



HP 1000 A700 Computer

Reference Manual

FEDERAL COMMUNICATIONS COMMISSION RADIO FREQUENCY INTERFERENCE STATEMENT

The Federal Communications Commission (in 47 CFR 15.805) has specified that the following notice be brought to the attention of the users of this product.

Warning: This equipment generates, uses, and can radiate radio frequency energy and if not installed and used in accordance with the instruction manual, may cause interference to radio communications. It has been tested and found to comply with the limits for Class A computing devices pursuant to Subpart J of Part 15 of FCC Rules, which are designed to provide reasonable protection against such interference. 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.



HEWLETT
PACKARD

HP Computer Museum
www.hpmuseum.net

For research and education purposes only.

PRINTING HISTORY

The Printing History below identifies the Edition of this Manual and any Updates that are included. Periodically, Update packages are distributed which contain replacement pages to be merged into the manual, including an updated copy of this Printing History page. Also, the update may contain write-in instructions.

Each reprinting of this manual will incorporate all past Updates, however, no new information will be added. Thus, the reprinted copy will be identical in content to prior printings of the same edition with its user-inserted update information. New editions of this manual will contain new information, as well as all Updates.

To determine what software manual edition and update is compatible with your current software revision code, refer to the appropriate Software Numbering Catalog, Software Product Catalog, or Diagnostic Configurator Manual.

First Edition Mar 1982
Update 1 Jul 1982

NOTICE

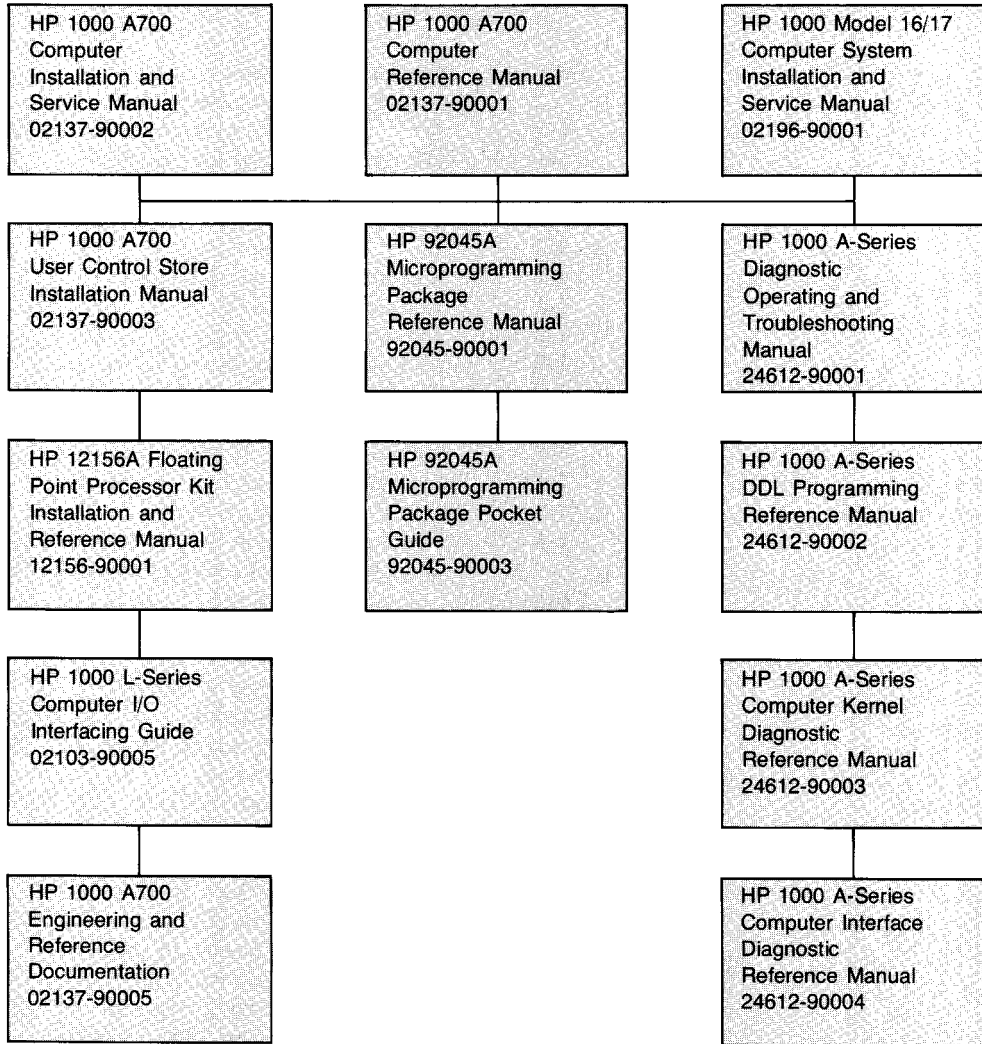
The information contained in this document is subject to change without notice.

HEWLETT-PACKARD MAKES NO WARRANTY OF ANY KIND WITH REGARD TO THIS MATERIAL, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE. Hewlett-Packard shall not be liable for errors contained herein or for incidental or consequential damages in connection with the furnishing, performance or use of this material.

Hewlett-Packard assumes no responsibility for the use or reliability of its software on equipment that is not furnished by Hewlett-Packard.

This document contains proprietary information which is protected by copyright. All rights are reserved. No part of this document may be photocopied, reproduced or translated to another program language without the prior written consent of Hewlett-Packard Company.

DOCUMENTATION MAP



CONTENTS

Section I	Page		
GENERAL FEATURES			
Architecture	1-1	Nonexistent Memory	3-4
Floating Point Processor	1-1	Base Set Instruction Formats	3-4
User Microprogramming	1-1	Memory Reference Instructions	3-4
Virtual Control Panel	1-2	Register Reference Instructions	3-4
Bootstrap Loaders	1-2	Input/Output Instructions	3-4
Self-Test Routines	1-2	Extended Arithmetic Memory	
Time Base Generator	1-2	Reference Instructions	3-5
Power Supply	1-2	Extended Arithmetic Register	
Input/Output	1-2	Reference Instructions	3-5
Memory	1-3	Extended Instructions	3-5
Software	1-3	Floating Point Instructions	3-5
HP Interface Bus	1-3	Language Instruction Set	3-5
Computer Network	1-4	Double Integer Instructions	3-5
Expansion and Enhancement	1-4	Virtual Memory Instructions	3-5
Specifications	1-4	Operating System Instructions	3-5
		Base Set Instruction Coding	3-5
		Memory Reference Instructions	3-5
		Register Reference Instructions	3-7
		Shift/Rotate Group	3-7
		Alter/Skip Group	3-10
		Input/Output Instructions	3-12
		Extended Arithmetic Memory	
		Reference Instructions	3-14
		Extended Arithmetic Register	
		Reference Instructions	3-15
		Extended Instruction Group	3-17
		Index Register Instructions	3-17
		Jump Instructions	3-20
		Byte Manipulation Instructions	3-21
		Bit Manipulation Instructions	3-22
		Word Manipulation Instructions	3-23
		Floating Point Instructions	3-24
		Single Precision Operations	3-24
		Double Precision Operations	3-25
		Language Instruction Set	3-27
		Double Integer Instructions	3-29
		Virtual Memory Instructions	3-31
		Operating System Instruction Set	3-33
		Execution Times	3-33
		Scientific Instruction Set	3-33
		Execution Times and Interrupts	3-38
		Vector Instruction Set	3-38
		Execution Times and Interrupts	3-45
		Assembly Language	3-45
		RTE Implementation	3-46
		Section IV	Page
		DYNAMIC MAPPING SYSTEM	
		Memory Addressing	4-1
		General Descriptions	4-2
		Page Mapping Register Instructions	4-2
		Working Map Instructions	4-2
		Cross-Map Instructions	4-2
		Detailed Descriptions	4-3
		DMS Instruction Execution Times	4-11
		Assembly Language and RTE Implementation	4-11
Section II	Page		
OPERATING FEATURES			
Hardware Registers	2-1		
A-Register	2-1		
B-Register	2-1		
P-Register	2-1		
Extend (E) Register	2-1		
Overflow (O) Register	2-1		
Central Interrupt Register	2-1		
Violation Register	2-1		
Parity Error Register	2-1		
Interrupt System Register	2-1		
X- and Y-Registers	2-1		
WMAP-Register	2-1		
Virtual Registers	2-2		
M-Register	2-2		
T-Register	2-2		
Controls and Indicators	2-2		
Self-Test	2-2		
Bootstrap Loaders	2-2		
Loader Selection for Auto-Boot	2-2		
Program Starts	2-3		
PROM Re-Entry and Sequential Execution	2-4		
Device Parameters and Media Formats	2-4		
Virtual Control Panel	2-4		
VCP Program Operation	2-4		
Loader Commands	2-6		
VCP User Considerations	2-6		
VCP Slave Functions	2-6		
Section III	Page		
PROGRAMMING INFORMATION			
Data Formats	3-1		
Addressing	3-1		
Paging	3-1		
Direct and Indirect Addressing	3-3		
Memory Mapping	3-3		
Reserved Memory Locations	3-3		

CONTENTS (Continued)

<p>Section V Page</p> <p>MICROPROGRAMMING</p> <p>The Microprogrammed Computer 5-1</p> <p>The Microprogrammable Computer 5-1</p> <p>Customized Instructions 5-1</p> <p>System Speed 5-1</p> <p>Memory Space and Security 5-1</p> <p>Developing Microprograms 5-2</p> <p>Support for the Microprogrammer 5-2</p> <p>FPP Microprogramming 5-2</p> <p>Conclusion 5-2</p> <p>Section VI Page</p> <p>INTERRUPT SYSTEM</p> <p>Power Fail Interrupt 6-1</p> <p>Parity Error Interrupt 6-3</p> <p>Memory Protect Interrupt 6-3</p> <p>Unimplemented Instruction Interrupt 6-3</p> <p>Time Base Generator Interrupt 6-3</p> <p>Input/Output Interrupt 6-4</p> <p>Interrupt Priority 6-4</p> <p>Central Interrupt Register 6-4</p> <p>Processor Status Register 6-4</p> <p>Interrupt Type Control 6-5</p> <p>Instruction Summary 6-5</p> <p>Section VII Page</p> <p>INPUT/OUTPUT SYSTEM</p> <p>Input/Output Addressing 7-1</p> <p>Input/Output Priority 7-1</p> <p>Interface Elements 7-3</p> <p style="padding-left: 20px;">Global Register 7-4</p>	<p>Control Bits 7-4</p> <p>Flag Bits 7-4</p> <p>Data Buffer Register 7-4</p> <p>Control Register 7-4</p> <p>Direct Memory Access 7-4</p> <p>Control Word 1 7-5</p> <p>Control Word 2 7-5</p> <p>Control Word 3 7-5</p> <p>DMA Transfer Initialization 7-5</p> <p>Self-Configured DMA 7-5</p> <p>DMA Data Transfer 7-5</p> <p>Non-DMA Data Transfer 7-7</p> <p>Input Data Transfer (Interrupt Method) 7-7</p> <p>Output Data Transfer (Interrupt Method) 7-7</p> <p>Non-Interrupt Data Transfer 7-7</p> <p style="padding-left: 20px;">Input 7-9</p> <p style="padding-left: 20px;">Output 7-9</p> <p>Diagnose Modes 7-9</p> <p style="padding-left: 20px;">Diagnose Mode 1 7-10</p> <p style="padding-left: 20px;">Diagnose Mode 2 7-10</p> <p style="padding-left: 20px;">Diagnose Mode 3 7-10</p> <p>Appendix A Page</p> <p>Character Codes A-3</p> <p>Octal Arithmetic A-4</p> <p>Octal/Decimal Conversions A-5</p> <p>Mathematical Equivalents A-6</p> <p>Octal Combining Tables A-8</p> <p>Instruction Codes in Octal A-9</p> <p>Base Set Instruction Codes in Binary A-10</p> <p>Extend and Overflow Examples A-15</p> <p>Interrupt and Control Summary A-16</p>
---	--

ILLUSTRATIONS

Title	Page	Title	Page
HP 1000 A700 Computers	1-0	Microprogramming Implementation Process	5-3
A700 Computer Simplified Block Diagram	1-2	Input/Output System	7-2
Loading Device Parameters and Media Formats ...	2-7	I/O Priority Assignments	7-2
Loader Command Format	2-9	Priority Linkage (Simplified)	7-2
Data Formats and Octal Notation	3-2	Interrupt Sequence	7-3
Base Set Instruction Formats	3-4	General Bit Definitions for Control Word 1	7-6
Shift and Rotate Functions	3-8	DMA Input Data Transfer	7-8
Examples of Double-Word Shifts and Rotates	3-16	Input Data Transfer (Interrupt Method)	7-8
Basic Logical Memory Addressing Scheme	4-1	Output Data Transfer (Interrupt Method)	7-9
Expanded Memory Addressing Scheme	4-1		

TABLES

Title	Page	Title	Page
Options and Accessories	1-5	SIS Instruction Error Codes	3-35
Specifications	1-6	Typical Scientific Instruction Set	
Start-Up Switch Settings	2-3	Execution Times	3-45
VCP Characters and Associated Registers	2-5	Typical Vector Instruction Set	3-45
VCP Commands	2-6	Instructions and Opcodes for RTE	
VCP Loader Command Errors	2-10	Implementation	3-47
Memory Paging	3-3	Dynamic Mapping Instructions Execution Times ...	4-11
Reserved Memory Locations	3-3	A700 Interrupt Assignments	6-1
Shift/Rotate Group Combining Guide	3-8	Sample Power Fail Subroutine	6-2
Alter/Skip Group Combining Guide	3-10	Instructions for Select Codes 00 through 07	6-5
Typical Base Set Instruction Execution Times	3-34	Noninterrupt Transfer Routines	7-9
Typical Execution Times for Floating Point		Diagnose Mode 1	7-10
Instructions with Optional Hardware FPP Card .	3-35	Diagnose Mode 2	7-10

ALPHABETICAL INDEX OF STANDARD INSTRUCTIONS

Instruction	Page	Instruction	Page		
ADA	Add to A	3-6	HLT	Halt	3-13
ADB	Add to B	3-6	INA	Increment A	3-11
ADX	Add Memory to X	3-17	INB	Increment B	3-11
ADY	Add Memory to Y	3-17	IOR	"Inclusive Or" to A	3-6
ALF	Rotate A Left Four	3-8	ISX	Increment X and Skip if Zero	3-18
ALR	A Left Shift, Clear Sign	3-8	ISY	Increment Y and Skip if Zero	3-18
ALS	A left Shift	3-8	ISZ	Increment and Skip if Zero	3-6
AND	"And" to A	3-6	JLA	Jump and Load A	3-20
ARS	A Right Shift	3-8	JLB	Jump and Load B	3-20
ASL	Arithmetic Shift Left	3-15	JLY	Jump and Load Y	3-20
ASR	Arithmetic Shift Right	3-15	JMP	Jump	3-6
BLF	Rotate B Left Four	3-8	JPY	Jump Indexed by Y	3-20
BLR	B Left Shift, Clear Sign	3-9	JSB	Jump to Subroutine	3-7
BLS	B Left Shift	3-9	LAX	Load A Indexed by X	3-18
BRS	B Right Shift	3-9	LAY	Load A Indexed by Y	3-18
CAX	Copy A to X	3-17	LBT	Load Byte	3-21
CAY	Copy A to Y	3-17	LBX	Load B Indexed by X	3-18
CBS	Clear Bits	3-22	LBY	Load B Indexed by Y	3-18
CBT	Compare Bytes	3-21	LDA	Load A	3-7
CBX	Copy B to X	3-17	LDB	Load B	3-7
CBY	Copy B to Y	3-17	LDMP	Load Map Set	4-3
CCA	Clear and Complement A	3-10	LDX	Load X from Memory	3-19
CCB	Clear and Complement B	3-10	LDY	Load Y from Memory	3-19
CCE	Clear and Complement E	3-11	LIA	Load Input to A	3-13
CLA	Clear A	3-11	LIB	Load Input to B	3-13
CLB	Clear B	3-11	LPMR	Load Page Map Register	4-3
CLC	Clear Control	3-12	LSL	Logical Shift Left (32)	3-15
CLE	Clear E	3-9,3-11	LSR	Logical Shift Right (32)	3-15
CLF	Clear Flag	3-12	LWD1	Load DATA1 Map	4-4
CLO	Clear Overflow	3-12	LWD2	Load DATA2 Map	4-4
CMA	Complement A	3-11	MB00	Cross Move Bytes, Execute to Execute	4-9
CMB	Complement B	3-11	MB01	Cross Move Bytes, Execute to DATA1	4-9
CME	Complement E	3-11	MB02	Cross Move Bytes, Execute to DATA2	4-9
CMW	Compare Words	3-23	MB10	Cross Move Bytes, DATA1 to Execute	4-9
CPA	Compare to A	3-6	MB11	Cross Move Bytes, DATA1 to DATA1	4-9
CPB	Compare to B	3-6	MB12	Cross Move Bytes, DATA1 to DATA2	4-10
CXA	Copy X to A	3-17	MB20	Cross Move Bytes, DATA2 to Execute	4-10
CXB	Copy X to B	3-17	MB21	Cross Move Bytes, DATA2 to DATA1	4-10
CYA	Copy Y to A	3-17	MB22	Cross Move Bytes, DATA2 to DATA2	4-10
CYB	Copy Y to B	3-17	MBT	Move Bytes	3-21
DIV	Divide	3-14	MIA	Merge Into A	3-13
DLD	Double Load	3-14	MIB	Merge Into B	3-13
DST	Double Store	3-14	MPY	Multiply	3-14
DSX	Decrement X and Skip if Zero	3-17	MVW	Move Words	3-23
DSY	Decrement Y and Skip if Zero	3-18	MW00	Cross Move Words, Execute to Execute	4-7
ELA	Rotate E Left with A	3-9	MW01	Cross Move Words, Execute to DATA1	4-7
ELB	Rotate E Left with B	3-9	MW02	Cross Move Words, Execute to DATA2	4-7
ERA	Rotate E Right with A	3-9	MW10	Cross Move Words, DATA1 to Execute	4-7
ERB	Rotate E Right with B	3-9	MW11	Cross Move Words, DATA1 to DATA1	4-8
FAD	Floating Point Add	3-24	MW12	Cross Move Words, DATA1 to DATA2	4-8
FDV	Floating Point Divide	3-24	MW20	Cross Move Words, DATA2 to Execute	4-8
FIX	Floating Point to Single Integer	3-25	MW21	Cross Move Words, DATA2 to DATA1	4-8
FLT	Single Integer to Floating Point	3-25	MW22	Cross Move Words, DATA2 to DATA2	4-8
FMP	Floating Point Multiply	3-24	NOP	No Operation	3-9
FSB	Floating Point Subtract	3-24	OTA	Output A	3-13

ALPHABETICAL INDEX OF STANDARD INSTRUCTIONS (Continued)

Instruction	Page	Instruction	Page
OTB	Output B	XLA1	Cross Load A through DATA1 Map
RAL	Rotate A Left	XLA2	Cross Load A through DATA2 Map
RAR	Rotate A Right	XLB1	Cross Load B through DATA1 Map
RBL	Rotate B Left	XLB2	Cross Load B through DATA2 Map
RBR	Rotate B Right	XOR	"Exclusive Or" to A
RRL	Rotate Left (32)	XSA1	Cross Store A through DATA1 Map
RRR	Rotate Right (32)	XSA2	Cross Store A through DATA2 Map
RSS	Reverse Skip Sense	XSB1	Cross Store B through DATA1 Map
SAX	Store A Indexed by X	XSB2	Cross Store B through DATA2 Map
SAY	Store A Indexed by Y	.CFER	Transfer Complex or Double Floating Point
SBS	Set Bits	.CPM	Single Integer Arithmetic Compare
SBT	Store Byte	.CPUID	Processor Identification
SBX	Store B Indexed by X	.DAD	Double Integer Add
SBY	Store B Indexed by Y	.DCO	Double Integer Compare
SEZ	Skip if E is Zero	.DDE	Double Integer Increment
SFB	Scan for Byte	.DDS	Double Integer Decrement and Skip if Zero
SFC	Skip if Flag Clear	.DFER	Transfer Extended Floating Point
SFS	Skip if Flag Set	.DIN	Double Integer Increment
SIMP	Save Interrupted Map	.DIS	Double Integer Increment and Skip if Zero
SLA	Skip if LSB of A is Zero	.DNG	Double Integer Negate
SLB	Skip if LSB of B is Zero	.DSB	Double Integer Subtract
SOC	Skip if Overflow Clear	.DSBR	Double Integer Subtract Reverse
SOS	Skip if Overflow Set	.ENTC	Transfer Parameter Addresses
SPMR	Store Page Mapping Register	.ENTN	Transfer Parameter Addresses
SSA	Skip if Sign of A is Zero	.ENTP	Transfer Parameter Addresses
SSB	Skip if Sign of B is Zero	.ENTR	Transfer Parameter Addresses
STA	Store A	.FWID	Firmware Identification
STB	Store B	.IMAP	16-Bit Subscript Mapping
STC	Set Control	.IRES	16-Bit Subscript Resolution
STF	Set Fla	.JMAP	32-Bit Subscript Mapping
STMP	Store Map Set	.JRES	32-Bit Subscript Resolution
STO	Set Overflow	.LBP	Mapping with Registers
STX	Store X to Memory	.LBPR	Mapping with DEF
STY	Store Y to Memory	.LPX	Indexed Mapping with Registers
SWMP	Save Working Map	.LPXR	Indexed Mapping with DEF
SZA	Skip if A is Zero	.PMAP	Map Specified Page
SZB	Skip if B is Zero	.SETP	Set A Table
TBS	Test Bits	.SIP	Skip if Interrupt Pending
XAX	Exchange A and X	.WFI	Wait for Interrupt
XAY	Exchange A and Y	.XFER	Transfer Extended Floating Point
XBX	Exchange B and X	.ZFER	Transfer Eight Words
XBY	Exchange B and Y	.FCM	Complement and Normalize Single Floating Point
XCA1	Cross Compare A through DATA1 Map		
XCA2	Cross Compare A through DATA2 Map		
XCB1	Cross Compare B through DATA1 Map		
XCB2	Cross Compare B through DATA2 Map		
XJMP	Cross Map Jump		

**ALPHABETICAL INDEX OF INSTRUCTIONS PROVIDED BY
THE OPTIONAL FLOATING POINT PROCESSOR CARD**

Instruction	Page
ALOG	Natural Logarithm 3-36
ALOGT	Common Logarithm 3-36
ATAN	Arctangent 3-36
COS	Cosine 3-36
DPOLY	Polynomial Evaluation 3-36
DVABS	Vector Absolute Value 3-41
DVADD	Vector Add 3-38
DVDIV	Vector Divide 3-39
DVDOT	Vector Dot Product 3-42
DVMAB	Vector Maximum Absolute Value 3-43
DVMAX	Vector Maximum Value 3-43
DVMIB	Vector Minimum Absolute Value 3-44
DVMIN	Vector Minimum Value 3-43
DVMOV	Vector Move 3-44
DVMPY	Vector Multiply 3-39
DVNRM	Vector Norm 3-42
DVPIV	Vector Pivot 3-41
DVSAD	Scalar-Vector Add 3-40
DVSDV	Scalar-Vector Divide 3-41
DVSMY	Scalar-Vector Multiply 3-40
DVSSB	Scalar-Vector Subtract 3-40
DVSUB	Vector Subtract 3-39
DVSUM	Vector Sum 3-42
DVSWP	Vector Swap 3-44
EXP	E to the Power X 3-36
SIN	Sine 3-36
SQRT	Square Root 3-33
TAN	Tangent 3-33
TANH	Hyperbolic Tangent 3-36
VABS	Vector Absolute Value 3-41
VADD	Vector Add 3-38
VDIV	Vector Divide 3-39
VDOT	Vector Dot Product 3-42
VMAB	Vector Maximum Absolute Value 3-43
VMAX	Vector Maximum Value 3-43

Instruction	Page
VMIB	Vector Minimum Absolute Value 3-44
VMIN	Vector Minimum Value 3-43
VMOV	Vector Move 3-44
VMPY	Vector Multiply 3-39
VNRM	Vector Norm 3-42
VPIV	Vector Pivot 3-41
VSAD	Scalar-Vector Add 3-40
VSDV	Scalar-Vector Divide 3-41
VSMY	Scalar-Vector Multiply 3-40
VSSB	Scalar-Vector Subtract 3-40
VSUB	Scalar Subtract 3-39
VSUM	Vector Sum 3-42
VSWP	Vector Swap 3-44
.BLE	Single Floating Point to Double Floating Point 3-27
.DDI	Double Integer Divide 3-30
.DDIR	Double Integer Divide Reverse 3-30
.DMP	Double Integer Multiply 3-29
.FIXD	Floating Point to Double Integer 3-25
.FLTD	Double Integer to Floating Point 3-25
.FPWR	Exponentiation 3-37
.NGL	Double Floating Point to Single Floating Point 3-27
.TADD	Double Floating Point Add 3-25
.TDIV	Double Floating Point Divide 3-26
.TFTD	Double Floating Point to Single Integer 3-26
.TFTS	Single Integer to Double Floating Point 3-26
.TFXD	Double Floating Point to Double Integer 3-26
.TFXS	Double Floating Point to Single Integer 3-26
.TMPY	Double Floating Point Multiply 3-26
.TSUB	Double Floating Point Subtract 3-25
.TPWR	Exponentiation 3-37
..TCM	Complement and Normalize Double Floating Point 3-29
/ATLG	(1-X)/(1+X) 3-37




Figure 1-1. HP 1000 A700 Computers

The HP 1000 A700 Computer and Computer System (hereafter referred to as A700 computers) are high performance members of the HP 1000 A-Series Computer family. The A700 computers deliver full minicomputer power to a wide variety of applications, and maintain software compatibility with previous HP 1000 Computers. The A700 hardware is available as HP 2137A Computers (boxes) and computer system processor units (HP 2197A/B). (See figure 1-1.)

1-1. ARCHITECTURE

The A700 computer architecture is based on a distributed intelligence concept that separates the processing of input/output (I/O) instructions from that of other instructions. The central processor unit (CPU) resides on two printed circuit boards and features a fully microprogrammed control processor. The CPU executes the powerful HP 1000 instruction set, including index instructions and a full complement of instructions for logical operations as well as bit and byte manipulation. The A700 computer base instruction set also includes double-integer and virtual memory addressing instructions, and a language instruction set that substantially increases program execution speed for such high-level languages as FORTRAN and Pascal. The CPU also performs several system level functions, including memory protect, power fail/auto restart, time base generation, parity error interrupt, and extensive self-tests.

All input/output instructions are executed by custom silicon-on-sapphire (SOS) input/output processor (IOP) integrated circuit chips that reside on the individual I/O interface cards. A common backplane links the processor, memory, and I/O cards. The instructions are fetched from memory and decoded by the processor card. When an instruction is decoded as being of the I/O type, it is broadcast on the backplane for execution by the appropriate I/O chip. Because each I/O card is capable of operating independently of the CPU, the A700 can perform direct memory access (DMA) I/O transfers very efficiently. An I/O card interacts with the CPU only on DMA initiation and completion; beyond that, the entire high-speed transfer is handled by the I/O card, leaving the CPU free to work on other tasks. This achieves high efficiency in CPU and I/O throughput. Figure 1-2 is a simplified block diagram of the A700 computer.

1-2. FLOATING POINT PROCESSOR

A hardware Floating Point Processor card can be plugged into the card cage of A700 computers to provide high-speed dedicated logic that performs floating point

arithmetic operations. The Floating Point Processor (FPP) has custom SOS chips and provides exceptionally fast execution of single precision (32-bit) and double precision (64-bit) floating point operations. The FPP includes a powerful Scientific Instruction Set (SIS) that performs trigonometric, logarithmic, and other transcendental functions. The FPP also includes a Vector Instruction Set (VIS) which uses the FPP as a computing resource to perform vector and matrix arithmetic and to process large data arrays. (The FPP is optional with A700 computers and standard in A700 systems.) The FPP achieves extremely high accuracy with computational speeds that are 6 to 30 times faster than comparable software routines.

1-3. USER MICROPROGRAMMING

The power and flexibility of microprogramming is made available to the A700 computer user through a microinstruction set of microorders. Microprogrammers have access to special scratch pad registers in addition to the other internal registers of the A700, and can address

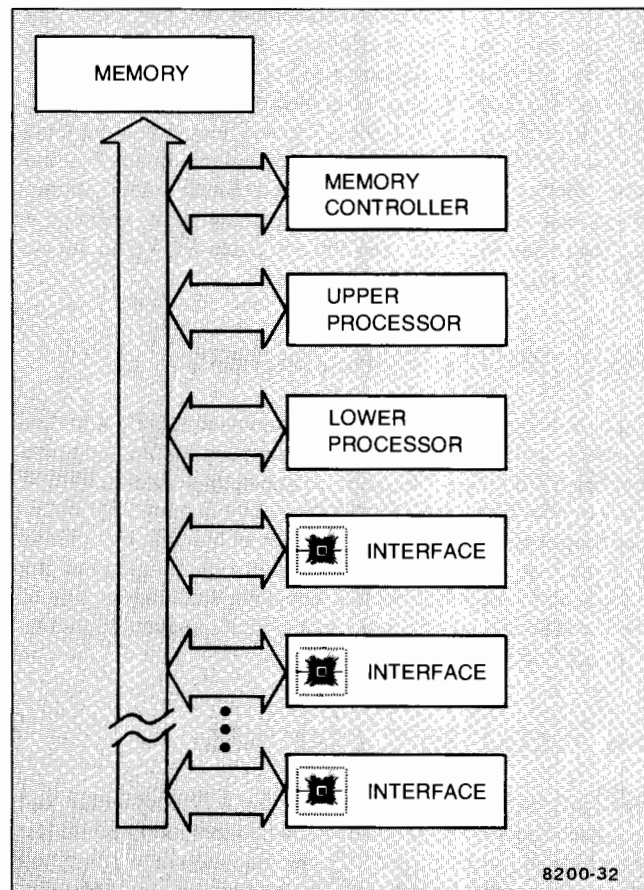


Figure 1-2. A700 Computer Simplified Block Diagram

up to 16K, 32-bit words of control store. A700 computers also support up to three levels of nested subroutines in microprograms. Microprogramming offers the advantages of speed and security as well as the ability to expand the instruction set to meet a variety of special computing needs.

Microprogramming is supported by Hewlett-Packard through a software package and customer training courses. A paraphraser microassembler allows the user to write microprograms in a versatile free-format style that greatly enhances program readability and documentation compared to traditional microprogramming techniques. User-developed microprograms can be dynamically loaded into optional Writable Control Store (WCS) cards for execution, and permanently fused into programmable read-only memory (PROM) chips for mounting on the optional PROM Control Store card. (Refer to Section V for more information on microprogramming.)

1-4. VIRTUAL CONTROL PANEL

The Virtual Control Panel (VCP) program is an interactive program that enables an external device (such as a terminal) to control the CPU in a manner similar to a conventional computer control panel and also provides additional features. That is, it allows the operator to access the various registers (A, B, P, etc.), examine or change memory, and control execution of a program. The VCP program is stored in PROM on the memory controller card. In a typical application, the VCP could be an HP 262x or HP 264x Terminal interfaced by an HP 12005A Asynchronous Serial Interface Card. When not being used as the VCP, the VCP-assigned terminal can be used in the same way as any other terminal connected to the system. When the A700 computer is operating as a node in a computer network via DS/1000-IV, the VCP can be an adjacent computer in the network.

1-5. BOOTSTRAP LOADERS

There are several bootstrap loaders stored in PROM on the memory controller card. The loaders provide program loading from several sources including disc drives, PROM storage modules, a DS/1000-IV network link, HP mini-cartridge tapes, and cartridge tapes of the HP 7908/11/12 Disc Drives. The first three loaders can be selected for auto-boot by switches on the computer frontplane; any of the loaders can be selected by operator commands via the Virtual Control Panel.

1-6. SELF-TEST ROUTINES

Self-test routines are standard in the A700 computer and are stored in PROM on a central processor card and on the memory controller card. These routines are executed whenever computer power is turned on, providing a convenient confidence-check of the processor card, memory cards, and part of the logic on each input/output card.

Execution of these routines can also be initiated by a switch on the computer frontplane or by operator command via the Virtual Control Panel.

1-7. TIME BASE GENERATOR

One of the processor cards includes a time base generator which can be used to time external events or to create a real-time clock in software. The time base generator (TBG) can generate an interrupt every 10 milliseconds. The TBG, which can be enabled and disabled by standard I/O instructions, is disabled at power up.

1-8. POWER SUPPLY

A700 computers have a power supply designed to continue normal operation in environments where ac line power may fluctuate widely. Input line voltages and frequencies may vary widely without affecting the operation of the computer. An optional battery backup card and battery pack can be installed in the power supply area to sustain memory for 30 or 60 minutes (depending on memory size) in the event of a complete power failure, thus providing an automatic restart capability. Another power supply option provides two 25-kHz voltages that can be rectified at the load and used to power accessory plug-in cards used for measurement and control applications.

1-9. INPUT/OUTPUT

The input/output system for A700 computers features a custom SOS chip on each I/O card, enabling each card to process its own I/O instructions and handle direct memory access (DMA) data transfers. The I/O system has a multilevel vectored priority interrupt structure with 53 distinct interrupt levels, each of which has a unique priority assignment. Any I/O device can be selectively enabled or disabled, or all I/O devices can be enabled or disabled under program control.

Data transfer between the computer and I/O devices can take place under DMA control or program control. The DMA capability provides a direct link between memory and I/O devices. The total bandwidth through multiple DMA channels is 4.0 million bytes (2.0 million words) per second.

The A700 computer backplane provides the link between the processor, memory, interface cards, and the power supply. The backplane has slots for 20 plug-in cards. Two slots must be used for the processor cards, one for the memory controller card, and one for each memory array card. Thus, depending on the number of memory array cards installed, there are up to 16 slots available for I/O cards in the box computer. In addition to the two processor cards and the two memory cards, the computer system includes as standard a Floating Point Processor card, a terminal interface card, and a disc drive interface card, leaving up to 13 slots available for I/O cards.

The A700 computer uses the HP L-Series I/O cards and an important feature of these cards is a common-content Global Register which can be loaded with the select code of a specific I/O card. When the Global Register is enabled all I/O instructions are executed only by the I/O card whose select code is in the Global Register. This not only facilitates setting up DMA transfers but also makes re-configuration of an I/O driver a simple matter of changing the Global Register to the appropriate select code. Also, since the Global Register can direct I/O instructions to a specific I/O card, the I/O-instruction address bits can be used to access registers on an I/O card. This feature is utilized in the design of the L-Series I/O cards to increase their capabilities.

About one-third of the area on all L-Series I/O cards is occupied by identical logic called the I/O Master, consisting of an I/O processor chip and its associated logic. The I/O Master is also available in breadboard form for users who wish to design their own I/O cards. The I/O Master is described in detail in the *HP 1000 L-Series Computer I/O Interfacing Guide*, part no. 02103-90005.

1-10. MEMORY

A700 computers are available with either of two semiconductor memory systems based on 64k-bit NMOS/RAM chips. The standard memory system consists of a memory controller card and up to four HP 12103A/B/C/D Memory Array Cards, each having 128k, 256k, 512k, or 1024k bytes of dynamic random-access memory (RAM).

The optional error correcting memory system provides fault-secure memory operation for the A700 computers. The system consists of the same memory controller card used for standard memory, and up to four HP 12104A 512k Byte Error Correcting Memory Array Cards. (Memory cycle time is the same for error correcting memory as it is for standard memory unless an error is being corrected.) The system is capable of correcting all single-bit errors and of detecting all double-bit errors and most multiple-bit errors. The memory controller card logs the error syndrome of single-bit errors, and LEDs on the memory array cards identify the memory chips causing these errors. The error correcting system is particularly valuable in computer systems with large amounts of memory, or where fault-secure operation is essential.

The maximum memory size available in A700 computers is four million bytes using standard memory, and two million bytes using error correcting memory. Addressing physical memory configurations larger than 64k bytes is made possible by the use of the Dynamic Mapping System (DMS), which is standard in the A700 and is described in Section IV. The DMS is a powerful memory management scheme that allows A700 computer users to address up to 32 megabytes of memory and provides either write or read-and-write protection of each individual 2048-byte page. For data integrity, memory parity checking is provided as a standard feature, and a parity-valid indi-

cator light is provided on each memory array card for quick fault isolation.

1-11. SOFTWARE

Software support for the A700 computers begins with RTE-A.1, a member of HP's family of Real-Time Executive (RTE) operating systems. RTE-A.1 is a real-time multiprogramming, multi-user system designed to take full advantage of the A700 I/O structure to enhance overall CPU and I/O throughput. RTE-A.1 offers a wide range of configurations, from a small, memory-based, execute-only system to a full disc-based system with on-line program development. Utilizing the A700 mapped memory system, RTE-A.1 supports user partitions of up to 64k bytes and memory sizes from 128k bytes to four megabytes. Memory can be divided into fixed and dynamically allocated partitions at system generation time. Critical programs can be made resident in fixed partitions to ensure fastest possible response to requests for their execution. Other programs can be assigned partitions from the dynamic memory pool according to need, using the smallest available block of memory.

RTE-A.1 also supports Virtual Memory Addressing (VMA) for access to data arrays much larger than main memory (up to 12.6 megabytes). The disc functions as an extension of main memory so far as data is concerned, in a manner that is transparent to the user and does not require any special programming. In addition, RTE-A.1 supports a special case of VMA, called Extended Memory Area (EMA). With EMA, up to two megabytes of a program's data can be in main memory at once, which affords faster processing of data arrays small enough to use the EMA capability. The programmer chooses the data array handling mode at program load time.

Disc-based RTE-A.1 systems support program development in FORTRAN 77, Pascal/1000, BASIC, and Macro/1000 Assembly Language. Program development for the A700 can be performed on an HP 1000 System under RTE-6/VM or RTE-IVB.

The diagnostic packages listed in table 1-1 may be used for testing and fault location.

1-12. HP INTERFACE BUS

Among the I/O interface cards available for the A700 computer is the HP 12009A HP-IB Interface Card which can interface the A700 computer to a variety of HP peripherals and other equipment compatible with the Hewlett-Packard Interface Bus (HP-IB). (HP-IB is the Hewlett-Packard implementation of IEEE standard 488-1978, "Digital Interface for Programmable Instrumentation".) A single HP 12009A can control up to 14 HP-IB instruments and several can be used to achieve concurrent operation of multiple HP-IB instrumentation clusters under the RTE-A.1 multiprogramming operating system.

1-13. COMPUTER NETWORK

The user can configure the A700 computer into an HP DS/1000-IV Distributed System by using either an HP 12007A or an HP 12044A HDLC Interface. Both of these interfaces support the high-level data link communications (HDLC) protocol, functioning as a preprocessor to handle low and medium levels of protocol processing. The A700 computers can be easily mixed with other members of the HP 1000 family in a single computer network. The HP 12042A Programmable Serial Interface allows the sophisticated OEM to design his own customized protocol for networks. Hewlett-Packard offers a customer training course on how to program the PSI card.

1-14. EXPANSION AND ENHANCEMENT

Table 1-1 lists accessory products available to expand or enhance the A700 computers.

1-15. SPECIFICATIONS

The HP 1000 A600/A700 Computational Products Technical Data handbook, part no. 5953-2898, provides complete specifications for the A700 computers and systems. Table 1-2 provides an abridged set of A700 specifications. Except where indicated, the specifications are applicable to both the HP 2137A Computer and the HP 2197A/2197B Computer System. Both the computer and the computer system have been product accepted by the Underwriters' Laboratories (UL) and the Canadian Standards Association (CSA). The A700 computer and system also meet the RFI standards of the Federal Communications Commission (FCC) and Verband Deutsches Electrotechniques (VDE).

Table 1-1. Options and Accessories

DESCRIPTION	PRODUCT NO.	OPTION NO.
Adds hardware floating point processor card	12156A	001
Removes standard memory array card	—	014
230 Vac Operation	—	015
128k Byte Memory Array Card	12103A	—
256k Byte Memory Array Card	12103B	—
512k Byte Memory Array Card	12103C	—
1024k Byte Memory Array Card	12103D	—
512k Byte Error Correcting Memory Array Card	12104A	—
Memory Connector for one memory array card	12038A	—
Memory Connector for two memory array cards	12038B	—
Memory Connector for three memory array cards	12038C	—
Memory Connector for four memory array cards	12038D	—
Asynchronous Serial Interface	12005B	—
Parallel Interface	12006A	—
HDLC Interface (modem operation)	12007A	—
PROM Storage Module	12008A	—
HP-IB Interface	12009A	—
Intelligent Breadboard	12010A	—
Extender Board	12011A	—
Priority Jumper Card	12012A	—
8-Channel Asynchronous Multiplexer	12040A	—
Programmable Serial Interface	12042A	—
HDLC Interface (hard-wired operation)	12044A	—
High-Level Analog Input Card	12060A	—
Expansion Multiplexer Card	12061A	—
Analog Output Card	12062A	—
16-In/16-Out Isolated Digital I/O Card	12063A	—
DS/1000-IV Data Link Slave Interface	12072A	—
DS/1000-IV Modem Interface to HP 3000	12073A	—
LAP-B Network Interface	12075A	—
DS/1000-IV Direct Connect Interface to HP 3000	12082A	—
Writable Control Store Card	12153A	—
PROM Control Store Card	12155A	—
Power Fail Recovery System	12157A	—
25 kHz Power Module	12158A	—
Diagnostic Package for A700 processor and interfaces	24612A*	—
Diagnostic Package for A700-compatible measurement and control interfaces	24613A	—
Diagnostic Package for A700-compatible hard disc drives and magnetic tape units	24398B*	—

*Included with the HP 2197A/B System Processor Units.



Table 1-2. Specifications

SPECIFICATIONS COMMON TO THE HP 2137A, 2197A, AND 2197B	
CENTRAL PROCESSOR	
Word Size:	16 bits
Instruction Set:	205 standard instructions. 270 with the Floating Point Processor card.*
Memory Reference:	14
Register Reference:	43
Input/Output:	13
Extended Arithmetic:	10
Index:	34
Bit, Byte, Word Manipulation:	10
Floating Point:	6 (16 with the Floating Point Processor card*)
Scientific:*	14
Language:	11 (14 with the Floating Point Processor card*)
Dynamic Mapping:	39
Vector Instructions:*	38
Double Integer:	12
Virtual Memory:	9
Operating System:	4
*Optional in the 2137A Computer; standard in the 2197A/2197B System.	
Registers:	
Accumulators:	Two (A and B), 16 bits each. Implicitly addressable, also explicitly addressable as memory locations.
Index:	Two (X and Y), 16 bits each.
Memory Register:	One (P), 15 bits.
Supplementary:	Two (overflow and extend), one bit each.
Power Fail Provisions:	When primary line power falls below a predetermined level while the computer is running, a power fail warning signal from the computer power supply causes an interrupt to memory location 00004. This location is intended to contain a jump-to-subroutine (JSB) instruction to a user-supplied power fail subroutine. A minimum of 5 milliseconds is available to execute the power fail subroutine.
Time Base Generator Interrupt:	A time base generator interrupt is provided for maintaining a real time clock. The interrupt request is made when the CPU signals, at 10-millisecond intervals, that its internal clock is ready to roll over. Timing accuracy of the time base generator is ± 2.16 seconds per 24-hour day.
CONTROL PROCESSOR	
Address Space:	16,384 words (256 blocks of 64 words each).
Microinstruction Word Size:	32 bits.
Word Types:	Six
Cycle Time:	250 nanoseconds.
Microorders:	198
Operations:	14
Special:	56
ALU and Conditional:	48
Store:	32
B-Bus:	32
A-Bus:	16

Table 1-2. Specifications (Continued)

SPECIFICATIONS COMMON TO THE HP 2137A, 2197A, AND 2197B (Continued)	
MEMORY	
Memory Structure:	64 pages minimum of 2048 bytes per page, with direct access to current page or base page (page 0), and indirect or mapped access to all pages.
Memory Size:	2137A: 128k bytes is standard, expandable to 3200k bytes. 2197A/2197B: 256k bytes is standard, expandable to 3328k bytes. Up to 4096k bytes if 1024k-byte memory array modules are installed. Up to 2048k bytes of optional fault control memory in 512k-byte increments.
Memory Cycle Time:	500 nanoseconds.
Memory Parity Checking:	Parity logic on the memory cards continuously generates correct parity for all words written into memory and monitors the parity of all words read out of memory. Either odd or even parity can be selected programmatically. A parity error generates an interrupt to memory location 00005, which can contain a JSB to a user-supplied parity error handling subroutine or a jump to a halt instruction.
INPUT/OUTPUT	
Determination of I/O Address:	I/O address select code is set for each interface card by select code switches on the card and is therefore independent of interface card position in the card cage.
I/O Device Interrupt Priority:	Depends upon I/O interface card position in the card cage with respect to the processor cards.
Interrupt Masking:	The I/O Master Logic includes an interrupt mask register which provides for selective inhibition of interrupts from specific interfaces under program control. This capability can be programmed to temporarily cut off undesirable interrupts from any combination of interfaces when they could interface with crucial transfers.
Interrupt Latency:	8.00 to 29.75 microseconds. 10.00 microseconds typical. (Interrupts cannot be serviced until a DMA cycle or an instruction in progress has completed execution.) The worst-case latency of 29.75 microseconds is based upon time to complete a floating point divide (FDV), the longest non-interruptible standard instruction.
Direct Memory Access (DMA):	The I/O processor chip supports DMA capability on each I/O interface, which reduces the number of interrupts from one per data item (byte or word) to one per complete DMA block.
DMA Latency:	Time interval from Service Request by an I/O device through completion of the DMA I/O data transfer to or from the I/O interface is 0.83 microseconds for input, 1.25 microseconds for output for the interface with highest hardware I/O priority.
Data Packing Under DMA:	When byte mode is specified in Control Word instructions, the I/O processor chip automatically manages byte packing or unpacking.
Maximum Achievable DMA Rate:	2.0 million words (4.0 megabytes) per second.
POWER SUPPLY	
Output:	DC voltages and tolerances. +5V ±2% +12V +6/-3% -12V ±6%
Optional AC Voltages and Tolerances:	27 Vrms ±8%, 25 kHz nominal, split phase from three pins on backplane-mating connector. Total harmonic distortion: <10%.
Maximum Output Current Ratings:	+5V +5M +12V -12V 25 kHz 68A 5.0A 5.6A 3.5A 2.5A
Short Circuit Protection:	All dc and ac power outputs are fault protected for short circuits. The power supply will shut down if any of the outputs are short circuited at turn on.

Table 1-2. Specifications (Continued)

SPECIFICATIONS COMMON TO THE HP 2137A, 2197A, AND 2197B (Continued)**+5V Output Overvoltage Protection:**

The +5V output is sensed for overvoltage and the +5V supply shuts down if its output voltage exceeds 5.5V. The ac power switch must be cycled to reset the +5V output.

DC REQUIRED:

MODULE	+5V dc	+5M dc	+12V dc	-12V dc
12152-60001 Upper CPU	5.0A			
12152-60002 Lower CPU	6.2A			
12152-60003 Memory Controller	4.4A	0.4A		
12103A 128 kb Memory	1.1A	1.0A		
(standby)*		0.5A		
12103B 256 kb Memory	1.1A	1.1A		
(standby)*		0.5A		
12103C 512 kb Memory	1.1A	1.1A		
(standby)*		0.6A		
12103D 1024 kb Memory	1.3A	1.4A		
(standby)*		0.9A		
12104A 512 kb E.C. Memory	1.5A	1.4A		
(standby)*		0.7A		
12153A WCS Card	4.1A		0.1A	
12155A PROM Control Store Card	1.2A			
— Add per 1k of PROM	0.6A			
12156A Floating Point Card	4.4A		0.1A	
12005A Asynchronous Serial Interface	1.6A		0.2A	0.1A
12006A Parallel Interface				
— With +5V logic	1.94A		0.18A	
— With +12V logic	1.61A		0.18A	
12007A HDLC Interface (modem)	2.6A		0.4A	0.2A
12008A PROM Storage Module	2.0A		0.1A	
12009A HP-IB Interface	2.1A		0.1A	
12010A Breadboard Interface**	0.8A		0.4A	
12040A Asynchronous Multiplexer	2.5A		0.1A	0.1A
12042B Programmable Serial Interface	2.6A		0.4A	0.2A
12044A HDLC Interface (direct)	2.4A		0.3A	0.1A
12060A Analog Input†	1.1A			
12061A Analog Multiplexert	0.1A			
12062A Analog Output†	1.2A			
12063A Digital I/O†	1.0A			
12072A Data Link Interface	1.5A		0.2A	0.1A
12073A Modem HP 3000 Interface	2.6A		0.4A	0.2A
12075A LAP-B Network Interface	2.6A		0.4A	0.2A
12082A 3000 Interface (direct)	2.4A		0.3A	0.1A

* Standby indicates the current supplied from the optional battery backup during power failure.

** 12010A current requirement listed here is for the I/O Master circuitry only; logic added by the user will require additional current.

† Requires 25 kHz power supply accessory.

Table 1-2. Specifications (Continued)

SPECIFICATIONS COMMON TO THE HP 2137A, 2197A, AND 2197B (Continued)**BATTERY BACKUP**

12157-60001 Battery Card.

1420-0304 Battery Pack.

When the battery pack is fully charged it sustains four megabytes of memory for at least 15 minutes, or sustains two megabytes for at least 30 minutes. Additional sustaining time can be achieved by connecting an external battery.

Recharge time:

12 hours for fully discharged battery pack.

Battery type:

Sealed lead acid.

SAFETY AND RFI

HP 1000 A700-Series products are UL listed and CSA certified to meet Underwriters' Laboratories (UL) and the Canadian Standards Association (CSA) standards for safety. The A700-Series also complies with the RFI standards of the Federal Communications Commission (FCC) and Verband Deutsches Elektrotechniken (VDE).

SPECIFICATIONS APPLICABLE ONLY TO THE HP 2137A COMPUTER**ELECTRICAL SPECIFICATIONS****AC Power Required****Line Voltage:**86-138V (115V -25%/+20%) standard;
178-276V (230V -23%/+20%) option 015.**Line Frequency:**

47.5 to 66 Hz.

Maximum Power Required:

700 Watts.

PHYSICAL CHARACTERISTICS**Dimensions****Height:**

266 mm (10.5 in).

Width:

483 mm (19 in).

Depth:

610 mm (24 in).

Weight:

26.8 kg (59 lb).

Ventilation:

Air intake is in through the front; exhaust is out through the rear.

Volume:

Approximately 10.1 cubic metres/min (360 cfm).

ENVIRONMENTAL SPECIFICATIONS**Temperature****Operating:**

0° to 55°C (32° to 131°F) up to 3.1 km (10,000 ft); 0° to 45°C (32° to 113°F) up to 4.6 km (15,000 ft).

Non-operating:

-40° to 75°C (-40° to 167°F).

Relative Humidity:

5% to 95% non-condensing.

Altitude**Operating:**

To 3.1 km (10,000 ft) at 0° to 55°C (32° to 131°F). To 4.6 km (15,000 ft) at 0° to 45°C (32° to 113°F).

Non-operating:

15.3 km (50,000 ft).

Vibration and Shock:

HP 1000 A700-Series products are type tested for normal shipping and handling shock and vibration. (Contact factory for review of any application that requires operation under continuous vibration.)

Table 1-2. Specifications (Continued)

SPECIFICATIONS APPLICABLE ONLY TO THE HP 2197A/2197B SYSTEMS	
ELECTRICAL SPECIFICATIONS	
Standard Line Voltage and Line Frequency	
Line Voltage (With 7908R):	88-127V (115V nominal).
Line Voltage (With 7911R or 7912R):	90-105V (100V nominal) or 108-126V (120V nominal).
Line Frequency:	With 7908R: 47.5 to 66 Hz. With 7911/12R: 54 to 66 Hz.
Option 015 Line Voltage and Line Frequency	
Line Voltage (With 7908R):	187-253V (230V nominal).
Line Voltage (With 7911R or 7912R):	198-231V (220V nominal) or 216-252V (240 nominal).
Line Frequency:	With 7908R: 47.5 to 66 Hz. With 7911/12R: 48 to 55 Hz.
Power Requirements:	Requires at least 20 Ampere grounded power receptacle for 115 Vac operation, or at least 10 Ampere grounded power receptacle for 230 Vac operation (option 015). The 2197A requires split-phase power; the 2197B requires single-phase power. An additional power receptacle is required for the system console.
Maximum Current Required:	2197A: 16 Amperes per phase. 2197B: 16 Amperes.
Maximum Power Dissipation in Upper Section of 2197A Cabinet:	250 Watts.
PHYSICAL CHARACTERISTICS	
Dimensions	
Height:	Model 17 (2197A): 1613 mm (63.4 in). Model 17 (2197B): 720 mm (28.3 in).
Width:	635 mm (25 in).
Depth:	813 mm (32 in).
Weight	
Without disc drive:	2197A: 139.7 kg (307.5 lb). 2197B: 94.3 kg (207.5 lb).
7908R Disc Drive adds:	37.0 kg (81.6 lb).
7911R/7912R Drive adds:	67.3 kg (148 lb).
Racking Limitations:	The additional space in the top half of the 2197A cabinet is intended for instrumentation installed on rails and not on slides. The maximum weight for this add-on equipment is limited to 125 kg (275 lbs) due to anti-tip safety requirements.
ENVIRONMENTAL SPECIFICATIONS	
Temperature	
Operating (SPU only):	0° to 55°C (32° to 131°F) up to 3.1 km (10,000 ft); 0° to 45°C (32° to 113°F) up to 4.6 km (15,000 ft).
Operating (79xxR Disc):	0° to 40°C (32° to 104°F), rate of change for 7908R <10°C (18°F) per hour; for 7911/12R <20°C (36°F) per hour.
Non-operating:	-40° to 60°C (-40° to 140°F).
Relative Humidity:	
SPU only:	5% to 95% with maximum wet bulb temperature not to exceed 25.6°C (78.1°F), excluding all conditions which cause condensation.

Table 2-1. Specifications (Continued)

SPECIFICATIONS APPLICABLE ONLY TO THE HP 2197A/2197B SYSTEMS (Continued)	
7908R/11R/12R Disc:	20% to 80% non-condensing.
Altitude	
Operating:	To 4.6 km (15,000 ft).
Non-operating:	To 15.3 km (50,000 ft).
Vibration and Shock:	HP 1000 A700-Series products are type tested for normal shipping and handling shock and vibration. (Contact factory for review of any application that requires operation under continuous vibration.)

This section describes the bootstrap loaders, the Virtual Control Panel (VCP) program, and the central processor registers accessible to the programmer.

2-1. HARDWARE REGISTERS

The processor cards have several working registers that can be selected for display and modification via the Virtual Control Panel program. (Interface card registers are described in Section VII of this manual and in the interface card reference manuals.) The functions of these processor card registers are described in the following paragraphs.

2-2. A-REGISTER

The A-register is a 16-bit accumulator that holds the results of arithmetic and logical operations performed by programmed instructions. This register can be addressed directly by any memory reference instruction as location 000000 (octal), thus permitting interrelated operations with the B-register (e.g., "add B to A," "compare B with A," etc.) using a single-word instruction.

2-3. B-REGISTER

The B-register is a second 16-bit accumulator which can hold the results of arithmetic and logical operations completely independent of the A-register. The B-register can be addressed directly by any memory reference instruction as location 000001 (octal) for interrelated operations with the A-register.

2-4. P-REGISTER

The 15-bit P-register holds the address of the next instruction to be fetched from memory.

2-5. EXTEND (E) REGISTER

The one-bit extend (E) register is used by rotate instructions to link the A- and B-registers or to indicate a carry from the most-significant bit (bit 15) of the A- or B-register by an add instruction or an increment instruction. This is of significance primarily for multiple-precision arithmetic operations. If already set (logic 1), the extend bit cannot be cleared by a carry. However, the extend bit can be selectively set, cleared, complemented, or tested by programmed instructions.

2-6. OVERFLOW (O) REGISTER

The one-bit overflow (O) register is used to indicate that an add instruction, divide instruction, or an increment instruction referencing the A- or B-register has caused (or will cause) the accumulators to exceed the maximum positive or negative number that can be contained in these registers. The overflow bit can be selectively set, cleared, or tested by programmed instructions.

2-7. CENTRAL INTERRUPT REGISTER

The central interrupt register is a six-bit register that holds the select code of the last interface card or internal condition whose interrupt request was serviced.

2-8. VIOLATION REGISTER

The violation register is a 15-bit register that records the logical address of any fetched instruction that violates memory protection rules.

2-9. PARITY ERROR REGISTER

The 24-bit parity error register stores the physical address of the last memory location that caused a parity error.

2-10. INTERRUPT SYSTEM REGISTER

The interrupt system register is a one-bit register that indicates the status of the interrupt system. When set (logic 1), the interrupt system is enabled; when cleared (0), the interrupt system is disabled.

2-11. X- AND Y-REGISTERS

These two 16-bit registers, designated X and Y, are accessed through the use of the 32 index register instructions and two jump instructions described in Section III.

2-12. WMAP-REGISTER

This 16-bit register holds the logical map numbers used for memory references by Dynamic Mapping System instructions. (The DMS is described in Section IV.)

There are two virtual registers, M and T, that are created by the Virtual Control Panel program and which can be accessed, via the VCP, to examine or change a program in memory or to manually create a program in memory.

2-14. M-REGISTER

The M-register holds the address of the memory cell currently being read from or written into by the Virtual Control Panel.

2-15. T-REGISTER

The T-register indicates the contents of the memory location currently pointed to by the M-register.

2-16. CONTROLS AND INDICATORS

Operator controls and indicators for an A700 computer

system are described in the appropriate system Getting

Started Manual.

On the HP 2137A A700 Computer there is only one operator control: a line-power switch on the rear panel. This two-position switch controls the application of ac line power to the computer power supply and ventilating fans. Light-emitting diodes (LEDs) on the processor frontplane provide indications for the computer self-test.

2-17. SELF-TEST

The self-test consists of two test programs (Test 1 and Test 2) that automatically execute each time the computer is powered up and which provide a quick, convenient check of basic computer operation. (Also, the self-test can be executed by pressing the Reset switch on the frontplane.) If either test program fails, the computer will not operate. Successful completion of the self-test is followed immediately by execution of either a bootstrap loader, the Virtual Control Panel program, or a program sustained in memory by an optional battery pack, as preselected by the user.

Test 1 is a microprogram stored in PROM on the Lower Processor card and executes immediately upon power up. It checks logic and registers on the processor cards and checks the PROM on the memory controller card. If the Floating Point Processor card is installed, it is also checked by Test 1. On successful completion Test 2 is started. If Test 1 detects a failure of a processor card or the memory controller, it stops executing and the frontplane LEDs indicate a failure code. If a floating point processor error is detected, it is indicated by the LEDs but self-test execution continues so that diagnostic testing can be performed. Test 1 execution time is negligible. (Each

processor card test and the memory controller test can be looped for troubleshooting. Refer to the computer installation and service manual for more information.)

Test 2 is an assembly language program stored in PROM on the memory controller card and executes upon successful completion of Test 1. Test 2 checks the computer's basic instruction set, several internal flags, and all the memory. If memory was sustained by the optional battery pack, Test 2 checks it in a non-destructive manner by reading each memory location, thus making a parity check on the data. If a parity error does occur, the location is read again. Soft errors (defined as a parity error only on the first of two reads of a memory location) are reported to the VCP (if present). If memory was not sustained, Test 2 writes all ones to each memory location, and reads back the data; and then writes all zeros and reads back. (The memory is cleared.) Test 2 also checks the I/O Master logic on each interface card to ensure that data transfer, flag, interrupt, and direct memory access (DMA) functions are processed correctly. If Test 2 detects a failure, it stops executing and the frontplane LEDs indicate a failure code. (If a VCP is in the system and the failure does not hinder VCP operation, the VCP program is entered and the failure code is displayed on the VCP.) The LED indication on

successful completion of Test 2 depends on the computer action selected by the Start-Up (BOOT SEL) switches on the frontplane. Test 2 has a maximum execution time of five seconds.

2-18. BOOTSTRAP LOADERS

Bootstrap loading of a program for the A700 computer is provided for by four loaders contained in PROMs on the memory controller card. These are the same PROMs that contain self-test Test 2 and the Virtual Control Panel (VCP) program. The loading devices are disc drive (via HP-IB), PROM storage module, DS/1000-IV network link, HP 264x mini-cartridge tape, and cartridge tape of the HP 7908/11/12 Disc Drive. There are two ways to invoke a loader: auto-boot when power comes up; and by VCP command. Auto-boot can only invoke three of the loaders: disc, PROM module, and DS/1000-IV; the VCP can invoke any of the loaders by a command from the operator. The VCP load commands are discussed later in this section.

2-19. LOADER SELECTION FOR AUTO-BOOT

The selection of an auto-boot is by means of four of the BOOT SEL switches located on the frontplane. These switches, the Start-Up switches, are set during installation and also provide options other than auto-boot selection. When a loader has been selected for auto-boot and the self-test completes, the boot loader executes if memory was lost; or the program in memory executes if memory was sustained by the optional battery backup pack. Refer to Table 2-1 for Start-Up switch settings.

Table 2-1. Start-Up Switch Settings

BOOT SEL SWITCHES*				COMPUTER ACTION
1	2	3	4	
C	C	C	Z	Loop on self-test regardless of error.
C	C	C	Z	Do not use. (Reserved.)
C	C	O	Z	Loop on self-test and stop on error.
C	C	O	Z	Do not use. (Reserved.)
C	O	C	Z	Start boot PROM program at location 30002B.
C	O	O	Z	Run VCP** routine on completion of self-test.
O	C	C	Z	If memory lost (not sustained), run VCP routine; otherwise, restart program (JMP 4B). (Note 2)
O	C	O	Z	If memory lost, load and execute program from PROM card; otherwise, restart program (JMP 4B). (Note 2) (In order to auto-boot from PROM, the card must have select code 22. Equivalent to loader command %BRM.)
O	O	C	Z	If memory lost, load and execute program via HDLC card; otherwise, restart program (JMP 4B). (Note 2) (In order to auto-boot via HDLC, the card must have select code 24. Equivalent to loader command %BDS.)
O	O	O	Z	If memory lost, load and execute program from first file of disc (via HP-IB); otherwise, restart program (JMP 4B). (Note 2) (In order to auto-boot via HP-IB, the HP-IB interface card must have select code 27 and the disc drive must have HP-IB address 2. Equivalent to loader command %BDC.)

*O = open (up); C = closed (down).
Z = C = Normal mode, break enabled.
= O = Break disabled.

**Virtual Control Panel.

Notes: 1. When a loader finishes an auto-boot, it starts execution of the loaded program at location 02.
2. If the auto-restart feature is disabled (frontplane switch M closed), the program cannot restart and the boot loader (or VCP routine) will execute.
3. Do not use any switch combination that is not shown above.

2-20. PROGRAM STARTS

When an auto-boot completes without error, the loaded program starts execution at memory location 02. The loader sets the contents of the A- and B-registers as follows:

a. Cold start (memory not sustained):

1. A = loader command parameters.
2. B = pointer to string area.

b. Auto-restart (memory sustained; execution starts at location 04):

1. A = zero.
 2. B = zero.
- c. %E command from VCP:
1. A = -1.
 2. B = zero.
- d. %B command from VCP:
1. A = loader command parameters.
 2. B = pointer to a string area where:
 - Word 1 = memory size.
 - Word 2 = string length.
 - Word 3 = first word of string.
 - Word n = n-2 word of string.

The memory controller card PROM can be re-entered from a program to boot load and execute the next sequential program from the loading device. The method to re-enter the PROM code, is as follows:

- a. With the disc loader, re-enter to boot load the specific program described by the "ABS" code in the following call back programming sequence.

CLA,CLE,INA	Indicate disc call back — do not suspend
HLT 3,C	Enable PROM
ABS...	HP-IB bus address
ABS...	Device unit number (head for 7906)
ABS...	Absolute starting sector (Vector 1 for 7908/11/12)
ABS...	Cylinder offset (Vector 2 for 7908/11/12)
ABS...	Vector 3 for 7908/11/12

This sequence assumes that the Global Register is set

prior to entry to the loader and that the absolute starting sector is the combined cylinder/head/sector for that drive.

When the load is completed, the loader will start execution in the standard JMP 2 manner. If a suspend after load was specified by the E-register being set when called, the program will halt after the load. In the case of the halt the operator can enter either a %E or a %R to continue. Any error will return to the VCP, if present, or start the original load over.

The 7906 will be accessed in the surface mode only, all other discs will be accessed in the cylinder mode.

2-22. DEVICE PARAMETERS AND MEDIA FORMATS

There is a specific data format for each combination of loader, interface card, loading device, and media. The data formats are described in figure 2-1.

2-23. VIRTUAL CONTROL PANEL

The Virtual Control Panel (VCP) program is an interactive program that enables an external device (such as a terminal) to control the CPU in a manner similar to a conventional computer control panel. That is, it allows the operator to load programs using the loaders, access the various registers (A, B, P, etc., plus I/O card registers), examine or change memory, and control execution of a program. There are two VCP programs stored in PROM on the memory controller card: one is for use with an HP 12005 Interface Card, and the other is for use with an HP 12007/12044 DS/1000-IV Card. Only one interface card in

2-24. VCP PROGRAM OPERATION

The VCP program is executed from PROM as a software program and uses the various machine registers (A, B, etc.) during its execution. Therefore, these registers are automatically saved upon entry to the VCP code. (The save area is in boot RAM on the memory controller card.) Thus, the response to an inquiry is the data that was saved at the time of entry to the program. The exceptions to this are indicated by the absence of an asterisk in table 2-2. When the operator enters the Run (%R) command, the VCP program restores the machine with the current data in the save area and starts execution as specified by the program execution address in the P-register.

The VCP program can be entered in three ways as follows:

- a. After a power-up, PROM execution is directed to the

VCP program instead of a boot load routine;

- b. When the VCP interface card requests a slave cycle to enable the VCP program (e.g., BREAK key pressed on VCP); or
- c. When a HLT (halt) instruction is fetched and one I/O card is enabled for break (otherwise the instruction has no effect).

After a power-up, the total memory size and the amount of error correcting memory are displayed on the VCP screen. The A-register is set to the number of I/O chips that were tested during the self-test. This enables the operator to verify that all installed memory and I/O cards were tested. (Also, except when the self-test detects a duplicate I/O select code and reports it in the B-register, the B-register contains the revision code of the VCP PROMs.) When entered, the VCP displays the basic set of registers (P, A, B, RW, M, and T) and issues the VCP prompt (VCP>) for an operator response. The operator can enter any of the characters or commands listed in Tables 2-2 and 2-3 and the VCP program will respond as indicated in the tables. A carriage return is used to terminate a VCP entry.

After a response to an inquiry the operator can change the data contained in that register or memory location by entering new data; for example:

```
A 001234 4321<cr>
```

operator inputs
<cr> is carriage return

```
A 004321
```

Table 2-2. VCP Characters and Associated Registers

CHARACTER ENTERED	RESPONSE†	MEANING
A*	xxxxxx	A-register contents
B*	xxxxxx	B-register contents
E*	x	E-register contents
G*	x000xx	Global Register (GR) contents and status (bit 15=1 if enabled, 0 if disabled)
I*	x	Interrupt system status (0=off, 1=on)
M*	0xxxxx	Memory address (pointer for T and Ln command)
O*	x	O-register contents
P*	0xxxxx	Program execution address
RS	xxxxxx	Switch register contents
T	0xxxxx xxxxxx	Memory contents pointed to by M
V	xxxxxx	Violation register (memory protect)
X*	xxxxxx	X-register contents
Y*	xxxxxx	Y-register contents
RC	xxxxxx	Central Interrupt Register contents
RD**	xxxxxx xxxxxx	Data for I/O diagnose modes 1 and 2 (refer to paragraph 7-6)
RF**	xxxxxx	I/O flags: Flags 20 thru 24, and Flag 30 (1 = flag set; 0 = flag clear)
RI**	xxxxxx	Interrupt mask register
RP	xxxxx xxxxxx	Parity violation register contents
RW*	xxxxxx	Working map set (WMAP)
R20**	xxxxxx	DMA self-configuration register
R21**	xxxxxx	DMA control register
R22**	xxxxxx	DMA address register
R23**	xxxxxx	DMA count register
R24**	xxxxxx	I/O scratch register
R25**	xxxxxx	I/O scratch register
R26**	xxxxxx	I/O scratch register
R30**	xxxxxx	I/O card data register
R31**	xxxxxx	Optional I/O card register
R32**	xxxxxx	Optional I/O card register
?		Output Help file

† x = octal data.

* Registers that are maintained in the VCP save area of boot RAM.

** Applies only to the I/O card whose select code equals the contents of the Global Register.

NOTE: When a register's contents are changed by the user the new value is returned; if the VCP does not accept a change, the VCP prompt is returned.

Table 2-3. VCP Commands

COMMAND*	MEANING
%B	Load and go (boot). Execute a specified loader routine and start program execution at completion of load. See Figure 2-2 for format.
%C	Clear memory. Set all memory to zero and perform a preset.
%E	Execute. Start execution of program at location P=2 (A-register equals -1 (all ones) and B-register equals 0).
%L	Load. Similar to %B except do not start execution. See Figure 2-2 for format. (%L followed by %R is equivalent to %B.)
%P	Preset. Generate a control reset (CRS) signal on the backplane to initialize all cards.
%R	Run. Set all registers to the appropriate values in the save area and start execution at address specified by the P-register.
%T	Test. Initiate the self-test and return to VCP (memory is sustained but the I/O system is reset).
%U	Start execution at address 30002B in boot memory, the optional user-supplied loader program.
%W	Write. Write to the selected device. (See Figure 2-2 for format.) When writing to a disc drive, the Count and Partial values defined in Figure 2-1 must be in memory locations 00000 and 00001.
D	Decrement. Decrement memory pointer and display the contents of the M- and T-registers. Valid only after T.
Ln	List. List n blocks of eight memory locations starting with location pointed to by the M-register.
N	Next. Same as D except increment the pointer. Valid only after T.
RMxx	List the 32 map registers in the DMS map set specified by xx.
RMxxPyy	Show the value of register yy in map set xx. If a number is input after this command, the register is changed to the new value.
?	Output Help file.

*Must be followed by a carriage return.

Data input is terminated by the operator entering a carriage return. If during an input the program cannot interpret a character, the program will output the characters "!" and then start a new line with the VCP prompt. Entry errors may be corrected by backspacing over them and entering the correct information. During any data input the operator can abort the input by entering a rub-out (DEL). The loader commands, %B, %L, and %W can also be aborted by a rub-out. When entering data into a register, leading zeros may be omitted. If the operator types a question mark, the VCP will output a "help" file that summarizes acceptable command entries.

2-25. LOADER COMMANDS

The loader commands can be entered via the VCP in either of two ways:

- Allow the parameter default values (given in figure 2-1) to be used; or
- Specify all necessary parameters.

The VCP loader command format is shown in figure 2-2. The loader command error codes and their meanings are listed in table 2-4.

2-26. VCP USER CONSIDERATIONS

When using the VCP to debug a program the user should be aware of the following conditions:

- The VCP program uses an interface card and modifies the characteristics of that card. When the VCP program exits, it sets Register 24 on the interface card to all ones to allow software detection of a VCP interaction and, thus, reinitialization of an operation. (This also causes an interrupt if the interrupt system is enabled.) Also, the VCP will leave the card in the output mode with both Flag 30 and Control 30 set.
- The status of the interrupt system (STC 4 (on) or CLC 4 (off)) is not indicated and will remain unchanged unless %P is executed to preset the computer.
- Memory protect is indicated by the sign bit of RW (WMAP register) and may be modified.

2-27. VCP SLAVE FUNCTIONS

The slave feature of an I/O processor chip is used in conjunction with the VCP program. The slave feature enable is read into the I/O chip of the VCP interface card on power-up and cannot be altered until the next power-up condition. After power-up a change in the state of the slave signal causes the I/O chip to generate a slave request on the next instruction fetch. When the request is granted, the I/O chip requests the CPU's current P-register con-

MINI-CARTRIDGE TAPE

Device: HP 264x Terminal

Interface: HP 12005A Asynchronous Serial Interface

Default
Parameters*: 000020

Format: Reads absolute binary file, writes 4k absolute binary block.

Loader: Transmits special escape sequence to invoke a read of a record and does checksum of the data. When writing to tape, a block number is used to specify which 4k-word memory area is to be dumped to tape (0 = 0 to 4k).

If a file number is specified then the program will issue a find file command; if not, the tape is read from where it stands. When writing to the tape, the program will not write a file mark; this allows sequential blocks to be written in a series. There are only two units (0 and 1) on the terminal; if a larger unit number is specified, the result will be unpredictable.

More than 32k words may be loaded into a system from a single cartridge tape.

PROM MODULE

Device: PROM (2k × 8 bits)

Interface: HP 12008A PROM Storage Module

Default
Parameters*: 000022

Format: Count-Partial-Data
Count = number of 64k byte blocks.
Partial = number of words of partial 64k byte block.
Data = 16-bit words, one word per location until Count and Partial are satisfied.

Loader: Uses STC-LIA process to transfer data. The PROM cannot be written to nor does it use the block number of unit number.

*See Figure 2-2 for loader command formats.

Figure 2-1. Loading Device Parameters and Media Formats (Part 1 of 2)

DISC DRIVE

Device: HP 9895, 7902, 7906, 7908, 7910, 7911, or 7912 Disc Drive, or cartridge tape drive of the 7908/11/12 Disc Drive.

Interface: HP 12009A HP-IB Interface

Default
Parameters*: 002027

Format: Count-Partial-Data
 Count† = number of 64k byte blocks.
 Partial† = number of words of partial 64k byte block.
 Data = 16-bit words, one word per location until Count and Partial are satisfied.

Loader: Uses HP-IB protocol to communicate with the disc. The load sequence is:

1. Device clear
2. Status check
3. Read/write 32k words via DMA
4. Status check

Note: Each file starts on a track boundary.

COMPUTER NETWORK

Device: HP 1000 Computer.

Interface: HP 12007A/12044A HDLC Interface.

Default
Parameters*: 000024

Format: Reads absolute binary or memory image files, writes a 32k memory image file.

Loader: Standard handshake using HP distributed system protocol. Block number and unit number are not used.

*See Figure 2-2 for loader command formats.

†The Count and Partial values are stored in memory locations 00000 and 00001, respectively.

Figure 2-1. Loading Device Parameters and Media Formats (Part 2 of 2)

LOADER COMMAND FORMAT

%B/L/W dv ffffbusc text

where:

dv = device type as follows:

DC = disc (cartridge or flexible) via HP-IB

CT = cartridge tape (HP 264x)

RM = PROM card

DS = DS computer network Link

ffff = file number (octal 0 to 77777 only)

b = 4k-word memory block number when writing to cartridge tape; HP-IB bus address of disc drive; otherwise, use 0.

u = unit number (0 to 7) only if used on device. For the HP 7906 Disc Drive, the unit number is the head number. For HP 7908, 7911, or 7912 Disc Drive that includes cartridge tape drive, unit 0 = disc drive and unit 1 = cartridge tape drive.

sc = select code of interface card to be used.

text = file name, or ASCII string to be passed to the program after it is loaded. This is only available with the %B and %L commands.

Note: See Figure 2-1 for default parameters for each loading device.

Note that spaces cannot be used in the command entry. The following formats are all acceptable:

%Bdvtext Device parameters are defaulted; text cannot start with a number.

%Bdvffffbusc No text passed.

%Bdvffffbusctext Text passed.

EXAMPLES:

%BDC Load and start execution of the default program on disc. (Disc parameters defaulted to 002027; see Figure 2-1).

%BDC30 Load and start execution of the default program on the disc at select code 30 and default other parameters.

%LDC27025 Load (but don't execute) and override parameter default values:
file number 2 (i.e., the third file)
HP-IB bus address 7
unit 0
select code 25

%WDC27025 Same as above except write to file 2.

Figure 2-2. Loader Command Format

Table 2-4. VCP Loader Command Errors

ERROR CODE	MEANING
2	Select code less than 20 octal.
3	No card with the select code you specified.
Cartridge Tape Loader Errors	
110	File forward error. Status in B-register.
111	Checksum error.
112	No data before EOF (end of file).
120	Write error. Status in B-register.
PROM Module Loader Errors	
211	End of programs.
212	Bad format.
213	System larger than 32k must start on card boundary.
DS/1000 Loader Errors	
310	Time out after CLC 0. Check select code specified.
311	Checksum error. P file not absolute binary.
312	Time out after download request.
313	Time out after file number.
314	Bad transfer (Central generated). Status in B-register.
315	Time out after buffer request.
316	Time out after count echo.
317	Time out waiting for data.
320	Time out after VCP mode requests a DS write.
321	Central will not accept data. Status in B-register.
Disc Loader Errors	
411	Time out reading disc type. Check HP-IB address.
412	Time out UDC (Universal Device Code) or reading status. Check disc.
413	Status error. Status in B-register.
414	Time out during file mask.
415	Time out during seek.
416	Time out during read or write command.
417	Time out during DMA of data.
420	Parity error during DMA transfer.
421	Time out during FIFO flush.
422	Time out during DSJ (Device Specified Jump) command.
423	Bad DSJ return. Returned value in B-register.
460	Disc not identifiable. Disc ID in B-register.

tents and saves these contents in a register in the I/O chip. The I/O chip then stores the starting address of the VCP program into the CPU's P-register, instructs the CPU to enable the boot PROM, and allows execution to start. The VCP program can be started in several other ways, as follows:

- a. On power-up and after the self-test the VCP program starts execution if it is selected in lieu of a boot loader. This selection may often be used because the loaders can be invoked individually from the VCP.
- b. When a HLT* (halt) instruction is executed the I/O processor chip interprets it in the same manner as a change in the slave enable signal. This allows a program to have breakpoints for debugging purposes. (Note that a HLT instruction is not executed but causes a memory protect interrupt if memory protect is enabled.)

During execution of the VCP program, access to the P-save register in the I/O chip is accomplished with LIA/B 3 and OTA/B 3 (without the instruction's Flag bit set). It should also be noted that the I/O chip will not execute a slave request until an STC 2 (enable break feature) instruction has been executed. This prevents re-entry of the VCP program once it has been entered.

During the self-test, the starting address of the VCP program is assigned to the break-enabled I/O card by an OTA/B 3,C* instruction with the A- or B-register set to the address. This address can also be read back with an LIA/B 3,C* instruction.

* If break is not enabled on any I/O card, then the instruction has no effect.

This section describes the software data formats and the base set machine-language instruction coding required to operate the computer and its associated input/output system. This section also describes the optional floating point instructions, Scientific Instruction Set instructions, and Vector Instruction Set instructions. Machine-language instruction coding for the Dynamic Mapping System is presented in Section IV.

3-1. DATA FORMATS

As shown in Figure 3-1, the basic data format is a 16-bit word in which bit positions are numbered from 0 through 15 in order of increasing significance. Bit position 15 of the data format is used for the sign bit; a logic 0 in this position indicates a positive number and a logic 1 in this position indicates a negative number. The data is assumed to be a whole number and the binary point is therefore assumed to be to the right of the number.

The basic word can also be divided into two 8-bit bytes or combined to form a 32-bit double integer. The byte format is used for character-oriented input/output devices; packing two bytes of data into one 16-bit word is accomplished by software drivers. In I/O operations, the higher-order byte (byte 1) is the first to be transferred.

The double integer format is used for extended arithmetic in conjunction with the extended arithmetic instructions described under paragraphs 3-21 and 3-22. Bit position 15 of the most-significant word is the sign bit and the binary point is assumed to be to the right of the least-significant word. The integer value is expressed by the remaining 31 bits.

The two floating point formats in Figure 3-1 are used with floating point software. Bit position 15 of the most-significant word is the mantissa sign and bit position 0 of the least-significant word is the exponent sign. Bits 1 through 7 of the least-significant word express the exponent and the remaining bits express the mantissa. A single precision floating point number is made up of a 23-bit mantissa (fraction) and sign and a 7-bit exponent and sign, thus providing six significant decimal digits in the mantissa. A double precision floating point number is made up of a 55-bit mantissa and a 7-bit exponent and sign, thus providing 16 significant decimal digits in the mantissa. If either the mantissa or the exponent is

negative, that part must be stored in two's complement form. The number must be in the approximate range of 10^{-38} to 10^{+38} . When loaded into the accumulators, the A-register contains the most-significant word and the B-register contains the least-significant word.

Figure 3-1 also illustrates the octal notation for both single-length (16-bit) and double-length (32-bit) words. Each group of three bits, beginning at the right, is combined to form an octal digit. A single-length (16-bit) word can therefore be fully expressed by six octal digits and a double-length (32-bit) word can be fully expressed by 11 octal digits. Octal notation is not shown for byte or floating point formats, since bytes normally represent characters and floating point numbers are given in decimal form.

The range of representable numbers for single integer data is +32,767 to -32,768 (decimal) or +77,777 to -100,000 (octal). The range of representable numbers for double integer data is +2,147,483,647 to -2,147,483,648 (decimal) or +17,777,777,777 to -20,000,000,000 (octal).

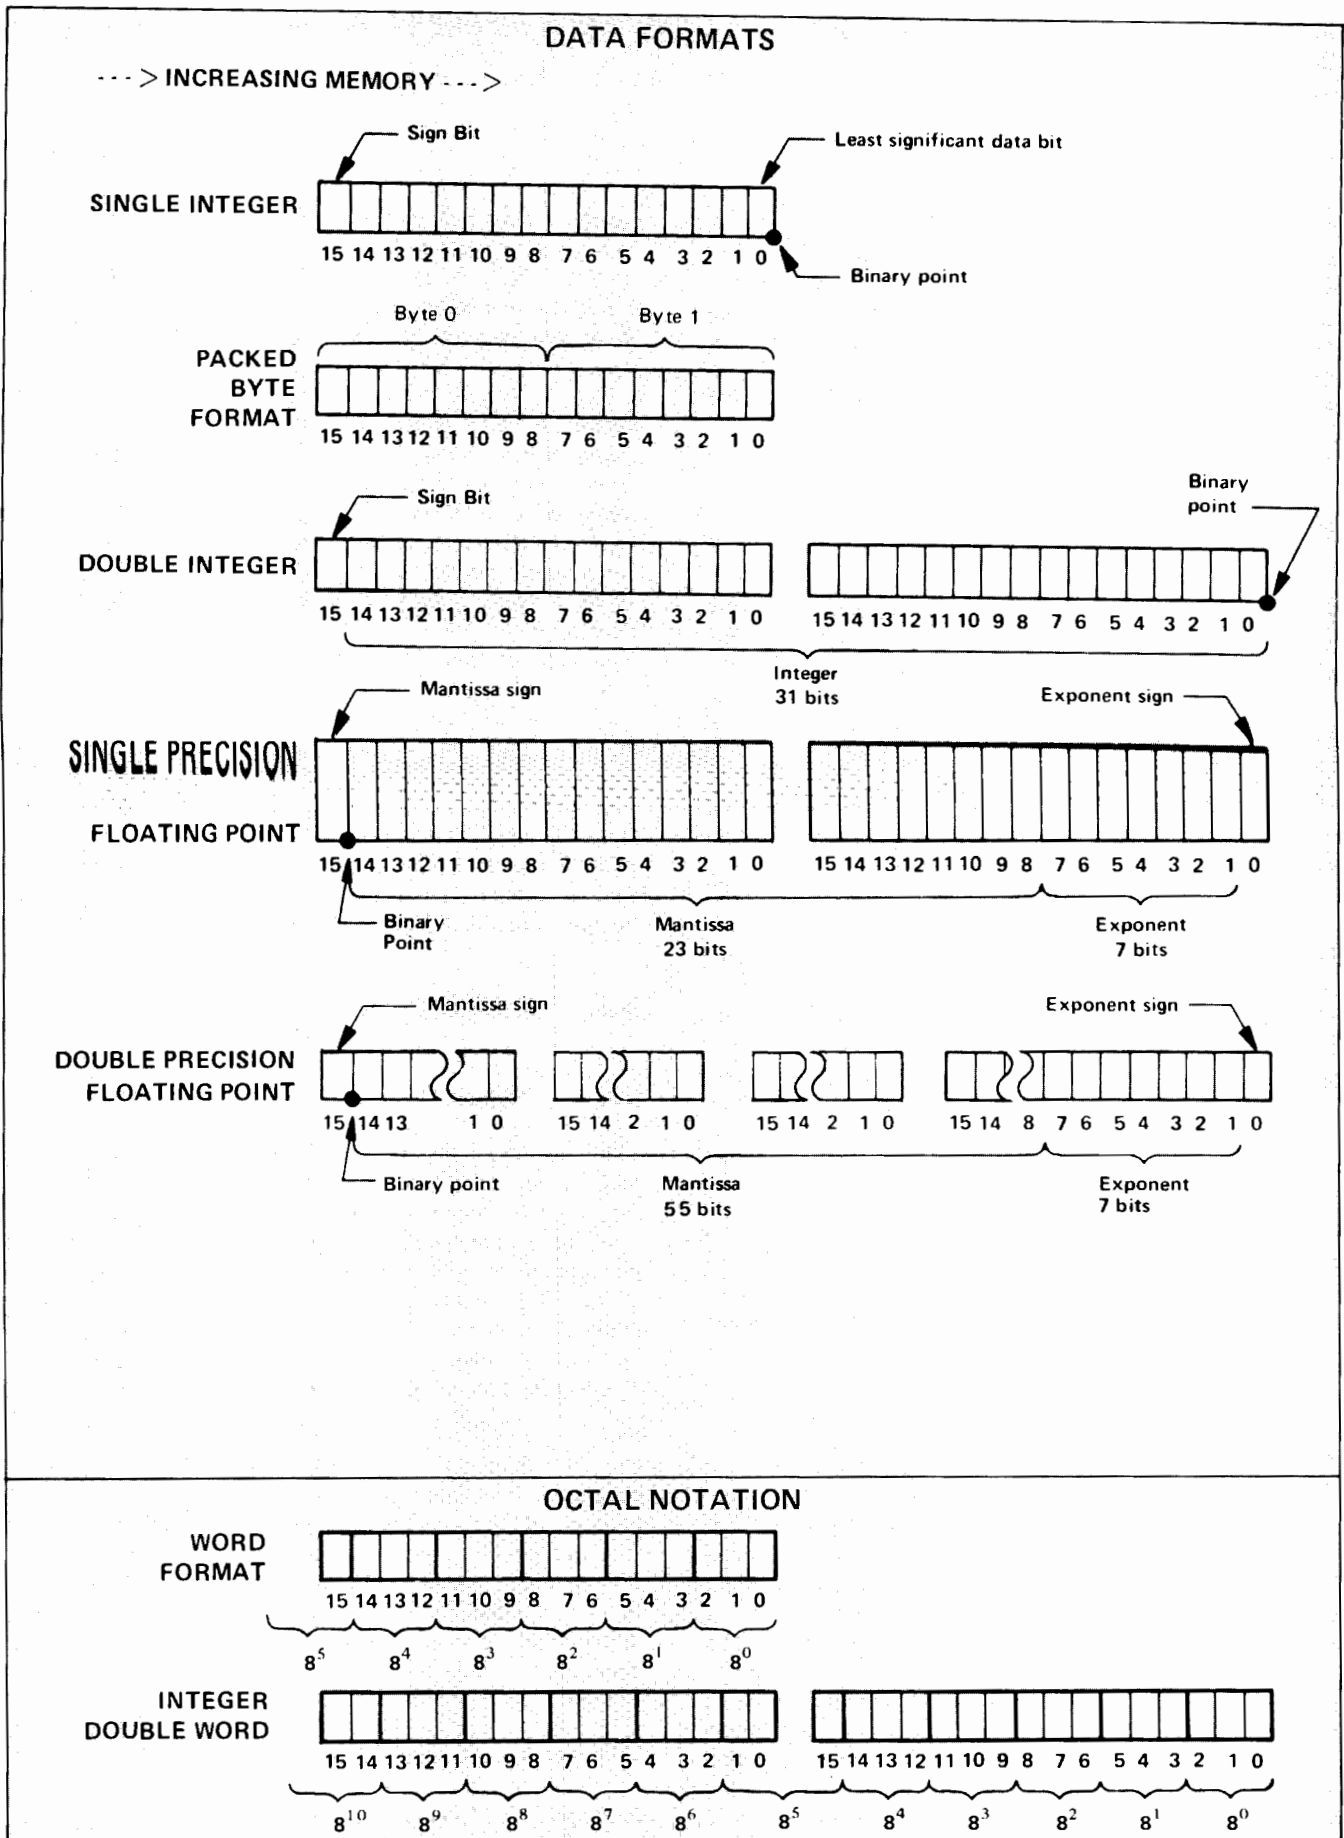
3-2. MEMORY ADDRESSING

3-3. PAGING

The computer memory is logically divided into pages of 1,024 words each. A page is defined as the largest block of memory that can be directly addressed by the address bits of a single-length memory reference instruction. (Refer to paragraph 3-8.) These memory reference instructions use 10 bits (bits 0 through 9) to specify a memory address; thus, the page size is 1,024 locations (2000 octal). Octal addresses for each page, up to a maximum memory size of 32k words, are listed in Table 3-1.

Provision is made to directly address one of two pages: page zero (the base page consisting of locations 00000 through 01777) and the current page (the page in which the instruction itself is located). Memory reference instructions reserve bit 10 to specify one or the other of these two pages. To address locations on any other page, indirect addressing is used as described in following paragraphs. Page references are specified by bit 10 as follows:

- a. Logic 0 = Page Zero (Z).
- b. Logic 1 = Current Page (C).



2270-2

Figure 3-1. Data Formats and Octal Notation

Table 3-1. Memory Paging

MEMORY SIZE	PAGE	OCTAL ADDRESSES
16K ↓	0	00000 to 01777
	1	02000 to 03777
	2	04000 to 05777
	3	06000 to 07777
	4	10000 to 11777
	5	12000 to 13777
	6	14000 to 15777
	7	16000 to 17777
	8	20000 to 21777
	9	22000 to 23777
	10	24000 to 25777
	11	26000 to 27777
	12	30000 to 31777
	13	32000 to 33777
	14	34000 to 35777
15	36000 to 37777	
	16	40000 to 41777
	17	42000 to 43777
	18	44000 to 45777
	19	46000 to 47777
	20	50000 to 51777
	21	52000 to 53777
	22	54000 to 55777
	23	56000 to 57777
	24	60000 to 61777
	25	62000 to 63777
	26	64000 to 65777
	27	66000 to 67777
	28	70000 to 71777
	29	72000 to 73777
	30	74000 to 75777
	31	76000 to 77777

3-4. DIRECT AND INDIRECT ADDRESSING

All memory reference instructions reserve bit 15 to specify either direct or indirect addressing. For single-length memory reference instructions, bit 15 of the instruction word is used; for extended arithmetic memory reference instructions, bit 15 of the address word is used. Indirect addressing uses the address part of the instruction to access another word in memory, which is taken as the new memory reference for the same instruction. This new address word is a full 16 bits long: 15 address bits plus another direct/indirect bit. The 15-bit length of the address permits access to any location in memory. If bit 15 again specifies indirect addressing, still another address is obtained; thus, multistep indirect addressing may be done to any number of levels. The first address obtained that

does not specify another indirect level becomes the effective address for the instruction. Direct or indirect addressing is specified by bit 15 as follows:

- a. Logic 0 = Direct (D).
- b. Logic 1 = Indirect (I).

3-5. MEMORY MAPPING

Memory mapping is a standard feature of the A700 computer and is used to access all locations of main memory. Memory mapping is provided by the Dynamic Mapping System described in Section IV.

3-6. RESERVED MEMORY LOCATIONS

The first 64 memory locations of the base page (octal address 00000 through 00077) are reserved as listed in Table 3-2. The first two locations are reserved as addresses for the two 16-bit accumulators (the A- and B-register). If options or input/output devices corresponding to locations 00020 through 00077 are not included in the system configuration, these locations can be used for programming purposes. The last 64 locations of the physical base page (octal addresses 1700 through 1777) are reserved for use by the Virtual Control Panel program for the string area.

Table 3-2. Reserved Memory Locations

MEMORY LOCATION	PURPOSE
00000	A-register address.
00001	B-register address.
00002-00003	Reserved.
00004	Power-fail interrupt.
00005	Memory parity interrupt.
00006	Time base generator interrupt.
00007	Memory protect interrupt.
00010	Unimplemented instruction interrupt.
00011-00017	Reserved.
00020-00077	Interrupt locations corresponding to interface card select codes.
01700-01777	VCP program string area.

3-7. NONEXISTENT MEMORY

Nonexistent memory is defined as those locations not physically implemented in the machine. Any attempt to write into a nonexistent memory location will be ignored (no operation). Any attempt to read from a nonexistent memory location will return an all-ones word (177777 octal); no parity error occurs.

3-8. BASE SET INSTRUCTION FORMATS

The base set of instructions are classified according to format. The six formats used are illustrated in Figure 3-2 and described in the following paragraphs except for the Dynamic Mapping System (DMS) instructions, which are described in Section IV. In all cases where a single bit is used to select one of two cases (e.g., D/I), the choice is made by coding a logic 0 or logic 1, respectively.

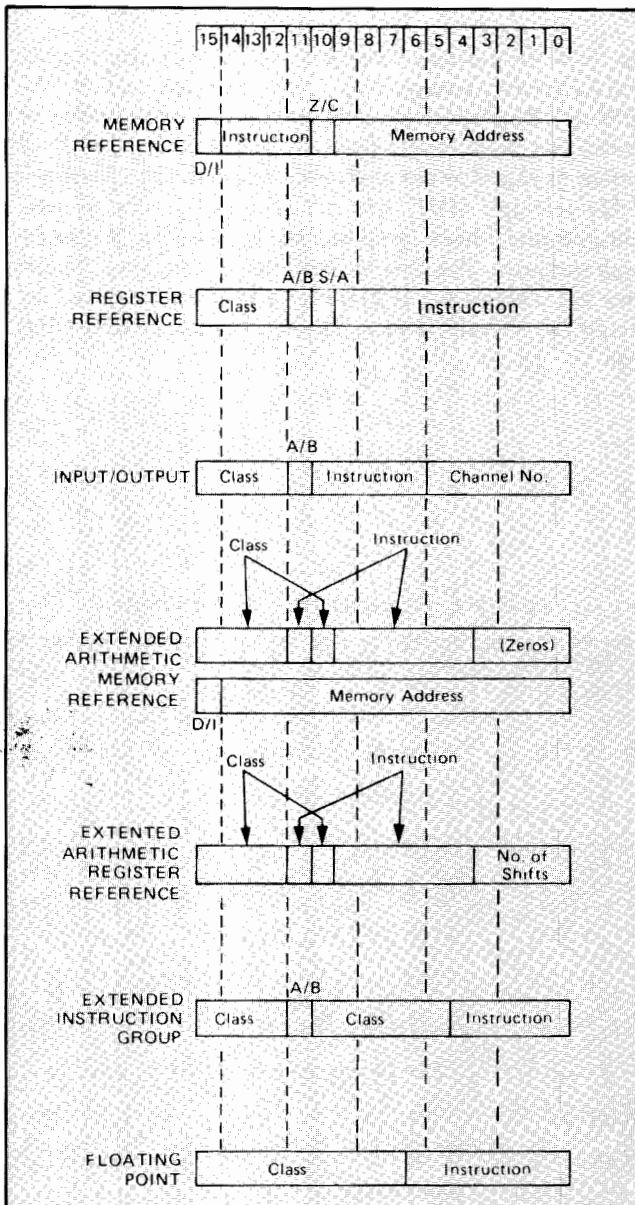


Figure 3-2. Base Set Instruction Formats

3-9. MEMORY REFERENCE INSTRUCTIONS

This class of instructions, which combines an instruction code and a memory address into one 16-bit word, is used to execute some function involving data in a specific memory location. Examples are storing, retrieving, and combining memory data to and from the accumulators (A- and B-registers) or causing the program to jump to a specified location in memory.

The memory cell referenced (i.e., the absolute address) is determined by a combination of 10 memory address bits (0 through 9) in the instruction word and 5 bits (10 through 14) assumed from the current contents of the M-register. This means that memory reference instructions can directly address any word in the current page; additionally, if the instruction is given in some location other than the base page (page zero), bit 10 (Z/C) of the instruction doubles the addressing range to 2,048 locations by allowing the selection of either page zero or the current page. (This causes bits 10 through 14 of the address contained in the M-register to be set to zero instead of assuming the current contents of the M-register.) This feature provides a convenient linkage between all pages of memory, since page zero can be reached directly from any other page.

As discussed under paragraph 3-4, bit 15 is used to specify direct or indirect memory addressing. Note also that since the A- and B-registers are addressable, any single-word memory reference instruction can apply to either of these registers as well as to memory cells. For example, an ADA 0001 instruction adds the contents of the B-register (address 0001) to the contents currently held in the A-register; specify page zero for these operations since the addresses of the A- and B-registers are on page zero.

3-10. REGISTER REFERENCE INSTRUCTIONS

In general, the register reference instructions manipulate bits in the A-register, B-register, and E-register; there is no reference to memory. This group includes 39 basic instructions which may be combined to form a one-word multiple instruction that can operate in various ways on the contents of the A-, B-, and E-registers. These 39 instructions are divided into two subgroups: the shift/rotate group (SRG) and the alter/skip group (ASG). The appropriate subgroup is specified by bit 10 (S/A). Typical operations are clear and/or complement a register, conditional skips, and register increment.

3-11. INPUT/OUTPUT INSTRUCTIONS

The input/output instructions use bits 6 through 11 for a variety of I/O instructions and bits 0 through 5 to apply the instructions either to a specific I/O channel (if the Global Register is disabled) or to an I/O card register. This provides the means of controlling all peripherals connected to the I/O channels and for transferring data to and from these peripherals. Included also in this group are

instructions to control the interrupt system, overflow bit, and computer halt.

3-12. EXTENDED ARITHMETIC MEMORY REFERENCE INSTRUCTIONS

As the single-word memory reference instruction described previously, the extended arithmetic memory reference instructions include an instruction code and a memory address. In this case, however, two words are required. The first word specifies the extended arithmetic class (bits 12 through 15 and 10) and the instruction code (bits 4 through 9 and 11); bits 0 through 3 are not needed and are coded with zeros. The second word specifies the memory address of the operand. Since the full 15 bits are used for the address, this type of instruction may directly address any location in memory. As with all memory reference instructions, bit 15 is used to specify direct or indirect addressing. Operations performed by this class of instructions are integer multiply and divide (using double-length product and dividend) and double load and double store.

3-13. EXTENDED ARITHMETIC REGISTER REFERENCE INSTRUCTIONS

This class of instructions provides long shifts and rotates on the combined contents of the A- and B-registers. Bits 12 through 15 and 10 identify the instruction class; bits 4 through 9 and 11 specify the direction and type of shift; and bits 0 through 3 control the number of shifts, which can range from 1 to 16 places.

3-14. EXTENDED INSTRUCTIONS

The extended instructions include index register instructions, bit and byte manipulation instructions, and move and compare instructions. Instructions comprising the extended instruction group are one, two, or three words in length. The first word is always the instruction code; operand addresses are given in the words following the instruction code or in the A- and B-registers. The operand addresses are 15 bits long, with bit 15 (most-significant bit) generally indicating direct or indirect addressing.

3-15. FLOATING POINT INSTRUCTIONS

The floating point instructions allow addition, subtraction, multiplication, and division of floating point quantities. Conversion routines are provided for transforming numerical integer representations to/from floating point representations.

3-16. LANGUAGE INSTRUCTION SET

The language instruction set performs several frequently used high-level language operations, including parameter passing, array address calculations, and floating point conversion, packing, rounding and normalizing.

3-17. DOUBLE INTEGER INSTRUCTIONS

The double integer instructions allow arithmetic and test operations on 32-bit quantities. The data format for double integer values is shown in Figure 3-1.

3-18. VIRTUAL MEMORY INSTRUCTIONS

The virtual memory instructions perform accesses to Virtual Memory and the Extended Memory Area, which are extensions of logical memory.

3-19. OPERATING SYSTEM INSTRUCTIONS

The operating system instructions provide instructions for ascertaining the CPU and firmware identification, and instructions for interrupt conditions.

3-20. BASE SET INSTRUCTION CODING

Machine language coding for the base set of instructions are provided in following paragraphs. Definitions for these instructions are grouped according to the instruction type: memory reference, register reference, input/output, extended arithmetic memory reference, and extended arithmetic register reference.

Directly above each definition is a diagram showing the machine language coding for that instruction. The gray shaded bits code the instruction type and the green shaded bits code the specific instruction. Unshaded bits are further defined in the introduction to each instruction type. The mnemonic code and instruction name are included above each diagram.

In all cases where an additional bit is used to specify a secondary function (D/I, Z/C, or H/C), the choice is made by coding a logic 0 or logic 1, respectively. In other words, a logic 0 codes D (direct addressing), Z (zero page), or H (hold flag); a logic 1 codes I (indirect addressing), C (current page), or C (clear flag).


3-21. MEMORY REFERENCE INSTRUCTIONS

The following 14 memory reference instructions execute a function involving data in memory. Bits 0 through 9 specify the affected memory location on a given memory page or, if indirect addressing is specified, the next address to be referenced. Indirect addressing may be continued to any number of levels; when bit 15 (D/I) is a logic 0 (specifying direct addressing), that location will be taken as the effective address. The A- and B-registers may be addressed as locations 00000 and 00001 (octal), respectively.


If bit 10 (Z/C) is a logic 0, the memory address is on page zero; if bit 10 is a logic 1, the memory address is on the




current page. If the A- or B-register is addressed, bit 10 must be a logic 0 to specify page zero, unless the current page is page zero.




Adds the contents of the addressed memory location to the contents of the A-register. The sum remains in the A-register and the contents of the memory cell are unaltered. The result of this addition may set the extend bit or the overflow bit. (Extend and overflow examples are illustrated on page A-13.)




Adds the contents of the addressed memory location to the contents of the B-register. The sum remains in the B-register and the contents of the memory cell are unaltered. The result of this addition may set the extend bit or the overflow bit. (Extend and overflow examples are illustrated on page A-13.)




Combines the contents of the addressed memory location and the contents of the A-register by performing a logical "and" operation. The contents of the memory cell are unaltered.




Compares the contents of the addressed memory location with the contents of the A-register. If the two 16-bit words are not identical, the next instruction is skipped; i.e., the P-register advances two counts instead of one count. If the two words are identical, the next sequential instruction is executed. Neither the A-register contents nor memory cell contents are altered.




Compares the contents of the addressed memory location with the contents of the B-register. If the two 16-bit words are not identical, the next instruction is skipped; i.e., the P-register advances two counts instead of one count. If the two words are identical, the next sequential instruction is executed. Neither the B-register contents nor memory cell contents are altered.




Combines the contents of the addressed memory location and the contents of the A-register by performing a logical "inclusive or" operation. The contents of the memory cell are unaltered.




Adds one to the contents of the addressed memory location. If the result of this operation is zero (memory contents incremented from 177777 to 000000), the next instruction is skipped; i.e., the P-register is advanced two counts instead of one count. If the result of this operation is not zero, the next sequential instruction is executed. In either case, the incremented value is written back into the memory cell.




Transfers control to the addressed memory location. That is, a JMP causes the P-register count to set according to the memory address portion of the JMP instruction so that the next instruction will be read from that location.




This instruction, executed in location P (P-register count), causes the computer control to jump unconditionally to the memory location (m) specified by the memory address portion of the JSB instruction. The contents of the P-register plus one (return address) is stored in memory location m, and the next instruction to be executed will be that contained in the next sequential memory location (m + 1). A return to the main program sequence at P + 1 will be effected by a JMP indirect through location m.




Loads the contents of the addressed memory location into the A-register. The contents of the memory cell are unaltered.




Loads the contents of the addressed memory location into the B-register. The contents of the memory cell are unaltered.



Stores the contents of the A-register in the addressed memory location. The previous contents of the memory cell are lost; the A-register contents are unaltered.



Stores the contents of the B-register in the addressed memory location. The previous contents of the memory cell are lost; the B-register contents are unaltered.



Combines the contents of the addressed memory location and the contents of the A-register by performing a logical "exclusive or" operation. The contents of the memory cell are unaltered.

3-22. REGISTER REFERENCE INSTRUCTIONS

The 39 register reference instructions execute functions on data contained in the A-register, B-register, and E-register. These instructions are divided into two groups: the shift/rotate group (SRG) and the alter/skip group (ASG). In each group, several instructions may be combined into one word. Since the two groups perform separate and distinct functions, instructions from the two groups cannot be mixed. Unshaded bits in the coding diagrams are available for combining other instructions.

3-23. SHIFT/ROTATE GROUP. The 20 instructions in the shift/rotate group (SRG) are defined first; this group is specified by setting bit 10 to a logic 0. A comparison of the various shift/rotate functions are illustrated in Figure 3-3. Rules for combining instructions in this group are as follows (refer to Table 3-3):

- a. Only one instruction can be chosen from each of the two multiple-choice columns.
- b. References can be made to either the A-register or B-register, but not both.
- c. Sequence of execution is from left to right.
- d. In machine code, use zeros to exclude unwanted microinstructions.
- e. Code a logic 1 in bit position 9 to enable shifts or rotates in the first position; code a logic 1 in bit position 4 to enable shifts or rotates in the second position.
- f. The extend bit is not affected unless specifically stated. However, if a "rotate-with-E" instruction (ELA, ELB, ERA, or ERB) is coded but disabled by a logic 0 in bit position 9 and/or position 4, the E-register will be updated even though the A- or B-register contents are not affected; to avoid this situation, code a "no operation" (three zeros) in the first and/or second positions.

Table 3-3. Shift/Rotate Group Combining Guide

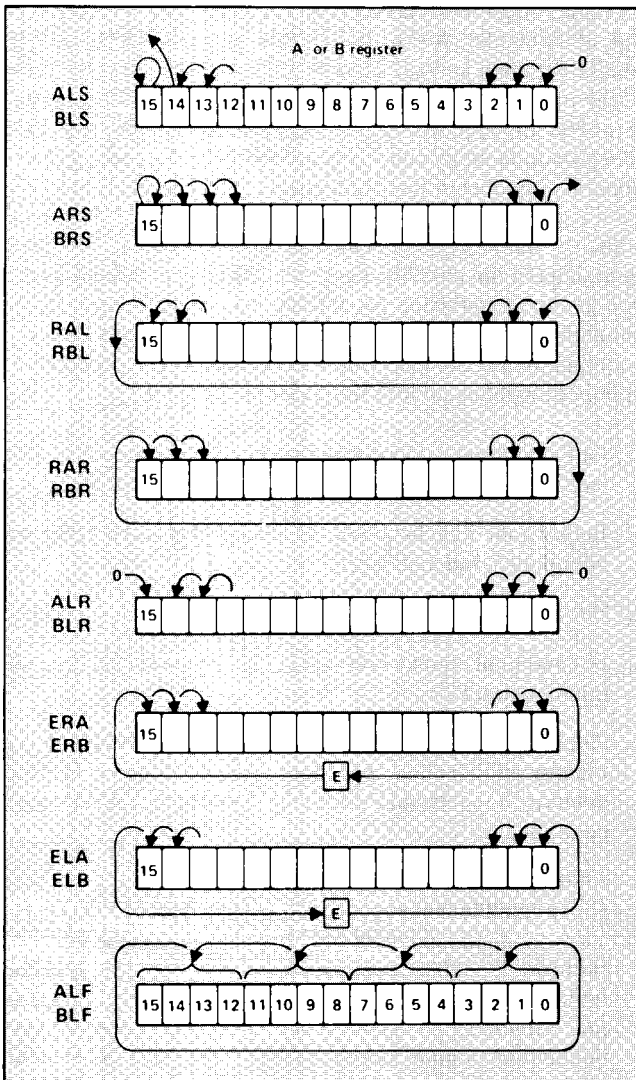
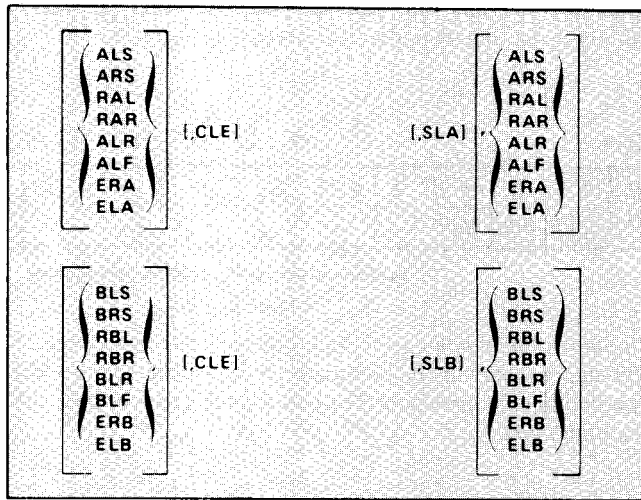





Figure 3-3. Shift and Rotate Functions




Rotates the A-register contents (all 16 bits) left four places. Bits 15, 14, 13, and 12 rotate around to bit positions 3, 2, 1, and 0, respectively. Equivalent to four successive RAL instructions.




Shifts the A-register contents left one place and clears sign bit 15.




Arithmetically shifts the A-register contents left one place, 15 magnitude bits only; bit 15 (sign) is not affected. The bit shifted out of bit position 14 is lost; a logic 0 replaces vacated bit position 0.




Arithmetically shifts the A-register contents right one place, 15 magnitude bits only; bit 15 (sign) is not affected. A copy of the sign bit is shifted into bit position 14; the bit shifted out of bit position 0 is lost.




Rotates the B-register contents (all 16 bits) left four places. Bits 15, 14, 13, and 12 rotate around to bit positions 3, 2, 1, and 0, respectively. Equivalent to four successive RBL instructions.




Shifts the B-register contents left one place and clears sign bit 15.




Arithmetically shifts the B-register contents left one place, 15 magnitude bits only; bit 15 (sign) is not affected. The bit shifted out of bit position 14 is lost; a logic 0 replaces vacated bit position 0.




Arithmetically shifts the B-register contents right one place, 15 magnitude bits only; bit 15 (sign) is not affected. A copy of the sign bit is shifted into bit position 14; the bit shifted out of bit position 0 is lost.




Clears the E-register; i.e., the extend bit becomes a logic 0.




Rotates the E-register content left with the A-register contents (one place). The E-register content rotates into bit position 0; bit 15 rotates into the E-register.




Rotates the E-register content left with the B-register contents (one place). The E-register content rotates into bit position 0; bit 15 rotates into the E-register.




Rotates the E-register content right with the A-register contents (one place). The E-register content rotates into bit position 15; bit 0 rotates into the E-register.




Rotates the E-register content right with the B-register contents (one place). The E-register content rotates into bit position 15; bit 0 rotates into the E-register.




This all-zeros instruction causes a no-operation cycle.




Rotates the A-register contents left one place (all 16 bits). Bit 15 rotates into bit position 0.




Rotates the A-register contents right one place (all 16 bits). Bit 0 rotates into bit position 15.




Rotates the B-register contents left one place (all 16 bits). Bit 15 rotates into bit position 0.



Rotates the B-register contents right one place (all 16 bits). Bit 0 rotates into bit position 15.



Skips the next instruction if the least-significant bit (bit 0) of the A-register is a logic 0.




Skips the next instruction if the least-significant bit (bit 0) of the B-register is a logic 0.

3-24. ALTER/SKIP GROUP. The 19 instructions comprising the alter/skip group (ASG) are defined next. This group is specified by setting bit 10 to a logic 1. Rules for combining instructions are as follows (refer to Table 3-4):


- Only one instruction can be chosen from each of the two multiple-choice columns.
- References can be made to either the A-register or B-register, but not both.
- Sequence of execution is from left to right.
- If two or more skip functions are combined, the skip function will occur if either or both conditions are met. One exception exists: refer to the RSS instruction.
- In machine code, use zeros to exclude unwanted instructions.

Table 3-4. Alter/Skip Group Combining Guide

<table border="1"> <tr><td>(CLA)</td></tr> <tr><td>(CMA)</td></tr> <tr><td>(CCA)</td></tr> </table>	(CLA)	(CMA)	(CCA)	.SEZ	<table border="1"> <tr><td>(CLE)</td></tr> <tr><td>(CME)</td></tr> <tr><td>(CCE)</td></tr> </table>	(CLE)	(CME)	(CCE)	.SSA .SLA .INA .SZA .RSS
(CLA)									
(CMA)									
(CCA)									
(CLE)									
(CME)									
(CCE)									
<table border="1"> <tr><td>(CLB)</td></tr> <tr><td>(CMB)</td></tr> <tr><td>(CCB)</td></tr> </table>	(CLB)	(CMB)	(CCB)	.SEZ	<table border="1"> <tr><td>(CLE)</td></tr> <tr><td>(CME)</td></tr> <tr><td>(CCE)</td></tr> </table>	(CLE)	(CME)	(CCE)	.SSB .SLB .INB .SZB .RSS
(CLB)									
(CMB)									
(CCB)									
(CLE)									
(CME)									
(CCE)									



Clears and complements the A-register contents; i.e., the contents of the A-register become 177777 (octal). This is the two's complement form of -1.



Clears and complements the B-register contents; i.e., the contents of the B-register become 177777 (octal). This is the two's complement form of -1.

CCE CLEAR AND COMPLEMENT E

15	14	13	12	11	10	9	8	7	6	5	4	3	2	1	0
0	0	0	0		1			1	1						

Clears and complements the E-register content (extend bit); i.e., the extend bit becomes a logic 1.

CLA CLEAR A

15	14	13	12	11	10	9	8	7	6	5	4	3	2	1	0
0	0	0	0	0	1	0	1								

Clears the A-register; i.e., the contents of the A-register become 000000 (octal).

CLB CLEAR B

15	14	13	12	11	10	9	8	7	6	5	4	3	2	1	0
0	0	0	0	1	1	0	1								

Clears the B-register; i.e., the contents of the B-register become 000000 (octal).

CLE CLEAR E

15	14	13	12	11	10	9	8	7	6	5	4	3	2	1	0
0	0	0	0		1			0	1						

Clears the E-register; i.e., the extend bit becomes a logic 0.

CMA COMPLEMENT A

15	14	13	12	11	10	9	8	7	6	5	4	3	2	1	0
0	0	0	0	0	1	1	0								

Complements the A-register contents (one's complement).

CMB COMPLEMENT B

15	14	13	12	11	10	9	8	7	6	5	4	3	2	1	0
0	0	0	0	1	1	1	0								

Complements the B-register contents (one's complement).

CME COMPLEMENT E

15	14	13	12	11	10	9	8	7	6	5	4	3	2	1	0
0	0	0	0		1			1	0						

Complements the E-register content (extend bit).

INA INCREMENT A

15	14	13	12	11	10	9	8	7	6	5	4	3	2	1	0
0	0	0	0	0	1									1	

Increments the A-register by one. The overflow bit will be set if an increment of the largest positive number (077777 octal) is made. The extend bit will be set if an all-ones word (177777 octal) is incremented.

INB INCREMENT B

15	14	13	12	11	10	9	8	7	6	5	4	3	2	1	0
0	0	0	0	1	1									1	

Increments the B-register by one. The overflow bit will be set if an increment of the largest positive number (077777 octal) is made. The extend bit will be set if an all-ones word (177777 octal) is incremented.

RSS REVERSE SKIP SENSE

15	14	13	12	11	10	9	8	7	6	5	4	3	2	1	0
0	0	0	0		1										1

Skip occurs for any of the following skip instructions, if present, when the non-zero condition is met. An RSS without a skip instruction in the word causes an unconditional skip. If a word with RSS also includes both SSA and SLA (or SSB and SLB), bits 15 and 0 must both be logic 1's for a skip to occur; in all other cases, a skip occurs if one or more skip conditions are met.

SEZ SKIP IF E IS ZERO

15	14	13	12	11	10	9	8	7	6	5	4	3	2	1	0
0	0	0	0		1					1					

Skips the next instruction if the E-register content (extend bit) is a logic 0.

SLA SKIP IF LSB OF A IS ZERO

15	14	13	12	11	10	9	8	7	6	5	4	3	2	1	0
0	0	0	0	0	1								1		

Skips the next instruction if the least-significant bit (bit 0) of the A-register is a logic 0; i.e., skips if an even number is in the A-register.

SLB SKIP IF LSB OF B IS ZERO

15	14	13	12	11	10	9	8	7	6	5	4	3	2	1	0
0	0	0	0	1	1								1		

of the B-register is a logic 0; i.e., skips if an even number is in the B-register.

SSA SKIP IF SIGN OF A IS ZERO

15	14	13	12	11	10	9	8	7	6	5	4	3	2	1	0
0	0	0	0	0	1						1				

Skips the next instruction if the sign bit (bit 15) of the A-register is a logic 0; i.e., skips if a positive number is in the A-register.

SSB SKIP IF SIGN OF B IS ZERO

15	14	13	12	11	10	9	8	7	6	5	4	3	2	1	0
0	0	0	0	1	1						1				

Skips the next instruction if the sign bit (bit 15) of the B-register is a logic 0; i.e., skips if a positive number is in the B-register.

SZA SKIP IF A IS ZERO

15	14	13	12	11	10	9	8	7	6	5	4	3	2	1	0
0	0	0	0	0	1									1	

Skips the next instruction if the A-register contents are zero (16 zeros).

SZB SKIP IF B IS ZERO

15	14	13	12	11	10	9	8	7	6	5	4	3	2	1	0
0	0	0	0	1	1									1	

Skips the next instruction if the B-register contents are zero (16 zeros).

3-25. INPUT/OUTPUT INSTRUCTIONS

The following input/output instructions provide the capability of setting, clearing or testing the flag and control bits associated with DMA, programmed I/O, interrupts, memory protect, time base generator, parity error, Global Register, and overflow. I/O instructions with select codes of seven or less have various functions. (Refer to Table 6-3 for further information regarding specific select-code functions.) I/O instructions permit data transfer between the A- and B-registers and either specific I/O devices or between registers associated with memory protect, parity error, or interrupts. The various registers and I/O devices are addressed by means of their register numbers and select codes.

Bit 11, where relevant, specifies the A- or B-register or distinguishes between set control and clear control; otherwise, bit 11 may be a logic 0 or a logic 1 without affecting the instruction (although the assembler will assign zeros in this case). In those instructions where bit position 9 includes the letters H/C, the programmer has the choice of holding (logic 0) or clearing (logic 1) the device flag after executing the instruction. (Exception: the H/C bit associated with instructions SOC and SOS holds or clears the overflow bit instead of the device flag.) Note that this H/C option is not supported on many of the I/O instructions with select code less than 10 octal.

Bits 8, 7, and 6, specify the appropriate I/O instruction. When the Global Register is enabled, bits 5 through 0 apply the instruction to a register on the I/O card whose select code is in the Global Register. (The Global Register is discussed further in paragraph 7-4.)

NOTE

Execution of I/O instructions is inhibited when the memory protect feature is enabled. Refer to paragraph 6-3.

The following instruction descriptions assume that the global register is disabled and, therefore, the instructions are addressed to a select code.

CLC CLEAR CONTROL

15	14	13	12	11	10	9	8	7	6	5	4	3	2	1	0
1	0	0	0	1	1	H/C	1	1	1						

Select Code or Register Number

Clears the control bit (Control 30) of the selected I/O channel or function. This turns off the specific device channel and prevents it from interrupting. A CLC 00 instruction clears the control bits from select code 20 upward, effectively turning off all I/O devices.

CLF CLEAR FLAG

15	14	13	12	11	10	9	8	7	6	5	4	3	2	1	0
1	0	0	0		1	1	0	0	1						


Select Code or Register Number

Clears the flag (Flag 30) of the selected I/O channel or function. A CLF 00 instruction disables the interrupt system for the time base generator and all interface cards; this does not affect the status of the individual channel flags.


CLO CLEAR OVERFLOW

15	14	13	12	11	10	9	8	7	6	5	4	3	2	1	0
1	0	0	0	0	1	1	0	0	1	0	0	0	0	0	1


Clears the overflow bit.




Halts the computer, holds or clears the flag of the selected I/O channel, and invokes the virtual control panel program. The HLT instruction will be contained in the T-register, which is displayed on the VCP when the VCP program starts executing. The P-register (also displayed) will normally contain the HLT location plus one. Note that if break is not enabled on any I/O card, the HLT instruction has no effect.




Loads the contents of the addressed I/O buffer or special function register into the A-register.




Loads the contents of the addressed I/O buffer or special function register into the B-register.




By executing a logical "inclusive or" function, merges the contents of the addressed I/O buffer or special function register into the A-register.




By executing a logical "inclusive or" function, merges the contents of the addressed I/O buffer or special function register into the B-register.




Outputs the contents of the A-register to the addressed I/O buffer or special function register. The contents of the A-register are not altered.




Outputs the contents of the B-register to the addressed I/O buffer or special function register. The contents of the B-register are not altered.




Skips the next programmed instruction if the flag (Flag 30) of the selected channel is clear (device busy).




Skips the next programmed instruction if the flag (Flag 30) of the selected channel is set (device ready).




Skips the next programmed instruction if the overflow bit is clear. Use the H/C bit (bit 9) to either hold or clear the overflow bit following the completion of this instruction (whether the skip is taken or not).




Skips the next programmed instruction if the overflow bit is set. Use the H/C bit (bit 9) to either hold or clear the overflow bit following the completion of this instruction (whether the skip is taken or not).



Sets the control bit (Control 30) of the selected I/O channel or function.



Sets the flag (Flag 30) of the selected I/O channel or function. An STF 00 instruction enables the interrupt system for the time base generator and all interface cards.




Sets the overflow bit.

3-26. EXTENDED ARITHMETIC MEMORY REFERENCE INSTRUCTIONS


The four extended arithmetic memory reference instructions provide for integer multiply and divide and for loading and storing double-length words to and from the A- and B-registers. The complete instruction requires two words: one for the instruction code and one for the address. When stored in memory, the instruction word is the first to be fetched; the address word is in the next sequential location.

Since 15 bits are available for the address, these instructions can directly address any location in memory.


As for all memory reference instructions, indirect addressing to any number of levels may also be used. A logic 0 in bit position 15 specifies direct addressing; a logic 1 specifies indirect addressing.




Divides a double-word integer in the combined B- and A-registers by a 16-bit integer in the addressed memory location. The result is a 16-bit integer quotient in the A-register and a 16-bit integer remainder in the B-register. Overflow can result from an attempt to divide by zero, or from an attempt to divide by a number too small for the dividend. In the former case (divide by zero), the division will not be attempted and the B- and A-register contents will be unchanged except that a negative quantity will be made positive. In the latter case (divisor too small), the execution will be attempted with unpredictable results left in the B- and A-registers. If there is no divide error, the overflow bit is cleared.



Loads the contents of addressed memory location m (and m + 1) into the A- and B-registers, respectively.



Stores the double-word quantity in the A- and B-registers into addressed memory locations m (and m + 1), respectively.




Multiplies a 16-bit integer in the A-register by a 16-bit integer in the addressed memory location. The resulting double-length integer product resides in the B- and A-registers, with the B-register containing the sign bit and the most-significant 15 bits of the quantity. The A-register may be used as an operand (i.e., memory address 0), resulting in an arithmetic square. The instruction clears the overflow bit.

3-27. EXTENDED ARITHMETIC REGISTER REFERENCE INSTRUCTIONS


The six extended arithmetic register reference instructions provide various types of shifting operations on the combined contents of the B- and A-registers. The B-register is considered to be to the left (most-significant word) and the A-register is considered to be to the right (least-significant word). An example of each type of shift operation is illustrated in Figure 3-4.

The complete instruction is given in one word and includes four bits (unshaded) to specify the number of shifts (1 to 16). By viewing these four bits as a binary-coded number, the number of shifts is easily expressed; i.e., binary-coded 1 = 1 shift, binary-coded 2 = 2 shifts . . . binary-coded 15 = 15 shifts. The maximum number of 16 shifts is coded with four zeros, which essentially exchanges the contents of the B- and A-registers.


The extend bit is not affected by any of the following instructions. Except for the arithmetic shifts, overflow also is not affected.




Arithmetically shifts the combined contents of the B- and A-registers left n places. The value of n may be any number from 1 through 16. Zeros are filled into vacated low-order positions of the A-register. The sign bit is not affected, and data bits are lost out of bit position 14 of the B-register. If any one of the lost bits is a significant data bit ("1" for positive numbers, "0" for negative numbers), the overflow bit will be set; otherwise, overflow will be cleared during execution. See ASL example in Figure 3-4. Note that two additional shifts in this example would cause an error by losing a significant '1'.




Arithmetically shifts the combined contents of the B- and A-registers right n places. The value of n may be any number from 1 through 16. The sign bit is unchanged and is extended into bit positions vacated by the right shift. Data bits shifted out of the least-significant end of the A-register are lost. Overflow cannot occur because the instruction clears the overflow bit.




Logically shifts the combined contents of the B- and A-registers left n places. The value of n may be any number from 1 through 16. Zeros are filled into vacated low-order bit positions of the A-register; data bits are lost out of the high-order bit positions of the B-register.



Logically shifts the combined contents of the B- and A-registers right n places. The value of n may be any number from 1 through 16. Zeros are filled into vacated high-order bit positions of the B-register; data bits are lost out of the low-order bit positions of the A-register.



Rotates the combined contents of the B- and A-registers left n places. The value of n may be any number from 1 through 16. No bits are lost or filled in. Data bits shifted out of the high-order end of the B-register are rotated around to enter the low-order end of the A-register.



Rotates the combined contents of the B- and A-registers

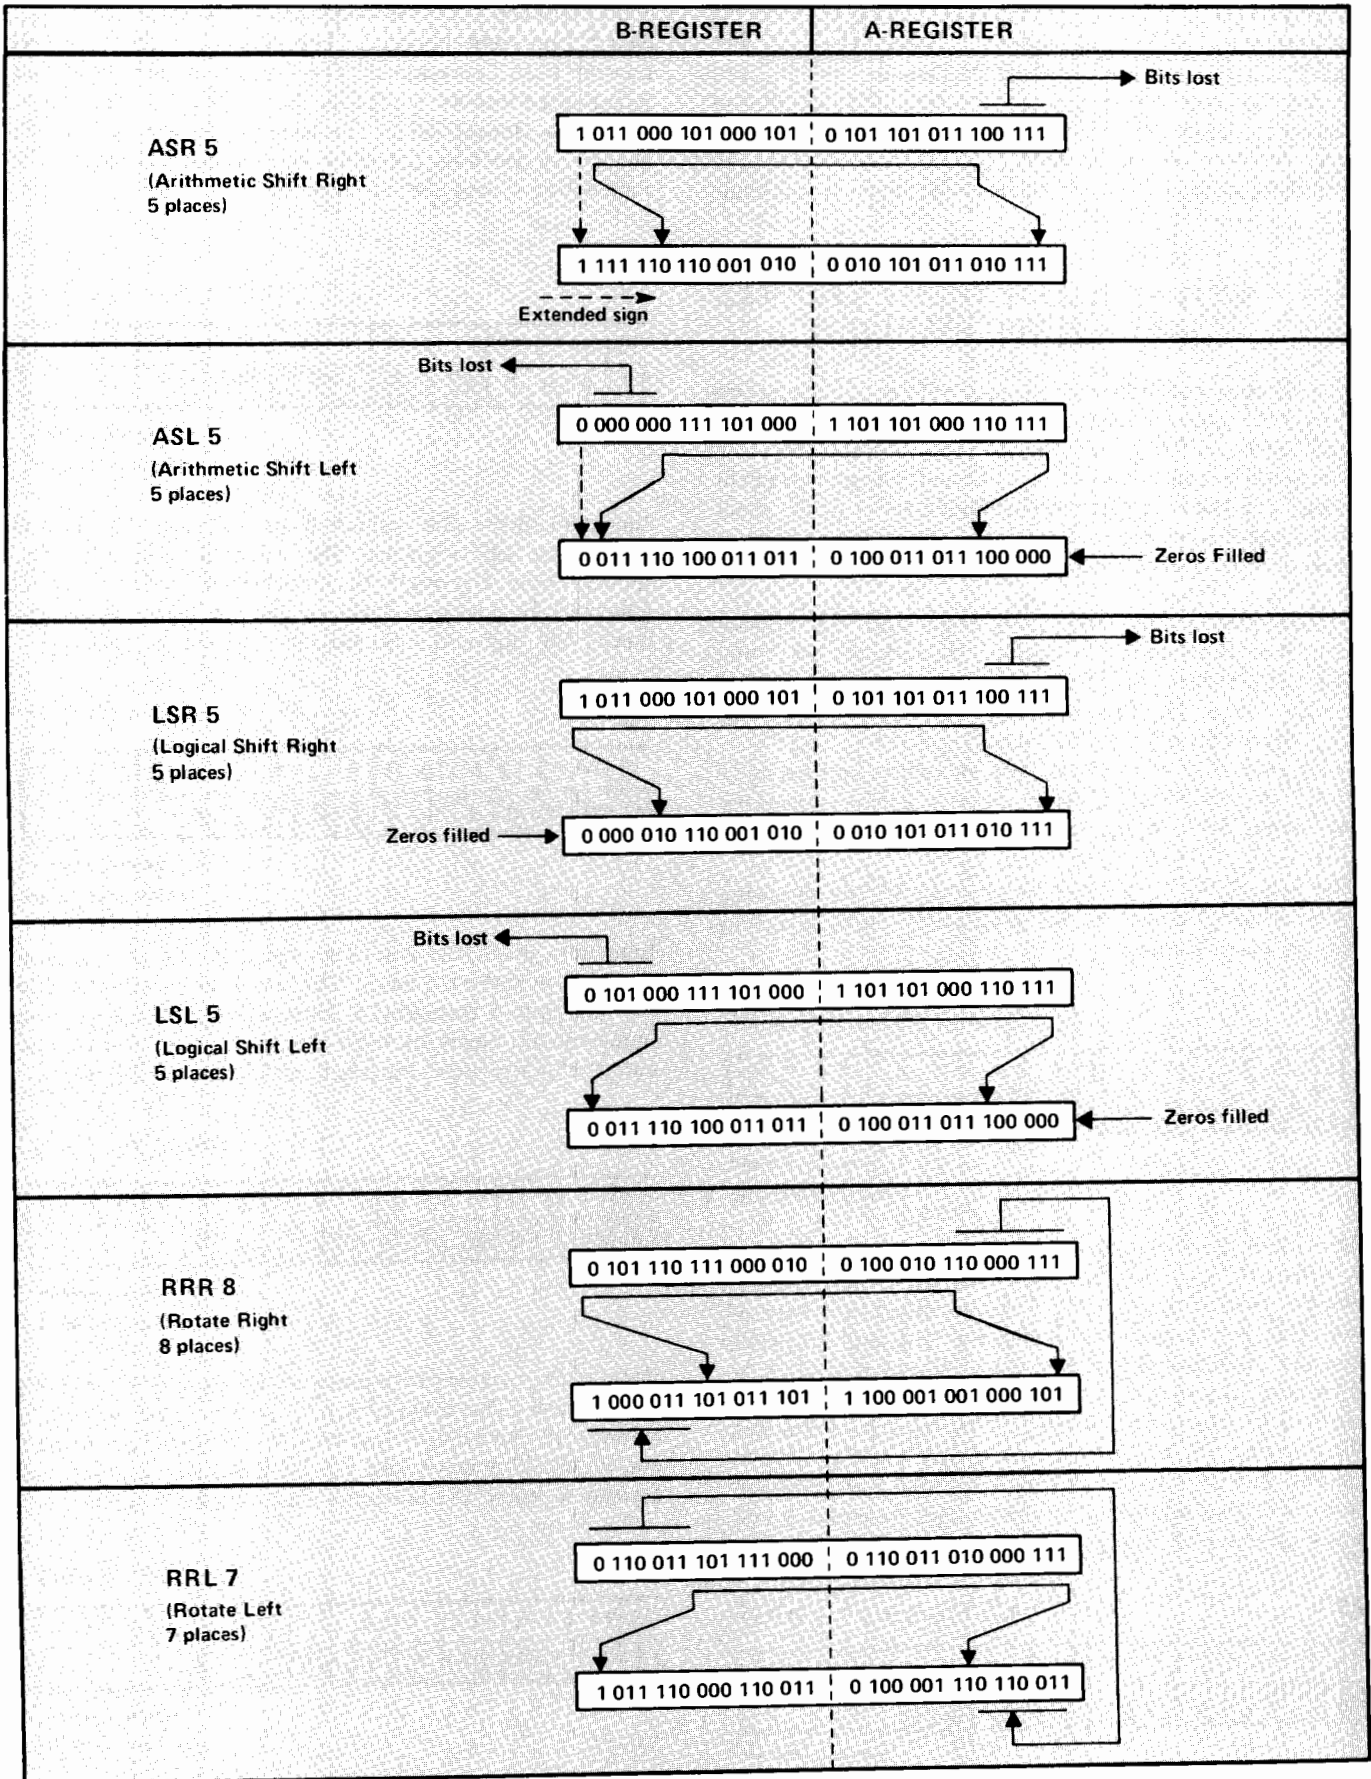


Figure 3-4. Examples of Double-Word Shifts and Rotates

right n places. The value of n may be any number from 1 through 16. No bits are lost or filled in. Data bits shifted out of the low-order end of the A-register are rotated around to enter the high-order end of the B-register.

3-28. EXTENDED INSTRUCTION GROUP

3-29. INDEX REGISTER INSTRUCTIONS. The index registers (X and Y) are two 16-bit registers accessible by the following instructions.

ADX		ADD MEMORY TO X													
15	14	13	12	11	10	9	8	7	6	5	4	3	2	1	0
1	0	0	0	1	0	1	1	1	1	1	0	0	1	1	0
\overline{D}/I															

Memory Address

Adds the contents of the addressed memory location to the contents of the X-register. The sum remains in the X-register and the contents of the memory cell are unaltered. The result of this addition may set the extend bit or the overflow bit.

ADY		ADD MEMORY TO Y													
15	14	13	12	11	10	9	8	7	6	5	4	3	2	1	0
1	0	0	0	1	0	1	1	1	1	1	0	1	1	1	0
\overline{D}/I															

Memory Address

Adds the contents of the addressed memory location to the contents of the Y-register. The sum remains in the Y-register and the contents of the memory cell are unaltered. The result of this addition may set the extend bit or the overflow bit.

CAX		COPY A TO X													
15	14	13	12	11	10	9	8	7	6	5	4	3	2	1	0
1	0	0	0	0	0	1	1	1	1	1	0	0	0	0	1

Copies the contents of the A-register into the X-register. The contents of the A-register are unaltered.

CAY		COPY A TO Y													
15	14	13	12	11	10	9	8	7	6	5	4	3	2	1	0
1	0	0	0	0	0	1	1	1	1	1	0	1	0	0	1

Copies the contents of the A-register into the Y-register. The contents of the A-register are unaltered.

CBX		COPY B TO X													
15	14	13	12	11	10	9	8	7	6	5	4	3	2	1	0
1	0	0	0	1	0	1	1	1	1	1	0	0	0	0	1

Copies the contents of the B-register into the X-register. The contents of the B-register are unaltered.

CBY		COPY B TO Y													
15	14	13	12	11	10	9	8	7	6	5	4	3	2	1	0
1	0	0	0	1	0	1	1	1	1	1	0	1	0	0	1

Copies the contents of the B-register into the Y-register. The contents of the B-register are unaltered.

CXA		COPY X TO A													
15	14	13	12	11	10	9	8	7	6	5	4	3	2	1	0
1	0	0	0	0	0	1	1	1	1	1	0	0	1	0	0

Copies the contents of the X-register into the A-register. The contents of the X-register are unaltered.

CXB		COPY X TO B													
15	14	13	12	11	10	9	8	7	6	5	4	3	2	1	0
1	0	0	0	1	0	1	1	1	1	1	0	0	1	0	0

Copies the contents of the X-register into the B-register. The contents of the X-register are unaltered.

CYA		COPY Y TO A													
15	14	13	12	11	10	9	8	7	6	5	4	3	2	1	0
1	0	0	0	0	0	1	1	1	1	1	0	1	1	0	0

Copies the contents of the Y-register into the A-register. The contents of the Y-register are unaltered.

CYB		COPY Y TO B													
15	14	13	12	11	10	9	8	7	6	5	4	3	2	1	0
1	0	0	0	1	0	1	1	1	1	1	0	1	1	0	0

Copies the contents of the Y-register into the B-register. The contents of the Y-register are unaltered.

DSX		DECREMENT X AND SKIP IF ZERO													
15	14	13	12	11	10	9	8	7	6	5	4	3	2	1	0
1	0	0	0	1	0	1	1	1	1	1	1	0	0	0	1

Subtracts one from the contents of the X-register. If the result of this operation is zero (X-register decremented from 000001 to 000000), the next instruction is skipped; i.e., the P-register count is advanced two counts instead of one count. If the result is not zero, the next sequential instruction is executed.

DSY **DECREMENT Y AND SKIP IF ZERO**

15	14	13	12	11	10	9	8	7	6	5	4	3	2	1	0
1	0	0	0	1	0	1	1	1	1	1	1	1	0	0	1

Subtracts one from the contents of the Y-register. If the result of this operation is zero (Y-register decremented from 000001 to 000000), the next instruction is skipped; i.e., the P-register count is advanced two counts instead of one count. If the result is not zero, the next sequential instruction is executed.

ISX **INCREMENT X AND SKIP IF ZERO**

15	14	13	12	11	10	9	8	7	6	5	4	3	2	1	0
1	0	0	0	1	0	1	1	1	1	1	1	0	0	0	0

Adds one to the contents of the X-register. If the result of this operation is zero (X-register rolls over to 000000 from 177777), the next instruction is skipped; i.e., the P-register count is advanced two counts instead of one count. If the result is not zero, the next sequential instruction is executed.

ISY **INCREMENT Y AND SKIP IF ZERO**

15	14	13	12	11	10	9	8	7	6	5	4	3	2	1	0
1	0	0	0	1	0	1	1	1	1	1	1	1	0	0	0

Adds one to the contents of the Y-register. If the result of this operation is zero (Y-register rolls over to 000000 from 177777), the next instruction is skipped; i.e., the P-register count is advanced two counts instead of one count. If the result is not zero, the next sequential instruction is executed.

LAX **LOAD A INDEXED BY X**

15	14	13	12	11	10	9	8	7	6	5	4	3	2	1	0
1	0	0	0	0	0	1	1	1	1	1	0	0	0	1	0
D/I															

Operand Address

Loads the A-register with the contents indicated by the effective address, which is computed by adding the contents of the X-register to the operand address. The effective address is loaded into the M-register; the X-register and memory contents are not altered. Indirect addressing is resolved before indexing; bit 15 of the effective address is ignored.

LAY **LOAD A INDEXED BY Y**

15	14	13	12	11	10	9	8	7	6	5	4	3	2	1	0
1	0	0	0	0	0	1	1	1	1	1	0	1	0	1	0
D/I															

Operand Address

Loads the A-register with the contents indicated by the effective address, which is computed by adding the contents of the Y-register to the operand address. The effective address is loaded into the M-register; the Y-register and memory contents are not altered. Indirect addressing is resolved before indexing; bit 15 of the effective address is ignored.

LBX **LOAD B INDEXED BY X**

15	14	13	12	11	10	9	8	7	6	5	4	3	2	1	0
1	0	0	0	1	0	1	1	1	1	1	0	0	0	1	0
D/I															

Operand Address


Loads the B-register with the contents indicated by the effective address, which is computed by adding the contents of the X-register to the operand address. The effective address is loaded into the M-register; the X-register and memory contents are not altered. Indirect addressing is resolved before indexing; bit 15 of the effective address is ignored.

LBY **LOAD B INDEXED BY Y**


15	14	13	12	11	10	9	8	7	6	5	4	3	2	1	0
1	0	0	0	1	0	1	1	1	1	1	0	1	0	1	0
D/I															

Operand Address


Loads the B-register with the contents indicated by the effective address, which is computed by adding the contents of the Y-register to the operand address. The effective address is loaded into the M-register; the X-register and memory contents are not altered. Indirect addressing is resolved before indexing; bit 15 of the effective address is ignored.




Loads the contents of the addressed memory location into the X-register. The A- and B-registers may be addressed as locations 00000 and 00001, respectively; however, if it is desired to load from the A- or B-register, copy instructions CAX or CBX should be used since they are more efficient.




Loads the contents of the addressed memory location into the Y-register. The A- and B-registers may be addressed as locations 00000 and 00001, respectively; however, if it is desired to load from the A- or B-register, copy instructions CAY or CBY should be used since they are more efficient.




Stores the contents of the A-register into the location indicated by the effective address, which is computed by adding the contents of the X-register to the operand address. The effective address is loaded into the M-register; the A- and X-register contents are not altered. Indirect addressing is resolved before indexing; bit 15 of the effective address is ignored.




Stores the contents of the A-register into the location indicated by the effective address, which is computed by adding the contents of the Y-register to the operand address. The effective address is loaded into the M-register; the A- and Y-register contents are not altered. Indirect addressing is resolved before indexing; bit 15 of the effective address is ignored.



Stores the contents of the B-register into the location indicated by the effective address, which is computed by adding the contents of the X-register to the operand address. The effective address is loaded into the M-register; the B- and X-register contents are not altered. Indirect addressing is resolved before indexing; bit 15 of the effective address is ignored.



Stores the contents of the B-register into the location indicated by the effective address, which is computed by adding the contents of the Y-register to the operand address. The effective address is loaded into the M-register; the B- and Y-register contents are not altered. Indirect addressing is resolved before indexing; bit 15 of the effective address is ignored.



Stores the contents of the X-register into the addressed memory location. The A- and B-registers may be addressed as locations 00000 and 00001, respectively. The X-register contents are not altered.

STY STORE Y TO MEMORY

15	14	13	12	11	10	9	8	7	6	5	4	3	2	1	0
1	0	0	0	1	0	1	1	1	1	1	0	1	0	1	1
D ₁															

Memory Address

Stores the contents of the Y-register into the addressed memory location. The A- and B-registers may be addressed as locations 00000 and 00001, respectively. The Y-register contents are not altered.

XAX EXCHANGE A AND X

15	14	13	12	11	10	9	8	7	6	5	4	3	2	1	0
1	0	0	0	0	0	1	1	1	1	1	0	0	1	1	1

Exchanges the contents of the A- and X-registers.

XAY EXCHANGE A AND Y

15	14	13	12	11	10	9	8	7	6	5	4	3	2	1	0
1	0	0	0	0	0	1	1	1	1	1	0	1	1	1	1

Exchanges the contents of the A- and Y-registers.

XBX EXCHANGE B AND X

15	14	13	12	11	10	9	8	7	6	5	4	3	2	1	0
1	0	0	0	1	0	1	1	1	1	1	0	0	1	1	1

Exchanges the contents of the B- and X-registers.

XBY EXCHANGE B AND Y

15	14	13	12	11	10	9	8	7	6	5	4	3	2	1	0
1	0	0	0	1	0	1	1	1	1	1	0	1	1	1	1

Exchanges the contents of the B- and Y-registers.

3-30. JUMP INSTRUCTIONS. The following four jump instructions allow a program to either jump to or exit from a subroutine.

JLY JUMP AND LOAD Y

15	14	13	12	11	10	9	8	7	6	5	4	3	2	1	0
1	0	0	0	1	0	1	1	1	1	1	1	0	0	1	0
D ₁															

Memory Address

This instruction is designed for entering a subroutine. The instruction, executed in location P, causes computer

control to jump unconditionally to the memory location specified in the memory address. Indirect addressing may be specified. The contents of the P-register plus two (return address) is loaded into the Y-register. A return to the main program sequence at P + 2 may be effected by a JPY instruction (described next).

JPY JUMP INDEXED BY Y

15	14	13	12	11	10	9	8	7	6	5	4	3	2	1	0
1	0	0	0	1	0	1	1	1	1	1	1	1	0	1	0
0															

Operand Address

Transfers control to the effective address, which is computed by adding the contents of the Y-register to the operand address. Indirect addressing is not allowed. The effective address is loaded into the P-register; the Y-register contents are not altered.

JLA JUMP AND LOAD A

15	14	13	12	11	10	9	8	7	6	5	4	3	2	1	0
1	0	0	0	0	0	0	1	1	0	0	0	0	0	0	0
D ₁															

Memory Address

This instruction, executed in location P, causes computer control to jump unconditionally to the memory location specified by the second word of the instruction. The contents of the program counter plus two are stored in the A-register. A return to the main program will be effected by a JMP indirect through location 00000 (the A-register).

JLB JUMP AND LOAD B

15	14	13	12	11	10	9	8	7	6	5	4	3	2	1	0
1	0	0	0	1	0	0	1	1	0	0	0	0	0	0	0
D ₁															

Memory Address

This instruction, executed in location P, causes computer control to jump unconditionally to the memory location specified by the second word of the instruction. The contents of the program counter plus two are stored in the B-register. A return to the main program will be effected by a JMP indirect through location 00001 (the B-register).

3-31. BYTE MANIPULATION INSTRUCTIONS.

A byte address is defined as two times the word address plus zero or one, depending on whether the byte is in the high-order position (bits 8 through 15) or low-order position (bits 0 through 7) of the word containing it. If the byte of interest is in bit positions 8 through 15 of memory location 100, for example, then the address of that byte is $2 * 100 + 0$, or 200; the address of the low-order byte in the same location is $201 (2 * 100 + 1)$. Because of the way byte addresses are defined, 16 bits are required to cover all possible byte addresses in a 32K-word memory configuration. Hence, for byte addressing, bit 15 does not indicate indirect addressing.

Byte addresses 000 through 003 reference bytes in the A- and B-registers. These addresses will not cause memory violations. The user should, however, be careful in referencing these byte addresses; for example, storing into byte address 002 or 003 would destroy the byte address originally contained in the B-register.

CBT				COMPARE BYTES											
15	14	13	12	11	10	9	8	7	6	5	4	3	2	1	0
1	0	0	0	1	0	1	1	1	1	1	1	0	1	1	0
D ₁															
0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0
Return if array 1 = array 2															
Return if array 1 < array 2															
Return if array 1 > array 2															

Compares the bytes in string 1 with those in string 2. This is a three-word instruction where

- Word 1 = Instruction code,
- Word 2 = Address of word containing the string count, and
- Word 3 = All-zeros word reserved for use by microcode.

The operand addresses are in the A- and B-registers. The A-register contains the first byte address of string 1 and the B-register contains the first byte address of string 2.

The number of bytes to be compared is given in the memory location addressed by Word 2 of the instruction; the number of bytes to be compared is restricted to a positive integer greater than zero. The strings are compared one byte at a time; the *i*th byte in string 1 is compared with the *i*th byte in string 2. The comparison is performed arithmetically; i.e., each byte is treated as a positive number. If all bytes in string 1 are identical with all bytes in string 2, the "equal" exit is taken. As soon as two bytes are compared and found to be different, the "less than" or "greater than" exit is taken, depending on whether the byte in string 1 is less than or greater than the byte in string 2. The three ways this instruction exits are as follows:

- a. No skip if string 1 is equal to string 2; the P-register advances one count from Word 3 of the instruction. The A-register contains its original value incremented by the count stored in the address specified in Word 2.
- b. Skips one word if string 1 is less than string 2; the P-register advances two counts from Word 3 of the instruction. The A-register contains the address of the byte in string 1 where the comparison stopped.
- c. Skips two words if string 1 is greater than string 2; the P-register advances three counts from Word 3 of the instruction. The A-register contains the address of the byte in string 1 where the comparison stopped.

For all three exits, the B-register will contain its original value incremented by the count stored in the address specified in Word 2. This instruction is interruptible. The interrupt routine is expected to save and restore the contents of the A- and B-registers. During the interrupt, the remaining count is stored in Word 3 of the instruction.

LBT				LOAD BYTE											
15	14	13	12	11	10	9	8	7	6	5	4	3	2	1	0
1	0	0	0	1	0	1	1	1	1	1	1	0	0	1	1
D ₁															
0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0

This one word instruction loads into the A-register the byte whose address is contained in the B-register. The byte is right-justified with leading zeros in the left byte. The B-register is incremented by one.

MBT				MOVE BYTES											
15	14	13	12	11	10	9	8	7	6	5	4	3	2	1	0
1	0	0	0	1	0	1	1	1	1	1	1	0	1	0	1
D ₁															
0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0

Moves bytes in a left-to-right manner; i.e., the byte having the lowest address from the source is moved first. This is a three word instruction where

- Word 1 = Instruction code,
- Word 2 = Address of word containing the byte count, and
- Word 3 = All-zeros word reserved for use by microcode.

The operand addresses are in the A- and B-registers. The A-register contains the first byte address source and the B-register contains the first byte address destination.

The number of bytes to be moved is given by a 16-bit positive integer greater than zero addressed by Word 2 of the instruction. The byte address in the A- and B-registers

are incremented as each byte is being moved. Thus, at the end of the operation, the A- and B-registers are incremented by the number of bytes moved. Wraparound of the byte address would result from a carry out of bit position 15; therefore, if the destination became 000, 001, 002, or 003, the next byte would be moved into the A- or B-register and destroy the proper byte addresses for the move operation. For each byte move, a memory protect check is performed.

This instruction is interruptible. The interrupt routine is expected to save and restore the contents of the A- and B-registers. During the interrupt, the remaining count is stored in Word 3 of the instruction.

SBT				STORE BYTE											
15	14	13	12	11	10	9	8	7	6	5	4	3	2	1	0
1	0	0	0	1	0	1	1	1	1	1	1	0	1	0	0

Stores the A-register low-order (right) byte in the byte address contained in the B-register. The B-register is incremented by one. A memory protect check is performed before the byte is stored. The left byte in the A-register does not have to be zeros. The other byte in the same word of the stored byte is not altered.

SFB				SCAN FOR BYTE											
15	14	13	12	11	10	9	8	7	6	5	4	3	2	1	0
1	0	0	0	1	0	1	1	1	1	1	1	0	1	1	1

This is a one word instruction with the operands in the A- and B-registers. The A-register contains a termination byte (high-order byte) and a test byte (low-order byte). The B-register contains the first byte address of the string to be scanned.

A string of bytes is scanned starting at the byte address given in the B-register. Scanning terminates when a byte in the string matches either the test byte or the termination byte in the A-register. The manner in which the instruction exits depends on which byte is matched first. If a byte in the string matches the test byte, the instruction will not skip upon exit; the B-register will contain the address of the byte matching the test byte. If a byte in the string matches the termination byte, the instruction will skip one word upon exit; the B-register will contain the address of the byte matching the termination byte *plus one*.

The scanning operation will not continue indefinitely even if neither the termination byte nor test byte exists in memory. These bytes are in the A-register with byte addresses 000 and 001, respectively. Thus, if no match is made by the time the B-register points to the last byte in memory, the B-register will roll over to zero and the next test will match the termination byte in the A-register with itself.

This instruction is interruptible. The interrupt routine is expected to save and restore the contents of the A- and B-registers.

3-32. BIT MANIPULATION INSTRUCTIONS.
The following three instructions allow any number of bits in a specified memory location to be cleared, set, or tested.

CBS				CLEAR BITS											
15	14	13	12	11	10	9	8	7	6	5	4	3	2	1	0
1	0	0	0	1	0	1	1	1	1	1	1	1	1	0	0
D ₁															
D ₁															

Memory Address

Clears bits in the addressed location. This is a three-word instruction where

Word 1 = Instruction code,

Word 2 = Address of a 16-bit mask, and

Word 3 = Address of word where bits are to be cleared.

The bits to be cleared correspond to logic 1's in the mask. The bits corresponding to logic 0's in the mask are not affected. A memory protect check is performed prior to modifying the word in memory.

SBS				SET BITS											
15	14	13	12	11	10	9	8	7	6	5	4	3	2	1	0
1	0	0	0	1	0	1	1	1	1	1	1	1	0	1	1
D ₁															
D ₁															

Memory Address

Sets bits in the addressed location. This is a three-word instruction where

Word 1 = Instruction code,

Word 2 = Address of a 16-bit mask, and

Word 3 = Address of word where bits are to be set.

The bits to be set correspond to logic 1's in the mask. The bits corresponding to logic 0's in the mask are not affected. A memory protect check is performed prior to modifying the word in memory.

TBS															TEST BITS			
15	14	13	12	11	10	9	8	7	6	5	4	3	2	1	0			
1	0	0	0	1	0	1	1	1	1	1	1	1	1	0	1			
D ₁																		
D _{1/2}																		

Memory Address

Tests (compares) bits in the addressed location. This is a three-word instruction where



- Word 1 = Instruction code,
- Word 2 = Address of a 16-bit mask, and
- Word 3 = Address of word in which bits are to be tested.

The bits in the addressed memory word corresponding to logic 1's in the mask are tested. If all the bits tested are 1's, the instruction will not skip; otherwise the instruction will skip one word (i.e., the P-register will advance two counts from Word 3 of the instruction).

3-33. WORD MANIPULATION INSTRUCTIONS.

The following instructions facilitate the comparing and moving of word arrays.

CMW															COMPARE WORDS			
15	14	13	12	11	10	9	8	7	6	5	4	3	2	1	0			
1	0	0	0	1	0	1	1	1	1	1	1	1	1	1	0			
D ₁																		
0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0			
Return if array 1 = array 2																		
Return if array 1 < array 2																		
Return if array 1 > array 2																		

Compares the words in array 1 with those in array 2. This is a three-word instruction where

- Word 1 = Instruction code,
- Word 2 = Address of word containing the word count, and
- Word 3 = All-zeros word reserved for use by microcode.

The operand addresses are in the A- and B-registers. The A-register contains the first word address of array 1 and the B-register contains the first word address of array 2. Bit 15 of the addresses in the A- and B-registers are ignored; i.e., no indirect addressing allowed.

The number of words to be compared is given in the memory location addressed by Word 2 of the instruction; the number of words to be compared is restricted to a positive integer greater than zero. The arrays are compared one word at a time; the *i*th word in array 1 is compared with the *i*th word in array 2. This comparison is performed arithmetically; i.e., each word is considered a two's complement number. If all words in array 1 are equal to all words in array 2, the "equal" exit is taken. As soon as two words are compared and found to be different, the "less than" or "greater than" exit is taken, depending on whether the word in array 1 is less than or greater than the word in array 2. The three ways this instruction exits are as follows:

- a. No skip if array 1 is equal to array 2; the P-register advances one count from Word 3 of the instruction. The A-register contains its original value incremented by the word count stored in the address specified in Word 2.
- b. Skips one word if array 1 is less than array 2; the P-register advances two counts from Word 3 of the instruction. The A-register contains the address of the word in array 1 where the comparison stopped.
- c. Skips two words if array 1 is greater than array 2; the P-register advances three counts from Word 3 of the instruction. The A-register contains the address of the word in array 1 where the comparison stopped.

For all three exits, the B-register will contain its original value incremented by the word count stored in the address specified in Word 2. This instruction is interruptible. The interrupt routine is expected to save and restore the contents of the A- and B-registers. During the interrupt, the remaining count is stored in Word 3 of the instruction.

MVW															MOVE WORDS			
15	14	13	12	11	10	9	8	7	6	5	4	3	2	1	0			
1	0	0	0	1	0	1	1	1	1	1	1	1	1	1	1			
D ₁																		
0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0			

Moves words in a left-to-right manner; i.e., the word having the lowest address in the source is moved first. This is a three-word instruction where

- Word 1 = Instruction code,
- Word 2 = Address of word containing the count, and
- Word 3 = All-zeros word reserved for use by microcode.

The operand addresses are in the A- and B-registers. The A-register contains the first word address source and the B-register contains the first word address destination. The

number of words to be moved is a 16-bit positive integer greater than zero addressed by Word 2 of the instruction. The word addresses in the A- and B-registers are incremented as each word is being moved. Thus, at the end of the operation, the A- and B-registers are incremented by the number of words moved.

Wraparound of the word address would result from a carry into bit position 15 (i.e., at 32767). If the destination address became 000 or 001, the next word would be moved into the A- or B-register and destroy the proper word addresses for the move operation. For each word move, a memory protect check is performed.

This instruction is interruptible. The interrupt routine is expected to save and restore the contents of the A- and B-registers. During the interrupt, the remaining count is stored in Word 3 of the instruction.

3-34. FLOATING POINT INSTRUCTIONS

The floating point instructions allow addition, subtraction, multiplication, and division of both single precision (32-bit) and double precision (64-bit) floating point quantities, and conversion of quantities from floating point format to integer format or vice versa. Data formats are shown in Figure 3-1. Except for zero, all floating point operands must be normalized (i.e., sign of mantissa differs from most significant bit of mantissa).

For multiple-word instructions, indirect addressing to any number of levels is permitted for the words indicated as memory address. A logic 0 in bit position 15 specifies direct addressing; a logic 1 specifies indirect addressing.

The execution times of the floating point instructions are specified in Tables 3-5 and 3-6. These instructions are non-interruptible; any attempted interrupt is held off for the full execution time of the currently active floating point instruction. However, data transfer via direct memory access is not held off.

Information required for direct user-microprogramming utilizing the Floating Point Processor card is provided in the *HP 92045A Microprogramming Package Reference Manual*, part no. 92045-90001.


NOTE

The optional Floating Point Processor (FPP) card is required for execution of the floating point instructions labeled "Optional" in the following paragraphs.

3-35. SINGLE PRECISION OPERATIONS. Overflow for single precision operations occurs if the result lies outside the range of representable single precision floating point numbers $[-2^{127}, (1-2^{-23}) 2^{127}]$. In such a case, the overflow flag is set and the result $(1-2^{-23}) 2^{127}$ is returned to the A- and B-registers. Underflow occurs if the result


lies inside the range $[-2^{-129}(1+2^{-22}), 2^{-129}]$. In such a case, the overflow flag is set and the result 0 is returned to the A- and B-registers.

FAD FLOATING POINT ADD




Adds the floating point quantity in the A- and B-registers to the floating point quantity in the specified memory locations. The floating point result is returned to the A- and B-registers.

FSB FLOATING POINT SUBTRACT




Subtracts the floating point quantity in the specified memory locations from the floating point quantity in the A- and B-registers. The floating point result is returned to the A- and B-registers.

FMP FLOATING POINT MULTIPLY



Multiplies the floating point quantity in the A- and B-registers by the floating point quantity in the specified memory locations. The floating point result is returned to the A- and B-registers.

FDV FLOATING POINT DIVIDE



Divides the floating point quantity in the A- and B-registers by the floating point quantity in the specified memory locations. The floating point result is returned to the A- and B-registers.

FIX FLOATING POINT TO SINGLE INTEGER

15	14	13	12	11	10	9	8	7	6	5	4	3	2	1	0
1	0	0	0	1	0	1	0	0	1	0	0	0	0	0	0

Converts the floating point quantity in the A- and B-registers to single integer format. The integer result is returned to the A-register. If the magnitude of the floating point number is <1, regardless of sign, the integer 0 is returned. If the magnitude of the exponent of the floating point number is ≥16, regardless of sign, the integer 32767 (077777 octal) is returned as the result and the overflow flag is set.

FLT SINGLE INTEGER TO FLOATING POINT

15	14	13	12	11	10	9	8	7	6	5	4	3	2	1	0
1	0	0	0	1	0	1	0	0	1	0	1	0	0	0	0

Converts the single integer quantity in the A-register to single precision floating point format. The floating point result is returned to the A- and B-registers.

.FIXD* FLOATING POINT TO DOUBLE INTEGER
(Optional)

15	14	13	12	11	10	9	8	7	6	5	4	3	2	1	0
1	0	0	0	1	0	1	0	0	1	0	0	0	1	0	0

Converts the floating point quantity in the A- and B-registers to double integer format. The integer result is returned to the A- and B-registers. (The A-register contains the most-significant word and the B-register contains the least-significant word.) If the magnitude of the floating point number is <1, regardless of sign, the integer 0 is returned. If the magnitude of the floating point number is ≥32, regardless of sign, the integer 2³¹-1 is returned as the result and the overflow flag is set.

.FLTD* DOUBLE INTEGER TO FLOATING POINT
(Optional)

15	14	13	12	11	10	9	8	7	6	5	4	3	2	1	0
1	0	0	0	1	0	1	0	0	1	0	1	0	1	0	0

Converts the double integer quantity in the A- and B-registers to single precision floating point format. The floating point result is returned to the A- and B-registers.

3-36. DOUBLE PRECISION OPERATIONS. Overflow for double precision operations occurs if the result lies outside the range of representable double precision floating point numbers $[-2^{127}, (1-2^{-55}) 2^{127}]$. In such a case, the overflow flag is set and $(1-2^{-55}) 2^{127}$ is returned as the result. Underflow occurs if the result lies inside the range $[-2^{-129} (1+2^{-54}), 2^{-129}]$. In such a case, the overflow flag is set and 0 is returned as the result.

.TADD* DOUBLE FLOATING POINT ADD
(Optional)

15	14	13	12	11	10	9	8	7	6	5	4	3	2	1	0
1	0	0	0	1	0	1	0	0	0	0	0	0	0	1	0
D ₁															
D ₁															
D ₁															

Memory Address

Adds two double precision floating point quantities (augend plus addend). This is a four-word instruction where

- Word 1 = Instruction code.
- Word 2 = Address of result.
- Word 3 = Address of augend.
- Word 4 = Address of addend.

.TSUB* DOUBLE FLOATING POINT SUBTRACT
(Optional)

15	14	13	12	11	10	9	8	7	6	5	4	3	2	1	0
1	0	0	0	1	0	1	0	0	0	0	1	0	0	1	0
D ₁															
D ₁															
D ₁															


Memory Address

Subtracts one double precision floating point quantity from another (minuend minus subtrahend). This is a four-word instruction where

- Word 1 = Instruction code.
- Word 2 = Address of result.
- Word 3 = Address of minuend.
- Word 4 = Address of subtrahend.

*For HP Assembly Language usage, refer to paragraph 3-46.


.TMPY* (Optional) DOUBLE FLOATING POINT MULTIPLY



Multiplies one double precision floating point quantity by another (multiplicand by multiplier). This is a four-word instruction where

- Word 1 = Instruction code.
- Word 2 = Address of result.
- Word 3 = Address of multiplicand.
- Word 4 = Address of multiplier.


.TDIV* (Optional) DOUBLE FLOATING POINT DIVIDE



Divides one double precision floating point quantity by another (dividend by divisor). This is a four-word instruction where

- Word 1 = Instruction code.
- Word 2 = Address of result.
- Word 3 = Address of dividend.
- Word 4 = Address of divisor.


.TFXS* (Optional) DOUBLE FLOATING POINT TO SINGLE INTEGER



Converts the double precision floating point quantity in the specified memory locations to single integer format. The integer result is returned to the A-register. If the magnitude of the floating point number is <1, regardless of sign, 0 is returned as the result. If the magnitude of the exponent of the floating point number is ≥16, regardless of sign, the integer 2¹⁵-1 is returned as the result and the overflow flag is set.


*For HP Assembly Language usage, refer to paragraph 3-46.

.TFTS* (Optional) SINGLE INTEGER TO DOUBLE FLOATING POINT




Converts the single integer quantity in the A-register to double precision floating point format. The floating point result is returned to the specified memory locations.

.TFXD* (Optional) DOUBLE FLOATING POINT TO DOUBLE INTEGER



Converts the double precision floating point quantity in the specified memory locations to double integer format. The integer result is returned to the A- and B-registers. (The A-register contains the most-significant word and the B-register contains the least-significant word.) If the magnitude of the floating point number is <1, regardless of sign, 0 is returned as the result. If the magnitude of the exponent of the floating point number is ≥32, regardless of sign, the integer 2³¹-1 is returned as the result and the overflow flag is set.

.TFTD* (Optional) DOUBLE INTEGER TO DOUBLE FLOATING POINT



Converts the double integer quantity in the A- and B-registers to double precision floating point format. The floating point result is returned to the specified memory locations.

3-37. LANGUAGE INSTRUCTION SET

The Language Instruction Set (LIS) instructions perform several frequently-used FORTRAN operations including parameter passing, array address calculations, and floating point conversion, packing, rounding and normalization operations.

For multiple-word instructions, indirect addressing to any number of levels is permitted for the words indicated as a memory address. A logic 0 in bit position 15 specifies direct addressing; a logic 1 specifies indirect addressing.

The following paragraphs provide machine language coding and definitions for the Language Instruction Set. Data formats are shown in Figure 3-1.

NOTE

The optional Floating Point Processor (FPP) card is required for execution of LIS instructions labeled "Optional" in the following paragraphs.

For a more detailed description of the instructions in the Language Instruction Set, refer to the Relocatable Library Reference Manual, HP part no. 24998-90001.

.BLE* SINGLE FLOATING POINT TO DOUBLE FLOATING POINT
(Optional)

15	14	13	12	11	10	9	8	7	6	5	4	3	2	1	0
1	0	0	0	1	0	1	0	1	0	0	0	0	1	1	1
D _I															
D _I															
D _I															

Memory Address

Converts the single precision floating point quantity in specified memory locations to a double precision floating point quantity. The result is returned to other specified memory locations. This is a four-word instruction where

- Word 1 = Instruction code.
- Word 2 = Return address.
- Word 3 = Address of result.
- Word 4 = Address of operand.

.NGL* DOUBLE FLOATING POINT TO SINGLE FLOATING POINT
(Optional)

15	14	13	12	11	10	9	8	7	6	5	4	3	2	1	0
1	0	0	0	1	0	1	0	1	0	0	0	1	1	0	0
D _I															
D _I															

Memory Address

Converts the double precision floating point quantity in the specified memory locations to a single precision floating point quantity. The result is placed in the A- and B-registers. Overflow is cleared unless, during execution, rounding results in overflow or underflow of the exponent, in which case overflow is set and the result is truncated to the greatest positive number. This is a three word instruction where

- Word 1 = Instruction code.
- Word 2 = Return address.
- Word 3 = Address of operand.

.XFER* TRANSFER EXTENDED FLOATING POINT

15	14	13	12	11	10	9	8	7	6	5	4	3	2	1	0
1	0	0	0	1	0	1	0	1	0	0	1	0	0	0	0

Transfers an extended precision floating point quantity (three consecutive words) from one memory location to another. The A-register must contain the source address and the B-register must contain the destination address. The source address +3 is returned to the A-register; the destination address +3 is returned to the B-register.

*For HP Assembly Language usage, refer to paragraph 3-46.

.DFER* TRANSFER EXTENDED FLOATING POINT

15	14	13	12	11	10	9	8	7	6	5	4	3	2	1	0
1	0	0	0	1	0	1	0	1	0	0	0	0	1	0	1
D ₁															
D ₁															

Memory Address

Transfers an extended precision floating point quantity (three consecutive words) from one memory location to another. The source address +3 is returned to the A-register; the destination address +3 is returned to the B-register. This is a three word instruction where

- Word 1 = Instruction code.
- Word 2 = Destination address.
- Word 3 = Source address.

.ZFER* TRANSFER EIGHT WORDS

15	14	13	12	11	10	9	8	7	6	5	4	3	2	1	0
1	0	0	0	1	0	1	0	1	0	0	1	1	1	1	1
D ₁															
D ₁															

Memory Address

Transfers eight consecutive words from one memory location to another. The source address +8 is returned to the A-register; the destination address +8 is returned to the B-register. This is a three word instruction where:

- Word 1 = Instruction code.
- Word 2 = Destination address.
- Word 3 = Source address.

.CFER* TRANSFER COMPLEX OR DOUBLE FLOATING POINT

15	14	13	12	11	10	9	8	7	6	5	4	3	2	1	0
1	0	0	0	1	0	1	0	1	0	0	1	1	0	0	1
D ₁															
D ₁															

Memory Address

Transfers a double precision floating point quantity (four consecutive words) from one memory location to another. The source address +4 is returned to the A-register; the destination address +4 is returned to the B-register. This is a three word instruction where

- Word 1 = Instruction code.
- Word 2 = Destination address.
- Word 3 = Source address.

.ENTN* TRANSFER PARAMETER ADDRESSES

15	14	13	12	11	10	9	8	7	6	5	4	3	2	1	0
1	0	0	0	1	0	1	0	1	0	0	1	1	1	0	0

Transfers the true addresses of parameters from a calling sequence into a subroutine; adjusts return address to the true return point. The return address stored in the SUB entry point references the word following the last parameter DEF in the calling routine. A true address is determined by eliminating all indirect references.

.ENTC* TRANSFER PARAMETER ADDRESSES

15	14	13	12	11	10	9	8	7	6	5	4	3	2	1	0
1	0	0	0	1	0	1	0	1	0	0	1	1	1	0	1

Transfers the true addresses of parameters from a calling sequence into a subroutine; adjusts return address to the true return point. The return address stored in the SUB entry point references the word following the last parameter DEF in the calling routine. There must be exactly two words between the subroutine entry point and the .ENTC instruction. A true address is determined by eliminating all indirect references. Used for privileged or re-entrant subroutines.

*For HP Assembly Language usage, refer to paragraph 3-46.

TRANSFER PARAMETER ADDRESSES

.ENTR*

15	14	13	12	11	10	9	8	7	6	5	4	3	2	1	0
1	0	0	0	1	0	1	0	1	0	0	1	0	0	1	1

Transfers the true addresses of parameters from a calling sequence into a subroutine; adjusts return address to the true return point. A true address is determined by eliminating all indirect references.

TRANSFER PARAMETER ADDRESSES

.ENTP*

15	14	13	12	11	10	9	8	7	6	5	4	3	2	1	0
1	0	0	0	1	0	1	0	1	0	0	1	0	1	0	0

Transfers the true addresses of parameters from a calling sequence into a subroutine; adjusts return address to the true return point. A true address is determined by eliminating all indirect references. Used for privileged or re-entrant subroutines.

.CPM* SINGLE INTEGER ARITHMETIC COMPARE

15	14	13	12	11	10	9	8	7	6	5	4	3	2	1	0
1	0	0	0	1	0	1	0	1	0	0	1	1	1	1	0
D ₁															
D ₁															
Return if operand 1 = operand 2															
Return if operand 1 < operand 2															
Return if operand 1 > operand 2															

Arithmetically compares operands addressed by second and third word. Does not skip if operands are equal; however, skips one instruction if the first operand is less than the second, or skips two instructions if the first operand is greater than the second.

.SETP*

SET A TABLE

15	14	13	12	11	10	9	8	7	6	5	4	3	2	1	0
1	0	0	0	1	0	1	0	1	0	0	0	0	1	1	1
0	Count ≥ 0														

Sets a table of increasing numbers in consecutive memory locations. The A-register must contain the initial number and the B-register must contain the initial memory address (direct only); the contents of the succeeding memory location must define the number of memory locations (count ≥ 0). Entries in the table are established by incrementing the initial address and number by one (1) for each successive entry until the last number, initial number +COUNT-1, is reached and the A-register equals the initial value+COUNT. Wraparound will produce undefined results. This instruction is interruptible. On return, the B-register equals the initial address +COUNT.

NOTE

If the initial address +COUNT -1 results in an address which is beyond the end of logical memory, addresses within the base page are destroyed.

..FCM* NEGATE SINGLE
FLOATING POINT

15	14	13	12	11	10	9	8	7	6	5	4	3	2	1	0
1	0	0	0	1	0	1	0	1	0	0	1	1	0	1	0

Negates a packed single precision floating point quantity located in the A- and B-registers. The result is returned to the A- and B-registers.

..TCM* NEGATE DOUBLE
FLOATING POINT
(Optional)

15	14	13	12	11	10	9	8	7	6	5	4	3	2	1	0
1	0	0	0	1	0	1	0	1	0	0	1	1	0	1	1
D ₁															

Memory Address

Negates a packed double precision floating point quantity located in the specified memory locations. The result is returned to the same specified memory locations.

3-38. DOUBLE INTEGER INSTRUCTIONS

The double integer instructions allow arithmetic and test operations on 32-bit integer quantities. The data format for double integer values is shown in figure 3-1. Double integer values contained in the (A,B) registers have the most significant bits in the A-register. Values stored in memory require two locations. The operand address in a double integer instruction points to the first memory location, which contains the most significant bits.

Instructions which do not return information in the extend or overflow bits will not alter the state of these flags. Operations which may return an overflow condition will clear overflow at entry.

The instructions .DMP, .DDI, and .DDIR utilize the Floating Point Processor card, if installed.

.DAD* DOUBLE INTEGER ADD

15	14	13	12	11	10	9	8	7	6	5	4	3	2	1	0
1	0	0	0	1	0	1	0	0	0	0	0	1	1	0	0
D ₁															

Memory Address

Performs the double integer operation:

$$(A,B) = (A,B) + \langle OPND \rangle$$

The contents of $\langle OPND \rangle$ are unaltered. In the event of overflow, the overflow bit is set and the returned result contains the lower 32-bits of the actual sum, in unsigned form. The extend bit will be set if an unsigned carry out of the A-register occurs.

.DSB* DOUBLE INTEGER SUBTRACT

15	14	13	12	11	10	9	8	7	6	5	4	3	2	1	0
1	0	0	0	1	0	1	0	0	0	0	1	1	1	0	0
D ₁															

Memory Address

Performs the double integer operation:

$$(A,B) = (A,B) - \langle OPND \rangle$$

The contents of $\langle OPND \rangle$ are unaltered. In the event of overflow, the overflow bit is set and the returned result contains the lower 32-bits of the actual difference, in unsigned form. The extend bit will be set if an unsigned borrow out of the A-register occurs.

.DSBR* DOUBLE INTEGER
SUBTRACT REVERSE

15	14	13	12	11	10	9	8	7	6	5	4	3	2	1	0
1	0	0	0	1	0	1	0	0	1	0	0	1	1	0	0
D ₁															

Memory Address

Performs the double integer operation:

$$(A,B) = \langle OPND \rangle - (A,B)$$

The contents of $\langle OPND \rangle$ are unaltered. In the event of overflow, the overflow bit is set and the returned result contains the lower 32-bits of the actual difference, in unsigned form. The extend bit will be set if an unsigned out of operand.

.DMP* DOUBLE INTEGER MULTIPLY

15	14	13	12	11	10	9	8	7	6	5	4	3	2	1	0
1	0	0	0	1	0	1	0	0	0	1	0	1	1	0	0
D ₁															

Memory Address

*For HP Assembly Language usage, refer to paragraph 3-46.

Performs the double integer operation:

$$(A,B) = (A,B) \times \langle OPND \rangle$$

The contents of $\langle OPND \rangle$ are unaltered. If overflow occurs, the result (077777, 177777) is returned and overflow is set.

.DDI* DOUBLE INTEGER DIVIDE

15	14	13	12	11	10	9	8	7	6	5	4	3	2	1	0
1	0	0	0	1	0	1	0	0	0	1	1	1	1	0	0
D/I															

Memory Address

Performs the double integer operation:

$$(A,B) = (A,B) \div \langle OPND \rangle$$

The contents of $\langle OPND \rangle$ are unaltered. If overflow or divide by zero occurs, the result (077777, 177777) is returned and overflow is set.

.DDIR* DOUBLE INTEGER DIVIDE REVERSE

15	14	13	12	11	10	9	8	7	6	5	4	3	2	1	0
1	0	0	0	1	0	1	0	0	1	0	1	1	1	0	0
D/I															

Memory Address

Performs the double integer operation:

$$(A,B) = \langle OPND \rangle \div (A,B)$$

The contents of $\langle OPND \rangle$ are unaltered. If overflow or divide by zero occurs, the result (077777, 177777) is returned and overflow is set.

.DNG* DOUBLE INTEGER NEGATE

15	14	13	12	11	10	9	8	7	6	5	4	3	2	1	0
1	0	0	0	1	0	1	0	1	0	0	0	0	0	1	1

Performs the double integer operation:

$$(A,B) = - (A,B)$$

An input value of (100000, 000000) is left unchanged and overflow is set. An input value of zero will cause the extend bit to be set.

.DCO* DOUBLE INTEGER COMPARE

15	14	13	12	11	10	9	8	7	6	5	4	3	2	1	0
1	0	0	0	1	0	1	0	1	0	0	0	0	1	0	0
D/I															

Memory Address

Compares the double integers (A,B) and $\langle OPND \rangle$

- If (A,B) = $\langle OPND \rangle$ Return to P+2
- If (A,B) < $\langle OPND \rangle$ Return to P+3
- If (A,B) > $\langle OPND \rangle$ Return to P+4

where P is the address of the .DCO instruction. The value of both double integers and the overflow bit are unaltered.

.DIN* DOUBLE INTEGER INCREMENT

15	14	13	12	11	10	9	8	7	6	5	4	3	2	1	0
1	0	0	0	1	0	1	0	1	0	0	0	1	0	0	0

Performs the double integer operation:

$$(A,B) = (A,B) + 1$$

An input value of (077777, 177777) will return a result of (100000, 000000) and set overflow. An input value of (177777, 177777) will return a result of zero and cause the extend bit to be set.

.DDE* DOUBLE INTEGER DECREMENT

15	14	13	12	11	10	9	8	7	6	5	4	3	2	1	0
1	0	0	0	1	0	1	0	1	0	0	0	1	0	0	1

Performs the double integer operation:

$$(A,B) = (A,B) - 1$$

*For HP Assembly Language usage, refer to paragraph 3-46.

An input value of (100000, 000000) will return the result (077777, 177777) and set overflow. An input value of zero will return the result (177777, 177777) and cause the extend bit to be set.

DOUBLE INTEGER INCREMENT AND SKIP IF ZERO

.DIS*

15	14	13	12	11	10	9	8	7	6	5	4	3	2	1	0
1	0	0	0	1	0	1	0	1	0	0	0	1	0	1	0
D/I															

Memory Address

Performs the double integer operation:

$$\langle \text{OPND} \rangle = \langle \text{OPND} \rangle + 1$$

If the new value of $\langle \text{OPND} \rangle$ equals zero, the next instruction will be skipped. The value in $\langle \text{OPND} \rangle$ is treated as an unsigned number, and a carry out of the $\langle \text{OPND} \rangle$ is ignored.

DOUBLE INTEGER DECREMENT AND SKIP IF ZERO

.DDS*

15	14	13	12	11	10	9	8	7	6	5	4	3	2	1	0
1	0	0	0	1	0	1	0	1	0	0	0	1	0	1	1
D/I															

Memory Address

Performs the double integer operation:

$$\langle \text{OPND} \rangle = \langle \text{OPND} \rangle - 1$$

If the new value of $\langle \text{OPND} \rangle$ equals zero, the next instruction will be skipped. The value in $\langle \text{OPND} \rangle$ is treated as an unsigned number, and a borrow out of the $\langle \text{OPND} \rangle$ is ignored.

3-39. VIRTUAL MEMORY INSTRUCTIONS

The Virtual Memory Instructions perform accesses to Virtual Memory and Extended Memory Area, which are extensions of logical memory. If an addressed data item is in physical memory, the instructions perform the required mapping, including modification of map registers and entry of the appropriate page numbers into the user's logical address space. If an addressed data item is not in physical memory, a macrocode routine first swaps the data item from disc to physical memory, then restarts the microcode routine that maps it.

For more information on VMA and EMA, refer to the *RTE-A.1 Programmer's Reference Manual*, HP part no. 92077-90007.

.IMAP* 16-BIT SUBSCRIPT MAPPING

15	14	13	12	11	10	9	8	7	6	5	4	3	2	1	0
1	0	0	0	1	0	1	0	1	0	1	0	1	0	0	0
D/I															
Word 2 = DEF dope vector															
Word 3 = Subscript N															
...															
Word N+2 = Subscript 1															

Performs a subscript calculation and maps the result into logical memory. Each of the subscripts and dimensions are 16-bit integers. However, the calculation uses 32-bit adds and multiplies. The subscript words cannot address the A- or B-register.

Word 2 points to a table that specifies the number of dimensions, dimension sizes, the number of words per element, and a two-word offset.

On a normal return, the A-register is undefined and the B-register contains the logical address.

.PMAP* MAP SPECIFIED PAGE

15	14	13	12	11	10	9	8	7	6	5	4	3	2	1	0
1	0	0	0	1	0	1	0	1	0	1	0	0	0	0	0
D/I															
Error return															
Normal return															

On entry, the A-register is loaded with the number of the user-map register to be altered and the B-register is loaded with the page ID, which are the parameters passed to the routine. If an attempt is made to map in the last+1 page, the last page of memory is mapped read and write protected and the E-register is set. When no error occurs, a normal return occurs to the second word after the instruction; mapping is complete, and the contents of the A- and B-registers are incremented. If a fault occurs and the sign bit is set in the A-register, an error return to the word following the instruction occurs. The O-register is undefined. The E-register is set if an attempt was made to map the last+1 page; otherwise it is cleared.

.IRES* 16-BIT SUBSCRIPT RESOLUTION

15	14	13	12	11	10	9	8	7	6	5	4	3	2	1	0
1	0	0	0	1	0	1	0	1	0	1	0	0	1	0	0
D/I															
Word 2 = DEF dope vector															
Word 3 = Subscript N															
...															
Word N+2 = Subscript 1															

*For HP Assembly Language usage, refer to paragraph 3-46.

Performs a subscript calculation. Each of the subscripts and dimensions are 16-bit integers. However, the calculation uses 32-bit adds and multiplies. The subscript words cannot address the A- or B-register.

Word 2 points to a table that specifies the number of dimensions, dimension sizes, the number of words per element, and a two-word offset.

On a normal return, the A- and B-registers contain the address of the array element in double-integer format (most significant word in the A-register).

.JMAP* 32-BIT SUBSCRIPT MAPPING

15	14	13	12	11	10	9	8	7	6	5	4	3	2	1	0
1	0	0	0	1	0	1	0	1	0	1	0	1	0	1	0
D ₁															
Word 2 = DEF dope vector															
Word 3 = Subscript N															
Word N+2 = Subscript 1															

Performs a subscript calculation and maps the result into logical memory. Each of the subscripts and dimensions are 32-bit integers, and the calculation uses 32-bit adds and multiplies. The subscript words cannot address the A- or B-register.

Word 2 points to a table that specifies the number of dimensions, dimension sizes, the number of words per element, and a two-word offset.

On a normal return, the A-register is undefined and the B-register contains the logical address.

.JRES* 32-BIT SUBSCRIPT RESOLUTION

15	14	13	12	11	10	9	8	7	6	5	4	3	2	1	0
1	0	0	0	1	0	1	0	1	0	1	0	0	1	0	1
D ₁															
Word 2 = DEF dope vector															
Word 3 = Subscript N															
Word N+2 = Subscript 1															

Performs a subscript calculation. Each of the subscripts and dimensions are 32-bit integers, and the calculation uses 32-bit adds and multiplies. The subscript words cannot address the A- or B-register.

Word 2 points to a table that specifies the number of dimensions, dimension sizes, the number of words per element, and a two-word offset.

On a normal return, the A- and B-registers contain the

address of the array element in double-integer format (most significant word in the A-register).

.LPXR* INDEXED MAPPING WITH DEF

15	14	13	12	11	10	9	8	7	6	5	4	3	2	1	0
1	0	0	0	1	0	1	0	1	0	1	0	1	1	0	0
D ₁															
D ₁															

Memory Address

On entry, the pointer specified by the second instruction word is resolved, and the double word it points to is loaded into the A- and B-registers. The offset specified in the third instruction word is resolved, and the double word it points to is added to the contents of the A- and B-registers. The result is treated as a 26-bit VMA pointer and is mapped. On exit, the B-register contains the logical address of the data item, and the A-register is undefined. The offset word cannot refer to the A- or B-register.

.LPX* INDEXED MAPPING WITH REGISTERS

15	14	13	12	11	10	9	8	7	6	5	4	3	2	1	0
1	0	0	0	1	0	1	0	1	0	1	0	1	1	0	1
D ₁															

Memory Address

On entry, the second instruction word either directly or indirectly points to a double integer in memory, which is to be added to the double integer in the A- and B-registers to form a double-word VMA pointer. The result is treated as a 26-bit VMA pointer and is mapped. On exit, the B-register contains the logical address of the data item, and the A-register is undefined.

.LBPR* MAPPING WITH DEF

15	14	13	12	11	10	9	8	7	6	5	4	3	2	1	0
1	0	0	0	1	0	1	0	1	0	1	0	1	1	1	0
D ₁															

Memory Address

On entry, the pointer specified by the second instruction word is resolved and the double word it points to is loaded into the A- and B-registers. This value is treated as a 26-bit VMA pointer and is mapped. On exit, the B-register contains the logical address of the data item, and the A-register is undefined.

*For HP Assembly Language usage, refer to paragraph 3-46.

.LBP* MAPPING WITH REGISTERS

15	14	13	12	11	10	9	8	7	6	5	4	3	2	1	0
1	0	0	0	1	0	1	0	1	0	1	0	1	1	1	1

On entry, the 26-bit VMA pointer is contained in the A-register (most significant word) and B-register. The data item is mapped. On exit, the B-register contains the logical address of the data item, and the A-register is undefined.

3-40. OPERATING SYSTEM INSTRUCTION SET

The operating system instructions provide instructions for ascertaining the CPU and firmware identification, and instructions for interrupt conditions.

.CPUID* PROCESSOR IDENTIFICATION

15	14	13	12	11	10	9	8	7	6	5	4	3	2	1	0
1	0	0	0	1	0	1	0	1	1	0	0	0	0	0	0

The A-register is loaded with a number that identifies the series of processor installed in the computer system, where:

Octal 2 = A600 Series.

Octal 3 = A700 Series.

.FWID* FIRMWARE IDENTIFICATION

15	14	13	12	11	10	9	8	7	6	5	4	3	2	1	0
1	0	0	0	1	0	1	0	1	1	0	0	0	0	0	1

On entry, the B-register holds a number indicating which bank of 1k microwords is identified. On exit, the A-register contains a number that identifies the specific ROM package (lower byte) and revision date code (upper byte). If no microcode exists in the selected block, the A-register is set to 177777 octal.

.WFI* WAIT FOR INTERRUPT

15	14	13	12	11	10	9	8	7	6	5	4	3	2	1	0
1	0	0	0	1	0	1	0	1	1	0	0	0	0	1	0

This instruction is equivalent to a **JMP *** except that the processor does not perform memory accesses, which would decrease the effective bandwidth of the memory backplane. This instruction is interruptible.

.SIP* SKIP IF INTERRUPT PENDING

15	14	13	12	11	10	9	8	7	6	5	4	3	2	1	0
1	0	0	0	1	0	1	0	1	1	0	0	0	0	1	1

The processor skips if an I/O interrupt is pending (INTRQ- is asserted on the A-Series backplane), which is independent of the Type 2 and Type 3 interrupt masks. (Refer to Table 6-1.)

3-41. EXECUTION TIMES

Table 3-5 lists the execution times required for the various base set instructions, and Table 3-6 lists the execution times for floating point instructions with the optional floating point processor card.

3-42. SCIENTIFIC INSTRUCTION SET

The Scientific Instruction Set (SIS) is included with the optional Floating Point Processor (FPP) card and performs nine trigonometric and logarithmic functions. The following paragraphs provide machine language coding and definitions for the SIS instructions. Error conditions and codes are given in Table 3-7. Note that except for zero, all floating point operands must be normalized (i.e., sign of mantissa differs from most significant bit of mantissa).

The following paragraphs provide machine language coding and definitions for the SIS instructions. Error conditions and codes are given in table 3-6. Note that except for zero, all floating point operands must be normalized (i.e., sign of mantissa differs from most significant bit of mantissa).

TAN* TANGENT

15	14	13	12	11	10	9	8	7	6	5	4	3	2	1	0
1	0	0	0	1	0	1	0	1	1	0	1	0	0	0	0

Calculates the tangent of the single precision floating point quantity (in radians) contained in the A- and B-registers. The result is returned to the A- and B-registers. A normal return will skip the next instruction. An error return will execute the next instruction, set the overflow bit, and return an ASCII error code in the A- and B-registers.

SQRT* SQUARE ROOT

15	14	13	12	11	10	9	8	7	6	5	4	3	2	1	0
1	0	0	0	1	0	1	0	1	1	0	1	0	0	0	1

*For HP Assembly Language usage, refer to paragraph 3-46.

Table 3-5. Typical Base Set Instruction Execution Times

INSTRUCTION	EXECUTION TIME (μ sec)	INSTRUCTION	EXECUTION TIME (μ sec)
Memory Reference Group			
(Direct)		OTA/B	3.00
LDA/B, STA/B	1.00	CLC	3.50
ADA/B, IOR, XOR, AND	1.00	STC	4.50
CPA/B	1.50	SC5: SFC, SFS without skip	3.00
ISZ without skip	1.50	with skip	3.25
with skip	1.75	STF	2.50
JSB	1.50	CLF	2.75
JMP	0.75	LIA/B	3.00
		CLC, STC	2.75
(One Indirect)		SC6: SFC, SFS without skip	3.50
LDA/B, STA/B	1.50	with skip	3.75
ADA/B, IOR, XOR, AND	1.50	STF	5.00
CPA/B	2.00	CLF	3.25
ISZ	2.00	CLC	4.50
JSB	1.50	STC	3.00
JMP	1.50	SC7: STC	3.00
(Each Additional Indirect)	0.50	LIA/B	3.25
Alter/Skip Group		0.75 to 2.25	
Shift/Rotate Group		0.75 to 3.00	
Extended Arithmetic Group			
DLD	2.50	SC20 and up:	
DST	2.25	CLC, CLF, STC, STF	3.50
MPY	6.00	SFC, SFS without skip	3.50
DIV	8.5 to 9.5	with skip	5.25
ASL	1.75 plus 0.50/shift	LIA/B	6.00
ASR, LSL, LSR, RRL, RRR	1.75 plus 0.25/shift	MIA/B	6.00
		OTA/B	5.25
Input/Output Group			
HLT	18.75	Extended Instruction Group	
By select code:		(Index Register Instructions)	
SC0: CLF	4.50	ADX, ADY	1.75
STF	5.75	CXA, CXB, CYA, CYB	0.75
SFC, SFS without skip	4.25	DSX, DSY	1.25
with skip	4.75	XAX, XBX, XAY, XBY	1.25
LIA/B	6.75	STX, STY	1.75
OTA/B	6.00	LDX, LDY	1.75
CLC	8.50	CAX, CBX, CAY, CBY	0.75
SC1: CLF, STF	2.00	LAX, LBX, LAY, LBY	2.25
SFC, SFS without skip	2.50	SAX, SBX, SAY, SBY	2.25
with skip	2.75	ISX, ISY	1.25
LIA/B	16.25	JLY, JPY	1.75
OTA/B	2.75	Per each indirect address level	0.50
SC2: STF	4.75	JLA, JLB	1.75
CLF	4.00	(Bit Manipulation Instructions)	
SFC, SFS with skip	4.50	CBS, SBS, TBS	3.50
without skip	4.75	(Word Manipulation Instructions)	
LIA/B	6.75	MVW	3.75 plus 1.00/word
OTA/B	6.00	CMW	3.75 plus 1.25/word
STC	3.50	(Byte Manipulation Instructions)	
SC3: LIA/B	6.75	LBT	1.75
OTA/B	6.00	SBT	2.50
SC4: SFC, SFS without skip	2.75	MBT, CBT	4.25 plus 2.33/byte
with skip	3.00	SFB (Exit on location of byte)	1.50 plus 1.50/byte
LIA/B	2.75	Floating Point Group	
		Faster execution if FPP card is installed.	
		(Single Precision)	
		FLT	1.75 to 6.50
		FIX	1.50 to 6.50
		FAD, FSB	7.75 to 26.00
		FMP	13.75 to 25.25

Table 3-5. Typical Base Set Instruction Execution Times (Continued)

INSTRUCTION	EXECUTION TIME (μsec)
FDV	18.25 to 29.75
Language Instruction Set	
.ENTR	4.75
.ENTP	5.25
.ENTN	3.5
.ENTC	4.25
per parameter (no indirect)	1.00
per indirect address level	1.00
.SETP (interruptible)	3.25
per table entry	0.50
.XFER	4.5
.DFER	5.75
.CFER	7.00
.ZFER	11.00
.CPM	3.50 to 4.25
..FCM	1.50 to 6.50
.NGL	6.75*
.BLE	6.25*
..TCM	7.00*
*With FPP card installed.	
Double Integer Instructions	
.DAD	2.50 to 3.50
.DSB, DSBR	3.25 to 3.50
.DNG, DIN, DDE	1.75
.DCO	3.25 to 3.50
.DIS	2.75 to 3.25
.DDS	3.00
.DMP standard	16.75 to 27.00
with FPP* card	5.13
.DDI standard	9.25 to 73.10
with FPP* card	7.29
.DDIR standard	9.5 to 73.50
with FPP* card	7.29
*Floating point processor.	
Virtual Memory Instructions	
.LBP	6.25
.LBPR	7.00
.LPX	7.75
.LPXR	9.00
.IMAP (Basic)	9.75
Per parameter (standard)	9.75 to 14.75
Per parameter (with FPP*)	6.5
.JMAP (Basic)	9.75
Per parameter (standard)	11.00 to 26.00
Per parameter (with FPP*)	7.5
.IRES (Basic)	5.0
Per parameter (standard)	9.75 to 14.75
Per parameter (with FPP*)	6.5
.JRES (Basic)	5.5
Per parameter (standard)	11.00 to 26.00
Per parameter (with FPP*)	7.50
.PMAP	5.00 to 6.00
*Floating point processor.	

INSTRUCTION	EXECUTION TIME (μsec)
Operating System Instructions	
.CPUID	1.25
.FWID	2.50
.SIP	1.50
with skip	1.36
.WFI	Until interrupted
Dynamic Mapping System Instruction Group	
Refer to Section IV for detailed descriptions and execution times.	

Table 3-6. Typical Execution Times for Floating Point Instructions with Optional Hardware FPP Card

INSTRUCTION	EXECUTION TIME (μsec)
FAD, FSB, FMP	4.25
FDV	6.25
FIX, FLT	1.75 to 3.25
FIXD	3.75
FLTD	3.00
.TADD, .TSUB, .TMPY	9.00
.TDIV	14.50
.TFXS	4.50
.TFTS	5.00
.TFXD	4.75
.TFTD	5.25

Table 3-7. SIS Instruction Error Codes

INSTRUCTION	ERROR CODE
TAN	09OR if $ x > 32768^*\pi/4$
SQRT	03UN if $x < 0$
ALOG	02UN if $x \leq 0$
ATAN	None
COS	05OR if $ x > 32768^*\pi/4$
SIN	05OR if $ x > 32768^*\pi/4$
EXP	07OF if $x > 88.029678$
ALOGT	02UN if $x \leq 0$
TANH	None
Where: OF = Integer or floating point overflow. OR = Out of range. UN = Floating point underflow.	

Note: Memory refresh during a processor memory access can make an instruction approximately 3% slower. Heavy DMA activity can also degrade instruction times due to contention for memory.

Calculates the square root of the single precision floating point quantity contained in the A- and B-registers. The result is returned to the A- and B-registers. A normal return will skip the next instruction. An error return will execute the next instruction, set the overflow bit, and return an ASCII error code in the A- and B-registers.

ALOG* NATURAL LOGARITHM

15	14	13	12	11	10	9	8	7	6	5	4	3	2	1	0
1	0	0	0	1	0	1	0	1	1	0	1	0	0	1	0

Calculates the natural logarithm of the single precision floating point quantity contained in the A- and B-registers. The result is returned to the A- and B-registers. A normal return will skip the next instruction. An error return will execute the next instruction, set the overflow bit, and return an ASCII error code in the A- and B-registers.

ATAN* ARCTANGENT

15	14	13	12	11	10	9	8	7	6	5	4	3	2	1	0
1	0	0	0	1	0	1	0	1	1	0	1	0	0	1	1

Calculates the arctangent of the single precision floating point quantity contained in the A- and B-registers. The result (in radians) is returned to the A- and B-registers.

COS* COSINE

15	14	13	12	11	10	9	8	7	6	5	4	3	2	1	0
1	0	0	0	1	0	1	0	1	1	0	1	0	1	0	0

Calculates the cosine of the single precision floating point quantity (in radians) contained in the A- and B-registers. The result is returned to the A- and B-registers. A normal return will skip the next instruction. An error return will execute the next instruction, set the overflow bit, and return an ASCII error code in the A- and B-registers.

SIN* SINE

15	14	13	12	11	10	9	8	7	6	5	4	3	2	1	0
1	0	0	0	1	0	1	0	1	1	0	1	0	1	0	1

Calculates the sine of the single precision floating point quantity (in radians) contained in the A- and B-registers. The result is returned to the A- and B-registers. A normal return will skip the next instruction. An error return will execute the next instruction, set the overflow bit, and return an ASCII error code in the A- and B-registers.

EXP* E TO THE POWER X

15	14	13	12	11	10	9	8	7	6	5	4	3	2	1	0
1	0	0	0	1	0	1	0	1	1	0	1	0	1	1	0

Calculates e to the power x of the single precision floating point quantity contained in the A- and B-registers. The result is returned to the A- and B-registers. A normal return will skip the next instruction. An error condition will execute the next instruction, set the overflow bit, and return an ASCII error code in the A- and B-registers.

ALOGT* COMMON LOGARITHM

15	14	13	12	11	10	9	8	7	6	5	4	3	2	1	0
1	0	0	0	1	0	1	0	1	1	0	1	0	1	1	1

Calculates the common logarithm of the single precision floating point quantity contained in the A- and B-registers. The result is returned to the A- and B-registers. A normal return will skip the next instruction. An error condition will execute the next instruction, set the overflow bit, and return an ASCII error code in the A- and B-registers.

TANH* HYPERBOLIC TANGENT

15	14	13	12	11	10	9	8	7	6	5	4	3	2	1	0
1	0	0	0	1	0	1	0	1	1	0	1	1	0	0	0

Calculates the hyperbolic tangent of the single precision floating point quantity contained in the A- and B-registers. The result is returned to the A- and B-registers.

DPOLY * POLYNOMIAL EVALUATION

15	14	13	12	11	10	9	8	7	6	5	4	3	2	1	0
1	0	0	0	1	0	1	0	1	1	0	1	1	0	0	1
F	S	X	X	X	X	X	X	X	X	X	X	X	X	X	T
D/I															
D/I															
D/I															
D/I															
D/I															

Memory Address

Evaluates a polynomial or quotient of polynomials using 64-bit floating-point. This is a seven-word instruction where:

Word 1 = Instruction code.

*For HP Assembly Language usage, refer to paragraph 3-46.

- Word 2 = Sub-opcode.
- Word 3 = Address of result Y (64-bit floating).
- Word 4 = Address of argument X (64-bit floating).
- Word 5 = Address of coefficient P_M (64-bit floating).
- Word 6 = Address of numerator order M (integer).
- Word 7 = Address of denominator order N (integer).

$$\text{Define } P(Z) = P_M Z^M + P_{M-1} Z^{M-1} + \dots + P_1 Z + P_0$$

$$Q(Z) = Z^N + Q_{N-1} + \dots + Q_1 Z + Q_0$$

The computation performed depends on the values of bits F, S, T of the sub-opcode:

- F=0: $Y = P(X)/Q(X)$
- F=1, S=0, T=0: $Y = P(X^2)/Q(X^2)$
- F=1, S=0, T=1: $Y = X * P(X^2)/Q(X^2)$
- F=1, S=1, T=0: $Y = P(X^2)/(P(X^2) - Q(X^2)) \quad (N > 0)$
- F=1, S=1, T=1: $Y = X * P(X^2)/(P(X^2) - Q(X^2)) \quad (N > 0)$

Horner's Rule is used to evaluate the polynomial(s). The coefficients must be stored sequentially in memory, starting with P_M , in the order:


$$P_M, P_{M-1}, \dots, P_1, P_0, Q_{N-1}, \dots, Q_1, Q_0$$

where $Q_N = 1.0$ is implied but not stored. If $N=0$, no coefficients are provided for Q and only P is evaluated. The case $N=0$ and $S=1$ is not allowed.

Any underflow or overflow which occurs invalidates the final result. M must be at least one. The A,B,X,Y and E registers are undefined after this instruction. The O register is set if any underflow or overflow occurs, else cleared.

This instruction is interruptible. Since it restarts after an interrupt, it is not recommended for very large values of (M+N).

Timing: Approximately $13.88 + 8.33M$ if $N=0$
 $18.5 + 8.33(M+N)$ if $N > 0$




Performs the computation $X = (1-X)/(1+X)$.

- Word 1 = Instruction code.
- Word 2 = Direct address of X (64-bit floating).

The A,B,X,Y,E and O registers are undefined after this instruction.

.FPWR* EXPONENTIATION



Raises a 32-bit floating-point number to an integer power. This is a two-word instruction, where:

- Word 1 = Instruction code.
- Word 2 = Address of base X (32-bit floating).

The power I is supplied in the A-register. It is unsigned and must be in the range [2,32768]. The left-to-right binary method is used to compute X^I , e.g. if $I=83_{10} = 123_8 = 1010011_2$ then


$X^2 = X * X$	10
$X^4 = X^2 * X^2$	100
$X^5 = X^4 * X$	101
$X^{10} = X^5 * X^5$	1010
$X^{20} = X^{10} * X^{10}$	10100
$X^{40} = X^{20} * X^{20}$	101000
$X^{41} = X^{40} * X$	101001
$X^{82} = X^{41} * X^{41}$	1010010
$X^{83} = X^{82} * X$	1010011

The X,Y and E registers are undefined. The O register is set if underflow or overflow occur else cleared. The A and B registers contain the result.

Timing: Approximately $18.5 + 2.5M + 3.24N$ microseconds

where $M = (\# \text{ bits in } I)$
 $N = (\# \text{ bits set in } I)$

.TPWR* EXPONENTIATION



Raises a 64-bit floating-point number to an integer power. This is a three-word instruction, where:

- Word 1 = Instruction code.
- Word 2 = Address of result (64-bit floating).
- Word 3 = Address of base X (64-bit floating).

*For HP Assembly Language usage, refer to paragraph 3-46.

The power I is supplied in the A-register. It is unsigned and must be in the range [2,32768]. The left-to-right binary method is used.

The A,B,X,Y and E registers are undefined. The O register is set if underflow or overflow occur else cleared.

Timing: Approximately $20.37 + 4.16M + 5.55N$ microseconds

where M = (# bits in I)
N = (# bits set in I)

3-43. EXECUTION TIMES AND INTERRUPTS

Table 3-8 lists the typical execution times required for the SIS instructions. Also listed is the maximum period of non-interruptible instruction execution. If an instruction is interrupted, its execution restarts from the beginning.

3-44. VECTOR INSTRUCTION SET

The Vector Instruction Set (VIS) is included with the optional Floating Point Processor (FPP) card and performs arithmetic operations on arrays of floating point numbers. The VIS utilizes the FPP as a computing resource and provides nineteen operations in both single and double precision formats, for a total of 38 instructions. For more information on the VIS instructions, refer to the VIS User's Manual, HP part no. 12824-90001.

Vector instructions require six to ten memory locations to specify parameters of the following type:

Opcode — Specifies the microcode entry point. Bit 4, or P-bit, indicates the precision of the operation. (P=0 for single precision, P=1 for double precision.)

Return address — Specifies the *direct* address of the next instruction.

The remaining parameters are addresses which may be direct or indirect, as indicated by bit 15. These include:

Vector — Specifies the address of the first vector element to be processed. Vector elements require two (single precision) or four (double precision) mem-

ory locations. All vectors in a given instruction must be of the same precision. Note that for instructions that contain two vector operands, these operands may both specify the same vector. Similarly, the result vector may replace one of the operands.

Scalar — Specifies the address of a single floating point quantity. Scalars are used for both operands or results. The precision of the scalar must match that of the associated vectors.

Integer — Specifies the address of an integer quantity in which a result is returned.

Increment — Specifies the address of an integer quantity associated with each vector. The increment indicates the spacing between vector elements to be processed. (An increment of 1 indicates that each element will be processed, an increment of 2 indicates every other element, etc.) An increment of zero will cause the first element of the vector to be used in all operations. Negative increments will step through the vector in reverse order, i.e., decreasing memory locations. Vector elements skipped over by the increment will not be modified.

#Elements — Specifies the address of an integer quantity indicating the number of vector elements to be processed. A value less than or equal to zero will result in a NOP operation.

VADD/DVADD * VECTOR ADD

	15	14	13	12	11	10	9	8	7	6	5	4	3	2	1	0
	1	0	0	0	1	0	1	0	0	0	0	P	0	0	0	1
	0	X	X	X	X	X	X	X	X	X	X	X	X	X	X	X
D/I																
D/I																
D/I																
D/I																
D/I																
D/I																
D/I																

Memory Address

*For HP Assembly Language usage, refer to paragraph 3-46.

Performs the vector operation:

$$V3 = V1 + V2$$

This is a nine-word instruction, where

- Word 1 = Instruction code.
- Word 2 = Return address.
- Word 3 = Address of vector V1.
- Word 4 = Address of increment INCR1.
- Word 5 = Address of vector V2.
- Word 6 = Address of increment INCR2.
- Word 7 = Address of vector V3.
- Word 8 = Address of increment INCR3.
- Word 9 = Address of # elements N.

VSUB/DVSUB * VECTOR SUBTRACT

15	14	13	12	11	10	9	8	7	6	5	4	3	2	1	0
1	0	0	0	1	0	1	0	0	0	0	P	0	0	1	1
0	X	X	X	X	X	X	X	X	X	X	X	X	X	X	X
D/I															
D/I															
D/I															
D/I															
D/I															
D/I															
D/I															

Memory Address

Performs the vector operation:

$$V3 = V1 - V2$$

This is a nine-word instruction, where

- Word 1 = Instruction code.
- Word 2 = Return address.
- Word 3 = Address of vector V1.
- Word 4 = Address of increment INCR1.
- Word 5 = Address of vector V2.
- Word 6 = Address of increment INCR2.
- Word 7 = Address of vector V3.
- Word 8 = Address of increment INCR3.
- Word 9 = Address of # elements N.

VMPY/DVMPY * VECTOR MULTIPLY

15	14	13	12	11	10	9	8	7	6	5	4	3	2	1	0
1	0	0	0	1	0	1	0	0	0	0	P	0	1	0	0
0	X	X	X	X	X	X	X	X	X	X	X	X	X	X	X
D/I															
D/I															
D/I															
D/I															
D/I															
D/I															
D/I															

Memory Address

Performs the vector operation:

$$V3 = V1 * V2$$

This is a nine-word instruction, where

- Word 1 = Instruction code.
- Word 2 = Return address.
- Word 3 = Address of vector V1.
- Word 4 = Address of increment INCR1.
- Word 5 = Address of vector V2.
- Word 6 = Address of increment INCR2.
- Word 7 = Address of vector V3.
- Word 8 = Address of increment INCR3.
- Word 9 = Address of # elements N.

VDIV/DVDIV * VECTOR DIVIDE

15	14	13	12	11	10	9	8	7	6	5	4	3	2	1	0
1	0	0	0	1	0	1	0	0	0	0	P	0	1	0	1
0	X	X	X	X	X	X	X	X	X	X	X	X	X	X	X
D/I															
D/I															
D/I															
D/I															
D/I															
D/I															
D/I															

Memory Address

*For HP Assembly Language usage, refer to paragraph 3-46.

Performs the vector operation:

$$V3 = V1 / V2$$

This is a nine-word instruction, where

- Word 1 = Instruction code.
- Word 2 = Return address.
- Word 3 = Address of vector V1.
- Word 4 = Address of increment INCR1.
- Word 5 = Address of vector V2.
- Word 6 = Address of increment INCR2.
- Word 7 = Address of vector V3.
- Word 8 = Address of increment INCR3.
- Word 9 = Address of # elements N.

VSAD/DVSAD * SCALAR-VECTOR ADD

15	14	13	12	11	10	9	8	7	6	5	4	3	2	1	0
1	0	0	0	1	0	1	0	0	0	0	P	0	1	1	0
0	X	X	X	X	X	X	X	X	X	X	X	X	X	X	X
D/I															
D/I															
D/I															
D/I															
D/I															

Memory Address

Performs the vector operation:

$$V2 = S + V1$$

This is an eight-word instruction, where

- Word 1 = Instruction code.
- Word 2 = Return address.
- Word 3 = Address of scalar S.
- Word 4 = Address of vector V1.
- Word 5 = Address of increment INCR1.
- Word 6 = Address of vector V2.
- Word 7 = Address of increment INCR2.
- Word 8 = Address of # elements N.

VSSB/DVSSB * SCALAR-VECTOR SUBTRACT

15	14	13	12	11	10	9	8	7	6	5	4	3	2	1	0
1	0	0	0	1	0	1	0	0	0	0	P	0	1	1	1
0	X	X	X	X	X	X	X	X	X	X	X	X	X	X	X
D/I															
D/I															
D/I															
D/I															
D/I															

Memory Address

Performs the vector operation:

$$V2 = S - V1$$

This is an eight-word instruction, where

- Word 1 = Instruction code.
- Word 2 = Return address.
- Word 3 = Address of scalar S.
- Word 4 = Address of vector V1.
- Word 5 = Address of increment INCR1.
- Word 6 = Address of vector V2.
- Word 7 = Address of increment INCR2.
- Word 8 = Address of # elements N.

VSMY/DVSMY * SCALAR-VECTOR MULTIPLY

15	14	13	12	11	10	9	8	7	6	5	4	3	2	1	0
1	0	0	0	1	0	1	0	0	0	0	P	1	0	0	0
0	X	X	X	X	X	X	X	X	X	X	X	X	X	X	X
D/I															
D/I															
D/I															
D/I															
D/I															

Memory Address

*For HP Assembly Language usage, refer to paragraph 3-46.

Performs the vector operation:

$$V2 = S * V1$$

This is an eight-word instruction, where

- Word 1 = Instruction code.
- Word 2 = Return address.
- Word 3 = Address of scalar S.
- Word 4 = Address of vector V1.
- Word 5 = Address of increment INCR1.
- Word 6 = Address of vector V2.
- Word 7 = Address of increment INCR2.
- Word 8 = Address of # elements N.

VPIV/DVPIV *

VECTOR PIVOT

15	14	13	12	11	10	9	8	7	6	5	4	3	2	1	0
1	0	0	0	1	0	1	0	0	1	0	P	0	0	0	1
0	X	X	X	X	X	X	X	X	X	X	X	X	X	X	X
D/I															
D/I															
D/I															
D/I															
D/I															
D/I															
D/I															
D/I															

Memory Address

VSDV/DVSDV *

SCALAR-VECTOR DIVIDE

15	14	13	12	11	10	9	8	7	6	5	4	3	2	1	0
1	0	0	0	1	0	1	0	0	0	0	P	1	0	0	1
0	X	X	X	X	X	X	X	X	X	X	X	X	X	X	X
D/I															
D/I															
D/I															
D/I															
D/I															
D/I															

Memory Address

Performs the vector operation:

$$V3 = S * V1 + V2$$

This is a ten-word instruction, where

- Word 1 = Instruction code.
- Word 2 = Return address.
- Word 3 = Address of scalar S.
- Word 4 = Address of vector V1.
- Word 5 = Address of increment INCR1.
- Word 6 = Address of vector V2.
- Word 7 = Address of increment INCR2.
- Word 8 = Address of vector V3.
- Word 9 = Address of increment INCR3.
- Word 10 = Address of # elements N.



Performs the vector operation:

$$V2 = S / V1$$

This is an eight-word instruction, where

- Word 1 = Instruction code.
- Word 2 = Return address.
- Word 3 = Address of scalar S.
- Word 4 = Address of vector V1.
- Word 5 = Address of increment INCR1.
- Word 6 = Address of vector V2.
- Word 7 = Address of increment INCR2.
- Word 8 = Address of # elements N.

VABS/DVABS *

VECTOR ABSOLUTE VALUE

15	14	13	12	11	10	9	8	7	6	5	4	3	2	1	0
1	0	0	0	1	0	1	0	0	1	0	P	0	0	1	1
0	X	X	X	X	X	X	X	X	X	X	X	X	X	X	X
D/I															
D/I															
D/I															
D/I															
D/I															

Memory Address

*For HP Assembly Language usage, refer to paragraph 3-46.

Performs the vector operation:

$$V2 = \text{ABS}(V1)$$

This is a seven-word instruction, where

- Word 1 = Instruction code.
- Word 2 = Return address.
- Word 3 = Address of vector V1.
- Word 4 = Address of increment INCR1.
- Word 5 = Address of vector V2.
- Word 6 = Address of increment INCR2.
- Word 7 = Address of # elements N.

VNRM/DVNRM *

VECTOR NORM

15	14	13	12	11	10	9	8	7	6	5	4	3	2	1	0
1	0	0	0	1	0	1	0	0	1	0	P	0	1	1	1
0	X	X	X	X	X	X	X	X	X	X	X	X	X	X	X
D/I															
D/I															
D/I															
D/I															

Memory Address

Performs the vector operation:

$$\text{SUM} = \sum \text{ABS}(V1)$$

Note that for VNRM the sum is internally accumulated in double precision; the answer is then truncated to single precision.

This is a six-word instruction, where

- Word 1 = Instruction code.
- Word 2 = Return address.
- Word 3 = Address of scalar SUM.
- Word 4 = Address of vector V1.
- Word 5 = Address of increment INCR1.
- Word 6 = Address of # elements N.

VSUM/DVSUM *

VECTOR SUM

15	14	13	12	11	10	9	8	7	6	5	4	3	2	1	0
1	0	0	0	1	0	1	0	0	1	0	P	0	1	0	1
0	X	X	X	X	X	X	X	X	X	X	X	X	X	X	X
D/I															
D/I															
D/I															
D/I															

Memory Address

Performs the vector operation:

$$\text{SUM} = \sum V1$$

Note that for VSUM the sum is internally accumulated in double precision; the answer is then truncated to single precision.

This is a six-word instruction, where

- Word 1 = Instruction code.
- Word 2 = Return address.
- Word 3 = Address of scalar SUM.
- Word 4 = Address of vector V1.
- Word 5 = Address of increment INCR1.
- Word 6 = Address of # elements N.

VDOT/DVDOT *

VECTOR DOT PRODUCT

15	14	13	12	11	10	9	8	7	6	5	4	3	2	1	0
1	0	0	0	1	0	1	0	0	1	0	P	1	0	0	0
0	X	X	X	X	X	X	X	X	X	X	X	X	X	X	X
D/I															
D/I															
D/I															
D/I															
D/I															
D/I															

Memory Address

*For HP Assembly Language usage, refer to paragraph 3-46.

Performs the vector operation:

$$DOT = \sum V1 * V2$$

Note that for VDOT the product and sum is internally accumulated in double precision; the answer is then truncated to single precision.

This is an eight-word instruction, where

- Word 1 = Instruction code.
- Word 2 = Return address.
- Word 3 = Address of scalar DOT.
- Word 4 = Address of vector V1.
- Word 5 = Address of increment INCR1.
- Word 6 = Address of vector V2.
- Word 7 = Address of increment INCR2.
- Word 8 = Address of # elements N.

VMAX/DVMAX* VECTOR MAXIMUM VALUE

15	14	13	12	11	10	9	8	7	6	5	4	3	2	1	0
1	0	0	0	1	0	1	0	0	1	0	P	1	0	0	1
0	X	X	X	X	X	X	X	X	X	X	X	X	X	X	X
D/I															
D/I															
D/I															
D/I															

Memory Address

Performs the vector operation:

$$IMAX = \text{Position} (\text{MAX}(V1))$$

Note that IMAX is the position of the maximum of those elements that were tested, as requested by INCR1 and N. If INCR1 # 1, the position is given by:

$$IPOS = 1 + \text{INCR1} * (\text{IMAX} - 1)$$

This is a six-word instruction, where

- Word 1 = Instruction code.
- Word 2 = Return address.
- Word 3 = Address of integer IMAX.
- Word 4 = Address of vector V1.
- Word 5 = Address of increment INCR1.
- Word 6 = Address of # elements N.

VMAB/DVMAB* VECTOR MAXIMUM ABSOLUTE VALUE

15	14	13	12	11	10	9	8	7	6	5	4	3	2	1	0
1	0	0	0	1	0	1	0	0	1	0	P	1	0	1	0
0	X	X	X	X	X	X	X	X	X	X	X	X	X	X	X
D/I															
D/I															
D/I															
D/I															

Memory Address

Performs the vector operation:

$$IMAB = \text{Position} (\text{MAX}(\text{ABS}(V1)))$$

Note that IMAB is the position of the maximum absolute value of those elements that were tested, as requested by INCR1 and N. If INCR1 # 1, the position is given by:

$$IPOS = 1 + \text{INCR1} * (\text{IMAB} - 1)$$

This is a six-word instruction, where

- Word 1 = Instruction code.
- Word 2 = Return address.
- Word 3 = Address of integer IMAB.
- Word 4 = Address of vector V1.
- Word 5 = Address of increment INCR1.
- Word 6 = Address of # elements N.

VMIN/DVMIN* VECTOR MINIMUM VALUE

15	14	13	12	11	10	9	8	7	6	5	4	3	2	1	0
1	0	0	0	1	0	1	0	0	1	0	P	1	0	1	1
0	X	X	X	X	X	X	X	X	X	X	X	X	X	X	X
D/I															
D/I															
D/I															
D/I															

Memory Address

*For HP Assembly Language usage, refer to paragraph 3-46.

Performs the vector operation:

$$IMIN = \text{Position}(\text{MIN}(V1))$$

Note that IMIN is the position of the minimum absolute value of those elements that were tested, as requested by INCR1 and N. If INCR1 # 1, the position is given by:

$$IPOS = 1 + INCR1 * (IMIN - 1)$$

This is a six-word instruction, where

- Word 1 = Instruction code.
- Word 2 = Return address.
- Word 3 = Address of integer IMIN.
- Word 4 = Address of vector V1.
- Word 5 = Address of increment INCR1.
- Word 6 = Address of # elements N.

VMIB/DVMIB* VECTOR MINIMUM ABSOLUTE VALUE

15	14	13	12	11	10	9	8	7	6	5	4	3	2	1	0
1	0	0	0	1	0	1	0	0	1	0	P	1	1	0	1
0	X	X	X	X	X	X	X	X	X	X	X	X	X	X	X
D/I															
D/I															
D/I															
D/I															

Memory Address

Performs the vector operation:

$$IMIB = \text{Position}(\text{MIN}(\text{ABS}(V1)))$$

Note that IMIB is the position of the minimum absolute value of those elements that were tested, as requested by INCR1 and N. If INCR1 # 1, the position is given by:

$$IPOS = 1 + INCR1 * (IMIB - 1)$$

This is a six-word instruction, where

- Word 1 = Instruction code.
- Word 2 = Return address.
- Word 3 = Address of integer IMIB.
- Word 4 = Address of vector V1.
- Word 5 = Address of increment INCR1.
- Word 6 = Address of # elements N.

*For HP Assembly Language usage, refer to paragraph 3-46.

VMOV/DVMOV* VECTOR MOVE

15	14	13	12	11	10	9	8	7	6	5	4	3	2	1	0
1	0	0	0	1	0	1	0	0	1	0	P	1	1	1	0
0	X	X	X	X	X	X	X	X	X	X	X	X	X	X	X
D/I															
D/I															
D/I															
D/I															
D/I															

Memory Address

Performs the vector operation:

$$V2 = V1$$

This is a seven-word instruction, where

- Word 1 = Instruction code.
- Word 2 = Return address.
- Word 3 = Address of vector V1.
- Word 4 = Address of increment INCR1.
- Word 5 = Address of vector V2.
- Word 6 = Address of increment INCR2.
- Word 7 = Address of # elements N.

VSWP/DVSWP* VECTOR SWAP

15	14	13	12	11	10	9	8	7	6	5	4	3	2	1	0
1	0	0	0	1	0	1	0	0	1	0	P	1	1	1	1
0	X	X	X	X	X	X	X	X	X	X	X	X	X	X	X
D/I															
D/I															
D/I															
D/I															
D/I															

Memory Address

Performs the vector operation:

$$V1 \Leftarrow V2$$

This is a seven-word instruction, where

- Word 1 = Instruction code.
- Word 2 = Return address.
- Word 3 = Address of vector V1.
- Word 4 = Address of increment INCR1.
- Word 5 = Address of vector V2.
- Word 6 = Address of increment INCR2.
- Word 7 = Address of # elements N.

Table 3-8. Typical Scientific Instruction Set Execution Times

Single-Precision Instructions	TIME (μ s) Min. - Max.	
Sine or Cosine	15 - 26	
Tangent	15 - 31	
Arc Tangent	3 - 36	
Hyperbolic Tangent	3 - 37	
Exponentiation	23 - 28	
Natural/Base 10 Logarithm	3 - 26	
Square Root	5 - 20	
Note: The maximum non-interruptible time is \leq 8 microseconds.		
Double-Precision Operations		
All of the single-precision instructions listed above are also available as double-precision operations. Although these had not been tested at the time of publication of this manual, they will probably take about 4 times as long to execute as the single-precision instructions.		
SIS Accuracy:	RMS Relative Error	
	Single-Precision	Double-Precision
Sine	9.2E-8	1.2E-16
Cosine	7.7E-8	1.3E-16
Tangent	1.5E-7	1.9E-16
Arc Tangent	1.5E-7	2.3E-16
Hyperbolic Tangent	2.2E-7	5.5E-16
Square Root	6.7E-8	1.6E-17
Exponentiation	3.2E-7	8.8E-17
Natural Logarithm	1.2E-7	1.3E-16
Base 10 Logarithm	1.6E-7	1.3E-16
Note: Memory refresh during a processor memory access can make an instruction approximately 3% slower. Heavy DMA activity can also degrade instruction times due to contention for memory.		

3-45. EXECUTION TIMES AND INTERRUPTS

Table 3-9 lists the typical execution times for the VIS instructions. VIS times are composed of two parts: a fixed time per instruction execution, and a per-element loop time. Thus the time to process 100 elements equals the fixed time plus 100 multiplied by loop time.

The maximum period of non-interruptible instruction execution for all VIS instructions is 24 microseconds. When a VIS instruction is interrupted, the current state of execution is stored in the first word of the result memory location, as well as the A, B, X, and Y registers. When re-entered following the interrupt processing, the VIS instruction will resume execution from the point of suspension.

Table 3-9. Typical Vector Instruction Set Execution Times

INSTRUCTION	EXECUTION TIME	
	FIXED	LOOP*
Single-Precision Instructions		
VADD, VSUB, VMPY	14.3	3.50
VDIV	14.3	6.25
VSAD, VSSB, VSMY	15.8	2.25
VSDV	15.8	5.00
VPIV	16.3	4.25
VABS	15.3	2.50
VSUM	15.3	2.50
VNRM	15.5	3.75
VDOT	18.5	5.75
VMAX, VMIN	14.5	3.50
VMAB, VMIB	14.5	3.75
VMOV	8.5	2.00
VSWP	8.5	4.00
Double-Precision Instructions		
DVADD, DVSUB, DVMPY	14.3	6.00
DVDIV	14.3	11.25
DVSAD, DVSSB, DVSMY	18.0	4.00
DVSDV	18.0	9.25
DVPIV	18.5	7.00
DVABS	15.8	3.00
DVSUM	15.3	3.00
DVNRM	15.5	3.25
DVDOT	18.5	6.50
DVMAX, DVMIN	15.0	4.25
DVMAB, DVMI B	15.3	4.50
DVMOV	9.50	4.00
DVSWP	9.50	8.00
NOTES: All times are in microseconds. Maximum non-interruptible time is \leq 24 microseconds for any VIS instruction.		
Memory refresh during a processor memory access can make an instruction approximately 3% slower. Heavy DMA activity can also degrade instruction times due to contention for memory.		
*Fixed time is instruction start-up time; loop time is the processing time per vector element. Total time equals fixed time plus the number of elements times loop time.		

3-46. ASSEMBLY LANGUAGE

New instructions not recognized by the HP Macroassembler require different handling in HP Assembly Language programming. These instructions are asterisked in the preceding paragraphs and must be used in the form: JSB x where x is the instruction. (The instruction, x, must be declared as an external at the beginning of the assembly language program.) Most of these

instructions correspond to library subroutines* and must be implemented into HP RTE systems (as described in the following paragraph) to enable their execution in hardware-firmware instead of in software.

3-47. RTE IMPLEMENTATION

New instructions can be implemented in an HP RTE-A.1 operating system simply by changing library entry points during the parameter input phase of system generation. (Refer to the appropriate RTE manual for the system generation procedure.) With the opcodes given in Table 3-10, the entry point changes would be as indicated below:

```
.JLA,RP,100600
.JLB,RP,104600
.
.
.MW21,RP,105736
.MW22,RP,105737
```

Alternatively, entry points may be changed by loading (via LOADR) a "replacement" program when user programs are loaded. Replacement programs RPL.7 and RPL.F are included in the RTE-A.1 system software.

*Refer to the Relocatable Library Reference Manual, part no. 24998-90001.

Table 3-10. Instructions and Opcodes for RTE Implementation

INSTRUCTION MNEMONIC	OCTAL OPCODE	INSTRUCTION MNEMONIC	OCTAL OPCODE	INSTRUCTION MNEMONIC	OCTAL OPCODE
JLA	100600	.DSB	105034	DVPIV	105121
JLB	104600	.DMP	105054	DVABS	105123
.FAD	105000	.DDI	105074	DVSUM	105125
.FSB	105020	.DSBR	105114	DVNRM	105127
.FMP	105040	.DDIR	105134	DVDOT	105130
.FDV	105060	.DNG	105203	DVMAX	105131
.FIX	105100	.DCO	105204	DVMAB	105132
.IFIX*	105100	.DIN	105210	DVMIN	105133
.FIXD	105104	.DDE	105211	DVMIB	105135
.FLT	105120	.DIS	105212	DVMOV	105136
FLOAT*	105120	.DDS	105213	DVSWP	105137
.FLTD	105124	.PMAP	105240	LPMR	105700
.TADD	105002	.IRES	105244	SPMR	105701
.TSUB	105022	.JRES	105245	LDMP	105702
.TMPY	105042	.IMAP	105250	STMP	105703
.TDIV	105062	.JMAP	105252	LWD1	105704
.TFXS	105102	.LPXR	105254	LWD2	105705
.TINT*	105102	.LPX	105255	SWMP	105706
.TFXD	105106	.LBPR	105256	SIMP	105707
.TFTS	105122	.LBP	105257	XJMP	105710
.ITBL*	105122	.CPUID	105300	XLA2	101721
.TFTD	105126	.FWID	105301	XSA2	101722
TAN	105320	.WFI	105302	XCA2	101723
SQRT	105321	.SIP	105303	XLA1	101724
ALOG	105322	VADD	105001	XSA1	101725
ATAN	105323	VSUB	105003	XCA1	101726
COS	105324	VMPY	105004	XLB2	105721
SIN	105325	VDIV	105005	XSB2	105722
EXP	105326	VSAD	105006	XCB2	105723
ALOGT	105327	VSSB	105007	XLB1	105724
TANH	105330	VSMY	105010	XSB1	105725
DPOLY	105331	VSDV	105011	XCB1	105726
/CMRT**	105332	VPIV	105101	MB00	101727
/ATLG	105333	VABS	105103	MB01	101730
.FPWR	105334	VSUM	105105	MB02	101731
.TPWR	105335	VNRM	105107	MB10	101732
.DFER	105205	VDOT	105110	MB11	101733
.BLE	105207	VMAX	105111	MB12	101734
.NGL	105214	VMAB	105112	MB20	101735
.XFER	105220	VMIN	105113	MB21	101736
.ENTR	105223	VMIB	105115	MB22	101737
.ENTP	105224	VMOV	105116	MW00	105727
.SETP	105227	VSWP	105117	MW01	105730
.CFER	105231	DVADD	105021	MW02	105731
..FCM	105232	DVSUB	105023	MW10	105732
..TCM	105233	DVMPY	105024	MW11	105733
.ENTN	105234	DVDIV	105025	MW12	105734
.ENTC	105235	DVSAD	105026	MW20	105735
.CPM	105236	DVSSB	105027	MW21	105736
.ZFER	105237	DVSMY	105030	MW22	105737
.DAD	105014	DVSDV	105031		

*Alternate mnemonic for the one preceding it.

**Not directly user callable. Used by HP 1000 software.

DYNAMIC MAPPING SYSTEM

SECTION

IV

The basic addressing space of the HP 1000 A700 computer is 32768 words, which is referred to as logical memory. The amount of memory actually installed in the computer system is referred to as physical memory. The Dynamic Mapping System (DMS) is standard logic in the HP 1000 A700 computer and provides an addressing capability for up to 16 million words of physical memory. The DMS allows logical memory to be mapped into physical memory through the use of dynamically-alterable memory maps.

4-1. MEMORY ADDRESSING

The basic memory addressing scheme provides for addressing 32 pages of logical memory, each of which consists of 1024 words. This memory is addressed through a 15-bit logical address bus as shown in Figure 4-1. The upper 5 bits of this bus provide the logical page address and the lower 10 bits provide the relative word offset within the page.

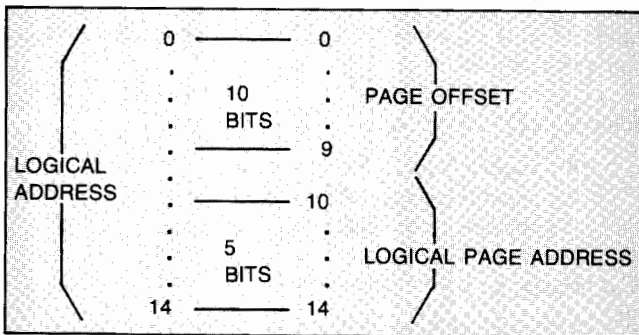


Figure 4-1. Basic Logical Memory Addressing Scheme

Also associated with any memory access is a 5-bit logical map number. The DMS converts the logical map number and the logical page address into a 14-bit physical page number, thereby allowing 16k (2^{14}) pages of physical memory to be addressed. This conversion is accomplished by having the 5-bit logical map number and the 5-bit logical page address access 1024 page mapping registers (PMRs), each of which is 16 bits wide. Each of these map registers contains the user-specified (privileged) 14-bit page address. This new page address is combined with the original 10-bit page offset to form a 24-bit memory address as shown in the figure 4-2.

The PMRs also contain two bits of memory protection information. Bit 15 indicates that the page is read-protected when privileged mode is disabled. Bit 14 indicates that the page is write-protected when privileged mode is disabled. Any attempt to read from a read-protected page will result in a read violation and the memory read will return an undefined result. Any attempt to write into a write-protected page will result in a write violation and the memory will not be altered.

If a read or write violation occurs, the DMS signals the memory protect logic (located on the memory controller card) that a violation has occurred, which causes the memory protect logic to generate an interrupt. As discussed in Section VI, memory protect violations are interrupted to select code 07.

The width of the PMRs is limited to a 16-bit word, of which two bits specify read/write protection, so the maximum width of the physical page address is 14 bits.

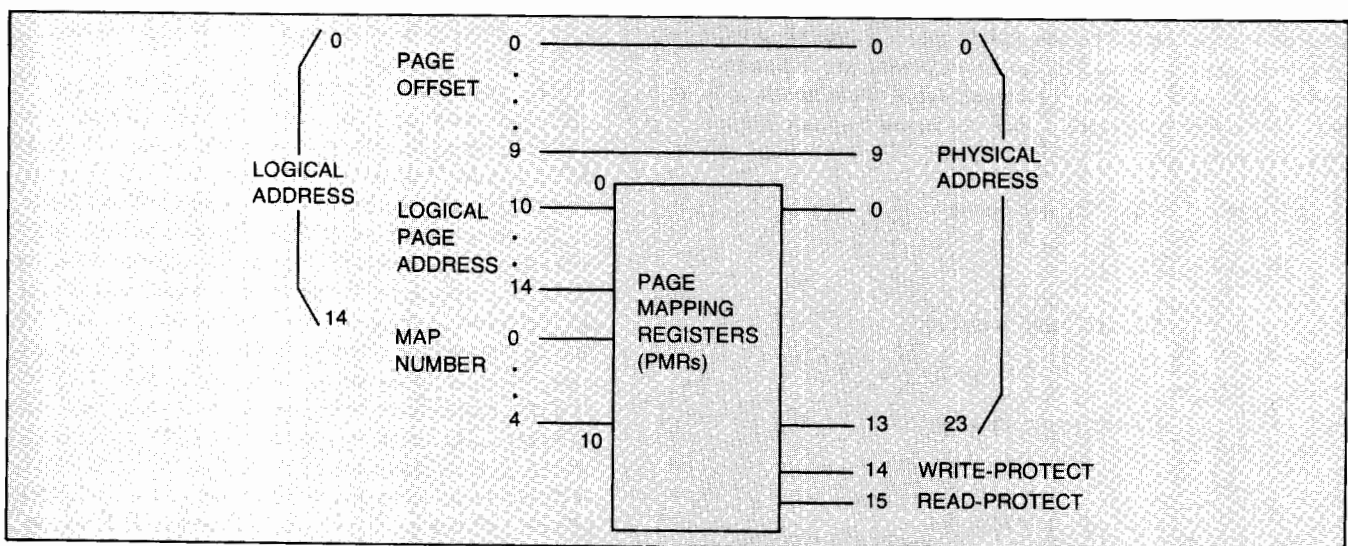


Figure 4-2. Expanded Memory Addressing Scheme

4-3. PAGE MAPPING REGISTER INSTRUCTIONS

The page mapping register instructions allow the privileged user to alter the PMRs, each of which have the following format:

PAGE MAPPING REGISTER FORMAT

0	}	physical page number
.		
.		
13		
14		— write protect this page
15		— read protect this page

The page mapping register instructions are:

- LPMR - load a PMR indexed by register A from register B
- SPMR - store a PMR indexed by register A to register B
- LDMP - load a map from memory
- STMP - store a map to memory

4-4. WORKING MAP INSTRUCTIONS

The computer will maintain three logical maps, cumulatively called the Working Map Set (WMAP). The working map instructions allow the system to alter the logical maps, and also to initiate a user program.

The Execute map is the map number used for instruction fetches and normal memory accesses. The data maps (DATA1 and DATA2) are the map numbers used in cross-map memory references. There are two data maps to allow the system to do cross-map moves from one area of memory to another without having to go through the system map. In addition, this feature allows the system to be able to quickly access one area of memory (such as a System Available Memory map) while being able to also access another (such as the user's map). A memory reference to locations 0 or 1 in the Execute map are defined to access the A- or B-registers, respectively. References to 0 or 1 in the data maps are defined to access physical memory locations.

The format of WMAP is as follows:

WMAP FORMAT:

0	}	Execute map number
.		
4		
5		
.	}	DATA1 map number
9		
10		
.		
14	}	DATA2 map number
15		

Upon servicing interrupts, the computer saves the currently executing WMAP in a register called IMAP, and loads WMAP with the following values:

- a. The DATA1 map is set to the old Execute map.
- b. The new Execute map is set to zero.
- c. The DATA2 map contains an undefined value.
- d. Memory protection is disabled.

The working map instructions are:

- XJMP - cross jump
- SWMP - store current WMAP into memory
- SIMP - store current IMAP into memory
- LWD1 - load WMAP field DATA1 from memory
- LWD2 - load WMAP field DATA2 from memory

4-5. CROSS-MAP INSTRUCTIONS

While the working map instructions provide a way to load the working map set, the cross-map instructions provide a means to use them.

These instructions are non-privileged. For all of these instructions, indirect DEF references are done through the Execute map, while the final reference is done through the specified map.

Abbreviations used are:

- "0" - means logical Execute map
- "1" - means logical DATA1 map
- "2" - means logical DATA2 map

The cross map instructions are:

- XLA1 - cross load A through the DATA1 map
- XLB1 - cross load B through the DATA1 map
- XLA2 - cross load A through the DATA2 map
- XLB2 - cross load B through the DATA2 map
- XSA1 - cross store A through the DATA1 map
- XSB1 - cross store B through the DATA1 map
- XSA2 - cross store A through the DATA2 map
- XSB2 - cross store B through the DATA2 map
- XCA1 - cross compare A through the DATA1 map
- XCB1 - cross compare B through the DATA1 map
- XCA2 - cross compare A through the DATA2 map
- XCB2 - cross compare B through the DATA2 map
- MW00 - cross move words from Execute to Execute
- MW01 - cross move words from Execute to DATA1
- MW02 - cross move words from Execute to DATA2
- MW10 - cross move words from DATA1 to Execute
- MW11 - cross move words from DATA1 to DATA1
- MW12 - cross move words from DATA1 to DATA2
- MW20 - cross move words from DATA2 to Execute
- MW21 - cross move words from DATA2 to DATA1
- MW22 - cross move words from DATA2 to DATA2
- MB00 - cross move bytes from Execute to Execute
- MB01 - cross move bytes from Execute to DATA1
- MB02 - cross move bytes from Execute to DATA2

- MB10 - cross move bytes from DATA1 to Execute
- MB11 - cross move bytes from DATA1 to DATA1
- MB12 - cross move bytes from DATA1 to DATA2
- MB20 - cross move bytes from DATA2 to Execute
- MB21 - cross move bytes from DATA2 to DATA1
- MB22 - cross move bytes from DATA2 to DATA2

4-6. DETAILED DESCRIPTIONS

The following paragraphs provide machine language coding and definitions for the DMS instructions.

LPMR LOAD PAGE MAPPING REGISTER

15	14	13	12	11	10	9	8	7	6	5	4	3	2	1	0
1	0	0	0	1	0	1	1	1	1	0	0	0	0	0	0

Loads the contents of the B-register into the page mapping register (PMR) addressed by the contents of the A-register. Any attempt to address a PMR outside the range of 0 to 1023 produces undefined results. The format for the PMR contents is: bit 15 = read protect; bit 14 = write protect; and bits 13 to 0 = physical page number. This instruction is privileged. After the operation, the A-register is incremented.

SPMR STORE PAGE MAPPING REGISTER

15	14	13	12	11	10	9	8	7	6	5	4	3	2	1	0
1	0	0	0	1	0	1	1	1	1	0	0	0	0	0	1

Loads the contents of the page mapping register (PMR) addressed by the value in the A-register into the B-register. Any attempt to address a PMR outside the range of 0 to 1023 produces undefined results. The format for the PMR contents is: bit 15 = read protect; bit 14 = write protect; and bits 13 to 0 = physical page number. This instruction is privileged. After the operation, the A-register is incremented.

LDMP LOAD A MAP

15	14	13	12	11	10	9	8	7	6	5	4	3	2	1	0
1	0	0	0	1	0	1	1	1	1	0	0	0	0	1	0
D _i															
D _i															

Loads the map number specified by Word 2 from the 32-word block of memory specified by Word 3, where:

- Word 1 = Instruction code.
- Word 2 = Pointer to Map number.
- Word 3 = Pointer to Map image.

There are 32 maps of 32 PMRs each; the beginning PMR number of a map is related to the map number as follows:

$$\text{PMR number} = \text{Map number} \times 32$$

Undefined results occur when a map number outside the range of 0 to 31 is addressed, when modification of a currently executing map is tried, or when the resolved address of the map image is outside the range of 2 to 77740 octal.

All memory references are done in the Execute map and may include the A- and B-registers. This instruction is privileged and is interruptible in that it may be restarted during indirect address resolution after three levels of indirection.

STMP STORE A MAP

15	14	13	12	11	10	9	8	7	6	5	4	3	2	1	0
1	0	0	0	1	0	1	1	1	1	0	0	0	0	1	1
D _i															
D _i															

Stores the map number specified by Word 2 to the 32-word block of memory specified by Word 3, where:

- Word 1 = Instruction code.
- Word 2 = Pointer to Map number.
- Word 3 = Pointer to Map image.

There are 32 maps of 32 PMRs each; the beginning PMR number of a map is related to the map number as follows:

$$\text{PMR number} = \text{Map number} \times 32$$

Undefined results occur when a map number outside the range of 0 to 31 is addressed, when modification of a currently executing map is tried, or when the resolved address of the map image is outside the range of 2 to 77740 octal.

All memory references are done in the Execute map and may include the A- and B-registers. This instruction is privileged and is interruptible in that it may be restarted during indirect address resolution after three levels of indirection.

XJMP CROSS MAP JUMP

15	14	13	12	11	10	9	8	7	6	5	4	3	2	1	0
1	0	0	0	1	0	1	1	1	1	0	0	1	0	0	0
D _i															
D _i															

Resolves indirect references, sets the program counter to the resolved address specified by Word 3, and loads WMAP with the contents of Word 2, where:

- Word 1 = Instruction code.
- Word 2 = Pointer to new WMAP number.
- Word 3 = Pointer to next instruction (new PC value).

All memory references (direct and indirect) are done in the Execute map and may include the A- and B-registers. The next instruction will be fetched using the new WMAP. This instruction is privileged and is interruptible in that it may be restarted during indirect address resolution after three levels of indirection.

SWMP **SAVE WORKING MAP**

15	14	13	12	11	10	9	8	7	6	5	4	3	2	1	0
1	0	0	0	1	0	1	1	1	1	0	0	0	1	1	0
D ₁															

Stores WMAP at the memory location pointed to by Word 2, where:

- Word 1 = Instruction code.
- Word 2 = Pointer to destination in memory.

All memory references are done in the Execute map and may include the A- and B-registers. This instruction is privileged and is interruptible in that it may be restarted during indirect address resolution after three levels of indirection.

SIMP **SAVE INTERRUPTED MAP**

15	14	13	12	11	10	9	8	7	6	5	4	3	2	1	0
1	0	0	0	1	0	1	1	1	1	0	0	0	1	1	1
D ₁															

Stores IMAP at the location pointed to by Word 2, where:

- Word 1 = Instruction code.
- Word 2 = Pointer to destination in memory.

All memory references are done in the Execute map and may include the A- and B-registers. This instruction is privileged and is interruptible in that it may be restarted during indirect address resolution after three levels of indirection.

LWD1 **LOAD DATA1 MAP**

15	14	13	12	11	10	9	8	7	6	5	4	3	2	1	0
1	0	0	0	1	0	1	1	1	1	0	0	0	1	0	0
D ₁															

Loads the DATA1 register from the memory location pointed to by Word 2, where:

- Word 1 = Instruction code.
- Word 2 = Pointer to new DATA1 map.

All memory references are done in the Execute map and may include the A- and B-registers. This instruction is privileged and is interruptible in that it may be restarted during indirect address resolution after three levels of indirection. Map numbers outside the range of 0-31 produce undefined results.

LWD2 **LOAD DATA2 MAP**

15	14	13	12	11	10	9	8	7	6	5	4	3	2	1	0
1	0	0	0	1	0	1	1	1	1	0	0	0	1	0	1
D ₁															

- Word 1 = Instruction code.
- Word 2 = Pointer to new DATA2 map.

All memory references are done in the Execute map and may include the A- and B-registers. This instruction is privileged and is interruptible in that it may be restarted during indirect address resolution after three levels of indirection. Map numbers outside the range of 0-31 produce undefined results.

Loads the DATA2 register from the memory location pointed to by Word 2, where:

XLA1 **CROSS LOAD A THROUGH DATA1 MAP**

15	14	13	12	11	10	9	8	7	6	5	4	3	2	1	0
1	0	0	0	0	0	1	1	1	1	0	1	0	1	0	0
D ₁															

Loads the A-register from the memory location pointed to by Word 2, where:

- Word 1 = Instruction code.
- Word 2 = Pointer to memory location in DATA1 map.

All indirect memory references are done in the Execute map and may include the A- and B-registers. The direct memory reference is done in the DATA1 map. Because A- and B-register addressing is disabled in the DATA1 map, direct addresses 0 and 1 refer to physical memory locations. This instruction is interruptible in that it may be restarted during indirect address resolution after three levels of indirection.

XLB1 CROSS LOAD B THROUGH DATA1 MAP

15	14	13	12	11	10	9	8	7	6	5	4	3	2	1	0
1	0	0	0	1	0	1	1	1	1	0	1	0	1	0	0
D ₁															

Loads the B-register from the memory location pointed to by Word 2, where:

Word 1 = Instruction code.

Word 2 = Pointer to memory location in DATA1 map.

All indirect memory references are done in the Execute map and may include the A- and B-registers. The direct memory reference is done in the DATA1 map. Because A- and B-register addressing is disabled in the DATA1 map, direct addresses 0 and 1 refer to physical memory locations. This instruction is interruptible in that it may be restarted during indirect address resolution after three levels of indirection.

XLA2 CROSS LOAD A THROUGH DATA2 MAP

15	14	13	12	11	10	9	8	7	6	5	4	3	2	1	0
1	0	0	0	0	0	1	1	1	1	0	1	0	0	0	1
D ₁															

Loads the A-register from the memory location pointed to by Word 2, where:

Word 1 = Instruction code.

Word 2 = Pointer to memory location in DATA2 map.

All indirect memory references are done in the Execute map and may include the A- and B-registers. The direct memory reference is done in the DATA2 map. Because A- and B-register addressing is disabled in the DATA2 map, direct addresses 0 and 1 refer to physical memory locations. This instruction is interruptible in that it may be restarted during indirect address resolution after three levels of indirection.

XLB2 CROSS LOAD B THROUGH DATA2 MAP

15	14	13	12	11	10	9	8	7	6	5	4	3	2	1	0
1	0	0	0	1	0	1	1	1	1	0	1	0	0	0	1
D ₁															

Loads the B-register from the memory location pointed to by Word 2, where:

Word 1 = Instruction code.

Word 2 = Pointer to memory location in DATA2 map.

All indirect memory references are done in the Execute map and may include the A- and B-registers. The direct memory reference is done in the DATA2 map. Because A- and B-register addressing is disabled in the DATA2 map, direct addresses 0 and 1 refer to physical memory locations. This instruction is interruptible in that it may be restarted during indirect address resolution after three levels of indirection.

XSA1 CROSS STORE A THROUGH DATA1 MAP

15	14	13	12	11	10	9	8	7	6	5	4	3	2	1	0
1	0	0	0	0	0	1	1	1	1	0	1	0	1	0	1
D ₁															

Stores the A-register contents in the memory location pointed to by Word 2, where:

Word 1 = Instruction code.

Word 2 = Pointer to memory location in DATA1 map.

All indirect memory references are done in the Execute map and may include the A- and B-registers. The direct memory reference is done in the DATA1 map. Because A- and B-register addressing is disabled in the DATA1 map, direct addresses 0 and 1 refer to physical memory locations. This instruction is interruptible in that it may be restarted during indirect address resolution after three levels of indirection.

XSB1 CROSS STORE B THROUGH DATA1 MAP

15	14	13	12	11	10	9	8	7	6	5	4	3	2	1	0
1	0	0	0	1	0	1	1	1	1	0	1	0	1	0	1
D ₁															

Stores the B-register contents in the memory location pointed to by Word 2, where:

Word 1 = Instruction code.

Word 2 = Pointer to memory location in DATA1 map.

All indirect memory references are done in the Execute map and may include the A- and B-registers. The direct memory reference is done in the DATA1 map. Because A- and B-register addressing is disabled in the DATA1 map, direct addresses 0 and 1 refer to physical memory locations. This instruction is interruptible in that it may be restarted during indirect address resolution after three levels of indirection.

XSA2 CROSS STORE A THROUGH DATA2 MAP

15	14	13	12	11	10	9	8	7	6	5	4	3	2	1	0
1	0	0	0	0	0	1	1	1	1	0	1	0	0	1	0
D ₁															

Stores the A-register contents in the memory location pointed to by Word 2, where:

Word 1 = Instruction code.

Word 2 = Pointer to memory location in DATA2 map.

All indirect memory references are done in the Execute map and may include the A- and B-registers. The direct memory reference is done in the DATA2 map. Because A- and B-register addressing is disabled in the DATA2 map, direct addresses 0 and 1 refer to physical memory locations. This instruction is interruptible in that it may be restarted during indirect address resolution after three levels of indirection.

XSB2 CROSS STORE B THROUGH DATA2 MAP

15	14	13	12	11	10	9	8	7	6	5	4	3	2	1	0
1	0	0	0	1	0	1	1	1	1	0	1	0	0	1	0
D ₁															

Stores the B-register contents in the memory location pointed to by Word 2, where:

Word 1 = Instruction code.

Word 2 = Pointer to memory location in DATA2 map.

All indirect memory references are done in the Execute map and may include the A- and B-registers. The direct memory reference is done in the DATA2 map. Because A- and B-register addressing is disabled in the DATA2 map, direct addresses 0 and 1 refer to physical memory locations. This instruction is interruptible in that it may be restarted during indirect address resolution after three levels of indirection.

XCA1 CROSS COMPARE A THROUGH DATA1 MAP

15	14	13	12	11	10	9	8	7	6	5	4	3	2	1	0
1	0	0	0	0	0	1	1	1	1	0	1	0	1	1	0
D ₁															

Compares the A-register contents with a value in the memory location pointed to by Word 2 and skips if the values are not equal, where:

Word 1 = Instruction code.

Word 2 = Pointer to memory location in DATA1 map.

All indirect memory references are done in the Execute map and may include the A- and B-registers. The direct memory reference is done in the DATA1 map. Because A- and B-register addressing is disabled in the DATA1 map, direct addresses 0 and 1 refer to physical memory locations. This instruction is interruptible in that it may be restarted during indirect address resolution after three levels of indirection.

XCB1 CROSS COMPARE B THROUGH DATA1 MAP

15	14	13	12	11	10	9	8	7	6	5	4	3	2	1	0
1	0	0	0	1	0	1	1	1	1	0	1	0	1	1	0
D ₁															

Compares the B-register contents with a value in the memory location pointed to by Word 2 and skips if the values are not equal, where:

Word 1 = Instruction code.

Word 2 = Pointer to memory location in DATA1 map.

All indirect memory references are done in the Execute map and may include the A- and B-registers. The direct memory reference is done in the DATA1 map. Because A- and B-register addressing is disabled in the DATA1 map, direct addresses 0 and 1 refer to physical memory locations. This instruction is interruptible in that it may be restarted during indirect address resolution after three levels of indirection.

XCA2 CROSS COMPARE A THROUGH DATA2 MAP

15	14	13	12	11	10	9	8	7	6	5	4	3	2	1	0
1	0	0	0	0	0	1	1	1	1	0	1	0	0	1	1
D ₁															

Compares the A-register contents with a value in the memory location pointed to by Word 2 and skips if the values are not equal, where:

Word 1 = Instruction code.

Word 2 = Pointer to memory location in DATA2 map.

All indirect memory references are done in the Execute map and may include the A- and B-registers. The direct memory reference is done in the DATA2 map. Because A- and B-register addressing is disabled in the DATA2 map, direct addresses 0 and 1 refer to physical memory locations. This instruction is interruptible in that it may be restarted during indirect address resolution after three levels of indirection.

XC82 CROSS COMPARE B THROUGH
DATA2 MAP

15	14	13	12	11	10	9	8	7	6	5	4	3	2	1	0
1	0	0	0	1	0	1	1	1	1	0	1	0	0	1	1
D ₁															

Compares the B-register contents with a value in the memory location pointed to by Word 2 and skips if the values are not equal, where:

Word 1 = Instruction code.

Word 2 = Pointer to memory location in DATA2 map.

All indirect memory references are done in the Execute map and may include the A- and B-registers. The direct memory reference is done in the DATA2 map. Because A- and B-register addressing is disabled in the DATA2 map, direct addresses 0 and 1 refer to physical memory locations. This instruction is interruptible in that it may be restarted during indirect address resolution after three levels of indirection.

MW00 CROSS MOVE WORDS,
EXECUTE TO EXECUTE

15	14	13	12	11	10	9	8	7	6	5	4	3	2	1	0
1	0	0	0	1	0	1	1	1	1	0	1	0	1	1	1

Moves a block of words from the Execute map to the Execute map. The A-register specifies the source address, the B-register specifies the destination address, and the X-register specifies the number of words to be moved (which must be a positive integer equal to or greater than zero). Address bit 15 must be zero, as indirect source and destination references are not allowed. On return, the A- and B-registers contain the original values incremented by the number of words moved, and the X-register is zero.

This instruction produces undefined results if the source or destination address rolls over. It is interruptible, with the context saved in the A-, B- and X-registers.

MW01 CROSS MOVE WORDS,
EXECUTE TO DATA1

15	14	13	12	11	10	9	8	7	6	5	4	3	2	1	0
1	0	0	0	1	0	1	1	1	1	0	1	1	0	0	0

Moves a block of words from the Execute map to the DATA1 map. The A-register specifies the source address in the Execute map, the B-register specifies the destina-

tion address in the DATA1 map, and the X-register specifies the number of words to be moved (which must be a positive integer equal to or greater than zero). Address bit 15 must be zero, as indirect source and destination references are not allowed. On return, the A- and B-registers contain the original values incremented by the number of words moved, and the X-register is zero.

This instruction produces undefined results if the source or destination address rolls over. It is interruptible, with the context saved in the A-, B- and X-registers.

MW02 CROSS MOVE WORDS,
EXECUTE TO DATA2

15	14	13	12	11	10	9	8	7	6	5	4	3	2	1	0
1	0	0	0	1	0	1	1	1	1	0	1	1	0	0	1

Moves a block of words from the Execute map to the DATA2 map. The A-register specifies the source address in the Execute map, the B-register specifies the destination address in the DATA2 map, and the X-register specifies the number of words to be moved (which must be a positive integer equal to or greater than zero). Address bit 15 must be zero, as indirect source and destination references are not allowed. On return, the A- and B-registers contain the original values incremented by the number of words moved, and the X-register is zero.

This instruction produces undefined results if the source or destination address rolls over. It is interruptible, with the context saved in the A-, B- and X-registers.

MW10 CROSS MOVE WORDS,
DATA1 TO EXECUTE

15	14	13	12	11	10	9	8	7	6	5	4	3	2	1	0
1	0	0	0	1	0	1	1	1	1	0	1	1	0	1	0

Moves a block of words from the DATA1 map to the Execute map. The A-register specifies the source address in the DATA1 map, the B-register specifies the destination address in the Execute map, and the X-register specifies the number of words to be moved (which must be a positive integer equal to or greater than zero). Address bit 15 must be zero, as indirect source and destination references are not allowed. On return, the A- and B-registers contain the original values incremented by the number of words moved, and the X-register is zero.

This instruction produces undefined results if the source or destination address rolls over. It is interruptible, with the context saved in the A-, B- and X-registers.

MW11 CROSS MOVE WORDS, DATA1 TO DATA1

15	14	13	12	11	10	9	8	7	6	5	4	3	2	1	0
1	0	0	0	1	0	1	1	1	1	0	1	1	0	1	1

Moves a block of words from one location in the DATA1 map to another in the DATA1 map. The A-register specifies the source address, the B-register specifies the destination address, and the X-register specifies the number of words to be moved (which must be a positive integer equal to or greater than zero). Address bit 15 must be zero, as indirect source and destination references are not allowed. On return, the A- and B-registers contain the original values incremented by the number of words moved, and the X-register is zero.

This instruction produces undefined results if the source or destination address rolls over. It is interruptible, with the context saved in the A-, B- and X-registers.

MW12 CROSS MOVE WORDS, DATA1 TO DATA2

15	14	13	12	11	10	9	8	7	6	5	4	3	2	1	0
1	0	0	0	1	0	1	1	1	1	0	1	1	1	0	0

Moves a block of words from the DATA1 map to the DATA2 map. The A-register specifies the source address in the DATA1 map, the B-register specifies the destination address in the DATA2 map, and the X-register specifies the number of words to be moved (which must be a positive integer equal to or greater than zero). Address bit 15 must be zero, as indirect source and destination references are not allowed. On return, the A- and B-registers contain the original values incremented by the number of words moved, and the X-register is zero.

This instruction produces undefined results if the source or destination address rolls over. It is interruptible, with the context saved in the A-, B- and X-registers.

MW20 CROSS MOVE WORDS, DATA2 TO EXECUTE

15	14	13	12	11	10	9	8	7	6	5	4	3	2	1	0
1	0	0	0	1	0	1	1	1	1	0	1	1	1	0	1

Moves a block of words from the DATA2 map to the Execute map. The A-register specifies the source address in the DATA2 map, the B-register specifies the destina-

tion address in the Execute map, and the X-register specifies the number of words to be moved (which must be a positive integer equal to or greater than zero). Address bit 15 must be zero, as indirect source and destination references are not allowed. On return, the A- and B-registers contain the original values incremented by the number of words moved, and the X-register is zero.

This instruction produces undefined results if the source or destination address rolls over. It is interruptible, with the context saved in the A-, B- and X-registers.

MW21 CROSS MOVE WORDS, DATA2 TO DATA1

15	14	13	12	11	10	9	8	7	6	5	4	3	2	1	0
1	0	0	0	1	0	1	1	1	1	0	1	1	1	1	0

Moves a block of words from the DATA2 map to the DATA1 map. The A-register specifies the source address in the DATA2 map, the B-register specifies the destination address in the DATA1 map, and the X-register specifies the number of words to be moved (which must be a positive integer equal to or greater than zero). Address bit 15 must be zero, as indirect source and destination references are not allowed. On return, the A- and B-registers contain the original values incremented by the number of words moved, and the X-register is zero.

This instruction produces undefined results if the source or destination address rolls over. It is interruptible, with the context saved in the A-, B- and X-registers.

MW22 CROSS MOVE WORDS, DATA2 TO DATA2

15	14	13	12	11	10	9	8	7	6	5	4	3	2	1	0
1	0	0	0	1	0	1	1	1	1	0	1	1	1	1	1

Moves a block of words from the DATA2 map to the DATA2 map. The A-register specifies the source address, the B-register specifies the destination address, and the X-register specifies the number of words to be moved (which must be a positive integer equal to or greater than zero). Address bit 15 must be zero, as indirect source and destination references are not allowed. On return, the A- and B-registers contain the original values incremented by the number of words moved, and the X-register is zero.

This instruction produces undefined results if the source or destination address rolls over. It is interruptible, with the context saved in the A-, B- and X-registers.

MB00 CROSS MOVE BYTES,
EXECUTE TO EXECUTE

15	14	13	12	11	10	9	8	7	6	5	4	3	2	1	0
1	0	0	0	0	0	1	1	1	1	0	1	0	1	1	1

Moves a block of bytes from one location in the Execute map to another in the Execute map. The A-register specifies the source address and the B-register specifies the destination address. The X-register specifies the number of bytes to be moved (which is an unsigned 16-bit number that may equal zero). Indirect addressing is not allowed because a byte address uses all 16 bits. A byte address is two times the word address plus zero or one, which specifies the high order (bits 15 to 8) or low order (bits 7 to 0) position, respectively. On return, the A- and B-registers contain the original values incremented by the number of bytes moved, and the X-register is zero.

This instruction produces undefined results if the source or destination address rolls over. It is interruptible, with the context saved in the A-, B- and X-registers.

MB01 CROSS MOVE BYTES,
EXECUTE TO DATA1

15	14	13	12	11	10	9	8	7	6	5	4	3	2	1	0
1	0	0	0	0	0	1	1	1	1	0	1	1	0	0	0

Moves a block of bytes from a location in the Execute map to one in the DATA1 map. The A-register specifies the source address in the Execute map, and the B-register specifies the destination address in the DATA1 map. The X-register specifies the number of bytes to be moved (which is an unsigned 16-bit number that may equal zero). Indirect addressing is not allowed because a byte address uses all 16 bits. A byte address is two times the word address plus zero or one, which specifies the high order (bits 15 to 8) or low order (bits 7 to 0) position, respectively. On return, the A- and B-registers contain the original values incremented by the number of bytes moved, and the X-register is zero.

This instruction produces undefined results if the source or destination address rolls over. It is interruptible, with the context saved in the A-, B- and X-registers.

MB02 CROSS MOVE BYTES,
EXECUTE TO DATA2

15	14	13	12	11	10	9	8	7	6	5	4	3	2	1	0
1	0	0	0	0	0	1	1	1	1	0	1	1	0	0	1

Moves a block of bytes from a location in the Execute map to one in the DATA2 map. The A-register specifies the source address in the Execute map, and the B-register

specifies the destination address in the DATA2 map. The X-register specifies the number of bytes to be moved (which is an unsigned 16-bit number that may equal zero). Indirect addressing is not allowed because a byte address uses all 16 bits. A byte address is two times the word address plus zero or one, which specifies the high order (bits 15 to 8) or low order (bits 7 to 0) position, respectively. On return, the A- and B-registers contain the original values incremented by the number of bytes moved, and the X-register is zero.

This instruction produces undefined results if the source or destination address rolls over. It is interruptible, with the context saved in the A-, B- and X-registers.

MB10 CROSS MOVE BYTES,
DATA1 TO EXECUTE

15	14	13	12	11	10	9	8	7	6	5	4	3	2	1	0
1	0	0	0	0	0	1	1	1	1	0	1	1	0	1	0

Moves a block of bytes from a location in the DATA1 map to one in the Execute map. The A-register specifies the source address in the DATA1 map, and the B-register specifies the destination address in the Execute map. The X-register specifies the number of bytes to be moved (which is an unsigned 16-bit number that may equal zero). Indirect addressing is not allowed because a byte address uses all 16 bits. A byte address is two times the word address plus zero or one, which specifies the high order (bits 15 to 8) or low order (bits 7 to 0) position, respectively. On return, the A- and B-registers contain the original values incremented by the number of bytes moved, and the X-register is zero.

This instruction produces undefined results if the source or destination address rolls over. It is interruptible, with the context saved in the A-, B- and X-registers.

MB11 CROSS MOVE BYTES, DATA1 TO DATA1

15	14	13	12	11	10	9	8	7	6	5	4	3	2	1	0
1	0	0	0	0	0	1	1	1	1	0	1	1	0	1	1

Moves a block of bytes from one location in the DATA1 map to another in the DATA1 map. The A-register specifies the source address and the B-register specifies the destination address. The X-register specifies the number of bytes to be moved (which is an unsigned 16-bit number that may equal zero). Indirect addressing is not allowed because a byte address uses all 16 bits. A byte address is two times the word address plus zero or one, which specifies the high order (bits 15 to 8) or low order (bits 7 to 0) position, respectively. On return, the A- and B-registers contain the original values incremented by the number of bytes moved, and the X-register is zero.

This instruction produces undefined results if the source or destination address rolls over. It is interruptible, with the context saved in the A-, B- and X-registers.

MB12 CROSS MOVE BYTES, DATA1 TO DATA2

15	14	13	12	11	10	9	8	7	6	5	4	3	2	1	0
1	0	0	0	0	0	1	1	1	1	0	1	1	1	0	0

Moves a block of bytes from a location in the DATA1 map to one in the DATA2 map. The A-register specifies the source address in the DATA1 map, and the B-register specifies the destination address in the DATA2 map. The X-register specifies the number of bytes to be moved (which is an unsigned 16-bit number that may equal zero). Indirect addressing is not allowed because a byte address uses all 16 bits. A byte address is two times the word address plus zero or one, which specifies the high order (bits 15 to 8) or low order (bits 7 to 0) position, respectively. On return, the A- and B-registers contain the original values incremented by the number of bytes moved, and the X-register is zero.

This instruction produces undefined results if the source or destination address rolls over. It is interruptible, with the context saved in the A-, B- and X-registers.

MB20 CROSS MOVE BYTES, DATA2 TO EXECUTE

15	14	13	12	11	10	9	8	7	6	5	4	3	2	1	0
1	0	0	0	0	0	1	1	1	1	0	1	1	1	0	1

Moves a block of bytes from a location in the DATA2 map to one in the Execute map. The A-register specifies the source address in the DATA2 map, and the B-register specifies the destination address in the Execute map. The X-register specifies the number of bytes to be moved (which is an unsigned 16-bit number that may equal zero). Indirect addressing is not allowed because a byte address uses all 16 bits. A byte address is two times the word address plus zero or one, which specifies the high order (bits 15 to 8) or low order (bits 7 to 0) position, respectively. On return, the A- and B-registers contain the original values incremented by the number of bytes moved, and the X-register is zero.

This instruction produces undefined results if the source or destination address rolls over. It is interruptible, with the context saved in the A-, B- and X-registers.

MB21 CROSS MOVE BYTES, DATA2 TO DATA1

15	14	13	12	11	10	9	8	7	6	5	4	3	2	1	0
1	0	0	0	0	0	1	1	1	1	0	1	1	1	1	0

Moves a block of bytes from a location in the DATA2 map to one in the DATA1 map. The A-register specifies the source address in the DATA2 map, and the B-register specifies the destination address in the DATA1 map. The X-register specifies the number of bytes to be moved (which is an unsigned 16-bit number that may equal zero). Indirect addressing is not allowed because a byte address uses all 16 bits. A byte address is two times the word address plus zero or one, which specifies the high order (bits 15 to 8) or low order (bits 7 to 0) position, respectively. On return, the A- and B-registers contain the original values incremented by the number of bytes moved, and the X-register is zero.

This instruction produces undefined results if the source or destination address rolls over. It is interruptible, with the context saved in the A-, B- and X-registers.

MB22 CROSS MOVE BYTES, DATA2 TO DATA2

15	14	13	12	11	10	9	8	7	6	5	4	3	2	1	0
1	0	0	0	0	0	1	1	1	1	0	1	1	1	1	1

Moves a block of bytes from one location in the DATA2 map to another in the DATA2 map. The A-register specifies the source address and the B-register specifies the destination address. The X-register specifies the number of bytes to be moved (which is an unsigned 16-bit number that may equal zero). Indirect addressing is not allowed because a byte address uses all 16 bits. A byte address is two times the word address plus zero or one, which specifies the high order (bits 15 to 8) or low order (bits 7 to 0) position, respectively. On return, the A- and B-registers contain the original values incremented by the number of bytes moved, and the X-register is zero.

This instruction produces undefined results if the source or destination address rolls over. It is interruptible, with the context saved in the A-, B- and X-registers.

4-7. DMS INSTRUCTION EXECUTION TIMES

Table 4-1 lists the execution times for the various DMS instructions.

4-8. ASSEMBLY LANGUAGE AND RTE IMPLEMENTATION

Refer to paragraphs 3-46 and 3-47 for information on implementing the DMS instructions in HP Assembly Language and in an HP RTE-A.1 operating system.

Table 4-1. Dynamic Mapping Instructions Execution Times

INSTRUCTION	EXECUTION TIME (μ s)
XLA1/XLB1, XLA2/XLB2	2.75
XSA1/XSB1, XSA2/XSB2	2.75
XCA1/XCB2, XCA2/XCB2	3.25
MW00, MW01, MW02, MW10, MW11, MW12, MW20, MW21, MW22	3.50 plus 1.00 per word moved
MB00, MB01, MB02, MB10, MB11, MB12, MB20, MB21, MB22	3.50 plus 2.25 per byte moved
LPMR	1.50
SPMR	1.25
LDMP, STMP	20.0
XJMP	5.75
LWD1, LWD2	2.50
SWMP	5.0
SIMP	2.0

Note: Memory refresh during a processor memory access can make an instruction approximately 3% slower. Heavy DMA activity can also degrade instruction times due to contention for memory.



This section contains an introductory discussion of the A700 computer microprogramming techniques and development. For additional information, refer to the *HP 92045A Microprogramming Package Reference Manual*, part no. 92045-90001.

5-1. THE MICROPROGRAMMED COMPUTER

The control section of a computer is the portion of the computer that directs and controls the other sections; i.e., the memory section, input-output section, and the arithmetic-logic section. In totally hardwired computers, the control section logic is normally "spread out" physically throughout the computer. This design approach makes it impossible to enhance the computer's instruction set without redesign. In contrast, A700 computers have a fully microprogrammed control section, which means that the sequence in which the control functions are performed are made programmable through the use of a technique called microprogramming.

The action taken when any one of the A700 base set of 205 assembly language instructions is executed is determined by a microprogram associated with the assembly language instruction (these microprograms reside in a special memory called control store); the control section oversees the translation and controls the execution of the microprogram. With this design approach, instruction set enhancements can be made by changing or adding to the set of microprograms that control the machine's execution. Many computers are microprogrammed; however, Hewlett-Packard has taken the concept one step further to offer the power of microprogramming to the user.

5-2. THE MICROPROGRAMMABLE COMPUTER

A700 computer users can more fully take advantage of the computer's power by utilizing microprogramming. The microprogrammer has more instructions, a more flexible word format, more registers, and faster execution times to work with than does the assembly language programmer. The microinstruction word length is 32 bits which enables concurrent operations to be performed in a single instruction. Microprogrammers can access 29 scratch pad registers in addition to those available to the assembly language programmer and have up to 16,384 32-bit words of memory (termed control store) in which to store microprograms. Up to three levels of nested subroutines are possible in A700 computers. The microprogrammer

works in a much faster environment than does the assembly language programmer for two reasons. One, since microinstructions have access to most of the internal parts of the computer's architecture, fewer memory fetches are required to accomplish most tasks. Two, the microinstruction execution time of 250 nanoseconds is much faster than the typical assembly instruction execution time of 1 to 2 microseconds.

These capabilities are easily taken advantage of by A700 computer users through the extensive support provided by Hewlett-Packard. Some of the more important benefits of Hewlett-Packard's microprogramming are given in the following paragraphs.

5-3. CUSTOMIZED INSTRUCTIONS

Through the use of microprogramming, the computer's assembly language instruction set can be expanded with instructions tailored for specific applications. By adding special purpose instruction sets, the general purpose computer can be uniquely adapted for a certain job and thus become very efficient at that job. Applications that may be profitably microcoded include arithmetic calculations, I/O device driver programs, and sorts and table searches.

Microprogramming is very similar to assembly language programming, although it is more powerful in many ways. Some knowledge of the internal structure of the computer is required, but once this knowledge is attained, the increased power and flexibility of microprogramming can ease the solution of many programming tasks. Microprograms are easily callable by assembly or higher level language programs.

5-4. SYSTEM SPEED

Microprogramming often-used routines will typically decrease program execution time by factors of two to ten and sometimes by as much as twenty or more. Software routines can be made to execute at the hardware speeds of the microprogram environment and the additional registers available to the microprogrammer can serve to eliminate many time-consuming memory fetches.

5-5. MEMORY SPACE AND SECURITY

By converting software routines into microprograms, space in main memory that would normally be required for time-critical routines can be freed for other uses. The

routines remain instantly callable, as opposed to routines stored in a peripheral device. Microprograms are also less accessible than conventional software which affords a higher degree of security to microcoded routines.

5-6. DEVELOPING MICROPROGRAMS

Developing microprograms is similar to developing Pascal language programs and is done with the aid of the HP paraphraser (MPARA). Since the user will not normally want to microcode all of a certain program, some analysis is required to determine which segments of the assembly language program can be most profitably converted to microcode. By substituting this section of code with a microprogrammed subroutine that is callable by the assembly or higher level program, overall execution time is reduced.

Once the microprogrammer has determined what segment to implement in microcode, the microprogram is developed as shown in Figure 5-1. The paraphraser program (in main memory) is used to assemble the source microprogram into an object program. Then, the object microprogram is loaded into writable control store (WCS).

When the microprogram is fully checked out, that user may choose to have his program reside permanently in programmable read-only memory (PROM) or in WCS where it may be altered programmatically. Implementation in ROM is accomplished by programming the PROMs with a PROM writer and installing the programmed PROMs in the computer. The mask tapes shown in figure 5-1 are required by the PROM writer and are generated by the software at the user's command. ROM-resident microprograms are permanent and do not have to be reloaded each time the computer is powered up; this implementation also prevents users from erroneously destroying the microprogram. The user who does not require such permanence for microprogram storage may execute his microcode from WCS. Microprograms used in this manner may be loaded with the WCS I/O utility routine and may be altered under program control to suit a variety of users. User-written microprograms are easily accessed by assembly or higher level programs. Once the microprogram is developed and loaded into control store, it may be called in a very similar manner to a software subroutine.

5-7. SUPPORT FOR THE MICROPROGRAMMER

Hewlett-Packard provides a comprehensive set of hardware manuals, software manuals, and training courses to make user microprogramming easy to learn and implement. For permanent implementation of microprograms, PROMs may be installed in the HP 12155A PROM Control Store Card. Up to 8,192 32-bit words of control store in the form of 8K bit PROMs may be

installed in the optional PROM Control Store Card which occupies a slot in the card cage of the computer mainframe.

The 4K writable control store (WCS) option provides a read-write control store module which can be used for the development and execution of user-supplied microprograms. Microprograms in WCS are executed at the same speed as those in the read-only control store. Each WCS module consists of a single card which plugs into the card cage, thus eliminating the need for extensive cabling or an additional power supply. A WCS card contains 4,096 32-bit locations of random-access-memory (RAM), including all necessary address and read/write circuits. WCS can be written into or read under computer control using standard input/output instructions. An I/O utility program makes it possible for FORTRAN and Pascal programs to write into or read from a WCS module using a conventional program call. A WCS module is read at full speed by way of a flat cable connecting it to the control section of the processor.

Available microprogramming software includes the paraphraser as well as a diagnostic, driver program, and I/O utility program for use with the writable control store module. These software aids operate under the Hewlett-Packard Real Time Executive (RTE) operating systems.

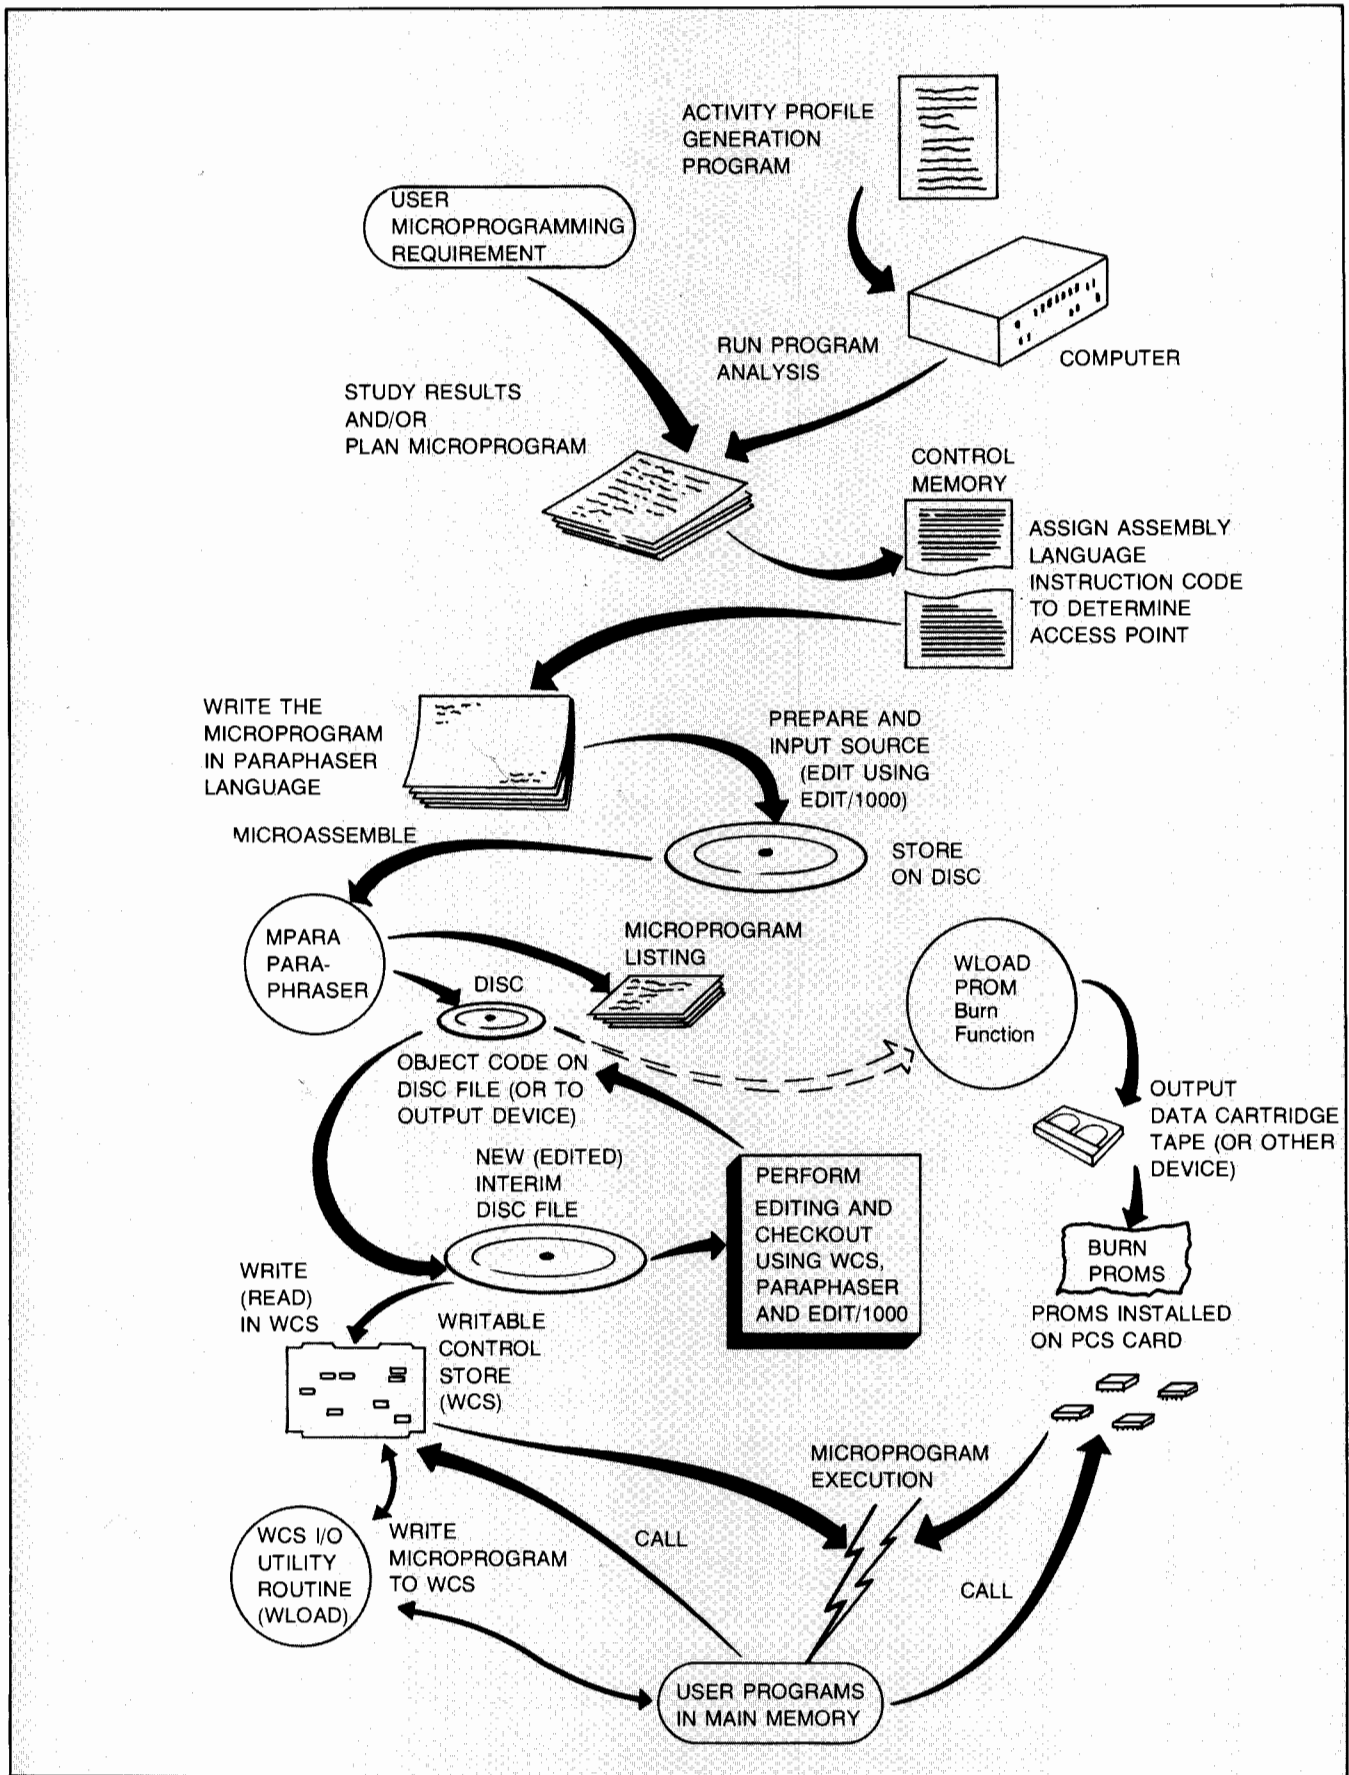
A course is offered at HP facilities in Cupertino, California for customer training. Requiring a knowledge of HP 1000 assembly language as a prerequisite, the course features in-depth coverage of microprogram development and implementation, and provides hands-on experience for the microprogrammer. The A700 microprogrammer may also take advantage of other user-written microprograms via the HP Contributed Library.

5-8. FPP MICROPROGRAMMING

Information on directly microprogramming the Floating Point Processor (FPP) card is given in the *HP 92045A Microprogramming Package Reference Manual*, part no. 92045-90001.

5-9. CONCLUSION

Microprogramming is a very powerful tool that gives the user many advantages in terms of speed, flexibility, and program security. Microprogramming does have its limitations however, and the potential user should examine very closely the extent of support provided by the computer manufacturer. Hewlett-Packard has by far sold and supported the greatest number of microprogrammable computers in the world, and provides world-wide customer support. Customer training courses and documentation have been refined from years of customer-contributed feedback and actual implementation is made easy through extensive software support packages and inexpensive hardware tools.



8200-8

Figure 5-1. Microprogramming Implementation Process

INTERRUPT SYSTEM

SECTION

VI

The vectored priority interrupt system has up to 53 distinct interrupt levels, each of which has a unique priority assignment. In the A700 computer, the interrupt priority of an I/O card is based on the card's proximity to the processor card and is independent of the card's select code. The I/O card in the slot nearest to the processor card has the highest interrupt priority. Each I/O card has higher interrupt priority than I/O cards farther from the processor card and lower priority than cards closer to the processor card. As shown in Table 6-1, the select code of an interrupt level is associated with an interrupt location in memory.

Any device can be selectively enabled or disabled under program control, thus switching the device into or out of the interrupt structure. In addition, the interrupt system is divided into types of interrupts (Table 6-1). Interrupt Type 3 can be enabled or disabled under program control using a single instruction, and interrupt Types 2 and 3 combined can be enabled or disabled using a single instruction.

Table 6-1. A700 Interrupt Assignments

SELECT CODE (OCTAL)	INTERRUPT LOCATION	ASSIGNMENT	INTERRUPT TYPE
04	00004	Power Fail Interrupt	2
05	00005	Memory Parity Interrupt	1
06	00006	Time Base Generator Interrupt	3
07	00007	Memory Protect Interrupt	2
10	00010	Unimplemented Instruction Interrupt	1
11-17	00011-00017	Reserved	
20-77	00020-00077	I/O Card Interrupts	3

6-1. POWER FAIL INTERRUPT

The computer power supply is equipped with power-sensing circuits. When primary line power fails or drops below a predetermined level while the computer is running, an interrupt to memory location 00004 is automatically generated. Memory location 00004 is intended to contain a jump-to-subroutine (JSB) instruction referencing the entry point of a user-supplied power fail

subroutine. The interrupt capability of lower-priority operations is automatically inhibited while a power fail subroutine is in process.

A minimum of five milliseconds is available between the detection of a power failure and the loss of usable power supply power to execute a power fail subroutine; the purpose of such a routine is to transfer the current state of the computer system into memory and then halt the computer. A sample power fail subroutine is given in Table 6-2. The optional battery backup module will supply enough power to preserve the contents of memory for a sustained line power outage of 30 or 60 minutes, depending on the amount of memory installed.

The user has a switch-selectable option of what action the computer will take upon restoration of primary power. When frontplane switch M is closed, the computer will execute either a loader or the Virtual Control Panel routine, depending on the setting of the Start-Up switches.

NOTE

Switch M is mounted on the processor frontplane and is not an operator control. The setting of this switch is normally determined by the System Manager prior to or during system installation.

When switch M is open, the automatic restart feature is enabled. After the self-test is executed following the return to normal power levels, an interrupt to location 00004 occurs. This time the power-down portion of the subroutine is skipped and the power-up portion begins. (Refer to Table 6-2.) Those conditions existing at the time of the power fail interrupt are restored and the computer continues the program from the point of the interruption.

Note that an auto-restart interrupt to location 00004 occurs only if that location's contents are not zero; otherwise, the system is re-booted. This is done so that if power fails and is restored during a boot, an attempt to restart a partially loaded program can be avoided. To enable this to happen the program being loaded should initially load location 00004 with zero and load the power-fail JSB instruction only when the load is otherwise complete.

If the computer memory does not contain a subroutine to service the power fail interrupt, location 00004 should contain a NOP instruction (00 octal).

At the end of a restart routine, consideration should be given to re-initializing the power-fail logic and to restoring the interrupt capability of the lower priority functions.

Table 6-2. Sample Power Fail Subroutine

LABEL	OPCODE	OPERAND	COMMENTS
PFAR DOWN	NOP		Power Fail/Auto Restart Subroutine.
	SFC	4B	Skip if interrupt was caused by power failure.
	JMP	UP	Power being restored; reset state of system.
	CLC	0B	Shut down any DMA or I/O.
	STA	SAVA	Save A-register contents.
	CCA		Set flag indicating that computer was running when power failed.
	STA	PFFLG	
	STB	SAVB	Save B-register contents.
	ERA,ALS		Transfer E-register content to A-register bit 15.
	SOC		Increment A-register if Overflow is set.
	INA		
	STA	SAVEO	Save E- and O-register contents.
	LDA	PFAR	Save contents of P-register at time of power failure.
	STA	SAVP	
	SIMP		
DEF	SAVI	Save IMAP contents.	
.			
.		Insert user-written routine to save I/O states.	
.			
SFS	4B		
JMP	*-1	Wait in case power comes back up.	
LDA	PFFLG	Was computer running when power failed?	
SZA,RSS			
HLT	4B	No, then halt.	
CLA		Yes, then reset computer Run flag to initial state.	
STA	PFFLG		
.			
.		Insert user-written routine to restore I/O devices.	
.			
LDA	SAVEO	Restore the contents of the E-register and O-register.	
CLO			
SLA,ELA			
STF	1B	Set O-register.	
LDA	SAVA	Restore A-register contents.	
LDB	SAVB	Restore B-register contents.	
STC	4B	Reset power fail logic for next power failure.	
XJMP		Cross jump to program executing at power failure.	
DEF	SAVI		
DEF	SAVP,I		
SAVEO	OCT	0	Storage for E and O.
SAVA	OCT	0	Storage for A.
SAVB	OCT	0	Storage for B.
SAVP	OCT	0	Storage for P.
PFFLG	OCT	0	Storage for Run flag.
SAVI	OCT	0	Storage for IMAP.

Note: The memory maps used must be saved and restored, as must (if used) the states of the interrupt mask register, memory protect (conditional restore), and Global Register.

6-2. PARITY ERROR INTERRUPT

Parity checking of memory is a standard feature in the A700 computer. The parity logic continuously generates correct parity for all words written into memory and monitors the parity of all words read out of memory. Parity can be programmatically set to even parity (STF 05) or cleared to odd parity (CLF 05). (Odd parity *must* be used with error correcting memory.) Correct odd parity is defined as having the total number of "1" bits in a 17-bit memory word (16 data bits plus the parity bit) equal to an odd value. If a "1" bit (or any odd number of "1" bits) is either dropped or added in the transfer process involving a standard memory array card, a Parity Error signal is generated when that word is read out of memory. In the case of data access from an error correcting memory (ECM) array card, the parity interrupt occurs only if a double-bit (non-correctable) error is detected; single-bit errors do not cause parity interrupts since the error is automatically corrected by the ECM array card.

The Parity Error signal will generate an interrupt to memory location 00005 if the parity system was previously enabled by an STC 05 instruction. Parity interrupts turn off the system. Location 00005 may contain either a JSB instruction referencing the entry point of a user-supplied parity error subroutine (included in RTE-A.1) or a JMP instruction pointing to a HLT instruction. (I/O instructions, including a HLT instruction, may not be in a trap cell). A parity error during a DMA transfer causes an interrupt to the memory location corresponding to the select code of the I/O card making the transfer if the proper bit has been set in the control word.

The memory address of the parity error will be loaded automatically into the parity register which is accessible to the user by a programmed LIA 05 or LIB 05 instruction for bits 0 through 15, and by an LIA 5,C or LIB 5,C for bits 16 through 23.

If a parity error occurs during a read of an instruction, that instruction is executed but memory writes are disabled. When a parity error occurs, it is recommended that the entire program or set of data containing the error location be reloaded.

6-3. MEMORY PROTECT INTERRUPT

The memory protect feature provides the capability of protecting selected pages of memory against alteration or entry by programmed instructions except those involving the A- and B-registers.

The memory protect logic, when enabled by an STC 07 instruction, also prohibits the execution of all I/O instructions except those referencing I/O select code 01 (the processor card status register and the overflow register). (Execution of all HLTs is prohibited.) Thus, an executive program residing in protected memory can have exclusive control of the I/O system.

The memory protect logic is disabled automatically by any interrupt and must be re-enabled by an STC 07 or XJMP instruction at the end of each interrupt subroutine.

Programming rules pertaining to the use of memory protect are as follows (assuming that an STC 07 instruction has been given):

- a. Locations 00000 and 00001 in the Execute map are the A- and B-registers and are not in protected memory. Locations 00000 and 00001 in the DATA1 and DATA2 maps are real memory locations (not the A- and B-registers) and may reside in protected memory.
- b. A user-specified 1024-word page of memory is read and/or write protected by Page Mapping Register instructions described in Section IV.
- c. Execution will be inhibited and an interrupt to location 07 will occur if any instruction addresses a location in protected memory, or if any privileged instruction is attempted (excluding those addressing select code 01 but not HLT 01). After three successive levels of indirect addressing, the logic will allow a pending interrupt.

Following a memory protect interrupt, the address of the illegal instruction will be present in the violation register. This address is made accessible to the programmer by an LIA 07 or LIB 07 instruction, which loads the address into the A- or B-register.

Note that DMA operation is not affected by memory protect.

6-4. UNIMPLEMENTED INSTRUCTION INTERRUPT

An unimplemented instruction interrupt (to memory location 00010) is requested when the CPU signals that the last instruction fetched was not recognized by itself or by any other system card. This interrupt provides a straightforward entry to software routines for the execution of instruction codes not recognized by the computer hardware. The unimplemented instruction interrupt must receive immediate service in order to recover the instruction code that caused it. For this reason, and because it is desirable to permit the use of unimplemented instructions anywhere, the unimplemented instruction interrupt is never inhibited.

6-5. TIME BASE GENERATOR INTERRUPT

A time base generator interrupt request is made when the CPU signals that its internal clock divider chain has rolled over. The clock divider is set to roll over at 10-

millisecond intervals for maintaining a real-time clock. The interrupt occurs through location 00006 and can be masked (inhibited) by using bit 1 of the interrupt mask register. (The interrupt mask register allows interrupts from the TBG and the I/O cards to be selectively masked. For details on the interrupt mask register, refer to the *HP 1000 L-Series Computer I/O Interfacing Guide*, part no. 02103-90005.) The TBG can be turned on by an STC 06 instruction, and turned off by a CLC 06 or CLC 00 instruction.

6-6. INPUT/OUTPUT INTERRUPT

Interrupt locations 20 through 77 (octal) are reserved for I/O devices. In a typical I/O operation, the computer issues a programmed command such as Set Control/Clear Flag (STC,C) to one or more external devices to initiate an input (read) or an output (write) operation, via either programmed I/O or DMA. While the I/O card is in the process of transferring data, the computer may be either running a program or looping, waiting for a flag to get set. Upon completion of the read or write operation, the interface flag is set. If the corresponding control bit is set, the interface will interrupt. Its request will be passed through a priority network so that only the highest priority interrupting device will receive service. The computer will acknowledge the interrupt and the highest priority device will receive service when the current instruction has finished executing, except under the following circumstances:

- a. Interrupt system disabled or interface card interrupt disabled (or masked).
- b. JMP indirect or JSB indirect instruction not sufficiently executed. These instructions inhibit all interrupts except power fail or memory protect until the succeeding instruction is executed. After three successive levels of indirect addressing, the logic will allow a pending I/O interrupt.
- c. A DMA (direct memory access) data transfer is in process.
- d. Current instruction is any I/O instruction. The interrupt in this case must wait until the succeeding instruction is executed.

After an interface card has been issued a Set Control (STC instruction) and its flag bit becomes set, all interrupt requests from lower-priority devices are inhibited until this flag bit is cleared by a Clear Flag (CLF) instruction, or until control is cleared by a Clear Control (CLC) instruction. A service subroutine in process for any device can be interrupted only by a higher-priority device; then, after the higher-priority device is serviced, the interrupted service subroutine can continue. In this way it is possible for several service subroutines to be in the interrupt state at one time; each of these service subroutines will be allowed to continue after the higher-

priority device is serviced. All such service subroutines normally end with a JMP indirect or XJMP instruction to return the computer to the point of interrupt.

Note that interrupt trap cells must contain a JMP or JSB instruction because maps change on interrupt.

6-7. INTERRUPT PRIORITY

The interrupt servicing priority among the system interrupts is as follows:

- a. Parity error (select code 5).
- b. Unimplemented instruction (select code 10).
- c. Memory protect (select code 7).
- d. Power fail (select code 4).
- e. Time base generator (select code 6).
- f. I/O interrupts (select codes 20 through 77).

6-8. CENTRAL INTERRUPT REGISTER

Each time an interrupt occurs, the address of the interrupt location is stored in the central interrupt register. The contents of this register are accessible at any time by executing an LIA 04 or LIB 04 instruction. This loads the address of the most recent interrupt into the A- or B-register.

6-9. PROCESSOR STATUS REGISTER

The processor status register is two registers: one for input and one for output. The input register shows the status of the frontplane BOOT SEL switches and is read into the upper eight bits of the A- or B-register by an LIA/B 01 instruction. The switch, bit, and function relationships are as follows:

<u>SWITCH</u>	<u>BIT</u>	<u>MEANING</u>
S1	8	Boot select
S2	9	Boot select
S3	10	Boot select
S4	11	Boot select
S5	12	VCP program select
S6	13	Not used
S7	14	Not used
S8	15	Auto-restart enabled (1)/ disabled (0)

The output register drives the frontplane LEDs. The output of the lower eight bits of the A- or B-register are sent to the LEDs by an OTA/B 01 instruction. A logic 1 in either register lights the corresponding LED.

6-10 INTERRUPT TYPE CONTROL

I/O address 00 is the master control address for Type 3 interrupts (TBG and I/O cards). An STF 00 instruction enables Type 3 interrupts and a CLF 00 disables Type 3 interrupts. (Type 3 interrupts are disabled when power is initially applied.) I/O address 04 is the master control address for Type 2 interrupts (power fail and memory protect) and Type 3 interrupts combined. An STC 04 instruction enables Type 2 and 3 interrupts, and a CLC 04 disables Type 2 and 3 interrupts.

6-11. INSTRUCTION SUMMARY

Table 6-3 is a summary of instructions for select codes 00 through 07. For a summary of instructions used with the I/O cards, refer to an I/O card reference manual.

The Type 2 and 3 interrupt mask from I/O address 04 is a different Type-3 mask than the Type-3 mask at I/O address 00. If either of these two masks are set, Type 3 interrupts will be disabled. In addition to these two interrupt masks, the Time Base Generator flag interrupt can also be masked by bit 0 of the Interrupt Mask Register. If any of these three masks are set then the TBG flag interrupt will be disabled.

Table 6-3. Instructions for Select Codes 00 through 07

INSTRUCTION	FUNCTION	INSTRUCTION	FUNCTION
STC 0	NOP	STC 4	Enable Type 2 and 3 interrupts
CLC 0	System reset	CLC 4	Disable Type 2 and 3 interrupts
STF 0	Enable Type 3 interrupts	STF 4	NOP
CLF 0	Disable Type 3 interrupts	CLF 4	NOP
SFS 0	Skip if Type 3 interrupts enabled	SFS 4	Skip if power is stable
SFC 0	Skip if Type 3 interrupts disabled	SFC 4	Skip if power going down
LI* 0	Load from interrupt mask register	LI* 4	Load from central interrupt register
MI* 0	NOP	MI* 4	NOP
OT* 0	Output to interrupt mask register	OT* 4	Output to central interrupt register
STC 1	NOP	STC 5	Enable parity error interrupts
CLC 1	NOP	CLC 5	Disable parity error interrupts
STF 1	Same as Set Overflow (STO)	STF 5	Set parity sense to even parity
CLF 1	Same as Clear Overflow (CLO)	CLF 5	Clear parity sense to odd parity
SFS 1	Same as Skip if Overflow Set (SOS)	SFS 5	Skip if parity sense is even
SFC 1	Same as Skip if Overflow Clear (SOC)	SFC 5	Skip if parity sense is odd
LI* 1	Load from processor status register	LI* 5	Load from parity register (bits 0-15)
MI* 1	Merge from processor status register	LI* 5,C	Load from parity register (bits 16-23)
OT* 1	Output to processor status register	MI* 5	NOP
STC 2	Enable break feature	OT* 5	NOP
CLC 2	NOP	STC 6	Turn on time base generator
STF 2	Disable Global Register	CLC 6	Turn off time base generator
CLF 2	Enable Global Register	STF 6	Set time base generator flag
SFS 2	Skip if Global Register disabled	CLF 6	Clear time base generator flag
SFC 2	Skip if Global Register enabled	SFS 6	Skip if time base generator flag set
LI* 2	Load from Global Register	SFC 6	Skip if time base generator flag clear
MI* 2	NOP	LI* 6	NOP
OT* 2	Output to Global Register (Note 1)	MI* 6	NOP
STC 3	NOP	OT* 6	NOP
CLC 3	NOP	STC 7	Enable memory protect
STF 3	NOP	CLC 7	NOP
CLF 3	NOP	STF 7	NOP
SFS 3	NOP	CLF 7	NOP
SFC 3	NOP	SFS 7	NOP
LI* 3	Load from P SAVE	SFC 7	NOP
MI* 3	NOP	LI* 7	Load from violation register
OT* 3	Output to P SAVE	MI* 7	NOP
LI* 3,C	Load from ROM P	OT* 7	NOP
OT* 3,C	Output to ROM P		

* = A or B.

Note 1. An OTA/B 2 with A/B equal to one through seven establishes a diagnose mode; refer to paragraph 7-22 for details.

The purpose of the input/output system is to transfer data between the computer and external devices. As shown in figure 7-1, data can be transferred either by a direct memory access (DMA) feature or through the A- or B-register in the CPU (non-DMA). Each L-series I/O card has DMA logic and DMA is normally used for most I/O data transfers. Once the DMA logic has been initialized, no programming is involved and the transfer occurs in two distinct steps as follows:

- a. Between the external device and its I/O interface card in the computer;
- b. Between the I/O card and memory via the backplane data bus. This two-step process also applies to a DMA output transfer except in reverse order.

As mentioned above, data may be transferred under program control without using the DMA feature. This type of transfer allows the computer to manipulate the data during the transfer process. A non-DMA input transfer is a three-step process as follows:

- a. Between the external device and its I/O card;
- b. Between the I/O card and the A- or B-register via the data bus and the processor card; and
- c. Between the A- or B-register and memory via the processor and the data bus.

Note that in the DMA transfer the processor card is bypassed. Since a DMA transfer eliminates programmed loading and storing via the accumulators, the time involved is very short. Further information on the DMA feature is given in paragraph 7-9.

7-1. INPUT/OUTPUT ADDRESSING

As shown in Figure 7-2, an external device is connected by cable directly to an interface card located in the computer mainframe. The interface card, in turn, plugs into one of the input/output slots, each of which is assigned a fixed interrupt priority. Note, however, that the select code of the L-Series interface cards is independent of the priority. The computer communicates with a specific device on the basis of its select code which is set by switches on the interface card.

Figure 7-2 shows an interface card inserted in the I/O slot having the highest priority. If it is decided that the associated device should have lower priority, its interface card and cable may simply be exchanged with those occupying some other I/O slot. This will change the priority but not the I/O address (select code). Due to priority

chaining, there can be no vacant slots from the highest priority slot to the lowest priority slot used. Only select codes 20 through 77 (octal) are available for input/output cards; the lower select codes (00 through 17) are reserved for other features.

7-2. INPUT/OUTPUT PRIORITY

The plug-in card slots of the A700 computers are numbered 1 through 20. Generally, slots 1 through 4 are used for the memory and processor cards and the remaining slots are available for I/O cards, with slot 5 having the highest I/O interrupt priority. An I/O channel consists of an I/O device (or devices) and its I/O card, and is assigned the number of the card slot.

When an input/output device is ready to be serviced, it causes its interface card to request an interrupt so that the computer will interrupt the current program and service the device. Since many device interface cards will be requesting service at random times, it is necessary to establish an orderly sequence for granting interrupts. Also, it is desirable that high-speed devices should not have to wait for low-speed device transfers. Both of these requirements are met by a series-linked priority structure illustrated by Figure 7-3. The bold line, representing a priority enabling signal, is routed in series through each card capable of causing an interrupt. The card cannot interrupt unless this enabling signal is present at its input.

Each device (or other interrupt function) can break the enabling line when it requests an interrupt. If two devices simultaneously request an interrupt, the device with the highest priority will be the first one that can interrupt because it has broken the enable line for the lower-priority device. The other device cannot begin its service routine until the first device is finished. However, a still higher-priority device (one interfaced through a lower-numbered slot) may interrupt the service routine of the first device. Figure 7-4 illustrates a hypothetical case in which several devices request service by interrupting a CPU program. Both simultaneous and time-separated interrupt requests are considered.

Assume that the computer is running a CPU program when an interrupt from I/O channel 7 occurs (at reference time t_1), and that the card in slot 7 is assigned select code 22. With the I/O card supplying the select code as the memory address, a JSB instruction in the interrupt location for select code 22 causes a program jump to the service routine for the channel-7 device (select code 22). The JSB instruction automatically saves the return address (in a location which the programmer must reserve in his

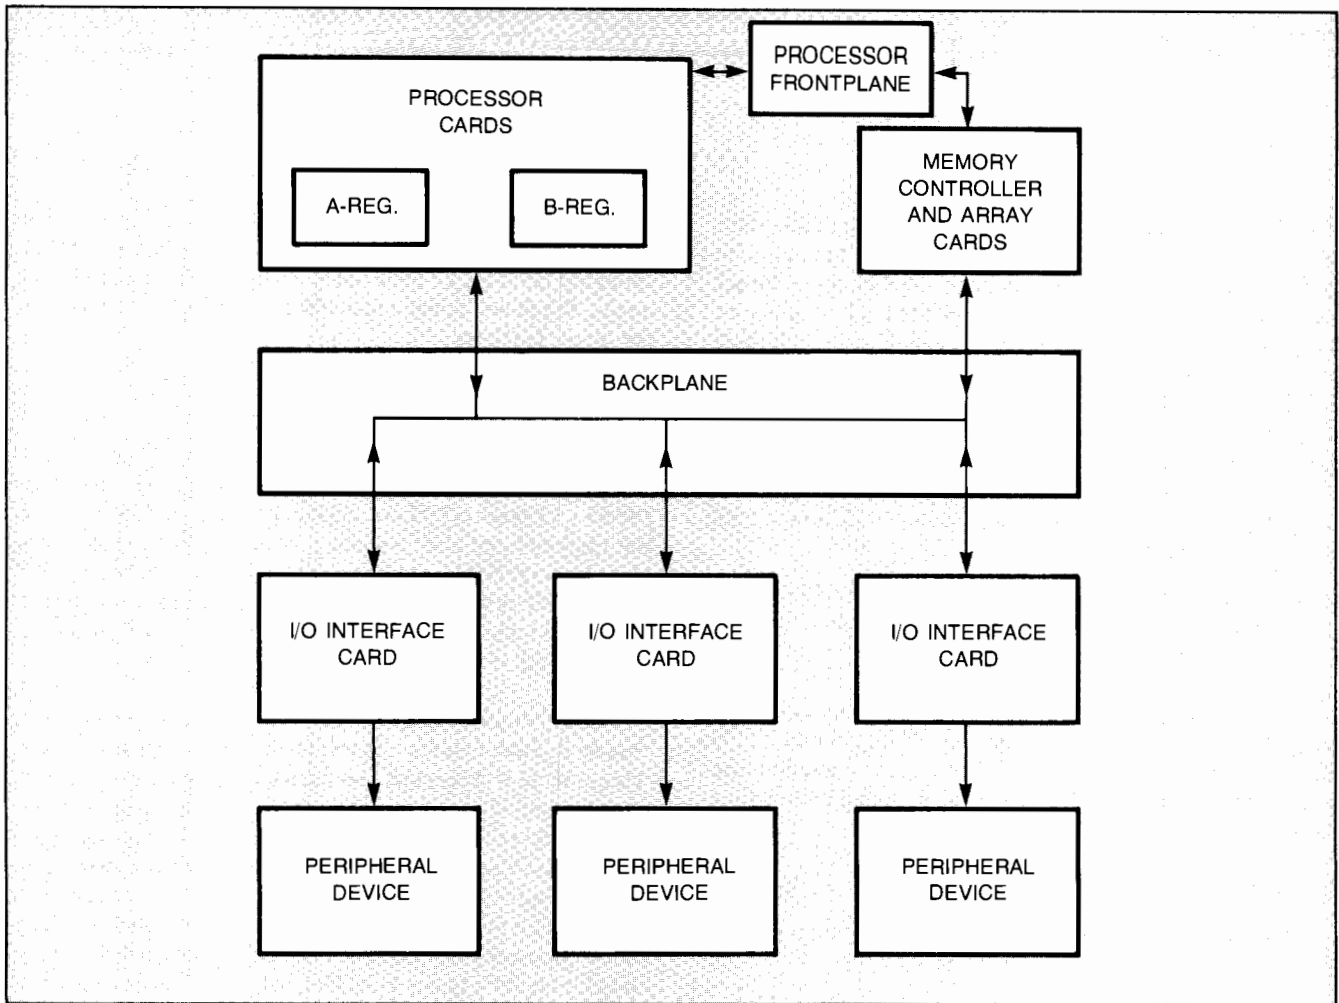



Figure 7-1. Input/Output System



8200-62

Figure 7-2. I/O Priority Assignments

routine) for a later return to the CPU program.

The routine for channel 7 (select code 22) 7 is still in progress when several other devices request service (set flag). First, channels 8 and 9 request service simultaneously at time t₂; however, since neither one has priority over channel 7, their flags are ignored and channel 7 continues transfer. But at t₃, a higher priority device on channel 5 requests service. This request interrupts the channel 7 transfer and causes the channel 5 transfer to begin. The JSB instruction saves the return address for return to the channel 7 routine.

During the channel 5 transfer, the channel 6 flag is set (t₄). Since it has lower priority than channel 5, channel 6 must wait until the end of the channel 5 routine. And since the channel 5 routine, when it ends, contains a return address to the channel 7 routine, program control temporarily returns to channel 7 (even though the waiting channel 6 has higher priority). The JMP,I instruction used for the return inhibits all interrupts until fully executed. At the end of this short interval, the channel 6 interrupt request is granted.

When channel 6 has finished its routine, control is returned to channel 7, which at last has sufficient priority to complete its routine. Since channel 7 has been saving a return address in the main CPU program, it returns control to this point.

The two waiting interrupt requests from channels 8 and 9 are now enabled. Channel 8 has the higher priority and goes first. At the end of the channel 8 routine control is temporarily returned to the CPU program. Then the lowest priority channel (channel 9) interrupts and completes its transfer. Finally, control is returned to the CPU program, which resumes processing.

7-3. INTERFACE ELEMENTS

The interface card provides the communication link between the computer and one or more external devices. The interface card includes several basic elements which either the computer or the device can control in order to effect the necessary communication. These basic elements are the Global Register, control bits, flag bits, data buffer

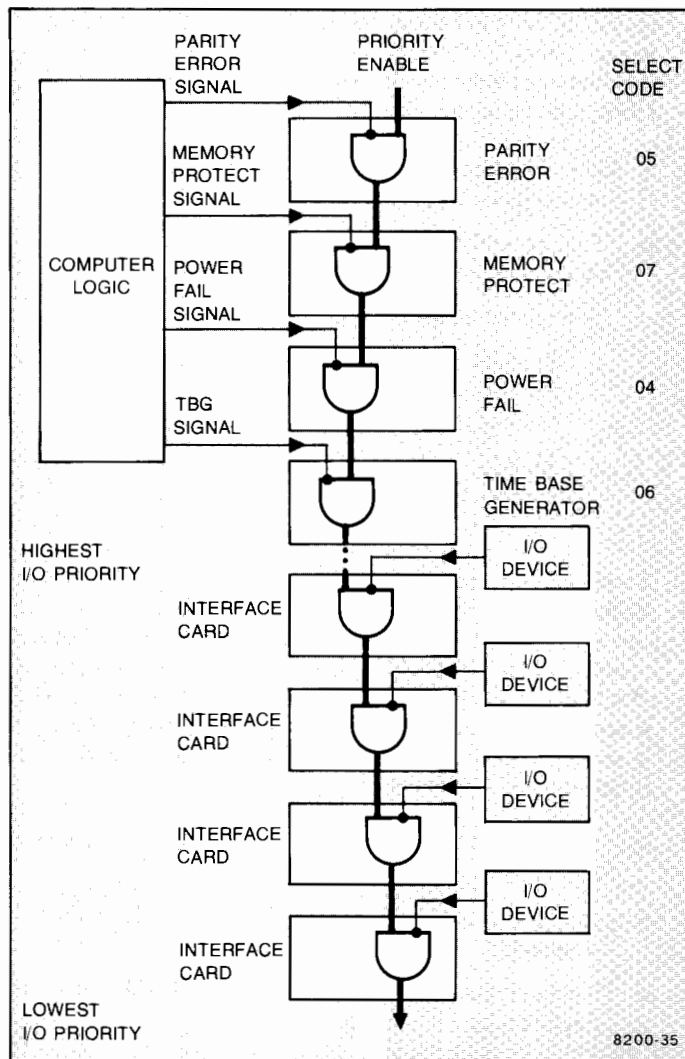


Figure 7-3. Priority Linkage (Simplified)

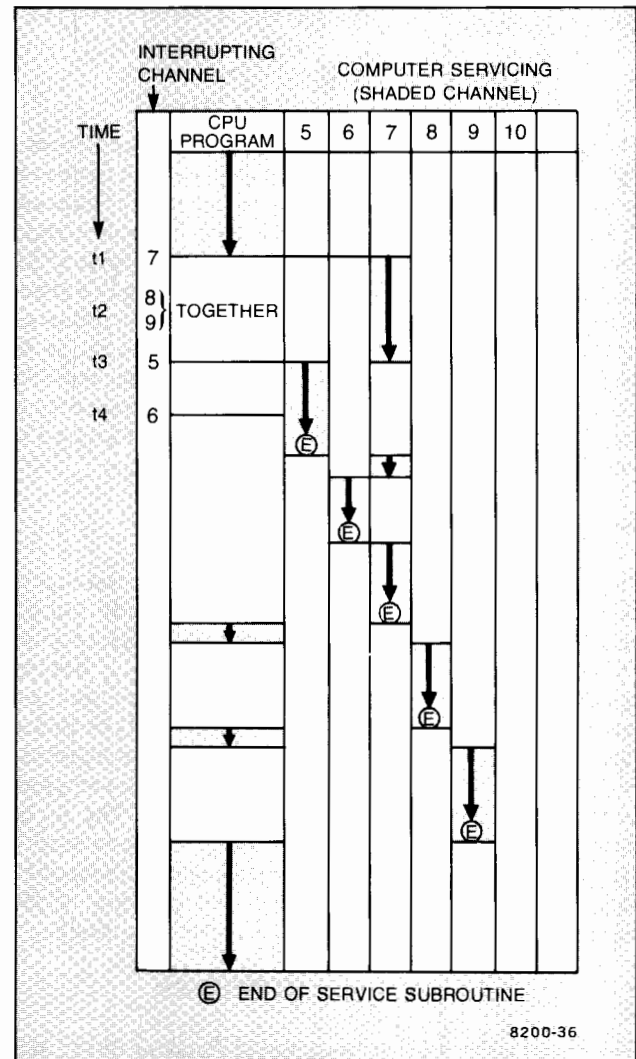


Figure 7-4. Interrupt Sequence

register, and control register. Other registers, associated only with DMA, are discussed in paragraph 7-9. The control and flag bits and the data buffer and control registers of an interface card can be addressed directly when the card's select code is in the Global Register (GR) and the GR is enabled. Refer to the interface card reference manuals for specific information on the data and control registers.

7-4. GLOBAL REGISTER

In the A-Series computers, the select code that is in the Global Register specifies which I/O card is enabled to execute I/O instructions. The Global Register (GR) is a register on each I/O card that can be loaded with the select code of any one of the I/O cards. (At any given time, the GR on all I/O cards is loaded with the same select code.) When the GR is enabled, an I/O instruction is executed only by the I/O card whose select code matches the select code in its GR. Also, the GR allows other registers on the selected I/O card to be accessed programmatically by I/O instructions. The Global Register on all I/O cards may be simultaneously loaded with an OTA/B 02 instruction, enabled with a CLF 02 instruction, and disabled with an STF 02 instruction.

7-5. CONTROL BITS

The control bits on an interface card are used to turn on a specific I/O function. In addition, a control bit must be set to allow the corresponding flag bit to interrupt. There are three control bits associated with each I/O select code: control 20, 21, and 30. Control 30 is the only control bit that can be accessed with or without the Global Register being enabled. When control 30 is set it generates an action command, allowing one word or character to be read or written. Control 20 and 21 can only be accessed when the Global Register is enabled. When control 20 is set it turns on DMA self-configuration. The setting of control 21 enables DMA transfers.

7-6. FLAG BITS

The flag bits (when set) are used primarily to interrupt or to signal completion of a task. Flag 30, the only flag bit accessible without using the Global Register, signals either one data element has been transferred or that an interrupting condition has been detected. There are three other flags, all of which must be accessed with the Global Register enabled. Flag 20 signals DMA self-configuring transfer complete; flag 21 signals DMA transfer complete; and flag 22 signals parity error during DMA. The device cannot clear the flag bit. If the corresponding control bit is set, priority is high, and the interrupt system is enabled, then setting the flag bit will cause an interrupt to the location corresponding to the I/O card's select code.

7-7. DATA BUFFER REGISTER

The data buffer register (designated Register 30) is used for the intermediate storage of data during an I/O transfer. Typically, the data capacity is 16 bits.

7-8. CONTROL REGISTER

The control register (designated Register 31) enables a general purpose interface card to be configured for compatibility with a specific I/O device or to be programmed for particular modes of operation. The control register must be programmatically set up for each particular application. Refer to the interface card manuals for specific information on the control register.

7-9. DIRECT MEMORY ACCESS

The direct memory access (DMA) capability of each L-Series interface card provides a direct data path between memory and a peripheral device, making it practical to use DMA for most data transfers. The use of DMA to perform I/O data transfers reduces the number of interrupts from one per byte or word to one per complete DMA block transfer. (Maximum DMA block size is 65,536 bytes.)

The maximum DMA transfer rate is 4.0 million bytes per second; this is also the combined limit for DMA transfers by two or more I/O cards. Except when the DMA feature is operating at full bandwidth, the central processor can interleave memory cycles with the DMA operation. The DMA feature is provided by the following elements:

- a. The common backplane that links the processor, memory, and I/O cards;
- b. The capability of the I/O cards to execute I/O instructions; and
- c. The Global Register which:
 1. Enables only the I/O card whose select code is in the Global Register to execute I/O instructions, freeing the address bits of the I/O instruction; and
 2. Enables the I/O-instruction address bits to be used to access registers on the I/O card specified by the Global Register.

Each I/O card has four registers associated with DMA. Three of them must be loaded with control words that specify the DMA operation. The fourth register is used for a special type of DMA operation called self-configured DMA which is discussed later. All of these registers can be accessed only when the select code of the desired I/O card is in the Global Register. The DMA registers and their functions are as follows:

- a. Register 20, DMA Self-Configuration Address Register;
- b. Register 21 (for Control Word 1), DMA Control Register;
- c. Register 22 (for Control Word 2), DMA Address Register; and
- d. Register 23 (for Control Word 3), Word/Byte Count Register.

7-10. CONTROL WORD 1

Control Word 1 (CW1) must be loaded into Register 21 of the desired I/O card as part of the DMA initialization process. The general definitions of the bits in Control Word 1 are given in figure 7-5. Note that the requirements of individual I/O cards may vary slightly from the general definitions and that it is necessary to refer to the I/O card reference manuals for specific programming information.

7-11. CONTROL WORD 2

Control Word 2 (CW2) loads into Register 22 the address of the first memory location to be read from or stored into when the DMA operation is initiated. The most significant bit, bit 15, is not used by the DMA control logic; when CW2 is read for status, bit 15 is the complement of bit 7 in CW1 (figure 7-5).

7-12. CONTROL WORD 3

Control Word 3 (CW3) loads into Register 23 the two's-complement number of data elements to be transferred by DMA. Data elements may be either words or bytes as specified by bit 13 of CW1 (figure 7-5). The end of a DMA data transfer is indicated by the transition from -1 to 0 of the value in Register 23 (the Word/Byte Count Register); this causes the I/O card to generate a completion interrupt. (A DMA transfer can also be terminated in other ways as described in the interface card manuals.)

7-13. DMA TRANSFER INITIALIZATION

A DMA data transfer is started by:

- a. Loading the Global Register with the select code of the desired I/O card;
- b. Loading the three DMA registers: DMA control into Register 21, DMA address into Register 22, and word/byte count into Register 23;
- c. Loading the control register (Register 31) of the I/O card (described in the individual interface card reference manuals); and
- d. Issuing an STC instruction to Register 21 (DMA Control Register).

A typical programming sequence to configure the DMA logic for a DMA transfer is as follows:

LDA SC	Load select code
OTA 2, C	Set up Global Register
CLC 21B	
LDA CW1	
OTA 21B	Output DMA control word
LDA CW2	
OTA 22B	Output DMA starting address
LDA CW3	
OTA 23B	Output DMA word/byte count

LDA CNTL	
OTA 31B	Output I/O card control word
STC 21B, C	Start DMA and device
<continue any other processing>	

7-14. SELF-CONFIGURED DMA

Each I/O card also has logic that can automatically load the DMA registers discussed previously with the DMA control words from sequential locations in memory. This process is performed by using the I/O card's Register 20, the Self-Configuration Register. The DMA self-configuration feature is initialized by setting the value of Register 20 to the memory address of a list of DMA "triplets" or "quadruplets".

A triplet is of the form: DMA control word, DMA transfer address, and word/byte count. The triplet words are the words to be loaded into Registers 21, 22, and 23, respectively. A quadruplet is of the form: DMA control word, I/O-card control word, transfer address, and word/byte count. Bit 8 of the DMA control word (Control Word 1) determines whether a triplet or quadruplet is loaded. (A quadruplet is used only when the I/O-card control word must be changed; refer to the interface card manuals for detailed information.) As each register is loaded, the contents of Register 20 are incremented, leaving it pointing to the memory location to be loaded into the next register.

DMA self-configuration can be chained to enable consecutive DMA transfers via the same I/O card with a minimum of interrupts. If bit 15 of Control Word 1 in a triplet (or quadruplet) is a logic 1, the DMA registers will be loaded with the next triplet or quadruplet in memory (as pointed to by Register 20) upon completion of the current DMA block transfer. When bit 15 (and bit 11) is a logic 0, the current DMA block transfer is followed by a completion interrupt.

7-15. DMA DATA TRANSFER

Figure 7-6 illustrates, in general, the sequence of operations for a DMA input data transfer (the minor differences for an output transfer are explained in text). Note that the Global Register has been enabled and loaded with the I/O card's select code.

The initialization routine sets up the DMA control registers on the I/O card (1) and issues the start command (STC 21,C) to the DMA Control Bit (Control 21). (If the operation is an output, the I/O card buffer is also loaded at this time.) The DMA logic is now turned on and the computer program continues with other instructions.

Setting the DMA Control bit (2) causes the I/O card to send a Start signal (with a data word if it is an output transfer) to the external device (3). The device goes through a read or write cycle and returns a Done signal (with a data word if it is an input transfer). The Done

15	14	13	12	11	10	9	8	7	6	5	4	0
CONT	DVCMD	BYTE	RES	CINT	REM	FOUR	AUTO	IN	Various		ADDR EXT BUS	

CONT (Continue), bit 15.

Bit 15 = 1: Enable a DMA re-configuration upon completion of a self-configured DMA transfer.

Bit 15 = 0: Stop DMA after current transfer.

DVCMD (Device Command), bit 14.

Bit 14 = 1: Issue a Device Command signal for each data element transferred.

Bit 14 = 0: No Device Command signal issued.

BYTE (Byte/word transfer), bit 13.

Bit 13 = 1: Conduct DMA transfer in byte mode.

Bit 13 = 0: Conduct DMA transfer in word mode.

RES (Residue), bit 12.

Bit 12 = 1: Write word/byte count back into memory.

Bit 12 = 0: Word/byte count is not written.

CINT (Completion Interrupt), bit 11.

Bit 11 = 1: Inhibit DMA completion interrupt.

Bit 11 = 0: Request completion interrupt when word/byte count goes from -1 to 0 and bit 15 equals 0.

REM (Remote), bit 10.

Bit 10 = 1: Enable remote (non-standard) memory for DMA transfer.

Bit 10 = 0: Remote memory not enabled.

FOUR (Fetch four control words), bit 9.

Bit 9 = 1: Causes DMA self-configuration to fetch four control words; i.e., three DMA control words and one I/O card control word.

Bit 9 = 0: Fetch three control words for DMA self-configuration.

AUTO (Automatic), bit 8. This bit is read only during self-configured DMA.

Bit 8 = 1: Initiate first data transfer once DMA is configured to output, without waiting for an SRQ. For input transfers, enable a Device Command signal after the last data element is transferred.

Bit 8 = 0: For output transfers, wait for a Service Request (SRQ) signal before performing the first transfer. For input transfers, the last data element is not followed by a Device Command.

IN (Transfer In), bit 7.

Bit 7 = 1: Perform DMA transfer from I/O device to memory.

Bit 7 = 0: Perform DMA transfer from memory to I/O device.

Various, bits 5 and 6, User definable.

ADDR EXT BUS, bits 4-0

These five bits allow DMA accesses to physical memory by referencing one map set of 32 registers each.

8200-53

Figure 7-5. General Bit Definitions for Control Word 1

signal (4) requests the DMA logic (5) to transfer a word into (or out of) memory (6). The process now loops back to step 3 to transfer the next word.

After the specified number of words has been transferred, the DMA logic generates a completion interrupt (7). The program control is now forced to a completion routine (8), the content of which is the programmer's responsibility.

For more detailed information on DMA, refer to the I/O interfacing guide, part no. 02103-90005.

7-16. NON-DMA DATA TRANSFER

The following paragraphs describe how data is transferred between memory and input/output devices without using DMA. The sequences presented are simplified in order to present an overall view without the involvement of software operating systems or device drivers.

7-17. INPUT DATA TRANSFER (INTERRUPT METHOD)

Figure 7-7 illustrates the sequence of events required to input data using the interrupt method. Note that some operations are under control of the computer program (programmer's responsibility) and some of the operations are automatic. Note also that the Global Register has been loaded and enabled and the I/O card's control register has been loaded.

The operations begin (1) with the programmed instruction STC 30,C which sets the Control bit (Control 30) and clears the Flag bit (Flag 30) on the I/O card. Since the next few operations are under control of the hardware, the computer program may continue the execution of other instructions. Setting the Control bit causes the card to output a Start signal (2) to the device, which reads out a data character and asserts the Done signal (3).

The device Done signal sets the Flag bit, which in turn generates an interrupt (4) provided that the interrupt conditions are met; i.e., the interrupt system must be on (STF 00 previously given), no higher priority interrupt is pending, and the Control bit is set (done in step 1).

The interrupt causes the current computer program to be suspended and control is transferred to a service subroutine (5). It is the programmer's responsibility to provide the linkage between the interrupt location (which agrees with the select code) and the service subroutine. It is also the programmer's responsibility to include in his service subroutine the instructions for processing the data (loading into an accumulator, manipulating if necessary, and storing into memory).

The subroutine may then issue further STC 30,C instructions to transfer additional data characters. One of the final instructions in the service subroutine must be CLC 30,C. This step (6) restores the interrupt capability to lower priority devices and returns the I/O card to its static

"reset" condition (Control clear and Flag clear). This condition is initially established by the computer at power turn-on and it is the programmer's responsibility to return the I/O card to the same condition on the completion of each data transfer operation. At the end of the subroutine, control is returned to the interrupted program via previously established linkages.

7-18. OUTPUT DATA TRANSFER (INTERRUPT METHOD)

Figure 7-8 illustrates the sequence of events required to output data using the interrupt method. Again note the distinction between programmed and automatic operations. Note also that the Global Register has been loaded and enabled and that the I/O card's control register has been loaded. It is assumed that the data to be transferred has been loaded into the A-register and is in a form suitable for output.

The output operation begins with a programmed instruction (OTA 30) to transfer the contents of the A-register to the I/O card buffer (1). This is followed (2) by the instruction STC 30,C which sets the Control bit (Control 30) and clears the Flag bit (Flag 30) on the I/O card. Since the next few operations are under control of the hardware, the computer program may continue the execution of other instructions. Setting the Control bit causes the card to output the buffered data and a Start signal (3) to the device, which writes (e.g., records, stores, etc.) the data character and asserts the Done signal (4).

The device Done signal sets the card's Flag bit, which in turn generates an interrupt (5) provided that the interrupt system is on, priority is high, and the Control bit is set (done in step 2). The interrupt causes the current computer program to be suspended and control is transferred to a service subroutine (6). It is the programmer's responsibility to provide the linkage between the interrupt location (which agrees with the select code) and the service subroutine. The detailed contents of the subroutine are also the programmer's responsibility and the contents will vary with the type of device.

The subroutine may then output further data to the I/O card and reissue the STC 30,C command for additional data character transfers. One of the final instructions in the service subroutine must be a clear control (CLC 30,C). This step (7) allows lower priority devices to interrupt and restores the I/O card to its static "reset" condition (Control clear and Flag clear). At the end of the subroutine, control is returned to the interrupted program via the previously established linkages.

7-19. NON-INTERRUPT DATA TRANSFER

It is also possible to transfer data without using the interrupt system. This involves a "wait-for-flag" method in which the computer commands the device to operate and then waits for the completion response. In using this method to transfer data, computer time is relatively

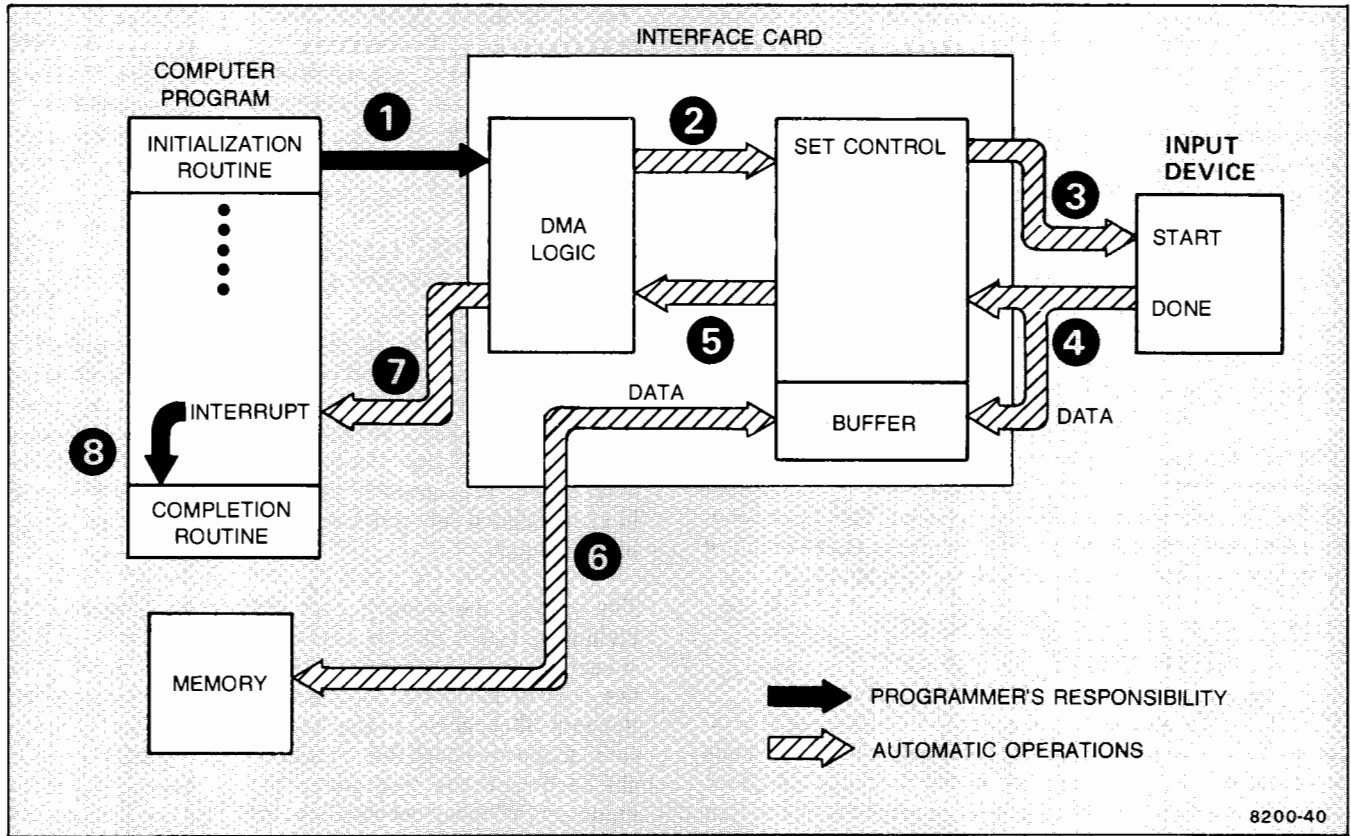


Figure 7-6. DMA Input Data Transfer

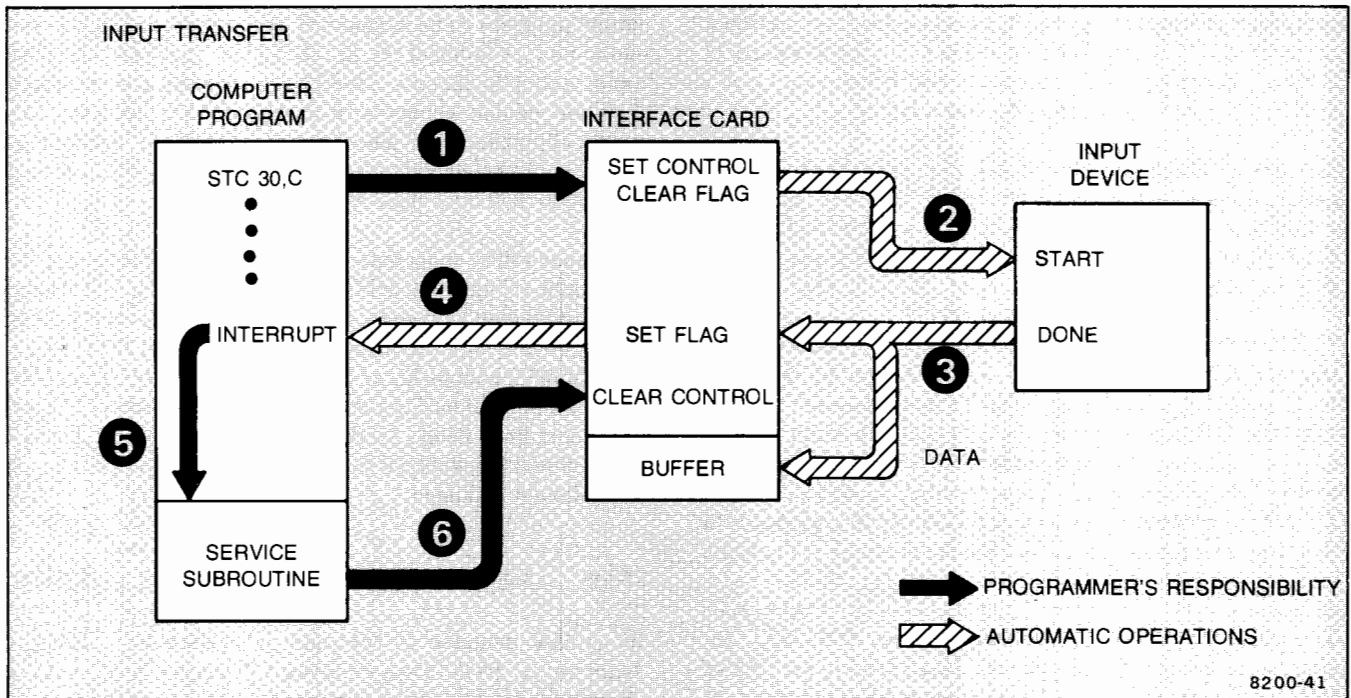


Figure 7-7. Input Data Transfer (Interrupt Method)

unimportant. It is assumed that the interrupt system is turned off (STF 00 not previously given). It is also assumed that the Global Register has been loaded and enabled and that the I/O card's control register has been loaded. As shown in table 7-1, the programming is very simple; each of the routines will transfer one word or character of data.

7-20. INPUT. As described in paragraph 7-17, an STC 30,C instruction begins the operation by commanding the device to read one word or character. The computer then goes into a waiting loop, repeatedly checking the status of the Flag bit (Flag 30). If the Flag bit is not set, the JMP *-1 instruction causes a jump back to the SFS instruction. (The *-1 operand is assembler notation for "this location minus one.") When the Flag bit is set, the skip condition for SFS is met and the JMP instruction is skipped. The computer thus exits from the waiting loop and the LIA 30 instruction loads the device input data into the A-register.

7-21. OUTPUT. The first step, which transfers the data to the I/O card buffer, is the OTA 30 instruction. Then STC 30,C commands the device to operate and accept the data. The computer then goes into a waiting loop as described in the preceding paragraph. When the Flag bit becomes set, indicating that the device has accepted the output data, the computer exits from the loop. (The final NOP is for illustration purposes only.)

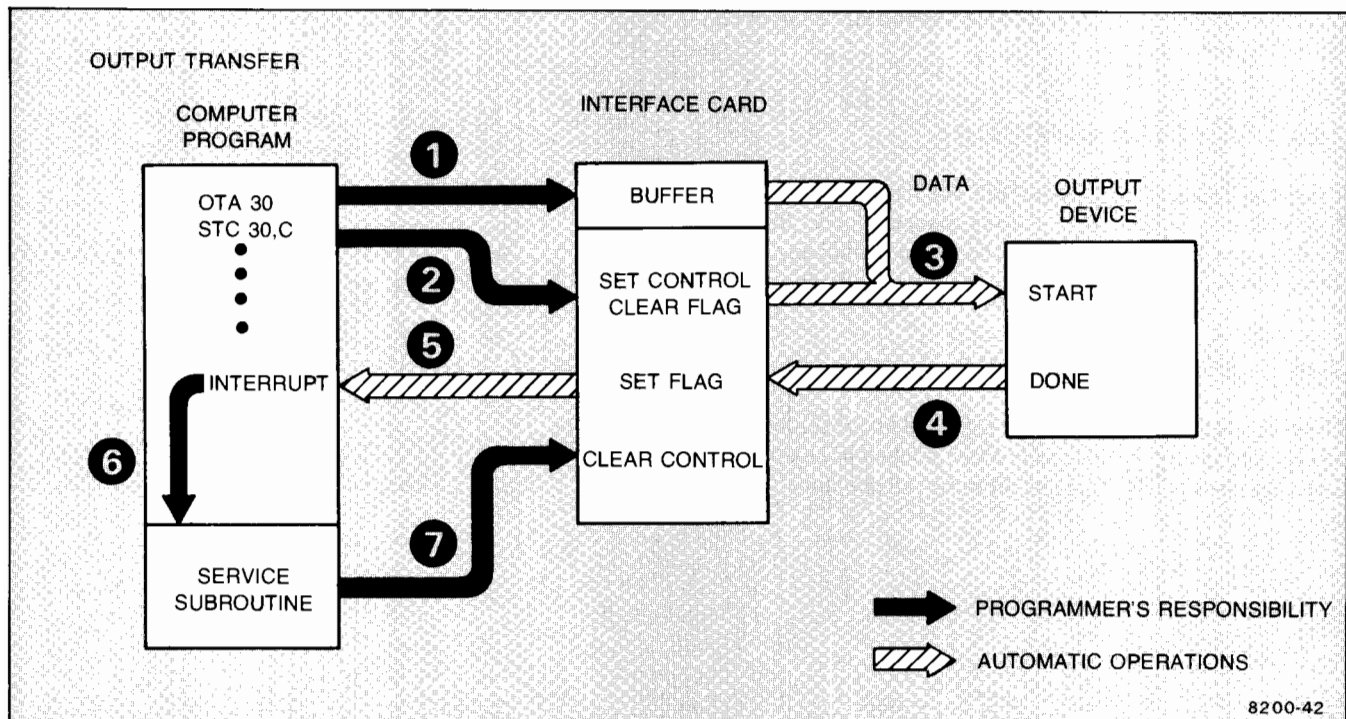
7-22. DIAGNOSE MODES

A diagnose mode allows the I/O cards to be accessed for

Table 7-1. Noninterrupt Transfer Routines

INSTRUCTIONS	COMMENTS
INPUT ROUTINE	
STC 30,C	Start device
SFS 30	Is input ready?
JMP *-1	No, repeat previous instruction
LIA 30	Yes, load input into A-register
OUTPUT ROUTINE	
OTB 30	Output data to I/O card's data register
STC 30,C	Start device
SFS 30	Has device accepted the data?
JMP *-1	No, repeat previous instruction
NOP	Yes, proceed

diagnostic or test purposes. A diagnose mode is established when an OTA/B 2 instruction (output to the Global Register) is executed with the A- or B-register value equal to one through seven. (The diagnose mode is terminated when an OTA/B 2 instruction is executed with the A- or B-register equal to zero.) When establishing a diagnose mode the current contents of the Global Register (GR) is not altered. The diagnose mode can be on an individual I/O card or on all I/O cards. If the GR is disabled then all I/O cards accept the diagnose mode. If the GR is enabled, only the I/O card whose select code is in the GR will accept the



diagnose mode. Diagnose Mode 7 is used to disable any service request (SRQ) signal coming into the I/O chip which may cause DMA to cycle during a test. (Mode 7 can be disabled only by a CRS signal (CLC 0).) Diagnose Modes 4 through 6 are reserved for future definition. Diagnose Modes 1 through 3 are described in the following paragraphs.

7-23. DIAGNOSE MODE 1

When an OTA/B 2 instruction is executed with the A- or B-register equal to one each I/O card responds by turning off priority to the next I/O card. When the instruction is complete the only I/O card receiving priority will be the highest priority I/O card (i.e., the one directly next to the processor card. When a subsequent LIA/B 2 instruction is executed, the I/O card receiving priority sets the A- or B-register equal to its select code and identification data (ID) and passes priority to the next I/O card. Having responded once it will not respond again unless Mode 1 is established again. The next LIA/B 2 executed sets the A- or B-register equal to the second I/O card's select code and ID. The second I/O card at completion of the instruction passes priority to the next I/O card. This process continues until the last I/O card responds. After the last I/O card responds the next LIA/B 2 will not affect the A- or B-register and therefore can be detected as a no response. (An OTA/B 2 with the A- or B-register equal to 0 terminates this sequence.)

Mode 1 can also be used to retrieve the select code and ID of a desired I/O card without going through the priority process. This is accomplished by establishing Mode 1 and then executing an LIA/B xx, where xx is the I/O card select code. This procedure will not modify a priority sequence already in process. The Mode 1 select code and ID format is shown in table 7-2.

7-24. DIAGNOSE MODE 2

Diagnose Mode 2 causes an I/O card to respond to an LIA/B 2 instruction in the same manner as in Mode 1 except that the data set into the A- or B-register is as shown in table 7-3.

7-25. DIAGNOSE MODE 3

Diagnose Mode 3 allows an I/O chip to do a DMA transfer without affecting the I/O card. When Mode 3 is entered the I/O chip does a DMA input transfer of the data in the

configuration address register to the location in memory pointed to by the DMA address register. The configuration address register is incremented after each transfer so that the data can be verified. The transfer continues until the DMA count is incremented to zero. Mode 3 also prevents any STC instructions from generating a device command to the I/O card.

Table 7-2. Diagnose Mode 1

A/B BITS	MEANING
15	Intelligent interface
14 } 13 } 12 } 11 } 10 } 9 }	Interface card type identification number
8 } 7 } 6 }	Interface card revision code
5 } 0 }	Interface card select code

Table 7-3. Diagnose Mode 2


A/B BITS	MEANING
15 } 14 } 13 }	Always zero
12	1 = Break feature is enabled
11	1 = Receiving interrupt priority
10	Always zero
9	Control bit
8	Flag bit
7	1 = Global register equals select code of interface card
6	Global register enabled/disabled
5 } 4 } 3 } 2 } 1 } 0 }	Current global register value



CHARACTER CODES

ASCII Character	First Character Octal Equivalent	Second Character Octal Equivalent
A	040400	000101
B	041000	000102
C	041400	000103
D	042000	000104
E	042400	000105
F	043000	000106
G	043400	000107
H	044000	000110
I	044400	000111
J	045000	000112
K	045400	000113
L	046000	000114
M	046400	000115
N	047000	000116
O	047400	000117
P	050000	000120
Q	050400	000121
R	051000	000122
S	051400	000123
T	052000	000124
U	052400	000125
V	053000	000126
W	053400	000127
X	054000	000130
Y	054400	000131
Z	055000	000132
a	060400	000141
b	061000	000142
c	061400	000143
d	062000	000144
e	062400	000145
f	063000	000146
g	063400	000147
h	064000	000150
i	064400	000151
j	065000	000152
k	065400	000153
l	066000	000154
m	066400	000155
n	067000	000156
o	067400	000157
p	070000	000160
q	070400	000161
r	071000	000162
s	071400	000163
t	072000	000164
u	072400	000165
v	073000	000166
w	073400	000167
x	074000	000170
y	074400	000171
z	075000	000172
0	030000	000060
1	030400	000061
2	031000	000062
3	031400	000063
4	032000	000064
5	032400	000065
6	033000	000066
7	033400	000067
8	034000	000070
9	034400	000071
NUL	000000	000000
SOH	000400	000001
STX	001000	000002
ETX	001400	000003
EOT	002000	000004
ENQ	002400	000005

ASCII Character	First Character Octal Equivalent	Second Character Octal Equivalent
ACK	003000	000006
BEL	003400	000007
BS	004000	000010
HT	004400	000011
LF	005000	000012
VT	005400	000013
FF	006000	000014
CR	006400	000015
SO	007000	000016
SI	007400	000017
DLE	010000	000020
DC1	010400	000021
DC2	011000	000022
DC3	011400	000023
DC4	012000	000024
NAK	012400	000025
SYN	013000	000026
ETB	013400	000027
CAN	014000	000030
EM	014400	000031
SUB	015000	000032
ESC	015400	000033
FS	016000	000034
GS	016400	000035
RS	017000	000036
US	017400	000037
SPACE	020000	000040
!	020400	000041
"	021000	000042
#	021400	000043
\$	022000	000044
%	022400	000045
&	023000	000046
'	023400	000047
(024000	000050
)	024400	000051
*	025000	000052
+	025400	000053
,	026000	000054
-	026400	000055
.	027000	000056
/	027400	000057
:	035000	000072
;	035400	000073
<	036000	000074
=	036400	000075
>	037000	000076
?	037400	000077
@	040000	000100
[054000	000133
\	056000	000134
]	056400	000135
^	057000	000136
_	057400	000137
`	060000	000140
{	075400	000173
	076000	000174
}	076400	000175
~	077000	000176
DEL	077400	000177



OCTAL ARITHMETIC

ADDITION

TABLE

0	01	02	03	04	05	06	07
1	02	03	04	05	06	07	10
2	03	04	05	06	07	10	11
3	04	05	06	07	10	11	12
4	05	06	07	10	11	12	13
5	06	07	10	11	12	13	14
6	07	10	11	12	13	14	15
7	10	11	12	13	14	15	16

EXAMPLE

```

Add:   3677  OCTAL
      + 1331  OCTAL
      -----
      (111-)  CARRIES
      5230  OCTAL

```

MULTIPLICATION

TABLE

1	02	03	04	05	06	07
2	04	06	10	12	14	16
3	06	11	14	17	22	25
4	10	14	20	24	30	34
5	12	17	24	31	36	43
6	14	22	30	36	44	52
7	16	25	34	43	52	61

EXAMPLE

```

Multiply: 657  OCTAL
          x 54  OCTAL
          -----
          3274
          4153
          -----
          45024  OCTAL

```

(Reminder: add in octal)

COMPLEMENT

To find the two's complement form of an octal number. (Same procedure whether converting from positive to negative or negative to positive.)

RULE

1. Subtract from the maximum representable octal value.
2. Add one.

EXAMPLE

Two's complement of 556_8

```

  177777
- 000556
-----
  177221
+ 1
-----
  1772228

```

OCTAL/DECIMAL CONVERSIONS

OCTAL TO DECIMAL

TABLE

OCTAL	DECIMAL
0-7	0-7
10-17	8-15
20-27	16-23
30-37	24-31
40-47	32-39
50-57	40-47
60-67	48-55
70-77	56-63
100	64
200	128
400	256
1000	512
2000	1024
4000	2048
10000	4096
20000	8192
40000	16384
77777	32767

EXAMPLE

Convert 463_8 to a decimal integer.

$$\begin{aligned}
 400_8 &= 256_{10} \\
 60_8 &= 48_{10} \\
 3_8 &= \underline{3}_{10} \\
 &307 \text{ decimal}
 \end{aligned}$$

DECIMAL TO OCTAL

TABLE

DECIMAL	OCTAL
1	1
10	12
20	24
40	50
100	144
200	310
500	764
1000	1750
2000	3720
5000	11610
10000	23420
20000	47040
32767	77777

EXAMPLE

Convert 5229_{10} to an octal integer.

$$\begin{aligned}
 5000_{10} &= 11610_8 \\
 200_{10} &= 310_8 \\
 20_{10} &= 24_8 \\
 9_{10} &= \underline{11}_8 \\
 &12155_8 \\
 &\uparrow \\
 &\text{(Reminder: add in octal)}
 \end{aligned}$$

NEGATIVE DECIMAL TO TWO'S COMPLEMENT OCTAL

TABLE

DECIMAL	2's COMP
-1	177777
-10	177766
-20	177754
-40	177730
-100	177634
-200	177470
-500	177014
-1000	176030
-2000	174080
-5000	166170
-10000	154360
-20000	130740
-32768	100000

EXAMPLE

Convert -629_{10} to two's complement octal.

$$\begin{aligned}
 -500_{10} &= 177014_8 \\
 -100_{10} &= 177634_8 \\
 -20_{10} &= 177754_8 \quad \text{(Add in octal)} \\
 -9_{10} &= \underline{177767}_8 \\
 &176613_8
 \end{aligned}$$

For reverse conversion (two's complement octal to negative decimal):

1. Complement, using procedure on facing page.
2. Convert to decimal, using OCTAL TO DECIMAL table.

MATHEMATICAL EQUIVALENTS

2 ± n IN DECIMAL

2 ⁿ	n	2 ⁻ⁿ			
1	0	1.0	65 536	16	0.00001 52587 89062 5
2	1	0.5	131 072	17	0.00000 76293 94531 25
4	2	0.25			
			262 144	18	0.00000 38146 97265 625
			524 288	19	0.00000 19073 48632 8125
8	3	0.125	1 048 576	20	0.00000 09536 74316 40625
16	4	0.0625			
32	5	0.03125			
			2 097 152	21	0.00000 04768 37158 20312 5
			4 194 304	22	0.00000 02384 18579 10156 25
64	6	0.01562 5	8 388 608	23	0.00000 01192 09289 55078 125
128	7	0.00781 25			
256	8	0.00390 625	16 777 216	24	0.00000 00596 04644 77539 0625
			33 554 432	25	0.00000 00298 02322 38769 53125
512	9	0.00195 3125	67 108 864	26	0.00000 00149 01161 19384 76562 5
1 024	10	0.00097 65625			
2 048	11	0.00048 82812 5	134 217 728	27	0.00000 00074 50580 59692 38281 25
			268 435 456	28	0.00000 00037 25290 29846 19140 625
4 096	12	0.00024 41406 25	536 870 912	29	0.00000 00018 62645 14923 09570 3125
8 192	13	0.00012 20703 125			
16 384	14	0.00006 10351 5625	1 073 741 824	30	0.00000 00009 31322 57461 54785 15625
			2 147 483 648	31	0.00000 00004 65661 28730 77392 57812 5
32 768	15	0.00003 05175 78125	4 294 967 296	32	0.00000 00002 32830 64365 38696 28906 25

10 ± n IN OCTAL

10 ⁿ	n	10 ⁻ⁿ	10 ⁿ	n	10 ⁻ⁿ
1	0	1.000 000 000 000 000 00	112 402 762 000	10	0.000 000 000 006 676 337 66
12	1	0.063 146 314 631 463 146 31	1 351 035 564 000	11	0.000 000 000 000 537 657 77
144	2	0.005 075 341 217 270 243 66	16 432 451 210 000	12	0.000 000 000 000 043 136 32
1 750	3	0.000 406 111 564 570 651 77	221 411 634 520 000	13	0.000 000 000 000 003 411 35
23 420	4	0.000 032 155 613 530 704 15	2 657 142 036 440 000	14	0.000 000 000 000 000 264 11
303 240	5	0.000 002 476 132 610 706 64	34 327 724 461 500 000	15	0.000 000 000 000 000 022 01
3 641 100	6	0.000 000 206 157 364 055 37	434 157 115 760 200 000	16	0.000 000 000 000 000 001 63
46 113 200	7	0.000 000 015 327 745 152 75	5 432 127 413 542 400 000	17	0.000 000 000 000 000 000 14
575 360 400	8	0.000 000 001 257 143 561 06	67 405 553 164 731 000 000	18	0.000 000 000 000 000 000 01
7 346 545 000	9	0.000 000 000 104 560 276 41			

MATHEMATICAL EQUIVALENTS

2^x IN DECIMAL

x	2 ^x	x	2 ^x	x	2 ^x
0.001	1.00069 33874 62581	0.01	1.00695 55500 56719	0.1	1.07177 34625 36293
0.002	1.00138 72557 11335	0.02	1.01395 94797 90029	0.2	1.14869 83549 97035
0.003	1.00208 16050 79633	0.03	1.02101 21257 07193	0.3	1.23114 44133 44916
0.004	1.00277 64359 01078	0.04	1.02811 38266 56067	0.4	1.31950 79107 72894
0.005	1.00347 17485 09503	0.05	1.03526 49238 41377	0.5	1.41421 35623 73095
0.006	1.00416 75432 38973	0.06	1.04246 57608 41121	0.6	1.51571 65665 10398
0.007	1.00486 38204 23785	0.07	1.04971 66836 23067	0.7	1.62450 47927 12471
0.008	1.00556 05803 98468	0.08	1.05701 80405 61380	0.8	1.74110 11265 92248
0.009	1.00625 78234 97782	0.09	1.06437 01824 53360	0.9	1.86606 59830 73615

η log₁₀ 2, η log₂ 10 IN OCTAL

η	η log ₁₀ 2	η log ₂ 10	η	η log ₁₀ 2	η log ₂ 10
1	0.30102 99957	3.32192 80949	6	1.80617 99740	19.93156 85693
2	0.60205 99913	6.64385 61898	7	2.10720 99696	23.25349 66642
3	0.90308 99870	9.96578 42847	8	2.40823 99653	26.57542 47591
4	1.20411 99827	13.28771 23795	9	2.70926 99610	29.89735 28540
5	1.50514 99783	16.60964 04744	10	3.01029 99566	33.21928 09489

MATHEMATICAL CONSTANTS IN OCTAL SCALE

π = (3.11037 552421) ₍₈₎	e = (2.55760 521305) ₍₈₎	γ = (0.44742 147707) ₍₈₎
π^{-1} = (0.24276 301556) ₍₈₎	e^{-1} = (0.27426 530661) ₍₈₎	$\ln \gamma$ = -(0.43127 233602) ₍₈₎
$\sqrt{\pi}$ = (1.61337 611067) ₍₈₎	\sqrt{e} = (1.51411 230704) ₍₈₎	$\log_2 \gamma$ = -(0.62573 030645) ₍₈₎
$\ln \pi$ = (1.11206 404435) ₍₈₎	$\log_{10} e$ = (0.33626 754251) ₍₈₎	$\sqrt{2}$ = (1.32404 746320) ₍₈₎
$\log_2 \pi$ = (1.51544 163223) ₍₈₎	$\log_2 e$ = (1.34252 166245) ₍₈₎	$\ln 2$ = (0.54271 027760) ₍₈₎
$\sqrt{10}$ = (3.12305 407267) ₍₈₎	$\log_2 10$ = (3.24464 741136) ₍₈₎	$\ln 10$ = (2.23273 067355) ₍₈₎

OCTAL COMBINING TABLES

MEMORY REFERENCE INSTRUCTIONS

INDIRECT ADDRESSING

Refer to octal instruction codes given on the following page.
To combine code for indirect addressing, merge "100000" with octal instruction code.

REGISTER REFERENCE INSTRUCTIONS

SHIFT-ROTATE GROUP (SRG)

1. select to operate A or B.
2. select 1 to 4 instructions, not more than one from each column.
3. combine octal codes (leading zeros omitted) by inclusive or.
4. order of execution is from column 1 to column 4.

A OPERATIONS

1	2	3	4
ALS (1000)	CLE (40)	SLA (10)	ALS (20)
ARS (1100)			ARS (21)
RAL (1200)			RAL (22)
RAR (1300)			RAR (23)
ALR (1400)			ALR (24)
ERA (1500)			ERA (25)
ELA (1600)			ELA (26)
ALF (1700)			ALF (27)

B OPERATIONS

1	2	3	4
BLS (5000)	CLE (4040)	SLB (4010)	BLS (4020)
BRS (5100)			BRS (4021)
RBL (5200)			RBL (4022)
RBR (5300)			RBR (4023)
BLR (5400)			BLR (4024)
ERB (5500)			ERB (4025)
ELB (5600)			ELB (4026)
BLF (5700)			BLF (4027)

ALTER-SKIP GROUP (ASG)

1. select to operate on A or B.
2. select 1 to 8 instructions, not more than one from each column.
3. combine octal codes (leading zeros omitted) by inclusive or.
4. order of execution is from column 1 to column 8.

A OPERATIONS

1	2	3	4
CLA (2400)	SEZ (2040)	CLE (2100)	SSA (2020)
CMA (3000)		CME (2200)	
CCA (3400)		CCE (2300)	
5	6	7	8
SLA (2010)	INA (2004)	SZA (2002)	RSS (2001)

B OPERATIONS

1	2	3	4
CLB (6400)	SEZ (6040)	CLE (6100)	SSB (6020)
CMB (7000)		CME (6200)	
CCB (7400)		CCE (6300)	
5	6	7	8
SLB (6010)	INB (6004)	SZB (6002)	RSS (6001)

INPUT/OUTPUT INSTRUCTIONS

CLEAR FLAG

Refer to octal instruction codes given on the following page.
To clear flag after execution (instead of holding flag), merge "001000" with octal instruction code.

INSTRUCTION CODES IN OCTAL

Memory Reference	SLB 006010	DSY 105771	.DCO 105204	VSUB 105003
ADA 04(0XX)—	SSA 002020	ISX 105760	.DDE 105211	VSUM 105105
ADB 04(1XX)—	SSB 006020	ISY 105770	.DDI 105074	VSWP 105117
AND 01(0XX)—	SZA 002002	JLY 105762	.DDIR 105134	DVABS 105123
CPA 05(0XX)—	SZB 006002	JPY 105772	.DDS 105213	DVADD 105021
CPB 05(1XX)—		LAX 101742	.DIN 105210	DVDIV 105025
IOR 03(0XX)—	Input/Output	LAY 101752	.DIS 105212	DVDOT 105130
ISZ 03(1XX)—	CLC 1067—	LBT 105763	.DNG 105203	DVMAB 105132
JMP 02(1XX)—	CLF 1031—	LBX 105742	.DMP 105054	DVMAX 105131
JSB 01(1XX)—	CLO 103101	LBY 105752	.DSB 105034	DVMIB 105135
LDA 06(0XX)—	HLT 1020—	LDX 105745	.DSBR 105114	DVMIN 105133
LDB 06(1XX)—	LIA 1025—	LDY 105755		DVMOV 105136
STA 07(0XX)—	LIB 1065—	MBT 105765	VMA/EMA	DVMPY 105024
STB 07(1XX)—	MIA 1024—	MVW 105777	.IMAP 105250	DVNRM 105127
XOR 02(0XX)—	MIB 1064—	SAX 101740	.IRES 105244	DVPIV 105121
	OTA 1026—	SAY 101750	.JMAP 105252	DVSAD 105026
	OTB 1066—	SBS 105773	.JRES 105245	DVSDV 105031
Binary	SFC 1022—	SBT 105764	.LBP 105257	DVSMY 105030
	SFS 1023—	SBX 105740	.LBPR 105256	DVSSB 105027
Shift-Rotate	SOC 102201	SBY 105750	.LPX 105255	DVSUB 105023
ALF 001700	SOS 102301	SFB 105767	.LPXR 105254	DVSUM 105125
ALR 001400	STC 1027—	STX 105743	.PMAP 105240	DVSWP 105137
ALS 001000	STF 1021—	STY 105753		
ARS 001100	STO 102101	TBS 105775	Oper. Syst. Set	Dynamic Map Syst.
BLF 005700		XAX 101747	.CPUID 105300	LDMP 105702
BLR 005400	Extended Arithmetic	XAY 101757	.FWID 105301	LPMR 105700
BLS 005000	ASL 1000(01X)—	XBX 105747	.SIP 105303	LWD1 105704
BRS 005100	ASR 1010(01X)—	XBY 105757	.WFI 105302	MB00 101727
CLE 000040	DIV 100400			MB01 101730
ELA 001600	JLA 100600	Floating Point	Scientific Inst. Set	MB02 101731
ELB 005600	DLD 104200	FAD 105000	ALOG 105322	MB10 101732
ERA 001500	DST 104400	FDV 105060	ALOGT 105327	MB11 101733
ERB 005500	JLB 104600	FIX 105100	ATAN 105323	MB12 101734
NOP 000000	LSL 1000(10X)—	FLT 105120	/ATLG 105333	MB20 101735
RAL 001200	LSR 1010(10X)—	FMP 105040	/CMRT* 105332	MB21 101736
RAR 001300	MPY 100200	FSB 105020	COS 105324	MB22 101737
RBL 005200	RRL 1001(00X)—	.FIXD 105104	DPOLY 105331	MW00 105727
RBR 005300	RRR 1011(00X)—	.FLTD 105124	EXP 105326	MW01 105730
SLA 000010		.TADD 105002	.FPWR 105334	MW02 105731
SLB 004010		.TDIV 105062	SIN 105325	MW10 105732
	Binary	.TFTD 105122	SQRT 105321	MW11 105733
	Ext. Inst. Group	.TFXD 105106	TAN 105320	MW12 105734
Alter-Skip	ADX 105746	.TFXS 105102	TANH 105330	MW20 105735
CCA 003400	ADY 105756	.TMPY 105042	.TPWR 105335	MW21 105736
CCB 007400	CAX 101741	.TSUB 105022		MW22 105737
CCE 002300	CAY 101751		Vector Inst. Set	SIMP 105707
CLA 002400	CBS 105774	Language Inst. Set	VABS 105103	STMP 105703
CLB 006400	CBT 105766	.BLE 105207	VADD 105001	SPMR 105701
CLE 002100	CBX 105741	.CFER 105231	VDIV 105005	SWMP 105706
CMA 003000	CBY 105751	.DFER 105205	VDOT 105110	XCA1 101726
CMB 007000	CMW 105776	.CPM 105236	VMAB 105112	XCA2 101723
CME 002200	CXA 101744	.ENTC 105235	VMAX 105111	XCB1 105726
INA 002004	CXB 105744	.ENTN 105234	VMIB 105115	XCB2 105723
INB 006004	CYA 101754	.ENTP 105224	VMIN 105113	XJMP 105710
RSS 002001	CYB 105754	.ENTR 105223	VMOV 105116	XLA1 101724
SEZ 002040	DSX 105761	.FCM 105232	VMPY 105004	XLA2 101721
SLA 002010		.NGL 105214	VNRM 105107	XLB1 105724
		.SETP 105227	VPIV 105101	XLB2 105721
		.TCM 105233	VSAD 105006	XSA1 101725
		.XFER 105220	VSDV 105011	XSA2 101722
		.ZFER 105237	VSMY 105010	XSB1 105725
			VSSB 105007	XSB2 105722
		Double Integer		
		.DAD 105014		
<p>Assuming: no indirect addressing. no combined instructions. shifts taken in first position only. hold flag after I/O execution.</p> <p>* Not directly user callable. Used by HP software.</p> <p>Refer to preceding page for octal combining tables.</p>				

BASE SET INSTRUCTION CODES IN BINARY

15	14	13	12	11	10	9	8	7	6	5	4	3	2	1	0
MEMORY REFERENCE INSTRUCTIONS															
D/I	AND	001		0	Z/C	← MEMORY ADDRESS →									
D/I	XOR	010		0	Z/C										
D/I	IOR	011		0	Z/C										
D/I	JSB	001		1	Z/C										
D/I	JMP	010		1	Z/C										
D/I	ISZ	011		1	Z/C										
D/I	AD*	100		A/B	Z/C										
D/I	CP*	101		A/B	Z/C										
D/I	LD*	110		A/B	Z/C										
D/I	ST*	111		A/B	Z/C										
SHIFT/ROTATE GROUP															
0		000		A/B	0	D/E	*LS		000	†CLE	D/E	‡SL*	*LS		000
				A/B	0	D/E	*RS		001		D/E		*RS		001
				A/B	0	D/E	R*L		010		D/E		R*L		010
				A/B	0	D/E	R*R		011		D/E		R*R		011
				A/B	0	D/E	*LR		100		D/E		*LR		100
				A/B	0	D/E	ER*		101		D/E		ER*		101
				A/B	0	D/E	EL*		110		D/E		EL*		110
				A/B	0	D/E	*LF		111		D/E		*LF		111
				NOP	000				000		000				000
ALTER/SKIP GROUP															
0		000		A/B	1	CL*	01	CLE	01	SEZ	SS*	SL*	IN*	SZ*	RSS
				A/B		CM*	10	CME	10						
				A/B		CC*	11	CCE	11						
INPUT/OUTPUT GROUP															
1		000			1	H/C	HLT		000	← SELECT CODE →					
					1	0	STF		001						
					1	1	CLF		001						
					1	0	SFC		010						
					1	0	SFS		011						
				A/B	1	H/C	MI*		100						
				A/B	1	H/C	LI*		101						
				A/B	1	H/C	OT*		110						
				0	1	H/C	STC		111						
				1	1	H/C	CLC		111						
					1	0	STO		001		000			001	
					1	1	CLO		001		000			001	
					1	H/C	SOC		010		000			001	
					1	H/C	SOS		011		000			001	
EXTENDED ARITHMETIC GROUP															
1		000		MPY**		000		010			000			000	
				DIV**		000		100			000			000	
				JLA		000		110			000			000	
				DLD**		100		010			000			000	
				DST**		100		100			000			000	
				JLB		100		110			000			000	
				ASR		001		000		0	1				
				ASL		000		000		0	1				
				LSR		001		000		1	0	← NUMBER OF BITS →			
				LSL		000		000		1	0				
				RRR		001		001		0	0				
				RRL		000		001		0	0				
FLOATING POINT INSTRUCTIONS															
1		000			101			00		FAD	000	0		000	
										FSB	001				
										FMP	010				
										FDV	011				
										FIX	100				
										FLT	101				
Notes: * = A or B, according to bit 11. D/I, A/B, Z/C, D/E, H/C coded 0/1. **Second word is Memory Address.															
										†CLE: Only this bit is required.					
										‡SL*: Only this bit and bit 11 (A/B as applicable) are required.					
														8200-56	

BASE SET INSTRUCTION CODES IN BINARY (Continued)

15	14	13	12	11	10	9	8	7	6	5	4	3	2	1	0
FLOATING POINT INSTRUCTION (Continued)															
1	000			101			00			.TADD 000 .TSUB 001 .TMPY 010 .TDIV 011 .TFXS 100 .TFTS 101 .FIXD 100 .FLTD 101 .TFXD 100 .TFTD 101		0	010 100 110		
DOUBLE INTEGER INSTRUCTIONS															
1	000			101			000 001 001 010			001 011 101 111 001 011 000 001		.DAD 100 .DSB 100 .DMP 100 .DDI 100 .DSBR 100 .DDIR 100 .DNG 011 .DCO 100 .DIN 000 .DDE 001 .DIS 010 .DDS 011			
LANGUAGE INSTRUCTION SET															
1	000			101			010			0	00 01 10 11		.DFER 101 .BLE 111 .NGL 100 .XFER 000 .ENTR 011 .ENTP 100 .SETP 111 .CFER 001 .FCM 010 .TCM 011 .ENTN 100 .ENTC 101 .CPM 110 .ZFER 111		
VIRTUAL MEMORY INSTRUCTIONS															
1	000			101			010			100 101		.PMAP 000 .IRES 100 .JRES 101 .JMAP 000 .JMAP 010 .LPXR 100 .LPX 101 .LBPR 110 .LBP 111			
OPERATING SYSTEM INSTRUCTION SET															
1	000			101			011			000		.CPUID 000 .FWID 001 .WFI 010 .SIP 011			

BASE SET INSTRUCTION CODES IN BINARY (Continued)


15	14	13	12	11	10	9	8	7	6	5	4	3	2	1	0
DMS INSTRUCTIONS															
1	000			1	01		111			000			LPMR 000		
				1									SPMR 001		
				1									LDMP 010		
				1									STMP 011		
				1									LWD1 100		
				1									LWD2 101		
				1									SWMP 110		
				1									SIMP 111		
				1									XJMP 000		
				A/B							001		XL*1 100		
				A/B							010		XS*1 101		
				A/B									XC*1 110		
				A/B									XL*2 001		
				A/B									XS*2 010		
				A/B									XC*2 011		
				B/W									M°00 111		
				B/W							011		M°01 000		
				B/W									M°02 001		
				B/W									M°10 010		
				B/W									M°11 011		
				B/W									M°12 100		
				B/W									M°20 101		
				B/W									M°21 110		
				B/W									M°22 111		
SCIENTIFIC INSTRUCTION SET															
1	000			101		011			010			TAN 000			
													SQRT 001		
													ALOG 010		
													ATAN 011		
													COS 100		
													SIN 101		
													EXP 110		
													ALOGT 111		
											011		TANH 000		
													DPOLY 001		
													/CMRT 010		
													/ATLG 011		
													.FPWR 100		
													.TPWR 101		
VECTOR INSTRUCTION SET															
1	000			101		000			000			VADD 001			
													VSUB 011		
													VMPY 100		
													VDIV 101		
													VSAD 110		
													VSSB 111		
											001		VSMY 000		
													VSDV 001		
											010		DVADD 001		
													DVSUB 011		
													DVMPY 100		
													DVDIV 101		
													DVSAD 110		
													DVSSB 111		
											011		DVSMY 000		
													DVSDV 001		
											000		VPIV 001		
													VABS 011		
													VSUM 101		
Notes: * = A (0) or B (1), according to bit 11. ° = B (0) or W (1), according to bit 11.															

BASE SET INSTRUCTION CODES IN BINARY (Continued)


15	14	13	12	11	10	9	8	7	6	5	4	3	2	1	0
VECTOR INSTRUCTION SET (Continued)															
1	000			101			001			011			VNRM 111		
													VDOT 000		
													VMAX 001		
													VMAB 010		
													VMIN 011		
													VMIB 101		
													VMOV 110		
													VSWP 111		
													DPIV 001		
													DVABS 011		
													DVSUM 101		
													DVNRM 111		
													DVDOT 000		
													DVMAX 001		
													DVMAB 010		
													DVMIN 011		
DVMIB 101															
DVMOV 110															
DVSWP 111															

EXTEND AND OVERFLOW EXAMPLES

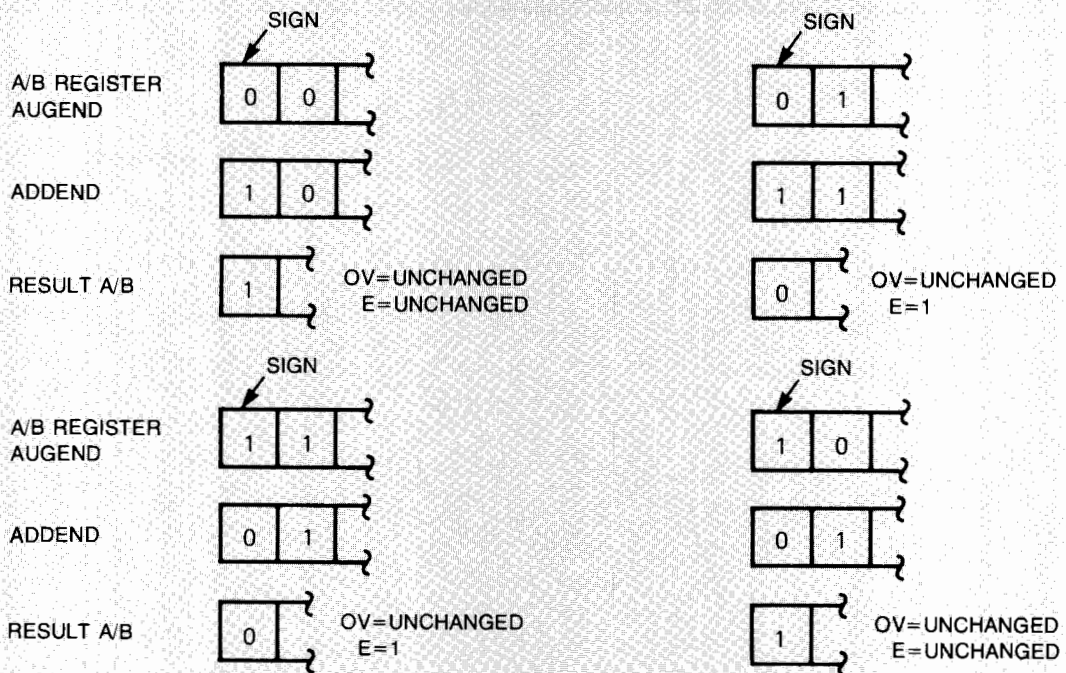
SAME SIGN (POSITIVE)



SAME SIGN (NEGATIVE)



DIFFERENT SIGNS



INTERRUPT AND CONTROL SUMMARY

INST	S.C. 00	S.C. 01	S.C. 02	S.C. 03	S.C. 04	S.C. 05	S.C. 06	S.C. 07
STC	NOP	NOP	Enable break mode.	NOP	Enable Type 2 and 3 interrupts.	Enable parity error interrupts.	Turn on Time Base Generator.	Turn on memory protect.
CLC	System reset.	NOP	NOP	NOP	Disable Type 2 and 3 interrupts.	Disable parity error interrupts.	Turn off Time Base Generator.	NOP
STF	Enable Type 3 interrupts.	STO	Disable Global Register.	NOP	NOP	Set parity sense to even parity.	Set Time Base Generator flag.	NOP
CLF	Disable Type 3 interrupts.	CLO	Enable Global Register.	NOP	NOP	Set parity sense to odd parity.	Clear Time Base Generator flag.	NOP
SFS	Skip if Type 3 interrupts are enabled.	SOS	Skip if Global Register is disabled.	NOP	Skip if power not going down	Skip if parity sense is even.	Skip if Time Base Generator flag is set.	NOP
SFC	Skip if Type 3 interrupts are disabled.	SOC	Skip if Global Register is enabled.	NOP	Skip if power is going down.	Skip if parity sense is odd.	Skip if Time Base Generator flag is clear.	NOP
LI*	Load from interrupt mask register.	Load from processor status register.	Load from Global Register.	Load from PSAVE or (with ,C) ROMP.	Load from central interrupt register.	Load bits 0-15 from parity error register, or (with ,C) bits 16-23.	NOP	Load from violation register.
MI*	NOP	Merge from processor status register.	NOP	NOP	NOP	NOP	NOP	NOP
OT*	Output to interrupt mask register.	Output to processor status register.	Output to Global Register. (Note 1)	Output to PSAVE or (with ,C) ROMP.	Output to central interrupt register.	NOP	NOP	NOP

Note 1: An OTA/B 2 with A/B equal to one through seven establishes a diagnose mode; refer to paragraph 7-22 for details.

