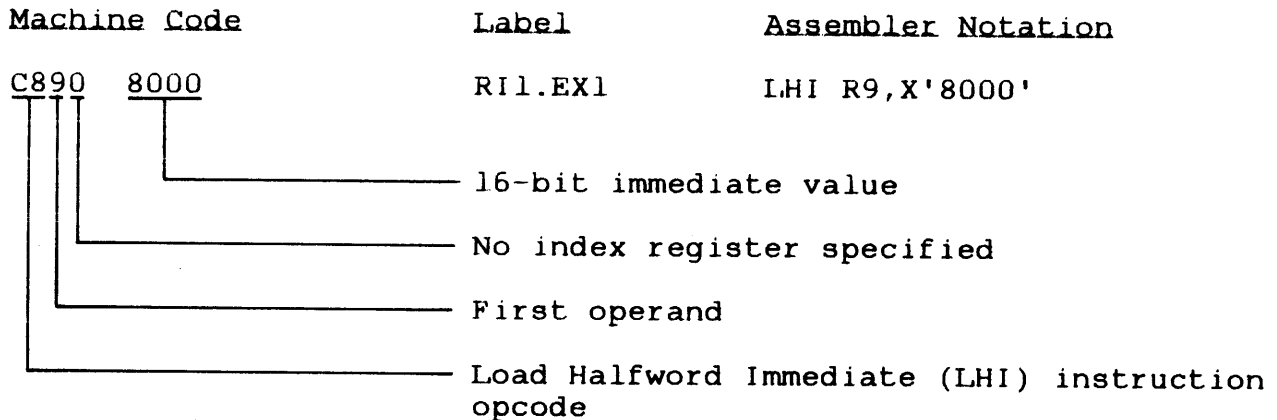
573




This format represents a 32-bit instruction word. Bits 0:7 contain the operand code; bits 8:11 contain the R1 specification; bits 12:15 contain the X2 specification; and bits 16:31 contain the 16-bit immediate value, I2.

In this format, the register specified by R1 contains the first operand. The 32-bit effective second operand is obtained by adding together the 32-bit representation of the signed 16-bit value contained in the I2 field and the contents of the register specified by X2.

Example 1:

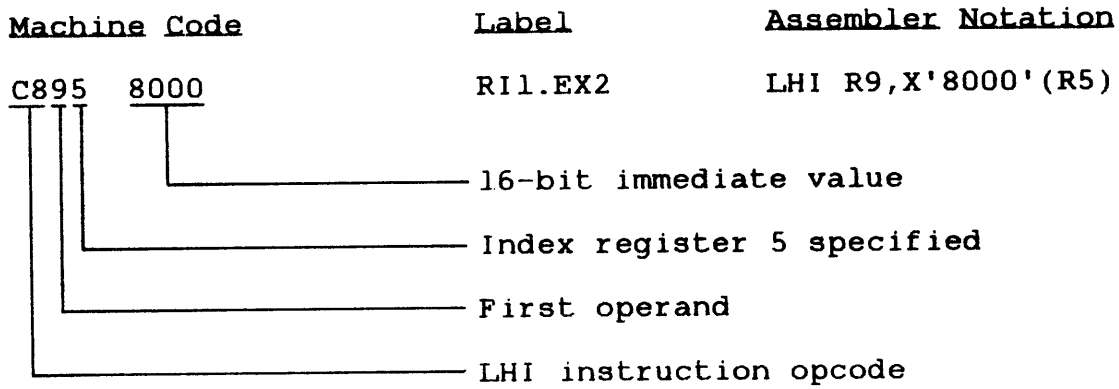


The second operand is calculated as follows:




Second operand = 32-bit representation of X'8000'
= Y'FFFF8000'

Example 2:




The second operand is calculated as follows:



Second operand = 32-bit representation of X'8000' + the contents of the index register 5 (see Figure 1-5)
 = Y'FFFF8000' + Y'0000000E'
 = Y'FFFF800E'

1.8.9 Register and Immediate Storage Two (RI2) Format

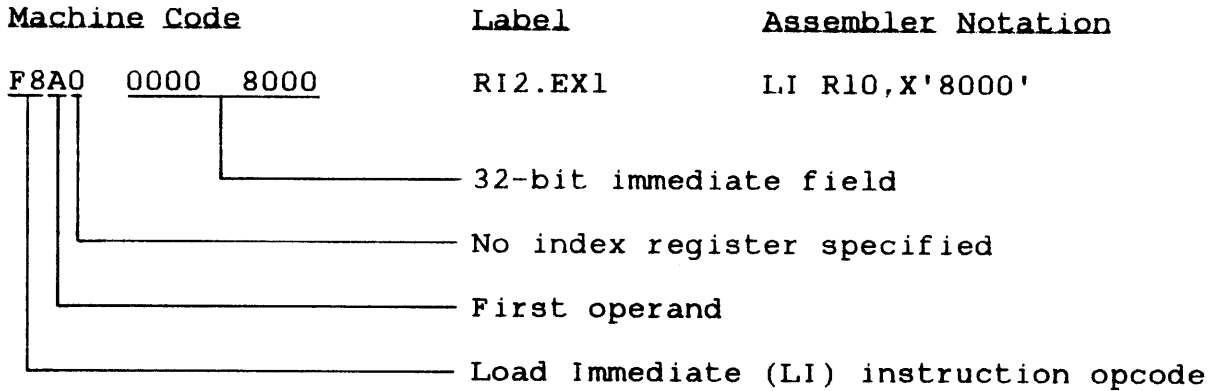
576



This is a 48-bit instruction format. Bits 0:7 contain the operation code; bits 8:11 contain the R1 specification; bits 12:15 contain the X2 specification; and bits 16:47 contain the 32-bit immediate value, I2.


The first operand is contained in the register specified by R1. The second operand is obtained by adding the contents of the index register, specified by X2, and the 32-bit immediate value contained in the I2 field.

Example 1:



The second operand is calculated as follows:

577-1



Second operand = contents of I2 field
 = Y'00008000'


Example 2:

Machine Code	Label	Assembler Notation
F8BA 0001 7FFE	RI2.EX2	LI R11,Y'17FFE' (R10)

32-bit immediate field
Specifies index register 10
First operand
LI instruction opcode

The second operand is calculated as follows:

578-1



Second operand = contents of I2 field + contents of index register 10 (see Figure 1-5)
= Y'00017FFE' + Y'00008000'
= Y'0001FFFE'

1.8.10 Register and Indexed Storage/Register and Indexed Storage (RXXR) Format

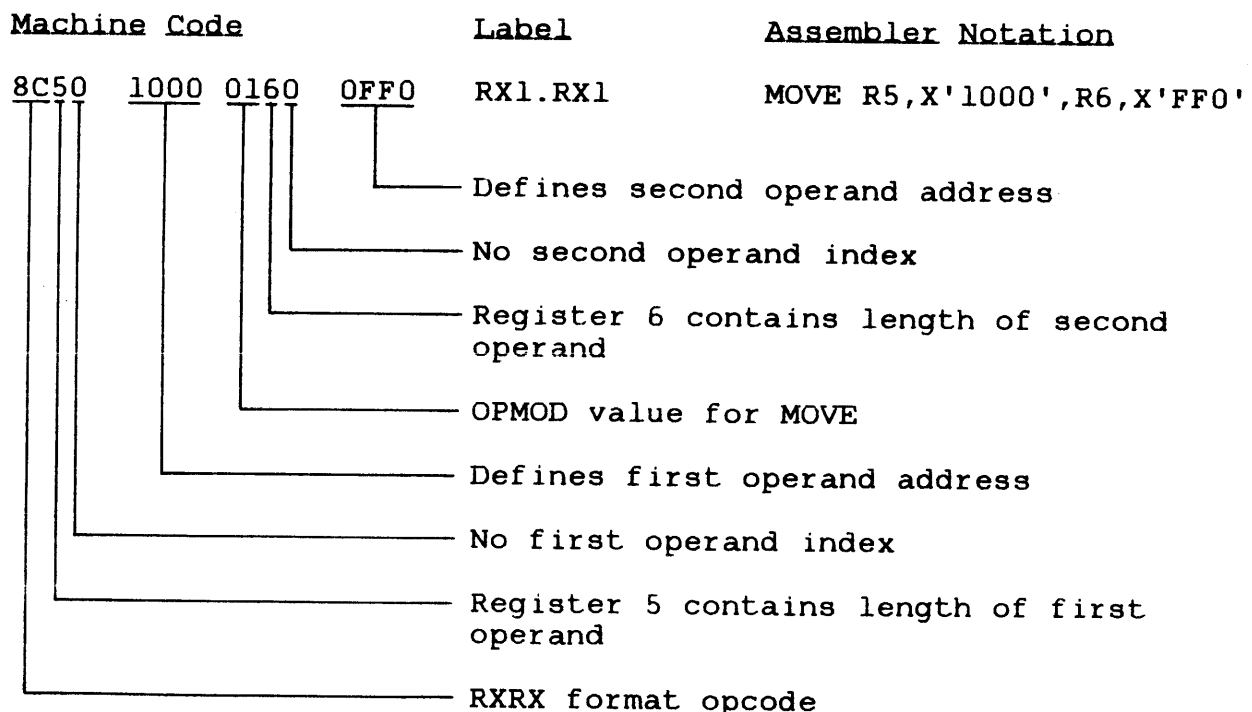
The RXXR format resembles a pair of adjacent RX format instructions, but represents only one instruction (see Figures 1-4 and 1-6). Each member of the instruction pair can have any one of the standard RX formats. For example, the first member might be RX1 and the second member might be RX3, resulting in a 10-byte instruction. The particular RX format chosen by the assembler for one member is independent of that chosen for the other; thus, the instruction can require 8, 10 or 12 bytes.

OP contains the operation code that defines the RXXR instruction class. The actual operation to be performed is defined by the OPMOD field.

The L1 field specifies the length of the first operand string. If bit 0 of OPMOD is set, L1 is the length with a maximum value of 15. If bit 0 of OPMOD is zero, the general register specified by L1 contains the length. The L2 field specifies the length of the second operand string. If bit 1 of OPMOD is set, this field contains the length with a maximum value of 15. If bit 1 of OPMOD is zero, the general register specified by L2 contains the length.

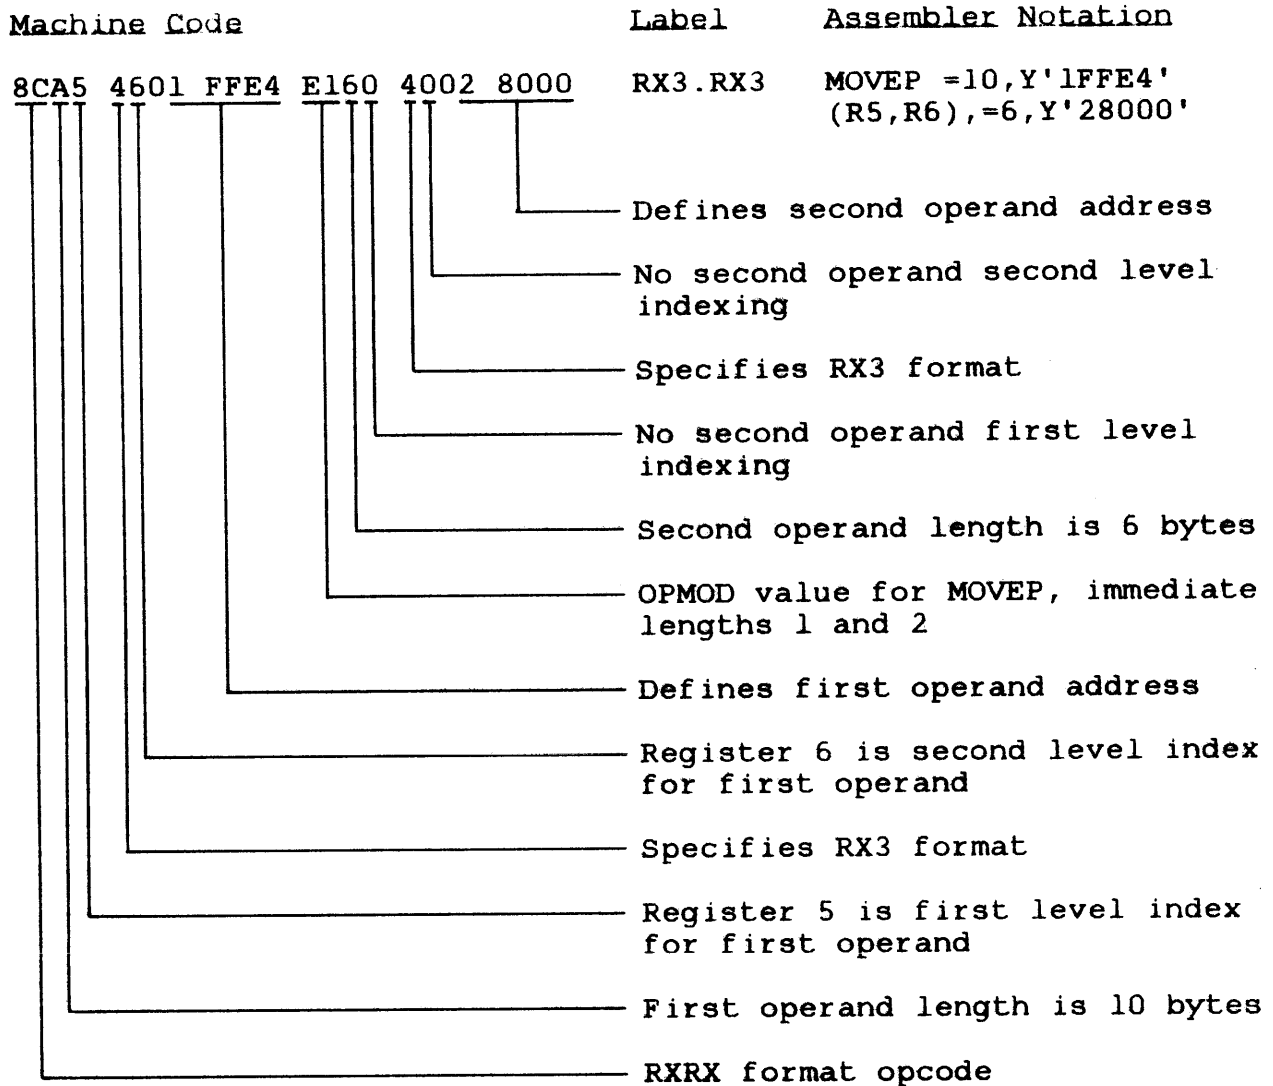
The effective address calculated for the first member is the address of the left-most (lowest address) byte of the first operand string. The effective address calculated for the second member is the address of the left-most byte of the second operand string. An RX2 displacement calculated for either member is with respect to the incremented LOC for that member.

Example 1:



In the above example, both members of the RXX instruction use the RX1 format. No indexing is specified for either member so the first operand address is X'1000', and the second operand address is X'0FF0'.

Example 2:



In this example, both members of the RXX instruction use the RX3 format. Double indexing is specified for the first member and no indexing is specified for the second member. The first operand address is X'1FFE4' plus the contents of index registers 6 and 5. The second operand address is X'28000'. The length of the first operand is 10 bytes and the second operand is 6 bytes.

579

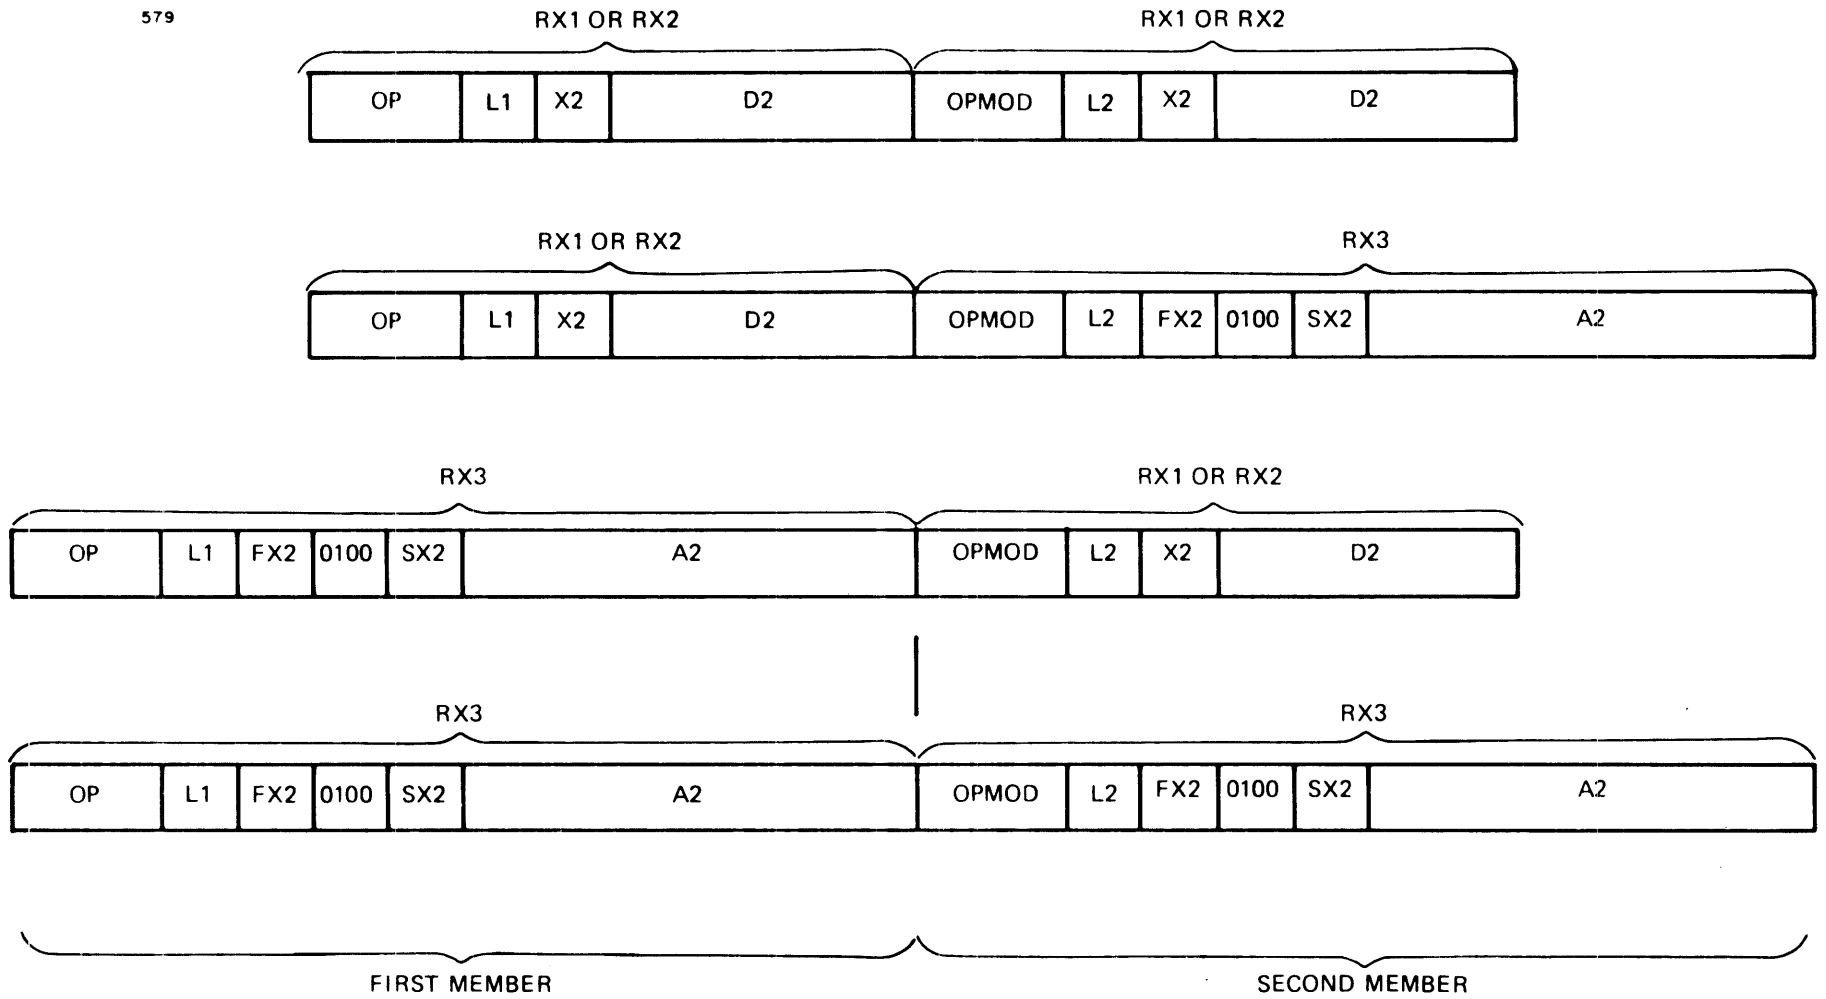


Figure 1-6 RRRX Formats

CHAPTER 2 SYSTEM CONTROL

2.1 INTRODUCTION

Operator control is provided by the console and the system terminal, a microcode-supported device interfaced to the system by an asynchronous line controller. The system terminal can be used as the operating system's console device and can be a video display unit (VDU) or a printing terminal. The asynchronous interface must be strapped as device numbers X'10' and X'11'.

2.2 CONFIGURATION

The console, shown in Figure 2-1, controls power to the system and initial program load (IPL). It also provides controls for system initialization, processor halt/run and single step. LEDs on the system console indicate the current state of the system.

580-2




Figure 2-1 Console

Keyboard commands through the system terminal allow the operator to examine and modify processor registers and main memory locations and begin program execution (see Figure 2-2). Hexadecimal characters and a number of special characters are recognized by the system terminal support microcode. Accepted characters and their meanings are shown in Table 2-1. No other characters are accepted; other characters cause a question mark (?) to be written to the system terminal. When not in use for operator control, the system terminal is available to a running program for use as an input/output (I/O) device. See Appendix F for a flowchart of the console service routine.

TABLE 2-1 SYSTEM TERMINAL SUPPORT COMMAND SUMMARY

581-2

KEY COMMAND SEQUENCE	MEANING	SYSTEM TERMINAL DISPLAY
@nnnnnnn CR	Select memory address and display halfword contents.	<@nnnnnn nnnnnnnn <u>YYYY</u> <
Rn CR	Select general register and display contents.	<Rn <u>YYYYYYYY</u> <
Fn CR	Select single precision floating point register and display contents.	<Fn <u>YYYYYYYY</u> <
Dn CR	Select double precision floating point register and display contents.	<Dn <u>YYYYYYYY</u> <u>YYYYYYYY</u> <
P CR	Select program status word (PSW) and display contents.	<P <u>YYYYYYYY</u> <u>YYYYYYYY</u> <
+	Increment memory location counter (LOC) to display next sequential halfword.	<+ <u>nnnnnnnn</u> <u>YYYY</u> <
-	Decrement memory LOC to display previous halfword.	<- <u>nnnnnnnn</u> <u>YYYY</u> <
=Y _X Y Y CR	Replace contents of currently selected memory location or register with new data.	<=YYYY for memory < <=YYYYYYYY for register <
<	Begin program execution at current memory location.	<<
#	Delete command.	<@10# <
>	Single step the instruction at current memory location.	<>

NOTES

1. Characters in boxes indicate operational key strokes are required for commands.
2. The character symbol of lower-case n is used to indicate the hexadecimal address of memory or register.
3. The character symbol of upper-case Y is used to indicate hexadecimal contents of memory or register.
4. Underlined characters are those output from the system. Characters not underlined are those typed by the operator.
5. A back arrow, underline (X'5F') or back space (X'08') character can be used to delete the previously input hexadecimal character.
6. Space characters can be entered as desired; they are ignored by the processor.




Figure 2-2 Model 6100 Keyboard Layout

2.3 CONSOLETTTE SWITCHES AND INDICATORS

The following sections detail the functions of the consolette switches and indicators.

2.3.1 Key-Operated Security Lock

This is a 3-position (STANDBY-ON-LOCK), key-operated switch that controls primary power to the system. It can also disable (LOCK) the initialize and console switches, thereby preventing any accidental manual input to the system. The power indicator lamp (POWER) is on when the security lock is in the ON or LOCK position.

2.3.2 Control Switches

All the control switches, with the exception of the IPL switch, are enabled only when the key-operated security lock is in the ON position, and primary AC power is applied.

SINGLE STEP


SINGLE



When in the UP position, control is automatically given to the system terminal support routine at the conclusion of each user level instruction. The PSW is displayed, including the address of the next sequential instruction (LOC). Execution of the next instruction is caused by pressing the HALT/RUN switch or by typing a less than character or greater than character (< or >) on the system terminal. To resume normal run mode execution, return the SINGLE STEP switch to the DOWN position and begin execution by pressing the HALT/RUN switch or by typing the less than character (<) on the system terminal. The SINGLE STEP switch is disabled when the security lock is in the LOCK position. Attempts to single step through instructions that I/O to the system terminal do not produce meaningful results.

HALT/RUN


HALT/RUN



This momentary contact switch causes program execution to be halted if the system was running or resumed if the system was halted. When halted, control is given to the system terminal support routine through which the memory or registers can be examined or modified and program execution restarted. If the processor was already in the system terminal support routine, program execution is started. This switch is disabled if the security lock is in the LOCK position.

IPL

ENABLE




This switch is not disabled by the security lock. When in the ENABLE position, an IPL from the loader storage unit (LSU) is performed after any of the following steps:

DISABLE

- Turn the security lock from the STANDBY to the ON position.
- Depress the initialize (INIT) switch.
- Return AC power to the system.

INITIALIZE

INIT



This momentary contact switch causes the system to be initialized. The initialization sequence clears all device controllers on the I/O bus and resets certain functions in the processor. The fault lamp (FAULT) comes on when the switch is depressed and is extinguished with the completion of the initialization sequence.

2.4 OPERATING INSTRUCTIONS

The following sections detail operating instructions for power up, entering the console service and IPL.

2.4.1 Power Up

To prevent IPL on power up, place the IPL switch in the DISABLE position. To power up the system, turn the key-operated security lock clockwise from the STANDBY to the ON position. The power lamp (POWER) lights up, and power is provided to the system. The fault lamp (FAULT) on the console also lights, and the microdiagnostic routine is entered. This routine exercises internal data paths and registers. If main memory power has fallen out of regulation since the system was last running, locations X'000000' to X'07FFFF' are initialized. The diagnostic routine tests the lowest 512kb of memory before extinguishing the FAULT lamp. This diagnostic is limited in scope, serving only to indicate a go/no-go condition. If an error is detected in any portion of the microdiagnostic, the microcode loops indefinitely, and the FAULT lamp remains on. If no errors are detected, the FAULT lamp is turned off.

2.4.2 Entering Console Service

If power was lost while the microcode was in the console service routine, control is returned to the console when the power-up sequence is complete, provided that IPL is not enabled. If the system was executing a program when power was lost, execution resumes when power returns, provided that IPL is not enabled. To enter console service in this case, depress the HALT/RUN switch.

2.4.3 Initial Program Load (IPL)

To perform IPL, place the IPL switch in the ENABLE position; then initialize the system by depressing the INIT switch momentarily. A power-down/power-up sequence is emulated, and diagnostics are performed. At the successful completion of the microdiagnostic sequence, an IPL from the LSU is performed. Control is transferred to the newly-loaded program.

2.5 SYSTEM TERMINAL COMMANDS

When the system terminal support routine is entered from power up or initialize, a carriage return (CR) and line feed (LF) sequence are output. The current value of the PSW status and LOC are output, followed by another CR and LF sequence. Finally, the less than operator prompt character (<) is output to indicate that the system is ready to receive operator commands. If memory power was lost, the LOC is set to X'00FFFFFFE', and the PSW is set to X'00008000'. In this case, the first 512kb of memory are written during power up to establish the error correction code (ECC) bits.

Space characters can be used as desired in any of the described system terminal commands. Spaces are ignored by the console routine.

2.5.1 Select an Address and Examine (@)

The commercial "at" sign (@) places the system terminal support routine in the address mode. This character can be followed by up to six hexadecimal digits of address. Leading zeros are not required. If more than six digits are input, only the least significant six are used. A CR is used to signal the end of the address; the address input is then copied into the LOC. A CR and LF sequence are output, followed by the new value of the LOC and the halfword contents of that location. Note that the data fetch is subject to memory relocation if enabled by the current PSW. After this display, a CR and LF sequence are output, followed by a new operator prompt.

If an invalid character is input by the operator, the system responds by outputting a question mark (?), CR, LF and an operator prompt.

2.5.2 Increment and Examine Next Location (+)

After examining a memory location, the plus character (+) can be used to advance the LOC by two. No other operator input is required. A CR and LF are output, followed by the new LOC value and the halfword contents of that location. This memory access is subject to the relocation defined by the current PSW. After outputting another CR and LF, the operator prompt character is output. This procedure can be repeated to examine sequential memory locations.

2.5.3 Decrement and Examine Prior Location (-)

After examining a memory location, the minus character (-) can be used to decrement the LOC by two. No other operation is required. A CR and LF are output, followed by the new LOC value and the halfword contents of that location. This memory access is subject to the relocation defined by the current PSW. After outputting another CR and LF, the operator prompt character is output. This procedure can be repeated to examine sequential memory locations.

2.5.4 Modify Current Location (=)

After examining a memory location, the equal sign (=) can be used to put the system terminal support routine in the memory write mode. This character can be followed by up to four hexadecimal digits of data to be written. Leading zeros are not required. If more than four digits are input, only the least significant four are used. A CR is used to signal the end of the data. At that time, the accumulated data is written into the memory location currently addressed by the LOC. This memory write is subject to the relocation defined by the current PSW. The current LOC is incremented by two and a CR, LF and an operator prompt are output. This procedure can be repeated to modify sequential memory locations.

2.5.5 Examine General Register (R)

The character R causes the system terminal support routine to interpret subsequent hexadecimal input as the number of a general register (in the set selected by the current PSW) to be displayed. A CR is used to signal the end of hexadecimal input. At that time, the least significant four bits of the accumulated hexadecimal data are taken as the desired register number. The fullword contents of that register are output followed by a CR, LF and an operator prompt. Plus and minus commands are invalid for general registers.

2.5.6 Modify General Register (=)

Immediately after examining a general register, the equal sign (=) can be used to change the contents of the currently selected register. The equal sign can be followed by up to eight hexadecimal digits of data. Leading zeros are not required. If more than eight digits are input, only the least significant eight are used. A CR is used to signal the end of the data input. At that time, the accumulated data is copied into the currently selected general register. A CR, LF and an operator prompt are then output.

2.5.7 Examine Single Precision Floating Point Register (F)

The character F causes the system terminal support routine to interpret subsequent hexadecimal input as the number of a single precision floating point register to be displayed. A CR is used to signal the end of hexadecimal input. At that time, the least significant four bits of the accumulated hexadecimal data are taken as the desired register number. If necessary, this number is rounded to the next lowest even number. The fullword contents of that register are output followed by a CR, LF and an operator prompt. Plus and minus commands are invalid for floating point registers.

2.5.8 Modify Single Precision Floating Point Register (=)

Immediately after examination of a single precision floating point register, that register is available for modification. Type an equal sign (=) followed by up to eight hexadecimal digits of data. Leading zeros are not required. If more than eight digits are input, only the least significant eight are used. A CR is used to signal the end of the data input. At that time, the accumulated data is copied into the currently selected single precision floating point register. This data is not tested for normalization; therefore, an unnormalized floating point number can be manually placed in the register. The system outputs a CR, LF and an operator prompt.

2.5.9 Examine Double Precision Floating Point Register (D)

The character D causes the system terminal support routine to interpret subsequent hexadecimal input as the number of a double precision floating point register to be displayed. A CR is used to signal the end of hexadecimal input. At that time, the least significant four bits of the accumulated hexadecimal data are taken as the desired register number. If necessary, this number is rounded to the next lowest even number. The doubleword contents of that register are output, followed by a CR, LF and an operator prompt. Plus and minus commands are invalid for floating point registers.

2.5.10 Modify Double Precision Floating Point Register (=)

Immediately after examining a double precision floating point register, that register is available for modification. Type an equal sign (=) followed by up to 16 hexadecimal digits. Leading zeros are not required. If more than 16 digits are input, only the last 16 digits are used. A CR is used to signal the end of the data input. At that time, the accumulated data is copied into the currently selected double precision register. The data is not tested for normalization; therefore, an unnormalized floating point number can be manually placed in a double precision register. The system outputs a CR, LF and an operator prompt.

2.5.11 Examine Program Status Word (PSW) (P)

The character P puts the system terminal support routine into the PSW display mode. A CR is required to complete this command input. Upon receipt of the CR, the contents of the PSW are output followed by a CR, LF and an operator prompt. The plus and minus commands are invalid for the PSW.

2.5.12 Modify Program Status Word (PSW) (=)

Immediately after examining the PSW, the equal sign (=) can be used to change the contents of the PSW status field. The equal sign can be followed by up to eight hexadecimal digits of data. Leading zeros are not required. If more than eight digits are input, only the least significant eight are used. A CR is used to signal the end of the data input. At that time, the accumulated data is copied into the PSW, which is then displayed. A CR, LF and an operator prompt are then output.

2.5.13 Execute Single Instruction (>)

Entering the greater than character (>) causes the processor to execute the instruction indicated by the LOC in single step mode. After this execution, the console service routine displays the PSW and LOC, followed by a CR, LF and an operator prompt.

2.5.14 Enter Run Mode (<)

Entering the less than character (<) causes the processor to begin program execution, starting with the instruction indicated by the LOC.

2.6 MEMORY INITIALIZATION

The following example shows how to set up dedicated low memory for loading either the 32-bit relocating loader or the diagnostic loader from magnetic tape.

Example:

8070

< [P] [CR]

007EF0 02073E

< [=] [2] [0] [0] [0] [CR]

002000 02073E

< [e] [3] [0] [CR]

000030 0000

< [=] [8] [8] [0] [0] [CR]

000032 8000

< [=] [2] [0] [0] [0] [CR]

000034 0000

< [+]

000036 0050

< [=] [3] [0] [CR]

000038 0000

< [=] [8] [8] [0] [0] [CR]

00003A 8000

< [=] [2] [0] [0] [0]

00003C 0000

< [+]

00003E 0050

< [=] [3] [8] [CR]

000040 4000

< [e] [5] [0] [CR]

000050 D500

Display PSW

Current PSW and LOC

Set PSW to X'002000', enabling machine malfunction interrupts

Current PSW and LOC

Select address X'30', the machine malfunction new PSW

Location X'30' contains X'0000'

Change to BRK instruction

Location X'32' contains X'8000'

Change contents of X'32' to X'2000'

Location X'34' contains X'0000', as desired

Advance the next location

Location X'36' contains X'0050'

Change contents of X'36' to X'0030'

Location X'38' contains X'0000'

Change to BRK instruction

Location X'3A' contains X'8000'

Change contents of X'3A' to X'2000'

Location X'3C' contains X'0000', as desired

Advance to next location

Location X'3E' contains X'0050'

Change contents of X'3E' to X'0038'

Location X'40' contains X'4000', which can be ignored

Select address X'50'

Location X'50' contains X'D500', the desired auto-load instruction

< <input type="checkbox"/> + <u>000052</u> 00CF	Advance to next location Location X'52' contains X'00CF', the usual auto-load ending address, which is desired
< <input type="checkbox"/> + <u>000054</u> 4300	Advance to next location Location X'54' contains X'4300', part of a branch instruction, which is desired
< <input type="checkbox"/> + <u>000056</u> 0080	Advance to next location Location X'56' contains X'0080', the desired branch address
< <input type="checkbox"/> @ <input type="checkbox"/> 7 <input type="checkbox"/> 8 <input type="checkbox"/> CR	Select address X'78'
<u>000078</u> C186	Location X'78' contains X'C186'
< <input type="checkbox"/> = <input type="checkbox"/> 8 <input type="checkbox"/> 5 <input type="checkbox"/> A <input type="checkbox"/> 1 <input type="checkbox"/> CR	Change contents of X'78' to X'85A1' the device number and command byte for the magnetic tape unit
<u>00007A</u> 0000	Location X'7A' contains X'0000'
< <input type="checkbox"/> + <u>00007C</u> 00F0	Advance to next location Location X'7C' contains the device address of the selector channel (SELCH), which might be used by the loader
< <input type="checkbox"/> \$ <input type="checkbox"/> 5 <input type="checkbox"/> 0 <input type="checkbox"/> CR	Select address X'50'
<u>000050</u> D500	Location X'50' contains X'D500'
< <input type="checkbox"/> <	Start program execution

After loading, the relocating loader relinquishes control of the processor to the loaded program.

2.7 SYSTEM TERMINAL PROGRAMMING INSTRUCTIONS

The system terminal uses a communication I/O board or multiperipheral controller (MPC). The MPC comprises the LSU, a universal clock (UCLOCK), a printer port and eight serial ports. Since the microprogram of the processor must communicate with the system terminal, the device address is fixed at X'010' and X'011'. The MPC supports only full-duplex operations.

The microprogram programs the system terminal interface for highest clock rate, two stop bits per character, seven data bits and even parity. Echoplex is not turned on.

CHAPTER 3
LOGICAL OPERATIONS

3.1 INTRODUCTION

The set of logical instructions provides a means for manipulating binary data. Many of the instructions grouped with the logical set can also be used in arithmetic and other operations. These instructions include loads, stores, compares, shifts, list processing, translation and cyclic redundancy checks (CRCs).

3.2 LOGICAL DATA FORMATS

Logical data can be organized as bytes, halfwords, fullwords or bit arrays of up to 2^{27} bits as shown in Figure 3-1.

585







Figure 3-1 Logical Data

3.3 OPERATIONS

In logical operations between the contents of a general register and a halfword operand, the halfword operand is expanded to a fullword before the operation starts. The halfword is expanded by propagating the most significant bits through bits 0:15 of the fullword. For example, the halfword 'A000' is expanded to 'FFFA000' before participating in the operation.

3.3.1 Boolean Operations

The Boolean operators AND, OR and Exclusive-OR (XOR) operate on halfword and fullword quantities. All bits in both operands participate individually. The Boolean functions are defined as follows:

0 AND 0 = 0	
0 AND 1 = 0	(logical product)
1 AND 0 = 0	
1 AND 1 = 1	
0 OR 0 = 0	
0 OR 1 = 1	(logical sum)
1 OR 0 = 1	
1 OR 1 = 1	
0 XOR 0 = 0	
0 XOR 1 = 1	(logical difference)
1 XOR 0 = 1	
1 XOR 1 = 0	

3.3.2 Translation

The Translate (TLATE) instruction is used to translate a character directly or to effect an unconditional branch to a special translate subroutine. Associated with the Translate instruction is a translation table. The entries in the table are halfwords, as shown in Figure 3-2.

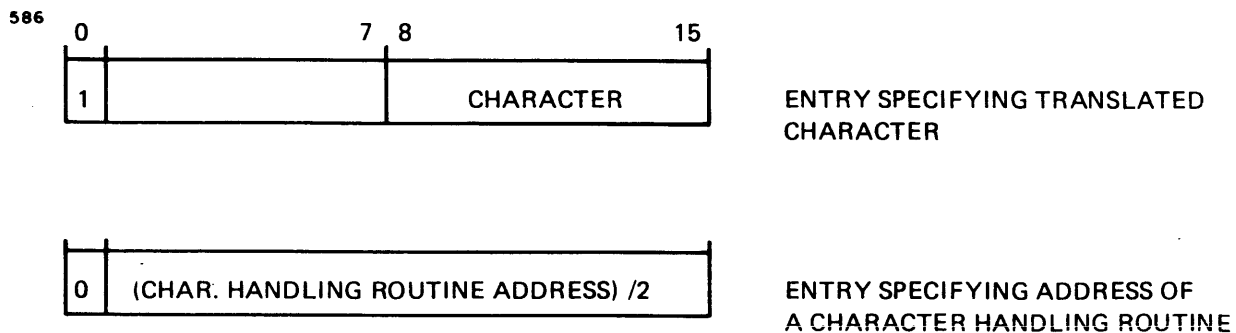


Figure 3-2 Translation Table Entry

The character to be translated is a byte of logical data. This unsigned quantity is doubled and used as an index into the translation table. If the corresponding table entry has a one in bit position zero, then bits 8:15 contain the character to be substituted for the data character. If there is a zero in bit position zero, bits 1:15 contain the address, divided by two, of the translation routine. When the Translate instruction results in a branch, this value is doubled to produce the address of the routine. Because this result is a 16-bit address, the software routine must be located in the first 64kb of the program address space. The program can reside anywhere in memory if it is relocated by the memory address translator (MAT). The translation table can contain up to 256 entries. However, if the data characters are always less than eight bits, fewer entries are required.

3.3.3 List Processing

The list processing instructions manipulate a circular list as defined in Figure 3-3.




Figure 3-3 Circular List Definition

The first four halfwords, called the list header, contain the list parameters. Immediately following the header is the list itself. The first fullword in the list is designated Slot 0. The remaining slots are designated 1, 2, 3, etc., up to a maximum slot number, which is equal to the number in the list minus one. An absolute maximum of 65,535 fullword slots can be specified. (Slots are designated 0 through X'FFFE'.)

The first halfword of the header indicates the number of slots (fullwords) in the entire list. The second halfword indicates the current number of slots being used. When this halfword equals zero, the list is empty. When this halfword equals the number of slots in the list, the list is full. Once initialized, this halfword is maintained automatically. It is incremented when elements are added to the list and decremented when elements are removed.

The third and fourth halfwords of the list header specify the current top of the list and the next bottom of the list, respectively. These pointers are also updated automatically (see Figure 3-4).

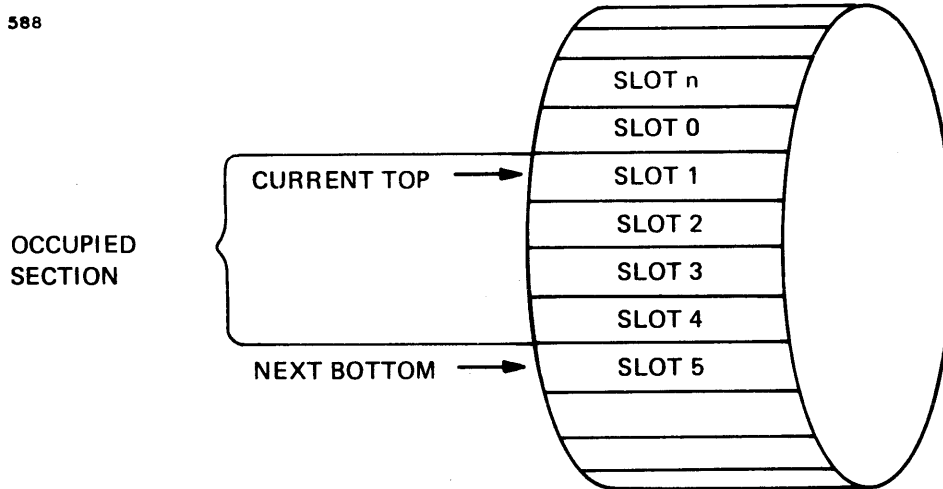


Figure 3-4 Circular List

3.4 LOGICAL INSTRUCTION FORMATS

The logical instructions use the Register to Register (RR), Short Form (SF), Register and Indexed Storage (RX), and Register and Immediate Storage (RI) instruction formats.

3.5 LOGICAL INSTRUCTIONS

The instructions described in this section are:

L	Load
LR	Load Register
LI	Load Immediate
LIS	Load Immediate Short
LCS	Load Complement Short

LH	Load Halfword
LHI	Load Halfword Immediate
LA	Load Address
LRA	Load Real Address
LHL	Load Halfword Logical
LM	Load Multiple
LB	Load Byte
LBR	Load Byte Register
EXHR	Exchange Halfword Register
EXBR	Exchange Byte Register
ST	Store
STH	Store Halfword
STM	Store Multiple
STB	Store Byte
STBR	Store Byte Register
CL	Compare Logical
CLR	Compare Logical Register
CLI	Compare Logical Immediate
CLH	Compare Logical Halfword
CLHI	Compare Logical Halfword Immediate
CLB	Compare Logical Byte
N	AND
NR	AND Register
NI	AND Immediate
NH	AND Halfword
NHI	AND Halfword Immediate
O	OR
OR	OR Register
OI	OR Immediate
OH	OR Halfword
OHI	OR Halfword Immediate
X	Exclusive-OR
XR	Exclusive-OR Register
XI	Exclusive-OR Immediate
XH	Exclusive-OR Halfword
XHI	Exclusive-OR Halfword Immediate

TI	Test Immediate
THI	Test Halfword Immediate
SLL	Shift Left Logical
SLLS	Shift Left Logical Short
SRL	Shift Right Logical
SRLS	Shift Right Logical Short
SLHL	Shift Left Halfword Logical
SLHLS	Shift Left Halfword Logical Short
SRHL	Shift Right Halfword Logical
SRHLS	Shift Right Halfword Logical Short
RLL	Rotate Left Logical
RRL	Rotate Right Logical
TS	Test and Set
TBT	Test Bit
SBT	Set Bit
RBT	Reset Bit
CBT	Complement Bit
CRC12	Cyclic Redundancy Check Modulo 12
CRC16	Cyclic Redundancy Check Modulo 16
TLATE	Translate
ATL	Add to Top of List
ABL	Add to Bottom of List
RTL	Remove from Top of List
RBL	Remove from Bottom of List

3.5.1 Load (L, LR, LI)

Load (L)

Load Register (LR)

Load Immediate (LI)

Assembler Notation	Opcode	Format
L R1,D2(X2)	58	RX1,RX2
L R1,A2(FX2,SX2)	58	RX3
LR R1,R2	08	RR
LI R1,I2(X2)	F8	RI2

Operation:

The second operand replaces the contents of the register specified in R1.

Condition Code:

C	V	G	L	
0	0	0	0	Value is zero
0	0	0	1	Value is not zero
0	0	1	0	Value is not zero

Programming Notes:

When the load instructions operate on fixed point data, the condition code indicates zero (no flags), negative (L flag) or positive (G flag) value.

In the RR format, if R1 equals R2, the Load instruction functions as a test on the contents of the register.

In the RX formats, the second operand must be located on a fullword boundary.

3.5.2 Load Immediate Short (LIS)

<u>Assembler Notation</u>	<u>Opcode</u>	<u>Format</u>
LIS R1,N	24	SF

Operation:

The 4-bit second operand is expanded to a 32-bit fullword with high order bits forced to zero. This fullword replaces the contents of the register specified by R1.

Condition Code:

C V G L	

0 0 0 0	Value is zero
0 0 1 0	Value is not zero

Programming Note:

When this instruction operates on fixed point data, the condition code indicates zero (no flags) or positive (G flag) value.

Example:

<u>Assembler Notation</u>	<u>Machine Code</u>	<u>Comments</u>
LIS REG4,15	244F	LOAD 15 INTO REG4

Result of LIS Instruction:

(REG4) = 0000000F
Condition code = 0010 (G=2)

3.5.3 Load Complement Short (LCS)

<u>Assembler Notation</u>	<u>Opcode</u>	<u>Format</u>
LCS R1,N	25	SF

Operation:

The 4-bit second operand is expanded to a 32-bit fullword with high order bits forced to zero. The two's complement value of this fullword then replaces the contents of the register specified by R1.

Condition Code:

C V G L	

0 0 0 0	Value is zero
0 0 0 1	Value is not zero

Programming Note:

When this instruction operates on fixed point data, the condition code indicates zero (no flags) or negative (L flag) value.

Example:

<u>Assembler Notation</u>	<u>Machine Code</u>	<u>Comments</u>
LCS REG8,7	2587	LOAD -7 INTO REG8

Result of LCS Instruction:

(REG8) = FFFF FFF9
Condition code = 0001 (L=1)

3.5.4 Load Halfword (LH, LHI)

Load Halfword (LH)
Load Halfword Immediate (LHI)

Assembler Notation	Opcode	Format
LH R1,D2(X2)	48	RX1,RX2
LH R1,A2(FX2,SX2)	48	RX3
LHI R1,I2(X2)	C8	R11

Operation:

The halfword second operand is expanded to a fullword by propagating the most significant bit through bits 0:15. This fullword replaces the contents of the register specified by R1.

Condition Code:

C	V	G	L	
0	0	0	0	Value is zero
0	0	0	1	Value is not zero
0	0	1	0	Value is not zero

Programming Notes:

When the load halfword instructions operate on fixed point data, the condition code indicates zero (no flags), negative (L flag) or positive (G flag) value.

In the RX formats, the second operand must be located on a halfword boundary.

In the R11 format, the 16-bit I2 field is extended to a fullword by propagating the sign bit through bits 0:15. The contents of the index register specified by X2 are then added to form the fullword second operand.

3.5.5 Load Address (LA)

Assembler Notation	Opcode	Format
LA R1,D2(X2)	E6	RX1,RX2
LA R1,A2(FX2,SX2)	E6	RX3

Operation:

The effective address of the second operand (24 bits) replaces bits 8:31 of the register specified by R1. Bits 0:7 of the register specified by R1 are forced to zero.

Condition Code:

Unchanged

Programming Note:

The length of the address quantity depends on the internal structure of the particular machine; in this processor the calculated address replaces bits 8:31 of the register specified by R1, and bits 0:7 are replaced by zero.

3.5.6 Load Real Address (LRA)

Assembler Notation	Opcode	Format
LRA R1,D2(X2)	63	RX1,RX2
LRA R1,A2(FX2,SX2)	63	RX3

Operation:

This instruction simulates the operation of the MAT. The register specified by R1 contains a program address (not relocated). The second operand address points to a relocation/protection module parameter block, in the format shown:

BYTE	OFFSET	0	1	14	15	31
+0	PSTD	(PST ENTRIES) -1			A(PROCESS SEGMENT TABLE)/128	
+4	SSTD	(SST ENTRIES) -1			A(SHARED SEGMENT TABLE)/128	

The address contained in the register specified by R1 is relocated, using the appropriate parameters. The relocated address replaces the contents of the register specified by R1.

Condition Code:

C	V	G	L	
1	0	0	0	Segment not mapped
0	1	0	0	Nonpresent segment
0	0	1	X	Write-protected segment
0	0	X	1	Read- or execute-protected segment
0	0	0	0	No restrictions

The condition code is determined on a priority basis with segment table size exceeded checked first, nonpresent segment second, segment limit exceeded third and all protect keys (as a group) last.

Programming Notes:

Segment tables must conform to the rules given in the section on memory management; otherwise, the results of the LRA instruction are undefined.

If the address is not mapped or not present, the register specified by R1 is unchanged.

Segment table size exceeded or segment limit exceeded results in condition code 1000 (unmapped).

The second operand must be located on a fullword boundary.

Example:

This example performs an address translation in the same manner as the MAT as implemented on this machine (4kb page size). The steps shown are not optimal and do not reflect the actual operation of the MAT.

To set up for this example, register R1 contains X'053147', the program address to be translated. RELOCBLK is the address of a relocation/protection module parameter block. This block contains two fullwords. The first of these is the process segment table descriptor (PSTD), with the value X'000E06BF'. The second is the shared segment table descriptor (SSTD), with the value X'000C06C0'. Memory location X'035FA8' contains the process segment table entry (PSTE) to be used with the value X'588A0028'. Memory location X'036028' contains the shared segment table entry (SSTE) to be used with the value X'58126800'. The instruction proceeds as follows:

```
LRA  R1,RELOCBLK          TRANSLATE ADDRESS IN R1
```

1. The PSTD is fetched from RELOCBLK and ANDED with X'FFFE0000' to extract the segment table size field. The result, X'000E0000', is shifted right 17 bit positions, yielding X'00000007'. This value is the number of entries in the process segment table (PST) minus one. Therefore, the PST has entries for segments 0 through 7.
2. The program address from register R1, X'053147', is shifted right 16 bit positions to yield the specified segment number, X'00000005'. The segment number is compared with the PST size. If the PST size were less than the segment number, this would mean that no entry existed in the PST for the specified segment, and that the segment was unmapped (condition code = 8). However, such is not the case, and the instruction proceeds.
3. The PSTD is ANDED with X'0001FFFF' to extract the segment table address field. The result, X'000006BF', is shifted left seven bit positions to multiply it by 128. This yields the address of the PST, X'35F80'.

4. The segment number specified by the program address in R1 (X'053147') is used as an index into the PST. Because each segment table entry (STE) requires eight bytes, the segment number, X'00000005', is shifted left three bit positions, to multiply it by eight. The result, X'00000028', and the address of the PST, X'035F80', are added. The result is the address X'035FA8', and the PSTE at that address is fetched. This PSTE has the value X'588A0028'.

The PSTE is ANDed with the value X'40000000' to test the presence bit in the STE. If the bit were zero, this would mean the segment was not present (condition code = 4). Such is not the case, however, and the instruction proceeds.

5. The PSTE is then ANDed with X'00800000' to test the shared segment bit. If the bit were zero, the LRA instruction would use the data in the PSTE as data in the SSTE also and perform the operations in Step 9 below; but such is not the case.

The shared segment bit in the PSTE is set, which means that an entry from the shared segment table (SST) must also be used in translating the program address. The SSTD (X'000C06C0') is ANDed with X'FFFE0000' to extract the segment table size field. The result, X'000C0000', is shifted right 14 bit positions to yield X'00000030'. This value is the maximum SST offset, which is the offset in bytes from the start of the SST to the beginning of the last entry.

6. The SSTD is ANDed with X'0001FFFF' to extract the segment table address field. The result, X'000006C0', is shifted left seven bit positions to yield the address of the SST, X'036000'.

7. The PSTE is now ANDed with X'0001FFFF' to extract the segment relocation field (SRF). This field has the value X'00000028'. If this value exceeded the maximum SST offset, this would mean that no entry existed in the SST for the specified segment, and that the segment was unmapped (condition code = 8); but such is not the case, so the instruction proceeds. The SRF is added with the PST address, X'036000'. The SSTE pointed to by the PSTE is located at the resulting address, X'036028'.

8. The SSTE is fetched, and its value found to be X'58126800'. This value is ANDed with X'40000000' to test the STE presence bit. If the bit were zero, this would mean the segment was not present (condition code = 4); but such is not the case, and the instruction proceeds.

9. The SSTE, with a value X'58126880', is ANDed with the value X'003C0000' to extract the segment limit field (SLF). The resulting value, X'00100000', is shifted right six bit positions, yielding an SLF value of X'00004000'.

The program address from R1, X'053147', is ANDed with X'0000F000'. The resulting value, X'00003000', is compared to the SLF value, X'00004000'. If the SLF value were the lesser of the two values, this would indicate that the program address was in an unreachable part of the segment (segment limit violation), and thus unmapped (condition code = 8); but such is not the case, and the instruction proceeds.

10. At this point, address translation can be performed. The SSTE, with value X'58126800', is ANDed with the value X'0001FFFF' to extract the SRF. This field has the value X'00006800'. The SRF is shifted left seven bit positions, giving the relocation value X'00340000'.

The program address from R1, X'053147', is ANDed with the value X'0000FFFF', giving the value X'00003147'. To this value is added the relocation value, X'00340000'. The result is the translated program address, X'343147', which replaces the contents of register R1.

11. The PSTE, with value X'588A0028', and the SSTE, with value X'58126800', are ANDed, yielding the value X'58020000'. This value contains the combined segment access keys. If ANDing the keys with X'08000000' yielded a zero result, the G flag would be set in the condition code to indicate a write-protected segment. If ANDing the keys with X'10000000' yielded a zero result, the L flag would be set in the condition code to indicate a read-protected segment; but neither is the case. ANDing the keys with X'04000000' does yield a zero result, and the L flag is set in the condition code to indicate that the segment is execute-protected. The LRA instruction terminates once these tests have been performed (see Figure 3-5).

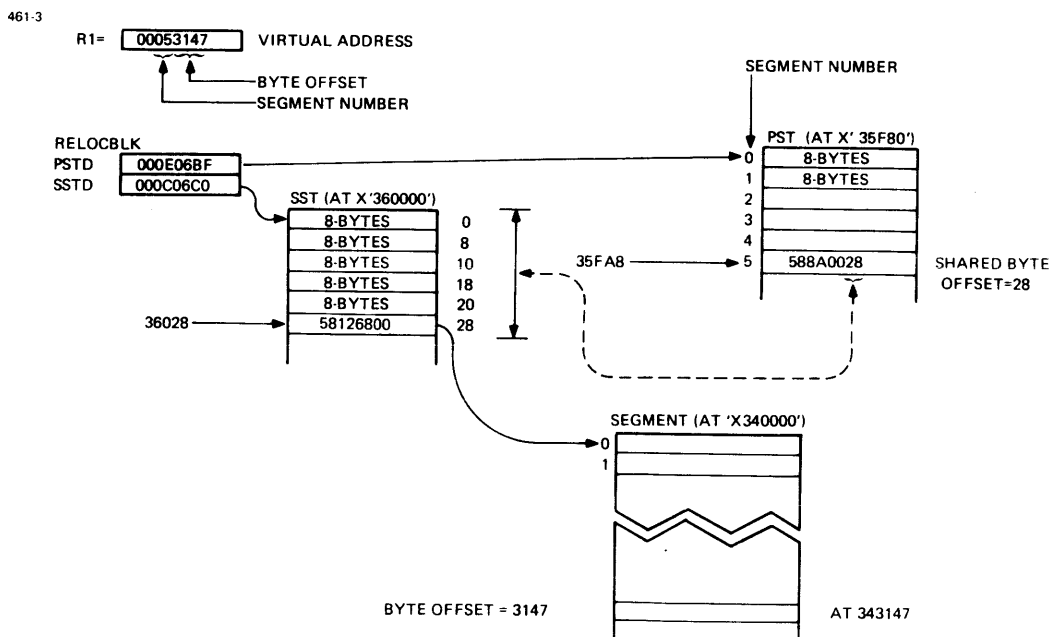


Figure 3-5 LRA Example

3.5.7 Load Halfword Logical (LHL)

<u>Assembler Notation</u>	<u>Opcode</u>	<u>Format</u>
LHL R1,D2(X2)	73	RX1,RX2
LHL R1,A2(FX2,SX2)	73	RX3

Operation:

The halfword second operand replaces bits 16:31 of the register specified by R1. Bits 0:15 of the register specified by R1 are replaced by zero.

Condition Code:

C	V	G	L	
0	0	0	0	Value is zero
0	0	1	0	Value is not zero

Programming Note:

The second operand must be located on a halfword boundary.

3.5.8 Load Multiple (LM)

<u>Assembler Notation</u>	<u>Opcode</u>	<u>Format</u>
LM R1,D2(X2)	D1	RX1,RX2
LM R1,A2(FX2,SX2)	D1	RX3

Operation:

Successive registers, starting with the register specified by R1, are loaded from successive memory locations, starting with the location specified as the effective address of the second operand. Each register is loaded with a fullword from memory. The process stops when register 15 has been loaded.

Condition Code:

Unchanged

Programming Notes:

The second operand must be located on a fullword boundary.

The second operand address is formed before any registers are loaded; therefore, X2, FX2 and SX2 can be among the registers loaded.

In the event of a machine malfunction due to a noncorrectable memory error or MAT fault, the effective address calculated at the beginning of the instruction is available if a retry is desired. For details, see Chapters 10 and 11.

3.5.9 Load Byte (LB, LBR)

Load Byte (LB)

Load Byte Register (LBR)

<u>Assembler Notation</u>	<u>Opcode</u>	<u>Format</u>
LB R1,D2(X2)	D3	RX1,RX2
LB R1,A2(FX2,SX2)	D3	RX3
LBR R1,R2	93	RR

Operation:

The 8-bit second operand replaces the least significant bits (bits 24:31) of the register specified by R1. Bits 0:23 of the register are forced to zero.

Condition Code:

Unchanged

Programming Note:

In the LBR instruction, the second operand is taken from the least significant eight bits (bits 24:31) of the register specified by R2.

3.5.10 Exchange Halfword Register (EXHR)

<u>Assembler Notation</u>	<u>Opcode</u>	<u>Format</u>
EXHR R1,R2	34	RR

Operation:

Bits 0:15 of the register specified by R2 replace bits 16:31 of the register specified by R1. Bits 16:31 of the register specified by R2 replace bits 0:15 of the register specified by R1.

Condition Code:

Unchanged

Programming Note:

If R1 equals R2, the two halfwords contained within the register are exchanged. If R1 does not equal R2, the contents of R2 are unchanged.

Example:

<u>Assembler Notation</u>	<u>Machine Code</u>	<u>Comments</u>
LI REG5,Y'0ABCDEF9'	F850 0ABC DEF9	(REG5) = 0ABCDEF9
LI REG7,Y'12345678'	F870 1234 5678	(REG7) = 12345678
EXHR REG5,REG7	3457	

Result of EXHR Instruction:

(REG5) = 56781234
(REG7) = 12345678
Condition code unchanged

3.5.11 Exchange Byte Register (EXBR)

<u>Assembler Notation</u>	<u>Opcode</u>	<u>Format</u>
EXBR R1,R2	94	RR

Operation:

The two bytes contained in bits 16:31 of the register specified by R2 are exchanged and loaded into bits 16:31 of the register specified by R1. Bits 0:15 of the register specified by R1 are unchanged. The register specified by R2 is unchanged.

Condition Code:

Unchanged

Programming Note:

R1 and R2 can specify the same register. In this case, the two bytes in bits 16:31 of the register specified by R2 are exchanged.

Example:

<u>Assembler Notation</u>	<u>Machine Code</u>	<u>Comments</u>
LI REG7,X'5A6B3C4D'	F870 5A6B 3C4D	(REG7) = 5A6B3C4D
LI REG3,Y'98761234'	F830 9876 1234	(REG3) = 98761234
EXBR REG7,REG3	9473	

Result of EXBR Instruction:

(REG7) = 5A6B3412

(REG3) = 98761234

Condition code unchanged

3.5.12 Store (ST)

Assembler Notation	Opcode	Format
ST R1,D2(X2)	50	RX1,RX2
ST R1,A2(FX2,SX2)	50	RX3

Operation:

The 32-bit contents of the register specified by R1 replace the contents of the fullword memory location specified by the effective address of the second operand.

Condition Code:

Unchanged

Programming Note:

The second operand must be located on a fullword boundary.

3.5.13 Store Halfword (STH)

<u>Assembler Notation</u>	<u>Opcode</u>	<u>Format</u>
STH R1,D2(X2)	40	RX1,RX2
STH R1,A2(FX2,SX2)	40	RX3

Operation:

Bits 16:31 of the register specified by R1 replace the contents of the halfword memory location specified by the effective address of the second operand.

Condition Code:

Unchanged

Programming Note:

The second operand must be located on a halfword boundary.

3.5.14 Store Multiple (STM)

Assembler Notation	Opcode	Format
STM R1,D2(X2)	D0	RX1,RX2
STM R1,A2(FX2,SX2)	D0	RX3

Operation:

The fullword contents of registers, starting with the register specified by R1, replace the contents of successive fullword memory locations, starting with the location specified by the effective address of the second operand. The process stops when register 15 has been stored.

Condition Code:

Unchanged

Programming Note:

The second operand must be located on a fullword boundary.

3.5.15 Store Byte (STB, STBR)

Store Byte (STB)

Store Byte Register (STBR)

<u>Assembler Notation</u>	<u>Opcode</u>	<u>Format</u>
STB R1,D2(X2)	D2	RX1,RX2
STB R1,A2(FX2,SX2)	D2	RX3
STBR R1,R2	92	RR

Operation:

The least significant eight bits (bits 24:31) of the register specified by R1 are stored in the byte second operand location.

Condition Code:

Unchanged

Programming Note:

In the STBR instruction, the 8-bit quantity is stored in bits 24:31 of the register specified by R2. Bits 0:23 of the register are unchanged.

Example:

The following example illustrates the use of the STBR instruction.

<u>Assembler Notation</u>	<u>Machine Code</u>	<u>Comments</u>
L.I REG4,Y'13577531'	F840 1357 7531	(REG4) = 13577531
L.I REG3,Y'24688642'	F830 2468 8642	(REG3) = 24688642
.		
.		
.		
STBR REG4,REG3	9243	

Result of STBR Instruction:

(REG4) = 13577531

(REG3) = 24688631

Condition code unchanged

3.5.16 Compare Logical (CL, CLR, CLI)

Compare Logical (CL)
 Compare Logical Register (CLR)
 Compare Logical Immediate (CLI)

Assembler Notation	Opcode	Format
CL R1,D2(X2)	55	RX1,RX2
CL R1,A2(FX2,SX2)	55	RX3
CLR R1,R2	05	RR
CLI R1,I2(X2)	F5	RI2

Operation:

The first operand, the contents of the register specified by R1, is compared logically to the second operand. The result is indicated by the condition code setting. Neither operand is changed.

Condition Code:

	C	V	G	L	
	0	X	0	0	First operand equal to second
<i>sk</i>	1	X	0	1	First operand less than second
	1	X	1	0	First operand less than second
	0	X	0	1	First operand greater than second
	0	X	1	0	First operand greater than second

Programming Notes:

In the RX formats, the second operand must be located on a fullword boundary.

The state of the V flag is undefined.

If the second operand is zero, the C flag cannot set.

It is helpful to check the following condition code mask (M1) after a logical comparison:

MASK	TRUE/FALSE*	INFERENCE
3	False	First operand equal to second
3	True	First operand not equal to second
8	False	First operand greater than or equal to second
8	True	First operand less than second

* See Chapter 4 for the true/false concept in branch instructions.

3.5.17 Compare Logical Halfword (CLH, CLHI)

Compare Logical Halfword (CLH)

Compare Logical Halfword Immediate (CLHI)

<u>Assembler Notation</u>	<u>Opcode</u>	<u>Format</u>
CLH R1,D2(X2)	45	RX1,RX2
CLH R1,A2(FX2,SX2)	45	RX3
CLHI R1,I2(X2)	C5	RI1

Operation:

The halfword second operand is expanded to a fullword by propagating the most significant bit through bits 0:15. The first operand, the contents of the register specified by R1, is compared to this fullword. The result is indicated by the condition code setting. Neither operand is changed.

Condition Code:

C	V	G	L	

0	X	0	0	First operand equal to second
1	X	0	1	First operand less than second
1	X	1	0	First operand less than second
0	X	0	1	First operand greater than second
0	X	1	0	First operand greater than second

Programming Notes:

In the RX formats, the second operand must be located on a halfword boundary.

In the RI1 format, the 16-bit I2 field is extended to a fullword by propagating the sign bit through bits 0:15. The contents of the index register specified by X2 are then added to form the fullword second operand.

The state of the V flag is undefined.

If the second operand is zero, the C flag cannot set.

It is helpful to check the following condition code mask (M1) after a logical comparison:

MASK	TRUE/FALSE*	INFERENCE
3	False	First operand equal to second
3	True	First operand not equal to second
8	False	First operand greater than or equal to second
8	True	First operand less than second

* See Chapter 4 for the true/false concept in branch instructions.

3.5.18 Compare Logical Byte (CLB)

Assembler Notation	Opcode	Format
CLB R1,D2(X2)	D4	RX1,RX2
CLB R1,A2(FX2,SX2)	D4	RX3

Operation:

The byte quantity, contained in bits 24:31 of the register specified by R1, is compared with the 8-bit second operand. The result is indicated by the condition code setting. Neither operand is changed.

Condition Code:

C	V	G	L	
0	X	0	0	First operand equal to second
1	X	0	1	First operand less than second
0	X	1	0	First operand greater than second

Programming Notes:

Both operands are treated as unsigned quantities.

If the second operand is zero, the C flag cannot set.

It is helpful to check the following condition code mask (M1) after a logical comparison:

MASK	TRUE/FALSE*	INFERENCE
2	False	First operand not greater than second
2	True	First operand greater than second
3	False	First operand equal to second
3	True	First operand not equal to second
8	False	First operand greater than or equal to second
8	True	First operand less than second

* See Chapter 4 for the true/false concept in branch instructions.

3.5.19 AND (N, NR, NI)

AND (N)
 AND Register (NR)
 AND Immediate (NI)

<u>Assembler Notation</u>	<u>Opcode</u>	<u>Format</u>
N R1,D2(X2)	54	RX1,RX2
N R1,A2(FX2,SX2)	54	RX3
NR R1,R2	04	RR
NI R1,I2(X2)	F4	RI2

Operation:

The logical product of the 32-bit second operand and the contents of the register specified by R1 replace the contents of the register specified by R1. The 32-bit logical product is formed on a bit-by-bit basis.

Condition Code:

C	V	G	L	
0	0	0	0	Result is zero
0	0	0	1	Result is not zero
0	0	1	0	Result is not zero

Programming Notes:

In the RX formats, the second operand must be located on a fullword boundary.

When operating on fixed point data, the condition code indicates zero (no flags), negative (L flag) or positive (G flag) result.

3.5.20 AND Halfword (NH, NHI)

AND Halfword (NH)

AND Halfword Immediate (NHI)

<u>Assembler Notation</u>	<u>Opcode</u>	<u>Format</u>
NH R1,D2(X2)	44	RX1,RX2
NH R1,A2(FX2,SX2)	44	RX3
NHI R1,I2(X2)	C4	RI1

Operation:

The halfword second operand is expanded to a fullword by propagating the most significant bit through bits 0:15. The logical product of this 32-bit quantity and the contents of the register specified by R1 replace the contents of the register specified by R1. The 32-bit logical product is formed on a bit-by-bit basis.

Condition Code:

C	V	G	L	
0	0	0	0	Result is zero
0	0	0	1	Result is not zero
0	0	1	0	Result is not zero

Programming Notes:

In the RX formats, the second operand must be located on a halfword boundary.

In the RI1 format, the 16-bit I2 field is extended to a fullword by propagating the sign bit through bits 0:15. The contents of the index register specified by X2 are then added to form the fullword second operand.

When operating on fixed point data, the condition code indicates zero (no flags), negative (L flag) or positive (G flag) result.

3.5.21 OR (O, OR, OI)

OR (O)
OR Register (OR)
OR Immediate (OI)

<u>Assembler Notation</u>	<u>Opcode</u>	<u>Format</u>
O R1,D2(X2)	56	RX1,RX2
O R1,A2(FX2,SX2)	56	RX3
OR R1,R2	06	RR
OI R1,I2(X2)	F6	RI2

Operation:

The logical sum of the 32-bit second operand and the contents of the register specified by R1 replace the contents of the register specified by R1. The 32-bit logical sum is formed on a bit-by-bit basis.

Condition Code:

C	V	G	L
=====			
0	0	0	0
0	0	0	1
0	0	1	0

Result is zero
Result is not zero
Result is not zero

Programming Notes:

In the RX formats, the second operand must be located on a fullword boundary.

When operating on fixed point data, the condition code indicates zero (no flags), negative (L flag) or positive (G flag) result.

3.5.22 OR Halfword (OH, OHI)

OR Halfword (OH)
OR Halfword Immediate (OHI)

<u>Assembler Notation</u>	<u>Opcode</u>	<u>Format</u>
OH R1,D2(X2)	46	RX1,RX2
OH R1,A2(FX2,SX2)	46	RX3
OHI R1,I2(X2)	C6	RI1

Operation:

The halfword second operand is expanded to a fullword by propagating the most significant bit through bits 0:15. The logical sum of this 32-bit quantity and the contents of the register specified by R1 replace the contents of the register specified by R1. The 32-bit logical sum is formed on a bit-by-bit basis.

Condition Code:

C	V	G	L	
0	0	0	0	Result is zero
0	0	0	1	Result is not zero
0	0	1	0	Result is not zero

Programming Notes:

In the RX formats, the second operand must be located on a halfword boundary.

In the RI1 format, the 16-bit I2 field is extended to a fullword by propagating the sign bit through bits 0:15. The contents of the index register specified by X2 are then added to form the fullword second operand.

When operating on fixed point data, the condition code indicates zero (no flags), negative (L flag) or positive (G flag) result.

3.5.23 Exclusive-OR (X, XR, XI)

Exclusive-OR (X)

Exclusive-OR Register (XR)

Exclusive-OR Immediate (XI)

<u>Assembler Notation</u>	<u>Opcode</u>	<u>Format</u>
X R1,D2(X2)	57	RX1,RX2
X R1,A2(FX2,SX2)	57	RX3
XR R1,R2	07	RR
XI R1,I2(X2)	F7	RI2

Operation:

The logical difference of the 32-bit second operand and the contents of the register specified by R1 replace the contents of the register specified by R1. The 32-bit logical difference is formed on a bit-by-bit basis.

Condition Code:

C	V	G	L	
0	0	0	0	Result is zero
0	0	0	1	Result is not zero
0	0	1	0	Result is not zero

Programming Notes:

In the RX formats, the second operand must be located on a fullword boundary.

When operating on fixed point data, the condition code indicates zero (no flags), negative (L flag) or positive (G flag) result.

3.5.24 Exclusive-OR Halfword (XH, XHI)

Exclusive-OR Halfword (XH)

Exclusive-OR Halfword Immediate (XHI)

<u>Assembler Notation</u>	<u>Opcode</u>	<u>Format</u>
XH R1,D2(X2)	47	RX1,RX2
XH R1,A2(FX2,SX2)	47	RX3
XHI R1,I2(X2)	C7	RI1

Operation:

The halfword second operand is expanded to a fullword by propagating the most significant bit through bits 0:15. The logical difference of this 32-bit quantity and the contents of the register specified by R1 replace the contents of the register specified by R1. The 32-bit logical difference is formed on a bit-by-bit basis.

Condition Code:

C	V	G	L	
=====				
0	0	0	0	Result is zero
0	0	0	1	Result is not zero
0	0	1	0	Result is not zero

Programming Notes:

In the RX formats, the second operand must be located on a halfword boundary.

In the RI1 format, the 16-bit I2 field is extended to a fullword by propagating the sign bit through bits 0:15. The contents of the index register specified by X2 are then added to form the fullword second operand.

When operating on fixed point data, the condition code indicates zero (no flags), negative (L flag) or positive (G flag) result.

3.5.25 Test Immediate (TI)

<u>Assembler Notation</u>	<u>Opcode</u>	<u>Format</u>
TI R1,I2(X2)	F3	RI2

Operation:

Each bit of the second operand is logically ANDed with the corresponding bit in the register specified by R1. Neither operand is changed.

Condition Code:

C	V	G	L	
0	0	0	0	Result is zero
0	0	0	1	Result is not zero
0	0	1	0	Result is not zero

Programming Notes:

When operating on fixed point data, the condition code indicates zero (no flags), negative (L flag) or positive (G flag) result.

This instruction works the same as the AND Immediate instruction (NI) except that the first operand is not changed.

Example:

This example tests if bit 16 of register 9 is set.

<u>Assembler Notation</u>	<u>Comments</u>
TI REG9,Y'00008000'	Test bit 16
BNZ LABEL	Branch if bit is set

Where:

(REG9) = 7EFBC230

Result of TI Instruction:

(REG9) unchanged
Condition code = 0010 (G=1)
The conditional branch is taken.

3.5.26 Test Halfword Immediate (THI)

<u>Assembler Notation</u>	<u>Opcode</u>	<u>Format</u>
THI R1,I2(X2)	C3	R11

Operation:

The halfword second operand is expanded to a fullword by propagating the most significant bit through bits 0:15. Each bit in this quantity is logically ANDed with the corresponding bit contained in the register specified by R1. Neither operand is changed.

Condition code:

C	V	G	L	
0	0	0	0	Result is zero
0	0	0	1	Result is not zero
0	0	1	0	Result is not zero

Programming Notes:

When operating on fixed point data, the condition code indicates zero (no flags), negative (L flag) or positive (G flag) result.

In the R11 format, the 16-bit I2 field is extended to a fullword by propagating the sign bit through bits 0:15. The contents of the index register specified by X2 are then added to form the fullword second operand.

This instruction works the same as the AND Halfword Immediate instruction (NHI) except that the first operand is not changed.

Example:

This example tests if any of bits 0:16 of register 9 are set.

<u>Assembler Notation</u>	<u>Comments</u>
THI REG9,X'8000'	Test bits 0:16
BNZ LABEL	Branch if any set

Where:

(REG9) = 80800000

Result of THI Instruction:

(REG9) unchanged
Condition code = 0001 (L=1)
The conditional branch is taken.

3.5.27 Shift Left Logical (SLL, SLLS)

Shift Left Logical (SLL)
Shift Left Logical Short (SLLS)

Assembler Notation	Opcode	Format
SLL R1,I2(X2)	ED	R11
SLLS R1,N	11	SF

Operation:

The first operand, the contents of the register specified by R1, is shifted left the number of places specified by the second operand. Bits shifted out of position 0 are shifted through the carry (C) flag of the condition code and then lost. The last bit shifted remains in the C flag. Zeros are shifted into position 31.

Condition Code:

C	V	G	L	
X	0	0	0	Result is zero
X	0	0	1	Result is not zero
X	0	1	0	Result is not zero
1	0	X	X	Carry

Programming Notes:

In the R11 format, the shift count is specified by the least significant five bits of the second operand. The maximum shift count is 31.

In the SF format, the maximum shift count is 15.

The state of the C flag indicates the state of the last bit shifted out of position 0.

If the second operand specifies a shift of zero places, the condition code is set in accordance with the value contained in the register. The C flag is zero in this case.

When the register specified by R1 contains fixed point data, the L flag set indicates a negative result and the G flag set indicates a positive result.

3.5.28 Shift Right Logical (SRL, SRLS)

Shift Right Logical (SRL)
Shift Right Logical Short (SRLS)

<u>Assembler Notation</u>	<u>Opcode</u>	<u>Format</u>
SRL R1,I2(X2)	EC	R11
SRLS R1,N	10	SF

Operation:

The first operand, the contents of the register specified by R1, is shifted right the number of places specified by the second operand. Bits shifted out of position 31 are shifted through the C flag of the condition code and then lost. The last bit shifted remains in the C flag. Zeros are shifted into position 0.

Condition Code:

C	V	G	L	
X	0	0	0	Result is zero
X	0	0	1	Result is not zero
X	0	1	0	Result is not zero
1	0	X	X	Carry

Programming Notes:

In the R11 format, the shift count is specified by the least significant five bits of the second operand. The maximum shift count is 31.

In the SF format, the maximum shift count is 15.

The state of the C flag indicates the state of the last bit shifted out of position 31.

When the register specified by R1 contains fixed point data, the L flag set indicates a negative result; the G flag set indicates a positive result.

If the second operand specifies a shift of zero places, the condition code is set in accordance with the value contained in the register. The C flag is zero in this case.

3.5.29 Shift Left Halfword Logical (SLHL, SLHLS)

Shift Left Halfword Logical (SLHL)
Shift Left Halfword Logical Short (SLHLS)

Assembler Notation	Opcode	Format
SLHL R1,I2(X2)	CD	R11
SLHLS R1,N	91	SF

Operation:

Bits 16:31 of the register specified by R1 are shifted left the number of places specified by the second operand. Bits shifted out of position 16 are shifted through the carry flag and lost. The last bit shifted remains in the C flag. Zeros are shifted into position 31. Bits 0:15 of the first operand remain unchanged.

Condition Code:

C	V	G	L	
X	0	0	0	Result is zero
X	0	0	1	Result is not zero
X	0	1	0	Result is not zero
1	0	X	X	Carry

Programming Notes:

The condition code setting is based on the halfword (bits 16:31) result.

In the R11 format, the shift count is specified by the least significant four bits of the second operand. The maximum shift count is 15.

In the SF format, the maximum shift count is 15.

The state of the C flag indicates the state of the last bit shifted out of position 16.

When the register specified by R1 contains fixed point data, the L flag set indicates a negative result; the G flag set indicates a positive result.

If the second operand specifies a shift of zero places, the condition code is set in accordance with the value contained in bits 16:31 of the register. The C flag is zero in this case.

3.5.30 Shift Right Halfword Logical (SRHL, SRHLS)

Shift Right Halfword Logical (SRHL)

Shift Right Halfword Logical Short (SRHLS)

<u>Assembler Notation</u>	<u>Opcode</u>	<u>Format</u>
SRHL R1,I2(X2)	CC	R11
SRHLS R1,N	90	SF

Operation:

Bits 16:31 of the register specified by R1 are shifted right the number of places specified by the second operand. Bits shifted out of position 31 are shifted through the C flag and lost. The last bit shifted remains in the C flag. Zeros are shifted into position 16. Bits 0:15 of the first operand remain unchanged.

Condition Code:

C	V	G	L	
X	0	0	0	Result is zero
X	0	0	1	Result is not zero
X	0	1	0	Result is not zero
1	0	X	X	Carry

Programming Notes:

The condition code setting is based on the halfword (bits 16:31) result.

In the R11 format, the shift count is specified by the least significant four bits of the second operand. The maximum shift count is 15.

In the SF format, the maximum shift count is 15.

The state of the C flag indicates the state of the last bit shifted out of position 31.

When the register specified by R1 contains fixed point data, the L flag set indicates a negative result; the G flag set indicates a positive result.

If the second operand specifies a shift of zero places, the condition code is set in accordance with the halfword value contained in bits 16:31 of the register. The C flag is zero in this case.

3.5.31 Rotate Left Logical (RLL)

<u>Assembler Notation</u>	<u>Opcode</u>	<u>Format</u>
RLL. R1,I2(X2)	EB	R11

Operation:

The 32-bit first operand, contained in the register specified by R1, is shifted left, end around, the number of positions specified by the second operand. Bits shifted out of position 0 are shifted into position 31.

Condition Code:

C	V	G	L	
0	0	0	0	Result is zero
0	0	0	1	Result is not zero
0	0	1	0	Result is not zero

Programming Notes:

The shift count is specified by the least significant five bits of the second operand. The maximum shift count is 31.

When the register specified by R1 contains fixed point data, the L flag set indicates a negative result; the G flag set indicates a positive result.

If the second operand specifies a shift of zero places, the condition code is set in accordance with the value contained in the register specified by R1.

Example 1:

<u>Assembler Notation</u>	<u>Machine Code</u>	<u>Comments</u>
LI REG9,Y'56789ABC'	F890 56789ABC	(REG9)=56789ABC
RLL REG9,X'0004'	EB90 0004	

Result of RLL Instruction:

(REG9) = 6789ABC5
Condition code = 0010 (G=1)

Example 2:

<u>Assembler Notation</u>	<u>Machine Code</u>	<u>Comments</u>
LI REG9,Y'88880000'	F890 8888 0000	(REG9)=88880000
RLL REG9,X'03'	EB90 0003	

Result of RLL Instruction:

(REG9) = 44400004
Condition code = 0010 (G=1)

3.5.32 Rotate Right Logical (RRL)

<u>Assembler Notation</u>	<u>Opcode</u>	<u>Format</u>
RRL R1,I2(X2)	EA	R11

Operation:

The 32-bit first operand, which is contained in the register specified by R1, is shifted right, end around, the number of positions specified by the second operand. Bits shifted out of position 31 are shifted into position 0.

Condition Code:

C	V	G	L	
0	0	0	0	Result is zero
0	0	0	1	Result is not zero
0	0	1	0	Result is not zero

Programming Notes:

The shift count is specified by the least significant five bits of the second operand. The maximum shift count is 31.

When the register specified by R1 contains fixed point data, the L flag set indicates a negative result; the G flag set indicates a positive result.

If the second operand specifies a shift of zero places, the condition code is set in accordance with the value contained in the register specified by R1.

Example 1:

<u>Assembler Notation</u>	<u>Machine Code</u>	<u>Comments</u>
LI REG4,Y'12345678'	F840 1234 5678	(REG4) = 12345678
RRL REG4,X'04'	EA40 0004	

Result of RRL Instruction:

(REG4) = 81234567
Condition code = 0001 (L=1)

Example 2:

<u>Assembler Notation</u>	<u>Machine Code</u>	<u>Comments</u>
LI REG4,Y'00001111'	F840 0000 1111	(REG4) = 00001111
RRL REG4,X'01'	EA40 0001	

Result of RRL Operation:

(REG4) = '800000888'
Condition code = 0001 (L=1)

3.5.33 Test and Set (TS)

Assembler Notation	Opcode	Format
TS D2(X2)	E0	RX1,RX2
TS A2(FX2,SX2)	E0	RX3

Operation:

The halfword operand is read from memory and, on the same cycle, written back with the most significant bit set. The other bits in the halfword are unchanged. On the read cycle, the most significant bit of the operand is tested. The condition code reflects the state of this bit at the time of the memory read.

Condition Code:

C V G L	

X X X 0	Most significant bit is zero
X X X 1	Most significant bit is set

Programming Notes:

The second operand must be located on a halfword boundary.

The TS instruction provides a mechanism for software synchronization and can be used in a single processor environment as follows: two or more user tasks (u-tasks) running under an operating system share a halfword. This halfword is located in a memory area referred to as task common. Each task can access the halfword using the TS instruction. The synchronization sequence can be as follows:

- TASK 1 sets the most significant bit using the TS instruction.
- TASK 2 senses the most significant bit using the TS instruction, sees that it is set, and performs the necessary software synchronization.

3.5.34 Test Bit (TBT)

<u>Assembler Notation</u>	<u>Opcode</u>	<u>Format</u>
TBT R1,D2(X2)	74	RX1,RX2
TBT R1,A2(FX2,SX2)	74	RX3

Operation:

The second operand address points to a bit array starting on a byte boundary. The value contained in the register specified by R1 is the bit displacement into the array. Bits in the array are counted from left to right starting with bit 0. The argument bit is located and tested. The test does not change the bit.

Condition Code:

C V G L	

0 0 0 0	Tested bit is zero
0 0 1 0	Tested bit is one

Programming Note:

For software compatibility with other processors, the bit array should start on a halfword boundary.

Example:

<u>Assembler Notation</u>	<u>Machine Code</u>	<u>Comments</u>
LIS REG8,3	2483	(REG8) = 3
TBT REG8,LABEL	7480 0BC4	LABEL = halfword in memory at location X'0BC4'. It contains X'B34A'.

Result of TBT Instruction:

Memory location X'BC4' unchanged

(REG8) unchanged

Condition code = 0010 (G=1)...Bit 3 of location X'BC4' is set.

3.5.35 Set Bit (SBT)

<u>Assembler Notation</u>	<u>Opcode</u>	<u>Format</u>
SBT R1,D2(X2)	75	RX1,RX2
SBT R1,A2(FX2,SX2)	75	RX3

Operation:

The second operand address points to a bit array starting on a byte boundary. The value contained in the register specified by R1 is the bit displacement into the array. Bits in the array are counted from left to right starting with bit 0. The argument bit is located and set to one.

Condition Code:

C V G L				

0 0 0 0	Previous state of bit was zero			
0 0 1 0	Previous state of bit was one			

Programming Note:

For software compatibility with other processors, the bit array should start on a halfword boundary.

Example:

<u>Assembler Notation</u>	<u>Machine Code</u>	<u>Comments</u>
LIS REG5,8	2458	(REG5) = 8
SBT REG5,LABEL	7550 1520	LABEL located at X'1520'. It contains X'2134'.

Result of SBT Instruction:

Contents of LABEL = 21B4
(REG5) unchanged
Condition code = 0000 (G=0)

3.5.36 Reset Bit (RBT)

<u>Assembler Notation</u>	<u>Opcode</u>	<u>Format</u>
RBT R1,D2(X2)	76	RX1,RX2
RBT R1,A2(FX2,SX2)	76	RX3

Operation:

The second operand address points to a bit array starting on a byte boundary. The value contained in the register specified by R1 is the bit displacement into the array. Bits in the array are counted from left to right starting with bit zero. The argument bit is located and forced to zero (reset).

Condition Code:

C	V	G	L	
0	0	0	0	Previous state of bit was zero
0	0	1	0	Previous state of bit was one

Programming Note:

For software compatibility with other processors, the bit array should start on a halfword boundary.

Example:

<u>Assembler Notation</u>	<u>Machine Code</u>	<u>Comments</u>
LIS REG2,3	2423	(REG2) = 3
RBT REG2,LABEL	7620 1A42	LABEL located at X'1A42' contains X'3143'.

Result of RBT Instruction:

Contents of LABEL = 2143
 (REG2) unchanged
 Condition code = 0010 (G=1)

3.5.37 Complement Bit (CBT)

<u>Assembler Notation</u>	<u>Opcode</u>	<u>Format</u>
CBT R1,D2(X2)	77	RX1,RX2
CBT R1,A2(FX2,SX2)	77	RX3

Operation:

The second operand address points to a bit array starting on a byte boundary. The value contained in the register specified by R1 is the bit displacement into the array. Bits in the array are counted from left to right starting with bit 0. The argument bit is located and complemented.

Condition Code:

C	V	G	L

0	0	0	0
0	0	1	0

Previous state of bit was zero
Previous state of bit was one

Programming Note:

For software compatibility with other processors, the bit array should start on a halfword boundary.

Example:

<u>Assembler Notation</u>	<u>Machine Code</u>	<u>Comments</u>
LIS REG9,3	2493	(REG9) = 3
CBT REG9,LABEL	7790 0C4A	LABEL located at X'C4A'. It contains X'2813'.

Result of CBT Instruction:

Contents of LABEL = 3813
(REG9) unchanged
Condition code = 0000 (G=0)

3.5.38 Cyclic Redundancy Check (CRC12, CRC16)

Cyclic Redundancy Check Modulo 12 (CRC12)

Cyclic Redundancy Check Modulo 16 (CRC16)

<u>Assembler Notation</u>	<u>Opcode</u>	<u>Format</u>
CRC12 R1,D2(X2)	5E	RX1,RX2
CRC12 R1,A2(FX2,SX2)	5E	RX3
CRC16 R1,D2(X2)	5F	RX1,RX2
CRC16 R1,A2(FX2,SX2)	5F	RX3

Operation:

These instructions are used to generate either a 12-bit or a 16-bit CRC residual halfword. The register specified by R1 contains, in bits 24:31, the data character to be included in the CRC residual. The second operand is the accumulated (old) CRC residual. The polynomial used for the 12-bit CRC generation is:

$$X^{12} + X^{11} + X^3 + X^2 + X + 1$$

The polynomial used for the 16-bit CRC generation is:

$$X^{16} + X^{15} + X^2 + 1$$

The halfword second operand is replaced by the generated CRC residual.

Condition Code:

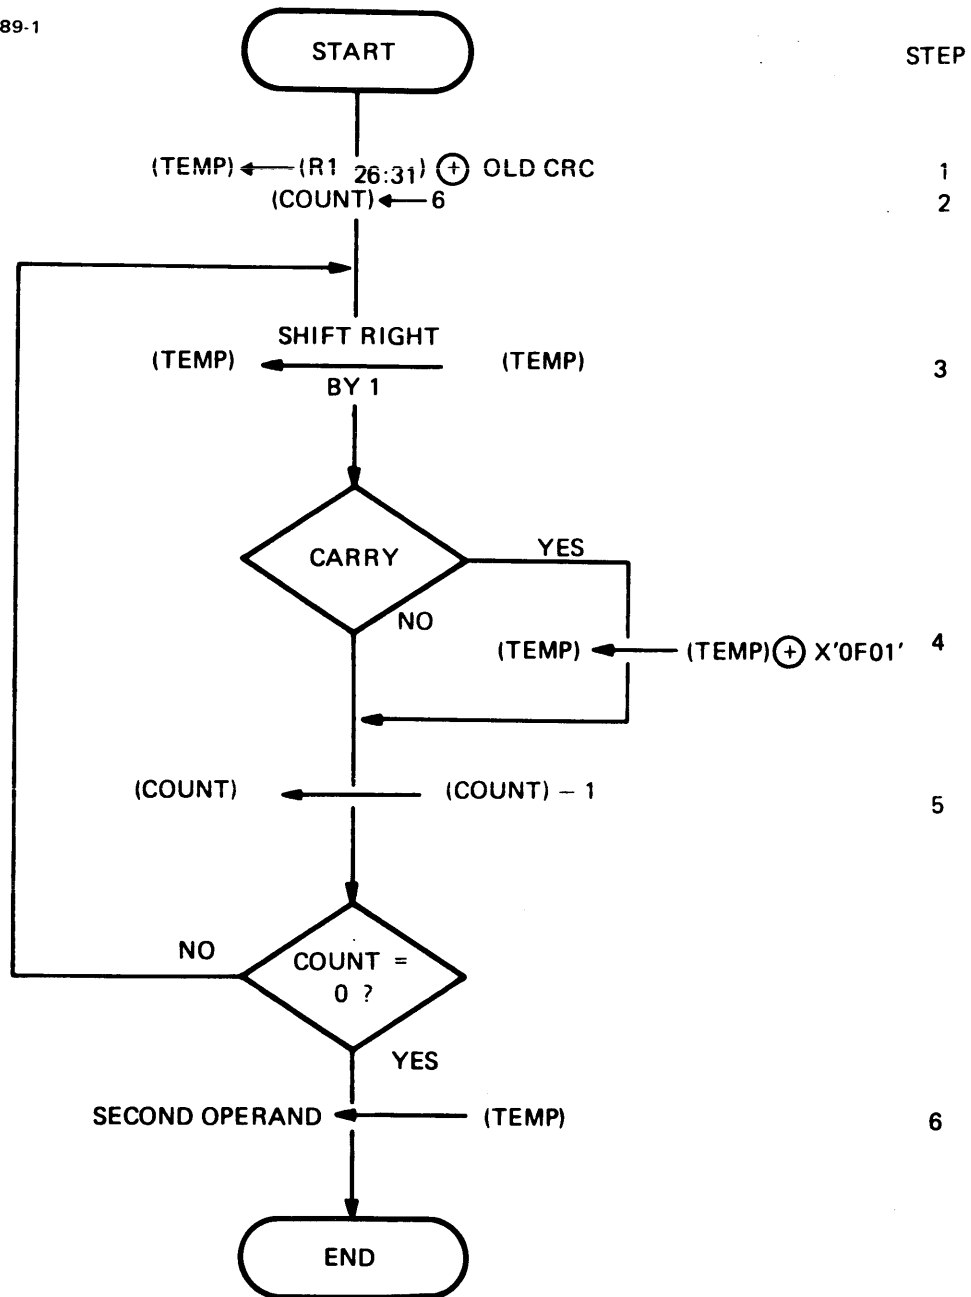
Unchanged

Programming Notes:

The register specified by R1 remains unchanged.

The second operand must be located on a halfword boundary.

Figure 3-6 illustrates a flowchart for CRC generation.



CRC12 ALGORITHM SHOWN

FOR CRC 16 ALGORITHM, USE: R1 24:31 INSTEAD OF R1 26:31 IN STEP 1
 8 INSTEAD OF 6 IN STEP 2
 X'A001' INSTEAD OF X'0F01' IN STEP 4

Figure 3-6 Flowchart for CRC Generation

3.5.39 Translate (TLATE)

Assembler Notation	Opcode	Format
TLATE R1,D2(X2)	E7	RX1,RX2
TLATE R1,A2(FX2,SX2)	E7	RX3

Operation:

The least significant eight bits (bits 24:31) of the register specified by R1 contain the character to be translated. The fullword location specified by the second operand address contains the address of a translation table. The table is made up of 256 halfwords. The character contained in the register specified by R1 is used as an index into the table.

If bit 0 of the table entry corresponding to the index character is one, bits 8:15 of the table entry replace the index character, and the next sequential instruction is executed.


If bit 0 of the table entry is zero, bits 1:15 of the table entry contain the address, divided by two, of a special character handling routine. In this case, no translation takes place. The address contained in bits 1:15 is shifted left by one (multiplied by two). This address replaces the current LOC, thereby effecting an unconditional branch to the special character handling routine. Translation of character string data may also be performed using the MVTU instruction (see Chapter 7).

Condition Code:

Unchanged

Programming Notes:

The second operand must be located on a fullword boundary.



Example:

This example illustrates the use of the TLATE instruction. The translation table must either be initialized or assembled to contain up to a total of 256 halfword entries. In this example, the table contains two entries.

<u>Label</u>	<u>Assembler Notation</u>	<u>Comments</u>
	LHI REG5, X'8052'	LOAD TABLE ENTRY INTO REG5
	STH REG5, TABLE	PUT ENTRY INTO TABLE
	LA REG7, TRANLAB	LOAD ANOTHER TABLE ENTRY
	SRLS REG7, 1	DIVIDE BY 2
	STH REG7, TABLE+4	PUT ENTRY INTO TABLE
	.	
	.	
	.	
TABADR	DC A(TABLE)	

Alternatively, this table can be assembled with the proper constant values. The T type constant can be used to assemble subroutine addresses in the proper format. For example:

```
        ALIGN      2
TABLE EQU      *
        DO         256
        DC         H'0'
        ORG       TABLE+4
        DC         T(TRANLAB)
        ORG       TABLE+512
```

Since a program is normally assembled as a relocatable program, the address of TRANLAB is not known, but for illustrative purposes, assume the address of TRANLAB is X'864'.

	0			15
TABLE+0				
TABLE+2				
TABLE+4	8	0	5	2
TABLE+6				
TABLE+8				
TABLE+10	0	4	3	2
TABLE+12				
TABLE+508				

TABLE+10 contains the address of TRANLAB divided by 2 (X'864'/2).

Example 1:

Using the above table, the following example translates the character in register 2.

<u>Label</u>	<u>Assembler Notation</u>	<u>Comments</u>
	LIS REG2,2	(REG2) = 0000 0002
	TLATE REG2,TABADR	

Result of TLATE Instruction:

(REG2) = 0000 0052
 Condition code unchanged

The entry used = Data at address of (2 times contents of REG2)
 + TABLE
 = Data at address TABLE + 4
 = X'8052'

Since the first bit of the entry is 1, direct translation is used and the contents of REG2 are replaced by X'0000 0052'.

Example 2:

Using the above table, the following example shows how the TLATE instruction can be used to branch to a special character handling routine.

<u>Label</u>	<u>Assembler Notation</u>	<u>Comments</u>
	LIS REG5,5	(REG5) = 0000 0005
	TLATE REG5,TABADR	
	.	
	.	
	.	
	.	
TRANLAB	LR R6,R5	THESE INSTRUCTIONS
	LB R3,0(R6)	OPERATE ON THE
	.	SPECIAL CHARACTER.
	.	
	.	
	.	
	.	

Result of TLATE Instruction:

(REG5) = 0000 0005
Condition code unchanged

Control is transferred to the subroutine at address TRANLAB (X'864').

The entry used = Data at address of (2 times contents of REG5)
+ TABLE
= Data at address TABLE + A
= X'0432'

Since the first bit of the entry is 0, the entry is multiplied by 2, a transfer occurs to TRANLAB (at address X'864'), and the processor executes instructions from the new address.

3.5.40 Add To List (ATL, ABL)

Add to Top of List (ATL)

Add to Bottom of List (ABL)

Assembler Notation	Opcode	Format
ATL R1,D2(X2)	64	RX1,RX2
ATL R1,A2(FX2,SX2)	64	RX3
ABL R1,D2(X2)	65	RX1,RX2
ABL R1,A2(FX2,SX2)	65	RX3

Operation:

The register specified by R1 contains the fullword element to be added to the list, which is located in memory at the address of the second operand. The tally of the number of slots used is compared with the number of slots in the list. If the number of slots used equals the number of slots in the list, an overflow condition exists. The element is not added to the list and the overflow flag in the condition code is set.

If the tally of the number of slots used is less than the number of slots in the list, it is incremented by one, the appropriate pointer is changed, and the element is added to the list. See Figure 3-4.

Condition Code:

C	V	G	L	
0	0	0	0	Element added successfully
0	1	0	0	List overflow

Programming Notes:

These instructions manipulate circular lists as described in the introduction to this chapter.

The second operand must be located on a fullword boundary.

The ATL instruction manipulates the current top pointer in the list. If no overflow occurs, the current top pointer, which points to the last element added to the top of the list, is decremented by one. The element is inserted in the slot pointed to by the new current top pointer. If the current top pointer was zero on entering this instruction, the current top pointer is set to the maximum slot number in the list. This condition is referred to as list wrap.

The ABL instruction manipulates the next bottom pointer. If no overflow occurs, the element is inserted in the slot pointed to by the next bottom pointer, and the next bottom pointer is incremented by one. If the incremented next bottom pointer is greater than the maximum slot number in the list, the next bottom pointer is set to zero. This condition is referred to as list wrap.

For the nonoverflow situation, pointer halfwords in the list header are not manipulated until after the element has been successfully added. This facilitates error recovery in the event of a memory fault.

See the examples in Section 3.5.41.

3.5.41 Remove From List (RTL, RBL)

Remove from Top of List (RTL)
Remove from Bottom of List (RBL)

<u>Assembler Notation</u>	<u>Opcode</u>	<u>Format</u>
RTL R1,D2(X2)	66	RX1,RX2
RTL R1,A2(FX2,SX2)	66	RX3
RBL R1,D2(X2)	67	RX1,RX2
RBL R1,A2(FX2,SX2)	67	RX3

Operation:

The element removed from the list replaces the contents of the register specified by R1. The list is located at the address of the second operand. If, at the start of the instruction execution, the tally of the number of slots used is zero, then the list is already empty and the instruction terminates with the overflow flag set in the condition code. This condition is referred to as list underflow; in this case, R1 is undefined. If underflow does not occur, the appropriate pointer is changed, the element is extracted and placed in the register specified by R1, and the number of slots used tally is decremented by one.

Condition Code:

C	V	G	L	
=====				
0	0	0	0	List now empty
0	0	1	0	List is not yet empty
0	1	0	0	List was already empty

Programming Notes:

These instructions manipulate circular lists as described in the introduction to this chapter.

The second operand must be located on a fullword boundary.

In the case of list underflow, the contents of the register specified by R1 are unchanged.

The RTL instruction manipulates the current top pointer. If no underflow occurs, the current top pointer points to the element to be extracted. The element is extracted and placed in the register specified by R1. The current top pointer is incremented by one and compared to the maximum slot number. If the current top pointer is greater than the maximum slot number, the current top pointer is set to zero. This condition is referred to as list wrap.

The RBL instruction manipulates the next bottom pointer. If no underflow occurs and the next bottom pointer is zero, it is set to the maximum slot number (list wrap); otherwise, it is decremented by one, and the element now pointed to is extracted and placed in the register specified by R1.

For the nonunderflow situation, pointer halfwords in the list header are not manipulated until after the element has been successfully removed. The register specified by R1 is not modified until the header has been updated. This facilitates error recovery in the event of a memory fault.

Examples:

The following are examples of the use of the four list processing instructions (ATL, ABL, RTL, RBL).

The original list is normally set up as shown in Figure 3-7.

590

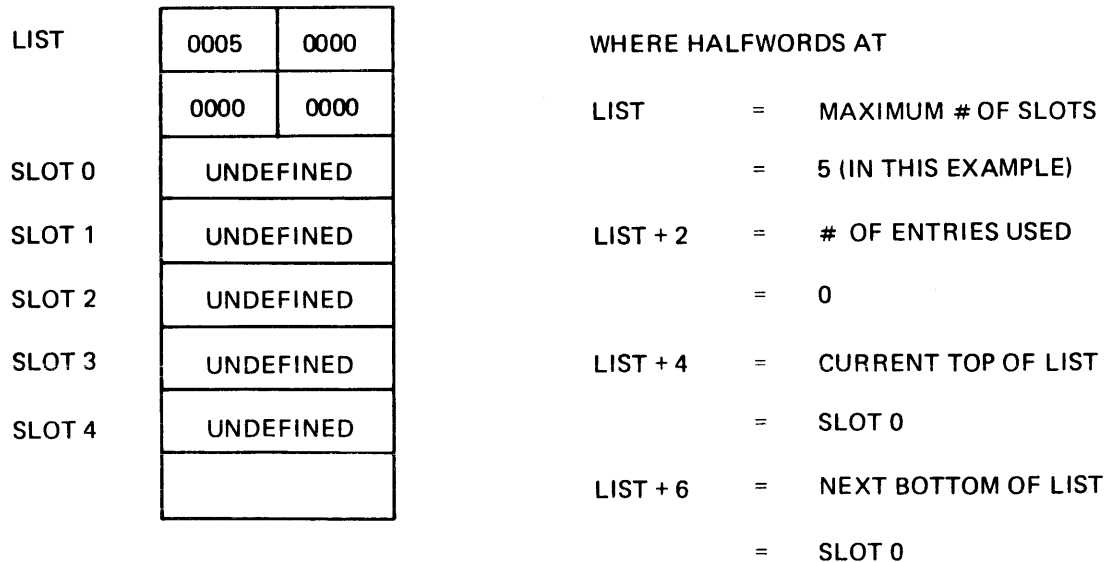


Figure 3-7 List Processing Instructions

Assembler Notation

Results and Comments

```

LIS  REG0,0
STH  REG0,LIST+2      INITIALIZE NUMBER OF ENTRIES
                        USED TO 0
ST   REG0,LIST+4      INITIALIZE POINTERS TO 0
LIS  REG1,1           REGISTERS 1 THROUGH 6 CONTAIN
LIS  REG2,2           1 THROUGH 6 RESPECTIVELY
LIS  REG3,3
LIS  REG4,4
LIS  REG5,5
LIS  REG6,6
STH  REG5,LIST        TOTAL NUMBER OF ENTRIES = 5

```

REF1 ATL REG1,LIST

	LIST	0005	0001
		0004	0000
SLOT 0		UNDEF INED	
SLOT 1		UNDEF INED	
SLOT 2		UNDEF INED	
SLOT 3		UNDEF INED	
SLOT 4		0000	0001

(List Wrap)

```

Condition code = 0000
Current top pointer = Slot 4
Next bottom pointer = Slot 0

```

REF2	ATL REG2,LIST	LIST	0005	0002
			0003	0000
	SLOT 0		UNDEF INED	
	SLOT 1		UNDEF INED	
	SLOT 2		UNDEF INED	
	SLOT 3		0000	0002
	SLOT 4		0000	0001

Condition code = 0000
 Current top pointer = Slot 3
 Next bottom pointer = Slot 0

REF3	ATL REG3,LIST	LIST	0005	0003
			0002	0000
	SLOT 0		UNDEF INED	
	SLOT 1		UNDEF INED	
	SLOT 2		0000	0003
	SLOT 3		0000	0002
	SLOT 4		0000	0001

Condition code = 0000
 Current top pointer = Slot 2
 Next bottom pointer = Slot 0

REF4	ABL REG4,LIST	LIST	0005	0004
			0002	0001
	SLOT 0		0000	0004
	SLOT 1	UNDEFINED		
	SLOT 2		0000	0003
	SLOT 3		0000	0002
	SLOT 4		0000	0001

Condition code = 0000
 Current top pointer = Slot 2
 Next bottom pointer = Slot 1

REF5	ABL REG5,LIST	LIST	0005	0005
			0002	0002
	SLOT 0		0000	0004
	SLOT 1		0000	0005
	SLOT 2		0000	0003
	SLOT 3		0000	0002
	SLOT 4		0000	0001

Condition code = 0000
 Current top pointer = Slot 2
 Next bottom pointer = Slot 2

REF6	ABL REG6,LIST	LIST	0005	0005
			0002	0002
	SLOT 0		0000	0004
	SLOT 1		0000	0005
	SLOT 2		0000	0003
	SLOT 3		0000	0002
	SLOT 4		0000	0001

Condition code = 0100 (List overflow)
 Current top pointer = Slot 2
 Next bottom pointer = Slot 2

REF7	RTL REG7,LIST	LIST	0005	0004
			0003	0002
	SLOT 0		0000	0004
	SLOT 1		0000	0005
	SLOT 2 X		0000	0003
	SLOT 3		0000	0002
	SLOT 4		0000	0001

(REG7) = 0000 0003
 Condition code = 0010
 Current top pointer = Slot 3
 Next bottom pointer = Slot 2

NOTE

X indicates an entry was removed from the list and is not accessible through further manipulation of list instructions.

REF8	RBL REG8,LIST	LIST	0005	0003
			0003	0001
	SLOT 0		0000	0004
	SLOT 1 X		0000	0005
	SLOT 2 X		0000	0003
	SLOT 3		0000	0002
	SLOT 4		0000	0001

(REG8) = 0000 0005
 Condition code = 0010
 Current top pointer = Slot 3
 Next bottom pointer = Slot 1

NOTE

X indicates an entry was removed from the list and is not accessible through further manipulation of list instructions.

REF9	RTL REG9,LIST	LIST	0005	0002
			0004	0001
	SLOT 0		0000	0004
	SLOT 1 X		0000	0005
	SLOT 2 X		0000	0003
	SLOT 3 X		0000	0002
	SLOT 4		0000	0001

(REG9) = 0000 0002
 Condition code = 0010
 Current top pointer = Slot 4
 Next bottom pointer = Slot 1

REF10	RBL	REG10,LIST	I,LIST	0005	0001
				0004	0000
		SLOT 0	X	0000	0004
		SLOT 1	X	0000	0005
		SLOT 2	X	0000	0003
		SLOT 3	X	0000	0002
		SLOT 4		0000	0001

(REG10) = 0000 0004
Condition code = 0010
Current top pointer = 4
Next bottom pointer = 0

NOTE

X indicates an entry was removed from the list and is not accessible through further manipulation of list instructions.

REF11	RTL	REG11,LIST	I,LIST	0005	0000
				0000	0000
		SLOT 0	X	0000	0004
		SLOT 1	X	0000	0005
		SLOT 2	X	0000	0003
		SLOT 3	X	0000	0002
		SLOT 4	X	0000	0001

(REG11) = 0000 0001
Condition code = 0000 (List is now empty)
Current top pointer = 0
Next bottom pointer = 0

REF12	RTL	REG12,LIST	LIST		
				0005	0000
				0000	0000
		SLOT 0	X	0000	0004
		SLOT 1	X	0000	0005
		SLOT 2	X	0000	0003
		SLOT 3	X	0000	0002
		SLOT 4	X	0000	0001

(REG12) = UNDEFINED
Condition code = 0100*
Current top pointer = 0
Next bottom pointer = 0

* List was already empty

NOTE

X indicates an entry was removed from the list and is not accessible through further manipulation of list instructions.

CHAPTER 4 BRANCHING

4.1 INTRODUCTION

In normal operations, the processor executes instructions in sequential order. The branch instructions allow this sequential mode of operation to be varied, so that programs can loop, transfer control to subroutines, or make decisions based on the results of previous operations.

4.2 OPERATIONS

The second operand of a branch instruction is the address of the memory location to which control is transferred. The address can be contained in a register, or it can be specified in the instruction as the second operand address or as a displacement.

4.2.1 Decision Making

The conditional branch instructions permit the program to make decisions based on some result. In these instructions, the R1 field contains a 4-bit mask, M1, which is tested by ANDing it with the condition code. The result of the test determines whether the branch is taken, or the next sequential instruction is executed.

Table 4-1 shows previous condition codes, masks specified in a branch instruction, and the results of the test on which the branch or no branch decision was made.

TABLE 4-1 DECISION TABLE

CONDITION CODE	MASK(M1)	RESULT OF TEST	(TRUE/FALSE)	BRANCH TRUE TAKEN	BRANCH FALSE TAKEN
0000	0010	0000	(False)	No	Yes
0001	1010	0000	(False)	No	Yes
1001	1000	1000	(True)	Yes	No

TABLE 4-1 DECISION TABLE (Continued)

CONDITION CODE	MASK(M1)	RESULT OF TEST	(TRUE/ FALSE)	BRANCH TRUE TAKEN	BRANCH FALSE TAKEN
0100	0100	0100	(True)	Yes	No
1010	0010	0010	(True)	Yes	No
0010	0011	0010	(True)	Yes	No
0010	0000	0000	(False)	No	Yes

4.2.2 Subroutine Linkage

The branch and link instructions allow branching to subroutines in such a way that a return address is passed to the subroutine. For these instructions, the address of the memory location immediately following the branch instruction is saved in the register specified by R1.

4.3 BRANCH INSTRUCTION FORMATS

The branch instructions use the Register-to-Register (RR), Short Form (SF), and Register and Indexed Storage (RX) formats.

4.4 BRANCH INSTRUCTIONS

The instructions described in this section are:

BTC	Branch on True Condition
BTCR	Branch on True Condition Register
BTBS	Branch on True Condition Backward Short
BTFS	Branch on True Condition Forward Short
BFC	Branch on False Condition
BFCR	Branch on False Condition Register
BFBS	Branch on False Condition Backward Short
BFFS	Branch on False Condition Forward Short
BAL	Branch and Link
BALR	Branch and Link Register
BXLE	Branch on Index Low or Equal
BXH	Branch on Index High

4.4.1 Branch on True (BTC, BTCR, BTBS, BTFS)

Branch on True Condition (BTC)

Branch on True Condition Register (BTCR)

Branch on True Condition Backward Short (BTBS)

Branch on True Condition Forward Short (BTFS)

<u>Assembler Notation</u>	<u>Opcode</u>	<u>Format</u>
BTC M1,D2(X2)	42	RX1,RX2
BTC M1,A2(FX2,SX2)	42	RX3
BTCR M1,R2	02	RR
BTBS M1,N	20	SF
BTFS M1,N	21	SF

Operation:

The condition code of the program status word (PSW) is tested for the conditions specified by the mask field, M1. If any conditions tested are found to be true, a branch is taken to the second operand location. If none of the conditions tested are found to be true, the next sequential instruction is executed.

Condition Code:

Unchanged

Programming Notes:

In the RR format, the branch address is contained in the register specified by R2.

In the SF format, the N field contains the number of halfwords to be added to or subtracted from the current location counter (LOC) to obtain the branch address.

In the RR and RX formats, the branch address must be located on a halfword boundary.

Example:

The following example illustrates the use of the BTC instruction.

<u>Assembler Notation</u>	<u>Machine Code</u>	<u>Comments</u>
LH R1,X'100'	4810 0100	Load halfword (X'1234') located at X'100'. Condition code is set to CVGL = 0010. Mask is 3 (i.e., M1 = 0011). Perform logical AND between CVGL and M1 (i.e., 0010 AND 0011). The result is 0010 (i.e., true); therefore, a branch is taken to LOC.
BTC 3,LOC	4230 ABC0	

4.4.2 Branch on False (BFC, BFCR, BFBS, BFFS)

Branch on False Condition (BFC)

Branch on False Condition Register (BFCR)

Branch on False Condition Backward Short (BFBS)

Branch on False Condition Forward Short (BFFS)

<u>Assembler Notation</u>	<u>Opcode</u>	<u>Format</u>
BFC M1,D2(X2)	43	RX1,RX2
BFC M1,A2(FX2,SX2)	43	RX3
BFCR M1,R2	03	RR
BFBS M1,N	22	SF
BFFS M1,N	23	SF

Operation:

The condition code of the PSW is tested for the conditions specified in the mask field, M1. If all conditions tested are found to be false, a branch is taken to the second operand location. If any of the conditions tested are found to be true, the next sequential instruction is executed.

Condition Code:

Unchanged

Programming Notes:

In the RR format, the branch address is contained in the register specified by R2.

In the SF format, the N field contains the number of halfwords to be added to or subtracted from the current LOC to obtain the branch address.

In the RR and RX formats, the branch address must be located on a halfword boundary.

Example:

The following example illustrates the use of the BFC instruction.

<u>Assembler Notation</u>	<u>Machine Code</u>	<u>Comments</u>
LCS R1,2	2512	(R1) = FFFFFFFE. Condition code is set to CVGL = 0001. Mask is 1001. Perform logical AND between M1 and CVGL (i.e., 1001 AND 0001). The result is 0001 (i.e., true); therefore, a branch is not taken to LOC.
BFC 9,LOC	4390 ABC0	

4.4.3 Branch and Link (BAL, BALR)

Branch and Link (BAL)

Branch and Link Register (BALR)

<u>Assembler Notation</u>	<u>Opcode</u>	<u>Format</u>
BAL R1,D2(X2)	41	RX1,RX2
BAL R1,A2(FX2,SX2)	41	RX3
BALR R1,R2	01	RR

Operation:

The address of the next sequential instruction is saved in the register specified by R1, and a branch is taken to the second operand address.

Condition Code:

Unchanged

Programming Notes:

The second operand must be located on a halfword boundary.

The branch address is calculated before the register specified by R1 is changed. R1 can specify the same register as X2, FX2, SX2 or R2.

Example:

The following example illustrates the use of the BAL instruction. This instruction causes control to be transferred to a subroutine called SUBROUT. After completion of the subroutine, the linking register is used to branch back to the next sequential instruction after the BAL; i.e., the instruction labeled RETURN.

	<u>Label</u>	<u>Assembler Notation</u>	<u>Comments</u>
MAIN	BEGIN	BAL REG4, SUBROUT	TRANSFER TO SUBROUT
	RETURN	XR R6, R6	
PROG		STH R6, LAB+4	
		.	
		.	
SUBROUTINE	SUBROUT	LHL R8, LOC	THE RETURN ADDRESS OF THE SUBROUTINE IS IN REG4
		AHI R8, 10	
		.	
		.	
	RTNEND	BR REG4	RETURN TO XR INST.

NOTE

The linking register (REG4 in the example) should not be used within the subroutine.


Result of BAL Instruction:

(REG4) = Address of instruction at SUBROUT
Condition code unchanged

4.4.4 Branch on Index Low or Equal (BXLE)

<u>Assembler Notation</u>	<u>Opcode</u>	<u>Format</u>
BXLE R1,D2(X2)	C1	RX1,RX2
BXLE R1,A2(FX2,SX2)	C1	RX3

Set Up:



Before execution of this instruction, the register specified by R1 must contain a starting index value. The register specified by R1+1 must contain an increment value. The register specified by R1+2 must contain a comparand (limit or final value). All values can be signed.

Operation:

Execution of this instruction causes the increment value to be added to the index value, creating a new index value. The result is compared logically to the limit or final value. If the new index value is less than or equal to the limit value, a branch is taken to the second operand location. If the new index value is greater than the limit value, the next sequential instruction is executed.

Condition Code:

Unchanged

Programming Notes:

The incremented index value replaces the contents of the register specified by R1.

Any three consecutive registers of the same set can be used by this instruction as specified by R1. These registers can be 6, 7, 8; or 14, 15, 0; or 15, 0, 1; etc.

The second operand must be located on a halfword boundary.

The branch address is calculated before incrementing the starting index value contained in the register specified by R1.

R1 can specify the same register as X2, FX2 or SX2.

Example:

Transfer 10 bytes in memory starting at the memory location labeled BUF0 to the memory location labeled BUF1.

<u>Label</u>	<u>Assembler Notation</u>	<u>Comments</u>
	LIS REG3,0	(REG3)=STARTING INDEX VALUE=0
	LIS REG4,1	(REG4)=INCREMENT VALUE
	LIS R5,9	(REG5)=FINAL VALUE=9
AGAIN	LB REG0,BUF0(R3)	(REG0)=1 BYTE FROM BUF0
	STB REG0,BUF1(R1)	COPY 1 BYTE TO BUF1
LABEL	BXLE R3,AGAIN	IF (REG3)>(REG5),DONE
	.	
	.	
	.	
BUF0	DS 10	
BUF1	DS 10	

Result of BXLE Instruction:

Code between the instructions labeled AGAIN and LABEL is executed ten times.

Condition code unchanged by BXLE instruction

(REG3) = 0000000A


(REG4) = 00000001

(REG5) = 00000009

4.4.5 Branch on Index High (BXH)

Assembler Notation	Opcode	Format
BXH R1,D2(X2)	C0	RX1,RX2
BXH R1,A2(FX2,SX2)	C0	RX3

Set Up:



Before execution of this instruction, the register specified by R1 must contain a starting index value. The register specified by R1+1 must contain an increment value. The register specified by R1+2 must contain a comparand (limit or final value). All values can be signed.

Operation:

Execution of this instruction causes the increment value to be added to the index value, creating a new index value. The result is logically compared to the limit or final value. If the new index value is greater than the limit value, a branch is taken to the second operand location. If the new index value is less than or equal to the limit value, the next sequential instruction is executed.

Condition Code:

Unchanged

Programming Notes:

The incremented index value replaces the contents of the register specified by R1.

Any three consecutive registers of the same set can be used by this instruction as specified by R1. These registers can be 6, 7, 8; or 14, 15, 0; or 15, 0, 1; etc.

The second operand must be located on a halfword boundary.

The branch address is calculated before incrementing the starting index value contained in the register specified by R1.

R1 can specify the same register as X2, FX2 or SX2.

Example:

The following example shows how to set up a counter (1-9) using the BXH instruction.

<u>Label</u>	<u>Assembler Notation</u>	<u>Comment</u>
	LIS REG1,1	(REG1)=0000 0001 (INDEX)
	LIS REG2,1	(REG2)=0000 0001 (INCREMENT)
	LIS REG3,9	(REG3)=0000 0009 (COMPARAND)
BEGIN	BXH REG1,LABEL	COMPARE INDEX WITH COMPARAND
	LH R6,COUNT	
	.	
	.	
	.	
	B BEGIN	BRANCH TO BXH INSTRUCTION
LABEL	LA R8,RTN	EXIT FROM BXH
	ST R8,MEM	

Result of BXH Instruction:

Code between the instructions labeled BEGIN and LABEL is executed nine times.

Condition code unchanged by BXH instruction

(REG1) = 0000 000A
(REG2) = 0000 0001
(REG3) = 0000 0009

4.5 EXTENDED BRANCH MNEMONICS

The common assembly language (CAL) assembler supports 47 extended branch mnemonics that generate the branch opcode (true or false conditional) and the condition code mask required. The programmer must supply the second operand address (symbolic or absolute). In the case of SF branch instructions, the second operand branch address must be within 15 halfwords of the LOC. The CAL assembler determines the backward or forward relationship of the second operand address and generates the appropriate operation code.

The instructions described in this section are:

BC	Branch on Carry	
BCR	Branch on Carry Register	
BCS	Branch on Carry Short	
BNC	Branch on No Carry	
BNCR	Branch on No Carry Register	
BNCS	Branch on No Carry Short	
BE	Branch on Equal	
BER	Branch on Equal Register	=
BES	Branch on Equal Short	
BNE	Branch on Not Equal	
BNER	Branch on Not Equal Register	≠
BNES	Branch on Not Equal Short	
BL	Branch on Low	
BLR	Branch on Low Register	<
BLS	Branch on Low Short	
BNL	Branch on Not Low	
BNLR	Branch on Not Low Register	≥
BNLS	Branch on Not Low Short	
BM	Branch on Minus	
BMR	Branch on Minus Register	
BMS	Branch on Minus Short	
BNM	Branch on Not Minus	
BNMR	Branch on Not Minus Register	
BNMS	Branch on Not Minus Short	
BP	Branch on Plus	
BPR	Branch on Plus Register	
BPS	Branch on Plus Short	
BNP	Branch on Not Plus	
BNPR	Branch on Not Plus Register	
BNPS	Branch on Not Plus Short	

BO	Branch on Overflow
BOR	Branch on Overflow Register
BOS	Branch on Overflow Short
BNO	Branch on No Overflow
BNOR	Branch on No Overflow Register
BNOS	Branch on No Overflow Short
BZ	Branch on Zero
BZR	Branch on Zero Register
BZS	Branch on Zero Short
BNZ	Branch on Not Zero
BNZR	Branch on Not Zero Register
BNZS	Branch on Not Zero Short
B	Branch (Unconditional)
BR	Branch Register (Unconditional)
BS	Branch Short (Unconditional)
NOP	No Operation
NOPR	No Operation Register

4.5.1 Branch on Carry (BC, BCR, BCS)

Branch on Carry (BC)
Branch on Carry Register (BCR)
Branch on Carry Short (BCS)

<u>Assembler Notation</u>	<u>Opcode+M1</u>	<u>Format</u>
BC D2(X2)	428	RX1,RX2
BC A2(FX2,SX2)	428	RX3
BCR R2	028	RR
BCS A	208 (Backward) 218 (Forward)	SF

Operation:

If the carry (C) flag in the condition code is set, a branch is taken to the second operand location. If the C flag is zero, the next sequential instruction is executed.

Condition Code:

Unchanged

Programming Notes:

The branch destination must be located on a halfword boundary.

In the RR format, the branch address is contained in the register specified by R2.

Example:

The following example illustrates the use of the BCS instruction.

<u>Assembler Notation</u>	<u>Machine Code</u>	<u>Comments</u>
SHIFT SLLS R9,1	1191	Register 9 is shifted left until the first zero bit is shifted out of position 0.
BCS SHIFT	2081	

4.5.2 Branch on No Carry (BNC, BNCR, BNCS)

Branch on No Carry (BNC)

Branch on No Carry Register (BNCR)

Branch on No Carry Short (BNCS)

<u>Assembler Notation</u>	<u>Opcode+M1</u>	<u>Format</u>
BNC D2(X2)	438	RX1,RX2
BNC A2(FX2,SX2)	438	RX3
BNCR R2	038	RR
BNCS A	228 (Backward)	SF
	238 (Forward)	

Operation:

If the C flag in the condition code is zero, a branch is taken to the second operand location. If the C flag is set, the next sequential instruction is executed.

Condition Code:

Unchanged

Programming Notes:

The branch destination must be located on a halfword boundary.

In the RR format, the branch address is contained in the register specified by R2.

4.5.3 Branch on Equal (BE, BER, BES)

Branch on Equal (BE)

Branch on Equal Register (BER)

Branch on Equal Short (BES)

<u>Assembler Notation</u>	<u>Opcode+M1</u>	<u>Format</u>
BE D2(X2)	433	RX1,RX2
BE A2(FX2,SX2)	433	RX3
BER R2	033	RR
BES A	223 (Backward)	SF
	233 (Forward)	

Operation:

If the greater than (G) and less than (L) flags are both zero in the condition code, a branch is taken to the second operand location. If either flag is set, the next sequential instruction is executed.

Condition Code:

Unchanged

Programming Notes:

The branch destination must be located on a halfword boundary.

In the RR format, the branch address is contained in the register specified by R2.

Example:

The following example illustrates the use of the BE instruction.

<u>Assembler Notation</u>	<u>Machine Code</u>	<u>Comments</u>
CLHI R4,X'23'	C540 0023	If R4 contains X'23', a branch is taken to location X'A00'. If not, the next sequential instruction is executed.
BE OPTIN	4330 0A00	

4.5.4 Branch on Not Equal (BNE, BNER, BNES)

Branch on Not Equal (BNE)
Branch on Not Equal Register (BNER)
Branch on Not Equal Short (BNES)

<u>Assembler Notation</u>	<u>Opcode+M1</u>	<u>Format</u>
BNE D2(X2)	423	RX1,RX2
BNE A2(FX2,SX2)	423	RX3
BNER R2	023	RR
BNES A	203 (Backward)	SF
	213 (Forward)	

Operation:

If the G or L flag is set in the condition code, a branch is taken to the second operand location. If both flags are zero, the next sequential instruction is executed.

Condition Code:

Unchanged

Programming Notes:

The branch destination must be located on a halfword boundary.

In the RR format, the branch address is contained in the register specified by R2.

4.5.5 Branch on Low (BL, BLR, BLS)

Branch on Low (BL)

Branch on Low Register (BLR)

Branch on Low Short (BLS)

<u>Assembler Notation</u>	<u>Opcode+Ml</u>	<u>Format</u>
BL D2(X2)	428	RX1,RX2
BL A2(FX2,SX2)	428	RX3
BLR R2	028	RR
BLS A	208 (Backward)	SF
	218 (Forward)	

Operation:

If the C flag in the condition code is set, a branch is taken to the second operand address. If the C flag is zero, the next sequential instruction is executed.

Condition Code:

Unchanged

Programming Notes:

The branch destination must be located on a halfword boundary.

In the RR format, the branch address is contained in the register specified by R2.

Example:

The following example illustrates the use of the BL instruction.

<u>Assembler Notation</u>	<u>Machine Code</u>	<u>Comments</u>
CLHI R1,X'FF'	C510 00FF	(R1) is compared to X'00FF'. If (R1) is less than X'00FF', a branch is taken to memory location X'0A00'.
BL RESTART	4280 0A00	

4.5.6 Branch on Not Low (BNL, BNLR, BNLS)

Branch on Not Low (BNL)
Branch on Not Low Register (BNLR)
Branch on Not Low Short (BNLS)

<u>Assembler Notation</u>	<u>Opcode+M1</u>	<u>Format</u>
BNL D2(X2)	438	RX1, RX2
BNL A2(FX2,SX2)	438	RX3
BNLR R2	038	RR
BNLS A	228 (Backward)	SF
	238 (Forward)	

Operation:

If the C flag in the condition code is zero, a branch is taken to the second operand address. If the C flag is set, the next sequential instruction is executed.

Condition Code:

Unchanged

Programming Notes:

The branch destination must be located on a halfword boundary.

In the RR format, the branch address is contained in the register specified by R2.

4.5.7 Branch on Minus (BM, BMR, BMS)

Branch on Minus (BM)

Branch on Minus Register (BMR)

Branch on Minus Short (BMS)

<u>Assembler Notation</u>	<u>Opcode+M1</u>	<u>Format</u>
BM D2(X2)	421	RX1,RX2
BM A2(FX2,SX2)	421	RX3
BMR R2	021	RR
BMS A	201 (Backward) 211 (Forward)	SF

Operation:

If the L flag in the condition code is set, a branch is taken to the second operand location. If the L flag is zero, the next sequential instruction is executed.

Condition Code:

Unchanged

Programming Notes:

The branch destination must be located on a halfword boundary.

In the RR format, the branch address is contained in the register specified by R2.

Example:

The following example illustrates the use of the BM instruction.

<u>Assembler Notation</u>	<u>Machine Code</u>	<u>Comments</u>
SIS R3,1	2631	If (R3) is less than 0 after the subtraction, a branch is taken to X'10A0'.
BM CONTINUE	4210 10A0	

4.5.8 Branch on Not Minus (BNM, BNMR, BNMS)

Branch on Not Minus (BNM)

Branch on Not Minus Register (BNMR)

Branch on Not Minus Short (BNMS)

<u>Assembler Notation</u>	<u>Opcode+M1</u>	<u>Format</u>
BNM D2(X2)	431	RX1,RX2
BNM A2(FX2,SX2)	431	RX3
BNMR R2	031	RR
BNMS A	221 (Backward)	SF
	231 (Forward)	

Operation:

If the L flag in the condition code is zero, a branch is taken to the second operand location. If the L flag is set, the next sequential instruction is executed.

Condition Code:

Unchanged

Programming Notes:

The branch destination must be located on a halfword boundary.

In the RR format, the branch address is contained in the register specified by R2.

4.5.9 Branch on Plus (BP, BPR, BPS)

Branch on Plus (BP)

Branch on Plus Register (BPR)

Branch on Plus Short (BPS)

<u>Assembler Notation</u>	<u>Opcode+M1</u>	<u>Format</u>
BP D2(X2)	422	RX1,RX2
BP A2(FX2,SX2)	422	RX3
BPR R2	022	RR
BPS A	202 (Backward)	SF
	212 (Forward)	

Operation:

If the G flag in the condition code is set, a branch is taken to the second operand location. If the G flag is zero, the next sequential instruction is executed.

Condition Code:

Unchanged

Programming Notes:

The branch destination must be located on a halfword boundary.

In the RR format, the branch address is contained in the register specified by R2.

4.5.10 Branch on Not Plus (BNP, BNPR, BNPS)

Branch on Not Plus (BNP)

Branch on Not Plus Register (BNPR)

Branch on Not Plus Short (BNPS)

<u>Assembler Notation</u>	<u>Opcode+M1</u>	<u>Format</u>
BNP D2(X2)	432	RX1,RX2
BNP A2(FX2,SX2)	432	RX3
BNPR R2	032	RR
BNPS A	222 (Backward)	SF
	232 (Forward)	

Operation:

If the G flag in the condition code is zero, a branch is taken to the second operand location. If the G flag is set, the next sequential instruction is executed.

Condition Code:

Unchanged

Programming Notes:

The branch destination must be located on a halfword boundary.

In the RR format, the branch address is contained in the register specified by R2.

4.5.11 Branch on Overflow (BO, BOR, BOS)

Branch on Overflow (BO)
Branch on Overflow Register (BOR)
Branch on Overflow Short (BOS)

<u>Assembler Notation</u>	<u>Opcode+M1</u>	<u>Format</u>
BO D2(X2)	424	RX1,RX2
BO A2(FX2,SX2)	424	RX3
BOR R2	024	RR
BOS A	204 (Backward)	SF
	214 (Forward)	

Operation:

If the overflow (V) flag in the condition code is set, a branch is taken to the second operand location. If the V flag is zero, the next sequential instruction is executed.

Condition Code:

Unchanged

Programming Notes:

The branch destination must be located on a halfword boundary.

In the RR format, the branch address is contained in the register specified by R2.

4.5.12 Branch on No Overflow (BNO, BNOR, BNOS)

Branch on No Overflow (BNO)
Branch on No Overflow Register (BNOR)
Branch on No Overflow Short (BNOS)

<u>Assembler Notation</u>	<u>Opcode+M1</u>	<u>Format</u>
BNO D2(X2)	434	RX1,RX2
BNO A2(FX2,SX2)	434	RX3
BNOR R2	034	RR
BNOS A	224 (Backward)	SF
	234 (Forward)	

Operation:

If the V flag in the condition code is zero, a branch is taken to the second operand location. If the V flag is set, the next sequential instruction is executed.

Condition Code:

Unchanged

Programming Notes:

The branch destination must be located on a halfword boundary.

In the RR format, the branch address is contained in the register specified by R2.

4.5.13 Branch on Zero (BZ, BZR, BZS)

Branch on Zero (BZ)
Branch on Zero Register (BZR)
Branch on Zero Short (BZS)

<u>Assembler Notation</u>	<u>Opcode+M1</u>	<u>Format</u>
BZ D2(X2)	433	RX1,RX2
BZ A2(FX2,SX2)	433	RX3
BZR R2	033	RR
BZS A	223 (Backward)	SF
	233 (Forward)	

Operation:

If the G and L flags are both zero in the condition code, a branch is taken to the second operand location. If the G or L flag is set, the next sequential instruction is executed.

Condition Code:

Unchanged

Programming Notes:

The branch destination must be located on a halfword boundary.

In the RR format, the branch address is contained in the register specified by R2.

4.5.14 Branch on Not Zero (BNZ, BNZR, BNZS)

Branch on Not Zero (BNZ)

Branch on Not Zero Register (BNZR)

Branch on Not Zero Short (BNZS)

<u>Assembler Notation</u>	<u>Opcode+M1</u>	<u>Format</u>
BNZ D2(X2)	423	RX1,RX2
BNZ A2(FX2,SX2)	423	RX3
BNZR R2	023	RR
BNZS A	203 (Backward)	SF
	213 (Forward)	

Operation:

If the G or L flag in the condition code is set, a branch is taken to the second operand address. If the G and L flags are both zero, the next sequential instruction is executed.

Condition Code:

Unchanged

Programming Notes:

The branch destination must be located on a halfword boundary.

In the RR format, the branch address is contained in the register specified by R2.

4.5.15 Branch (Unconditional) (B, BR, BS)

Branch (Unconditional) (B)
Branch Register (Unconditional) (BR)
Branch Short (Unconditional) (BS)

<u>Assembler Notation</u>	<u>Opcode+M1</u>	<u>Format</u>
B D2(X2)	430	RX1,RX2
B A2(FX2,SX2)	430	RX3
BR R2	030	RR
BS A	220 (Backward) 230 (Forward)	SF

Operation:

A branch is unconditionally taken to the second operand address.

Condition Code:

Unchanged

Programming Notes:

The branch destination must be located on a halfword boundary.

In the RR format, the branch address is contained in the register specified by R2.

This instruction is assembled as a Branch on False Condition instruction with no condition specified (M1 = 0); therefore, the branch test is always false and the branch is always taken.

Example:

The following example illustrates the use of the B instruction.

<u>Assembler Notation</u>	<u>Machine Code</u>	<u>Comments</u>
B OPTIN	4300 0A00	An unconditional branch is taken to location X'0A00'.

4.5.16 No Operation (NOP, NOPR)

No Operation (NOP)
No Operation Register (NOPR)

<u>Assembler Notation</u>	<u>Opcode+M1</u>	<u>Format</u>
NOP D2(X2)	420	RX1,RX2
NOP A2(FX2,SX2)	420	RX3
NOPR R2	020	RR

Operation:

The next sequential instruction is executed.

Condition Code:

Unchanged

Programming Notes:

D2(X2) or A2(FX2,SX2) and R2 are ignored and usually equal zero.

This instruction is assembled as a Branch on True Condition instruction with no condition specified (M1 = 0); therefore, no branch is taken and the next instruction is fetched and executed.

Example:

The following example illustrates the use of the NOP and NOPR instructions.

<u>Assembler Notation</u>	<u>Machine Code</u>	<u>Comments</u>
NOP 0(0,0)	4200 4000 0000	No operation
NOP 0	4200 0000	No operation
NOPR	0200	No operation

CHAPTER 5 FIXED POINT ARITHMETIC

5.1 INTRODUCTION

Fixed point arithmetic instructions provide a complete set of operations for calculating addresses and indices, counting, and general-purpose fixed point arithmetic.

5.2 FIXED POINT DATA FORMATS

There are three formats for fixed point data: the halfword, the fullword and the doubleword. In each of these formats, the most significant bit (bit 0) is the sign bit. The remaining 15, 31 or 63 bits represent the magnitude. See Figure 5-1.

599




Figure 5-1 Fixed Point Data Formats

Positive values are represented in true binary form with a sign bit of zero. Negative values are represented in two's complement form with a sign bit of one. To change the sign of a number, the two's complement of the number can be produced by subtracting the number from zero. Other ways would be to:

- change all zeros to ones, and all ones to zeros, or
- add one.

5.3 FIXED POINT NUMBER RANGE

Fixed point numbers represent integers. Table 5-1 shows relations between different formats, along with decimal values.

TABLE 5-1 FIXED POINT FORMAT RELATIONS

DOUBLEWORD	FULLWORD	HALFWORD	DECIMAL
8000000000000000 (most negative)			-9 223 372 036 854 775 808
	80000000 (most negative)		-2 147 483 648
		8000 (most negative)	-32 768
FFFFFFFFFFFFFFFF	FFFFFFF	FFFF (least negative)	-1
0000000000000000	00000000	0000	0
0000000000000001	00000001	0001 (least negative)	1
		7FFF (most negative)	32 767
	7FFFFFFF (most positive)		2 147 483 647
7FFFFFFFFFFFFFFFF			9 223 372 036 854 775 807

5.4 OPERATIONS

Fixed point instructions include both fullword and halfword operations. Fullword operations take place between the contents of two general registers, between the contents of a general register and a fullword stored in memory, or between the contents of a general register and a fullword obtained from the instruction stream. Fullword multiply produces a doubleword result that is contained in two adjacent registers. Fullword divide operates on doubleword data contained in two adjacent registers.

Halfword operations take place between a fullword contained in one of the general registers and a halfword contained in memory. Before the operation is started, the halfword in memory is expanded to a fullword by propagating the most significant bit (sign bit) into the high-order bits of the fullword. The halfword multiply and divide instructions are exceptions to this rule.

5.5 CONDITION CODE

All fixed point arithmetic instructions, except multiply and divide, affect the condition code to indicate the outcome of the operation on the 32-bit result.

In fixed point add and subtract operations, the arguments are represented in two's complement form; therefore, all bits, including sign, participate in forming the result. Consequently, the occurrence of a carry or borrow has no real arithmetic significance.

For example, an add operation between a minus one (FFFF FFFF) and a plus two (0000 0002) produces the correct result of plus one (0000 0001) and a Carry (C). The condition code is set to 1010 (C=1 and G=1). C means that the complete result, which in this case would have been 1 0000 0001, would not fit in 32 bits.

An Overflow (V) occurs when the result does not fit in 31 bits. Note that bit zero must be reserved for the sign of the result. For example, adding one to the largest positive fixed point value produces a V:

```
  7FFF FFFF
+0000 0001
-----
  8000 0000
```

The resulting condition code is 0101 (V=1 and L=1).

The result, 8000 0000, is logically correct, but because the sign bit is negative when the result should be positive, the V condition exists.

The columns of the condition code table given for each instruction description show the state of the C, V, greater than (G) and less than (L) flags for the possible results.

An 'X' in a condition code column means that the particular flag is not defined and can be either 0 or 1. Hence, no inference should be drawn by testing that particular flag.

5.6 FIXED POINT INSTRUCTION FORMATS

The fixed point instructions use Register-to-Register (RR), Short Form (SF), Register and Indexed Storage (RX), and Register and Immediate (RI) instruction formats.

5.7 FIXED POINT INSTRUCTIONS

The fixed point instructions described in this section are:

A	Add
AR	Add Register
AI	Add Immediate
AIS	Add Immediate Short
AH	Add Halfword
AHI	Add Halfword Immediate
AM	Add to Memory
AHM	Add Halfword to Memory
S	Subtract
SR	Subtract Register
SI	Subtract Immediate
SIS	Subtract Immediate Short
SH	Subtract Halfword
SHI	Subtract Halfword Immediate
C	Compare
CR	Compare Register
CI	Compare Immediate
CH	Compare Halfword
CHI	Compare Halfword Immediate
M	Multiply
MR	Multiply Register
MH	Multiply Halfword
MHR	Multiply Halfword Register
D	Divide
DR	Divide Register
DH	Divide Halfword
DHR	Divide Halfword Register
SLA	Shift Left Arithmetic
SLHA	Shift Left Halfword Arithmetic

SRA	Shift Right Arithmetic
SRHA	Shift Right Halfword Arithmetic
CHVR	Convert to Halfword Value Register

5.7.1 Add (A, AR, AI, AIS)

Add (A)
Add Register (AR)
Add Immediate (AI)
Add Immediate Short (AIS)

Assembler Notation	Opcode	Format
A R1,D2(X2)	5A	RX1,RX2
A R1,A2(FX2,SX2)	5A	RX3
AR R1,R2	0A	RR
AI R1,I2(X2)	FA	RI2
AIS R1,N	26	SF

Operation:

The second operand is added algebraically to the contents of the register specified by R1. The result of this 32-bit addition replaces the contents of the register specified by R1.

Condition Code:

C	V	G	L	
X	0	0	0	Result is zero
X	0	0	1	Result is less than zero
X	0	1	0	Result is greater than zero
X	1	X	X	Arithmetic overflow
1	X	X	X	Carry

Programming Notes:

The second operand for the AIS instruction is obtained by expanding the 4-bit data field, N, to a 32-bit fullword by forcing the high-order bits to zero.

In the RI2 format, the contents of the index register specified by X2 are added to the 32-bit I2 field to form the fullword second operand.

In the RX formats the second operand must be located on a fullword boundary.

Example 1:

This example of the A instruction adds the contents of memory location LAB to the contents of register 4.

<u>Assembler Notation</u>	<u>Comments</u>
A REG4,LAB	ADD (LAB) TO (REG4)

Where:

REG4 contains X'7F341234'
Fullword in memory at LAB contains X'7124321'

Result of A Instruction:

(REG4) = X'FE465555'
(LAB) unchanged by this instruction
Condition code = 0101 (V=1, L=1)

Example 2:

This example of the A instruction adds the contents of memory location LAB to the contents of register 5.

<u>Assembler Notation</u>	<u>Comments</u>
A REG5,LAB	ADD (LAB) TO (REG5)

Where:

REG5 contains X'8000 0001'
Fullword in memory at LAB contains X'80000002'

Result of A Instruction:

(REG5) = X'00000003'
(LAB) unchanged by this instruction
Condition code = 1110 (C=1, V=1, G=1)

5.7.2 Add Halfword (AH, AHI)

Add Halfword (AH)

Add Halfword Immediate (AHI)

<u>Assembler Notation</u>	<u>Opcode</u>	<u>Format</u>
AH R1,D2(X2)	4A	RX1,RX2
AH R1,A2(FX2,SX2)	4A	RX3
AHI R1,I2(X2)	CA	RI1

Operation:

The 16-bit second operand is expanded to a 32-bit fullword by propagating the most significant bit through bits 0:15 of the fullword. The fullword operand is added to the fullword contents of the register specified by R1. The result replaces the contents of the register specified by R1.

Condition Code:

C	V	G	L	
=====				
X	0	0	0	Result is zero
X	0	0	1	Result is less than zero
X	0	1	0	Result is greater than zero
X	1	X	X	Arithmetic overflow
1	X	X	X	Carry

Programming Notes

In the RX formats, the second operand must be located on a halfword boundary.

In the RI1 format, the 16-bit I2 field is extended to a fullword by propagating the sign bit through bits 0:15. The contents of the index register specified by X2 are then added to form the fullword second operand.

Example 1:

This example of the AH instruction adds the halfword at memory location LAB to the contents of register 4.

<u>Assembler Notation</u>	<u>Comments</u>
AH REG4,LAB	ADD (LAB) TO (REG4)

Where:

REG4 contains X'00230002'
Halfword at memory location LAB contains X'FFFF'

Result of AH Instruction:

(REG4) = X'00230001'
(LAB) unchanged by this instruction
Condition code = 1010 (C=1, G=1)

Example 2:

This example of the AH instruction adds the halfword at memory location LAB to the contents of register 5.

<u>Assembler Notation</u>	<u>Comments</u>
AH REG5,LAB	ADD (LAB) TO (REG5)

Where:

REG5 contains X'FFFF FFF5'
LAB contains X'FFF2'

Result of AH Instruction:

(REG5) = 'FFFF FFE7'
(LAB) unchanged by this instruction
Condition code = 1001 (C=1, L=1)

5.7.3 Add to Memory (AM)

<u>Assembler Notation</u>	<u>Opcode</u>	<u>Format</u>
AM R1,D2(X2)	51	RX1,RX2
AM R1,A2(FX2,SX2)	51	RX3

Operation:

The first operand contained in the register specified by R1 is added algebraically to the fullword second operand. The result replaces the fullword second operand in memory. The contents of the register specified by R1 are not changed.

Condition Code:

C	V	G	L	
X	0	0	0	Result is zero
X	0	0	1	Result is less than zero
X	0	1	0	Result is greater than zero
X	1	X	X	Arithmetic overflow
1	X	X	X	Carry

Programming Note:

The second operand must be located on a fullword boundary.

Example 1:

This example of the AM instruction adds the contents of register 8 to memory location LOC.

<u>Assembler Notation</u>	<u>Comments</u>
AM REG8,LOC	ADD (REG8) TO (LOC)

Where:

REG8 contains X'00000008'
 Fullword in memory at LOC contains X'034289AB'

Result of AM Instruction:

(REG8) unchanged by this instruction
(LOC) = X'034289B3'
Condition code = 0010 (G=1)

Example 2:

This example of the AM instruction adds the contents of register 7 to memory location LOC.

<u>Assembler Notation</u>	<u>Comments</u>
AM REG7,LOC	ADD (REG7) TO (LOC)

Where:

REG7 contains X'7F341234'
Fullword in memory at LOC contains X'7F124321'

Result of AM Instruction:

(REG7) unchanged by this instruction
(LOC) = X'FE465555'
Condition code = 0101 (V=1, L=1)

5.7.4 Add Halfword to Memory (AHM)

<u>Assembler Notation</u>	<u>Opcode</u>	<u>Format</u>
AHM R1,D2(X2)	61	RX1,RX2
AHM R1,A2(FX2,SX2)	61	RX3

Operation:

The halfword second operand is added algebraically to the least significant 16 bits (bits 16:31) of the register specified by R1. The 16-bit result replaces the contents of the memory location specified by the effective address of the second operand. The contents of the register specified by R1 are not changed.

Condition Code:

C	V	G	L	
X	0	0	0	Result is zero
X	0	0	1	Result is less than zero
X	0	1	0	Result is greater than zero
X	1	X	X	Arithmetic overflow
1	X	X	X	Carry

Programming Notes:

The second operand must be located on a halfword boundary.

The condition code settings are based on the halfword result.

Example 1:

This example of the AHM instruction adds the contents of register 5 to the contents of memory location LAB.

<u>Assembler Notation</u>	<u>Comments</u>
AHM REG5,LAB	ADD (REG5) TO (LAB)

Where:

REG5 contains X'00230002'
 Halfword in memory at LAB contains X'FFFF'

Result of AHM Instruction:

(REG5) unchanged by this instruction
(LAB) = 0001
Condition code = 1010 (C=1, G=1)

Example 2:

This example of the AHM instruction adds the contents of register 6 to the contents of memory location LAB.

<u>Assembler Notation</u>	<u>Comments</u>
AHM REG6,LAB	ADD (REG6) TO (LAB)

Where:

REG6 contains X'FFFF FFF5'
LAB contains X'FFF2'

Result of AHM Instruction:

(REG6) unchanged by this instruction
(LAB) = FFE7
Condition code = 1001 (C=1, L=1)

5.7.5 Subtract (S, SR, SI, SIS)

Subtract (S)
 Subtract Register (SR)
 Subtract Immediate (SI)
 Subtract Immediate Short (SIS)

Assembler Notation	Opcode	Format
S R1,D2(X2)	5B	RX1,RX2
S R1,A2(FX2,SX2)	5B	RX3
SR R1,R2	0B	RR
SI R1,I2(X2)	FB	RI2
SIS R1,N	27	SF

Operation:

The fullword second operand is subtracted algebraically from the contents of the register specified by R1. The result replaces the contents of the register specified by R1.

Condition Code:

C	V	G	L	
X	0	0	0	Result is zero
X	0	0	1	Result is less than zero
X	0	1	0	Result is greater than zero
X	1	X	X	Arithmetic overflow
1	X	X	X	Borrow

Programming Notes:

The second operand for the SIS instruction is obtained by expanding the 4-bit data field, N, to a 32-bit fullword by forcing the high-order bits to zero.

In the RI2 format, the contents of the index register specified by X2 are added to the 32-bit I2 field to form the fullword second operand.

In the RX formats, the second operand must be located on a fullword boundary.

Example 1:

This example of the S instruction subtracts the fullword at memory location LOC from the contents of register 9.

<u>Assembler Notation</u>	<u>Comments</u>
S REG9,LOC	SUBTRACT (LOC) FROM (REG9)

Where:

REG9 contains X'44444444'
LOC contains X'44444444'

Result of S Instruction:

(REG9) = 0
(LOC) unchanged by this instruction
Condition code = 0000

Example 2:

This example of the S instruction subtracts the fullword at memory location LOC from the contents of register 9.

<u>Assembler Notation</u>	<u>Comments</u>
S REG9,LOC	SUBTRACT (LOC) FROM (REG9)

Where:

REG9 contains X'23456789'
LOC contains X'FFFF4321'

Result of S Instruction:

(REG9) = 23462368
(LOC) unchanged by this instruction
Condition code = 1010 (C=1, G=1)

5.7.6 Subtract Halfword (SH, SHI)

Subtract Halfword (SH)

Subtract Halfword Immediate (SHI)

<u>Assembler Notation</u>	<u>Opcode</u>	<u>Format</u>
SH R1,D2(X2)	4B	RX1,RX2
SH R1,A2(FX2,SX2)	4B	RX3
SHI R1,I2(X2)	CB	R11

Operation:

The 16-bit second operand is expanded to a 32-bit fullword by propagating the most significant bit through bits 0:15. This fullword is subtracted from the contents of the register specified by R1. The result replaces the contents of the register specified by R1.

Condition Code:

C	V	G	L	
X	0	0	0	Result is zero
X	0	0	1	Result is less than zero
X	0	1	0	Result is greater than zero
X	1	X	X	Arithmetic overflow
1	X	X	X	Borrow

Programming Notes:

In the RX formats, the second operand must be located on a halfword boundary.

In the R11 format, the 16-bit I2 field is extended to a fullword by propagating the sign bit through bits 0:15. The contents of the index register specified by X2 are then added to form the fullword second operand.

Example 1:

This example of the SH instruction subtracts the halfword at memory location LOC from the contents of register 9.

<u>Assembler Notation</u>	<u>Comments</u>
SH REG9,LOC	SUBTRACT (LOC) FROM (REG9)

Where:

REG9 contains X'00123456'
LOC contains X'FFF4'

Result of SH Instruction:

(REG9) = 00123462
(LOC) unchanged by this instruction
Condition code = 1010

Example 2:

This example of the SH instruction subtracts the halfword at memory location LOC from the contents of register 9.

<u>Assembler Notation</u>	<u>Comments</u>
SH REG9,LOC	SUBTRACT (LOC) FROM (REG9)

Where:

REG9 contains X'FFFF4567'
LOC contains X'2345'

Result of SH Instruction:

(REG9) = FFFF2222
(LOC) unchanged by this instruction
Condition code = 0001

5.7.7 Compare (C, CR, CI)

Compare (C)
Compare Register (CR)
Compare Immediate (CI)

<u>Assembler Notation</u>	<u>Opcode</u>	<u>Format</u>
C R1,D2(X2)	59	RX1,RX2
C R1,A2(FX2,SX2)	59	RX3
CR R1,R2	09	RR
CI R1,I2(X2)	F9	RI2

Operation:

The first operand contained in the register specified by R1 is compared algebraically to the 32-bit second operand. The result is indicated by the condition code setting. Neither operand is changed.

Condition Code:

C	V	G	L	
0	X	0	0	First operand is equal to second
1	X	0	1	First operand is less than second
0	x	1	0	First operand is greater than second

Programming Notes:

In the RX formats, the second operand must be located on a fullword boundary.

The state of the V flag is undefined.

Example:

This example of the C instruction compares the contents of register 3 to the contents of the fullword in memory location LAB.

Assembler Notation

Comments

C REG3,LAB

COMPARE (REG3) TO (LAB)

Where:

REG3 contains X'44567894'

Fullword at LAB contains X'04321243'

Result of C Instruction:

(REG3) unchanged by this instruction

(LAB) unchanged by this instruction

Condition code = 0010 (G=1)

5.7.8 Compare Halfword (CH, CHI)

Compare Halfword (CH)

Compare Halfword Immediate (CHI)

Assembler Notation	Opcode	Format
CH R1,D2(X2)	49	RX1,RX2
CH R1,A2(FX2,SX2)	49	RX3
CHI R1,I2(X2)	C9	R11

Operation:

The halfword second operand is expanded to a fullword by propagating the most significant bit through bits 0:15. The first operand, the contents of the register specified by R1, is compared algebraically to the effective second operand. The result is indicated by the condition code setting. Neither operand is changed.

Condition Code:

C	V	G	L	
0	X	0	0	First operand is equal to second
1	X	0	1	First operand is less than second
0	X	1	0	First operand is greater than second

Programming Notes:

In the RX formats, the second operand must be located on a halfword boundary.

In the R11 format, the 16-bit I2 field is extended to a fullword by propagating the sign bit through bits 0:15. The contents of the index register specified by X2 are then added to form the fullword second operand.

Condition code settings are based on the fullword comparison. The state of the V flag is undefined.

Example:

This example of the CH instruction compares the contents of register 8 to the halfword at memory location LAB.

<u>Assembler Notation</u>	<u>Comments</u>
CH REG8,LAB	COMPARE (REG8) TO (LAB)

Where:

REG8 contains X'F4567891'
Halfword at LAB contains X'3123'

Result of CH Instruction:

(REG8) unchanged by this instruction
(LAB) unchanged by this instruction
Condition code = 1001 (C=1, V=1)

5.7.9 Multiply (M, MR)

Multiply (M)

Multiply Register (MR)

<u>Assembler Notation</u>	<u>Opcode</u>	<u>Format</u>
M R1,D2(X2)	5C	RX1,RX2
M R1,A2(FX2,SX2)	5C	RX3
MR R1,R2	1C	RR

Operation:

The fullword first operand contained in the register specified by R1+1 is multiplied by the fullword second operand. The 64-bit result is stored in the registers specified by R1 and R1+1. The sign of the result is determined by the rules of algebra.

Condition Code:

Unchanged

Programming Notes:

The R1 field of these instructions must specify an even-numbered register. If the R1 field of these instructions is odd, the result is undefined.

In the RX formats, the second operand must be located on a fullword boundary.

The most significant bits of the result are placed in the register specified by R1; the least significant bits are placed in the register by R1+1.

Example 1:

This example of the M instruction multiplies the contents of register 9 by the contents of memory location LOC and places the result in registers 8 and 9 (64 bits).

<u>Assembler Notation</u>	<u>Comments</u>
M REG8,LOC	MULTIPLY (REG9) BY (LOC)

Where:

REG8 contains unknown data
REG9 contains X'00002431'
Fullword at location LOC contains X'43120000'

Result of M Instruction:

REG8 and REG9 together contain the result
(REG8, REG9) = 0000 097B, 5E72 0000
(IOC) unchanged by this instruction
Condition code unchanged by this instruction

Example 2:

This example of the MR instruction multiplies the contents of register 9 by the contents of register 8 and places the result in registers 8 and 9 (64 bits).

<u>Assembler Notation</u>	<u>Comments</u>
MR REG8,REG8	MULTIPLY (REG9) BY (REG8)

Where:

REG8 contains X'00010000'
REG9 contains X'12345678'

Result of MR Instruction:

REG8 and REG9 together contain the result
(REG8, REG9) = 0000 1234, 5678 0000
Condition code unchanged by this instruction

5.7.10 Multiply Halfword (MH, MHR)

Multiply Halfword (MH)
Multiply Halfword Register (MHR)

<u>Assembler Notation</u>	<u>Opcode</u>	<u>Format</u>
MH R1,D2(X2)	4C	RX1,RX2
MH R1,A2(FX2,SX2)	4C	RX3
MHR R1,R2	0C	RR

Operation:

The first operand, contained in bits 16:31 of the register specified by R1, is multiplied by the 16-bit second operand, taken from memory or from bits 16:31 of the register specified by R2. Both operands are 16-bit signed two's complement values. The 32-bit result replaces the contents of the register specified by R1. The sign of the result is determined by the rules of algebra.

Condition Code:

Unchanged

Programming Note:

In the RX formats, the second operand must be located on a halfword boundary.

Example 1:

This example of the MH instruction multiplies the halfword contents of register 8 by the halfword in memory location LAB.

<u>Assembler Notation</u>	<u>Comments</u>
MH REG8,LAB	MULTIPLY LEAST SIGNIFICANT HALF OF (REG8) BY (LAB)

Where:

REG8 contains X'ABCD 0045'
Halfword at memory location LAB contains X'8674'

Result of MH Instruction:

(REG8) = FFDF3D44
(LAB) unchanged by this instruction
Condition code unchanged by this instruction

Example 2:

This example of the MHR instruction multiplies the contents of register 11 by the halfword contents of register 4.

<u>Assembler Notation</u>	<u>Comments</u>
MHR REG11,REG4	MULTIPLY LS HALF OF (REG11) BY LS HALF OF (REG4)

Where:

REG11 contains X'37210004'
REG4 contains X'FFFF0307'

Result of MHR Instruction:

(REG11) = 00000C1C
(REG4) unchanged by this instruction
Condition code unchanged by this instruction

5.7.11 Divide (D, DR)

Divide (D)

Divide Register (DR)

<u>Assembler Notation</u>	<u>Opcode</u>	<u>Format</u>
D R1,D2(X2)	5D	RX1,RX2
D R1,A2(FX2,SX2)	5D	RX3
DR R1,R2	1D	RR

Operation:

The 64-bit signed dividend contained in the two registers specified by R1 and R1+1 is divided by the signed fullword second operand. The 32-bit signed remainder replaces the contents of the register specified by R1. The signed 32-bit quotient replaces the contents of the register specified by R1+1.

The sign of the quotient is determined by the rules of algebra. The sign of the remainder is the same as the sign of the dividend.

Condition Code:

Unchanged

Programming Notes:

The R1 field of these instructions must specify an even-numbered register. If the R1 field of these instructions is odd, the result is undefined.

The most significant bits of the dividend must be contained in the register specified by R1. The least significant bits of the dividend must be contained in the register specified by R1+1.

In the RX formats, the second operand must be located on a fullword boundary.

If the divisor is equal to zero, the instruction is not executed, the operand registers remain unchanged, and the arithmetic fault interrupt is taken.

If the value of the quotient is more positive than X'7FFFFFFF' or more negative than X'80000000', quotient overflow is said to occur. If quotient overflow occurs, the operand registers remain unchanged, and the arithmetic fault interrupt is taken.

Example 1:

This example of the D instruction divides the contents of registers 8 and 9 by the fullword contents of memory location LOC.

<u>Assembler Notation</u>	<u>Comments</u>
D REG8,LOC	DIVIDE (REG8,9) BY (LOC)

Where:

REG8 contains X'12345678' = Most significant half of dividend
REG9 contains X'98765432' = Least significant half of dividend
LOC contains X'34343434' = Divisor

Result of D Instruction:

(REG8) = 1E1E1E1E = Remainder
(REG9) = 59455459 = Quotient
(LOC) unchanged by this instruction
Condition code unchanged by this instruction

Example 2:

This example of a D instruction divides the contents of registers 8 and 9 by the fullword contents of memory location LOC.

<u>Assembler Notation</u>	<u>Comments</u>
D REG8,LOC	DIVIDE (REG8,9) BY (LOC)

Where:

REG8 contains X'FFFF1234' = Most significant half of dividend
REG9 contains X'00000000' = Least significant half of dividend
LOC contains X'12345678' = Divisor

Result of D Instruction:

(REG8) = F250D9E0 = Remainder
(REG9) = FFF2EFFC = Quotient
LOC unchanged by this instruction
Condition code unchanged by this instruction

Example 3:

This example of a D instruction divides the contents of registers 8 and 9 by the fullword contents of memory location LOC.

<u>Assembler Notation</u>	<u>Comments</u>
D REG8,LOC	DIVIDE (REG8,9) BY (LOC)

Where:

REG8 contains X'43657898' = Most significant half of dividend
REG9 contains X'12123456' = Least significant half of dividend
LOC contains X'00000000' = Divisor

Result of D Instruction:

Division by zero causes arithmetic fault to be taken. Operands and condition code remain unchanged by this instruction.

Example 4:

This example of a D instruction divides the contents of registers 8 and 9 by the fullword contents of memory location LOC.

<u>Assembler Notation</u>	<u>Comments</u>
D REG8,LOC	DIVIDE (REG8,9) BY (LOC)

Where:

REG8 contains X'80000000' = Most significant half of dividend
REG9 contains X'00000001' = Least significant half of dividend
LOC contains X'00000001' = Divisor

Result of D Instruction:

Quotient overflow causes arithmetic fault to be taken. Operands and condition code remain unchanged by this instruction.

Example 5:

This example of the DR instruction divides the contents of register 8 and 9 by the contents of register 2.

Assembler Notation

Comments

DR REG8,REG2 DIVIDE (REG8,9) BY (REG2)

Where:

REG8 contains X'FFFFFFFF' = Most significant half of dividend
REG9 contains X'FFFFFFFFD' = Least significant half of dividend
REG2 contains X'FFFFFFFE' = Divisor

Result of DR Instruction:

(REG8) = FFFFFFFF = Remainder
(REG9) = 00000001 = Quotient
(REG2) unchanged by this instruction
Condition code unchanged by this instruction

5.7.12 Divide Halfword (DH, DHR)

Divide Halfword (DH)

Divide Halfword Register (DHR)

<u>Assembler Notation</u>	<u>Opcode</u>	<u>Format</u>
DH R1,D2(X2)	4D	RX1,RX2
DH R1,A2(FX2,SX2)	4D	RX3
DHR R1,R2	0D	RR

Operation:

The 32-bit signed dividend contained in the register specified by R1 is divided by the 16-bit signed second operand. The 16-bit signed remainder is copied to R1 (bits 16:31) and the halfword value is converted to a fullword value. The 16-bit signed quotient is copied to the register specified by R1 + 1 after conversion to a fullword value.

The sign of the quotient is determined by the rules of algebra. The sign of the remainder is the same as the sign of the dividend.

Condition Code:

Unchanged

Programming Notes:

In the RX formats, the second operand must be located on a halfword boundary. In the RR format, the second operand is taken from bits 16:31 of the register specified by R2.

If the divisor is equal to zero, the instruction is not executed, the operand registers remain unchanged, and the arithmetic fault interrupt is taken.

If the value of the quotient is more positive than X'7FFF' or more negative than X'8000', quotient overflow is said to occur. If quotient overflow occurs, the operand registers remain unchanged, and the arithmetic fault interrupt is taken.

Example 1:

This example of the DH instruction divides the contents of register 7 by the halfword contents of memory location LOC.

<u>Assembler Notation</u>	<u>Comments</u>
DH REG7,LOC	DIVIDE (REG7) BY (LOC)

Where:

REG7 contains X'0000 0054' = Dividend
LOC contains X'0008' = Divisor

Result of DH Instruction:

(REG7) = 0000 0004 = Remainder
(REG8) = 0000 000A = Quotient
(LOC) unchanged by this instruction
Condition code unchanged by this instruction

Example 2:

This example of the DH instruction divides the contents of register 7 by the halfword contents of memory location LOC.

<u>Assembler Notation</u>	<u>Comments</u>
DH REG7,LOC	DIVIDE (REG7) BY (LOC)

Where:

REG7 contains X'1234 5678' = Dividend
LOC contains X'0000' = Divisor

Result of DH Instruction:

Division by zero causes arithmetic fault to be taken. Operands and condition code remain unchanged by this instruction.

Example 3:

This example of the DH instruction divides the contents of register 7 by the halfword contents of memory location LOC.

<u>Assembler Notation</u>	<u>Comments</u>
DH REG7,LOC	DIVIDE (REG7) BY (LOC)

Where:

REG7 contains X'8000 0002' = Dividend
LOC contains X'0001'

Result of DH Instruction:

Quotient overflow causes arithmetic fault to be taken. Operands and condition code remain unchanged by this instruction.

5.7.13 Shift Left Arithmetic (SLA)

Assembler Notation	Opcode	Format
SLA R1, I2(X2)	EE	RI1

Operation:

Bits 1:31 of the first operand, contained in the register specified by R1, are shifted left the number of places specified by the second operand. The sign bit (bit 0) remains unchanged. Bits shifted out of position 1 are shifted through the carry flag and then lost. The last bit shifted remains in the carry flag. Zeros are shifted into position 31.

Condition Code:

C	V	G	L	
X	0	0	0	Result is zero
X	0	0	1	Result is less than zero
X	0	1	0	Result is greater than zero
1	0	X	X	Carry

Programming Notes:

The state of the C flag indicates the state of the last bit shifted.

The shift count is specified by the least significant five bits of the second operand. The maximum shift count is 31.

A shift of zero places causes the condition code to be set in accordance with the value contained in the register specified by R1. The C flag is zero in this case.

Example:

This example of the SLA instruction shifts the bits in register 5 left by the number specified by the second operand.

Assembler Notation

Comments

SLA REG5,4

SHIFT (REG5) LEFT 4 PLACES

Where:

REG5 contains X'80005647'

Result of SLA Instruction:

(REG5) = 80056470

Condition code = 0001 (L=1)

5.7.14 Shift Left Halfword Arithmetic (SLHA)

Assembler Notation	Opcode	Format
SLHA R1,I2(X2)	CF	R11

Operation:

Bits 17:31 of the register specified by R1 are shifted left the number of places specified by the second operand. Bit 16 of the register, the halfword sign bit, remains unchanged. Bits shifted out of position 17 are shifted through the C flag and then lost. The last bit shifted remains in the C flag. Zeros are shifted into position 31. Bits 0:15 of the first operand register remain unchanged.

Condition Code:

C	V	G	L	
X	0	0	0	Result is zero
X	0	0	1	Result is less than zero
X	0	1	0	Result is greater than zero
1	0	X	X	Carry

Programming Notes:

The condition code settings are based on the halfword (bits 16:31) result.

The state of the C flag indicates the state of the last bit shifted.

The shift count is specified by the least significant four bits of the second operand. The maximum shift count is 15.

A shift of zero places causes the condition code to be set in accordance with the halfword value contained in bits 16:31 of the register specified by R1. The C flag is zero in this case.

5.7.15 Shift Right Arithmetic (SRA)

<u>Assembler Notation</u>	<u>Opcode</u>	<u>Format</u>
SRA R1,I2(X2)	EE	R11

Operation:

Bits 1:31 of the first operand, contained in the register specified by R1, are shifted right the number of places specified by the second operand. The sign bit (bit 0) remains unchanged and is propagated right as many positions as specified by the second operand. Bits shifted out of position 31 are shifted through the C flag and lost. The last bit shifted remains in the C flag.

Condition Code:

C	V	G	L	
X	0	0	0	Result is zero
X	0	0	1	Result is less than zero
X	0	1	0	Result is greater than zero
1	0	X	X	Carry

Programming Notes:

The state of the C flag indicates the state of the last bit shifted.

The shift count is specified by the least significant five bits of the second operand. The maximum shift count is 31.

A shift of zero places causes the condition code to be set in accordance with the value contained in the register specified by R1. The C flag is zero in this case.

Example:

This example of the SRA instruction shifts the contents of register 9 right the number of places specified by the second operand.

Assembler Notation

Comments

SRA REG9,8

SHIFT (REG9) RIGHT 8 PLACES

Where:

REG9 contains X'800004256'

Result of SRA Instruction:

(REG9) = X'FF800042'

Condition code = 0001 (L=1)

5.7.16 Shift Right Halfword Arithmetic (SRHA)

<u>Assembler Notation</u>	<u>Opcode</u>	<u>Format</u>
SRHA R1,I2 (X2)	CE	R11

Operation:

Bits 17:31 of the register specified by R1 are shifted right the number of places specified by the second operand. Bit 16 of the register, the halfword sign bit, remains unchanged and is propagated right the number of positions specified by the second operand. Bits shifted out of position 31 are shifted through the C flag and lost. The last bit shifted remains in the C flag. Bits 0:15 of the first operand register remain unchanged.

Condition Code:

C	V	G	L	
X	0	0	0	Result is zero
X	0	0	1	Result is less than zero
X	0	1	0	Result is greater than zero
1	0	X	X	Carry

Programming Notes:

The condition code settings are based on the halfword (bits 16:31) result.

The state of the C flag indicates the state of the last bit shifted.

The shift count is specified by the least significant four bits of the second operand. The maximum shift count is 15.

A shift of zero places causes the condition code to be set in accordance with the halfword value contained in bits 16:31 of the register specified by R1. The C flag is zero in this case.

5.7.17 Convert to Halfword Value Register (CHVR)

Assembler Notation	Opcode	Format
CHVR R1,R2	12	RR

Operation:

The halfword second operand, bits 16:31 of the register specified by R2, is expanded to a fullword by propagating the most significant bit (bit 16) through bits 0:15. This fullword replaces the contents of the register specified by R1.

Condition Code:

C	V	G	L	
X	X	0	0	Result is zero
X	X	0	1	Result is less than zero
X	X	1	0	Result is greater than zero
X	1	X	X	Source operand cannot be represented by a 16-bit signed number
1	X	X	X	Carry flag was set in previous condition code
0	X	X	X	Carry flag was zero in previous condition code

Programming Notes:

The V flag is set when bit 15 of the second operand is not the same as bit 16 of the second operand. The G and L flags reflect the algebraic value of bits 16:31 of the second operand.

Execution of this instruction following halfword operations guarantees the same results as those obtained if the program were run on a 16-bit machine. For example, if location A in memory contains the halfword value of X'7FFF' (decimal 32767) then:

LH	R1,A	R1 contains X'00007FFF'
AIS	R1,1	R1 contains X'00008000'

Following the add operation, the condition code is:

```
-----  
| C | V | G | L |  
|-----|  
| 0 | 0 | 1 | 0 |  
|-----|
```

indicating a result greater than zero, which is correct for fullword operations. If the same sequence were executed on a 16-bit processor:

```
LH      R1,A      R1 contains X'7FFF'  
AIS     R1,1      R1 contains X'8000'
```

Following this, the condition code in the halfword processor is:

```
-----  
| C | V | G | L |  
|-----|  
| 0 | 1 | 0 | 1 |  
|-----|
```

indicating overflow and a negative result. Going back to the original sequence and adding the convert to halfword value register instruction produces the following:

```
LH      R1,A      R1 contains X'00007FFF'  
AIS     R1,1      R1 contains X'00008000'  
CHVR    R1,R1     R1 contains X'FFFF8000'
```

Following this sequence, the condition code is:

```
-----  
| C | V | G | L |  
|-----|  
| 0 | 1 | 0 | 1 |  
|-----|
```

which is identical to that of the 16-bit processor and can be tested in the same manner.

CHAPTER 6 FLOATING POINT ARITHMETIC

6.1 INTRODUCTION

Floating point arithmetic instructions provide a means for rapid handling of scientific data expressed as floating point numbers. Single and double precision floating point instructions, as well as mixed mode floating point instructions, are described in this chapter. The comprehensive set of instructions includes load and store floating point numbers; add, subtract, multiply, divide and compare two floating point numbers; convert fixed point to floating point and vice versa; and mixed mode operations that translate single precision to double precision and vice versa.

Floating point is a means of representing a quantity in any numbering system. For example, the decimal number 123 (base 10), can be represented in the following forms:

$$\begin{array}{ll} 123.0 & \times 10^0 \\ 1.23 & \times 10^2 \\ 0.123 & \times 10^3 \\ 0.0123 & \times 10^4 \end{array}$$

In this example, the decimal point moved; this is called a floating point. In actual floating point representation, the significant digits are always fractional and are collectively referred to as fractions. The power to which the base number is raised is called the exponent. For example, in the number $.45678 \times 10^2$, 45678 is the fraction and 2 is the exponent. Both the fraction and the exponent can be signed. If there is a floating point representation such as:

$$(\text{sign of fraction}) \times (\text{exponent}) \times (\text{fraction})$$

the following representation applies.

NUMBER	FLOATING POINT	
+ 32.94	= +.3294 x 10 ²	+ +2 3294
-23760000.0	= -.2376 x 10 ⁸	- +8 2376
+0.000059	= +.59 x 10 ⁻⁴	+ -4 59
-0.0000000092073	= -.92073 x 10 ⁻⁸	- -8 92073

Large or small numbers can be easily expressed in floating point, making it ideally suitable for scientific computation. Note the compactness of floating point notation in the above examples.


Floating point representation in the processor is similar to the above representation. The differences are:

- Hexadecimal, instead of decimal, numbering system is used.
- Physical size of the number is limited; therefore, the magnitude and precision are limited.


6.2 FLOATING POINT DATA FORMATS

Floating point numbers occur in one of two formats: single and double precision. The single precision format requires a fullword (32 bits). When such a value is contained in memory, it must exist on a fullword address boundary. The sign (S), exponent (X) and fraction (consisting of the digits F1, F2, F3, F4, F5 and F6) fields are designated as follows:

602



The double precision format requires a doubleword (64 bits). When two general registers hold a double precision value, an even/odd pair of general registers must be used. The even-numbered register contains the most significant 32 bits, and the next sequential odd register contains the least significant 32 bits. The sign (S), exponent (X) and fraction (consisting of digits F1 through F14) fields are designated as follows.



NOTE

Floating point uses sign/magnitude notation rather than the two's complement notation used for integers.

6.3 FLOATING POINT NUMBER

In the processor, a floating point number is represented in the following form:

```
-----
| Sign | Exponent | Fraction |
-----
```

Sign is the most significant bit of a floating point number. The sign bit is zero for positive numbers and one for negative numbers. The floating point value of zero always has a positive sign.

Exponent is the 7-bit field, bits 1:7, that is designated as the exponent field. The exponent is expressed in excess-64 notation. The number in this field contains the true value of the exponent plus X'40' (decimal 64). This helps to represent very small magnitudes between 0 and 1. Examples of the exponent values follow.

EXPONENT IN EXCESS-64 NOTATION	TRUE EXPONENT IN HEXADECIMAL	TRUE EXPONENT IN DECIMAL	MULTIPLY FRACTION BY
00	-40	-64	16^{-64}
3F	-1	-1	16^{-1}
40	0	0	16^0
41	1	1	16^1
7F	3F	63	16^{63}

The exponent field for true zero is always 00.

Fraction is the fraction field that contains six hexadecimal digits for single precision floating point numbers and 14 hexadecimal digits for double precision floating point numbers. As in any other fraction, the floating point fraction is expressed with greatest precision when the most significant hexadecimal digit (not necessarily the most significant bit) is nonzero. The floating point number with such a fraction is called a normalized floating point number. In the Perkin-Elmer Series 3200 Processors, normalized numbers are always used to obtain the maximum possible precision. See Appendix D for hexadecimal fraction conversion.

The following examples illustrate the sign, exponent and fraction concept of a floating point number.

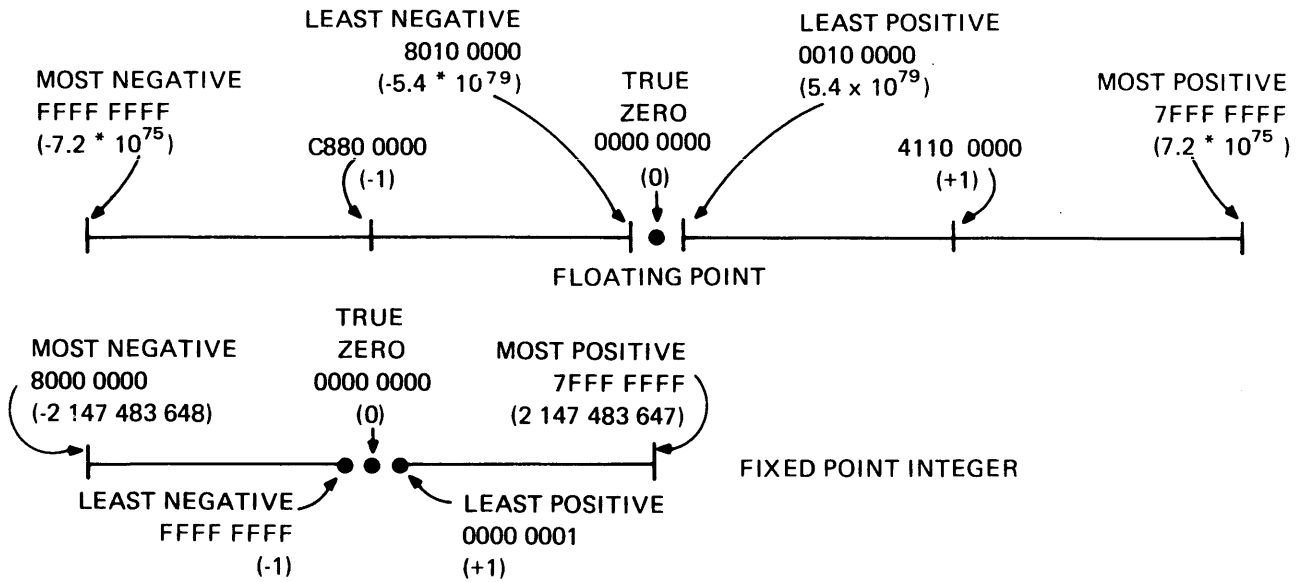
NUMBER IN HEXADECIMAL INTEGER-FRACTION NOTATION	SIGN EXPONENT/ FRACTION SHOWN FOR CLARITY	SINGLE PRECISION FLOATING POINT NUMBER
+1.3A25678	0 41 13A25678	4113A256
-6.89F2C	1 41 689F2C	C1689F2C
+1A.C39D21	0 42 1AC39021	421AC39D
-3C1DF.82A3	1 45 3C1DF82A3	C53C1DF8
+ABCDEF12.9AC	0 48 ABCDEF129A	48ABCDEF
+0.0032A9CF2	0 3E 32A9CF2	3E32A9CF
-0.000002C7B5	1 3B 2C7B5	BB2C7B50

6.3.1 Floating Point Number Range

The range of magnitude (M) of a normalized floating point number is as follows:


$$\begin{array}{l}
 \text{Single precision:} \quad 16^{-65} \leq M \leq (1 - 16^{-6}) * 16^{63} \\
 \text{Double precision:} \quad 16^{-65} \leq M \leq (1 - 16^{-14}) * 16^{63} \\
 \text{Approximately for both:} \quad 5.4 * 10^{-79} \leq M \leq 7.2 * 10^{79}
 \end{array}$$

The following diagram shows the floating point range in relation to the fixed point range along with the decimal values.



6.3.2 Normalization

Normalization is a process of making the most significant digit (F1) of the fraction of a floating point number nonzero. In the normalization process, the floating point fraction is shifted left hexadecimally (i.e., four bits at a time), and its exponent is decremented by one for each hexadecimal shift until the most significant digit (not necessarily the most significant bit) of the fraction is nonzero.



Except for the load instructions, all floating point operations assume and require normalized operands for consistent results. The load instructions normalize an unnormalized operand.

Examples:

	OPERAND	AFTER NORMALIZATION
1.	42012345	41123450
2.	21000ABC	1EABC000
3.	C900FE12	C7FE1200
4.	6C000000	00000000 (true zero)
5.	82000A67	00000000 (exponent underflow)


In Example 4, the fraction of the operand is zero. During the normalization process, such a fraction is detected, and the floating point number is set to true zero.

In Example 5, the exponent of the operand is very small. During the normalization process, the exponent is decremented from 00 to 7F. Such a transition results in exponent underflow, and the floating point number is set to true zero.

Normalized results are always produced in floating point operations, provided that the operands are normalized. Results of operations between unnormalized numbers are undefined.

6.3.3 Equalization

Equalization is a process of equalizing exponents of two floating point numbers. The fraction of the floating point number with the smaller exponent is shifted right hexadecimally (i.e., four bits at a time), and its exponent is incremented by one for each hexadecimal shift until the two exponents are equal.



During floating point addition and subtraction, the two floating point operands are equalized.

Examples:

	OPERAND	AFTER EQUALIZATION
1.	43123456 3F789ABC	43123456 43000078
2.	C7FE1234 4956789A	C900FE12 4956789A

In this example, normalized floating point numbers are shown because addition and subtraction require normalization. If the exponents differ by more than six for single precision or more than 14 for double precision, the representable significance of the lower exponent floating point number is lost in the process of equalization. Digits shifted out are shifted through the guard digits and can still have an effect on the result, sum or difference.

6.3.4 True Zero

A floating point number is true zero when the exponent and the fraction fields are all zeros; therefore, all data bits must be zero. A zero value always has a positive sign. In general, zero values participate as normal operands in all floating point operations.

A true zero can be used as an operand. It can also result from an arithmetic operation that caused an exponent underflow, in which case the entire number may be forced to true zero. If an arithmetic operation produces a result in which the fraction digits are all zeros (sometimes referred to as loss of significance), the entire number is forced to true zero.

Examples:

	NUMBER	OPERATION	RESULT	REASON
1.	030000AB	Normalization	0000 0000	Exponent underflow
2.	41ABCDEF	Subtraction	0000 0000	Loss of significance

6.3.5 Exponent Overflow

In floating point operations, exponent overflow occurs when a resulting exponent is greater than +63. If overflow occurs, the result register is unchanged. The condition code is set to reflect the overflow situation and the resulting sign. An arithmetic fault interrupt is also taken. Exponent overflow interrupts cannot be disabled. Figure 6-1 illustrates exponent overflow using a line representation of numbers.

609




Figure 6-1 Exponent Overflow

6.3.6 Exponent Underflow

The normalization process, during a floating point operation, may produce an exponent underflow. This underflow occurs when a result exponent is less than -64. Figure 6-2 illustrates exponent underflow using a line representation of numbers.

610




Figure 6-2 Exponent Underflow

If underflow occurs, an arithmetic fault interrupt is taken if enabled by the current program status word (PSW). Both operands remain unchanged. If underflow is disabled by the current PSW, the result is forced to zero (the closest possible answer), the V flag in the condition code is set, and the next sequential instruction is executed.

6.3.7 Guard Digits and R* Rounding

When an intermediate floating point result has been formed, it consists of a sign, an exponent and a fraction field. The fraction field is extended by a number of guard digits containing the least significant fraction digits of the intermediate result. Before the result is copied to a destination, it is rounded to compensate for the loss of the guard digits in the final result.

Quotients are simple-rounded rather than R* rounded. R* (or nonbiased) differs from simple-rounded only when the truncated fraction is precisely one-half. Nonbiased rounding is statistically important for the accuracy of additions and subtractions: it is not important for division since the truncated fraction is hardly ever exactly one-half.

The following are rules for the R* Rounding scheme:

- If the most significant guard digit is hexadecimal 7 or less, no rounding is performed (see Example 1).
- If the most significant guard digit is hexadecimal 8, and all other guard digits are 0, the least significant bit of the final result is forced to 1 (see Example 2).
- If the most significant guard digit is hexadecimal 8, and another guard digit is nonzero, or if the most significant guard digit is hexadecimal 9 or greater, 1 is added to the fraction field of the final result (see Example 3). If this addition produces a carry out of the fraction field (i.e., fraction field was all ones), the result exponent is incremented by 1, the most significant fraction digit (F1) is set to hexadecimal 1, and all other fraction digits are set to 0 (see Example 4). Note that exponent overflow could occur as the result of rounding.

Examples:

	INTERMEDIATE RESULT		FINAL SINGLE PRECISION
	DATA	GUARD DIGITS	RESULT
1.	42ABCD12	32680000	42ABCD12
2.	C1183756	80000000	C1183757
3.	3E265739	80100000	3E26573A
4.	41FFFFFF	F0000000	42100000

6.3.8 Conversion from Decimal

To convert a decimal number into the excess-64 notation used internally by the processor, the following steps must be taken.

1. Separate the decimal integer from the decimal fraction.

$$182.375_{10} = (182 + .375)_{10}$$

2. Convert each part to hexadecimal by referring to the integer conversion table and the fraction conversion table in Appendix D.

$$182_{10} = B6_{16} \quad .375_{10} = .6_{16}$$

3. Combine the hexadecimal integer and fraction.

$$B6.6_{16} = (B6.6 \times 16^0)_{16}$$

4. Shift the radix point.

$$(B6.6 \times 16^0)_{16} = (.B66 \times 16^2)_{16}$$

5. Add 64 (X'40') to the exponent.

$$40_{16} + 2_{16} = 42_{16}$$

6. Convert the exponent field and fractions to binary allowing 1 bit for the sign, 7 bits for exponent field, and 24 or 56 bits for the fraction.

42B66 = 0100 0010 1011 0110 0110 0000 0000 0000

6.4 CONDITION CODE

Most floating point operations affect the condition code. For each instruction description, the possible condition code settings are shown.

6.5 FLOATING POINT INSTRUCTIONS

All floating point instructions are illegal when PSW bit 13 (FLM) is set. Floating point instructions cannot be executed when the processor is in the floating point masked (FLM) mode.

Floating point instructions use the Register-to-Register (RR) and the Register and Indexed Storage (RX) instruction formats. In all of the RR formats, except for the fix and float instructions, the R1 and R2 fields each specify one of the floating point registers. There are eight single precision floating point registers and eight double precision floating point registers numbered 0, 2, 4, 6, 8, 10, 12 and 14. Floating point instructions must specify even-numbered floating point registers, or the results of the instructions are undefined. Except for the FXR, FXDR, LGER and LGDR instructions, the R1 field always specifies a floating point register.

Floating point arithmetic operations, excluding loads and stores, require normalized operands to ensure correct results. If the operands are not normalized, the results of these operations are undefined. Floating point results are normalized. The floating point load instructions normalize the floating point data presented as the second operand.

The single precision floating point instructions described in this section are:

LU	Load Unnormalized Floating Point
LUR	Load Unnormalized Floating Point Register
LE	Load Floating Point
LER	Load Floating Point Register
LEGR	Load Floating Point from General Register
LPER	Load Positive Floating Point Register
LCER	Load Complement Floating Point Register
LME	Load Floating Point Multiple

LGER	Load General Register from Floating Point Register
STE	Store Floating Point
STME	Store Floating Point Multiple
AE	Add Floating Point
AER	Add Floating Point Register
SE	Subtract Floating Point
SER	Subtract Floating Point Register
CE	Compare Floating Point
CER	Compare Floating Point Register
ME	Multiply Floating Point
MER	Multiply Floating Point Register
DE	Divide Floating Point
DER	Divide Floating Point Register
FXR	Fix Register
FLR	Float Register

The double precision floating point instructions described in this section are:

LW	Load Unnormalized Double Precision
LWR	Load Unnormalized Double Precision Register
LD	Load Double Precision Floating Point
LDR	Load Register Double Floating Point
LDGR	Load Double Precision Floating Point from General Register
LPDR	Load Positive Register Double Precision Floating Point
LCDR	Load Complement Register Double Point Multiple
LMD	Load Double Precision Floating Point Multiple
LGDR	Load General Register from Double Precision Floating Point Register
STD	Store Double Precision Floating Point
STMD	Store Multiple Double Precision Floating Point
AD	Add Double Precision Floating Point
ADR	Add Register Double Precision Floating Point

SD	Subtract Double Precision Floating Point
SDR	Subtract Register Double Precision Floating Point
CD	Compare Double Precision Floating Point
CDR	Compare Register Double Precision Floating Point
MD	Multiply Double Precision Floating Point
MDR	Multiply Register Double Precision Floating Point
DD	Divide Double Floating Point
DDR	Divide Register Double Precision Floating Point
FXDR	Fix Register Double Precision Floating Point
FLDR	Float Register Double Precision Floating Point

The mixed mode floating point instructions described in this section are:

LED	Load Single Precision Floating Point from Double Precision Point
LEDR	Load Register Double Precision Floating Point from Single Precision Floating Point
LDE	Load Double Precision Floating Point from Single Precision Floating Point
LDER	Load Register Single Precision Floating Point from Double Precision Floating Point
STDE	Store Double Precision Floating Point in Single Precision Floating Point

6.5.1 Load Unnormalized Floating Point (LU, LUR)

Load Unnormalized Floating Point (LU)

Load Unnormalized Floating Point Register (LUR)

<u>Assembler Notation</u>	<u>Opcode</u>	<u>Format</u>
LU R1,D2(X2)	4E	RX1,RX2
LU R1,A2(FX2,SX2)	4E	RX3
LUR R1,R2	1E	RR

Operation:

The fullword second operand is placed in the single precision floating point register specified by R1. No normalization is performed.

Condition Code:

C	V	G	L	

0	0	0	0	Result is zero
0	0	0	1	Result is less than zero
0	0	1	0	Result is greater than zero

Programming Notes:

In the RX formats, the second operand must be located on a fullword boundary. This instruction is intended for data manipulation only. Floating point operations using data in a register loaded in this manner may not produce predictable results.

6.5.2 Load Floating Point (LE, LER, LEGR)

Load Floating Point (LE)

Load Floating Point Register (LER)

Load Floating Point from General Register (LEGR)

Assembler Notation	Opcode	Format
LE R1,D2(X2)	68	RX1,RX2
LE R1,A2(FX2,SX2)	68	RX3
LER R1,R2	28	RR
LEGR R1,R2	A5	RR

Operation:

The floating point second operand is normalized, if necessary, and placed in the single precision floating point register specified by R1.

Condition Code:

C	V	G	L	
0	0	0	0	Floating point result is zero
0	0	0	1	Floating point result is less than zero
0	0	1	0	Floating point result is greater than zero
0	1	0	0	Exponent underflow

Programming Notes:

If the argument fraction is zero, the entire result is forced to zero, X'0000 0000'.

Normalization can produce exponent underflow. If PSW bit 19 is set, an arithmetic fault interrupt is taken, and the register specified by R1 is unchanged. If an exponent underflow occurs, and bit 19 of the current PSW is zero, no arithmetic fault occurs. Zeros replace the contents of the register specified by R1.

In the RX formats, the second operand must be located on a fullword boundary.

Example:

This example of the LE instruction normalizes data taken from the fullword at memory location LOC and places it in floating point register 8.

Assembler Notation

Comments

LE REG8,LOC

LOAD FROM LOC AND NORMALIZE

Where:

Floating point REG8 contains unknown data.
LOC contains X'4200 1000'

Result of LE Instruction:

(REG8) = X'4010 0000'
(LOC) unchanged by this instruction
Condition code = 0010

6.5.3 Load Positive Floating Point Register (LPER)

Assembler Notation	Opcode	Format
LPER R1,R2	13	RR

Operation:

The floating point second operand data from the single precision floating point register specified by R2 is forced positive normalized, if necessary, and placed in the single precision floating point register specified by R1.

Condition Code:

C	V	G	L	
0	0	0	0	Floating point result is zero
0	0	1	0	Floating point result is greater than zero
0	1	0	0	Exponent underflow

Programming Notes:

If the argument fraction is zero, the entire result is forced to zero, X'0000 0000'.

Normalization can produce exponent underflow. If PSW bit 19 is set, an arithmetic fault interrupt is taken, and the register specified by R1 is unchanged. If an exponent underflow occurs, and bit 19 of the current PSW is zero, no arithmetic fault occurs. Zeros replace the contents of the register specified by R1.

Example:

Assembler Notation	Comments
LPER REG6,REG8	LOAD REG6 WITH POSITIVE OF (REG8)

Where:

Floating point REG6 contains unknown data
Floating point REG8 contains X'C11921FB'

Result of LPER Instruction:

(REG6) = X'411921FB'
(REG8) unchanged by this instruction
Condition code = 0010

6.5.4 Load Complement Floating Point Register (LCER)

<u>Assembler Notation</u>	<u>Opcode</u>	<u>Format</u>
LCER R1,R2	17	RR

Operation:

The sign of the floating point second operand data from the single precision floating point register specified by R2 is complemented. The resulting floating point number is normalized, if necessary, and placed in the single precision floating point register specified by R1.

Condition Code:

C	V	G	L	
0	0	0	0	Floating point result is zero
0	0	0	1	Floating point result is less than zero
0	0	1	0	Floating point result is greater than zero
0	1	0	0	Exponent underflow

Programming Notes:

If the argument fraction is zero, the entire result is forced to zero, X'0000 0000'.

Normalization can produce exponent underflow. If PSW bit 19 is set, an arithmetic fault interrupt is taken, and the register specified by R1 is unchanged. If an exponent underflow occurs, and bit 19 of the current PSW is zero, no arithmetic fault occurs. Zeros replace the contents of the register specified by R1.

6.5.5 Load Multiple Floating Point (LME)

<u>Assembler Notation</u>	<u>Opcode</u>	<u>Format</u>
LME R1,D2(X2)	72	RX2,RX2
LME R1,A2(FX2,SX2)	72	RX3

Operation:

Successive single precision floating point registers, starting with the register specified by R1, are loaded from successive fullword memory locations starting with the address of the second operand. The process stops when floating point register 14 has been loaded.

Condition Code:

Unchanged

Programming Notes:

Values loaded into the floating point registers are assumed to be normalized, and no test or adjustment is performed.

The second operand must be located on a fullword boundary. Loading a register with a "dirty zero" using this instruction will result in a load of true zero.

6.5.6 Load General Register from Floating Point Register (LGER)

Assembler Notation	Opcode	Format
LGER R1,R2	15	RR

Operation:

The floating point second operand, contained in the single precision floating point register specified by R2, is placed in the general register specified by R1. The second operand is unchanged.

Condition Code:

C	V	G	L	
0	0	0	0	Result is zero
0	0	0	1	Result is less than zero
0	0	1	0	Result is greater than zero

6.5.7 Store Floating Point (STE)

<u>Assembler Notation</u>	<u>Opcode</u>	<u>Format</u>
STE R1,D2(X2)	60	RX1,RX2
STE R1,A2(FX2,SX2)	60	RX3

Operation:

The floating point first operand, contained in the single precision floating point register specified by R1, is placed in the fullword memory location specified by the second operand address. The first operand is unchanged.

Condition Code:

Unchanged

Programming Note:

The second operand must be located on a fullword boundary.

6.5.8 Store Multiple Floating Point (STME)

<u>Assembler Notation</u>	<u>Opcode</u>	<u>Format</u>
STME R1,D2(X2)	71	RX1,RX2
STME R1,A2(FX2,SX2)	71	RX3

Operation:

The contents of successive single precision floating point registers, starting with the even-numbered register specified by R1, are stored in successive fullword memory locations, starting with the address of the second operand. The operation stops when the contents of floating point register 14 have been stored.

Condition Code:

Unchanged

Programming Note:

The second operand must be located on a fullword boundary.

6.5.9 Add Floating Point (AE, AER)

Add Floating Point (AE)

Add Floating Point Register (AER)

<u>Assembler Notation</u>	<u>Opcode</u>	<u>Format</u>
AE R1,D2(X2)	6A	RX1,RX2
AE R1,A2(FX2,SX2)	6A	RX3
AER R1,R2	2A	RR

Operation:

The two operand exponents are compared. If the exponents differ, the fraction with the smaller exponent is shifted right hexadecimally (four bits at a time), and its exponent is incremented by one for each hexadecimal shift, until the two exponents are equal. The hexadecimal digits (of four bits each) are shifted through the guard digits for additional precision. If no equalizing shifts are required, the guard digits remain zero. The fractions are then algebraically added. The guard digits participate in this addition.

If the addition of fractions produces a carry out of F1, the exponent of the result is incremented by one, and the fraction of the result is shifted right one hexadecimal digit. The carry bit is shifted back into the most significant hexadecimal digit of the fraction, producing a normalized result. This result is then R*-rounded and replaces the contents of the single precision floating point register specified by R1.

If the addition of fractions does not produce a carry, the result is normalized, if necessary, and R*-rounded. This result replaces the contents of the single precision floating point register specified by R1.

Condition Code:

C	V	G	L	
0	0	0	0	Floating point result is zero
0	0	0	1	Floating point result is less than zero
0	0	1	0	Floating point result is greater than zero
0	1	0	1	Exponent overflow, result is less than zero
0	1	1	0	Exponent overflow, result is greater than zero
0	1	0	0	Exponent underflow

Programming Notes:

If an exponent overflow is detected, an arithmetic fault interrupt is taken and the contents of the register specified by R1 remain unchanged.

Normalization of the result can produce exponent underflow. If PSW bit 19 is set, an arithmetic fault interrupt is taken, and the register specified by R1 is unchanged. If exponent underflow occurs and bit 19 of the current PSW is zero, no arithmetic fault occurs. Zeros replace the contents of the register specified by R1.

In the RX formats, the second operand must be located on a fullword boundary.

Example:

This example of the AE instruction adds the contents of LOC to the contents of LOC floating point register 8 and places the result in floating point register 8.

<u>Assembler Notation</u>	<u>Comments</u>
AE REG8,LOC	ADD (LOC) TO (REG8)

Where:

Floating point REG8 contains X'7EFF FFFF'.
LOC contains X'7EFF FFFF'

Result of AE Instruction:

(Floating point REG8) = 7F1F FFFF
(LOC) unchanged by this instruction
Condition code = 0010

6.5.10 Subtract Floating Point (SE, SER)

Subtract Floating Point (SE)

Subtract Floating Point Register (SER)

<u>Assembler Notation</u>	<u>Opcode</u>	<u>Format</u>
SE R1,D2(X2)	6B	RX1,RX2
SE R1,A2(FX2,SX2)	6B	RX3
SER R1,R2	2B	RR

Operation:

The two operand exponents are compared. If the exponents differ, the fraction with the smaller exponent is shifted right hexadecimally (four bits at a time), and its exponent is incremented by one for each hexadecimal shift until the two exponents are equal. The hexadecimal digits (of four bits each) are shifted through the guard digits for additional precision. If no equalizing shifts are required, the guard digits remain zero. The second operand fraction is then subtracted algebraically from the first operand fraction. The guard digits participate in this subtraction.

If the subtraction of fractions produces a carry out of F1, the exponent of the result is incremented by one, and the fraction of the result is shifted right one hexadecimal digit. The carry bit is shifted back into the most significant hexadecimal digit of the fraction, producing a normalized result. This result is then R*-rounded and replaces the contents of the single precision floating point register specified by R1.

If the subtraction of fractions does not produce a carry, the result is normalized, if necessary, then R*-rounded. This result replaces the contents of the single precision floating point register specified by R1.

Condition Code:

C	V	G	L	
0	0	0	0	Floating point result is zero
0	0	0	1	Floating point result is less than zero
0	0	1	0	Floating point result is greater than zero
0	1	0	1	Exponent overflow, result is less than zero
0	1	1	0	Exponent overflow, result is greater than zero
0	1	0	0	Exponent underflow

Programming Notes:

If an exponent overflow is detected, an arithmetic fault interrupt is taken and the contents of R1 remain unchanged.

Normalization of the result can produce exponent underflow. If PSW bit 19 is set, an arithmetic fault interrupt is taken, and the register specified by R1 is unchanged. If exponent underflow occurs and bit 19 of the current PSW is zero, no arithmetic fault occurs. Zeros replace the contents of the register specified by R1.

In the RX formats, the second operand must be located on a fullword boundary.

Example:

This example of the SE instruction subtracts the contents of LOC from the contents of floating point register 8 and places the result in floating point register 8.

Assembler Notation

Comments

```
SE REG8,LOC          SUBTRACT (LOC) FROM (REG8)
```

Where:

Floating point REG8 contains X'7EFF FFFF'
LOC contains X'7A10 0000'

Result of SE Instruction:

(Floating point REG8) = 7EFF FFEF
(LOC) unchanged by this instruction
Condition code = 0010

6.5.11 Compare Floating Point (CE, CER)

Compare Floating Point (CE)
Compare Floating Point Register (CER)

<u>Assembler Notation</u>	<u>Opcode</u>	<u>Format</u>
CE R1,D2(X2)	69	RX1,RX2
CE R1,D2(FX2,SX2)	69	RX3
CER R1,R2	29	RR

Operation:

The first and second operands are compared. Comparison is algebraic, and the sign, fraction and exponent of each number must be considered. The result is indicated by the condition code setting. Neither operand is changed.

Condition Code:

C	V	G	L	
=====				
0	X	0	0	First operand is equal to second
1	X	0	1	First operand is less than second
0	X	1	0	First operand is greater than second

Programming Notes:

The state of the V flag is undefined.

In the RX formats, the second operand must be located on a fullword boundary.

6.5.12 Multiply Floating Point (ME, MER)

Multiply Floating Point (ME)

Multiply Floating Point Register (MER)

Assembler Notation	Opcode	Format
ME R1,D2(X2)	6C	RX1,RX2
ME R1,A2(FX2,SX2)	6C	RX3
MER R1,R2	2C	RR

Operation:

The exponents of each operand, as derived from the excess-64 notation used in floating point representation, are added to produce the exponent of the result. This exponent is converted back to excess-64 notation, and the fractions are then multiplied.

If the product is zero, the entire floating point value is forced to zero, X'0000 0000'. If the product is not zero, the result is normalized. The sign of the result is determined by the rules of algebra. The R*-rounded result replaces the contents of the single precision floating point register specified by R1.

Condition Code:

C	V	G	L	
0	0	0	0	Floating point result is zero
0	0	0	1	Floating point result is less than zero
0	0	1	0	Floating point result is greater than zero
0	1	0	1	Exponent overflow, result is less than zero
0	1	1	0	Exponent overflow, result is greater than zero
0	1	0	0	Exponent underflow

Programming Notes:

Multiplication of two 6-hexadecimal digit fractions effectively produces a result of six hexadecimal digits and six guard digits. The guard digits participate in the R*-rounding of the final result.

The addition of exponents can produce exponent overflow. In this case, an arithmetic fault interrupt is taken, and both operands remain unchanged.

The addition of exponents or the normalization process can produce exponent underflow. If PSW bit 19 is set, an arithmetic fault interrupt is taken and the register specified by R1 is unchanged. If exponent underflow occurs and bit 19 of the current PSW is zero, no arithmetic fault occurs. Zeros replace the contents of the register specified by R1.

In the RX formats, the second operand must be located on a fullword boundary.

Example:

This example of the ME instruction multiplies the contents of floating point register 8 by the contents of memory location LOC and places the result in floating point register 8.

<u>Assembler Notation</u>	<u>Comments</u>
ME REG8,LOC	MULTIPLY (REG8) BY (LOC)

Where:

Floating point REG8 contains X'5FFF FFFF'
LOC contains X'60FF FFFF'

Result of ME Instruction:

(Floating point REG8) = 7FFF FFFE
(LOC) unchanged by this instruction
Condition code = 0010

6.5.13 Divide Floating Point (DE, DER)

Divide Floating Point (DE)

Divide Floating Point Register (DER)

Assembler Notation	Opcode	Format
DE R1,D2 (X2)	6D	RX1,RX2
DE R1,A2 (FX2,SX2)	6D	RX3
DER R1,R2	2D	RR

Operation:

The exponents of each operand, as derived from the excess-64 notation used in floating point representation, are subtracted to produce the exponent of the result. This exponent is converted back to excess-64 notation.

The first operand fraction is then divided by the second operand fraction. Division continues until the quotient is normalized, adjusting the exponent for each additional division required.

No remainder is returned. The sign of the quotient is determined by the rules of algebra. The simple-rounded quotient replaces the contents of the single precision floating point register specified by R1.

Condition Code:

C	V	G	L	
0	0	0	0	Floating point result is zero
0	0	0	1	Floating point result is less than zero
0	0	1	0	Floating point result is greater than zero
0	1	0	1	Exponent overflow, result is less than zero
0	1	1	0	Exponent overflow, result is greater than zero
0	1	0	0	Exponent underflow
1	1	0	0	Divisor equal to zero

Programming Notes:

Before starting the divide operation, the divisor is checked. If it is equal to zero, the operation is aborted, and the arithmetic fault interrupt is taken. Neither operand is changed.

Subtraction of exponents can produce exponent overflow. In this case, an arithmetic fault interrupt is taken, and both operands remain unchanged.

The subtraction of exponents or the division process can produce exponent underflow; normalization of the result can produce exponent underflow. If PSW bit 19 is set, an arithmetic fault interrupt is taken, and the register specified by R1 is unchanged. If exponent underflow occurs and bit 19 of the current PSW is zero, no arithmetic fault occurs. Zeros replace the contents of the register specified by R1.

The 6-hexadecimal digit first operand fraction is divided by the 6-hexadecimal digit second operand, effectively producing the 6-hexadecimal digit quotient along with a number of guard digits. The guard digits participate in the rounding of the final result.

In the RX formats, the second operand must be located on a fullword boundary.

Example:

This example of the DE instruction divides the contents of floating point register 4 by the contents of memory location LOC and places the result in floating point register 4.

<u>Assembler Notation</u>	<u>Comments</u>
DE REG4,LOC	DIVIDE (REG4) BY (LOC)

Where:

Floating point REG4 contains X'44FF FFFF' = Dividend
LOC contains X'0611 1111' = Divisor

Result of DE Instruction:

(Floating point REG4) = 7FF0 0000
(LOC) unchanged by this instruction
Condition code = 0010

6.5.14 Fix Register (FXR)

Assembler Notation	Opcode	Format
FXR R1,R2	2E	RR

Operation:

R1 and R2 specify a general-purpose register and a floating point register, respectively. The normalized floating point number contained in the floating point register is converted to a two's complement notation integer value by shifting and truncating. The result is stored in the general register specified by R1.

Condition Code:

C	V	G	L	
X	0	0	0	Result is zero or underflow
X	0	0	1	Result is less than zero
X	0	1	0	Result is greater than zero
X	1	0	1	Overflow, result is less than zero
X	1	1	0	Overflow, result is greater than zero

Programming Notes:

The range of floating point magnitudes (M) that produces a nonzero integral result is:

$$+X'4110\ 0000' \leq M \leq +X'4880\ 0000'$$

Floating point magnitudes greater than +X'487F FFFF' or -X'4880 0000' cause overflow. The result is forced to X'7FFF FFFF' if positive or to X'8000 0000' if negative. The V flag is set in the condition code along with either the G or L flag, depending on the sign of the result.

Floating point magnitudes less than +X'4110 0000' cause underflow, and the result is forced to zero.

In the event of overflow or underflow, no arithmetic fault interrupt is taken, even if enabled in the current PSW.

Example:

This example of the FXR instruction converts the contents of floating point register 8 to a fixed point number and places it in register 3.

<u>Assembler Notation</u>	<u>Comments</u>
FXR REG3,REG8	CONVERT (REG8) TO FIXED POINT

Where:

Floating point REG8 contains X'46FF FF00'
REG3 contains unknown data

Result of FXR Instruction:

(REG3) = 00FFFF00
(Floating point REG8) unchanged by this instruction
Condition code = 0010

6.5.15 Float Register (FLR)

Assembler Notation	Opcode	Format
FLR R1,R2	2F	RR

Operation:

R1 and R2 specify a floating point register and a general-purpose register, respectively. The integer value contained in the general register specified by R2 is converted to a floating point number and stored in the single precision floating point register specified by R1.

Condition Code:

C	V	G	L	
X	0	0	0	Floating point result is zero
X	0	0	1	Floating point result is less than zero
X	0	1	0	Floating point result is greater than zero

Programming Note:

The full range of fixed point integer values can be converted to floating point. The fixed point value X'7FFF FFFF', the largest positive integer, converts to the floating point value X'487F FFFF'. The fixed point value X'8000 0000', the most negative integer, converts to the floating point value X'C880 0000'. The result in R1 is normalized and truncated, if necessary, to fit in the six fraction digits.

Example:

This example of the FLR instruction converts the fixed point contents of register 4 to a floating point number and places it in floating point register 8.

Assembler Notation	Comments
FLR REG8,REG4	CONVERT (REG4) TO FLOATING POINT

Where:

REG4 contains X'7FFF FFF0'
 Floating point REG8 contains unknown data

Result of FLR Instruction:

(Floating point REG8) = 487FFFFF
(REG4) unchanged by this instruction
Condition code = 0010

6.5.16 Load Unnormalized Double Precision Floating Point (LW, LWR)

Load Unnormalized Double Precision Floating Point (LW)

Load Unnormalized Double Precision Floating Point Register (LWR)

<u>Assembler Notation</u>	<u>Opcode</u>	<u>Format</u>
LW R1,D2(X2)	4F	RX1,RX2
LW R1,A2(FX2,SX2)	4F	RX3
LWR R1,R2	1F	RR

Operation:

The doubleword second operand is placed in the double precision floating point register specified by R1. No normalization is performed.

Condition Code:

C	V	G	L	

0	0	0	0	Result is zero
0	0	0	1	Result is less than zero
0	0	1	0	Result is greater than zero

Programming Notes:

In the RX formats, the second operand must be located on a fullword boundary. This instruction is intended for data manipulation only. Floating point operations using data in a register loaded in this manner may not produce predictable results.

6.5.17 Load Double Precision Floating Point (LD, LDR, LDGR)

Load Double Precision Floating Point (LD)

Load Register Double Precision Floating Point (LDR)

Load Double Precision Floating Point Registers from General Registers (LDGR)

<u>Assembler Notation</u>	<u>Opcode</u>	<u>Format</u>
LD R1,D2(X2)	78	RX1,RX2
LD R1,A2(FX2,SX2)	78	RX3
LDR R1,R2	38	RR
LDGR R1,R2	A6	RR

Operation:

The floating point second operand is normalized, if necessary, and placed in the double precision floating point register specified by R1.

Condition Code:

C	V	G	L	
0	0	0	0	Double precision result is zero
0	0	0	1	Double precision result is less than zero
0	0	1	0	Double precision result is greater than zero
0	1	0	0	Exponent underflow

Programming Notes:

If the argument fraction is zero, the entire result is forced to zero, X'0000 0000 0000 0000'.

Normalization can produce exponent underflow. If PSW bit 19 is set, the arithmetic fault interrupt is taken, and the register specified by R1 remains unchanged. If exponent underflow occurs, and bit 19 of the current PSW is zero, no arithmetic fault occurs. Zeros replace the contents of the register specified by R1.

In the RX formats, the second operand must be located on a fullword boundary.

The R2 field for LDGR must specify the even member of an even/odd pair of general registers. The register specified by R2 contains the most significant 32 bits, and R2+1 contains the least significant 32 bits. If R2 does not specify an even-numbered register, unpredictable results occur.

6.5.18 Load Positive Double Precision Register (LPDR)

<u>Assembler Notation</u>	<u>Opcode</u>	<u>Format</u>
LPDR R1,R2	33	RR

Operation:

The double precision floating point second operand contained in the double precision floating point register specified by R2 is forced positive. The result is normalized if necessary and placed in the double precision floating point register specified by R1.

Condition Code:

C	V	G	L	
0	0	0	0	Double precision result is zero
0	0	1	0	Double precision result is greater than zero
0	1	0	0	Exponent underflow

Programming Notes:

If the argument fraction is zero, the entire result is forced to zero, X'0000 0000 0000 0000'.

Normalization of the result can produce exponent underflow. If PSW bit 19 is set, the arithmetic fault interrupt is taken, and the register specified by R1 remains unchanged. If exponent underflow occurs, and bit 19 of the current PSW is zero, no arithmetic fault occurs. Zeros replace the contents of the register specified by R1.

6.5.19 Load Complement Double Precision Register (LCDR)

<u>Assembler Notation</u>	<u>Opcode</u>	<u>Format</u>
LCDR R1,R2	37	RR

Operation:

The sign of the double precision floating point second operand contained in the double precision floating point register specified by R2 is complemented. The result is normalized, if necessary, and placed in the double precision floating point register specified by R1.

Condition Code:

C	V	G	L	
0	0	0	0	Double precision result is zero
0	0	0	1	Double precision result is less than zero
0	0	1	0	Double precision result is greater than zero
0	1	0	0	Exponent underflow

Programming Notes:

If the argument fraction is zero, the entire result is forced to zero, X'0000 0000 0000 0000'.

Normalization can produce exponent underflow. If PSW bit 19 is set, the arithmetic fault interrupt is taken and the register specified by R1 remains unchanged. If an exponent underflow occurs and bit 19 of the current PSW is zero, no arithmetic fault occurs. Zeros replace the contents of the register specified by R1 in this case.

6.5.20 Load Multiple Double Precision Floating Point (LMD)

<u>Assembler Notation</u>	<u>Opcode</u>	<u>Format</u>
LMD R1,D2(X2)	7F	RX1,RX2
LMD R1,A2(FX2,SX2)	7F	RX3

Operation:

Successive double precision floating point registers, starting with the register specified by R1, are loaded from successive fullword memory location pairs, starting with the address of the second operand. The process stops when double precision floating point register 14 has been loaded.

Condition Code:

Unchanged

Programming Notes:

Values loaded into the double precision floating point registers are assumed to be normalized, and no test or adjustment is performed.

The second operand must be located on a fullword boundary. Loading a register with a "dirty zero" using this instruction will result in a load of true zero.

6.5.21 Load General Registers from Double Precision Floating Point Register (LGDR)

Assembler Notation	Opcode	Format
LGDR R1,R2	16	RR

Operation:

The double precision floating point second operand, contained in the double precision register specified by R2, is placed in the general register pair specified by R1. The second operand is unchanged.

Condition Code:

C	V	G	L	
0	0	0	0	Result is zero
0	0	0	1	Result is less than zero
0	0	1	0	Result is greater than zero

Programming Notes:

The R1 field must specify the even member of the even/odd pair of general registers receiving the result. The even-numbered register receives the most significant 32 bits while the next sequential odd numbered register receives the least significant 32 bits.

If R1 and R2 do not specify even-numbered registers, unpredictable results occur.

6.5.22 Store Double Precision Floating Point (STD)

Assembler Notation	Opcode	Format
STD R1,D2(X2)	70	RX1,RX2
STD R1,A2(FX2,SX2)	70	RX3

Operation:

The floating point first operand, contained in the double precision floating point register specified by R1, is placed in the doubleword memory location specified by the second operand address. The first operand is unchanged.

Condition Code:

Unchanged

Programming Note:

The second operand must be located on a fullword boundary.

6.5.23 Store Multiple Double Precision Floating Point (STMD)

<u>Assembler Notation</u>	<u>Opcode</u>	<u>Format</u>
STMD R1,D2(X2)	7E	RX1,RX2
STMD R1,A2(FX2,SX2)	7E	RX1,RX2

Operation:

The contents of successive double precision floating point registers, starting with the even-numbered register specified by R1, are stored in successive fullword memory location pairs, starting with the address of the second operand. The operation stops when the contents of double precision floating point register 14 have been stored.

Condition Code:

Unchanged

Programming Note:

The second operand must be located on a fullword boundary.

6.5.24 Add Double Precision Floating Point (AD, ADR)

Add Double Precision Floating Point (AD)

Add Register Double Precision Floating Point (ADR)

Assembler Notation	Opcode	Format
AD R1,D2(X2)	7A	Format
AD R1,A2(FX2,SX2)	7A	RX3
ADR R1,R2	3A	RR

Operation:

The two operand exponents are compared. If the exponents differ, the fraction with the smaller exponent is shifted right hexadecimally (four bits at a time), and its exponent is incremented by one for each hexadecimal shift until the two exponents are equal. Hexadecimal digits are shifted through the guard digits to retain precision. The fractions are then added algebraically.

If the addition of fractions produces a carry out of F1, the exponent of the result is incremented by one and the fraction of the result is shifted right one hexadecimal position. The carry bit is shifted back into the most significant hexadecimal digit of the fraction, producing a normalized result. This result is R*-rounded and replaces the contents of the double precision floating point register specified by R1.

If the addition of fractions does not produce a carry, the result is normalized, if necessary, and placed in the double precision floating point register specified by R1.

Condition Code:

C	V	G	L	
0	0	0	1	Double precision result is less than zero
0	0	1	0	Double precision result is greater than zero
0	1	0	1	Exponent overflow, result is less than zero
0	1	1	0	Exponent overflow, result is greater than zero
0	1	0	0	Exponent underflow

Programming Notes:

If an exponent overflow is detected, an arithmetic fault interrupt is taken and both operands remain unchanged.

Normalization of the result can produce exponent underflow. If PSW bit 19 is set, an arithmetic fault interrupt is taken, and the register specified by R1 is unchanged. If exponent underflow occurs and bit 19 of the current PSW is zero, no arithmetic fault occurs. Zeros replace the contents of the register specified by R1.

In the RX formats, the second operand must be located on a fullword boundary.

6.5.25 Subtract Double Precision Floating Point (SD, SDR)

Subtract Double Precision Floating Point (SD)

Subtract Register Double Precision Floating Point (SDR)

<u>Assembler Notation</u>	<u>Opcode</u>	<u>Format</u>
SD R1,D2(X2)	7B	RX1,RX2
SD R1,A2(FX2,SX2)	7B	RX3
SDR R1,R2	3B	RR

Operation:

The two operand exponents are compared. If the exponents differ, the fraction with the smaller exponent is shifted right hexadecimally (four bits at a time), and its exponent is incremented by one for each hexadecimal shift until the two exponents are equal. Hexadecimal digits are shifted through the guard digits to retain precision. The second operand fraction is then subtracted algebraically from the first operand fraction.

If the subtraction of fractions produces a carry out of F1, the exponent of the result is incremented by one and the fraction of the result is shifted right one hexadecimal position. The carry bit is shifted back into the most significant hexadecimal digit of the fraction producing a normalized result. This result is R*-rounded and replaces the contents of the double precision floating point register specified by R1.

If the subtraction of fractions does not produce a carry, the result is normalized, if necessary, then R*-rounded and placed in the double precision floating point register specified by R1.

Condition Code:

C	V	G	L	
0	0	0	0	Double precision result is zero
0	0	0	1	Double precision result is less than zero
0	0	1	0	Double precision result is greater than zero
0	1	0	1	Exponent overflow, result is less than zero
0	1	1	0	Exponent overflow, result is greater than zero
0	1	0	0	Exponent underflow

Programming Notes:

If an exponent overflow is detected, an arithmetic fault is taken and the contents of R1 remain unchanged.

Normalization of the result can produce exponent underflow. If PSW bit 19 is set, an arithmetic fault interrupt is taken, and the register specified by R1 is unchanged. If exponent underflow occurs and bit 19 of the current PSW is zero, no arithmetic fault occurs. Zeros replace the contents of the register specified by R1.

In the RX formats, the second operand must be located on a fullword boundary.

6.5.26 Compare Double Precision Floating Point (CD, CDR)

Compare Double Precision Floating Point (CD)

Compare Register Double Precision Floating Point (CDR)

<u>Assembler Notation</u>	<u>Opcode</u>	<u>Format</u>
CD R1,D2(X2)	79	RX1,RX2
CD R1,A2(FX2,SX2)	79	RX3
CDR R1,R2	39	RR

Operation:

The first and second operands are compared. Comparison is algebraic, taking into account the sign, exponent and fraction of each number. The result is indicated by the condition code setting. Neither operand is changed.

Condition Code:

C	V	G	I.	
0	X	0	0	First operand is equal to second
1	X	0	1	First operand is less than second
0	X	1	0	First operand is greater than second

Programming Notes:

The state of the overflow flag is undefined.

In the RX formats, the second operand must be located on a fullword boundary.

6.5.27 Multiply Double Precision Floating Point (MD, MDR)

Multiply Double Precision Floating Point (MD)

Multiply Register Double Precision Floating Point (MDR)

<u>Assembler Notation</u>	<u>Opcode</u>	<u>Format</u>
MD R1,D2(X2)	7C	RX1,RX2
MD R1,A2(FX2,SX2)	7C	RX3
MDR R1,R2	3C	RR

Operation:

The exponents of the two operands, as derived from the excess-64 notation used in floating point representation, are added to produce the exponent of the result. This exponent is converted back to excess-64 notation. The fractions are then multiplied.

If the product is zero, the entire double precision value is forced to zero, X'0000 0000 0000 0000'. If the product is not zero, the result is normalized, if necessary. The sign of the result is determined by the rules of algebra. The R*-rounded result replaces the contents of the double precision floating point register specified by R1.

Condition Code:

C	V	G	L	
0	0	0	0	Double precision result is zero
0	0	0	1	Double precision result is less than zero
0	0	1	0	Double precision result is greater than zero
0	1	0	1	Exponent overflow, result is less than zero
0	1	1	0	Exponent overflow, result is greater than zero
0	1	0	0	Exponent underflow

Programming Notes:

Multiplication of two 14-hexadecimal digit fractions effectively produces a result of 14 hexadecimal digits and 14 guard digits. The guard digits participate in the R*-rounding of the final result.

The addition of exponents can produce exponent overflow. In this case, an arithmetic fault interrupt is taken and both operands remain unchanged.

Normalization of the result can produce exponent underflow. If PSW bit 19 is set, an arithmetic fault interrupt is taken, and the register specified by R1 is unchanged. If exponent underflow occurs and bit 19 of the current PSW is zero, no arithmetic fault occurs. Zeros replace the contents of the register specified by R1.

In the RX formats, the second operand must be located on a fullword boundary.

6.5.28 Divide Double Precision Floating Point (DD, DDR)

Divide Double Precision Floating Point (DD)

Divide Register Double Precision Floating Point (DDR)

Assembler Notation	Opcode	Format
DD R1,D2(X2)	7D	RX1,RX2
DD R1,A2(FX2,SX2)	7D	RX3
DDR R1,R2	7D	RR

Operation:

The exponents of the two operands, as derived from the excess-64 notation used in floating point representation, are subtracted to produce the exponent of the result. This exponent is converted back to excess-64 notation.

The first operand fraction is then divided by the second operand fraction. Division continues until the quotient is normalized, adjusting the exponent for each additional division required.

No remainder is returned. The sign of the result is determined by the rules of algebra. The simple-rounded quotient replaces the contents of the double precision floating point register specified by R1.

Condition Code:

C	V	G	L	
0	0	0	0	Double precision result is zero
0	0	0	1	Double precision result is less than zero
0	0	1	0	Double precision result is greater than zero
0	1	0	1	Exponent overflow, result is less than zero
0	1	1	0	Exponent overflow, result is greater than zero
0	1	0	0	Exponent underflow
1	1	0	0	Divisor equal to zero

Programming Notes:

Before starting the divide operation, the divisor is checked. If it is equal to zero, the operation is aborted, and the arithmetic fault interrupt is taken. Neither operand is changed.

The subtraction of exponents can produce exponent overflow. In this case, an arithmetic fault interrupt is taken and both operands remain unchanged.

Subtraction of exponents or the division process can produce exponent underflow. Normalization of the result can produce exponent underflow. If PSW bit 19 is set, an arithmetic fault interrupt is taken, and the register specified by R1 is unchanged. If exponent underflow occurs and bit 19 of the current PSW is zero, no arithmetic fault occurs. Zeros replace the contents of the register specified by R1.

The 14-hexadecimal digit first operand fraction is divided by the 14-hexadecimal digit second operand fraction, effectively producing the 14-hexadecimal digit quotient along with a number of guard digits. The guard digits participate in the rounding of the final result.

In the RX formats, the second operand must be located on a fullword boundary.

6.5.29 Fix Register Double Precision (FXDR)

Assembler Notation	Opcode	Format
FXDR R1,R2	3E	RR

Operation:

R1 and R2 specify a general-purpose register and a double precision floating point register, respectively. The normalized floating point number contained in the floating point register is converted to an integer value by shifting and truncating. The result is placed in the general register specified by R1.

Condition Code:

C	V	G	L	
X	0	0	0	Result is zero or underflow
X	0	0	1	Result is less than zero
X	0	1	0	Result is greater than zero
X	1	0	1	Overflow, result is less than zero
X	1	1	0	Overflow, result is greater than zero

Programming Notes:

The range of the floating point magnitude (M) that produces a nonzero integral result is:

$$+ X'4110\ 0000\ 0000\ 0000' \leq M \leq + X'4880\ 0000\ 0000\ 0000'$$

Double precision floating point magnitudes greater than +X'487F FFFF FFFF' or -X'4880 0000 0000 0000' cause overflow. The result is forced to X'7FFF FFFF' if positive or to X'8000 0000' if negative. The V flag is set in the condition code along with either the G or L flag, depending on the sign of the result.

Double precision floating point magnitudes less than +X'4110 0000 0000' cause underflow, and the result is forced to zero.

In the event of overflow or underflow, no arithmetic fault interrupt is taken even if enabled in the current PSW.

6.5.30 Float Register Double Precision (FLDR)

<u>Assembler Notation</u>	<u>Opcode</u>	<u>Format</u>
FLDR R1,R2	3F	RR

Operation:

R1 and R2 specify a double precision floating point register and a general-purpose register, respectively. The integer value contained in the general register specified by R2 is converted to a floating point number and placed in the double precision floating point register specified by R1.

Condition Code:

C	V	G	L	

X	0	0	0	Double precision result is zero
X	0	0	1	Double precision result is less than zero
X	0	1	0	Double precision result is greater than zero

Programming Notes:

The full range of fixed point integer values can be converted to double precision floating point. The fixed point value X'7FFF FFFF', the largest positive integer, converts to a double precision floating point value of X'487F FFFF FF00 0000'. The fixed point value X'8000 0000', the most negative integer, converts to a double precision floating point value of X'C880 0000 0000 0000'.

The result in R1 is normalized.

6.5.31 Load Single Precision Floating Point Register from Double (LED, LEDR)

Load Single Precision Floating Point Register from Double Precision Memory (LED)

Load Single Precision Floating Point Register from Double Precision Register (LEDR)

<u>Assembler Notation</u>	<u>Opcode</u>	<u>Format</u>
LED R1,D2(X2)	84	RX1,RX2
LED R1,A2(FX2,SX2)	84	RX3
LEDR R1,R2	A4	RR

Operation:

Double precision floating point data from the second operand location is R*-rounded to single precision accuracy, and placed in the single precision floating point register specified by R1.

Condition Code:

C	V	G	L	
0	0	0	0	Floating point result is zero
0	0	0	1	Floating point result is less than zero
0	0	1	0	Floating point result is greater than zero
0	1	0	0	Exponent underflow
0	1	0	1	Exponent overflow, result is less than zero
0	1	1	0	Exponent overflow, result is greater than zero

Programming Notes:

R1 and R2 must specify even-numbered registers.

Rounding of the result can cause exponent overflow. In this case, the register specified by R1 is unchanged, and the arithmetic fault interrupt is taken.

Normalization of the result can produce exponent underflow. If enabled by PSW bit 19, the arithmetic fault interrupt is taken, and the register specified by R1 remains unchanged. If bit 19 of the current PSW is zero, zeros replace the contents of the register specified by R1.

In the RX formats, the second operand must be located on a fullword boundary.

6.5.32 Load Double Precision Floating Point Register from Single (LDE, LDER)

Load Double Precision Floating Point Register from Single Precision Memory (LDE)

Load Double Precision Floating Point Register from Single Precision Register (LDER)

<u>Assembler Notation</u>	<u>Opcode</u>	<u>Format</u>
LDE R1,D2(X2)	87	RX1,RX2
LDE R1,A2(FX2,SX2)	87	RX3
LDER R1,R2	A7	RR

Operation:

Single precision floating point data from the second operand location is converted to double precision data by appending trailing zeros. The result replaces the contents of the double precision floating point register specified by R1.

Condition Code:

C	V	G	L	
0	0	0	0	Double precision result is zero
0	0	0	1	Double precision result is less than zero
0	0	1	0	Double precision result is greater than zero
0	1	0	0	Exponent underflow

Programming Notes:

The registers specified by R1 and R2 must be even-numbered registers.

Normalization of the result can produce exponent underflow. If enabled by PSW bit 19, the arithmetic fault interrupt is taken, and the register specified by R1 remains unchanged. If bit 19 of the current PSW is zero, no arithmetic fault occurs. Zeros replace the contents of the register specified by R1.

In the RX formats, the second operand must be located on a fullword boundary.

6.5.33 Store Double Precision Floating Point Register in Single Precision Memory (STDE)

<u>Assembler Notation</u>	<u>Opcode</u>	<u>Format</u>
STDE R1,D2(X2)	82	RX1,RX2
STDE R1,A2(FX2,SX2)	82	RX3

Operation:

Data from the double precision floating point register specified by R1 is R*-rounded to single precision accuracy and stored in the fullword second operand location.

Condition Code:

Unchanged

Programming Notes:

The register specified by R1 must be an even-numbered register.

Normalization of the rounded result can produce exponent underflow. In this case, zero, X'0000 0000', replaces the contents of the second operand location.

Rounding of the result can cause exponent overflow. In this case, the contents of the second operand location remain unchanged, and the arithmetic fault interrupt is taken.

The second operand must be located on a fullword boundary.

CHAPTER 7 STRING OPERATIONS

7.1 INTRODUCTION

String operations deal with operands that are strings of consecutive bytes in memory beginning and ending on byte boundaries. Information contained in such a string can represent packed decimal data or ASCII character information including unpacked decimal data.

7.2 DECIMAL DATA FORMAT DEFINITIONS

Decimal operands can be in either packed or unpacked (zoned) format. The decimal operands are considered as right-aligned integers. The address of a decimal operand specifies the address of the left-most or most significant byte of the operand.

7.2.1 Packed Decimal

A number represented in packed decimal format is a fixed point, signed integer and consists of from 1 to 16 consecutive bytes (see Figure 7-1). Each byte is divided into two digit fields; thus each byte, except for the right-most in the string, contains two decimal digits represented in binary code. The only values allowed in a decimal digit field are 0 through 9. The right-most byte in the string contains the least significant decimal digit and the sign digit.




Figure 7-1 Packed Decimal Format

7.3 DECIMAL AND ALPHANUMERIC STRING INSTRUCTION FORMATS

The two binary/decimal conversion instructions use the standard RX format. The remaining string operations use the RXX format.

In the instruction descriptions, the RXX format is diagrammed as follows:

$$\text{OP} \quad \left\{ \begin{array}{l} \text{R1} \\ =\text{L1} \end{array} \right\}, \left\{ \begin{array}{l} \text{D2(X2)} \\ \text{A2(FX2, SX2)} \end{array} \right\}, \left\{ \begin{array}{l} \text{R2} \\ =\text{L2} \end{array} \right\}, \left\{ \begin{array}{l} \text{D2(X2)} \\ \text{A2(FX2, SX2)} \end{array} \right\}$$

where any field can have either one of the options shown in the braces. R1/=L1 refers to the first operand length and R2/=L2 refers to the second operand length. Length of operand strings is always expressed as a number of bytes. These can vary from 0 to 15 for immediate length formats, and from 0 to maximum memory for register length. See Section 1.8.10 for further details of the RXX instruction format.

7.4 STRING INSTRUCTIONS

The string instructions are interruptible and use the scratchpad registers. If an interrupt occurs during the execution of a string instruction, bit 14 (IIP) is set by the processor in the old program status word (PSW) to indicate that the scratchpad registers contain information pertinent to the interrupted instruction. See Section 10.3.4 for further information.

The instructions described in this section are:

LPB	Load Packed Decimal String as Binary (convert from decimal to binary)
STBP	Store Binary as Packed Decimal String (convert from binary to decimal)
MVTU	Move Translated Until
MOVE	Move and Pad
MOVEP	Move and Pad with Default Pad
CPAN	Compare Alphanumeric
CPANP	Compare Alphanumeric with Default Pad
PMV	Pack and Move (convert unpacked decimal string to packed decimal string)
PMVA	Pack and Move Absolute (force positive result)
UMV	Unpack and Move (convert packed decimal string to unpacked decimal string)
UMVA	Unpack and Move Absolute (force positive result)

7.4.1 Load Packed Decimal String as Binary (LPB)

<u>Assembler Notation</u>	<u>Opcode</u>	<u>Format</u>
LPB R1,D2(X2)	6F	RX1,RX2
LPB R1,A2(FX2,SX2)	6F	RX3

Operation:

The second operand address points to the left-most byte of a packed decimal string of length 16 bytes (31 packed decimal digits plus sign). Digits of the operand are checked for validity as the operand is converted to a 64-bit, two's complement binary number. The result replaces the contents of the even/odd general register pair specified by R1 and R1+1.

Condition Code:

C	V	G	L	
0	0	0	0	Result is zero
0	0	0	1	Result is less than zero
0	0	1	0	Result is greater than zero
0	1	0	0	Overflow

Programming Notes:

This instruction is interruptible.

R1 must specify an even-numbered register, or unpredictable results will occur.

If an illegal decimal digit or sign digit is detected during conversion, the registers specified by R1 and R1+1 remain unchanged, and a data format fault interrupt is taken.

The largest positive number that can be processed without overflow is 9,223,372,036,854,775,807.

7.4.2 Store Binary As Packed Decimal String (STBP)

<u>Assembler Notation</u>	<u>Opcode</u>	<u>Format</u>
STBP R1,D2(X2)	6E	RX1,RX2
STBP R1,A2(FX2,SX2)	6E	RX3

Operation:

The contents of the even/odd general register pair specified by R1 and R1+1 are converted and stored in memory as a packed decimal string of length 16 bytes (31 packed decimal digits plus sign). The left-most byte is stored at the address specified by the second operand.

Condition Code:

C	V	G	L	

0	0	0	0	Result is zero
0	0	0	1	Result is less than zero
0	0	1	0	Result is greater than zero

Programming Notes:

This instruction is interruptible.

R1 must specify an even-numbered register, or unpredictable results will occur.

7.4.3 Move Translated Until (MVTU)

<u>Assembler Notation</u>	<u>Op- code</u>	<u>Function Code</u>	<u>Format</u>
MVTU { R1 } { D2(X2) } , { R2 } { D2(X2) } { =L1 } , { A2(FX2, SX2) } , { =L2 } , { A2(FX2, SX2) }	8C	00	RXRX

Operation:

General register 0 contains the escape character whose occurrence causes the instruction to terminate. General register 2 contains the address of a translation table. This translation table is a simple list of 256 single byte entries, not to be confused with the table used by the Translate instruction. The first operand string begins at the address specified by the first operand address. The length of this string is equal to either the contents of the register specified by R1 or the value of L1. The second operand string begins at the address specified by the second operand address. The length of this string is equal to either the contents of the register specified by R2 or the value of L2.

Successive bytes from the second operand string are moved to the first operand string, as follows:

1. A byte is fetched from the second operand string (this is the argument byte). The contents of general register 2 are tested. If general register 2 contains zero, no translation occurs. If general register 2 does not contain zero, it contains the address of a translation table of maximum size 256 bytes. In this case, the argument byte fetched from the second operand string is used as an index into the translation table, and the byte at the resulting address is fetched and used as the argument byte.
2. The argument byte is compared with the escape character contained in bits 24:31 of general register 0. If the bytes are the same, the C flag is set in the condition code, and the instruction terminates. Otherwise, the argument byte is stored in the first operand string, and the next successive byte is processed. This operation is repeated until either the escape character is encountered, the first operand string has been filled, or the second operand string has been exhausted.
3. When the instruction terminates, the address of the next byte to be moved from the second operand string is returned in general register 1.

Condition Code:

C	V	G	L	
0	0	0	0	Escape string moved
0	1	0	0	First operand filled before entire string moved
1	0	0	0	Escape character encountered

Programming Notes:

This instruction is interruptible.

The contents of general register 1 can change during instruction execution, but are not valid until instruction termination.

Bytes are moved from the second operand string to the first operand string in a left-to-right sequence. If the strings overlap such that the source is to the left of the destination, unpredictable results occur.

7.4.4 Move (MOVE, MOVEP)

Move and Pad (MOVE)

Move and Pad with Default Pad (MOVEP)

<u>Assembler Notation</u>	<u>Op-Code</u>	<u>Function Code</u>	<u>Format</u>
MOVE { R1 } { D2(X2) } { R2 } { D2(X2) } { =L1 }, { A2(FX2, SX2) }, { =L2 }, { A2(FX2, SX2) }	8C	01	RXR X
MOVEP { R1 } { D2(X2) } { R2 } { D2(X2) } { =L1 }, { A2(FX2, SX2) }, { =L2 }, { A2(FX2, SX2) }	8C	01	RXR X

Operation:

The first operand string begins at the address specified by the first operand address and has a length equal either to the contents of the register specified by R1 or to the value of L1. The second operand string begins at the address specified by the second operand address and has a length equal either to the contents of the register specified by R2 or to the value of L2.

Successive bytes from the second operand string are moved to the first operand string. If the second operand string is exhausted before the first operand string is filled, the remaining bytes in the first operand string are filled using the pad character. If MOVE is specified, the pad character is contained in bits 24:31 of general register 0. If MOVEP is specified, the remainder of the first operand is filled with ASCII space characters (X'20'). If the first operand string is filled before the second operand string is exhausted, overflow results, and the operation is terminated.

When the instruction terminates, the address of the next byte to be moved from the second operand string is returned in general register 1.

Condition Code:

C	V	G	L
0	0	0	0
0	1	0	0

Entire string moved

First operand filled before entire string moved

Programming Notes:

These instructions are interruptible.

The contents of general register 1 can change during instruction execution, but are not valid until instruction termination.

If MOVEP is specified, the contents of general register 0 are ignored.

Bytes are moved from the second operand string to the first operand string in a left to right sequence. If the strings overlap such that the source is to the left of the destination, unpredictable results occur.

7.4.5 Compare (CPAN, CPANP)

Compare Alphanumeric (CPAN)

Compare Alphanumeric with Default Pad (CPANP)

<u>Assembler Notation</u>	<u>Op- code</u>	<u>Function Code</u>	<u>Format</u>
CPAN { R1 } { D2(X2) } , { R2 } { D2(X2) } { =L1 } , { A2(FX2, SX2) } , { =L2 } , { A2(FX2, SX2) }	8C	02	RXRX
CPANP { R1 } { D2(X2) } , { R2 } { D2(X2) } { =L1 } , { A2(FX2, SX2) } , { =L2 } , { A2(FX2, SX2) }	8C	02	RXRX

Operation:

The first operand string begins at the address specified by the first operand address and has a length equal either to the contents of the register specified by R1 or to the value of L1. The second operand string begins at the address specified by the second operand address and has a length equal either to the contents of the register specified by R2 or to the value of L2.

The two strings are compared a byte at a time until the first unequal byte pair is found, or until the length of both strings is exhausted.

If the strings are of unequal length, the shorter string is logically extended to the length of the longer string. If CPAN is specified, this is done by using the pad character contained in bits 24:31 of general register 0. If CPANP is specified, the ASCII space character (X'20') is used as the default pad character.

Upon termination, general register 1 is set equal to one less than the number of second operand bytes in memory that successfully matched corresponding bytes in the first operand string. This count (or offset) includes pad characters if the second operand string was longer than the first.

For example, a first operand string of three bytes in length contains the characters ABC. A second operand string of six bytes in length contains the characters ABCDDD.

A CPANP instruction returns a condition code of 0001 (first operand string less than second operand string) and general register 1 is set equal to 2. The first nonmatching character was the character 'D' in the second operand string. Given the same operand strings, a CPAN instruction with general register 0 set equal to a pad character of 'D' returns a condition code of 0000 (strings are equal including pad characters), and general register 1 is set equal to 5.

Condition Code:

C	V	G	L
0	0	0	0
0	0	1	0
1	0	0	1

Strings are equal

First operand string greater than second

First operand string less than second

Programming Notes:

If CPANP is specified, the contents of general register 0 are ignored. If CPAN is specified, bits 0:23 of general register 0 are ignored.

These instructions are interruptible.

7.4.6 Pack and Move (PMV, PMVA)

Pack and Move (PMV)

Pack and Move Absolute (PMVA)

<u>Assembler Notation</u>	<u>Op- code</u>	<u>Function Code</u>	<u>Format</u>
PMV { R1 } { D2(X2) } { L2 } { D2(X2) } { =L1 }, { A2(FX2, SX2) } , { =L2 }, { A2(FX2, SX2) }	8C	03	RXRX
PMVA { R1 } { D2(X2) } { R2 } { D2(X2) } { =L1 }, { A2(FX2, SX2) } , { =L2 }, { A2(FX2, SX2) }	8C	03	RXRX

Operation:

The first operand string begins at the address specified by the first operand address. The length of this string in bytes is one greater than either the contents of the register specified by R1 or the value of L1. The second operand string begins at the address specified by the second operand address. The length of this string in bytes is one greater than either the contents of the register specified by R2 or the value of L2.

The second operand string consists of unpacked decimal data digits with a sign digit. Data in this string is packed and replaces the first operand string. Leading zeros are supplied as required to fill the higher order positions of the first operand string. The sign of the first operand string is forced to a standard value (C or D).

Condition Code:

C	V	G	L	
0	0	0	0	Result is zero
0	X	0	1	Result is less than zero
0	X	1	1	Result is greater than zero
0	1	1	1	Overflow
1	X	X	X	Invalid digit in second operand string

Programming Notes:

PMVA causes the sign digit of the first operand string to be forced positive.

Overflow occurs if the length of the first operand string is not sufficient to contain the packed representation of the second operand string. The V flag is set in the condition code, and the specified number of digits in the first operand string receive packed data from the second operand string. Higher order digits of packed data are lost in this case.

Leading zero digits do not cause overflow. They are truncated if necessary.

These instructions are interruptible instructions.

Since packing is done conceptually from right to left with any overlapping allowed, the instruction PMV can be used to check the validity of decimal data.

If the destination string is to the left of the source string such that the signed byte of the destination string is taken as data from the source string, the sign digit is found to be an illegal data digit, and the C flag is set at completion of the instruction.

7.4.7 Unpack and Move (UMV, UMVA)

Unpack and Move (UMV)

Unpack and Move Absolute (UMVA)

<u>Assembler Notation</u>	<u>Op- code</u>	<u>Function Code</u>	<u>Format</u>
UMV { R1 } { D2(X2) } { R2 } { D2(X2) } { =L1 }, { A2(FX2, SX2) }, { =L2 }, { A2(FX2, SX2) }	8C	04	RXR
UMVA { R1 } { D2(X2) } { R2 } { D2(X2) } { =L1 }, { A2(FX2, SX2) }, { =L2 }, { A2(FX2, SX2) }	8C	24	RXR

Operation:

The first operand string begins at the address specified by the first operand address. The length of this string in bytes is one greater than either the contents of the register specified by R1 or the value of L1. The second operand string begins at the address specified by the second operand address. The length of this string in bytes is one greater than either the contents of the register specified by R2 or the value of L2.

The second operand string consists of packed decimal data digits with a sign digit. Data in this string is unpacked and replaces the first operand string. Leading zeros are supplied as required to fill the higher order positions of the first operand string. The sign of the first operand string is forced to a standard value (C or D).

Condition Code:

C	V	G	L	
0	0	0	0	Result is zero
0	X	0	1	Result is zero
0	X	1	0	Result is greater than zero
0	1	X	X	Overflow
1	X	X	X	Invalid digit in second operand string

Programming Notes:

UMVA causes the sign digit of the first operand string to be forced positive.

Overflow occurs if the length of the first operand string is not sufficient to contain the unpacked representation of the second operand string. The V flag is set in the condition code, and the specified number of digits in the first operand string receive unpacked data from the second operand string. Higher order digits of unpacked data are lost in this case.

Leading zero digits do not cause overflow. They are truncated if necessary.

These instructions are interruptible instructions.

Since unpacking is done conceptually from right to left with any overlapping allowed, the instruction UMV can be used to check the validity of decimal data.

If the destination string is to the left of the source string such that the signed byte of the destination string is taken as data from the source string, the sign digit is found to be an illegal data digit, and the C flag is set at the completion of the instruction.

CHAPTER 8 HIGH-SPEED DATA HANDLING INSTRUCTIONS

8.1 INTRODUCTION

The data handling instructions are used to compute polynomial error check redundancy characters, as used by most data communications protocols. Communications protocols supported include, but are not limited to, the following:

- Binary Synchronous Communications (BISYNC or BSC) - IBM's* widely accepted half-duplex protocol uses the cyclic redundancy check (CRC) BISYNC error check polynomial ($X^{16} + X^{15} + X^2 + 1$).
- Synchronous Data Link Control (SDLC) - IBM's full-duplex protocol uses the CRC SDLC error check polynomial ($X^{16} + X^{12} + X^5 + 1$).
- Advanced Data Communications Control Procedure (ADCCP) - ANSI's proposed National Standard full-duplex protocol uses CRC SDLC.
- High Level Data Link Control (HDLC) - The International Standard Organization's (ISO) full-duplex protocol uses CRC SDLC.

8.2 DATA HANDLING INSTRUCTION FORMATS

The data handling instructions use the Register-to-Register (RR) and Register and Indexed Storage (RX) formats.

8.3 DATA HANDLING INSTRUCTIONS

The instructions described in this section are:


PB	Process Byte
PBR	Process Byte Register

* IBM is a registered trademark of International Business Machines Corporation

8.3.1 Process Byte (PB)

<u>Assembler Notation</u>	<u>Opcode</u>	<u>Format</u>
PB R1,D2(X2)	62	RX1, RX2
PB R1,A2(FX2,SX2)	62	RX3

Set Up:



Bits 24:31 of the register specified by R1 contain the data byte to be processed. Bits 8:15 of the register specified by R1 contain a check code to indicate the type of processing. This byte is interpreted as follows:

X'00'	Cumulative check zero (CRC BISYNC)
X'01'	Cumulative check one (CRC SDLC)
X'02'	Cumulative check two longitudinal redundancy check (LRC)

The second operand address points to a halfword residual checksum to be included in the cumulative check.

Operation:

If CRC BISYNC is specified, the data byte and the old residual checksum participate in the generation of a new residual checksum based on the evaluation of the polynomial $(X^{16} + X^{15} + X^2 + 1)$.

If CRC SDLC is specified, a similar operation is performed, using the polynomial $(X^{16} + X^{12} + X^5 + 1)$.

In both of these cases, the new residual checksum replaces the old residual checksum at the second operand location.

If LRC is specified, the Exclusive-OR of the data byte with the old residual checksum replaces the old residual checksum at the second operand location.

Condition Code:

Unchanged

Programming Notes:

Bits 0:7 and 16:23 of the register specified by R1 are ignored.

The register specified by R1 remains unchanged.

The second operand must be located on a halfword boundary.

Undefined check codes should not be used. If they are used, the results are undefined.

Example:

This example performs a process byte instruction and stores the residue in RESIDUE.

Assembler Notation

Comments

PB	R1,RESIDUE	RESIDUE on halfword boundary
----	------------	------------------------------

Where:

Register 1 contains X'0001007A' where 01 = CRC SDLC,
and 7A = DATA BYTE
Residue contains X'D053' = old residue


Result of PB Instruction:

(R1) unchanged by this instruction
(RESIDUE) = X'BC13' = new residue
Condition code unchanged by this instruction

8.3.2 Process Byte Register (PBR)

Assembler Notation	Opcode	Format
PBR R1,R2	32	RR

Set Up:



Bits 24:31 of the register specified by R1 contain the data byte to be processed. Bits 8:15 of the register specified by R1 contain a check code indicating the type of processing. This byte is interpreted as follows:

X'00'	Cumulative check zero (CRC BISYNC)
X'01'	Cumulative check one (CRC SDLC)
X'02'	Cumulative check two (LRC)

and is a fullword contained in the register specified by R2. Bits 16:31 of the second operand contain the residual checksum to be included in the processing.

Operation:

If CRC BISYNC is specified, the data byte and the old residual checksum participate in the generation of a new residual checksum, based on the evaluation of the polynomial $(X^{16} + X^{15} + X^2 + 1)$.

If CRC SDLC is specified, a similar operation is performed, using the polynomial $(X^{16} + X^{12} + X^5 + 1)$.

In both of these cases, the new residual checksum replaces the contents of bits 16:31 of the register specified by R2.

If LRC is specified, the Exclusive-OR of the data byte with the old residual checksum replaces the old residual checksum in the second operand.

Condition Code:

Unchanged

Programming Notes:

Bits 0:7 and 16:23 of the register specified by R1 are ignored. The register specified by R1 remains unchanged. Bits 0:15 of the register specified by R2 are not used and must be zero.

Undefined check codes should not be used. If they are used, the results are undefined.

CHAPTER 9 INPUT/OUTPUT (I/O) OPERATIONS

9.1 INTRODUCTION AND CONFIGURATION OF INPUT/OUTPUT (I/O) SYSTEM

I/O operations, as defined for the processor, provide a versatile means for the exchange of information between the processor, memory and external devices. Communication between the processor and external devices is accomplished over the I/O bus. Data transfers over the I/O bus require processor intervention, either programmed or automatic, for each item transferred.

The Model 3205 processor board includes one selector channel (SELCH). There is no provision for additional SELCHs.

Direct data transfers between external devices and memory are accomplished over the private I/O bus, and other program processing can proceed concurrently.

9.2 DEVICE CONTROLLERS

The basic functions of a device controller are to:

- Provide synchronization with the processor.
- Provide device address recognition.
- Transmit operational commands from processor to device.
- Translate device status into meaningful information for the processor.
- Request processor attention when required.

In addition, a controller can generate parity, convert serial data to parallel, buffer incoming or outgoing data, or perform other device-dependent functions.

9.2.1 Device Addressing

The system design allows as many as 1,022 external devices. Each device must have its own address or device number, ranging from X'001' to X'3FF' but not including X'0F0' (reserved for SELCH) (device number X'000' is not assigned). The minimum system provides for 255 device numbers. Larger systems can have either 511 or 1,022.

9.2.2 Processor/Controller Communication

Device controllers can communicate with the processor either directly, using the I/O bus, or indirectly through the SELCH. Communication between the processor and controller is a bidirectional, request/response operation.

The processor can initiate communication by sending the device address out onto the I/O bus. When a controller recognizes its address, it returns a synchronization signal to the processor and remains ready to accept commands from the processor. (All other devices become deselected.) The processor waits up to 50ms for the synchronization signal. If no signal is received within this period, the processor aborts the operation and notifies the controlling program. In this case, the status returned is X'04', known as False Sync. The condition code in the program status word (PSW) is also set to X'4' (V flag=1). Controller malfunction and software failure (incorrect device address) are the most common causes of this type of time-out.

A controller can initiate communication with the processor by generating an attention signal. If the processor is in an interruptible state as defined by PSW bit 17, this signal causes the processor to temporarily suspend the normal "fetch instruction/execute/fetch next instruction" operation at the end of the execute phase, and to transmit an acknowledge signal over the I/O bus. The controller requesting attention responds with a synchronization signal and transmits its device number to the processor.

9.2.3 Interrupt Queuing

Any device controller attempting to interrupt the processor activates the attention line and holds that line active until the processor acknowledges the interrupt. Requests for attention are asynchronous, and more than one request may be pending at any time. The system resolves these conflicts according to device priority, which is determined by the physical placement of the device controller on the I/O bus. When two or more device controllers request attention at the same time, the controller nearest to the processor in the RACK0/TACK0 priority wiring pattern captures the acknowledge signal from the processor and is serviced first. All other interrupting controllers of lower priority must wait for the next acknowledge signal from the processor.

9.3 INTERRUPT SERVICE POINTER TABLE (ISPT)

Device requests for service can result in either an immediate interrupt or an auto driver channel operation. The processor chooses one of these options according to information contained in the ISPT.

The ISPT is an ordered list containing one entry for each possible device number in the system. The table starts at memory location X'0000D0' and contains a halfword entry for each device number in the system. For a minimum system (255 device numbers), the table extends through memory location X'0002CF'; for a maximum system (1,022 device numbers), the table extends through memory location X'0008CF'. The software controlling I/O operations must set up the table.

When the processor receives the device address after acknowledging a request for service, it adds twice the device address to X'0000D0'. The result is the address, within the table, of the entry reserved for the device requesting attention.

If the entry in the table is even (bit 15 equals 0), the processor takes an immediate interrupt and transfers control to the software interrupt service routine (ISR) at the address contained in the table. If the entry in the table is odd (bit 15 equals 1), the processor transfers control to the auto driver channel, without interrupting the currently running program.

At the time the processor transfers control to the software ISR, the old PSW (current at the time of the device request) is saved in registers 0 and 1 of the new register set. The device number is saved in register 2 and the status in register 3. The status portion of the current PSW is replaced by the value X'0000280X', where X is the least significant four bits of the device status. Machine malfunction interrupts are enabled and all other interrupts are disabled. The entry in the ISPT is now the new location counter (LOC).

9.4 CONTROL OF INPUT/OUTPUT (I/O) OPERATIONS

The I/O structure allows several data transfers depending on the particular application and on the characteristics of the external devices. Primary methods of data transfer between the processor and external devices are listed below.

- One byte or one halfword to or from any of the general registers
- One byte or one halfword to or from memory
- A block of data to or from memory under control of the integrated SELCH
- Multiplexed blocks of data to or from memory under control of the auto driver channel

Standard device controllers require a predetermined sequence of commands to effect data transfers. These commands address the device, put it in the correct mode, and cause data to be transferred. Because all I/O instructions are privileged operations, I/O control programs must run in the supervisor mode, i.e., with PSW bit 23 of the current PSW zero. I/O control programs should disable immediate interrupts or enable only higher level interrupts, as controlled by PSW bits 17 and 20.

9.5 STATUS MONITORING INPUT/OUTPUT (I/O)

The simplest form of I/O programming is status monitoring I/O. In this mode of operation, only one device is handled at a time, and the processor cannot overlap other operations with the data transfer. The sequence of operations in this type of programming is shown below.

1. Address the device and set the proper mode (output command instruction).
2. Test the device status (sense status instruction).
3. Loop back to the sense status instruction until the status byte indicates that the device is ready (conditional branch instruction).
4. When the device is ready, transfer the data (read or write instruction).
5. If the transfer is not complete, branch back to the sense status instruction. If it is complete, terminate.

9.6 INTERRUPT DRIVEN INPUT/OUTPUT (I/O)

Interrupt driven I/O allows the processor to cope with the disparity in speed between itself and the external devices being controlled. With status monitoring, the processor spends time waiting for the device. With interrupt driven programming, the processor can use this time performing other functions. This kind of programming establishes at least two levels of operation. On one level are the interrupt service programs. On the other level are interruptible programs that run with the immediate interrupt enabled.

Before starting interrupt driven operations, the ISPT must be set up. This table starts at memory location X'0000D0' and must contain a halfword address entry for every possible device. The table is ordered according to device addresses so that X'0000D0' plus two times the device address equals the memory address of the table entry reserved for that device. The value placed in the location reserved for a device is the address of the ISR for the device.

For example, if a terminal is connected at an address of X'10' and the interrupt routine resides in memory at address X'3000', the setup involves writing X'3000' at memory location X'F0'. Note that X'F0' = X'D0' + 2 times the terminal address.

Although there may be gaps in device address assignments, the ISPT should be completely filled. Entries for nonexistent devices should point to an error recovery routine. This precaution prevents system failure in the event of spurious interrupts caused by hardware malfunction or by improper use of the simulate interrupt instruction.

The next step is to prepare the device for the transfer, preferably with the immediate interrupts disabled. Once the table pointer has been set up and the device prepared, the processor can move on to an interruptible program.

The sequence of operations in this type of program is listed below.

1. Set up the ISPT to vector to error addresses for undefined devices.
2. Store the address of the software ISR at two times the device number plus X'D0' (X'D0' is starting address of service pointer table).
3. Set up the software ISR.
4. Set up the device and enable device interrupts.
5. Enable I/O interrupts in the PSW.

When the device signals a need for service, the processor saves its current state and transfers control to the ISR at the location specified in the ISPT. At this time, the current PSW has a status that indicates running state, machine malfunction interrupt enabled, higher level I/O interrupts enabled and all other interrupts disabled. The condition code contains bits 4:7 of the device status. Registers 0 and 1 of the new set contain the old PSW, indicating the status and location of the interrupted program. Register 2 of that set contains the device address. Register 3 contains the device status.

The ISR should:

- check the device status in register 3, and if satisfactory,
- make the transfer, and
- return to the interrupted program by reloading the old PSW from registers 0 and 1 (LPSWR R0).

The ISR should not enable I/O interrupts. This would allow other interrupt requests to be acknowledged, and the contents of registers 0:4 could be lost. If it is necessary to enable immediate interrupts, the routine should save the register set, switch to a different register set, save it if necessary, and then enable immediate interrupts.

9.7 SELECTOR CHANNEL (SELCH) INPUT/OUTPUT

The integrated SELCH controls the transfer of data directly between high-speed devices and memory. As many as 16 devices can be attached to the SELCH, only one of which can be operating at any one time. The advantage in using the SELCH is that other program processing can proceed simultaneously with the transfer of data between the external device and memory. This is possible because the SELCH accesses memory through the processor, permitting the processor and the channel to share memory. Execution time of the program in progress may be affected, depending on the rate at which the SELCH and processor compete for memory cycles.

In the Model 3205 System, a single SELCH is integrated on the processor board. It has a simulated I/O bus device number X'F0', which the processor recognizes. Like I/O device controllers, it can request processor attention through the immediate interrupt.

9.7.1 Selector Channel (SELCH) Devices

The SELCH has a private bus similar to the processor's I/O bus. Controllers for the devices associated with the SELCH are attached to this bus. When the SELCH is idle, its private bus is connected directly to the I/O bus. If this condition exists, the processor can address, command and accept interrupt requests from the devices attached to the SELCH. When the SELCH is busy, this connection is broken. All communication between the processor and devices on the SELCH is cut off. Any attempt by the processor to address a device on the channel when the channel is busy results in instruction time-out.

9.7.2 Selector Channel (SELCH) Operation

Two registers in the SELCH hold the current memory address and the final memory address. With the use of write instructions, the control software places the address of the first byte of the data buffer into the current address register and the address of the last byte into the final address register. This is done before starting a SELCH operation. During the data transfer, the channel increments the current address register by two for each halfword transferred. When the transfer count indicates the last byte has been transferred, the channel terminates.

The SELCH accesses memory a minimum of one halfword at a time. The starting address of the data buffer must always be on an even byte (halfword) boundary. The starting address must be less than the final address.

Upon termination, the software should issue a STOP command to the SELCH, then read back from the SELCH the address contained in the current address register. If this address is not equal to the final address specified for the transfer, and if the buffer limits were properly checked before the transfer, this condition indicates a device malfunction or an unusual condition within the device. For example, crossing a cylinder boundary on a disk can result in an abnormal termination. The reason for the termination is indicated in the SELCH status or the device status.

9.7.3 Selector Channel (SELCH) Programming

The usual method of programming with the SELCH uses the immediate interrupt. The first step in the operation is to check the status of the SELCH. If the SELCH is not busy, the address of the termination ISR is placed in the location within the ISPT reserved for the SELCH. The program should then proceed as follows:

1. Give the SELCH a command to stop. This command initializes the SELCH registers and assures an idle condition with the private bus connected to the I/O bus, so that the device can be set up for data transfer.
2. Give the SELCH the starting and final addresses.
3. Prepare the device for the transfer with the required commands and information.
4. Give the SELCH the command to start.

With the start command, the SELCH breaks the connection between its private bus and the processor's I/O bus and provides a direct path between memory and the last device addressed over its bus. When the device becomes ready, the channel starts the transfer, which proceeds to completion without further programmed intervention on a cycle-steal basis. Once the start command has been given, the processor can proceed with the execution of concurrent programs.

Upon termination, the channel signals the processor that it requires service. The processor subsequently takes an immediate interrupt, transferring control to the SELCH ISR. At this time, registers Q:3 of the new set are set up as for any other immediate interrupt.

If a power fail/restore sequence occurs while using the SELCH, the contents of the SELCH's internal registers are undefined.

9.8 INPUT/OUTPUT (I/O) INSTRUCTION FORMATS

I/O instructions use the Register to Register (RR) and the Register and Indexed Storage (RX) instruction formats.

9.9 INPUT/OUTPUT (I/O) INSTRUCTIONS

Following most I/O instructions, the V flag in the condition code indicates instruction time-out. This means that the operation was not completed, either because the device did not respond at all, or because it responded incorrectly.

In the sense status (SS) and autoloading (AL) instructions, the V flag can also mean examine status. To distinguish between these two conditions, the program should test bits 0:3 of the device status byte. If all of these bits are zero, device time-out has occurred.

The instructions described in this section are:

OC	Output Command
OCR	Output Command Register
SS	Sense Status
SSR	Sense Status Register
RD	Read Data
RDR	Read Data Register
RH	Read Halfword
RHR	Read Halfword Register
WD	Write Data
WDR	Write Data Register
WH	Write Halfword
WHR	Write Halfword Register
AL	Autoload
SCP	Simulate Channel Program

9.9.1 Output Command (OC, OCR)

Output Command (OC)

Output Command Register (OCR)

<u>Assembler Notation</u>	<u>Opcode</u>	<u>Format</u>
OC R1,D2(X2)	DE	RX1,RX2
OC R1,A2(FX2,SX2)	DE	RX3
OCR R1,R2	9E	RR

Operation:

Bits 22:31 of the register specified by R1 contain the 10-bit device address. The processor addresses the device and transfers an 8-bit command byte from the second operand location to the device. Neither operand is changed.

Condition Code:

C V G L	

0 0 0 0	Operation successful
0 1 0 0	Instruction time-out (FALSE SYNC)

Programming Notes:

In the RR format, bits 24:31 of the register specified by R2 contain the device command.

These instructions are privileged operations.

9.9.2 Sense Status (SS, SSR)

Sense Status (SS)

Sense Status Register (SSR)

<u>Assembler Notation</u>	<u>Opcode</u>	<u>Format</u>
SS R1,D2(X2)	DD	RX1,RX2
SS R1,A2(FX2,SX2)	DD	RX3
SSR R1,R2	9D	RR

Operation:

Bits 22:31 of the register specified by R1 contain the 10-bit device address. The device is addressed, and the 8-bit device status is transferred to the second operand location. The condition code is set equal to the least significant four bits of the device status byte. The first operand is unchanged.

Condition Code:

Bits 4:7 of the device status byte are copied into the condition code. See the appropriate device manual for a description of this status.

If the device is not in the system, the condition code is set to 0100 (false sync). In this case, the status byte returned is X'04'.

Programming Notes:

In the RR format, the device status byte replaces bits 24:31 of the register specified by R2. Bits 0:23 are forced to zero.

These instructions are privileged operations.

9.9.3 Read Data (RD, RDR)

Read Data (RD)

Read Data Register (RDR)

Assembler Notation	Opcode	Format
RD R1,D2(X2)	D8	RX1,RX2
RD R1,A2(FX2,SX2)	D8	RX3
RDR R1,R2	9B	RR

Operation:

Bits 22:31 of the register specified by R1 contain the 10-bit device address. The processor addresses the device and transfers an 8-bit data byte from the device to the second operand location.

Condition Code:

C	V	G	L
0	0	0	0
0	1	0	0

Operation successful
Instruction time-out (FALSE SYNC)

Programming Notes:

In the RR format, the 8-bit data byte replaces bits 24:31 of the register specified by R2. Bits 0:23 of the register are forced to zero.

These instructions are privileged operations.

Instruction time-out does not prevent the second operand location from being modified.

9.9.4 Read Halfword (RH, RHR)

Read Halfword (RH)

Read Halfword Register (RHR)

<u>Assembler Notation</u>	<u>Opcode</u>	<u>Format</u>
RH R1,D2(X2)	D9	RX1,RX2
RH R1,A2(FX2,SX2)	D9	RX3
RHR R1,R2	99	RR

Operation:

Bits 22:31 of the register specified by R1 contain the 10-bit device address. The processor addresses the device. If the device is halfword-oriented, the processor transfers 16 bits of data from the device to the second operand location. If the device is byte-oriented, the processor transfers two 8-bit bytes in successive operations.

Condition Code:

C V G L	

0 0 0 0	Operation successful
0 1 0 0	Instruction time-out (FALSE SYNC)

Programming Notes:

If the device is byte-oriented, it must be capable of supplying both bytes without intervening status checks. This instruction does not perform status checking between the two byte transfers.

In the RR format, the data transferred from a halfword device replaces bits 16:31 of the register specified by R2. Bits 0:15 are forced to zero. The first byte of data from a byte device replaces bits 16:23 of the register specified by R2 and the second byte replaces bits 24:31. Bits 0:15 of the register specified by R2 are forced to zero.

In the RX format, the second operand must be located on a halfword boundary. The first byte of data from a byte device replaces bits 0:7 of the halfword operand in memory and the second byte replaces bits 8:15.

These instructions are privileged operations.

Instruction time-out does not prevent the second operand location from being modified.

9.9.5 Write Data (WD, WDR)

Write Data (WD)

Write Data Register (WDR)

Assembler Notation	Opcode	Format
WD R1,D2(X2)	DA	RX1,RX2
WD R1,A2(FX2,SX2)	DA	RX3
WDR R1,R2	9A	RR

Operation:

Bits 22:31 of the register specified by R1 contain the 10-bit device address. The processor addresses the device and transfers an 8-bit data byte from the second operand location to the device. Neither operand is changed.

Condition Code:

C V G L	

0 0 0 0	Operation successful
0 1 0 0	Instruction time-out (FALSE SYNC)

Programming Notes:

In the RR format, the 8-bit data byte is transferred from bits 24:31 of the register specified by R2.

These instructions are privileged operations.

9.9.6 Write Halfword (WH, WHR)

Write Halfword (WH)

Write Halfword Register (WHR)

<u>Assembler Notation</u>	<u>Opcode</u>	<u>Format</u>
WH R1,D2(X2)	D8	RX1,RX2
WH R1,A2(FX2,SX2)	D8	RX3
WHR R1,R2	98	RR

Operation:

Bits 22:31 of the register specified by R1 contain the 10-bit device address. The processor addresses the device. If the device is halfword-oriented, the processor transfers 16 bits of data from the second operand location to the device. If the device is byte-oriented, the processor transfers two 8-bit data bytes in successive operations.

Condition Code:

C V G L	

0 0 0 0	Operation successful
0 1 0 0	Instruction time-out (FALSE SYNC)

Programming Notes:

If the device is byte-oriented, it must be capable of accepting both bytes without intervening status checks. This instruction does not perform status checking between the two byte transfers.

In the RR format, data is transferred to a halfword device from bits 16:31 of the register specified by R2. The first byte of data is transferred to a byte device from bits 16:23 of the register specified by R2; the second byte comes from bits 24:31.

In the RX format, the second operand must be located on a halfword boundary. The first byte of data is transferred to a byte device from bits 0:7 of the halfword operand in memory and the second byte is transferred from bits 8:15.

These instructions are privileged operations.

9.9.7 Autoload (AL)

Assembler Notation	Opcode	Format
AL D2(X2)	D5	RX1,RX2
AL R1,D2(X2)	D5	RX1,RX2
AL A2(FX2,SX2)	D5	RX3
AL R1,A2(FX2,SX2)	D5	RX3

Operation:

The AL instruction loads memory with a block of data from an input device. The 8-bit input device address is specified by memory location X'000078'. The device command byte is specified by memory location X'000079'.

If the R1 field of this instruction is not specified or contains zero, the default value X'000080' is used for the start address of the data block in memory and the second operand address is used for the end of the data block. If the R1 field of this instruction contains a value other than zero, then the contents of the general registers specified by R1 and R1+1 are used for the start and end of the data block, respectively. If the start address is greater than the end address, the instruction is aborted.

The address of a SELCH is specified by memory location X'00007D'. If the byte at this location contains zero, the SELCH is not used by this instruction. In this case, data is transferred a byte at a time from the input device to successive memory locations, beginning with the specified block start address. If any blank or zero bytes are input before the first nonzero byte, these bytes are considered to be leader and are ignored. All other zero bytes are stored as data. When a data byte has been stored at the specified block end address, the instruction terminates.

If the SELCH address specified by memory location X'00007D' is X'F0', the SELCH is used to transfer data from the input device to successive memory locations, beginning with the specified block start address. All data bytes are transferred; no checking for leading zero bytes can be made. The instruction terminates when data has been stored at the specified block end address.

Condition Code:

C	V	G	L	
0	0	0	0	Operation successful or aborted
X	1	X	X	Examine status or time-out
X	X	1	X	End of medium (EOM)
X	X	X	1	Device unavailable

Programming Notes:

This instruction may be used only with devices whose addresses are less than or equal to X'FF'.

This instruction is a privileged operation.

Bad status termination results if any of the least significant three bits of the device status are set.

If the R1 field of this instruction is used, it must specify the even member of an even/odd register pair.

9.9.8 Simulate Channel Program (SCP)

Assembler Notation	Opcode	Format
SCP R1,D2(X2)	E3	RX1,RX2
SCP R1,A2(FX2,SX2)	E3	RX3

Operation:

The second operand address is the address of a channel control block (CCB). The buffer switch bit of the channel command word (CCW) specifies the buffer to be used for the data transfer. If this bit is set, buffer 1 is used. If it is zero, buffer 0 is used. If the byte count field of the current buffer is greater than zero, the V flag in the condition code is set, and the next sequential instruction is executed. If the byte count field is not greater than zero, the following data transfer operation is performed.

If the CCW specifies read, a byte of data is moved from bits 24:31 of the register specified by R1 to the appropriate buffer location. If the CCW specifies write, a byte of data is moved from the appropriate buffer location to bits 24:31 of the register specified by R1. Bits 0:23 are forced to zero.

After a byte has been transferred, the count field of the appropriate buffer is incremented by one. If the count field is now greater than zero, and if the fast bit of the CCW is zero, the buffer switch bit of the CCW is complemented.

Condition Code:

C	V	G	L	
0	0	0	0	Count field is now zero
0	0	0	1	Count field is now less than zero
0	0	1	0	Count field is now greater than zero
0	1	0	0	Count field was greater than zero

Programming Notes:

If the CCW specifies fast mode, buffer 1 can be used, but the buffer bit is not switched when the count field becomes greater than zero.

The second operand must be located on a fullword boundary.

This instruction is a privileged operation.

9.10 AUTO DRIVER CHANNEL

The auto driver channel provides a means for multiplexing block data transfers between memory and low- or medium-speed I/O devices. The channel operation is similar, in some respects, to interrupt driven I/O. The channel is activated as a result of a service request from a device on the I/O bus. Upon receipt of such a request, the processor uses the device number to index into the ISPT. If the value contained in the table is even, the processor transfers control to the interrupt service routine. If the value is odd, it transfers control to the auto driver channel.

To the auto driver channel, the address in the ISPT is the address plus one (making it odd) of a CCB. The CCB is a channel program consisting of a description of the operation to be performed and a list of parameters associated with the operation. In addition to the functions of read and write, the channel can also:

- translate characters,
- test device status,
- chain buffers,
- calculate longitudinal and cyclic redundancy check (CRC) values, and
- transfer control to software routines to take care of unusual situations.

9.11 CHANNEL COMMAND BLOCK (CCB)

The CCB, as shown in Figure 9-1, consists of a CCW (16 bits) that describes the function, count fields (16 bits each) for two buffers, final addresses (32 bits each) for two buffers, a check word (16 bits) for the longitudinal redundancy check (LRC) or CRC, the address (32 bits) of a translation table, and the address (16 bits) of a software routine. The CCB requires 22 bytes of memory.

Many ISRs can be available at any time to service device requests. There can also be many CCBs in the system ready to handle data transfers as required. Each CCB must be aligned on a fullword boundary. The CCB address plus one must be placed in the ISPT location for the device involved in the transfer.

	0	15
0	CHANNEL COMMAND WORD (HALFWORD)	
2	BUFFER BYTE COUNT (HALFWORD)	
4	BUFFER 0 END ADDRESS (FULLWORD)	
8	CHECK WORD (HALFWORD)	
10	BUFFER 1 BYTE COUNT (HALFWORD)	
12	BUFFER 1 END ADDRESS (FULLWORD)	
16	TRANSLATION TABLE ADDRESS (FULLWORD)	
20	SUBROUTINE ADDRESS (HALFWORD)	

Figure 9-1 Channel Command Block

9.11.1 Subroutine Address

To handle special situations, channel control is transferred to the software subroutine, whose address is contained in the CCB. When this occurs, registers 0:4 of the appropriate set have already been set up by the processor to contain the old PSW, the device number, the device status and the address of the CCB. The current PSW status specifies run state, machine malfunction interrupt enabled, higher level I/O interrupts enabled and all other interrupts disabled.

The channel transfers control to the subroutine unconditionally (controlled by a bit in the CCW) because of bad device status, special character translation, or because it has reached the limit of a buffer. It indicates its reason for transferring control by adjusting the condition code as follows:

C	V	G	L	
0	0	0	0	Unconditional transfer of special character
0	0	0	1	Bad status
0	0	1	0	Buffer limit

The subroutine address in the CCB is a 16-bit physical address. For this reason, the subroutine at that address, or at least the first instruction of the subroutine, must reside in the 64kb of memory.

9.11.2 Buffers

There is a space in the CCB to describe two data buffer areas. The data areas can be located anywhere in memory. The limits of each data area are described by an address field and a count field. The address field contains the physical address of the last byte in the data area. This address is right-justified in the fullword provided. If the device being controlled is a halfword-oriented device, the final address must be odd. If the device is a byte-oriented device, the address can be either odd or even. The active buffer is selected by a bit in the CCW. When one buffer has been exhausted, the channel may reverse the state of this bit and switch to the alternate buffer. Automatic buffer switching is available only for byte-oriented devices and if the fast bit of the CCW is zero. If the fast bit is set, buffer 0 is always used.

The count field, in most operations, contains a negative number whose absolute value is equal to one less than the number of bytes to be transferred. The one exception is the case of a single data transfer, for which the count field contains zero.

During data transfers, the channel adds the value contained in the count field to the final address in order to obtain the current address. It makes the transfer, using the current address, then increments the value in the count field by one for a byte device or by two for a halfword device. When the count field becomes greater than zero, the channel sets the G flag in the condition code and transfers control to the specified software subroutine. If the count field is greater than zero upon channel activation, the channel makes no transfer and relinquishes control of the processor.

9.11.3 Translation

The translation feature is available only for byte-oriented devices or if the fast bit in the CCW is zero. If translation is specified, the fullword provided in the CCB must contain the address, which is right-justified, of a translation table. This table, which must be aligned to a halfword boundary, can contain up to 256 halfword entries. The format of this table is identical to that used by the Translate (TLATE) instruction (see Section 3.3.2). During data transfers, the channel multiplies the data byte by two and adds this value to the translation table address. The result is the address within the translation table of the halfword entry corresponding to the data byte.

The channel tests this entry, and, if bit 0 of the halfword is set, it substitutes bits 8:15 of the halfword for the data byte and proceeds with the operation. If bit 0 of the halfword is a zero, the channel:

- does not increment the byte count for the appropriate buffer;

- puts the data byte, untranslated, in bits 24:31 of register 3 of the appropriate set and forces bits 0:23 of register 3 to zero; and
- multiplies the value contained in the translation table by two and transfers control to the software special character translation routine located at the resulting address.

Upon transfer to the translation subroutine, registers 0 and 1 contain the old PSW; register 2 contains the device number; register 3 contains the untranslated character; and register 4 contains the address of the CCB. The current PSW indicates run state, machine malfunction interrupt enabled, higher level I/O interrupts enabled and all other interrupts disabled. The condition code is zero.

9.11.4 Check Word

The check word in the CCB contains the accumulated residual for LRC or CRC. The initial value for the check word is usually zero. (There are data-dependent exceptions, e.g., where initial characters are not to be included in the check.)

The longitudinal check is an Exclusive-OR of the character with the check word.

The CRC uses the formula for CRC16:

$$X^{16} + X^{15} + X^2 + 1$$

If the data communication option is equipped, the CRC can optionally use the formula for CRC SDLC:

$$X^{16} + X^{12} + X^5 + 1$$

On input, if both redundancy checking and translation are required, the character is translated first, then the CRC is done using the original character input rather than the translated character. On output, the translated character participates in the redundancy check. Redundancy checking can be used only with byte devices and is only performed if the fast bit of the CCW is zero.

9.11.5 Channel Command Word (CCW)

The CCW, as shown in Figure 9-2, consists of two parts. Bits 0:7 contain a status mask. Bits 8:15 describe the channel operation.




Figure 9-2 Channel Command Word

Status Mask

On every channel operation, if the execute bit is set, the status mask is ANDed with the device status. This operation does not change the status mask. If the result is zero, the channel proceeds with the operation. If the result is nonzero, the channel sets the L flag in the condition code and transfers control to the specified software subroutine.

Execute Bit (E)

If this bit is zero, the channel unconditionally transfers control to the specified subroutine without taking any other action. The condition code is zero. If this bit is set, the channel continues with the operation as specified in the CCW.

Fast Bit (F)

If this bit is set, the channel performs the I/O transfer in the fast mode. In this mode, buffer switching, redundancy checking and translation are not allowed. This bit must be set for halfword devices. If this bit is set, buffer 0 is always used.

Read/Write Bit (R/W)

This bit indicates the type of operation. If this bit is zero, a byte or a halfword is input from the device. If this bit is set, a byte or a halfword is output to the device.

Translate Bit (T)

If this bit is set and the fast bit is zero, the channel translates the data byte using the translation table defined in the CCB.

Redundancy Check Type Bits (RC)

These two encoded bits specify the type of redundancy check required. No check is performed if the fast bit is set. Table 9-1 contains the valid types of checks.

TABLE 9-1 VALID REDUNDANCY CHECKS

BIT 10	BIT 11	REDUNDANCY CHECK TYPE
0	0	LRC
0	1	CRC BISYNC
1	0	Reserved; must not be specified
1	1	CRC SDLC

Buffer Switch Bit (B)

When zero, this bit specifies that buffer 0 is to be used for the transfer. If it is set, buffer 1 is used. The channel chains buffers when the count field becomes greater than zero by complementing the buffer switch bit before transferring control to the specified software routine. Buffer 0 is always used if the fast bit in the CCW is set.

9.11.6 Valid Channel Command Codes

Table 9-2 is a list of valid codes for the CCW, bits 8:15. Note that only the first three can be used with halfword devices.

TABLE 9-2 CHANNEL COMMAND WORD

HEXADECIMAL	BINARY	MEANING
00	00000000	Transfer to subroutine
81	10000001	Read fast mode
85	10000101	Write fast mode
80	10000000	LRC, Buffer 0, read
82	10000010	LRC, Buffer 0, read, translate
84	10000100	LRC, Buffer 0, write
86	10000110	LRC, Buffer 0, write, translate
88	10001000	LRC, Buffer 1, read
8A	10001010	LRC, Buffer 1, read, translate
8C	10001100	LRC, Buffer 1, write
8E	10001110	LRC, Buffer 1, write, translate
90	10010000	CRC BISYNC, Buffer 0, read
92	10010010	CRC BISYNC, Buffer 0, read, translate
94	10010100	CRC BISYNC, Buffer 0, write
96	10010110	CRC BISYNC, Buffer 0, write, translate
98	10011000	CRC BISYNC, Buffer 1, read
9A	10011010	CRC BISYNC, Buffer 1, read, translate
9C	10011100	CRC BISYNC, Buffer 1, write
9E	10011110	CRC BISYNC, Buffer 1, write, translate
B0	10110000	CRC SDLC, Buffer 0, read
B2	10110010	CRC SDLC, Buffer 0, read, translate
B4	10110100	CRC SDLC, Buffer 0, write
B6	10110110	CRC SDLC, Buffer 0, write, translate
B8	10111000	CRC SDLC, Buffer 1, read
BA	10111010	CRC SDLC, Buffer 1, read, translate
BC	10111100	CRC SDLC, Buffer 1, write
BE	10111110	CRC SDLC, Buffer 1, write, translate

9.11.7 General Auto Driver Channel Programming Procedure

The following steps describe the general auto driver channel programming procedure. See Figure 9-3 for a flowchart of the procedure.

1. Set up ISPT to vector to error routines for undefined devices.
2. Set up address of CCW + 1 (odd) in table at two times device number plus X'D0' (start of ISPT).
3. Set up complete CCB.
4. Set up device and enable device interrupt.
5. Enable I/O interrupts in PSW (auto driver channel performs I/O operation).

6. Check for good termination of auto driver channel operation when the subroutine defined in the CCB is entered.

620

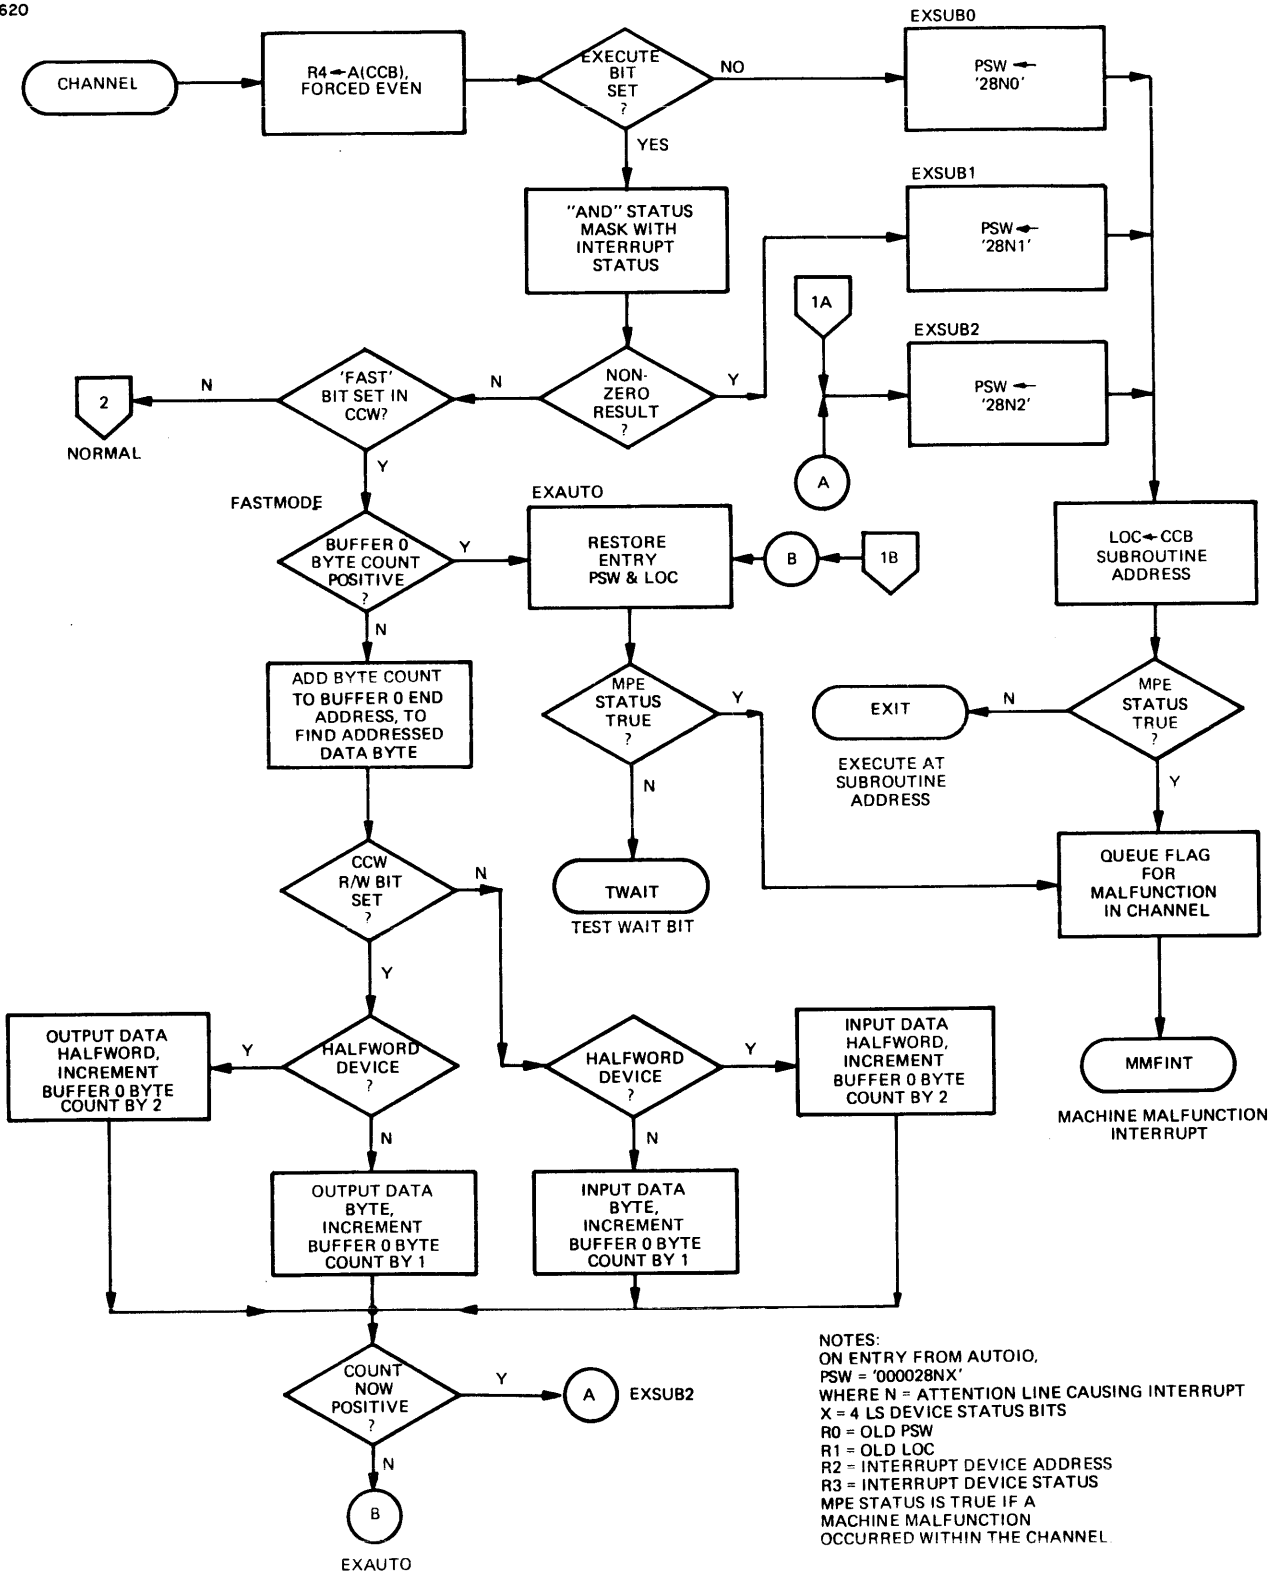


Figure 9-3 Auto Driver Channel Flowchart

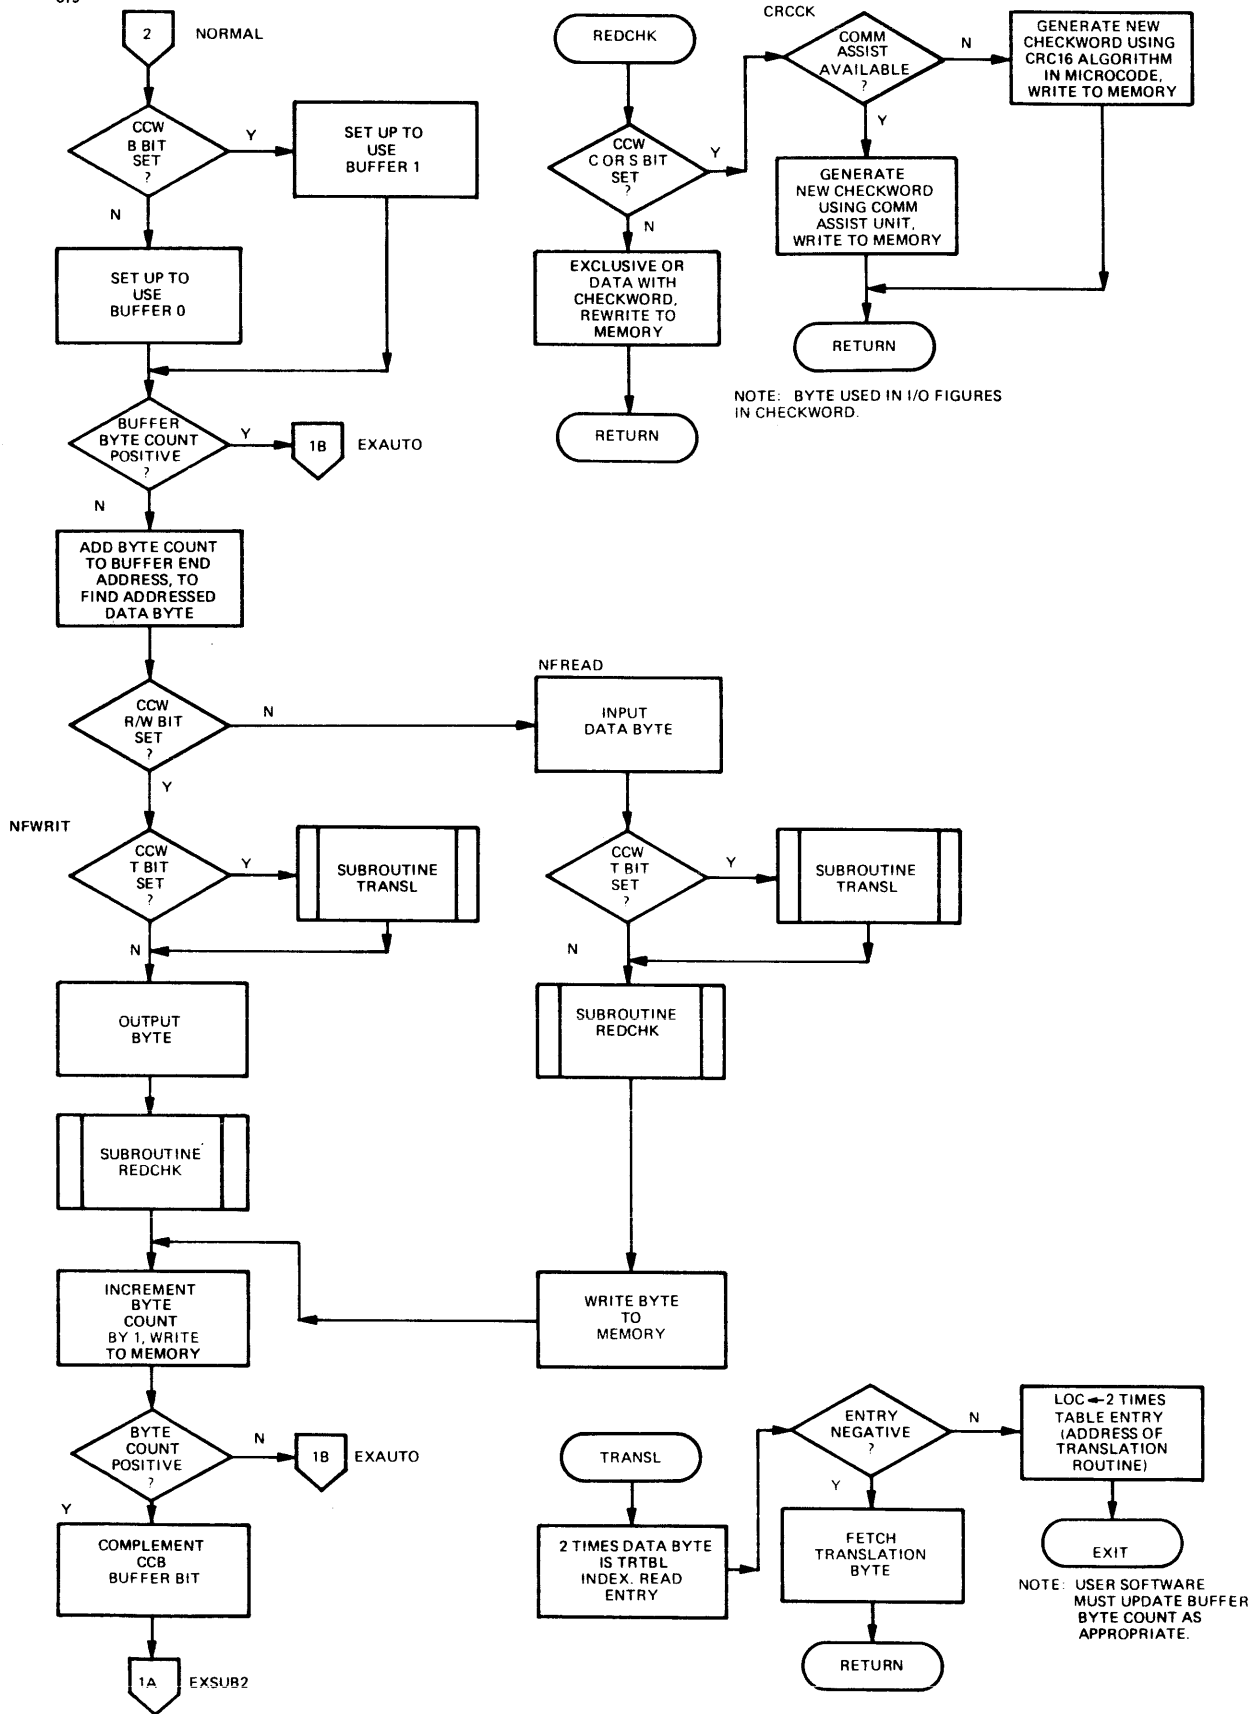


Figure 9-3 Auto Driver Channel Flowchart (Continued)

CHAPTER 10 STATUS SWITCHING AND INTERRUPTS

10.1 INTRODUCTION

The processor's interrupt system provides a mechanism for escape from the normal processing sequence to handle external and internal events. The software routine that is executed in response to an interrupt is called an interrupt service routine (ISR). Before transferring control to a service routine, the current state of the processor is preserved so that, upon completion of the service routine, the execution of an interrupted program can be resumed.

Interrupts can be classified as being synchronous or asynchronous, depending on whether they occur in fixed relationship to the execution of instructions or at random times due to events external to the processor, respectively. Examples of asynchronous interrupts include power fail, console attention and peripheral device interrupts.

Synchronous interrupts occur due to fault conditions or, in the case of software interrupts, may be programmed to occur. Examples of fault conditions that cause synchronous interrupts include noncorrectable memory errors, illegal instructions and arithmetic faults.

Software interrupts occur when the Supervisor Call (SVC) or Simulate Interrupt (SINT) instructions are executed, or as a result of adding an entry to the system queue. The Breakpoint (BRK) instruction causes program execution to be suspended so that the system console terminal may be activated. See Chapter 2.

Each interrupt condition is reset when the corresponding interrupt is serviced by the processor.

10.2 PROGRAM STATUS WORD (PSW) AND RESERVED MEMORY LOCATIONS

The PSW shown in Figure 10-1 is a 64-bit quantity that controls the operation of the processor. The PSW provides information about various states and conditions affecting the operation of the processor. The PSW is composed of two fullwords: bits 0:31 are the status word, and bits 32:63 are the location counter (LOC). The various PSW fields are described in Table 10-1.

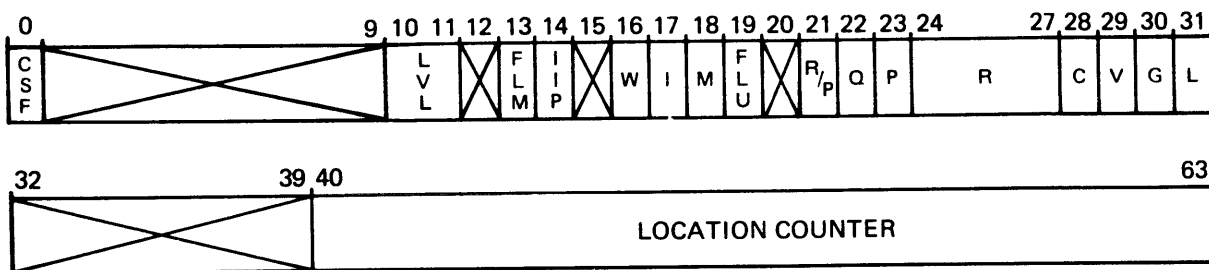


Figure 10-1 Program Status Word

TABLE 10-1 PSW BITS

BIT	MNEMONIC	MEANING
0	CSF	Must be zero; IF SET, CATASTROPHIC SYSTEM FAILURE
1:9		Unused; must be zero
10:11	LVL	Memory access level
12		Reserved; must be zero
13	FLM	Floating point arithmetic masked mode
14	IIP	Interruptible instruction in progress
15		Reserved; must be zero
16	W	Wait state
17	I	I/O interrupt mask
18	M	Machine malfunction interrupt mask
19	FLU	Floating point arithmetic underflow mask
20		Reserved; must be zero
21	R/P	Relocation/protection interrupt mask
22	Q	System queue interrupt mask
23	P	Protect mode
24:27	R	Register set select bits

TABLE 10-1 PSW BITS (Continued)

BIT	MNEMONIC	MEANING
28:31	C,V,G,L	Condition code
32:39		Reserved; must be zero
40:63		Program address (LOC)

10.2.1 Program Status Word (PSW)

Bits 0:31 of the PSW are called the status word. This word controls interrupts, defines the status of the processor, and contains the condition code. The following sections provide detailed definitions of various states of the processor and how the status word controls them. Unused bits of the status word must always be zero, with the exception of bit 20.

10.2.1.1 Catastrophic System Failure (CSF)

Bit 0 of the PSW is known as the CSF indicator and must not be set by the user. If the hardware detects a failure previously identified as impossible to recover from reliably, the system will pull a break to the console service routine and set this bit.

10.2.1.2 Memory Access Level Field (LVL)

When PSW bit 21 (R/P) is set, PSW bits 10 and 11 participate in an access level check for any memory access attempted by the current program. The LVL field of PSW is compared numerically to the access level field of the appropriate segment table entry. If the LVL field contains a lesser value than the access level field, a MAT fault interrupt occurs.

When PSW bit 21 is zero, PSW bits 10 and 11 are ignored, and no access level check is performed.

10.2.1.3 Floating Point Masked Mode (FLM)

When bit 13 of the current PSW is zero, a program can execute any legal floating point instruction.

When bit 13 of the current PSW is set, the processor is in the FLM mode. A program running in this mode is not allowed to execute floating point arithmetic instructions. If execution of any floating point arithmetic instruction is attempted in FLM mode, an illegal instruction interrupt occurs. If the processor is in FLM mode when a context switch is made by the system program and the processor state must be saved, the contents of the floating point registers need not be saved. This results in a faster context switch.

10.2.1.4 Interruptible Instruction in Progress (IIP)

PSW bit 14 is set by the processor while an interruptible instruction is in progress and is zero when the interruptible instruction terminates. This bit is set by the processor to indicate that the scratchpad registers contain valid parameters for the interruptible instruction and that these parameters need not be recalculated before resuming the interrupted instruction.

If bit 14 of the current PSW is set when the processor transfers control to a software ISR, that routine must not allow the contents of the scratchpad registers to be modified before the interruptible instruction is resumed. The STPS, LDPS, ISSV and ISRST instructions provide the means for saving and restoring these registers if they must be used by the ISR.

10.2.1.5 Wait State (W)

When PSW bit 16 is set, the processor is in the wait state. In the wait state, the normal fetch instruction/execute instruction/fetch next instruction sequence is suspended. While in the wait state, the processor is responsive to console attention interrupts and primary power fail (PPF), as well as any interrupts specifically enabled by the current PSW.

PSW bit 16 is zero when the processor is executing instructions. This bit is forced to zero whenever the single step, run switch or system console terminal is used to initiate instruction execution. This bit is not forced set by entry to the console mode.

If an interrupt occurs, PSW bit 16 is set according to the new PSW defined for servicing the interrupt. Bit 16 of the new PSW for any I/O interrupt is zero.

Except for an I/O interrupt, the state of bit 16 of the new PSW is tested as the PSW is loaded. If bit 16 of the newly loaded PSW is set, the processor enters the wait state, provided that no interrupt is still pending. All pending interrupts are serviced before the processor enters the wait state.

10.2.1.6 Input/Output (I/O) Interrupt Mask (I)

PSW bit 17 is used to enable or disable recognition of interrupt requests generated by peripheral devices on any of the interrupt levels, as detailed below:

BIT 17	MEANING
0	Disabled
1	I/O interrupts enabled

*system queue?
(hardware)*

An I/O interrupt request is queued until the processor acknowledges the interrupt unless the request is programmed reset, or power fail occurs. The state of PSW bit 17 is ignored by the Simulate Interrupt (SINT) instruction.

10.2.1.7 Machine Malfunction Interrupt Enable (M)

PSW bit 18 is used to enable and disable detection of various malfunction conditions within the processor and the resulting machine malfunction interrupt. When this bit is set, any of the following conditions results in a machine malfunction interrupt.

- Early power failure (EPF)
- Power restore
- Noncorrectable memory data error
- Nonconfigured memory address
- Register Set (REX) or MAT parity failure

The processor is designed with the concept that all software must enable the machine malfunction interrupt for maximum data integrity. Unlike other processors, this does not require that this interrupt ever be disabled. The processor resets each detected interrupt condition as it occurs.

While performing a machine malfunction interrupt PSW swap, the processor sets PSW bit 18 to allow error detection for the new PSW data fetched from memory. If the new PSW cannot be fetched correctly, the processor effectively stops by entering the console mode. This prevents a runaway situation in the event of a double fault.

If PSW bit 18 is zero, any noncorrectable memory data error is logged by the error logger. Accesses to memory using a nonconfigured memory address result in undefined data, with no error indication. No machine malfunction interrupt occurs for any of the reasons given above. A machine malfunction due to EPF is queued until PSW bit 18 is set by software, or until automatic shutdown occurs. The interrupt is not queued for any other reason.

10.2.1.8 Floating Point Underflow Interrupt Enable (FLU)

PSW bit 19 controls response of the processor to an arithmetic underflow resulting from a single or double precision floating point arithmetic operation.

If this bit is set when the underflow occurs, an arithmetic fault interrupt occurs, and the participating floating point registers remain unchanged.

If this bit is zero when the underflow occurs, the result of the operation is replaced by zero, and the condition code is set to 0100 (V flag only), as defined in the description of the specific floating point instruction.

10.2.1.9 Relocation/Protection Enable (R/P)

PSW bit 21 is used to enable and disable the relocation and protection programmed into the MAT. When this bit is set, relocation, protection and the MAT fault interrupt are enabled. When this bit is zero, relocation, protection and the MAT fault interrupt are disabled.

10.2.1.10 System Queue Service (SQS) Interrupt Enable (Q)

If bit 22 of the new PSW loaded by any of the instructions listed below is set, the state of the system queue is tested. If the system queue is not empty, an SQS interrupt occurs. If the system queue is empty, the next instruction is fetched and executed, according to the newly loaded PSW.

If bit 22 of the newly loaded PSW is zero, the SQS interrupt is disabled.

The following instructions test the state of the system queue:

MNEMONIC	MEANING
EPSR	Exchange Program Status Register
LDPS	Load Process State
LPSW	Load Program Status Word
LPSWR	Load Program Status Word Register

10.2.1.11 Protect Mode Enable (P)

When PSW bit 23 is set, the processor is in the protect mode. Any attempt by a program running in this mode to execute a privileged instruction causes an illegal instruction interrupt to occur. The processor does not attempt to execute the offending instruction. The BRK instruction is a privileged instruction.

When PSW bit 23 is zero, the processor is in privileged mode. A program running in privileged mode can execute any legal instruction within the constraints imposed by the system configuration and the state of PSW bit 13 (FLM).

10.2.1.12 Register Set Select Field (R)

Bits 24, 25, 26 and 27 of the current PSW select the active general register set. Although 16 different sets can be specified by using the four bits of this field, only eight sets of general registers are implemented in this processor. The implemented sets are numbered 0, 1, 2, 3, 4, 5, 6 and 15.

Set 0 is automatically selected by the processor in handling an I/O interrupt. Registers 0 through 4 of that set are used to maintain information pertaining to the I/O interrupt request. Therefore, set 0 should not be used for general-purpose processing. This set can, however, be used for processing internal interrupts, which use registers 11 through 15 of the selected set to maintain information pertaining to the interrupt.

Sets 1, 2, 3, 4, 5, 6 and 15 can be allocated according to processing needs without special consideration. Sets 7 through 14 are not implemented. (An attempt to select a set that is not implemented can result in the selection of any set without any special indication of the error.)

When a new PSW is loaded, the specified register set becomes the active set for the next instruction executed.

24	PSW BITS			SELECTED REGISTER SET
	25	26	27	
0	0	0	0	0
0	0	0	1	1
0	0	1	0	2
0	0	1	1	3
0	1	0	0	4
0	1	0	1	5
0	1	1	0	6
1	1	1	1	15

10.2.1.13 Condition Code (C, V, G, L)

PSW bits 28:31 contain the condition code. As part of the execution of certain instructions, the state of the condition code can be updated to reflect the nature of the result. Not all instructions affect the condition code.

For most interrupts, bits 28:31 of the new PSW are simply copied to the condition code. For immediate interrupts, the least significant four bits of the status byte for the interrupting device are copied to the condition code after the new PSW has been loaded. No restrictions are imposed on the condition code field of a new PSW contained in a memory location or register. Any condition code value can be specified.

The condition code of the current PSW can be tested by the conditional branch instructions described in Chapter 4.

10.2.2 Program Status Word (PSW) Location Counter (LOC)

PSW bits 32:63 comprise the LOC, which contains the address following the instruction currently being executed by the processor. When the current instruction is successfully completed, the value contained in the LOC is used, and the instruction at the resulting address is fetched.

An instruction which results in a branch being taken causes the contents of the LOC to be replaced with the effective branch address; i.e., with the address of the instruction to which control is to be transferred. The instruction at the new address is the next instruction to be fetched and executed.

When an interrupt occurs, the entire PSW, bits 0:63, is replaced. If bit 16 of the new PSW (the wait bit) is set, the instruction indicated by the new contents of the LOC is not fetched. Manual intervention is required to cause the wait bit to be zero, and the instruction to be fetched and executed. If an interrupt causes the PSW with the wait bit set to be replaced by another new PSW that has the wait bit zero, the instruction indicated by the LOC of that new PSW is fetched and executed.

If an instruction has not been successfully completed when an interrupt PSW swap occurs, the 64-bit PSW in effect for the instruction being executed at the time of the interrupt is saved before the interrupt handler is entered. The LOC in the saved PSW points to the instruction being executed at the time the interrupt occurred. If the interrupt occurs after the successful completion of one instruction and before beginning another, the LOC in the saved PSW points to the next instruction to be executed.

See Section 10.5 for an explanation of old, current and new PSWs, and of the use of these PSWs by the processor in scheduling ISRs.

10.2.3 Reserved Memory Locations

Physical memory locations X'000000' through X'0002CF' are reserved memory locations. For systems with expanded I/O interrupt service pointer tables (ISPTs), physical memory locations X'0002D0' through X'0004CF' or X'0002D0' through X'0008CF' are also reserved memory locations. These locations contain assorted information used in servicing interrupts, as shown in Table 10-2. Use of data in these locations as the result of an interrupt is detailed in the section describing the interrupt.

TABLE 10-2 RESERVED MEMORY LOCATIONS

LOCATION	MEANING
X'000000' - X'00001F'	Reserved; must be zero
X'000020' - X'000027'	Machine malfunction interrupt old PSW
X'000028' - X'00002B'	Used by console service microcode
X'00002C' - X'00002F'	LM effective address word
X'000030' - X'000037'	Illegal instruction interrupt new PSW
X'000038' - X'00003F'	Machine malfunction interrupt new PSW
X'000040' - X'000043'	Machine malfunction status word
X'000044' - X'000047'	Machine malfunction virtual (program) address word
X'000048' - X'00004F'	Arithmetic fault interrupt new PSW
X'000050' - X'00007F'	Bootstrap loader and device definition table
X'000080' - X'000083'	System queue pointer
X'000084' - X'000087'	INT/ Power fail save area pointer
X'000088' - X'00008F'	System queue service interrupt new PSW
X'000090' - X'000097'	Relocation/protection (MAT fault) new PSW
X'000098' - X'00009B'	SVC new PSW status word
X'00009C' - X'0000BB'	SVC new PSW LOC values
X'0000BC' - X'0000C7'	Reserved - must be zero

TABLE 10-2 RESERVED MEMORY LOCATIONS (Continued)

LOCATION	MEANING
X'0000C8' - X'0000CF'	Data format fault new PSW
X'0000D0' - X'0002CF'	Interrupt service pointer table
X'0002D0' - X'0004CF'	Expanded interrupt service pointer table
X'0004D0' - X'0008CF'	Expanded interrupt service pointer table

10.3 INTERRUPT TIMING AND PRIORITY

The following sections discuss interrupt timing and priority features.

10.3.1 Maskable and Nonmaskable Interrupts

Maskable interrupt conditions are controlled by bits in the PSW. When a request to interrupt due to a maskable condition occurs, the corresponding control bit in the PSW is examined. If the control bit indicates that the interrupt is enabled, an interrupt is taken and control is transferred to the appropriate service routine. The section describing each interrupt provides details about the control bit(s), how the interrupt is enabled or disabled, and the effects of enabling or disabling an interrupt.

Nonmaskable interrupts are those which have no corresponding control bits in the PSW. Examples of nonmaskable interrupts are SVC, SINT Illegal instruction, and console attention. Sections describing each interrupt provide further details.

Figure 10-2 shows the various maskable and nonmaskable interrupts.

622 9

NOTES

- a) NUMBERS IN CIRCLES INDICATE THE PRIORITY OF INTERRUPTS. 1 REPRESENTS THE HIGHEST PRIORITY.
- b) FAULTS ABORT THE CURRENT INSTRUCTION. THE OLD PSW POINTS TO THE FAULTING INSTRUCTION. OTHER INTERRUPTS ARE RECOGNIZED AT THE END OF THE CURRENT INSTRUCTION AND OLD PSW POINTS TO THE FOLLOWING INSTRUCTION.

- (c) SYNCHRONOUS INTERRUPTS ARE RECOGNIZED AS THEY OCCUR. ASYNCHRONOUS INTERRUPTS ARE RECOGNIZED BETWEEN THE COMPLETION OF CURRENT INSTRUCTION AND THE INITIATION OF THE NEXT INSTRUCTION.
- (d) SOS MAY OCCUR ONLY AS PART OF THE LPSW, LPSWR, EPSR, AND LDPS INSTRUCTIONS.

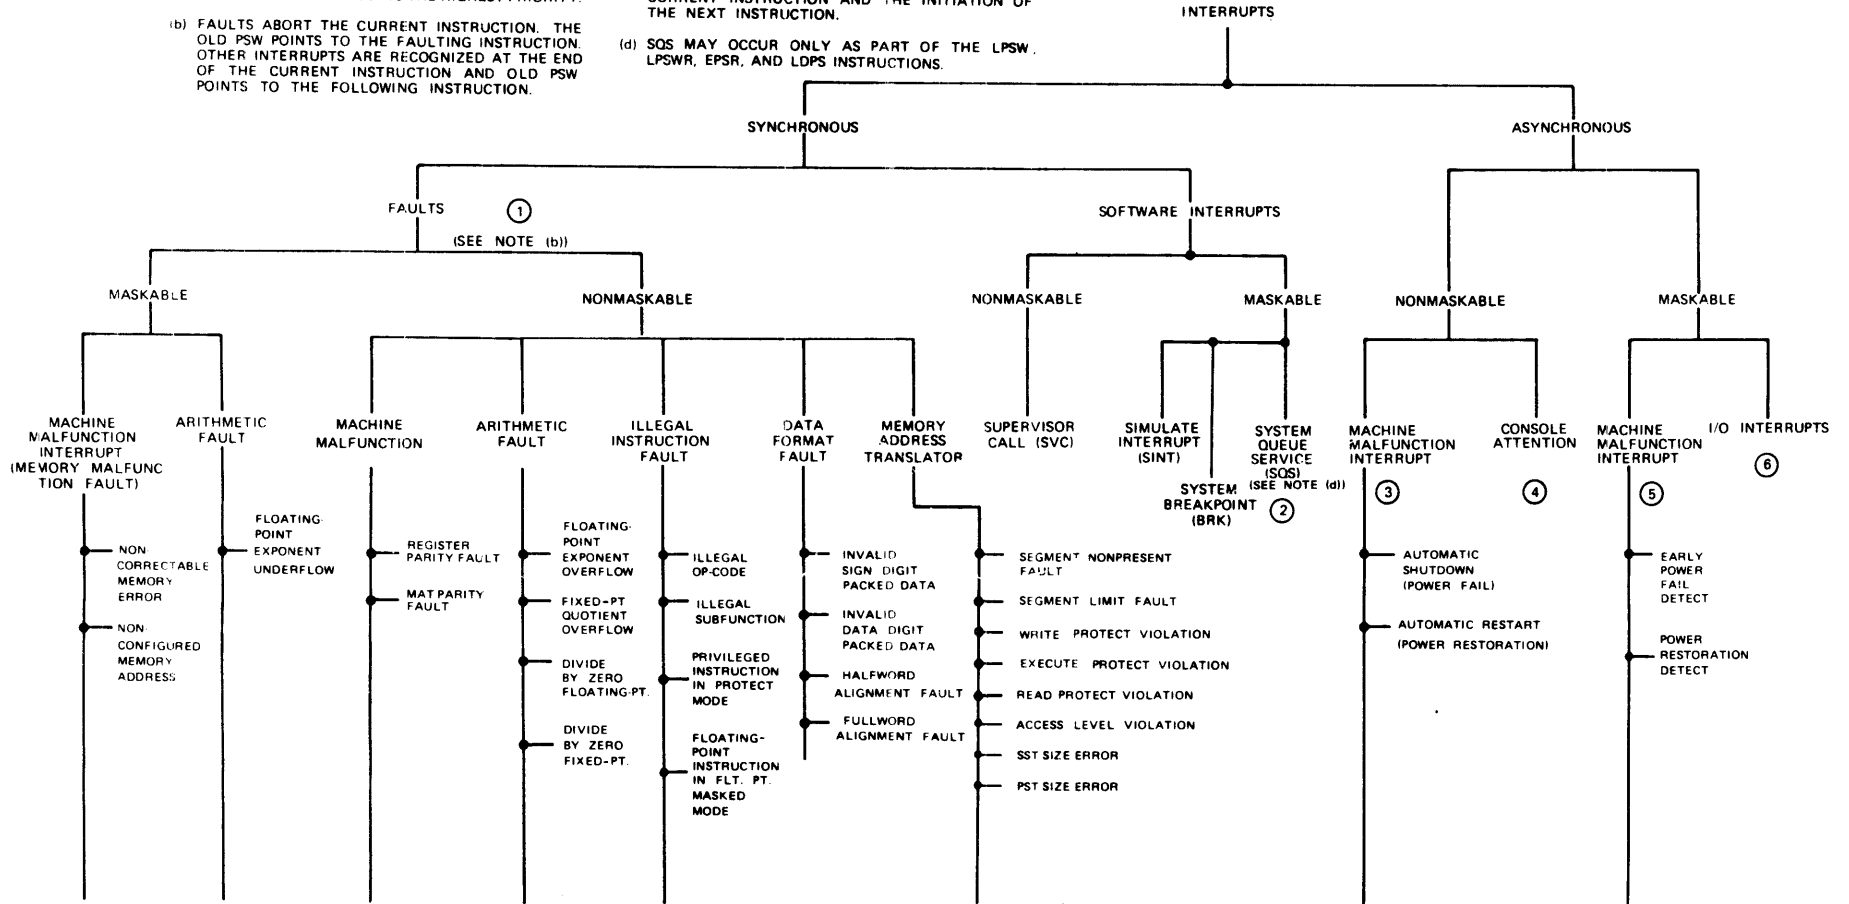


Figure 10-2 Schematic Diagram of Interrupt System Architecture

10.3.2 Interrupt Timing

Asynchronous interrupts are normally permitted to occur only after execution of an instruction has been completed, and before execution of the next instruction begins. However, asynchronous interrupts are permitted to occur at the end of any iteration while an interruptible instruction is being executed.

A synchronous interrupt is permitted to occur at the time the condition causing the interrupt is detected. The SQS interrupt, which occurs at some indefinite time following addition of an entry to the system queue, is called a deferred synchronous interrupt. A synchronous interrupt due to a fault causes the offending instruction to be aborted with no modification of the contents of registers or memory locations generally resulting from execution of that instruction. Fixed and floating point Load/Store Multiple and Store Double Precision are exceptions to this rule. In the case of an interruptible instruction, the current iteration of the instruction is aborted by such an interrupt without modification of the contents of registers or memory as a result of the faulted iteration.

For all interrupts, the old PSW LOC presented to the interrupt handler points to the next logically executed instruction in the interrupted program. If the interrupt is caused by a fault, the instruction causing the fault was not completed and is logically the next instruction to be executed. The old PSW LOC presented to the fault ISR, therefore, always points to the instruction that caused the fault.

Multiple memory accesses are required for the manipulation of a circular list structure using the ATL, ABL, RTL or RBL instruction. For each of these instructions, the list header is not updated until the body of the list has been successfully accessed. For the RTL and RBL instructions, no registers are modified unless the list element has been successfully accessed and the list header has been successfully updated.

10.3.3 Interrupt Precedence

Considering the instant of instruction fetch request as the time of reference, interrupts have the following precedence (highest to lowest):

INTERRUPT PRECEDENCE TABLE

Synchronous Interrupts	{	Fault interrupts System queue service
Asynchronous Interrupts	{	Primary power fail/restore Console attention Early power fail I/O interrupts

Fault interrupts are caused by various conditions that have the following logical precedence in descending priority order.

- Register parity error
- MAT parity error
- Relocation/protection fault on an instruction fetch
- Machine malfunction fault due to memory malfunction on an instruction fetch
- Illegal instruction fault
- Illegal subfunction fault
- Data format fault due to alignment error on a data read/write operation
- Relocation/protection fault on a data read/write operation
- Machine malfunction fault due to memory malfunction on a data read/write operation
- Data format fault for other than boundary alignment error
- Arithmetic fault

For a memory malfunction, a nonconfigured memory address fault takes precedence over a noncorrectable memory data fault.

Since any fault interrupt causes execution of an instruction to be aborted at the point of the fault interrupt condition, no more than one fault interrupt condition can occur at a time. However, other interrupts in the synchronous and asynchronous interrupt classes given in the preceding interrupt precedence table can occur simultaneously. In such a case, the order given in the list above governs the servicing sequence for the interrupts.

10.3.4 Interruptible Instructions

For any interruptible instruction, execution consists of the following phases: instruction fetch, instruction decode, an iterative loop and termination. An interrupt during any phase of an interruptible instruction does not affect the operation of the instruction. It can simply be reexecuted once the interrupt has been serviced. An interrupt during the iterative phase of the instruction causes the processor to resume the iterative phase when the instruction is reexecuted, as though the interrupt never occurred. If the interrupt was caused by a fault, the iteration that resulted in the interrupt is repeated when the instruction is reexecuted.

To abort an interruptible instruction when it is interrupted, PSW bit 14 must be forced to zero before any subsequent interruptible instruction is attempted.

CAUTION

SOFTWARE MUST NEVER SET PSW BIT 14 UNLESS RESUMING EXECUTION OF THE INTERRUPTIBLE INSTRUCTION THAT CAUSED BIT 14 OF THE PSW TO BE SET. RESUMPTION OF ANY INTERRUPTIBLE INSTRUCTION MUST NEVER BE ATTEMPTED IF THE CONTENTS OF THE SCRATCHPAD REGISTERS ARE NOT KNOWN TO HAVE BEEN PRESERVED BETWEEN INSTRUCTION INTERRUPTION AND RESUMPTION.

10.4 PROCESSOR MODES

At any given time, the processor can be in the console mode or run mode. The single step mode provides a means for alternating between the console and run modes. Wait and run states only have meaning for the run mode.

10.4.1 Console Mode

While the processor is dedicated to communicating with the system console terminal, it is said to be in the console mode. In this mode, program execution is suspended so that the user can examine and modify the data contained in certain registers and memory locations.

Appendix F provides a flowchart for the console service routine. The console mode can be entered in any of the ways listed below.

- The BRK instruction is executed by a running program when PSW bit 23 is zero.
- Execution of an instruction is completed while in the single step mode.
- The HALT/RUN or SINGLE switch is depressed momentarily while the processor is in the run mode.
- Following a system initialization sequence, back-up power to memory is found not to have been maintained within regulation, and the loader storage unit (LSU) is not enabled when the sequence is complete.

- Following a system initialization sequence, if back-up power to memory was maintained within regulation but the LSU is not enabled, then the contents of physical memory location X'000028' indicate that the processor was in the console mode when system initialization occurred.
- An attempt to fetch a machine malfunction interrupt new PSW results in a noncorrectable memory error. In this case, the console service routine will display the PSW with the CSF bit set and LOC at the time of the failure.

Note that system initialization occurs when the power supply detects that AC line voltage is failing; when the initialize (INIT) switch on the console is momentarily depressed; or when the key-operated LOCK/ON/STANDBY switch is moved to the STANDBY position. The initialization sequence completes when power is restored to the processor. System initialization resets all pending interrupts for the system console and other I/O devices in the system. Direct memory access (DMA) operations are also terminated.

While the processor is in the console mode, interrupt conditions are not handled in the same manner as they are when detected during execution of a program.

Interrupt requests for the system console terminal and all other I/O devices remain queued until the run mode is entered. DMA operations are not affected by changing processor modes.

PSW bit 16 is always forced to zero before the run mode is entered from the console mode.

Fault conditions caused by memory accesses while in the console mode are reset when they occur and do not cause interrupts when the run mode is entered. If a fault condition occurs while attempting to modify a memory location, that location cannot be changed. If a fault occurs while attempting to examine a memory location, the console service routine is aborted and restarted.

System initialization while in the console mode results in automatic shutdown with no machine malfunction interrupt due to power failure.

10.4.2 Run Mode

When the processor is not dedicated to communicating with the system console terminal, it is in the run mode. In this mode, program execution is controlled by the contents of the 64-bit PSW. While the processor is in the run mode, it can be in either the wait state (PSW bit 16 is set) or the run state (PSW bit 16 is zero). In the run state, the processor performs a repetitive fetch instruction/execute instruction/fetch next instruction sequence. In the wait state, this sequence is suspended.

The run mode can be entered in any of the following ways:

- The less than (<) prompt character is entered from the system console terminal when the processor is in the console mode.
- The HALT/RUN switch is depressed momentarily while the processor is in the console mode.
- The LSU is installed and enabled when a system initialization sequence is completed. In this case, the program loaded from the LSU is given control of the processor.
- The greater than (>) single step character is entered from the system console terminal when the processor is in the console mode. This causes the instruction to be executed in single step mode, regardless of the position of the SINGLE switch.

Interrupt conditions cannot cause the processor to enter the run mode from the console mode, with the following two exceptions:

- An initialization sequence performed while the processor is in the console mode causes a program to be loaded from the enabled LSU. Control of the processor is given to that program.
- The HALT/RUN switch is depressed momentarily while the processor is in the console mode.

10.4.3 Single Step Mode

When the SINGLE switch is in the SINGLE position, the processor is in the single step mode. In this mode, whenever execution of an instruction is completed, the processor leaves the run mode and enters the console mode. Manual intervention is normally required to execute the next instruction.

Interrupts are handled according to the methods detailed in the previous sections. If the processor is in the single step mode and the run state when an interrupt request occurs, the processor completes the current instruction (or iteration) and then performs the interrupt PSW swap. The first instruction of the ISR is not executed.

If system initialization occurs while in the single step mode, any instruction in progress (or the current iteration of an interruptible instruction) completes. When the initialization sequence is complete, a maximum of one instruction is executed before the processor again enters the console mode.

If the processor is in the run state when the SINGLE switch is placed in the SINGLE position, the console mode is entered. Note that in the single step mode PSW bit 16 is always forced to zero before entering the run mode to fetch a user instruction.

NOTE

If interrupts are enabled at the system control terminal interface by software, entering the console mode causes interrupts to be queued from device X'011' (the write side). Depression of any key at the console may cause an interrupt to be queued from device X'010' (the read side).

10.5 STATUS SWITCHING

The PSW that is loaded in the processor at any given time is called the current PSW. The register set selected by this PSW, the data contained in the general, floating point or scratchpad registers accessible by the user program, and the machine status defined by the PSW collectively constitute the process state. If the status word or both the LOC and status word are changed, a status switch has occurred. A status switch can be caused explicitly by executing a status switching instruction or may be forced to occur by an interrupt. When the value of the PSW that was current at the time of a status switch is saved, that value is called the old PSW.

The scheduling of ISRs is based upon the concepts of old PSW, current PSW and new PSW. When an interrupt occurs, the following status switch takes place: the current PSW becomes the old PSW; the new PSW defined for the interrupt is loaded and becomes the current PSW.

For a status switch resulting from an interrupt, the old PSW is stored in dedicated registers of the set specified by the new PSW defined for the interrupt. The machine malfunction interrupt is the exception to this rule; for this interrupt, the old PSW is stored in dedicated memory locations.

For meaningful processor response to multiple interrupts, it is important that the new PSW defined for a particular interrupt class does not enable interrupts of the same class.

The various interrupts that can occur and the response of the processor to each interrupt are described in the following sections.

10.5.1 Illegal Instruction Interrupt

The illegal instruction interrupt occurs if an attempt is made to execute an instruction whose operation code is not one of those permitted by the system. This interrupt can occur for any of the reasons listed below.

- The operation code is undefined for the system.
- The operation code has several possible subfunction specifications, and the subfunction specified is undefined.
- The instruction is a privileged instruction, and PSW bit 23 is set.
- The instruction is a floating point instruction, and PSW bit 13 is set.

The illegal instruction interrupt cannot be disabled. No attempt is made by the processor to execute an illegal instruction.

When an illegal instruction interrupt occurs, the following actions are taken:

1. The current PSW is stored in registers 14 and 15 of the set selected by the illegal instruction interrupt new PSW, found in memory at physical address X'000030'.
2. The illegal instruction interrupt new PSW becomes the current PSW.

The old PSW LOC presented to the ISR in register 15 points to the illegal instruction.

10.5.2 Data Format Fault Interrupt

The data format fault interrupt occurs if the required halfword or fullword alignments are violated for memory accesses, or if it is otherwise determined that data is not properly aligned to the specified fields. The data format fault interrupt cannot be disabled. When a data format fault interrupt occurs, the following actions are taken:

1. The current PSW is stored in registers 14 and 15 of the set selected by the data format fault new PSW, found in memory at physical address X'0000C8'.

2. Register 13 of the selected set is loaded with a code to indicate the reason for the interrupt, as shown in the following list:

CODE	REASON FOR INTERRUPT
0	Reserved code
1	Reserved code
2	Invalid sign digit, packed data
3	Invalid data digit, packed data
4	Reserved code
5	Reserved code
6	Fullword or halfword alignment fault

3. If the interrupt was caused by a halfword or fullword alignment fault, register 12 of the selected set is loaded with the nonaligned virtual address causing the fault.
4. The data format fault interrupt new PSW becomes the current PSW.

The old PSW LOC presented to the ISR in register 15 points to the instruction being executed when the fault occurred. A data format fault causes the current instruction, or the current iteration of an interruptible instruction, to be aborted immediately.

10.5.2.1 Alignment Faults

An attempt to fetch a fullword of data from memory, or to write a fullword of data to memory, using a virtual address (VA) that does not have zeros as its two least significant bits causes a fullword alignment fault.

An attempt to read a halfword of data from memory using a program address that does not have zero as its least significant bit causes a halfword alignment fault. The processor does not distinguish between fullword and halfword alignment faults.

If an alignment fault occurs while attempting to write to memory, the fullword or halfword at the next lower aligned address can be modified.

10.5.2.2 Invalid Digit Faults

If an invalid sign or data digit is encountered while processing numeric string data, it is presumed that the data is not aligned to the specified fields. Additional information can be found in the description of the instruction used to process the numeric string.

10.5.3 Relocation/Protection (MAT) Fault Interrupt

This fault interrupt occurs if an executing program violates any of the relocation and protection conditions programmed into the MAT. MAT error checking and the MAT fault interrupt are enabled when PSW bit 21 is set. MAT faults are not queued.

When a MAT fault interrupt occurs, the following actions are taken:

1. The current PSW is stored in registers 14 and 15 of the set selected by the MAT fault interrupt new PSW, found in memory at physical address X'000090'.
2. Register 13 of the selected set is loaded with a code to indicate the reason for the interrupt. This code is copied from the MAT status register while simultaneously resetting the fault.

CODE	REASON FOR INTERRUPT
0	Reserved code
1	Execute protect violation
2	Write protect violation
3	Read protect violation
4	Access level fault
5	Segment limit fault
6	Nonpresent segment
7	Shared segment table (SST) size exceeded
8	Process segment table (PST) size exceeded

3. Register 12 of the selected set is loaded with the program address that caused the fault.
4. If the fault occurred on a data fetch while attempting to load the general registers using the Load Multiple (LM) instruction, register 11 of the selected set is loaded with the effective second operand address calculated at the start of the LM instruction.
5. The MAT fault interrupt new PSW becomes the current PSW. The old PSW LOC presented to the ISR in register 15 points to the instruction being executed when the fault occurred. Further information on memory management may be found in Chapter 11.

10.5.4 Machine Malfunction Interrupt

The machine malfunction interrupt occurs when any one of the following conditions is detected:

- Register or MAT parity failure

- EPF
- Power restore
- Noncorrectable memory data error
- Nonconfigured memory address

Detection of the listed conditions and the machine malfunction interrupt are enabled when PSW bit 18 is set. An EPF interrupt is queued until PPF occurs if PSW bit 18 is zero. All other malfunction conditions are ignored, and the interrupts are lost.

When a machine malfunction interrupt occurs, the following actions are taken:

1. The current PSW is stored in memory beginning at physical address X'000020'.
2. The machine malfunction status word at physical address X'000040' is loaded with a code to indicate the reason for the interrupt. Only one bit is set in this code.

BITS	MNEMONIC	REASON FOR INTERRUPT
0	PF	Power failure
1	PR	Power restoration
2	NCD	Noncorrectable memory error during data fetch
3	NCI	Noncorrectable memory error during instruction fetch
4	NCA	Noncorrectable memory error during auto driver channel operation
5	NVD	Nonconfigured memory address during data fetch
6	NVI	Nonconfigured memory address during instruction fetch
7	NVA	Nonconfigured memory address during auto driver channel operation
8	RPF	Register parity failure
9	MPF	MAT parity failure

3. If the interrupt was caused by a noncorrectable memory error or nonconfigured memory address, the VA used for the memory access is stored in the machine malfunction VA word at physical address X'000044'. Otherwise, the contents of this word are undefined.

4. If the interrupt was caused by a noncorrectable memory error or nonconfigured memory address, and the fault occurred on a data fetch while attempting to load the general registers using the LM instruction, the effective second operand address calculated at the start of that instruction is stored in the LM effective address word at physical address X'00002C'. Otherwise, the contents of this word are undefined.
5. The machine malfunction interrupt new PSW, found at physical address X'000038', becomes the new PSW.

If the interrupt was caused by executing an instruction, the old PSW LOC presented to the ISR points to the offending instruction. Otherwise, the old PSW LOC presented to the ISR points to the instruction to be executed once the interrupt has been serviced.

If the interrupt was caused by executing the LM instruction, bits 2 and 5 of the machine malfunction status word can be used to determine if any registers were modified before the interrupt occurred. If the old PSW LOC points to an LM instruction, and if bits 2 and 5 of the machine malfunction status word (MMSW) are both zero, no registers were modified. If bit 2 or bit 5 of the machine malfunction status word is set, then:

- If the data stored at physical addresses X'000044' and X'00002C' are equal to one another, no registers were modified by the instruction before the fault occurred.
- If the data stored at physical addresses X'000044' and X'00002C' are not equal to one another, at least one register was modified by the instruction before the fault occurred. The number of registers modified may be determined by taking the difference of the data stored at physical addresses X'000044' and X'00002C' and dividing the result by four.

1322-3




Figure 10-3 Machine Malfunction Status Word

10.5.4.1 Early Power Fail (EPF) Detect and Automatic Shutdown

EPF detect occurs when the PPF sensor detects a low voltage, the power switch is turned from the ON to STANDBY position, or the INIT switch is depressed.

At the end of execution of the current instruction or the current iteration of the current interruptible instruction, a machine malfunction interrupt is taken if PSW bit 18 is set.

Following EPF detect, software has one millisecond before the automatic shutdown procedure of the processor takes control as a result of PPF. During this procedure, the following actions occur:

1. The fullword power fail save area pointer is fetched from location X'000084'.
2. The following information is saved by firmware in the power fail save area:

DATA	OFFSET IN SAVE AREA (IN BYTES)
Current PSW	0-7
The eight general register sets (in order, 0 through F)	8-519
Interruptible instruction state (scratchpad registers)	520-583
Optional floating point registers, single and double	584-679

3. The processor waits for power restore.

NOTES

1. If the pointer found in location X'000084' does not specify a save area aligned to a fullword boundary, the processor forces correct alignment by replacing the two least significant bits of the pointer with zeros.
2. The floating point masked mode bit in the PSW has no effect on the saving of the floating point registers.
3. The IIP bit in the PSW has no effect on the saving of the scratchpad registers.

10.5.4.2 Power Restore

When power restore occurs, a simple go/no go self-test of various internal buses and registers is performed. If the back-up supply voltages to memory were not maintained within margins between shutdown and power restore, the first 512kb of memory are filled with a data pattern to prevent spurious noncorrectable memory error indications, and the general registers, scratchpad registers and floating point registers are loaded with predetermined data.

The first 512kb of memory are then tested to see if data can be held. This test does not modify the data contained in memory. Failure of self-test or the memory test causes that test to loop as long as the failure persists. During the test, the processor is responsive only to a PPF, which results in an automatic shutdown, and the FAULT lamp on the console switch panel is on.

When memory testing is complete, the FAULT lamp is turned off, and the state of the LSU is tested. In all cases, bit 1 of the machine malfunction status word at physical address X'000040' is set to indicate power restore.

10.5.4.2.1 If the Loader Storage Unit (LSU) Is Disabled

If the back-up voltages to memory were not maintained within margins between shutdown and power restore, then memory is assumed not to contain valid data. In this case, a PSW status of '00008000' (wait bit only) and LOC of '00FFFFFFE' are loaded and displayed on the system console terminal. Manual intervention is required to restart the processor. The memory voltage failure (MVF) indication is reset in this case. MVF is discussed in Section 10.5.4.2.2.

If the back-up voltages to memory were maintained, the data saved in the power fail save area by the automatic shutdown procedure is reloaded.

If the data in memory at physical address X'000028' indicates that the processor was in console mode when power failed, the reloaded PSW is displayed, and communication with the system console terminal resumes.

If the processor was not in console mode when power failed, bit 18 of the reloaded PSW is tested. If the bit is set, a machine malfunction interrupt occurs.

If bit 18 of the reloaded PSW is zero, program execution is resumed using the reloaded PSW. Note that the state of the wait bit (bit 16) of the PSW is tested before executing any instruction.

NOTE

Data in the MAT and selector channel (SELCH) control registers is volatile and must be considered invalid following any power fail/restore sequence.

10.5.4.2.2 If the Loader Storage Unit (LSU) Is Enabled

After the FAULT lamp is turned off, the program in the LSU is loaded, and control is transferred to it using the PSW specified in the program. If the memory start address is greater than the memory end address specified for the LSU program, the program is not loaded, and the console mode is entered.

An MVF indication is available to the processor if memory voltages are not maintained within margins between shutdown and power restore. MVF is reset when the console mode is entered or the Reset Memory Voltage Failure (RMVF) instruction is executed.

If MVF is indicated following power restore, it is assumed that memory does not contain an executable program. The MVF indication is retained until reset as described above, even if multiple shutdown/power restore sequences occur. Software loaded via the optional LSU should execute the RMVF instruction once the load is complete and all interrupt new PSWs have been established. Proper use of the RMVF instruction prevents a potential runaway condition in the event of multiple power failures.

10.5.4.3 Noncorrectable Memory Error

During write operations to memory, an error correction code (ECC) is generated. This code enables the memory system to correct any single bit error detected on a subsequent read operation in each halfword of memory. If the operation is only a byte or halfword write to memory, the memory system reads and updates the ECC for the halfword of memory that contains the byte or halfword that is being written.

Each time data is read from memory, the ECC is recreated and compared to the code generated when data was last written to any part of the halfword memory location. If a data error is detected and the error is a single bit error, it is corrected transparent to the processor. If, however, a multiple bit error is detected, a memory malfunction fault is generated, since multiple bit errors cannot be corrected.

Note that data with three or more bits in error may not result in a fault. Detection of any error causes a bit to be set in the error logger for subsequent readouts using the REL instruction.

If PSW bit 18 is zero when the error occurs, the error is ignored but is logged in the error logger.

If PSW bit 18 is set, occurrence of a noncorrectable memory error causes the current instruction (or the current iteration of an interruptible instruction) to be aborted immediately, and a machine malfunction interrupt occurs. Bit 2, 3 or 4 of the machine malfunction status word at physical address X'000040' is set to indicate the reason for the interrupt. The program address used for the memory access is stored in the machine malfunction address word at physical address X'000044'.

If the error occurs on a data fetch while attempting to load the general registers using the LM instruction, the effective second operand address calculated at the start of the LM instruction is stored in the LM effective address word at physical address X'00002C'. This data allows the instruction to be simulated in the event that the specified index registers were modified.

If the error occurs while fetching an instruction, the old PSW LOC, presented to the ISR, points to the first halfword of the instruction being fetched.

If the error occurs during an auto driver channel operation, registers 0 and 1 of the set indicated by the old PSW presented to the ISR contain the PSW for the instruction interrupted by the I/O interrupt that activated the channel. Register 4 of the set indicated contains the address of the channel command block (CCB) that was being executed when the error occurred.

10.5.4.4 Nonconfigured Memory Address

The processor tests the physical address used for each memory access, if PSW bit 18 is set. When access to memory assigned to a memory controller not physically in the system is attempted, a machine malfunction interrupt occurs. The current instruction (or the current iteration of an interruptible instruction) is immediately aborted. Bit 5, 6 or 7 of the machine malfunction status word at physical address X'000040' is set to indicate the reason for the interrupt. The program address used for the memory access is stored in the machine malfunction address word at physical address X'000044'.

If the error occurs on a data fetch while attempting to load the general registers using the LM instruction, the effective second operand address calculated at the start of the LM instruction is stored in the LM effective address word at physical address X'00002C'. This data allows the instruction to be simulated in the event specified index registers were modified.

If the error occurs while fetching an instruction, the old PSW LOC, presented to the ISR, points to the first halfword of the instruction being fetched.

If the error occurs during an auto driver channel operation, registers 0 and 1 of the set indicated by the old PSW presented to the ISR contain the PSW for the instruction interrupted by the I/O interrupt that activated the channel. Register 4 of the indicated set contains the address of the CCB that was being executed when the error occurred.

10.5.5 Input/Output (I/O) Device Interrupts

The following sections detail I/O device interrupts.

10.5.5.1 Priority Levels

Interrupt requests from I/O devices can occur on only one priority level. Acknowledgement of interrupt requests is enabled by PSW bit 17, as shown below.

PSW BIT 17	MEANING
0	Disabled
1	I/O interrupts enabled

A unique register set is selected for I/O interrupt requests acknowledged on each priority level 0. For example, when an interrupt request is acknowledged at priority level 0, register set 0 is selected by the processor for handling the interrupt request. If the request results in entry to a software ISR, register set 0 is selected by the PSW in effect at the time the routine is entered, and information pertaining to the interrupt is contained in registers 0 to 3 or 0 to 4 of that set.

Enabling of interrupts is dependent upon the state of PSW bit 17. When an interrupt request occurs but is not acknowledged by the processor, the request remains queued until one of the following occurs:

- The interrupt request is acknowledged by the processor when enabled by the current PSW.
- The interrupt request is programmed reset by the software.
- System initialization occurs.

When the processor acknowledges an I/O interrupt request, the result can be either an auto driver channel operation or an immediate interrupt. In either case, register set 0 is used in processing the interrupt.

For further information on programming a device interrupt request reset, see the programming manual for the specific device. This feature is not available for all I/O devices.

10.5.5.2 Immediate Interrupt - Auto Driver Channel Operation

An interrupt request by an I/O device is acknowledged only when interrupts are enabled as defined by the state of PSW bit 17.

The processor recognizes I/O interrupts between the execution of instructions or at the end of an iteration of an interruptible instruction. When an I/O interrupt is recognized, the following actions occur:

1. The current PSW is saved in registers 0 and 1 of set 0 (PSW bits 0:31 are saved in register 0 and bits 32:63 in register 1).
2. The PSW status word is loaded with the value Y'00002800'. This status enables machine malfunction interrupts. Also note that the MAT is disabled.
3. The I/O interrupt request is acknowledged and reset. The address of the interrupting device is placed in register 2 of set 0. The status byte from the interrupting device replaces the contents of register 3. The device number and status are placed in the least significant bit positions in the register; the most significant bits are forced to zero. The four least significant bits of the status of the interrupting device are placed in the condition code.
4. The device number is added twice to X'0000D0', the start of the ISPT, to obtain the address within the table that corresponds to the interrupting device. The contents of this halfword of memory are fetched and examined to see if the interrupt is to be treated as an immediate interrupt or as an auto driver channel operation. If bit 15 of the halfword is zero, an immediate interrupt is required. If bit 15 of the halfword is one (the halfword is odd), an auto driver channel operation is required. If the interrupt is an immediate interrupt, the value in the table becomes the LOC portion of the current PSW. If the interrupt is an auto driver channel operation, then the least significant bit of the halfword is replaced by zero and the resulting value is placed in register 4 of set 0. The auto driver channel is then activated.

10.5.6 Simulated Interrupt (SINT)

The SINT results from executing a SINT instruction when PSW bit 23 is zero. SINT is a privileged instruction and cannot be executed when PSW bit 23 is set.

Execution of the SINT instruction causes the processor to simulate acknowledgement of an enabled I/O interrupt request from an external device. The device address for the SINT is specified by the operands of the SINT instruction.

The state of PSW bit 17 is ignored by the SINT instruction. For purposes of the SINT, I/O interrupts are assumed to be enabled. No pending device interrupt request is actually acknowledged by the processor as a result of executing the SINT instruction. With the exception of the differences described here, the SINT request is handled as detailed in Section 10.5.5.

CAUTION

DUE TO THE FACT THAT THE SINT INSTRUCTION IGNORES THE STATE OF PSW BIT 17, IT SHOULD BE USED CAREFULLY BY PROGRAMS THAT RUN IN REGISTER SET 0. FOR EXAMPLE, IF A PROGRAM EXECUTING IN REGISTER SET 0 DISABLES INTERRUPTS, DATA IN THE REGISTERS OF SET 0 ARE NOT NORMALLY SUBJECT TO CHANGE AS A RESULT OF AN I/O INTERRUPT. HOWEVER, IF THE PROGRAM EXECUTING IN REGISTER SET 2 DOES A SINT, AN INTERRUPT OCCURS REGARDLESS OF THE STATE OF PSW BIT 17. IF AN I/O INTERRUPT REQUEST OCCURRED, IT WOULD BE HONORED, CAUSING REGISTERS 0, 1, 2 AND 3 (AND POSSIBLY 4) OF SET 0 TO BE OVERWRITTEN.

IF THESE REGISTERS ARE NOT STORED BEFORE THE SINT INSTRUCTION IS EXECUTED, DATA IN THE REGISTERS IS LOST, AND SYSTEM SOFTWARE COULD BE LEFT IN AN INDETERMINATE STATE.

The SINT is a software interrupt.

10.5.7 System Queue Service (SQS) Interrupt

When any of the instructions listed below are executed, as the instruction completes, bit 22 of the new PSW loaded by the instruction is tested. If the bit is zero, the SQS interrupt is disabled, and program execution continues according to the new PSW loaded.

MNEMONIC

MEANING

EPSR	Exchange Program Status Register
LDPS	Load Process State
LPSW	Load Program Status Word
LPSWR	Load Program Status Word Register

If bit 22 of the new PSW loaded by any of these instructions is set, the state of the system queue (the physical address of which is found at physical location X'000080') is tested. The system queue is assumed to be maintained according to the circular list format. The number used field is fetched from the list header. If this field contains zero, the system queue is assumed to be empty, and program execution continues according to the new PSW loaded.

If the number used field for the system queue is not zero when it is tested, the following actions are taken to cause an SQS interrupt.

1. The current PSW, which was loaded by execution of one of the listed instructions, is stored in registers 14 and 15 of the set selected by the SQS interrupt new PSW, found in memory at physical address X'000088'.
2. Register 13 of the selected set is loaded with the address of the system queue.
3. The SQS interrupt new PSW becomes the current PSW.

If the SQS interrupt occurs as a result of executing an EPSR instruction, the old PSW LOC presented to the ISR in register 15 points to the instruction following the EPSR instruction. If the interrupt occurs as a result of executing any of the other listed instructions, the old PSW LOC contains the value loaded by the instruction causing the interrupt.

Items can be added to the system queue while the SQS interrupt is enabled or disabled. The Add to Top of List (ATL) and Add to Bottom of List (ABL) instructions are normally used for this purpose. The fact that the items have been added to the system queue is recorded in the list header. Only when a new PSW is loaded that enables the SQS interrupt is the state of the queue tested and an interrupt allowed.

The system queue has a maximum size, as determined by the list header established by system software. If an attempt is made to add an item to the queue when it is already full, the data can be lost. This could result in system software being left in an indeterminate state.

Note that the address of the system queue contained in the system queue pointer must be aligned to a fullword boundary.

See Section 10.6 for a description of the EPSR, LDPS, LPSW and LPSWR instructions.

The SQS interrupt is a deferred synchronous software interrupt.

10.5.8 Supervisor Call (SVC) Interrupt

The SVC interrupt occurs when the SVC instruction is executed. This instruction and the resulting interrupt provide a means for any program to communicate with system software.

When the SVC instruction is executed, the processor takes the following actions:

1. The current PSW is saved in registers 14 and 15 of the set selected by the SVC interrupt new PSW, found in memory at physical address X'000098'.
2. Register 13 of the selected set is loaded with the effective second operand address calculated for the SVC instruction executed. This is normally the address of an SVC parameter block, aligned to a fullword boundary.
3. The SVC interrupt new PSW becomes the current PSW, with a new LOC value chosen from the ordered list of halfwords at physical location X'9C'.

The old PSW LOC presented to the ISR in register 15 points to the instruction following the SVC instruction.

The SVC interrupt is a software interrupt and cannot be disabled.

10.5.9 System Breakpoint Interrupt

A system breakpoint results if a BRK instruction is executed when PSW bit 23 is zero. BRK is a privileged instruction and cannot be executed when PSW bit 23 is set.

Execution of the BRK instruction causes the processor to enter the console mode. In this mode, the processor is dedicated to communication with the system console terminal. Various registers and memory locations can be examined or modified by the user from the system console terminal while in this mode.

When the BRK instruction is executed, no registers or memory locations are modified. The PSW status and LOC are not modified by the BRK instruction. The LOC, at entry to the console mode, points to the BRK instruction.

When the run mode is entered from the console mode, PSW bit 16 is forced to zero, so that an instruction is fetched and executed. If the run mode is entered immediately after a BRK instruction is executed, the same BRK instruction results in another system breakpoint.

The system breakpoint interrupt is a software interrupt.

10.5.10 Arithmetic Fault Interrupt

The arithmetic fault interrupt results from either a fixed point or a floating point arithmetic operation when the magnitude of the result is too large to be represented within the required number of bits. Division by zero is a special case and always results in an arithmetic fault interrupt. Interrupts for any of these reasons cannot be disabled.

Floating point underflow occurs when the normalized result of a floating point load, conversion or other arithmetic operation is not zero, but is so small that it cannot be represented within the floating point number system defined for the processor.

If PSW bit 19 is zero when floating point underflow occurs, no arithmetic fault interrupt results. In this case, the result of the operation is set to true zero. This means that every bit of the result is forced to zero as the result is copied to its destination. If PSW bit 19 is set when floating point underflow occurs, an arithmetic fault interrupt does occur.

When an arithmetic fault interrupt occurs, the following actions are taken:

1. The instruction causing the interrupt is aborted before the data in any register or memory location is modified.
2. The current PSW is stored in registers 14 and 15 of the set selected by the arithmetic fault interrupt new PSW, found in memory at physical address X'000048'.
3. Register 13 of the selected set is loaded with a code to indicate the reason for the interrupt.

CODE	REASON FOR INTERRUPT
0	Fixed point division by zero
1	Fixed point quotient overflow
2	Floating point division by zero
3	Floating point exponent underflow
4	Floating point exponent overflow

4. Register 12 of the selected set is loaded with the address of the instruction following the instruction causing the interrupt.
5. The arithmetic fault interrupt new PSW becomes the current PSW.

The old PSW LOC presented to the ISR in register 15 points to the instruction that caused the interrupt.

10.6 STATUS SWITCHING INSTRUCTIONS

Status switching instructions provide for software control of the system's interrupt structure. They also allow user-level programs to communicate efficiently with control software. All status switching instructions, except the SVC instruction, are privileged operations; therefore, all interrupt handling routines must run in the supervisor mode.

The status switching instructions described in this section are:

LPSW	Load Program Status Word
LPSWR	Load Program Status Word Register
EPSR	Exchange Program Status Register
SINT	Simulate Interrupt
SVC	Supervisor Call
BRK	System Breakpoint
PSF	Privileged System Function

10.6.1 Load Program Status Word (LPSW)

<u>Assembler Notation</u>	<u>Opcode</u>	<u>Format</u>
LPSW D2(X2)	C2	RX1,RX2
LPSW A1(FX2,SX2)	C2	RX3

Operation:

The 64-bit second operand replaces the current PSW.

Condition Code:

Determined by the new PSW (bits 28:31).

Programming Notes:

The R1 field of this instruction must be zero.

The second operand must be aligned to a fullword boundary.

This instruction is a privileged operation.

This instruction may be used to change register sets. The new set becomes active for execution of the next instruction.

If bit 22 of the new PSW is set, the state of the system queue is tested. If the queue is not empty, an SQS interrupt occurs. In this case, the newly loaded PSW is saved as the old PSW when the SQS interrupt occurs.

10.6.2 Load Program Status Word Register (LPSWR)

<u>Assembler Notation</u>	<u>Opcode</u>	<u>Format</u>
LPSWR R2	18	RR

Operation:

The contents of the register specified by R2 replace bits 0:31 of the current PSW. The contents of the register specified by R2+1 replace bits 32:63 of the current PSW.

Condition Code:

Determined by the new PSW (bits 28:31).

Programming Notes:

The R1 field of this instruction must be zero.

The R2 field of this instruction must specify an even-numbered register.

This instruction can be used to change register sets. The new set becomes active for execution of the next instruction.

This instruction is a privileged operation.

If bit 22 of the new PSW is set, the state of the system queue is tested. If the queue is not empty, an SQS interrupt occurs. In this case, the newly loaded PSW is saved as the old PSW when the SQS interrupt occurs.

10.6.3 Exchange Program Status Register (EPSR)

<u>Assembler Notation</u>	<u>Opcode</u>	<u>Format</u>
EPSR R1,R2	95	RR

Operation:

Bits 0:31 of the current PSW replace the contents of the register specified by R1. The contents of the register specified by R2 then replace bits 0:31 of the current PSW.

Condition Code:

Determined by the new PSW (bits 28:31).

Programming Notes:

R1 and R2 can specify any general-purpose registers.

If R1 and R2 specify the same register, bits 0:31 of the current PSW are copied into the register specified by R2, but otherwise remain unchanged.

This instruction can be used to change register sets. The new set becomes active for execution of the next instruction.

This instruction is a privileged operation.

If bit 22 of the new PSW is set, the state of the system queue is tested. If the queue is not empty, an SQS interrupt occurs. In this case, the newly loaded PSW is saved as the old PSW when the SQS interrupt occurs.

10.6.4 Simulate Interrupt (SINT)

<u>Assembler Notation</u>	<u>Opcode</u>	<u>Format</u>
SINT I2(X2)	E2	R11
SINT R1,I2(X2)	E2	R11

Operation:

The least significant 10 bits of the second operand are presented to the interrupt handler as a device number. The device number is used to index into the ISPT, simulating an interrupt request from an external device. The result is either an immediate interrupt or an auto driver channel operation.

Condition Code:

The condition code is determined by the status of the addressed device in the case of the immediate interrupt, or set by the auto driver channel at termination.

Programming Notes:

It is assumed that an interrupt from level 0 is required, and register set 0 is selected.

This instruction is a privileged operation.

This instruction causes the processor to load registers 0 through 3 or 0 through 4 of the new set as for a real interrupt request.

During the execution of this instruction, the device is addressed and the status byte is returned in register 3 of the new set.

If the specified device does not respond to the status request, register 3 of the new set contains X'00000004' due to time-out. If an immediate interrupt is being simulated, the V flag is also set in the condition code as a result of the time-out.

The SINT instruction does not cause any pending interrupt to be acknowledged.

10.6.5 Supervisor Call (SVC)

<u>Assembler Notation</u>	<u>Opcode</u>	<u>Format</u>
SVC N,D2(X2)	E1	RX1, RX2
SVC N,A2(FX2,SX2)	E1	RX3

Operation:

The second operand (normally the program address of an SVC parameter block) replaces bits 8:31 of register 13 of the set designated by the SVC new PSW status. Bits 0:7 of this register are forced to zero. The current PSW replaces the contents of registers 14 and 15 of that set. The fullword quantity located at X'000098' in memory replaces bits 0:31 of the current PSW. The 4-bit N field is doubled and added with X'00009C'. The halfword quantity located at the resultant address becomes the current LOC.

Condition Code:

Determined by the new PSW (bits 28:31).

Programming Note:

This instruction provides a means to switch from the protect mode to the supervisor mode. It is used by a program running under an operating system to initiate certain functions in the supervisor program. The second operand address is normally a pointer to the memory location of parameters needed by the supervisor program to perform the specified function. Such a pointer must indicate a parameter block aligned to a fullword boundary. The type of SVC is specified in the N field of the instruction. Sixteen different calls are provided for. Return from the supervisor is made by executing an LPSWR instruction specifying the stored old PSW in registers 14 and 15 of the set selected by the SVC interrupt new PSW (LPSWR R14).

10.6.6 System Breakpoint (BRK)

<u>Assembler Notation</u>	<u>Opcode</u>	<u>Format</u>
BRK	88	SF

Operation:

The BRK instruction causes the processor to enter the console mode.

Programming Notes:

The LOC is not incremented.

This instruction is a privileged instruction.

10.6.7 Privileged System Function (PSF)

<u>Assembler Notation</u>	<u>Opcode</u>	<u>Format</u>
PSF N,D2(X2)	DF	RX1,RX2
PSF N,A2(FX2,SX2)	DF	RX3

Operation:

The PSF instruction can perform any one of 16 functions, as specified by the value contained in the N field. The assembler recognizes extended mnemonics that cause the proper value to be specified in the N field of this instruction. The nature of the specified function can vary from processor to processor. The following paragraphs detail PSF operations performed by this processor.

VALUE OF N	EXTENDED PSF MNEMONIC	MEANING
0	REL	Read Error Logger
1	LPSTD	Load Process Segment Table Descriptor
2	LSSTD	Load Shared Segment Table Descriptor
3	STPS	Store Process State
4	LDPS	Load Process State
5	ISSV	Save Interruptible State
6	ISRST	Restore Interruptible State
7	XSTB	Store Byte, no ECC
8	RMVF	Reset Memory Voltage Failure

Programming Notes:

This instruction is a privileged instruction.

PSF functions selected by values of N other than those listed above are undefined for this processor and result in an illegal instruction interrupt.


10.6.7.1 Read Error Logger (REL)

Assembler Notation	Opcode	Format
REL R2	DF0	RX1 (see programming notes)

Operation:

Data read from the error logger replaces the contents of the register specified by R2+1. Bit 16 of the data is copied to the L flag in the condition code. Once the data has been read from the error logger location, the status flag (bit 16) is set to zero.

The format of the data read from the error logger is shown below.



BITS	MEANING
0:15	are always zero.
16	is the status bit. If this bit is set to one, an error has been logged.
17:20	are always set to one.
21:25	are the SIP row address. In a multiprocessor system, bits 21 and 22 are zero, and bits 23, 24 and 25 point to the row that had a failure. If the status flag is not set, the SIP row address is undefined and has no meaning.
26:31	are the syndrome bits. These bits describe the SIP within the row which last had a failure. (See the Model 3205 System Theory of Operation and Maintenance Manual for table.) If the status flag is not set, the syndrome bits are undefined and have no meaning.

Condition Code (after reading error logger status):

C	V	G	L	
X	X	X	0	No new error bits in error logger
X	X	X	1	New error bit in error logger

Programming Notes:

This instruction is a privileged instruction.

The R2 field of this instruction must specify an even-numbered register.

Reading error logger status sets the error bit to zero, but does not necessarily zero the error logger bits at any syndrome address.

REL is assembled as an RX1 format instruction in which the displacement field is always zero.


REL is an extended PSF mnemonic.

10.6.7.3 Load Shared Segment Table Descriptor (LSSTD)

Assembler Notation	Opcode	Format
LSSTD D2(X2)	DF2	RX1,RX2
LSSTD A2(FX2,SX2)	DF2	RX3

Operation:

The second operand address points to a fullword shared segment table descriptor (SSTD), which has the following format:



BITS	MEANING
0:1	are reserved and must be zero.
2:14	contain the number of doubleword entries in the shared segment table (SST), minus one.
15:31	contain the absolute address of the SST in main memory, divided by 128.

The data in the SST is used in translation of program addresses from program to physical address space when PSW bit 21 is set if the PST entry specifies that the segment is shared.

Condition Code:

Unchanged

Programming Notes:

The operand address must be aligned to a fullword boundary.

The LSSTD instruction can be executed regardless of the state of PSW bit 21.

The new SST is available for execution of the next instruction that is executed with PSW bit 21 set.

Following an LSSTD instruction, the PSTD must be loaded, using the LPSTD or LDPS instruction, before attempting MAT translation with the newly defined SST.

This instruction is a privileged instruction.

LSSTD is an extended PSF mnemonic.

10.6.7.4 Store Process State (STPS)

<u>Assembler Notation</u>	<u>Opcode</u>	<u>Format</u>
STPS D2(X2)	DF3	RX1,RX2
STPS A2(FX2,SX2)	DF3	RX3

Operation:

The process state, defined by the old PSW in registers 14 and 15 of the current set, is saved in the area of memory starting at the address specified by the operand. The area has the following format:

<u>NORMAL OFFSET (BYTES)</u>	<u>STORED DATA</u>
0-7	Process PSW
8-11	Reserved; not used
12-75	Process general registers
76-139	Process interruptible state
140-235	Single and double precision floating point registers

Condition Code:

Unchanged

Programming Notes:

The operand address must be aligned to a fullword boundary.

This instruction is a privileged instruction.

STPS is an extended PSF mnemonic.

The process general register set is specified by the old PSW in register 14 when this instruction is executed.

If bit 14 of the process PSW in register 14 is zero, the process interruptible state is not saved, and the save area is compacted accordingly. In this case, the process' floating point registers are saved beginning at an offset of 76 bytes from the specified operand address.

If bit 13 of the process PSW in register 14 is set, the floating point registers are not saved, and the save area is compacted accordingly.

10.6.7.5 Load Process State (LDPS)

<u>Assembler Notation</u>	<u>Opcode</u>	<u>Format</u>
LDPS D2(X2)	DF4	RX1,RX2
LDPS A2(FX2,SX2)	DF4	RX3

Operation:

Data from the area of memory specified by the operand replaces the current process state. The area has the following format:

<u>NORMAL OFFSET (BYTES)</u>	<u>STORED DATA</u>
0-7	Process PSW
8-11	PSTD
12-75	Process general registers
76-139	Process interruptible state (if bit 14 in saved PSW is set)
140-235	Process single and double precision floating point registers (if bit 13 in saved PSW is zero)

The new PSW at the operand address specifies the general register set, which is loaded from the save area. If bit 14 of the new PSW is set, the interruptible state is loaded from the save area. If bit 13 of the new PSW is zero, then the single and double precision floating point registers are loaded from the save area. If bit 21 of the new PSW is set, the PSTD is loaded. Finally, the new PSW at the operand address becomes the current PSW.

Programming Notes:

The operand address must be aligned to a fullword boundary.

This instruction is a privileged instruction.

LDPS is an extended PSF mnemonic.

If bit 14 of the new PSW is zero, the process interruptible state is not loaded, and the save area is assumed to be compacted accordingly. In this case, the process's floating point registers are loaded from memory beginning at an offset of 76 bytes from the specified operand address.

If bit 13 of the new PSW is set, the process's floating point registers are not loaded, and the save area is assumed to be compacted accordingly.

If bit 22 of the new PSW is set, the state of the system queue is tested before testing the wait bit (bit 16). If the queue is not empty, an SQS interrupt occurs. In this case, the newly loaded PSW is saved as the old PSW when the SQS interrupt occurs.

The state of the wait bit (PSW bit 16) is tested before the next instruction is executed.

The process register set is selected in order to load the process general registers. All data is fetched from the save area before the process PSW is loaded. If a fault occurs during the execution of this instruction, one or more of the specified registers may have been modified. The old PSW presented to the fault ISR in register 14 can select the general register set specified by the process PSW in the save area, but is otherwise the same as the PSW in effect when this instruction is fetched and executed. The old PSW LOC presented to the ISR in register 15 points to the LDPS instruction.

10.6.7.6 Save Interruptible State (ISSV)

<u>Assembler Notation</u>	<u>Opcode</u>	<u>Format</u>
ISSV D2(X2)	DF5	RX1,RX2
ISSV A2(FX2,SX2)	DF5	RX3

Operation:

The contents of the interruptible instruction scratchpad registers are stored in the 16 fullwords of memory starting at the address specified by the operand.

Condition Code:

Unchanged

Programming Notes:

The operand address must be aligned to a fullword boundary.

This instruction is a privileged instruction.

ISSV is an extended PSF mnemonic.

10.6.7.7 Restore Interruptible State (ISRST)

<u>Assembler Notation</u>	<u>Opcode</u>	<u>Format</u>
ISRST D2(X2)	DF6	RX1,RX2
ISRST A2(FX2,SX2)	DF6	RX3

Operation:

The interruptible instruction scratchpad registers are loaded from the 16 fullwords in memory starting at the address specified by the operand.

Condition Code:

Unchanged

Programming Notes:

The operand address must be aligned to a fullword boundary.

This instruction is a privileged instruction.

ISRST is an extended PSF mnemonic.

10.6.7.8 Store Byte, No Error Correction Code (ECC) (XSTB)

<u>Assembler Notation</u>	<u>Opcode</u>	<u>Format</u>
XSTB D2(X2)	DF7	RX1,RX2
XSTB A2(FX2,SX2)	DF7	RX3

Operation:

The contents of bits 24:31 of general register 0 are stored in memory at the address specified by the operand without changing the ECC bits for the specified memory location.

Condition Code:

Unchanged

Programming Notes:

This instruction is a privileged instruction.

XSTB is an extended PSF mnemonic.

This instruction can be used in conjunction with the REL instruction to test the operation of the ECC.

10.6.7.9 Reset Memory Voltage Failure (RMVF)

<u>Assembler Notation</u>	<u>Opcode</u>	<u>Format</u>
RMVF	DFB	RX1 (See programming notes)

Operation:

The processor's internal MVF indication is reset. The MVF indication is set only as a result of the voltages to main memory not being maintained within acceptable margins during a power fail/restore sequence.

Condition Code:

Unchanged

Programming Notes:

This instruction should be executed by software loaded via the LSU after all interrupt new PSWs have been established. Proper use of this instruction prevents a potential runaway condition in the event of multiple power fail/restore sequences.

MVF is reset by the processor when the console mode is entered.

This instruction is a privileged instruction.

RMVF is an extended PSF mnemonic.

RMVF generates an RX1 format instruction in which the displacement field is always zero.

CHAPTER 11 MEMORY MANAGEMENT

11.1 INTRODUCTION

The memory address translator (MAT) supports:

- 4Mb physical address space
- 16Mb virtual address (VA) space
- Segmentation
- Shared segments
- Read, write and execute protection
- Four levels of hardware controlled access to segments

The purpose of the MAT is to translate a VA (used by the program) into a real address (RA) (used by physical memory). This translation frees programs from many of the limitations imposed by the hardware's configuration. This allows the operating system to respond efficiently to the changing memory needs of user programs. This is particularly important in multitasking environments where several independent programs are run concurrently.

Virtual memory has several significant advantages over other memory systems. They include the following:

- Each program runs in its own address space, thereby protecting memory from access or modification by other programs.
- A program can be loaded anywhere in physical memory without the appearance of requiring relocation. The memory assigned to a program does not need to be contiguous.

Figure 11-1 provides a simplified model of the translation from a VA to an RA and should aid in an understanding of the MAT process.

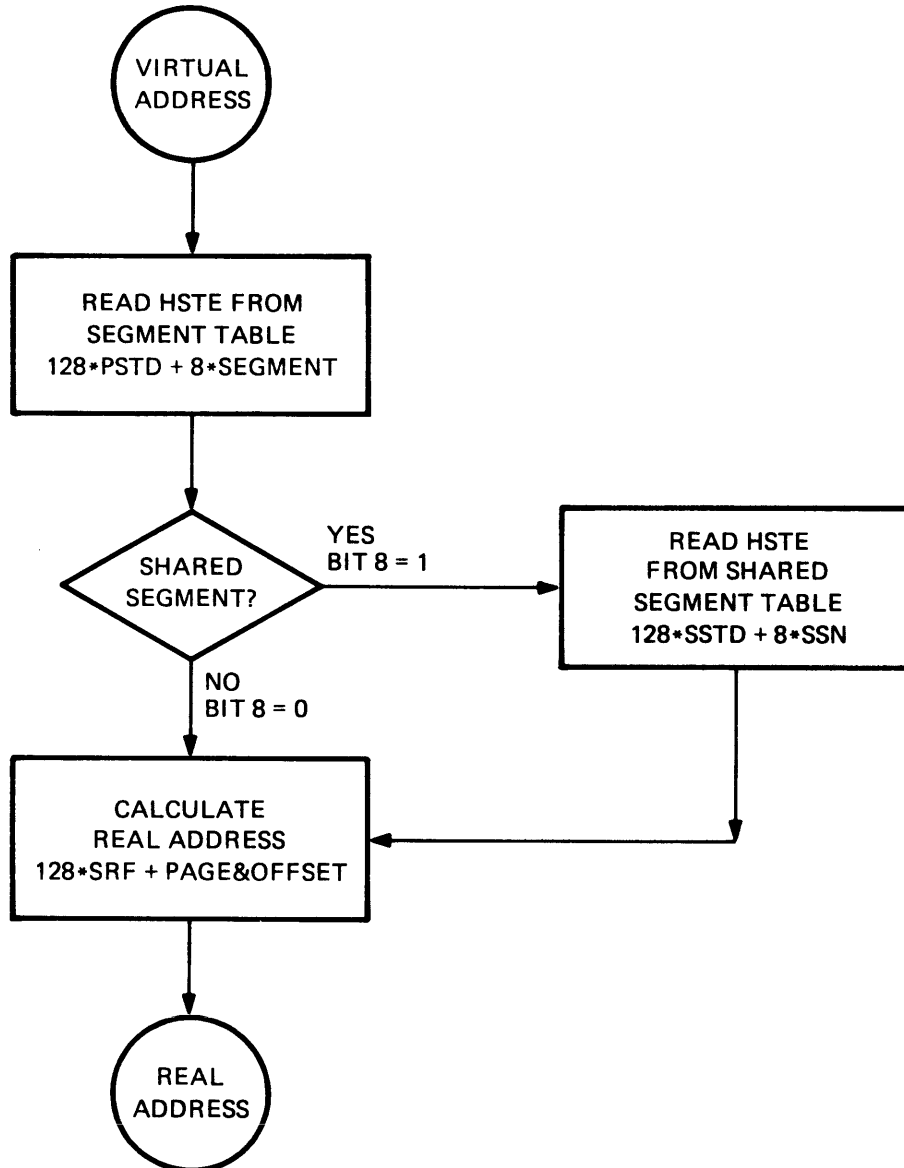


Figure 11-1 Flowchart of MAT Process


11.2 TRANSLATION FROM VIRTUAL TO REAL ADDRESS

The mapping of VA space to RA space is accomplished by using information supplied in a segment table. The table must be aligned with a 128-byte boundary in physical memory and can contain from 1 to 256 doubleword entries. Each doubleword entry, called a segment table entry (STE), is indexed by the segment number field of the VA.

The translation is controlled by two tables in main memory: the process segment table (PST) and the shared segment table (SST). The central processing unit (CPU) locates these tables using two special descriptor registers: the process segment table descriptor (PSTD), which points to the PST, and the shared segment table descriptor (SSTD), which points to the SST. The PSTD and SSTD are initialized at the beginning of each task by the operating system. A description of the PSTD and SSTD registers is shown in Figure 11-2.

8064

PSTD: PROCESS SEGMENT TABLE DESCRIPTOR REGISTER



SSTD: SHARED SEGMENT TABLE DESCRIPTOR REGISTER




Figure 11-2 PSTD and SSTD Registers

The PST contains two types of entries: shared and private. Depending on whether the segment is shared or private, varying processes are implemented. In either case, the eight most significant bits of the VA are added to the PSTD, giving the resultant address of the hardware segment table entry (HSTE). The HSTE is read by the operating system to determine whether the segment is shared or private. If bit 8 is set, then the segment is shared and the entry points to an entry in the SST. Otherwise, the segment is private.

In order to obtain a real address using a private segment, the least significant sixteen bits of the virtual address are added to the segment relocation field (SRF) (see Figure 11-3). However, if bit 8 is set and the segment is shared, then two further operations occur. In the first operation, the SSN is multiplied by 8 and added to the SSTD in order to obtain the shared hardware segment table entry. In the second operation, the SRF is added to the offset. The final result is a translated real address, as shown in Figure 11-4.

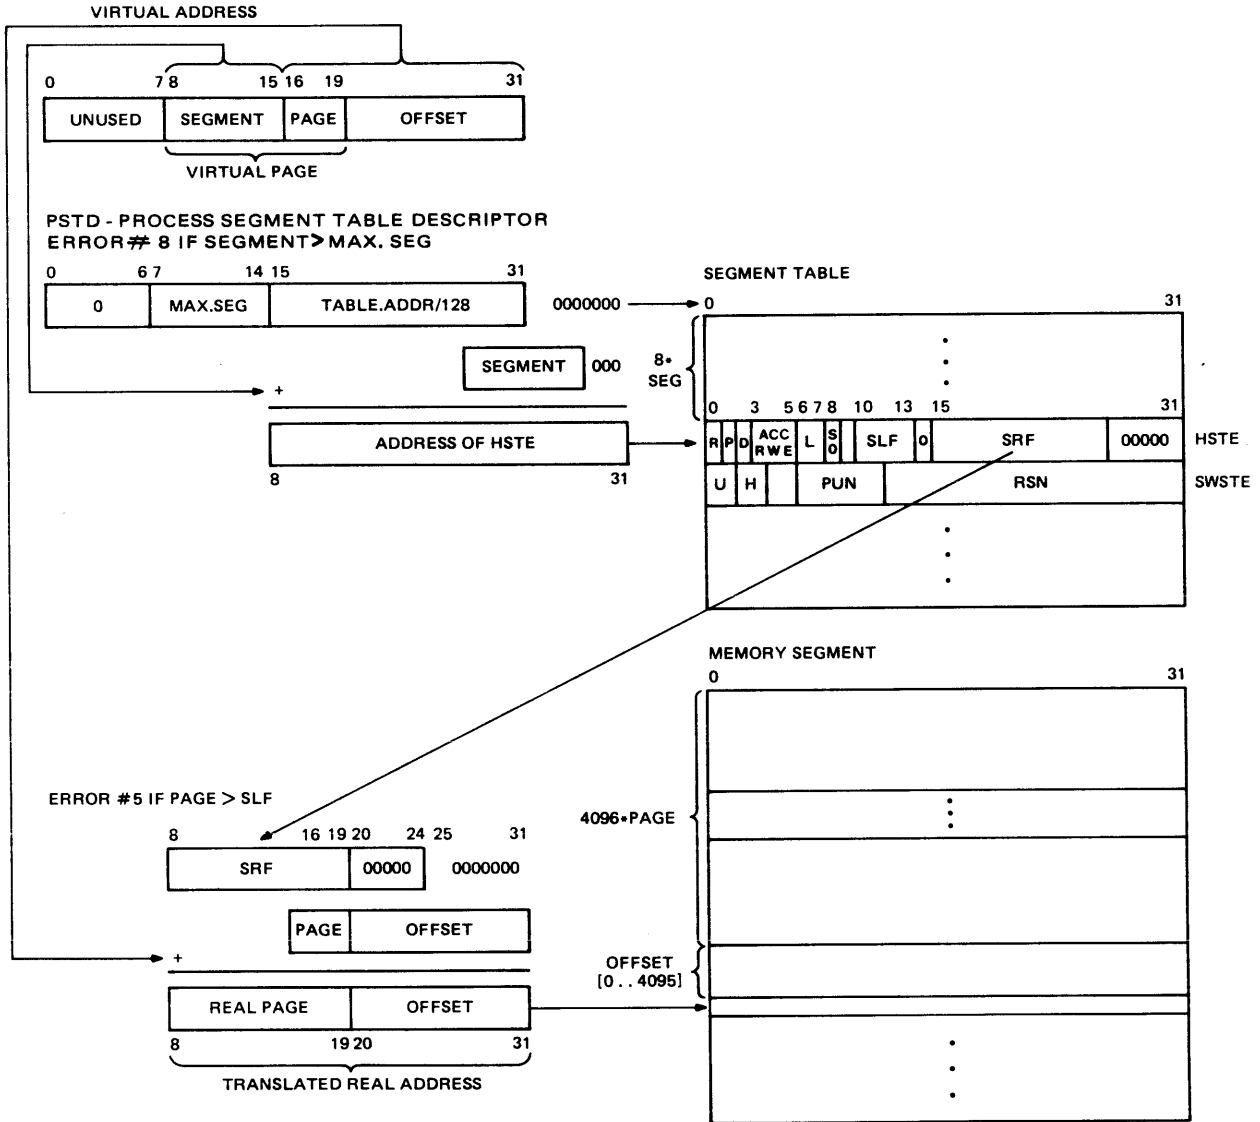


Figure 11-3 MAT Translation, Private Segment

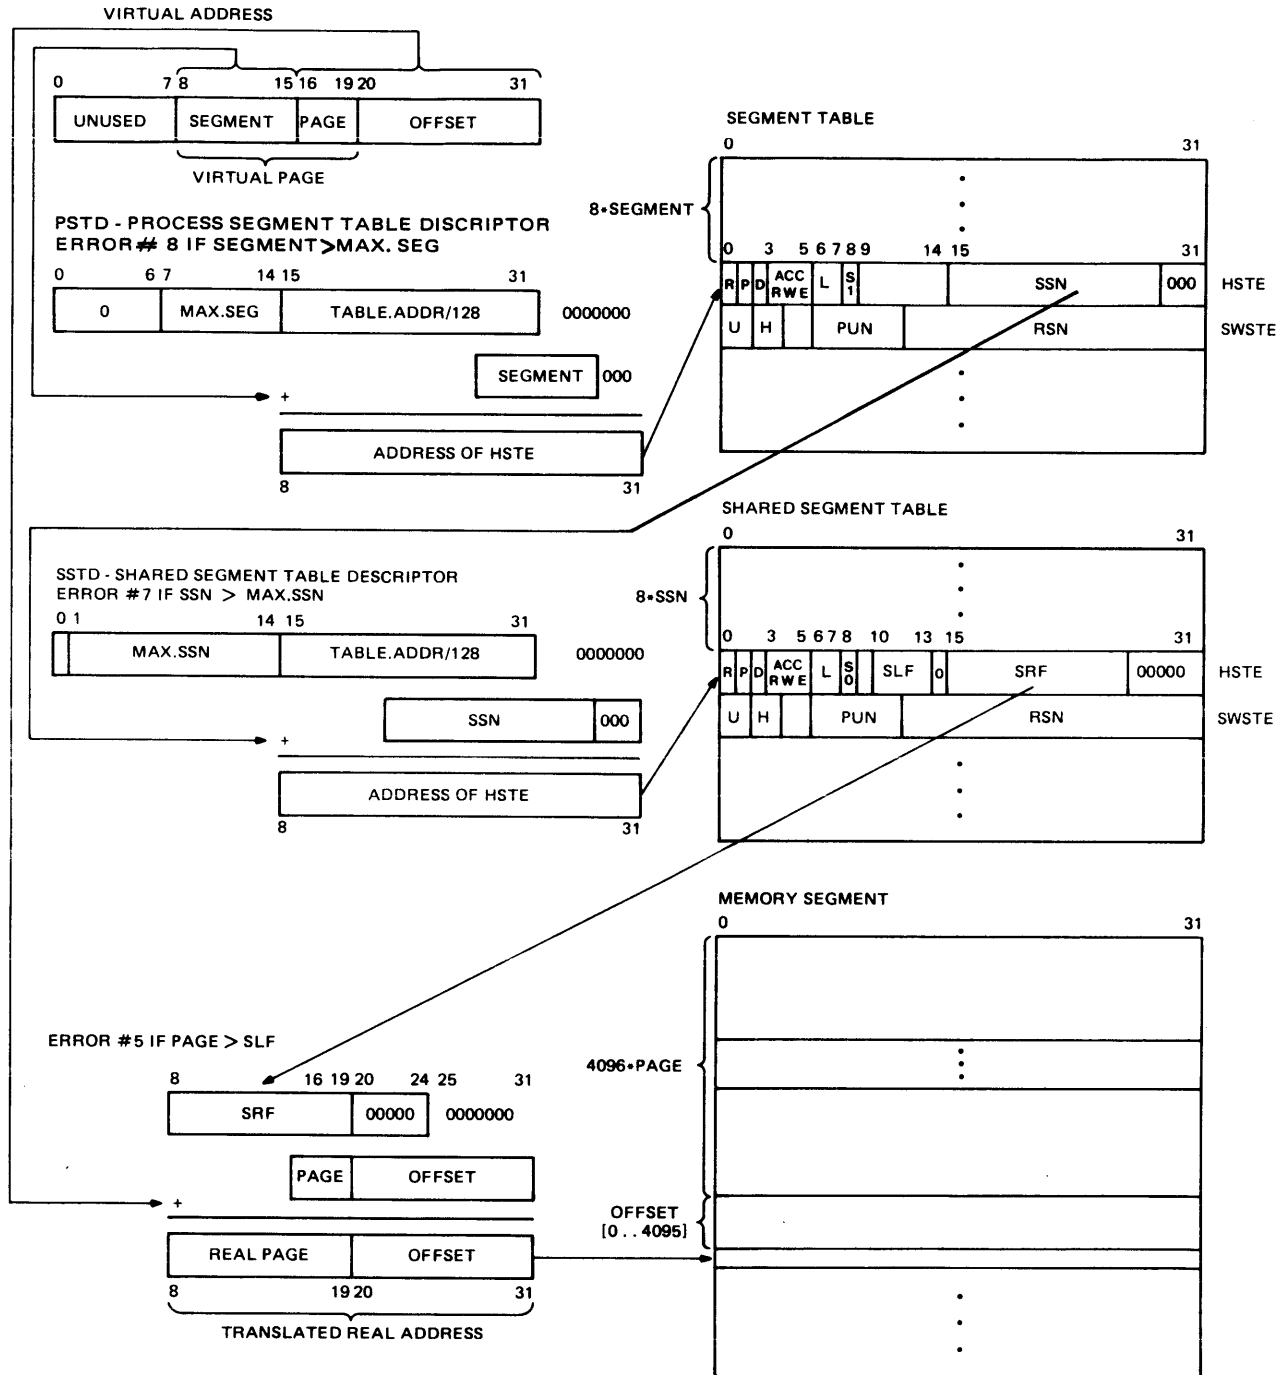


Figure 11-4 MAT Translation, Shared Segment

11.3 ADDRESS SPACE

This processor supports management of a 4Mb RA or 16Mb VA space. When RA or VA are manipulated, they are treated as 24-bit quantities. In general, 32-bit quantities are available to the processor for address calculation. When intermediate calculations are complete, bits 0 through 7 of the 32-bit effective result are forced to zero or discarded, giving a calculated address 24 bits in length, which occupies bits 8 through 31 of the 32-bit effective result.

In some instances, an address consisting of less than 24 bits can be used by the processor. Such an address is extended to 24 bits in length by forcing the most significant bits to zero.

11.3.1 Virtual Address (VA)

The VA consists of three fields: segment, page and offset. The segment and page comprise the virtual page address, which is translated into a real page address by the MAT hardware. The offset portion is not affected by the translation. See Figure 11-5 for a diagram of the VA.




Figure 11-5 Virtual Address

11.3.1.1 Segment Field

The VA contains a maximum program address (PA) space of 16Mb that is divided into 256 segments of 65,536 bytes each. A particular 64kb segment is selected by the most significant eight bits of the VA to make up what is known as the segment field.

For example, VA in the range Y'000000' to Y'00FFFF' select segment 0, as the most significant eight bits of each address are zero. VA in the range Y'4F0000' to Y'4FFFFFF' select segment 4F₁₆ (79₁₀), as the most significant eight bits of the address are 4F₁₆.

11.3.1.2 Offset and Page Field

The offset and page fields are comprised of the least significant 16 bits of the VA, and this value is used as a byte offset into the selected segment. The offset field of the VA has no special significance to the MAT except with respect to segment limit checking.

11.3.2 Selection of Virtual or Physical Addressing

Program status word (PSW) bit 21, the relocation/protection bit, controls selection of virtual or real addressing. When bit 21 is zero, the MAT is disabled. In this mode, all addresses generated are physical addresses. No segment table is used; hence, no level checking, access mode checking, etc., is performed. Bits 10 and 11 of the PSW (the access level bits) are ignored in the physical mode.

The user of the physical mode must be careful when modifying memory. The fact that a data area has been modified is not recorded by hardware. If it is desired to reflect the modification information in the segment tables, this must be done explicitly by the program running in the physical mode.

When PSW bit 21 is set, the MAT is enabled. All addresses generated are VAs, which are translated to RAs using the segment tables. System software must ensure that segment table addresses have been specified via the Load Process Segment Table Descriptor (LPSTD) and Load Shared Segment Table Descriptor (LSSTD) instructions.

When the MAT is enabled, bits 10 and 11 of the PSW indicate the level at which the program is running. When a VA is generated, the access level specified in the segment table entry (STE) is compared to the contents of bits 10 and 11 of the PSW. If the value of bits 10 and 11 is greater than or equal to the access level specified in the STE, then access to the segment is permitted; otherwise, a MAT fault occurs. System software should set bits 10 and 11 of the PSW according to the level at which the process is running to ensure protection of segments.

11.4 SHARED AND PRIVATE SEGMENTS

There can be a number of processes resident in the system at any given time. Each of these processes has its own VA space requirements, reflected in the PST associated with that process. Consequently, there can be several PSTs in memory concurrently, although only one, the segment table for the active process, may be known to the MAT at any given time.

Segments of the VA space of a process that are used only by that process are called private segments. Other segments of the VA space that can be shared with other processes may exist; these segments are consequently called shared segments. Although the STE describing a shared segment may be replicated in the segment tables associated with each process using the segment, it is preferable to maintain a separate SST. For a shared segment, the private STE has an indication that the segment's description is not found in the PST, but rather in the SST.

The data contained in a segment must be stored in contiguous locations in physical memory. This is called unpagged allocation. For unpagged allocation, each segment must be aligned to a 4,096 byte boundary in physical memory.

11.4.1 Segment Table Descriptors (STDs) and Their Use

The MAT is enabled only when PSW bit 21 is set. Prior to enabling the MAT, the locations and sizes of the PST and SST to be used must be identified to the system by loading the appropriate descriptor registers. These registers can be changed while the MAT is enabled. To specify the address of the PST to the system, the LPSTD instruction is used; to specify the address of the SST, the LSSTD instruction is used.

11.4.1.1 Format of a Segment Table Descriptor (STD)

Bits 0 and 1 of the STD are reserved and must always be zero. Bits 2 through 14 specify the segment table size, minus 1. For example, if the segment table size were 4, this field would have a value of 3. For a PSTD, this field has a maximum value of 255 (Y'FF'). For an SSTD, this field has a maximum value of 8,191 (Y'1FFF').

Bits 15 through 31 of the STD specify the segment table RA, divided by 128. A segment table must be aligned to a 2-byte boundary in physical memory. For example, if a segment table starts at location Y'035F80', then bits 15:31 of the STD contain Y'06BF' (Y'035F80'; divided by 128). The value of 0 for this field is a reserved value and, therefore, no segment table can start at physical address 0.

11.4.2 Setting the Virtual Address Space Size

When a PSTD is loaded, its segment table size field determines the maximum valid VA. The maximum valid virtual address (MVVA) is defined by the following formula:

$$\text{MVVA} = (\text{number of segment table entries} - 1) * (\text{Y}'10000') + \text{X}'\text{FFFF}'$$

The MVVA includes address space for the zeroth STE. For example, if the specified segment table size in the STD is 5 (requiring six segment table entries), then VAs in the range Y'000000' to Y'05FFFF' are considered valid. If a VA is generated that is greater than the MVVA, a MAT fault occurs (see Section 11.6).


Within the valid VA space, there may be segments that are not used. For example, a VA space of a process uses segments 0, 1, 2 and 5, while segments 3 and 4 are unused in that process. In this case, the segment table must contain entries for segments 3 and 4. To indicate that each of these segments is unused, its STE indicates that it is nonpresent and unused (see Section 11.5.4).

11.5 SEGMENT TABLE ENTRIES (STEs)

Figure 11-6 represents an STE. The STE consists of 64 bits; bits 0:31 describe the hardware portion and bits 32:63 comprise the software entry. Entries in the SST and the PST have the same format with minor differences, which are detailed in the following sections.

8068

HSTE: HARDWARE SEGMENT TABLE ENTRY
 S = 1: SHARED SEGMENT, INDEX INTO SST



- R – REFERENCED
- P – PRESENCE
- D – DIRTY (WRITTEN)
- RD – READ ENABLE
- WR – WRITE ENABLE
- EX – EXECUTE ENABLE
- ACCESS – SWSTE. ACCESS AND PSTE. ACCESS
- LEVEL – MINIMUM (SWSTE. LEVEL, PSTE.LEVEL)

SWSTE: SOFTWARE SEGMENT TABLE ENTRY



Figure 11-6 STE and SWSTE

11.5.1 Segment Table Entry (STE) Size

An STE is a 64-bit quantity. Bits 0:31 are the HSTE and bits 32:63 are the SWSTE. The HSTE contains the necessary information to enable hardware to perform VA to RA translation.

The SWSTE contains information used by system software to manage the private address space and keep track of segment history.

11.5.2 Segment Tables

The tables contain two words for each segment: the HSTE and the SWSTE. The HSTE is used by the microcode to translate the VA. The SWSTE is used only by the memory management program.

11.5.3 Hardware Segment Table Entry (HSTE)

Bits 0:31 of the STE contain the HSTE, which is comprised of several fields as shown in Figure 11-7. There are eight bits in the HSTE for memory management. These bits allow the operating system to:


- specify which tasks have access to a segment
- specify acceptable operations on a segment, and
- keep a record of whether a segment was accessed, written to or existing.

Reserved fields of the HSTE must always contain zero. A list of the bits and their function follows.

8069

HSTE: HARDWARE SEGMENT TABLE ENTRY

S = 0: POINTER TO SEGMENT. (IN PST OR SST.)



R – REFERENCED (SEGMENT MUST BE ALIGNED TO PAGE BOUNDARY.)
P – PRESENCE
D – DIRTY (WRITTEN)
RD – READ ENABLE
WR – WRITE ENABLE
EX – EXECUTE ENABLE

S=0 PRIVATE SEGMENT
S=1 SHARED SEGMENT (INDEX INTO SSN)
SLF – SEGMENT LENGTH FIELD
(MAXIMUM PAGE NUMBER IN SEGMENT)

Figure 11-7 Bit Representation of HSTE

Reference Bit (R)

Bit 0 of the HSTE is called the reference (R) bit. This bit of the HSTE is set by the hardware when a reference is made to the segment by software; i.e., when an attempt is made to read, write or execute the contents of the segment. DMA I/O does not set the reference bit for the participating segment.

Presence Bit (P)

Bit 1 of the HSTE is the presence (P) bit. The presence bit is set when the segment described by the STE is in memory; it is zero when the segment is not present in memory. A segment is said to be present when any portion of the segment is in memory. When P=0, hardware ignores the contents of the rest of the HSTE (bits 2 through 31), but the R bit is set as a result of the attempted access. Referencing a segment that is not present (P=0) causes a fault (see Section 11.6.1.2).

Dirty Bit (D)

Bit 2 of the HSTE is called the dirty (D) bit. This bit of the HSTE is set by the hardware when a program modifies (writes into) a segment. This bit is not set by DMA I/O operations that modify the segment. If P=0, attempts to modify the segment do not affect the dirty bit. If the PST STE indicates that the segment is shared, the D bit is set in both the SST STE and the PST STE for a write.

Access Mode Bits (A)

Bits 3 through 5 of the HSTE are called the access mode (A) bits. These bits determine the allowed modes of access to the segment. The three modes of access to a segment are: read access, write access and execute access. If an attempt is made to access a segment in a manner not permitted by the setting of the access mode bits, a MAT fault occurs. For example, instructions cannot be fetched from a segment unless execute access is enabled. See Section 11.6 for a detailed definition of all MAT faults. Table 11-1 defines access field settings and the types of access that are enabled.

TABLE 11-1 SEGMENT ACCESS FIELD SETTINGS

BIT	ACCESS MODE	FIELD SETTING
3	Read enabled	0 = Read protected
4	Write enable	0 = Write protected
5	Execute enable	0 = Execute protected

Access Level Bits (L)

Bits 6 to 7 of the HSTE are called the access level (L) bits. The L field is used in conjunction with bits 10 and 11 of the PSW to determine if a program can access a segment. If the value represented by the contents of PSW bits 10 and 11 is greater than or equal to the L bits, then access to the selected segment is allowed; otherwise, a fault occurs. The L bits are checked before the A bits. See Section 11.6 for a detailed description of MAT faults.

Shared Segment Bit (S)

Bit 8 of the HSTE is called the shared segment (S) bit. If this bit is zero, the MAT performs protection and relocation functions as defined for a private segment. The S bit must be zero for all STEs in the SST.

If the S bit is set in a PSTE, the selected segment is a shared segment. In this case, the segment relocation field (SRF) of the PST STE is used as a byte offset into the SST. The SST STE found at the resulting address is used by the MAT in performing protection and relocation functions, as follows: the values of the A fields in the PST STE and the SST STE are ANDed to determine the allowed access modes; all other MAT protection and relocation functions are performed using data from the SST STE.

Segment Limit Field (SLF)

Bits 10 through 14 of the HSTE are called the SLF. The SLF indicates the size of a segment, according to the following formula:

$$SLF = (\text{size of segment}) \text{ divided by } (4\text{kb}) - 1$$

For example, for a segment of size 4kb, the SLF would contain 0. When a VA is generated, the contents of bits 8:12 of the VA is compared to the SLF. If the SLF is less than this number, a MAT fault occurs (see Section 11.6).

The granularity of memory allocation is 4kb (4,096). This means that memory must be allocated in 4kb units.

Segment Relocation Field (SRF)

Bits 15 through 31 of the HSTE are called the SRF. The interpretation of the SRF depends upon the setting of the S bit. If S is set in the PST, the PST SRF contains a byte offset into the SST at which the STE for the segment can be found. If the value contained in the PST SRF is greater than the size in bytes of the SST, a MAT fault occurs. See Section 11.4 for a detailed description of MAT faults.

If S is zero in the PST, the SRF contains the RA of the segment in memory, divided by 128. For example, if the segment starts at RA Y'146000', the SRF of the HSTE should contain X'28C0'.

Note that for a shared segment, the least significant three bits of the SRF in a PST HSTE must be zero, as the indicated SST HSTE is aligned to a doubleword boundary. The least significant five bits must be zero for the SRF in all other cases, as the SRF is the address of a segment aligned to a 4kb boundary divided by 28. If the MAT attempts to use an SRF that has a one in any of these trailing bits, the results are undefined.

11.5.4 Software Segment Table Entry (SWSTE)

Bits 32:63 of the STE are called the SWSTE. These bits are used by software in managing the VA space and have no hardware significance to the processor.

The information presented in this section details one possible scheme for software management of memory. The fullword SWSTE is available for any software memory management scheme.

1. Usage mode bits - Bits 0 and 1 of the SWSTE (bits 32 and 33 of the STE) are called the usage mode (U) bits. In conjunction with the P bit in the HSTE, these bits indicate the state of each segment.

Possible states of a segment are:

BIT STATUS	P	U
Unused	0	00
Used	0	01
Loading	0	10
Unloading	0	11
Active	1	00
I/O ongoing	1	01
Unload pending	1	10

- Unused state - If a segment is logically an invalid portion of the VA space of a process, then it is said to be in the unused state. For example, if a process has data in segments 0, 1, 2 and 5, but has no data in segments 3 and 4, references to segments 3 and 4 are invalid. Since segments 3 and 4 must each have an STE, the fact that these segments represent an invalid portion of the VA space of the process is indicated by setting P=0 and U=0 for the unused state. Since P=0 in the HSTE, any references to such a segment cause a MAT fault. The fault handler, using the U field, may determine that the fault was caused by a reference to an invalid portion of the private VA space and can take appropriate action.
- Used state - If a segment is logically a valid part of the private VA space, but is not physically present in memory, that segment is said to be in the used state. When a reference is made to such a segment, a MAT fault occurs since P=0 in the HSTE. The fault handler, using the U field, may determine that the fault was caused by a reference to a used segment, and then take action to load the segment.

When a used segment is to be loaded, the segment state is changed by software from used to loading as explained in the next paragraph.

- Loading state - If a segment that is logically part of the VA space of the process is being moved from backing store into physical memory, it is said to be in the loading state. If a reference is made to a segment that is in the loading state, a fault occurs since P=0 in the HSTE. The fault handler normally places the process that made the reference into a wait state, pending completion of the loading operation.

When a segment has successfully been loaded, software normally changes its state from loading to active. At this point, any process that faulted because it made reference to the segment in the loading state and was consequently put into a wait state, may resume execution.

- Unloading state - When a segment that is logically part of the VA space of the process is being temporarily removed from physical memory and copied to back-up store, it is said to be in the unloading state. Any reference to such a segment causes a MAT fault, because P=0 in the HSTE. When a MAT fault occurs and the U field indicates that it was caused by a segment in an unloading state, the fault handler has a choice of actions it can take. For example, if the segment was unloaded to make physical memory space available, the process that made the reference can be put into a wait state. When the unload completes, the physical memory that the segment occupies may be retained, and the segment put into an active state. If it is desired to give the physical memory that the segment occupied to another segment, then the unloaded segment should be put into the used state and the fault should be treated in the same manner as faults that occur in a used state.

In some systems, a segment might be unloaded because it is being removed from physical memory and is also being removed logically from the VA space of the process. A dynamically attached and detached buffer segment is an example of such a segment. If a segment in an unloading state is being logically removed from the VA of the process, the segment enters an unused state at the completion of the unloading operation. Faults generated by references to such segments are normally errors.

- Active state - When a segment is logically in the VA space of a process, physically resident in memory, and free to be used by the process within the restrictions placed by the A and L fields, it is said to be in the active state. The active state is the normal state for a segment that is being used by a process.
- I/O ongoing state - When I/O operations are being performed upon the contents of an active segment, the segment is put into the I/O ongoing state. The physical memory occupied by the segment cannot be deallocated and reassigned to another segment. The segment should not be unloaded until all I/O operations terminate.

A segment should be in I/O ongoing state until all I/O operations performed upon the contents of the segment have been completed. At this point, the segment can be returned to the active state.

- Unload pending state - If it is determined that a segment is to be unloaded and the segment is in the I/O ongoing state, the segment must be placed in an unload pending state, which indicates that it is to be unloaded when all I/O operations terminate. When the last I/O operation completes, the segment can be placed in an unloading state and can then be unloaded. If the decision to unload the segment is changed while the segment is in an unload pending state, the state should be changed back to either I/O ongoing, if there are still outstanding I/O operations, or active, if all I/O operations have completed.

2. Reference history bits - Bits 34 and 35 of the STE (bits 2 and 3 of the SSTE) are called the reference history (R) bits. The H field is used to manage replacement algorithms. At given intervals, the state of the R and D bits in the HSTE are recorded in the H field and are reset in the HSTE.

The state of the R bit is retained only between intervals. For example, if the R bit is reset at the time it is examined, the H field will indicate that the last state of the R bit was reset (0). In contrast to this, once D has been set in the HSTE, that fact is retained in the H field until either the segment is unloaded or a copy of the modified state of the segment is made.

The H field is comprised of two bits. The most significant bit of the H field will always be set equal to the value of R at the time the HSTE was last scanned and reset.

The least significant bit of the H field will be an OR of its previous contents and the setting of D in the HSTE the last time D was scanned. This results in four possible values for the H field:

- No reference to the segment in last interval, unmodified (H=00)
- Reference to the segment in last interval, unmodified (H=10)
- Reference to the segment in last interval, modified (H=11)
- No reference to the segment in last interval, modified (at some previous time) (H=01)

3. Reserved field - Bits 36 and 37 of the STE (bits 4 and 5 of the SWSTE) are reserved. These bits must be set to 0.

4. Disk address - Bits 38 through 63 of the STE (bits 6 through 31 of the SWSTE) contain the disk address (DA) field. This field contains two subfields: the paging unit number (PUN), which is contained in bits 38 through 43 of the STE (bits 6 through 11 of the SWSTE), and the relative sector number (RSN), which is contained in bits 44 through 63 of the STE (bits 12 through 31 of the SWSTE). A paging unit is a randomly accessible device that can be read from or written to. This unit is used to load and unload segments. The PUN is used as an index into a page device table (PDT), which is used to translate the PUN into a physical device. The PDT entry contains a physical device address and a device starting sector. The RSN in the SWSTE is added to the starting sector specified in the PDT entry to compute the actual sector number at which segment can be found.

There can be up to 32 paging units used in a system at any given time. The PDT allows independence of the logical paging unit from the physical paging unit. For example, a given physical device could be divided into multiple paging units or several physical devices could be combined to be a single paging unit.

11.6 MEMORY ADDRESS TRANSLATOR (MAT) FAULTS

Previous sections of this manual have stated that certain conditions result in MAT faults. A fault is an indication that some exception condition has occurred and that system software should take some action in response. Some faults (such as access violation) are indicative of error, while other faults (such as presence fault) can be used for management of the software system.

11.6.1 Conditions that Cause Memory Address Translator (MAT) Faults

The conditions that cause MAT faults to occur are described in the following sections.

11.6.1.1 Process Segment Table (PST) or Shared Segment Table (SST) Size Exceeded Fault

The LPSTD or LDPS instruction defines the MVVA (see Section 11.7.1). If an address is generated that is greater than the MVVA, a PST size exceeded fault occurs.

The LSSTD instruction defines a size for the SST. If an STE in the PST specifies an SST offset greater than the size of the SST, an SST size exceeded fault occurs.

If the MVVA is exceeded for the PST, then no reference to memory is made. If the fault is caused by exceeding the valid size of the SST, then only the PST will have been referenced.

11.6.1.2 Nonpresence Fault

The nonpresence fault occurs when reference is made to an STE having its presence bit reset (0). The VA that caused the fault is returned to systems software. The R bit of the STE is set, but the contents of the segment and the D bit in the STE are not modified as a result of a reference to a nonpresent segment.

If the nonpresent segment can be loaded, the instruction that caused the fault may be reexecuted after the segment is loaded. For certain instructions, software intervention may be required to allow correct reexecution. Section 11.6.4 contains a detailed description of how to recover from a nonpresence fault.

11.6.1.3 Access Level Fault

An access level fault occurs when the access level specified by bits 10 and 11 of the PSW is less than the access level specified in an STE to which a reference is made. The R bit of the referenced STE is set; the contents of the segment and the D bit in the STE are not modified as a result of a reference to the segment, which causes an access level fault.

If system software can correct the fault, the faulting instruction may be reexecuted with certain restrictions (see Section 11.6.4 for details).

11.6.1.4 Access Mode Faults

There are three access mode faults: read access, write access and execute access fault. Each fault occurs when a mode of access is attempted to a segment that does not allow the attempted mode of access.

The R bit of the referenced STE is set; but the contents of the segment and the D bit in the STE are not modified as the result of an attempted access resulting in the access mode fault.

If system software can correct the fault, the instruction can be reexecuted with certain restrictions (see Section 11.6.4 for details).

11.6.1.5 Segment Limit Fault

A segment limit fault occurs when the value contained in bits 8:12 of a VA is greater than the value specified in the SLF field of the HSTE. The R bit of the STE is set; but the contents of the segment and the D bit in the STE are not modified as the result of an attempted access resulting in a segment limit fault.

If the system software can correct the fault, then the instruction that caused the fault may be reexecuted with certain restrictions (see Section 11.6.4 for details).

11.6.2 Fault Precedence

While some faults can be physically checked for in parallel by the hardware, there is a logical priority in which faults are checked (descending order):

- Segment table size exceeded
- Nonpresent segment
- Segment limit violation
- Access level violation
- Access mode violation

Detection of any of the listed MAT faults causes the user instruction to be aborted immediately. The reason for the abort is reported to system software as detailed in Section 11.4.3. Only one MAT fault can occur for a single memory operation.

11.6.3 Memory Address Translator (MAT) Fault Handling Routine

When a MAT fault occurs, the MAT fault handling routine pointed to by the MAT fault handler new PSW (physical location X'000090') is entered.

The PSW in effect at the time the fault occurs is placed in registers 14 and 15 of the set selected by the MAT fault handler new PSW. The LOC of the old PSW (register 15) contains the address of the instruction that caused the fault.

Register 13 of the selected set is loaded with a value to indicate the reason for the fault. The possible values are:

VALUE	MEANING
0	Reserved code
1	Access mode fault - execute attempted
2	Access mode fault - write attempted
3	Access mode fault - read attempted
4	Access level fault
5	SLF
6	Nonpresent segment fault
7	SST size exceeded
8	PST size exceeded

Register 12 of the selected set is loaded with the VA that caused the fault.

If the fault occurred during execution of the Load Multiple (LM) instruction, the calculated address of the start of the data block is placed in register 11 of the selected set.

11.6.4 Reexecution of Faulting Instructions

In general, an instruction that caused a correctable MAT fault can simply be reexecuted after the fault is corrected.

The LM instruction in some cases cannot simply be reexecuted, but must be simulated. When an LM instruction faults, register 11 of the set specified by the MAT interrupt new PSW is loaded with the VA calculated by the hardware as the effective second operand address of the instruction. If that address is the same as the VA that caused the fault (contained in register 12), then the instruction can be reexecuted once the fault has been corrected; no registers were modified by the LM instruction.

If the addresses in register 11 and register 12 are not equal, at least one register was modified by the LM instruction. Once the fault has been corrected, system software should build and execute an instruction sequence to load the required registers, using the calculated VA in register 11.

If the addresses are not equal, then the difference in the addresses, D , should be computed. The last register modified, $M = (D/4) - 1 + R1$, should be calculated. If M is less than the contents of the X2 field in an RX1 or RX2 instruction, or is less than the contents of both the FX2 and SX2 fields in an RX3, the instruction may be reexecuted. If this is not the case, then system software must build an instruction sequence to load the remaining registers from the appropriate memory locations. The location portion of the old PSW should then be incremented by the length of the faulted instruction. At this point, normal execution can be resumed by loading the old PSW.

11.6.5 Effect of System Initialization on the Memory Address Translator (MAT)

When the initialize switch (INIT) on the display panel is depressed or the processor is powered up, all segmentation, relocation, protection and MAT interrupts are disabled regardless of the state of bit 21 in the current PSW. The contents of the SST and PST descriptor registers must be restored by software after power fail.

The MAT remains disabled until an LPSTD instruction is issued. At this time, the MAT is enabled or remains disabled, according to the state of bit 21 of the current PSW.

11.7 MEMORY MANAGEMENT INSTRUCTIONS

Instructions are provided to control the MAT. These instructions are:

LPSTD	Load Process Segment Table Descriptor
LSSTD	Load Shared Segment Table Descriptor

11.7.1 Load Process Segment Table Descriptor (LPSTD)

<u>Assembler Notation</u>	<u>Opcode</u>	<u>Format</u>
LPSTD D2(X2)	DF1	RX3
LPSTD A2(FX2,SX2)	DF1	RX3

Operation:

The operand specifies the address of the fullword PSTD. This descriptor is loaded and its contents define the PST to be used in VA to RA translation when bit 21 of the PSW is set.

Condition Code:

Unchanged

Programming Notes:

This instruction is a privileged instruction.

The operand address must be aligned to a fullword boundary.

A PSTD can be loaded while PSW bit 21 is set or zero.

LPSTD is an extended PSF mnemonic.

11.7.2 Load Shared Segment Table Descriptor (LSSTD)

<u>Assembler Notation</u>	<u>Opcode</u>	<u>Format</u>
LSSTD D2(X2)	DF2	RX1,RX2
LSSTD A2(FX2,SX2)	DF2	RX3

Operation:

The operand specifies the address of the fullword SSTD. This descriptor is loaded and its contents define the SST to be used in virtual to physical address translation when bit 21 of the PSW is set.

Condition Code:

Unchanged

Programming Notes:

This instruction is a privileged instruction.

The operand address must be aligned to a fullword boundary.

An SSTD can be loaded while PSW bit 21 is set or zero.

LSSTD is an extended PSF mnemonic.

Following an LSSTD instruction, the PSTD must be loaded, using the LPSTD or LDPS instruction, before attempting MAT translation with the newly defined shared segment table.

APPENDIX A OPCODE MAP

635-6

MSD →

	0	1	2	3	4	5	6	7	8	9	A	B	C	D	E	F
LSD																
0		SRLS	BTBS		STH ⁴	ST ³	STE ³ ₁	STD ³ ₁		SRHLS			BXH ⁴	STM ³	TS ⁴	
1	BALR ⁴	SLLS	BTFS		BAL ⁴	AM ³	AHM ⁴	STME ³ ₁		SLHLS			BXLE ⁴	LM ³	SVC ³	
2	BTCR ⁴	CHVR	BFBS	PBR ²	BTC ⁴		PB ² ₄	LME ³ ₁	STDE ³ ₁	STBR			LPSW ³ _*	STB	SINT _*	
3	BFCR ⁴	LPER ¹	BFBS	LPDR ¹	BFC ⁴		LRA ³	LHL ⁴		LBR			THI	LB	SCP ³ _*	TI
4	NR		LIS	EXHR	NH ⁴	N ³	ATL ³	TBT	LED ³ ₁	EXBR	LEDR ₁		NHI	CLB		NI
5	CLR	LGER ₁	LCS		CLH ⁴	CL ³	ABL ³	SBT		EPSR _*	LEGR ₁		CLHI	AL _*		CLI
6	OR	LGDR ₁	AIS		OH ⁴	O ³	RTL ³	RBT			LDGR ₁		OHI		LA	OI
7	XR	LCER ₁	SIS	LCDR ₁	XH ⁴	X ³	RBL ³	CBT	LDE ³ ₁		LDER ₁		XHI		TLATE ³	XI
8	LR	LPSWR _*	LER ₁	LDR ₁	LH ⁴	L ³	LE ³ ₁	LD ³ ₁	BRK _*	WHR			LHI	WH ⁴ _*		LI
9	CR		CER ₁	CDR ₁	CH ⁴	C ³	CE ³ ₁	CD ³ ₁		RHR _*			CHI	RH ⁴ _*		CI
A	AR		AER ₁	ADR ₁	AH ⁴	A ³	AE ³ ₁	AD ³ ₁		WDR _*			AHI	WD _*	RRL	AI
B	SR		SER ₁	SDR ₁	SH ⁴	S ³	SE ³ ₁	SD ³ ₁		RDR _*			SHI	RD _*	RLL	SI
C	MHR	MR	MER ₁	MDR ₁	MH ⁴	M ³	ME ³ ₁	MD ³ ₁	RXR ⁵				SRHL		SRL	
D	DHR	DR	DER ₁	DDR ₁	DH ⁴	D ³	DE ³ ₁	DD ³ ₁		SSR _*			SLHL	SS _*	SLL	
E		LUR ₁	FXR ₁	FXDR ₁	LU ₁	⁴ CRC12	⁵ STBP	³ STMD ₁		OCR _*			SRHA	OC _*	SRA	
F		LWR ₁	FLR ₁	FLDR ₁	LW ₁	⁴ CRC16	⁵ LPB	³ LMD ₁					SLHA	PSF _*	SLA	

- 1. FLOATING POINT INSTRUCTION
- 2. HIGH-SPEED DATA HANDLING INSTRUCTION
- 3. SECOND OPERAND ADDRESS MUST BE FULLWORD ALIGNED
- 4. SECOND OPERAND ADDRESS MUST BE HALFWORD ALIGNED
- 5. USES SCRATCHPAD REGISTERS
- * PRIVILEGED INSTRUCTION

RXRX SUB FUNCTIONS

		MSD →				
		0	1	2	3	
		4	5	6	7	IMMEDIATE LENGTH SECOND OPERAND
		8	9	A	B	IMMEDIATE LENGTH FIRST OPERAND
		C	D	E	F	IMMEDIATE LENGTH BOTH OPERANDS
LSD						
	0	MVTU				} USE SCRATCHPAD REGISTERS
	1	MOVE		MOVEP		
	2	CPAN		CPANP		
	3	PMV		PMVA		
	4	UMV		UMVA		

PRIVILEGED SYSTEM FUNCTIONS (PSF)

<u>OP-CODE</u>	<u>MNEMONIC</u>	<u>MEANING</u>
DF0	REL	READ ERROR LOGGER
DF1	LPSTD	LOAD PROCESS SEGMENT TABLE DESCRIPTOR
DF2	LSSTD	LOAD SHARED SEGMENT TABLE DESCRIPTOR
DF3	STPS	SAVE PROCESS STATE
DF4	LDPS	LOAD PROCESS STATE
DF5	ISSV	SAVE INTERRUPTIBLE STATE
DF6	ISRST	RESTORE INTERRUPTIBLE STATE
DF7	XSTB	STORE BYTE WITHOUT ECC
DF8	RMVF	RESET MEMORY VOLTAGE FAILURE

APPENDIX B
INSTRUCTION SUMMARY - ALPHABETICAL BY MNEMONIC

MNEMONIC	OPCODE	INSTRUCTION
A	5A	Add
ABL	65	Add to Bottom of List
AD	7A	Add DPFP
ADR	3A	Add DPFP Register
AE	6A	Add SPFP
AER	2A	Add SPFP Register
AH	4A	Add Halfword
AHI	CA	Add Halfword Immediate
AHM	61	Add Halfword to Memory
AI	FA	Add Immediate
AIS	26	Add Immediate Short
AL	D5	Autoload
AM	51	Add to Memory
AR	0A	Add Register
ATL	64	Add to Top of List
B	430	Branch Unconditional
BAL	41	Branch and Link
BALR	01	Branch and Link Register
BC	428	Branch on Carry
BCR	028	Branch on Carry Register
BCS	208	Branch on Carry Short (Backward)
BCS	218	Branch on Carry Short (Forward)
BE	433	Branch on Equal
BER	033	Branch on Equal Register
BES	223	Branch on Equal Short (Backward)
BES	233	Branch on Equal Short (Forward)
BFBS	22	Branch on False Condition Backward Short
BFC	43	Branch on False Condition
BFCR	03	Branch on False Condition Register
BFFS	23	Branch on False Condition Forward Short
BL	428	Branch on Low
BLR	028	Branch on Low Register
BLS	208	Branch on Low Short (Backward)
BLS	218	Branch on Low Short (Forward)
BM	421	Branch on Minus
BMR	021	Branch on Minus Register
BMS	201	Branch on Minus Short (Backward)
BMS	211	Branch on Minus Short (Forward)
BNC	438	Branch on No Carry
BNCR	038	Branch on No Carry Register
BNCS	228	Branch on No Carry Short (Backward)
BNCS	238	Branch on No Carry Short (Forward)
BNE	423	Branch on Not Equal

MNEMONIC	OPCODE	INSTRUCTION
BNER	023	Branch on Not Equal Register
BNES	203	Branch on Not Equal Short (Backward)
BNES	213	Branch on Not Equal Short (Forward)
BNL	438	Branch on Not Low
BNLR	038	Branch on Not Low Register
BNLS	228	Branch on Not Low Short (Backward)
BNLS	238	Branch on Not Low Short (Forward)
BNM	431	Branch on Not Minus
BNMR	031	Branch on Not Minus Register
BNMS	221	Branch on Not Minus Short (Backward)
BNMS	231	Branch on Not Minus Short (Forward)
BNO	434	Branch on No Overflow
BNOR	034	Branch on No Overflow Register
BNOS	224	Branch on No Overflow Short (Backward)
BNOS	234	Branch on No Overflow Short (Forward)
BNP	432	Branch on Not Plus
BNPR	032	Branch on Not Plus Register
BNPS	222	Branch on Not Plus Short (Backward)
BNPS	232	Branch on Not Plus Short (Forward)
BNZ	423	Branch on Not Zero
BNZR	023	Branch on Not Zero Register
BNZS	203	Branch on Not Zero Short (Backward)
BNZS	213	Branch on Not Zero Short (Forward)
BO	424	Branch on Overflow
BOR	024	Branch on Overflow Register
BOS	204	Branch on Overflow Short (Backward)
BOS	214	Branch on Overflow Short (Forward)
BP	422	Branch on Plus
BPR	022	Branch on Plus Register
BPS	202	Branch on Plus Short (Backward)
BPS	212	Branch on Plus Short (Forward)
BR	030	Branch Unconditional Register
BRK	88	Breakpoint
BS	220	Branch Unconditional Short (Backward)
BS	230	Branch Unconditional Short (Forward)
BTBS	20	Branch on True Condition Backward Short
BTC	42	Branch on True Condition
BTCR	02	Branch on True Condition Register
BTFS	21	Branch on True Condition Forward Short
BXH	C0	Branch on Index High
BXLE	C1	Branch on Index Low or Equal
BZ	433	Branch on Zero
BZR	033	Branch on Zero Register
BZS	223	Branch on Zero Short (Backward)
BZS	233	Branch on Zero Short (Forward)
C	59	Complement Bit
CBT	77	Compare
CD	79	Compare Double Floating Point
CDR	39	Compare Double Floating Point Register
CE	69	Compare Floating Point
CER	29	Compare Floating Point Register
CH	49	Compare Halfword
CHI	C9	Compare Halfword Immediate

MNEMONIC	OPCODE	INSTRUCTION
CHVR	12	Convert Halfword Value Register
CI	F9	Compare Immediate
CL	55	Compare Logical
CLB	D4	Compare Logical Byte
CLH	45	Compare Logical Halfword
CLHI	C5	Compare Logical Halfword Immediate
CLI	F5	Compare Logical Immediate
CLR	05	Compare Logical Register
CPAN	8C/02	Compare Alphanumeric
CPANP	8C/22	Compare Alphanumeric and Pad
CR	09	Compare Register
CRC12	5E	Cyclic Redundancy Check Modulo 12
CRC16	5F	Cyclic Redundancy Check Modulo 16
D	5D	Divide
DD	7D	Divide Double Precision Floating Point
DDR	3D	Divide Double Floating Point Register
DE	6D	Divide Floating Point
DER	2D	Divide Floating Point Register
DH	4D	Divide Halfword
DHR	0D	Divide Halfword Register
DR	1D	Divide Register
EPSR	95	Exchange Program Status Register
EXBR	94	Exchange Byte Register
EXHR	34	Exchange Halfword Register
FLR	2F	Float Register
FLDR	3F	Float Register Double Precision
FXDR	3E	Fix Register Double Precision Floating Point
FXR	2E	Fix Register
ISRST	DF6	Interruptible State Restore
ISSV	DF5	Interruptible State Save
LA	E6	Load Address
LB	D3	Load Byte
LBR	93	Load Byte Register
LCDR	37	Load Complement Double Floating Register
LCER	17	Load Complement Floating Point Register
LCS	25	Load Complement Short
LD	78	Load Double Precision Floating Point
LDE	87	Load Double Floating Point From Single
LDER	A7	Load Double From Single Register
LDGR	A6	Load Double From General Register
LDPS	DF4	Load Process State
LDR	38	Load Double Precision Register
LE	68	Load Floating Point
LED	84	Load Floating From Double Precision
LEDR	A4	Load Floating From Double Register
LEGR	A5	Load Floating From General Register
LER	28	Load Floating Point Register
LH	48	Load Halfword
LHI	C8	Load Halfword Immediate

MNEMONIC	OPCODE	INSTRUCTION
LHL	73	Load Halfword Logical
LI	F8	Load Immediate
LIS	24	Load Immediate Short
LM	D1	Load Multiple
LMD	7F	Load Multiple Double Precision Floating Point
LME	72	Load Multiple Floating Point
LPB	6F	Load Packed Binary
LPDR	33	Load Positive Double Floating Register
LPER	13	Load Positive Floating Register
LPSTD	DF1	Load Process Segment Table Description
LPSW	C2	Load Program Status Word
LPSWR	18	Load Program Status Word Register
LR	08	Load Register
LRA	63	Load Real Address
LSSTD	DF2	Load Shared Segment Table Descriptor
LU	4E	Load Unnormalized Floating Point
LUR	1E	Load Unnormalized Floating Point Register
LW	4F	Load Unnormalized Double Precision Floating Point
LWR	1F	Load Unnormalized Double Precision Floating Point Register
M	5C	Multiply
MD	7C	Multiply Double Floating Point
MDR	3C	Multiply Double Floating Register
ME	6C	Multiply Floating Point
MER	2C	Multiply Floating Point Register
MR	4C	Multiply Halfword
MHR	0C	Multiply Halfword Register
MOVE	8C/01	Move
MOVEP	8C/21	Move and Pad
MR	1C	Multiply Register
N	54	AND
NH	44	AND Halfword
NHI	C4	AND Halfword Immediate
NI	F4	AND Immediate
NOP	420	No Operation
NOPR	020	No Operation Register
NR	04	AND Register
O	56	OR
OC	DE	Output Command
OCR	9E	Output Command Register
OH	46	OR Halfword
OHI	C6	OR Halfword Immediate
OI	F6	OR Immediate
OR	06	OR Register
PB	62	Process Byte
PBR	32	Process Byte Register
PMV	8C/03	Pack and Move
PMVA	8C/23	Pack and Move Absolute

MNEMONIC	OPCODE	INSTRUCTION
RBL	67	Remove from Bottom of List
RBT	76	Reset Bit
RD	DB	Read Data
RDR	9B	Read Data Register
REL	DF0	Read Error Logger
RH	D9	Read Halfword
RHR	99	Read Halfword Register
RLL	EB	Rotate Left Logical
RMVF	DF8	Reset Memory Voltage Fault
RRL	EA	Rotate Right Logical
RTL	66	Remove from Top of List
S	5B	Subtract
SBT	75	Set Bit
SCP	E3	Simulate Channel Program
SD	7B	Subtract Double Precision Floating Point
SDR	3B	Subtract Register Double Precision Floating Point
SE	6B	Subtract Floating Point
SER	2B	Subtract Floating Point Register
SH	4B	Subtract Halfword
SHI	CB	Subtract Halfword Immediate
SI	FB	Subtract Immediate
SINT	E2	Simulate Interrupt
SIS	27	Subtract Immediate Short
SLA	EF	Shift Left Arithmetic
SLHA	CF	Shift Left Halfword Arithmetic
SLHL	CD	Shift Left Halfword Logical
SLHL	91	Shift Left Halfword Logical Short
SLL	ED	Shift Left Logical
SLLS	11	Shift Left Logical Short
SR	0B	Subtract Register
SRA	EE	Shift Right Arithmetic
SRHA	CE	Shift Right Halfword Arithmetic
SRHL	CC	Shift Right Halfword Logical
SRHLS	90	Shift Right Halfword Logical Short
SRL	EC	Shift Right Logical
SRLS	10	Shift Right Logical Short
SS	DD	Sense Status
SSR	9D	Sense Status Register
ST	50	Store
STB	D2	Store Byte
STBP	6E	Store Binary as Packed
STBR	92	Store Byte Register
STD	70	Store Double Precision Floating Point
STDE	82	Store Double Precision in Single Precision
STE	60	Store Floating Point
STH	40	Store Halfword
STM	D0	Store Multiple
STMD	7E	Store Multiple Double Precision Floating Point
STME	71	Store Multiple Floating Point
STPS	DF3	Save Process State
SVC	E1	Supervisor Call

MNEMONIC	OPCODE	INSTRUCTION
TBT	74	Test Bit
THI	C3	Test Halfword Immediate
TI	F3	Test Immediate
TLATE	E7	Translate
TS	E0	Test and Set
UMV	8C/04	Unpack and Move
UMVA	8C/24	Unpack and Move Absolute
WD	DA	Write Data
WDR	9A	Write Data Register
WH	D8	Write Halfword
WHR	98	Write Halfword Register
X	57	Exclusive-OR
XH	47	Exclusive-OR Halfword
XHI	C7	Exclusive-OR Halfword Immediate
XI	F7	Exclusive-OR Immediate
XR	07	Exclusive-OR Register
XSTB	DF7	Store Byte, No ECC

APPENDIX C
INSTRUCTION SUMMARY - NUMERICAL BY OPCODE

OPCODE	MNEMONIC	INSTRUCTION
01*	BALR	Branch and Link Register
02*	BTCR	Branch on True Condition Register
03*	BFCR	Branch on False Condition Register
04	NR	AND Register
05	CLR	Compare Logical Register
06	OR	OR Register
07	XR	Exclusive OR Register
08	LR	Load Register
09	CR	Compare Register
0A	AR	Add Register
0B	SR	Subtract Register
0C*	MHR	Multiply Halfword Register
0D*	DHR	Divide Halfword Register
10	SRLS	Shift Right Logical Short
11	SLLS	Shift Left Logical Short
12	CHVR	Convert to Halfword Value Register
13+	LPER	Load Positive Floating Point
15+	LGFR	Load General Register from Floating
16+	LGDR	Load General from Double Floating
17+	LCFR	Load Complement Floating Register
18	LPSWR	Load Program Status Word Register
1C*	MR	Multiply Register
1D*	DR	Divide Register
1E+	LUR	Load Unnormalized Floating Point Register
1F+	LWR	Load Unnormalized Double Precision Floating Point
20*	BTBS	Branch on True Condition Backward Short
21*	BTFS	Branch on True Condition Forward Short
22*	BFBS	Branch on False Condition Backward Short
23*	BFFS	Branch on False Condition Forward Short
24	LIS	Load Immediate Short
25	LCS	Load Complement Short
26	AIS	Add Immediate Short
27	SIS	Subtract Immediate Short
28+	LER	Load
29+	CER	Compare Floating Point

* Does not change condition code

+ Floating point instruction

OPCODE	MNEMONIC	INSTRUCTION
2A+	AER	Add Floating Point Register
2B+	SER	Subtract Floating Point Register
2C+	MER	Multiply Floating Point Register
2D+	DER	Divide Floating Point Register
2E+	FXR	Fix Register
2F+	FLR	Float Register
32*	PBR	Process Byte Register
33+	LPDR	Load Positive Double Register
34*	EXHR	Exchange Halfword Register
37+	LCDR	Load Complement Double Register
38+	LDR	Load Register Double Precision Floating Point
39+	CDR	Compare Register Double Precision Floating Point
3A+	ADR	Add Register Double Precision Floating Point
3B+	SDR	Subtract Register Double Precision Floating Point
3C+	MDR	Multiply Register Double Precision Floating Point
3D+	DDR	Divide Register Double Precision Floating Point
3E+	FXDR	Fix Register Double Precision Floating Point
3F+	FLDR	Float Register Double Precision Floating Point
40*	STH	Store Halfword
41*	BAL	Branch and Link
42*	BTC	Branch on True Condition
43*	BFC	Branch on False Condition
44	NH	AND Halfword
45	CLH	Compare Logical Halfword
46	OH	OR Halfword
47	XH	Exclusive-OR Halfword
48	LH	Load Halfword
49	CH	Compare Halfword
4A	AH	Add Halfword
4B	SH	Subtract Halfword
4C*	MH	Multiply Halfword
4D*	DH	Divide Halfword
4E+	LU	Load Unnormalized Floating Point
4F+	LW	Load Unnormalized Double Precision Floating Point
50*	ST	Store
51	AM	Add to Memory

* Does not change condition code
+ Floating point instruction

OPCODE	MNEMONIC	INSTRUCTION
54	N	AND
55	CL	Compare Logical
56	O	OR
57	X	Exclusive-OR
58	L	Load
59	C	Compare
5A	A	Add
5B	S	Subtract
5C*	M	Multiply
5D*	D	Divide
5E*	CRC12	Cyclic Redundancy Check Modulo 12
5F*	CRC16	Cyclic Redundancy Check Modulo 16
60*+	STE	Store Floating Point
61	AHM	Add Halfword to Memory
62*	PB	Process Byte
63	LRA	Load Read Address
64	ATL	Add to Top of List
65	ABL	Add to Bottom of List
66	RTL	Remove from Top of List
67	RBL	Remove from Bottom of List
68+	LE	Load Floating Point
69+	CE	Compare Floating Point
6A+	AE	Add Floating Point
6B+	SE	Subtract Floating Point
6C+	ME	Multiply Floating Point
6D+	DE	Divide Floating Point
6E	STBP	Store Binary as Packed
6F	LPB	Load Packed Binary
70*+	STD	Store Double Precision Floating Point
71*+	STME	Store Floating Point Multiple
72*+	LME	Load Floating Point Multiple
73	LHL	Load Halfword Logical
74	TBT	Test Bit
75	SBT	Set Bit
76	RBT	Reset Bit
77	CBT	Complement Bit
78+	LD	Load Double Precision Floating Point
79+	CD	Compare Double Precision Floating Point
7A+	AD	Add Double Precision Floating Point
7B+	SD	Subtract Double Precision Floating Point
7C+	MD	Multiply Double Precision Floating Point
7D+	DD	Divide Double Precision Floating Point
7E*+	STMD	Store Multiple Double Precision Floating Point
7F*+	LMD	Load Multiple Double Precision Floating Point

* Does not change condition code

+ Floating point instruction

OPCODE	MNEMONIC	INSTRUCTION
82*+	STDE	Store Double Precision in Single
84+	LED	Load Floating from Double Precision
87+	LDE	Load Double from Floating Point
88*	BRK	Breakpoint
8C	(RXX)	RXX class designator
8C/00	MVTU	Move Translated Until
8C/01	MOVE	Move
8C/02	CPAN	Compare Alphanumeric
8C/03	PMV	Pack and Move
8C/04	UMV	Unpack and Move
8C/21	MOVEP	Move and Pad
8C/22	CPANP	Compare Alphanumeric and Pad
8C/23	PMVA	Pack and Move Absolute
8C/24	UMVA	Unpack and Move Absolute
90	SRHLS	Shift Right Halfword Logical Short
91	SLHLS	Shift Left Halfword Logical Short
92*	STBR	Store Byte Register
93*	LBR	Load Byte Register
94*	EXBR	Exchange Byte Register
95	EPSR	Exchange Program Status Word
98	WHR	Write Halfword Register
99	RHR	Read Halfword Register
9A	WDR	Write Data Register
9B	RDR	Read Data Register
9D	SSR	Sense Status Register
9E	OCR	Output Command Register
A4+	LEDR	Load Floating from Double Register
A5+	LEGR	Load Floating from General Register
A6+	LDGR	Load Double from General Register
A7+	LDER	Load Double from Floating Register
C0*	BXH	Branch on Index High
C1*	BXLE	Branch on Index Low or Equal
C2	LPSW	Load Program Status Word
C3	THI	Test Halfword Immediate
C4	NHI	AND Halfword Immediate
C5	CLHI	Compare Logical Halfword Immediate
C6	OHI	OR Halfword Immediate
C7	XHI	Exclusive-OR Halfword Immediate
C8	LHI	Load Halfword Immediate
C9	CHI	Compare Halfword Immediate

* Does not change condition code
+ Floating point instruction

OPCODE	MNEMONIC	INSTRUCTION
CA	AHI	Add Halfword Immediate
CB	SHI	Subtract Halfword Immediate
CC	SRHL	Shift Right Halfword Logical
CD	SLHL	Shift Left Halfword Logical
CE	SRHA	Shift Right Halfword Arithmetic
CF	SLHA	Shift Left Halfword Arithmetic
D0*	STM	Store Multiple
D1*	LM	Load Multiple
D2*	STB	Store Byte
D3*	LB	Load Byte
D4	CLB	Compare Logical Byte
D5	AL	Autoload
D8	WH	Write Halfword
D9	RH	Read Halfword
DA	WD	Write Data
DB	RD	Read Data
DD	SS	Sense Status
DE	OC	Output Command
DF	(PSF)	PSF Class Designator
DF0	REL	Read Error Logger
DF1*	LPSTD	Load Process Segment Table Descriptor
DF2*	LSSTD	Load Shared Segment Table Descriptor
DF3*	STPS	Save Process State
DF4	LDPS	Load Process State
DF5*	ISSV	Interruptible State Save
DF6*	ISRST	Interruptible State Restore
DF7*	XSTB	Store Byte, No ECC
DF8*	RMVF	Reset Memory Voltage Fault
E0	TS	Test and Set
E1	SVC	Supervisor Call
E2	SINT	Simulate Interrupt
E3	SCP	Simulate Channel Program
E6*	LA	Load Address
E7*	TLATE	Translate
EA	RRL	Rotate Right Logical
EB	RLL	Rotate Left Logical
EC	SRL	Shift Right Logical
ED	SLL	Shift Left Logical
EE	SRA	Shift Right Arithmetic
EF	SLA	Shift Left Arithmetic

- * Does not change condition code
- + Floating point instruction

OPCODE	MNEMONIC	INSTRUCTION
F3	TI	Test Immediate
F4	NI	AND Immediate
F5	CLI	Compare Logical Immediate
F6	OI	OR Immediate
F7	XI	Exclusive-OR Immediate
F8	LI	Load Immediate
F9	CI	Compare Immediate
FA	AI	Add Immediate
FB	SI	Subtract Immediate

- * Does not change condition code
- + Floating point instruction

**APPENDIX D
ARITHMETIC REFERENCES**

TABLE D-1 POWERS OF TWO

637-1

2^n	n	2^{-n}																								
1	0	1.0																								
2	1	0.5																								
4	2	0.25																								
8	3	0.125																								
16	4	0.062	5																							
32	5	0.031	25																							
64	6	0.015	625																							
128	7	0.007	812	5																						
256	8	0.003	906	25																						
512	9	0.001	953	125																						
1 024	10	0.000	976	562	5																					
2 048	11	0.000	488	281	25																					
4 096	12	0.000	244	140	625																					
8 192	13	0.000	122	070	312	5																				
16 384	14	0.000	061	035	156	25																				
32 768	15	0.000	030	517	578	125																				
65 536	16	0.000	015	258	789	062	5																			
131 072	17	0.000	007	629	394	531	25																			
262 144	18	0.000	003	814	697	265	625																			
524 288	19	0.000	001	907	348	632	812	5																		
1 048	576	20	0.000	000	953	674	316	406	25																	
2 097	152	21	0.000	000	476	837	158	203	125																	
4 194	304	22	0.000	000	238	418	579	101	562	5																
8 388	608	23	0.000	000	119	209	289	550	781	25																
16 777	216	24	0.000	000	059	604	644	775	390	625																
33 554	432	25	0.000	000	029	802	322	387	695	312	5															
67 108	864	26	0.000	000	014	901	161	193	847	656	25															
134 217	728	27	0.000	000	007	450	580	596	923	828	125															
268 435	456	28	0.000	000	003	725	290	298	461	914	062	5														
536 870	912	29	0.000	000	001	862	645	149	230	957	031	25														
1 073	741	824	30	0.000	000	000	931	322	574	615	478	515	625													
2 147	483	648	31	0.000	000	000	465	661	287	307	739	257	812	5												
4 294	967	296	32	0.000	000	000	232	830	643	653	869	628	906	25												
8 589	934	592	33	0.000	000	000	116	415	321	826	934	814	453	125												
17 179	869	184	34	0.000	000	000	058	207	660	913	467	407	226	562	5											
34 359	738	368	35	0.000	000	000	029	103	830	456	733	703	613	281	25											
68 719	476	736	36	0.000	000	000	014	551	915	228	366	851	806	640	625											
137 438	953	472	37	0.000	000	000	007	275	957	614	183	425	903	320	312	5										
274 877	906	944	38	0.000	000	000	003	637	978	807	091	712	951	660	156	25										
549 755	813	888	39	0.000	000	000	001	818	989	403	545	856	475	830	078	125										
1 099	511	627	776	40	0.000	000	000	000	909	494	701	772	928	237	915	039	062	5								

TABLE D-2 POWERS OF SIXTEEN

638

16^n						n
					1	0
					16	1
					256	2
				4	096	3
				65	536	4
			1	048	576	5
			16	777	216	6
			268	435	456	7
		4	294	967	296	8
		68	719	476	736	9
	1	099	511	627	776	10
	17	592	186	044	416	11
	281	474	976	710	656	12
	4	503	599	627	370	13
	72	057	594	037	927	14
1	152	921	504	606	846	15

DECIMAL VALUES

TABLE D-3 HEXADECIMAL ADDITION AND SUBTRACTION

EXAMPLES: 5 + A = F; 18 - D = B; A + B = 15

639

	1	2	3	4	5	6	7	8	9	A	B	C	D	E	F	
1	2	3	4	5	6	7	8	9	A	B	C	D	E	F	10	1
2	3	4	5	6	7	8	9	A	B	C	D	E	F	10	11	2
3	4	5	6	7	8	9	A	B	C	D	E	F	10	11	12	3
4	5	6	7	8	9	A	B	C	D	E	F	10	11	12	13	4
5	6	7	8	9	A	B	C	D	E	F	10	11	12	13	14	5
6	7	8	9	A	B	C	D	E	F	10	11	12	13	14	15	6
7	8	9	A	B	C	D	E	F	10	11	12	13	14	15	16	7
8	9	A	B	C	D	E	F	10	11	12	13	14	15	16	17	8
9	A	B	C	D	E	F	10	11	12	13	14	15	16	17	18	9
A	B	C	D	E	F	10	11	12	13	14	15	16	17	18	19	A
B	C	D	E	F	10	11	12	13	14	15	16	17	18	19	1A	B
C	D	E	F	10	11	12	13	14	15	16	17	18	19	1A	1B	C
D	E	F	10	11	12	13	14	15	16	17	18	19	1A	1B	1C	D
E	F	10	11	12	13	14	15	16	17	18	19	1A	1B	1C	1D	E
F	10	11	12	13	14	15	16	17	18	19	1A	1B	1C	1D	1E	F
	1	2	3	4	5	6	7	8	9	A	B	C	D	E	F	

TABLE D-4 HEXADECIMAL MULTIPLICATION AND DIVISION

EXAMPLES: $5 \times 6 = 1E$; $75 \div D = 9$; $58 \div 8 = B$; $9 \times C = 6C$

640

	1	2	3	4	5	6	7	8	9	A	B	C	D	E	F	
1	1	2	3	4	5	6	7	8	9	A	B	C	D	E	F	1
2	2	4	6	8	A	C	E	10	12	14	16	18	1A	1C	1E	2
3	3	6	9	C	F	12	15	18	1B	1E	21	24	27	2A	2D	3
4	4	8	C	10	14	18	1C	20	24	28	2C	30	34	38	3C	4
5	5	A	F	14	19	1E	23	28	2D	32	37	3C	41	46	4B	5
6	6	C	12	18	1E	24	2A	30	36	3C	42	48	4E	54	5A	6
7	7	E	15	1C	23	2A	31	38	3F	46	4D	54	5B	62	69	7
8	8	10	18	20	28	30	38	40	48	50	58	60	68	70	78	8
9	9	12	1B	24	2D	36	3F	48	51	5A	63	6C	75	7E	87	9
A	A	14	1E	28	32	3C	46	50	5A	64	6E	78	82	8C	96	A
B	B	16	21	2C	37	42	4D	58	63	6E	79	84	8F	9A	A5	B
C	C	18	24	30	3C	48	54	60	6C	78	84	90	9C	A8	B4	C
D	D	1A	27	34	41	4E	5B	68	75	82	8F	9C	A9	B6	C3	D
E	E	1C	2A	38	46	54	62	70	7E	8C	9A	A8	B6	C4	D2	E
F	F	1E	2D	3C	4B	5A	69	78	87	96	A5	B4	C3	D2	E1	F
	1	2	3	4	5	6	7	8	9	A	B	C	D	E	F	

TABLE D-5 MATHEMATICAL CONSTANTS

641-2

CONSTANT	DECIMAL VALUE				HEXADECIMAL VALUE		FLOATING POINT VALUE			
							DOUBLE PRECISION			
							SINGLE PRECISION			
π	3.14159	26535	89793	23846	3.243F	6A89	4132	43F6	A888	5A31
$\pi-1$	03.1830	98861	83790	67153	0.517C	C1B7	4051	7CC1	B727	220B
$\sqrt{\pi}$	1.77245	38509	05516	02729	1.C5BF	891C	411C	5BF8	91B4	EF6B
$\text{Ln } \pi$	1.4472	98858	49400	17414	1.250D	048F	4112	50D0	48E7	A1BD
$\sqrt{3}$	1.73205	08075	68877	29353	1.B67A	E858	411B	67AE	8584	CAA7
e	2.71828	18284	59045	23536	2.B7E1	5163	412B	7E15	1628	AED3
e^{-1}	0.36787	94411	71442	32159	0.5E2D	58D9	405E	2D58	D8B3	BCDF
\sqrt{e}	1.64872	12707	00128	14680	1.A612	98E2	411A	6129	8E1E	069C
$\log_{10} e$	0.43429	44819	03251	82765	0.6F2D	EC55	406F	2DEC	549B	9439
$\log_2 e$	1.44269	50408	88963		1.7154	7653	4117	1547	652	----
γ	0.57721	56649	01532	86060	0.93C4	67E4	4093	C467	E37D	B0C8
$\text{Ln } \gamma$	-0.54953	93129	81644	82233	-0.8CAE	9BC1	C08C	AE9B	C11F	5A60
$\sqrt{2}$	1.41421	35623	73095	04880	1.6A09	E668	4116	A09E	667F	3BCD
$\text{Ln} 2$	0.69314	71805	59945	30941	0.B172	17F8	40B1	7217	F7D1	CF7A
$\log_{10} 2$	0.30102	99956	63981	19521	0.4D10	4D42	404D	104D	427D	E7FC
$\sqrt{10}$	3.16227	76601	68379	33200	3.298B	075C	4132	98B0	75B4	B6A5
$\text{Ln} 10$	2.30258	50929	04945	68401	2.4D76	3777	4124	D763	776A	AA2B

TABLE D-6 FRACTION CONVERSION

642

Hexadecimal and Decimal Fraction Conversion Table

HALFWORD													
BYTE 0						BYTE 1							
BITS		4567				0123		4567					
Hex	Decimal	Hex	Decimal			Hex	Decimal			Hex	Decimal Equivalent		
.0	.0000	.00	.0000	0000	.000	.0000	0000	0000	.0000	.0000	0000	0000	0000
.1	.0625	.01	.0039	0625	.001	.0002	4414	0625	.0001	.0000	1525	8789	0625
.2	.1250	.02	.0078	1250	.002	.0004	8828	1250	.0002	.0000	3051	7578	1250
.3	.1875	.03	.0117	1875	.003	.0007	3242	1875	.0003	.0000	4577	6367	1875
.4	.2500	.04	.0156	2500	.004	.0009	7656	2500	.0004	.0000	6103	5156	2500
.5	.3125	.05	.0195	3125	.005	.0012	2070	3125	.0005	.0000	7629	3945	3125
.6	.3750	.06	.0234	3750	.006	.0014	6484	3750	.0006	.0000	9155	2734	3750
.7	.4375	.07	.0273	4375	.007	.0017	0898	4375	.0007	.0001	0681	1523	4375
.8	.5000	.08	.0312	5000	.008	.0019	5312	5000	.0008	.0001	2207	0312	5000
.9	.5625	.09	.0351	5625	.009	.0021	9726	5625	.0009	.0001	3732	9101	5625
.A	.6250	.0A	.0390	6250	.00A	.0024	4140	6250	.000A	.0001	5258	7890	6250
.B	.6875	.0B	.0429	6875	.00B	.0026	8554	6875	.000B	.0001	6784	6679	6875
.C	.7500	.0C	.0468	7500	.00C	.0029	2968	7500	.000C	.0001	8310	5468	7500
.D	.8125	.0D	.0507	8125	.00D	.0031	7382	8125	.000D	.0001	9836	4257	8125
.E	.8750	.0E	.0546	8750	.00E	.0034	1796	8750	.000E	.0002	1362	3046	8750
.F	.9375	.0F	.0585	9375	.00F	.0036	6210	9375	.000F	.0002	2888	1835	9375
1	2		3		4								

TO CONVERT .ABC HEXADECIMAL TO DECIMAL

- Find .A in position 1 .6250
- Find .0B in position 2 .0429 6875
- Find .00C in position 3 .0029 2968 7500
- .ABC Hex is equal to .6708 9843 7500

TO CONVERT .13 DECIMAL TO HEXADECIMAL

1. Find .1250 next lowest to subtract

$$\begin{array}{r} .1300 \\ - .1250 \\ \hline \end{array} = .2 \text{ Hex}$$
2. Find .0039 0625 next lowest to subtract

$$\begin{array}{r} .0050 \ 0000 \\ - .0039 \ 0625 \\ \hline \end{array} = .01$$
3. Find .0009 7656 2500

$$\begin{array}{r} .0010 \ 9375 \ 0000 \\ - .0009 \ 7656 \ 2500 \\ \hline \end{array} = .004$$
4. Find .0001 0681 1523 4375

$$\begin{array}{r} .0001 \ 1718 \ 7500 \ 0000 \\ - .0001 \ 0681 \ 1523 \ 4375 \\ \hline \end{array} = .0007$$

$$\begin{array}{r} .0000 \ 1037 \ 5976 \ 5625 \\ \hline \end{array} = .2147 \text{ Hex}$$
5. 13 Decimal is approximately equal to $\xrightarrow{\hspace{10em}}$

To convert fractions beyond the capacity of table, use techniques below:

HEXADECIMAL FRACTION TO DECIMAL

Convert the hexadecimal fraction to its decimal equivalent using the same technique as for integer numbers. Divide the results by 16^n (n is the number of fraction positions).
 Example: $.8A7_{16} = .540771_{10}$

$$\begin{array}{r} 8A7_{16} = 2215_{10} \\ 16^3 = 4096 \quad \underline{4096} \ 2215.000000 \\ \hline \end{array} \begin{array}{l} .540771 \\ \hline \end{array}$$

DECIMAL FRACTION TO HEXADECIMAL

Collect integer parts of product in the order of calculation.

Example: $.5408_{10} = .8A7_{16}$

$$\begin{array}{r} .5408 \\ \times 16 \\ \hline 8 \leftarrow \boxed{8} \ 6528 \\ \times 16 \\ \hline A \leftarrow \boxed{A} \ 4448 \\ \times 16 \\ \hline 7 \leftarrow \boxed{7} \ 1168 \end{array}$$

TABLE D-7 INTEGER CONVERSION

643

Hexadecimal and Decimal Integer Conversion Table

HALFWORD								HALFWORD							
BYTE 0				BYTE 1				BYTE 2				BYTE 3			
BITS: 0123		4567		0123		4567		0123		4567		0123		4567	
Hex	Decimal	Hex	Decimal	Hex	Decimal	Hex	Decimal	Hex	Decimal	Hex	Decimal	Hex	Decimal	Hex	Decimal
0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0
1	268,435,456	1	16,777,216	1	1,048,576	1	65,536	1	4,096	1	256	1	16	1	1
2	536,870,912	2	33,554,432	2	2,097,152	2	131,072	2	8,192	2	512	2	32	2	2
3	805,306,368	3	50,331,648	3	3,145,728	3	196,608	3	12,288	3	768	3	48	3	3
4	1,073,741,824	4	67,108,864	4	4,194,304	4	262,144	4	16,384	4	1,024	4	64	4	4
5	1,342,177,280	5	83,886,080	5	5,242,880	5	327,680	5	20,480	5	1,280	5	80	5	5
6	1,610,612,736	6	100,663,296	6	6,291,456	6	393,216	6	24,576	6	1,536	6	96	6	6
7	1,879,048,192	7	117,440,512	7	7,340,032	7	458,752	7	28,672	7	1,792	7	112	7	7
8	2,147,483,648	8	134,217,728	8	8,388,608	8	524,288	8	32,768	8	2,048	8	128	8	8
9	2,415,919,104	9	150,994,944	9	9,437,184	9	589,824	9	36,864	9	2,304	9	144	9	9
A	2,684,354,560	A	167,772,160	A	10,485,760	A	655,360	A	40,960	A	2,560	A	160	A	10
B	2,952,790,016	B	184,549,376	B	11,534,336	B	720,896	B	45,056	B	2,816	B	176	B	11
C	3,221,225,472	C	201,326,592	C	12,582,912	C	786,432	C	49,152	C	3,072	C	192	C	12
D	3,489,660,928	D	218,103,808	D	13,631,488	D	851,968	D	53,248	D	3,328	D	208	D	13
E	3,758,096,384	E	234,881,024	E	14,680,064	E	917,504	E	57,344	E	3,584	E	224	E	14
F	4,026,531,840	F	251,658,240	F	15,728,640	F	983,040	F	61,440	F	3,840	F	240	F	15
B		7		6		5		4		3		2		1	

TO CONVERT HEXADECIMAL TO DECIMAL

1. Locate the column of decimal numbers corresponding to the left-most digit or letter of the hexadecimal; select from this column and record the number that corresponds to the position of the hexadecimal digit or letter.
2. Repeat step 1 for the next (second from the left) position.
3. Repeat step 1 for the units (third from the left) position.
4. Add the numbers selected from the table to form the decimal number.

EXAMPLE	
Conversion of Hexadecimal Value	D34
1. D	3328
2. 3	48
3. 4	4
4. Decimal	3380

To convert integer numbers greater than the capacity of table, use the techniques below:

HEXADECIMAL TO DECIMAL

Successive cumulative multiplication from left to right, adding units position.

Example: $D34_{16} = 3380_{10}$ $D = 13$

$$\begin{array}{r} \times 16 \\ 208 \\ 3 = + 3 \\ 211 \\ \times 16 \\ 3376 \\ 4 = + 4 \\ 3380 \end{array}$$

TO CONVERT DECIMAL TO HEXADECIMAL

1. (a) Select from the table the highest decimal number that is equal to or less than the number to be converted.
(b) Record the hexadecimal of the column containing the selected number.
(c) Subtract the selected decimal from the number to be converted.
2. Using the remainder from step 1(c) repeat all of step 1 to develop the second position of the hexadecimal (and a remainder).
3. Using the remainder from step 2 repeat all of step 1 to develop the units position of the hexadecimal.
4. Combine terms to form the hexadecimal number.

EXAMPLE	
Conversion of Decimal Value	3380
1. D	<u>-3328</u> 52
2. 3	<u>-48</u> 4
3. 4	<u>-4</u> 0
4. Hexadecimal	D34

DECIMAL TO HEXADECIMAL

Divide and collect the remainder in reverse order.

Example: $3380_{10} = X_{16}$

$$\begin{array}{r} 16 \overline{) 3380} \\ 16 \overline{) 211} \\ 16 \overline{) 13} \end{array} \begin{array}{l} \text{remainder} \\ \uparrow \\ 4 \\ 3 \\ D \end{array} \quad 3380_{10} = D34_{16}$$

APPENDIX E
INPUT/OUTPUT (I/O) REFERENCES

TABLE E-1 ASCII/HEXADECIMAL CONVERSION

644-1

BITS				b ₆ b ₅ b ₄	0 0 0	0 0 1	0 1 0	0 1 1	1 0 0	1 0 1	1 1 0	1 1 1
b ₃	b ₂	b ₁	b ₀	MSD LSD	0	1	2	3	4	5	6	7
0	0	0	0	0	NUL	DLE	SP	0	@	P	`	p
0	0	0	1	1	SOH	DC1	!	1	A	Q	a	q
0	0	1	0	2	STX	DC2	"	2	B	R	b	r
0	0	1	1	3	ETX	DC3	#	3	C	S	c	s
0	1	0	0	4	EOT	DC4	\$	4	D	T	d	t
0	1	0	1	5	ENQ	NAK	%	5	E	U	e	u
0	1	1	0	6	ACK	SYN	&	6	F	V	f	v
0	1	1	1	7	BEL	ETB	'	7	G	W	g	w
1	0	0	0	8	BS	CAN	(8	H	X	h	x
1	0	0	1	9	HT	EM)	9	I	Y	i	y
1	0	1	0	A	LF	SUB	*	:	J	Z	j	z
1	0	1	1	B	VT	ESC	+	;	K	[k	{
1	1	0	0	C	FF	FS	,	<	L	\	l	
1	1	0	1	D	CR	GS	-	=	M]	m	}
1	1	1	0	E	SO	RS	.	>	N	^	n	~
1	1	1	1	F	SI	US	/	?	O	_	o	DEL

NUL	Null	DLE	Data link escape
SOH	Start of heading	DC1-4	Device control
STX	Start of text	NAK	Negative acknowledge
ETX	End of text	SYN	Synchronous idle
EOT	End of transmission	ETB	End of transmission block
ENQ	Enquiry	CAN	Cancel
ACK	Acknowledge	EM	End of medium
BEL	Audible signal	SUB	Start of special sequence
BS	Backspace	ESC	Escape
HT	Horizontal tabulation	FS	File separator
LF	Line feed	GS	Group separator
VT	Vertical tabulation	RS	Record separator
FF	Form feed	US	Unit separator
CR	Carrier return	SP	Space
SO	Shift out	DEL	Delete/Idle
SI	Shift in		

TABLE E-2 ASCII/CARD CODE CONVERSION

645-1

GRAPHIC	7-BIT ASCII CODE	CARD CODE	GRAPHIC	7-BIT ASCII CODE	CARD CODE
SPACE	20	BLANK	@	40	8-4
!	21	11-8-2	A	41	12-1
"	22	8-7	B	42	12-2
#	23	8-3	C	43	12-3
\$	24	11-8-3	D	44	12-4
%	25	0-8-4	E	45	12-5
&	26	12	F	46	12-6
'	27	8-5	G	47	12-7
(28	12-8-5	H	48	12-8
)	29	11-8-5	I	49	12-9
*	2A	11-8-4	J	4A	11-1
+	2B	12-8-6	K	4B	11-2
,	2C	0-8-3	L	4C	11-3
-	2D	11	M	4D	11-4
.	2E	12-8-3	N	4E	11-5
/	2F	0-1	O	4F	11-6
0	30	0	P	50	11-7
1	31	1	Q	51	11-8
2	32	2	R	52	11-9
3	33	3	S	53	0-2
4	34	4	T	54	0-3
5	35	5	U	55	0-4
6	36	6	V	56	0-5
7	37	7	W	57	0-6
8	38	8	X	58	0-7
9	39	9	Y	59	0-8
:	3A	8-2	Z	5A	0-9
;	3B	11-8-6	[5B	12-8-2
<	3C	12-8-4	\	5C	0-8-2
=	3D	8-6]	5D	12-8-7
>	3E	0-8-6	↑	5E	11-8-7
?	3F	0-8-7	←	5F	0-8-5

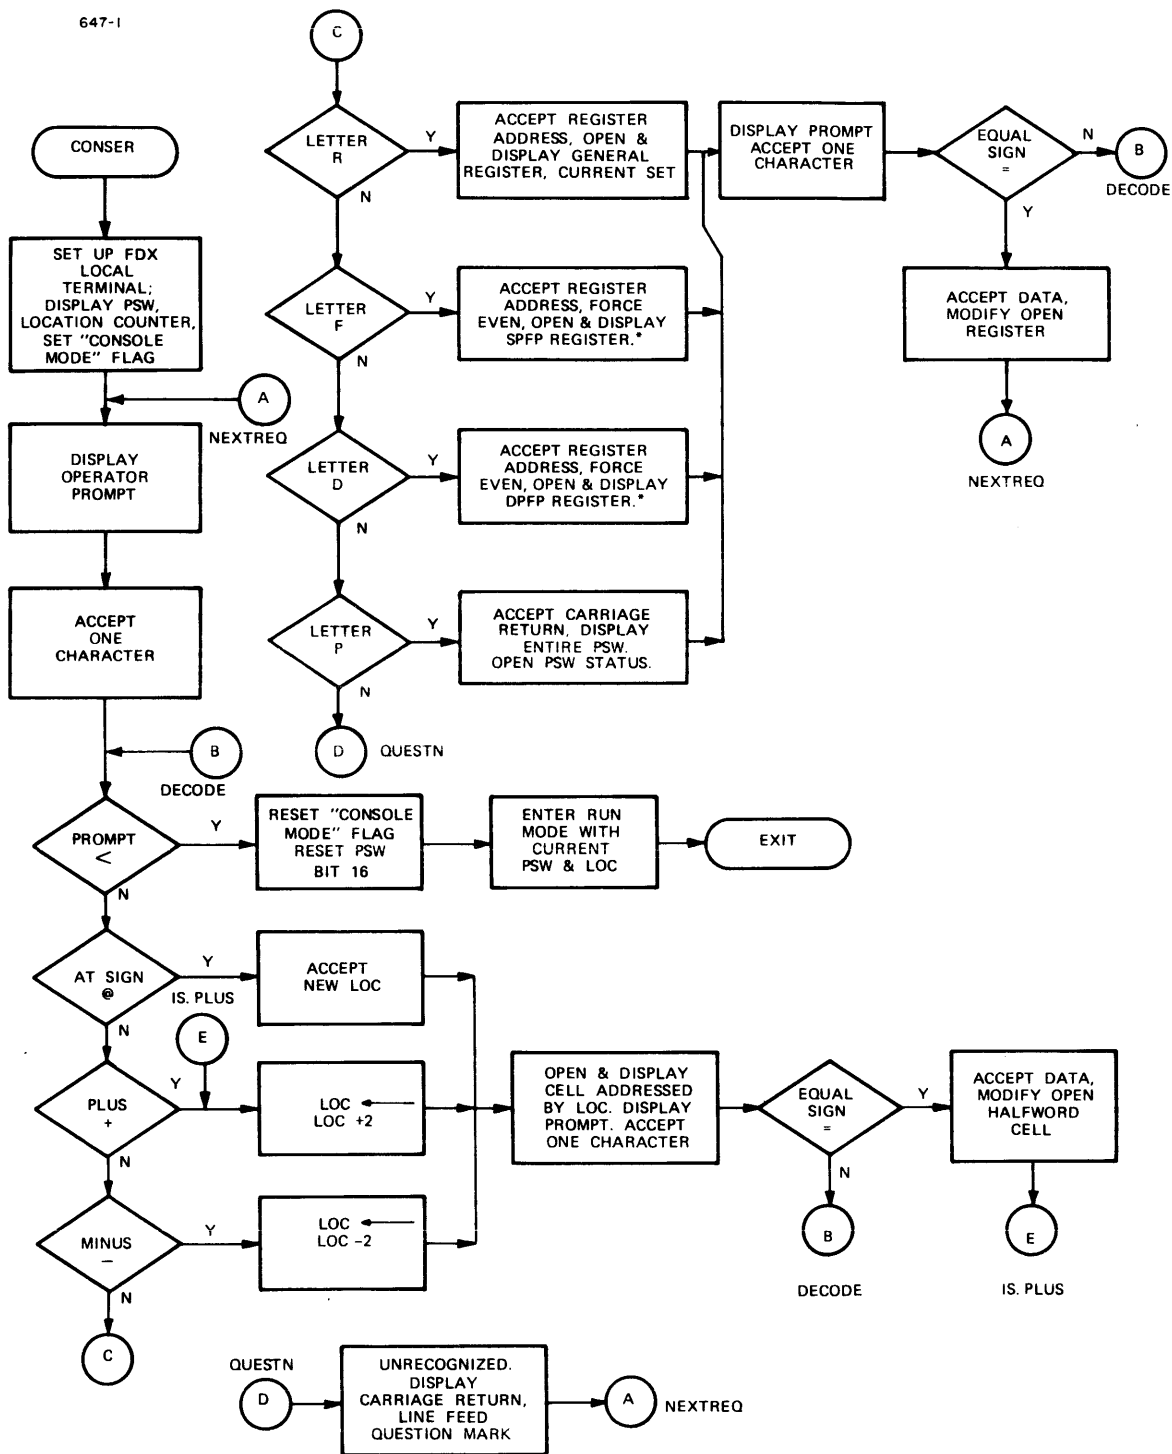
TABLE E-3 STANDARD-PREFERRED ADDRESS TABLE

646-1	0	1	2	3	4	5	6	7	8	9	A	B	C	D	E	F
MSD 0		RESERVED	TTY CAROUSEL 15, 30 CRT ON CLI		CARD READER	LOADER STORAGE UNIT	RESERVED		MDIO						201/301 DATA SET HDX	201/301 DATA SET FDX
1	COMM MUX															
2	8-LINE INTERRUPT MODULE (ADRS 20 TO 27)							SECOND 8-LINE INTERRUPT MODULE (ADRS 28 TO 2F)								
3	CONTACT CLOSURE MODULE	I/O BUS SWITCH													360/370 AUX. INF	360/370 INF
4												DIGITAL MUX				
5																
6			LINE PRINTERS										UNIVERSAL CLOCK VARIABLE, 60 Hz			
7	RELAY DRIVER MODULE															801 DIALER
8				CONVERSION EQUIPMENT		556/800 BPI MAG TAPE			AIC			ULI				
9									AOC							
A									DIO							
B									REMOVABLE CARTRIDGE DISK CONT	QSA						
C	MICROBUS ADAPTER	FLOPPY DISK				1600 BPI MAG TAPE		DISK 0	FIXED DISK 0							
D								DISK 1	FIXED DISK 1							
E								DISK 2	FIXED DISK 2							
F	SELECTOR CHANNELS							DISK 3	FIXED DISK 3			MSM DISK SYSTEM	DRIVE 0	DRIVE 1	DRIVE 2	DRIVE 3

AIC = ANALOG INPUT CONTROLLER
 AOC = ANALOG OUTPUT CONTROLLER
 DIO = DIGITAL I/O CONTROLLER

QSA = QUAD SYNCHRONOUS ADAPTER
 ULI = UNIVERSAL LOGIC INTERFACE
 MDIO = MEMORY DISPATCHED I/O

APPENDIX F CONSOLE SERVICE ROUTINE FLOWCHART



NOTES:

1. ALL RECEIVED CHARACTERS ECHOED BY PROCESSOR.
2. LOWER-CASE CHARACTERS INTERPRETED AS UPPER-CASE.
3. SPACE CHARACTERS IGNORED.
4. BACKSPACE, UNDERLINE, DELETE CAUSE PREVIOUS NUMERIC CHARACTER TO BE IGNORED.
- * IF HARDWARE FLOATING POINT IS NOT AVAILABLE, DISPLAYS ARE ALL "F"s AS CONTENTS.

INDEX

A			
Access level bits (L)	11-12	Branch instruction formats	
Access level fault	11-18	programming examples	1-14
Access mode bits (A)	11-11	RI1	1-25
Access mode faults	11-18	RI2	1-27
Add (A)	5-6	RR	1-14
Add double precision			1-16
floating point (AD)	6-45	RX	1-14
Add floating point (AE)	6-24	RX1	1-18
Add floating point register		RX2	1-20
(AER)	6-24	RX3	1-23
Add halfword (AH)	5-8	RXRX	1-29
Add halfword immediate (AHI)	5-8	SF	1-14
Add halfword to memory (AHM)	5-12		1-17
Add immediate (AI)	5-6	Branch instructions	4-2
Add immediate short (AIS)	5-6	BAL	4-7
Add register (AR)	5-6	BALR	4-7
Add register double precision		BFBS	4-5
floating point (ADR)	6-45	BFC	4-5
Add to bottom of list (ABL)	3-58	BFCCR	4-5
Add to memory (AM)	5-10	BFPS	4-5
Add to top of list (ATL)	3-58	BTBS	4-3
Address space		BTC	4-3
offset and page field	11-7	BTCR	4-3
page field	11-7	BTFS	4-3
segment field	11-6	BXH	4-11
Alignment faults	10-19	BXLE	4-9
Alphanumeric		Branch on carry (BC)	4-15
byte string	1-9	Branch on carry register	
string data	1-11	(BCR)	4-15
string instruction		Branch on carry short (BCS)	4-15
formats	7-3	Branch on equal (BE)	4-17
AND (N)	3-30	Branch on equal register	
AND halfword	3-31	(BER)	4-17
AND halfword (NH)	3-31	Branch on equal short (BES)	4-17
AND immediate (NI)	3-30	Branch on false condition	
AND register (NR)	3-30	(BFC)	4-5
Arithmetic		Branch on false condition	
fault interrupt	10-32	backward short (BFBS)	4-5
operands	1-10	Branch on false condition	
references	D-1	forward short (BFPS)	4-5
Array		Branch on false condition	
bit	1-9	register (BFCCR)	4-5
byte	1-9	Branch on index high (BXH)	4-11
Auto driver channel	9-18	Branch on index low or equal	
CRC	1-2	(BXLE)	4-9
data buffer chaining	1-2	Branch on low (BL)	4-19
operation	10-28	Branch on low register (BLR)	4-19
programming procedure	9-24	Branch on low short (BLS)	4-19
Autoload (AL)	9-15	Branch on minus (BM)	4-21
		Branch on minus register	
		(BMR)	4-21
		Branch on minus short (BMS)	4-21
		Branch on no carry (BNC)	4-16
		Branch on no carry register	
		(BNCR)	4-16
		Branch on no carry short	
		(BNCS)	4-17
		Branch on no overflow (BNO)	4-26
		Branch on no overflow	
		register (BNOR)	4-26
B			
Bit array	1-9		
Branch (unconditional) (B)	4-29		
Branch and link (BAL)	4-7		
Branch and link register			
(BALR)	4-7		

Extended branch instructions
(Continued)

BNOR	4-26
BNOS	4-26
BNP	4-24
BNPR	4-24
BNPS	4-24
BNZ	4-28
BNZR	4-28
BNZS	4-28
BO	4-25
BOR	4-25
BOS	4-25
BP	4-23
BPR	4-23
BPS	4-23
BR	4-29
BS	4-29
BZ	4-27
BZR	4-27
BZS	4-27
NOP	4-30
NOPR	4-30

F

Faults	11-17
access level	11-18
access mode	11-18
MAT	11-17
nonpresence	11-19
precedence	11-21
PST or SST size exceeded	11-18
reexecution of faulting	11-19
instructions	11-17
segment limit fault	11-20
Fix register (FXR)	11-19
Fix register double	6-33
precision (FXDR)	6-54
Fixed point	
arithmetic	5-1
data	1-10
data formats	5-1
number	1-9
number range	5-2
operations	1-10
Fixed point instructions	5-2
A	5-4
AH	5-6
AHI	5-8
AHI	5-8
AHM	5-12
AI	5-6
AIS	5-6
AM	5-10
AR	5-6
C	5-18
CH	5-18
CHI	5-20
CHVR	5-39
CI	5-18
CR	5-18
D	5-26

Fixed point instructions
(Continued)

DH	5-30
DHR	5-30
DR	5-26
formats	5-4
M	5-22
MH	5-24
MHR	5-24
MR	5-22
S	5-14
SH	5-16
SHI	5-16
SI	5-14
SIS	5-14
SLA	5-33
SLHA	5-35
SR	5-14
SRA	5-36
SRHA	5-38
FLM	10-3
Float register (FLR)	6-35
Float register double	
precision (FLDR)	6-55
Floating point	
arithmetic	6-1
condition code	6-11
data	1-10
data formats	6-2
equalization	6-6
exponent overflow	6-8
exponent underflow	6-8
guard digits and R*	
rounding	6-9
instructions	6-1
normalization	6-5
number	1-9
number range	6-3
registers	6-4
true zero	1-7
Floating point instructions	6-7
AD	6-11
ADR	6-45
AE	6-45
AER	6-24
CD	6-24
CDR	6-49
CE	6-49
CE	6-28
CER	6-28
DD	6-52
DDR	6-52
DE	6-31
DER	6-31
FLDR	6-55
FLR	6-35
FXDR	6-54
FXR	6-33
LCDR	6-40
LCER	6-19
LD	6-38
LDER	6-38
LDGR	6-38
LDR	6-38
LDR	6-57
LDR	6-38
LDR	6-57
LE	6-15

Floating point instructions
(Continued)

LED	6-56
LEDR	6-56
LEGR	6-15
LER	6-15
LGDR	6-42
LGER	6-21
LMD	6-41
LME	6-20
LPDR	6-39
LPER	6-17
LU	6-14
LUR	6-14
LW	6-37
LWR	6-37
MD	6-50
MDR	6-50
ME	6-29
MER	6-29
SD	6-47
SDR	6-47
SE	6-26
SER	6-26
STD	6-43
STDE	6-58
STE	6-21
STMD	6-44
STME	6-23
Floating point masked mode. See FLM.	
Floating point underflow interrupt enable (FLU)	10-6
Flowchart of MAT process	11-2
Fullwords	1-11

G

General register	1-7
examine (R)	2-7
modify (=)	2-8
Guard digits	6-9

H

Halfwords	1-11
Hardware segment table entry. See HSTE.	
HSTE	11-3
access level bits (L)	11-10
access mode bits (A)	11-12
dirty bit (D)	11-11
presence bit (P)	11-11
reference bit (B)	11-11
segment limit field	11-12
shared segment bits (S)	11-12
SRF	11-13
SWSTE	11-13

I,J,K

I/O device interrupts	
immediate	10-28
priority levels	10-27
I/O instructions	9-8
AL	9-15
OC	9-9
OCR	9-9
RD	9-11
RDR	9-11
RH	9-12
RHR	9-12
SCP	9-17
SS	9-10
SSR	9-10
WD	9-13
WDR	9-13
WH	9-14
WHR	9-14
I/O interrupt mask (I)	10-5
I/O operations	
control of	9-3
device controllers	9-1
I/O bus	9-1
instruction formats	9-8
instructions	9-8
interrupt driven I/O	9-4
ISPT	9-2
SELCH I/O	9-6
status monitoring I/O	9-4
I/O references	E-1
IIP	10-4
Increment and examine next location (+)	2-7
Initial program load. See IPL.	
Input/output. See I/O.	
Instruction formats	1-12
alphanumeric string	7-3
branch instructions	1-14
decimal string	7-3
Instruction summary	
alphabetical by mnemonic	B-1
numerical by opcode	C-1
Instructions	
alignment	1-11
BRK	10-1
SINT	10-1
SVC	10-1
user level	1-11
Interrupt	
arithmetic fault	10-32
data format faults	10-18
driven	9-4
illegal instruction	10-18
instructions	10-13
machine malfunction	10-20
maskable	10-10
nonmaskable	10-10
precedence	10-12
queuing	9-2
relocation/protection (MAT) fault	10-20
simulated	10-29

Logical instructions
(Continued)

CLI	3-25
CLR	3-25
CRC12	3-52
CRC16	3-52
EXBR	3-20
EXHR	3-19
formats	3-4
L	3-7
LA	3-11
LB	3-18
LBR	3-18
LCS	3-9
LH	3-10
LHI	3-10
LHL	3-16
LI	3-7
LIS	3-8
LM	3-17
LR	3-7
LRA	3-12
N	3-30
NH	3-31
NHI	3-31
NI	3-30
NR	3-30
O	3-32
OH	3-33
OHI	3-33
OI	3-32
OR	3-32
RBL	3-60
RBT	3-50
RLL	3-43
RRL	3-45
RTL	3-60
SBT	3-49
SLHL	3-41
SLHLS	3-41
SLL	3-39
SLLS	3-39
SRHL	3-42
SRHLS	3-42
SRL	3-40
SRLS	3-40
ST	3-21
STB	3-24
STBR	3-24
STH	3-22
STM	3-23
TBT	3-48
THI	3-37
TI	3-36
TLATE	3-54
TS	3-47
X	3-34
XH	3-35
XHI	3-35
XI	3-34
XR	3-34
LSU	2-5 10-24 10-25
LVL	10-3

M

Machine malfunction interrupt	10-20
Machine malfunction interrupt enable (M)	10-5
Machine malfunction status word	10-22
MAT	1-2 3-13
faults	11-17
process flowchart	11-2
real address	11-1
MAT (Continued)	
virtual address	11-1
Memory	
initialization	2-9
reserved locations	1-8
Memory access level field.	
See LVL.	
Memory address translator.	
See MAT.	
Memory management	
MAT	11-1
Memory management instructions	11-21
LPSTD	11-22
LSSTD	11-23
Mode	
console	10-14
run	10-15
single step	10-16
Modify	
current location (=)	2-7
double precision floating point register (=)	2-9
general register (=)	2-8
PSW (=)	2-9
single precision floating point register (=)	2-8
Move and pad (MOVE)	7-8
Move and pad with default pad (MOVEP)	7-8
Move translated until (MVTU)	7-6
Multiple double precision floating point (MD)	6-50
Multiply (M)	5-22
Multiply floating point (ME)	6-29
Multiply floating point register (MER)	6-29
Multiply halfword (MH)	5-24
Multiply halfword register (MHR)	5-24
Multiply register (MR)	5-22
Multiply register double precision floating point (MDR)	6-50

N

No operation (NOP)	4-30
No operation register (NOPR)	4-30
Nonconfigured memory address	10-26
Noncorrectable memory error	10-25
Nonpresence fault	11-18
Normalization	6-5

O	
Opcode map	A-1
Operands	
arithmetic	1-10
Operations	
Boolean	3-2
list processing	3-3
translation	3-2
OR (O)	3-32
OR halfword (OH)	3-33
OR halfword immediate (OHI)	3-33
OR immediate (OI)	3-32
OR register (OR)	3-32
Output command (OC)	9-9
Output command register (OCR)	9-9

P,Q	
Pack and move (PMV)	7-12
Pack and move absolute (PMVA)	7-12
Packed	
decimal data	1-11
format	7-1
Physical address	
selection of	11-7
Power restore	
LSU	10-24
	10-25
Presence bit (P)	11-11
Private segments	11-7
Privileged system function (PSF)	10-40
Process byte (PB)	8-2
Process byte register (PBR)	8-4
Process segment table descriptor. See PSTD.	
Process segment table. See PST.	
Processor block diagram	1-3
Processor interrupts	
PSW	1-7
status switch	1-7
status word	1-7
Processor modes	
console	10-14
run	10-15
single step	10-16
Processor/controller communication	9-2
Program status word. See PSW.	
Protect mode enable (P)	10-7
PST	11-3
PSTD	11-3
PSW	1-2
	10-3
condition code	1-6
	10-8
CSF	10-3
examine (P)	2-9
FLM	10-3
FLU	10-6
I/O interrupt mask (I)	10-5
IIP	10-4
interrupt masks	1-4
LOC	1-4
	10-1
	10-8
LVL	10-3

PSW (Continued)	
machine malfunction	
interrupt enable (M)	10-5
modify (=)	2-9
protect mode enable (P)	10-7
register set select	1-5
register set select	
field (R)	10-7
relocation/protection	
enable (R/P)	10-6
reserved memory locations	10-9
SQS interrupt enable (Q)	10-6
status information	1-4
status word	10-1
wait state (W)	10-4

R	
R* rounding	6-9
Read data (RD)	9-11
Read data register (RDR)	9-11
Read error logger (REL)	10-41
Read halfword (RH)	9-12
Read halfword register (RHR)	9-12
Real address	11-1
Reference bit (B)	11-11
Register and indexed storage/	
register and indexed stor-	
age. See RXXR.	
Register and immediate	
storage one. See RI1.	
Register and immediate	
storage two. See RI2.	
Register and indexed storage	
one. See RX1.	
Register and indexed storage	
three. See RX3.	
Register and indexed storage	
two. See RX2.	
Register and indexed	
storage. See RX.	
Register set select field (R)	10-7
Register to register. See	
RR.	
Registers	
floating point	1-7
general	1-7
Relocation/protection (MAT)	
fault interrupt	10-20
Relocation/protection enable	
(R/P)	10-6
Remove from bottom of list	
(RBL)	3-60
Remove from top of list (RTL)	3-60
Reserved memory locations	1-8
	10-9
Reset bit (RBT)	3-50
Reset memory voltage failure	
(RMVF)	10-51
Restore interruptible state	
(LSRST)	10-49
RI1 format	1-25
RI2 format	1-27
Rotate left logical (RLL)	3-43
Rotate right logical (RRL)	3-45
RR	
format	1-14
	1-16
instruction	1-2
Run mode (<)	2-9
RX format	1-14

RX1 format	1-18	Single precision floating	
RX2 format	1-20	point register	
RX3 format	1-23	examine (F)	2-8
RXR format	1-29	modify (=)	2-8
		SINT	10-29
		Software segment table	
		entry. See SWSTE.	
		SQS interrupt	10-29
		SQS interrupt enable (Q)	10-6
		SRF	11-3
			11-13
		SST	11-3
		SSTD	11-3
		Status monitoring	9-4
		Status switching	10-17
		data format fault	
		interrupt	10-18
		illegal instruction	
		interrupt	10-18
		Status switching instructions	10-33
		BRK	10-39
		ECC XSTB	10-50
		EPSR	10-36
		ISRST	10-49
		ISSV	10-48
		LDPS	10-46
		LPSTD	10-43
		LPSW	10-34
		LPSWR	10-35
		LSSTD	10-44
		PSF	10-40
		REL	10-41
		RMVF	10-51
		SINT	10-37
		STPS	10-45
		SVC	10-38
		STD format	11-8
		STE	11-9
		HSTE	11-10
		segment table	11-10
		size	11-10
		SWSTE	11-13
		Store (ST)	3-21
		Store binary as packed	
		decimal string (STBP)	7-5
		Store byte, no ECC (XSTB)	10-50
		Store byte (STB)	3-24
		Store byte register (STBR)	3-24
		Store double precision float-	
		ing point register in single	
		precision memory (STDE)	6-58
		Store double precision	
		floating point (STD)	6-43
		Store floating point (STE)	6-22
		Store halfword (STH)	3-22
		Store multiple (STM)	3-23
		Store multiple double pre-	
		cision floating point (STMD)	6-44
		Store multiple floating	
		point (STME)	6-23
		Store process state (STPS)	10-45
		String	
		alphanumeric byte	1-9
		decimal byte	1-9
		String instructions	7-3
		CPAN	7-10
		CPANP	7-10

String instructions (Continued)	
LPB	7-4
MOVE	7-8
MOVEP	7-8
MVTU	7-6
PMV	7-12
PMVA	7-12
STBP	7-5
UMV	7-14
UMVA	7-14
String operations	
packed decimal data	7-1
unpacked decimal data	7-1
Subroutine address	9-19
Subroutine linkage	
branch	4-2
link	4-2
Subtract (S)	5-14
Subtract double precision	
floating point (SD)	6-47
Subtract floating point (SE)	6-26
Subtract floating point	
register (SER)	6-26
Subtract halfword (SH)	5-16
Subtract halfword immediate	
(SHI)	5-16
Subtract immediate (SI)	5-14
Subtract immediate short	
(SIS)	5-14
Subtract register (SR)	5-14
Subtract register double pre-	
cision floating point (SDR)	6-47
Supervisor call (SVC)	10-1
	10-38
Supervisor call. See SVC.	
SVC interrupt	10-31
SWSTE	
disk address	11-17
reference history bits	11-16
reserved field	11-16
usage mode bits	11-13
System	
breakpoint interrupt	10-31
initialization on MAT	11-21
terminal	2-1
terminal commands	2-6
terminal support command	
summary	2-2
System breakpoint (BRK)	10-39
System queue service. See	
SQS.	

T

Terminal system commands	2-6
-----------------------------	-----

Terminal (Continued)	
system support command	
summary	2-2
Test and set (TS)	3-47
Test bit (TBT)	3-48
Test immediate (TI)	3-36
Test immediate halfword (THI)	3-37
Translate (TLATE)	3-54
Translation	3-2
	9-20
True zero	6-7

U

Unpack and move (UMV)	7-14
Unpack and move absolute	
(UMVA)	7-14
Unpacked	
decimal data	1-11
format	7-1
zoned decimal	7-2
Usage mode bits	11-13
active state	11-15
I/O ongoing state	11-15
loading state	11-14
unload pending state	11-16
unloading state	11-15
unused state	11-14
used state	11-14
User level instructions	1-11

V

VDU	2-1
Video display unit. See VDU.	2-1
Virtual address	11-1
	11-6
offset field	11-7
page field	11-7
segment field	11-6
selection of	11-7
setting space size	11-8
translation to real	
address	11-2
	11-3
Virtual memory	11-1

W-Z

Wait state (W)	10-4
Write data (WD)	9-13
Write data register (WDR)	9-13
Write halfword (WH)	9-14
Write halfword register (WHR)	9-14