

digital

processor
handbook

pdp11/45

digital

pdp-11/45

processor
handbook

digital equipment corporation

Copyright 1973 by
Digital Equipment Corporation
PDP, UNIBUS, DIGITAL are registered trademarks
of Digital Equipment Corporation

The PDP-11 is a family of upward-compatible computer systems. We believe that these systems represent a significant departure from traditional methods of computer design.

The initial design step was the development of a totally new language, notation, and theory of computers called the Instruction Set Processor (ISP). This language provides a concise and powerful generalized method for defining an arbitrary computer system and its operation. Along with the development of ISP, a PDP-10 program was written for simulating the operation of any computer system on the basis of its ISP description. With the aid of ISP and the machine simulation program, benchmark comparison tests were run on a large number of potential computer designs. In this manner it was possible to evaluate a variety of design choices and compare their features and advantages, without the time and expense of actually constructing physical prototypes.

Since the main design objective of the PDP-11 was to optimize total system performance, the interaction of software and hardware was carefully considered at every step in the design process. System programmers continually evaluated the efficiency of the code which would be produced by the system software, the ease of coding a program, the speed of real-time response, the power and speed that could be built into a system executive, the ease of system resource management, and numerous other potential software considerations.

The current PDP-11 Family is the result of this design effort. We believe that its general purpose register and UNIBUS organization provides unparalleled power and flexibility. This design is the basis for our continuing commitment to further PDP-11 product development.

Thus the PDP-11 Family is at once a new concept in computer systems, and a tested and tried system. The ultimate proof of this new design approach has come from the large and rapidly increasing number of PDP-11 users all around the world.



Kenneth H. Olsen
President,
Digital Equipment Corporation

CONTENTS

CHAPTER 1 - INTRODUCTION	1
1.1 PDP-11 FAMILY	1
1.2 GENERAL CHARACTERISTICS	1
1.3 PERIPHERALS AND OPTIONS	5
1.4 SOFTWARE	6
CHAPTER 2 - SYSTEM ARCHITECTURE	9
2.1 INTRODUCTION	9
2.2 THE UNIBUS	10
2.3 CENTRAL PROCESSOR	11
2.4 FLOATING POINT PROCESSOR	14
2.5 MEMORY	15
2.6 SYSTEM INTERACTION	19
2.7 PROCESSOR TRAPS	22
2.8 MULTI-PROGRAMMING	24
CHAPTER 3 - ADDRESSING MODES	25
3.1 SINGLE OPERAND ADDRESSING	26
3.2 DOUBLE OPERAND ADDRESSING	26
3.3 DIRECT ADDRESSING	27
3.4 DEFERRED ADDRESSING	32
3.5 USE OF PC	35
3.6 USE OF STACK POINTER	38
CHAPTER 4 - INSTRUCTION SET	41
4.1 INTRODUCTION	41
4.2 INSTRUCTION FORMATS	42
4.3 BYTE INSTRUCTIONS	43
4.4 SINGLE OPERAND INSTRUCTIONS	45
4.5 DOUBLE OPERAND INSTRUCTIONS	65
4.6 PROGRAM CONTROL INSTRUCTIONS	77
4.7 MISCELLANEOUS INSTRUCTIONS	118
4.8 CONDITION CODE OPERATORS	126


CHAPTER 5 - ADVANCED PROGRAMMING TECHNIQUES	127
5.1 THE STACK	127
5.2 SUBROUTINE LINKAGE	131
5.3 INTERRUPTS	135
5.4 REENTRANCY	137
5.5 POSITION INDEPENDENT CODE	140
5.6 RECURSION	140
5.7 CO-ROUTINES	141
CHAPTER 6 - MEMORY MANAGEMENT	143
6.1 BASIC ADDRESSING LOGIC	143
6.2 VIRTUAL ADDRESSING	144
6.3 INTERRUPT CONDITIONS UNDER MANAGEMENT CONTROL	145
6.4 CONSTRUCTION OF A PHYSICAL ADDRESS	145
6.5 MANAGEMENT REGISTERS	147
6.6 FAULT RECOVERY REGISTERS	150
6.7 EXAMPLES	154
6.8 TRANSPARENCY	159
6.9 MANAGEMENT REGISTER MAP	160
CHAPTER 7 - FLOATING POINT PROCESSOR	163
7.1 INTRODUCTION	163
7.2 OPERATION	163
7.3 ARCHITECTURE	164
7.4 FLOATING POINT DATA FORMATS	165
7.5 FPP STATUS REGISTER	166
7.6 FEC REGISTER	168
7.7 FPP INSTRUCTION ADDRESSING	168
7.8 INSTRUCTION TIMING	169
7.9 FPP INSTRUCTIONS	169
CHAPTER 8 - SYSTEM OPERATOR'S CONSOLE	195
8.1 CONSOLE ELEMENTS	195
8.2 SYSTEM POWER SWITCH	196
8.3 CPU STATE INDICATORS	196
8.4 ADDRESS DISPLAY REGISTER	197
8.5 ADDRESSING ERROR DISPLAY	198
8.6 DATA DISPLAY REGISTER	198
8.7 SWITCH REGISTER	198
8.8 CONTROL SWITCHES	199

APPENDICES

APPENDIX A - INSTRUCTION SET PROCESSOR (ISP)	203
APPENDIX B - INSTRUCTION TIMING	217
APPENDIX C - MEMORY MAP	229
APPENDIX D - PROGRAM INTERRUPT REQUESTS	235
APPENDIX E - MEMORY PARITY	237

INDEXES

GENERAL INDEX	239
INSTRUCTION INDEX	240



CHAPTER 1

INTRODUCTION

The PDP-11/45 is a powerful 16-bit computer representing the large computer end of the PDP-11 family of computers. It is designed as a powerful computational tool for high-speed real-time applications and for large multi-user, multi-task applications requiring up to 124K words of addressable memory space. It will operate with solid state and core memories, and includes many features not normally associated with 16-bit computers. Among its major features are a fast central processor with choices of 300 or 450 nanosecond memory, an advanced Floating Point Processor, and a sophisticated memory management scheme.

1.1 THE PDP-11 FAMILY

The PDP-11 family includes several processors, a large number of peripheral devices and options, and extensive software. PDP-11 machines are architecturally similar and hardware and software upwards compatible, although each machine has some of its own characteristics. New PDP-11 systems will be compatible with existing family members. The user can choose the system which is most suitable to his application, but as needs change or grow he can easily add or change hardware. The major characteristics of PDP-11 family computers are summarized in Table 1-1 at the end of this chapter.

1.2 GENERAL CHARACTERISTICS

1.2.1 The UNIBUS

All computer system components and peripherals connect to and communicate with each other on a single high-speed bus known as the UNIBUS—the key to the PDP-11's many strengths. Since all system elements, including the central processor, communicate with each other in identical fashion via the UNIBUS, the processor has the same easy access to peripherals as it has to memory.

With bidirectional and asynchronous communication on the UNIBUS, devices can send, receive and exchange data independently without processor intervention. For example, a CRT display can refresh itself from a disk file while the central processor unit (CPU) attends to other tasks. Because it is asynchronous the UNIBUS is compatible with devices operating over a wide range of speeds.

Device communications on the UNIBUS are interlocked. For each command issued by a "master" device, a response signal is received from a "slave" completing the data transfer. Device-to-device communication is completely independent of physical bus length and the response times of master and slave devices. Interfaces to the UNIBUS are not time-dependent; there are no pulse-width or rise-time restrictions to worry

about. The maximum transfer rate on the UNIBUS is one 16-bit word every 400 nanoseconds, or 2,500,000 words per second.

Input/output (I/O) devices transferring directly to or from memory are given highest priority, and may request bus mastership and "steal" bus cycles during instruction operations. The processor resumes operation immediately after the memory transfer. Multiple devices can operate simultaneously at maximum direct memory access (DMA) rates by stealing bus cycles. The UNIBUS is further explained in Paragraph 2.2, Chapter 2, and is covered in considerable detail in the PDP-11 Peripherals and Interfacing Handbook.

1.2.2 Central Processor

The central processor, connected to the UNIBUS as a subsystem, controls the allocation of the UNIBUS for peripherals and performs arithmetic and logic operations and instruction decoding. It contains multiple high-speed general-purpose registers which can be used as accumulators, pointers, index registers, or as autoindexing pointers in autoincrement or autodecrement modes. The processor can perform data transfers directly between I/O devices and memory without disturbing the registers; does both single- and double-operand addressing; handles both 16-bit word and 8-bit byte data; and, by using its dynamic stacking technique, allows nested interrupts and automatic reentrant subroutine calling.

Instruction Set

The instruction complement uses the flexibility of the general-purpose registers to provide over 400 powerful hard-wired instructions—the most comprehensive and powerful instruction repertoire of any computer in the 16-bit class. Unlike conventional 16-bit computers, which usually have three classes of instructions—memory reference instructions, operate or AC control instructions, and I/O instructions—all operations in the PDP-11 are accomplished with one set of instructions. Since peripheral device registers can be manipulated as flexibly as core memory by the central processor, instructions that are used to manipulate data in core memory may be used equally well for data in peripheral device registers. For example, data in an external device register can be tested or modified directly by the CPU without bringing it into memory or disturbing the general registers. One can add data directly to a peripheral device register, or compare contents with a mask and branch. Thus all PDP-11 instructions can be used to create a new dimension in the treatment of computer I/O and the need for a special class of I/O instructions is eliminated. PDP-11/45 instructions are described in Chapter 4.

The following example contrasts the rotate operation in the PDP-11 with a similar operation in a conventional computer:

	PDP-11 Approach		Conventional Approach
ROR A	;rotate contents of memory location A right one place	LDA A	;load contents of memory location A into AC
		ROT	;rotate contents of AC right one place
		STA A	;store contents of AC in location A

The basic order code of the PDP-11 uses both single and double operand address instructions for words or bytes. The PDP-11 therefore performs very efficiently in one step such operations as adding or subtracting two operands, or moving an operand from one location to another:

PDP-11 Approach		Conventional Approach	
ADD A,B	;add contents of location A to location B	LDA A	;load contents of memory location A into AC
		ADD B	;add contents of memory location B
		STA B	;store results at location B to AC

Priority Interrupts

A multi-level automatic priority interrupt system permits the processor to respond automatically to conditions outside the system, or in the processor itself. Any number of separate devices can be attached to each level.

Each peripheral device in the PDP-11 system has a hardware pointer to its own pair of memory words (one points to the device's service routine, and the other contains the new status information). This unique identification eliminates the need for polling of devices to identify an interrupt, since the interrupt servicing hardware selects and begins executing the appropriate service routine, after having automatically saved the status of the interrupted program segment.

The devices interrupt priority and service routine priority are independent. This allows adjustment of system behavior in response to real-time conditions, by dynamically changing the priority level of the service routine.

The interrupt system allows the processor to continually compare its own priority level with the level of any interrupting devices and to acknowledge the device with the highest level above the processors priority level. Servicing an interrupt for a device can be interrupted for servicing a higher priority device. Service to the lower priority device is resumed automatically upon completion of the higher level servicing. Such a process, called nested interrupt servicing, can be carried out to any level, without requiring the software to worry about the saving and restoring of processor status at each level.

The interrupt scheme is explained in Paragraph 2.6, Chapter 2.

Reentrant Code

Both the interrupt handling hardware and the subroutine call hardware are designed to facilitate writing reentrant code for the PDP-11. This type of code allows use of a single copy of a given subroutine or program to be shared by more than one process or task. This reduces the amount of core needed for multi-task applications such as concurrent servicing of many peripheral devices.

Addressing

Much of the power of the PDP-11 is derived from its wide range of addressing capabilities. PDP-11 addressing modes include list sequential addressing, full address indexing, full 16-bit word addressing, 8-bit byte addressing, and stack addressing. Variable length instruction formatting allows a minimum number of bits to be used for each addressing mode. This results in efficient use of program storage space. Addressing modes are described in Chapter 3.

Stacks

In a PDP-11, a stack is a temporary data storage area which allows a program to make efficient use of frequently accessed data. The stack is used automatically by program interrupts, subroutine calls, and trap instructions. When the processor is interrupted, the central processor status word and the program counter are saved (pushed) into the stack area, while the processor services the interrupting device. A new status word is then automatically acquired from an area in core memory which is reserved for interrupt instructions (vector area). A return from the interrupt instruction restores the original processor status and returns control to the interrupted program without software intervention. Stacks are explained in Chapter 5.

Direct Memory Access

All PDP-11s provide for direct access to memory. Any number of DMA devices may be attached to the UNIBUS. Maximum priority is given to DMA devices thus allowing memory data storage or retrieval at memory cycle speeds. Latency is minimized by the organization and logic of the UNIBUS which samples requests and priorities in parallel with data transfers.

Power Fail and Restart

Power fail and restart, not only protects memory when power fails, but also allows the user to save the existing program location and status (including all dynamic registers), thus preventing harm to devices, and eliminating the need for reloading programs. Automatic restart is accomplished when power returns to safe operating levels, enabling remote or unattended operations of PDP-11 systems.

1.2.3 Memories

Memories with different ranges of speeds and various characteristics can be freely mixed and interchanged in a single PDP-11 system. Thus as memory needs expand and as memory technology grows, a PDP-11 can evolve, with none of the growing pains and obsolescence associated with conventional computers. See Paragraph 2.5, Chapter 2.

1.2.4 Floating Point Processor

An advanced-design Floating Point Processor functions as an integral part of the PDP-11/45 central processor. Floating point instructions overlap CPU instructions and can continue without CPU intervention, leaving the CPU free to execute other instructions. Floating Point Processor instructions are described in Chapter 7.

1.2.5 Memory Management

PDP-11/45 memory management is an advanced memory extension, relocation and protection feature which will:

- extend memory space from 28K to 124K words

- provide effective protection of memory pages in multi-user environments.

Memory Management is explained in Chapter 6.

1.3 PERIPHERALS OPTIONS

Digital Equipment Corporation designs and manufactures many of the peripheral devices offered with PDP-11s. As a designer and manufacturer of peripherals, DIGITAL can offer extremely reliable equipment, lower prices, more choices, and quantity discounts.

Many processor, input/output, memory, bus, and storage options are available. These devices are explained in detail in the Peripherals and Interfacing Handbook.

1.3.1 I/O Devices

All PDP-11 systems are available with Teletypes as standard equipment. However, their I/O capabilities can be increased with high-speed paper tape readers—punches, line printers, card readers or alphanumeric display terminals. The LA30 DECwriter, a totally DIGITAL designed and built teleprinter, can serve as an alternative to the Teletype. It has several advantages over standard electromechanical typewriter terminals, including higher speed, fewer mechanical parts and very quiet operation.

PDP-11 terminals include:

- DECterminal alphanumeric display

- DECwriter teleprinter

- High-speed line printers

- High-speed paper tape reader punch

- Teletypes

- Card readers

1.3.2 Storage Devices

Storage devices range from convenient, small-reel magnetic tape units to mass storage magnetic tapes and disk memories. With the UNIBUS, a large number of storage devices, in any combination, may be connected to a PDP-11 system. TU56 DECTapes, highly reliable tape units with small tape reels, designed and built by DIGITAL, are ideal for applications with modest storage requirements. Each DECTape provides storage for 174K 16 bit words. For applications which require handling of large volumes of data, DIGITAL offers the industry compatible TU10 Magtape.

Disk storage devices include fixed head disk units and moving-head re-

movable cartridge and disk pack units. These devices range from the 65K RS64 DECdisk memory, to the RP02 Disk Pack system which can store up to 93.6 million words. PDP-11 storage devices include:

DECtape

Magtape

RS64 64K-256K word fixed head disk

RF11 256K-2M word fixed head disk

RK03 1-2M word moving head disk

RP02 10M word moving head disk

1.3.3 Bus Options

Several options (bus switches, bus extenders) are available for extending the UNIBUS or for configuring multi-processor or shared-peripheral systems.

1.4 SOFTWARE

Extensive software, consisting of disk and paper tape systems, is available for PDP-11 Family systems. The larger the PDP-11 configuration, the larger and more comprehensive the software package that comes with it.

1.4.1 Disk Operating System Software

The Disk Operating System software includes:

Text Editor (ED-11)

Relocatable Assembler (PAL-11R)

Linker (LINK-11)

File Utilities Packages (PIP)

On Line Debugging Technique (ODT-11R)

Librarian (LIBR-11)

1.4.2 Higher Level Languages

PDP-11 users needing an interactive conversational language can use BASIC which can be run on the paper tape software system with only 4.096 words of core memory. A multi-user extension of BASIC is available so up to eight terminal users can access a PDP-11 with only 8K of core.

RSTS/E

The PDP-11 Resource Timesharing System (RSTS/E) with BASIC-PLUS, an enriched version of BASIC, is available for up to 32 terminal users.

FORTRAN

PDP-11 FORTRAN is an ANSI-standard Fortran IV compiler with elements that provide easy compatibility with IBM 1130 FORTRAN.

Paper tape software is available on systems without disks.

1.5 DATA COMMUNICATIONS

The advanced architecture of PDP-11 family machines makes them ideal for use in data communications applications. For example the UNIBUS performs like a multiplexer and multiple single-line interfaces can be added without special multiplexing hardware: byte handling, the key to communications applications, is accomplished easily and efficiently by the PDP-11. To provide total systems capability in the communications area DIGITAL has developed a full line of communications hardware and communications-oriented software.

1.6 DATA ACQUISITION AND CONTROL

The PDP-11, modular process interfaces, and special state-of-the art software (RSX-11D real-time executive) combine to provide efficient, low-cost and reliable systems for industrial data acquisition and control (IDACS) applications. IDACS-11 hardware is described in the Peripherals and Interfacing Handbook.

TABLE 1-1
PDP-11 Family Computers

	PDP-11/05	PDP-11/15	PDP-11/20	PDP-11/45
Central Processor	KD11	KC11	KA11	KB11
General Purpose Registers	8	8	8	16
Instructions	Basic Set	Basic Set	Basic Set	Basic Set and: MUL, DIV, XOR. ASH, ASHC, MARK, SXT, SOB, SPL, RTT MTPX, MFPX
Memory Management Hardware	No	No	No	Optional
Stacks	Yes	Yes	Yes	Yes
Stack overflow Detection	Fixed	Fixed	Fixed	Variable
Automatic Priority Interrupts	Single-Line, Multi-Level	Single-Line Multi-Level (four line optional)	Multi-Line Multi-Level	Multi-Line Multi-Level Plus 7 Software Levels
Overlapped Instructions	No	No	No	Yes
Extended Arithmetic	Option	Option	Option	Standard
Floating Point	Software	Software	Software	Internal to CPU
Basic Memory	Core	Core	Core	Core, MOS or Bipolar

SYSTEM ARCHITECTURE

2.1 INTRODUCTION

The PDP-11/45 is a medium scale general purpose computer designed around the basic architecture of all PDP-11 family machines.

The Central Processing Unit has a cycle time of 300 nsec and performs all arithmetic and logical operations required in the system. A Floating Point Processor (described in Chapter 7) mounts integrally into the Central Processor as does a Memory Management Unit which provides a full memory management facility through relocation and protection (described in Chapter 6).

The PDP-11/45 hardware has been optimized towards a multi-programming environment and the processor therefore operates in three modes (Kernel, Supervisor, and User) and has two sets of General Registers.

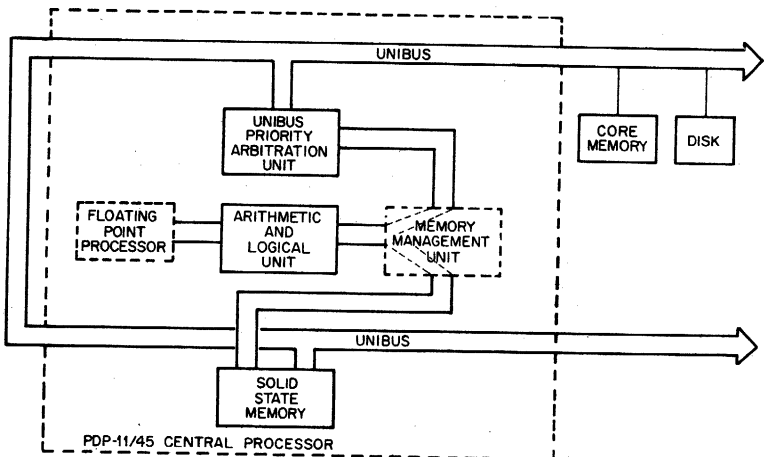


Figure 2-1 PDP-11/45 System Block Diagram

The PDP-11/45 communicates with its options through a bidirectional, asynchronous bus, the UNIBUS.

2.2 THE UNIBUS

All devices are connected through hardware registers to the UNIBUS.




Figure 2-2 The UNIBUS

Any device (except memory) can dynamically request the UNIBUS to transfer information to another using a scheme based on real and simulated core locations. All device registers are located in the uppermost 4K words of address space (124K-128K). Thus, the Central Processor can look on its peripherals as if they were locations in memory with special properties, and operate on them using the same set of instructions it uses to operate on memory.




Figure 2-3 Location of Device Registers

The UNIBUS provides the communications path for address, data, and control information for all devices on the bus through its bidirectional lines. Therefore the same device registers can be used for both input and output functions.

Devices communicate on the UNIBUS in a master-slave relationship. During any bus operation, one device has control of the bus at any given instant of time. This device, called the master, controls the bus when communicating with another, called the slave.

The relationship is dynamic, thus the Central Processor as master could send control information to a disk (slave) which then could obtain the bus as a master to communicate with memory, the slave.

The UNIBUS is used by the processor and all I/O devices. A priority structure determines which device has control of the bus at any given instant of time. Therefore, every device capable of becoming bus master has an assigned priority; and when two devices request the bus at the same time, the device with the higher priority will receive control first.

Communication on the UNIBUS is interlocked between devices. For each control signal issued by the master, there is a response from the slave; thus, communication is independent of physical bus length and the response time of the master and slave devices. The maximum transfer rate on the UNIBUS is one 16-bit word every 400 nsec or 2.5 million 16-bit

words per second. The UNIBUS is fully described in the PDP-11 Peripherals and Interfacing Handbook.

2.3 CENTRAL PROCESSOR

The PDP-11/45 performs all arithmetic and logical operations required in the system. It also acts as the arbitration unit for UNIBUS control by regulating bus requests and transferring control of the bus to the requesting device with the highest priority.

The central processor contains arithmetic and control logic for a wide range of operations. These include high-speed fixed point arithmetic with hardware multiply and divide, extensive test and branch operations, and other control operations. It also provides room for the addition of the high-speed Floating Point Processor, and Memory Management Unit.

The machine operates in three modes: Kernel, Supervisor, and User. When the machine is in Kernel mode a program has complete control of the machine; when the machine is in any other mode the processor is inhibited from executing certain instructions and can be denied direct access to the peripherals on the system. This hardware feature can be used to provide complete executive protection in a multi-programming environment.

The central processor contains 16 general registers which can be used as accumulators, index registers, or as stack pointers. Stacks are extremely useful for nesting programs, creating re-entrant coding, and as temporary storage where a Last-In First-Out structure is desirable. A special instruction "MARK" is provided to further facilitate re-entrant programming and is described in Chapter 5. One of the general registers is used as the PDP-11/45's program counter. Three others are used as Processor Stack Pointers, one for each operational mode.

The CPU is directly connected to the high-speed memories as well as to the general purpose registers and the UNIBUS and UNIBUS Priority Arbitration Unit.

Figure 2-4 illustrates the data paths in the CPU.

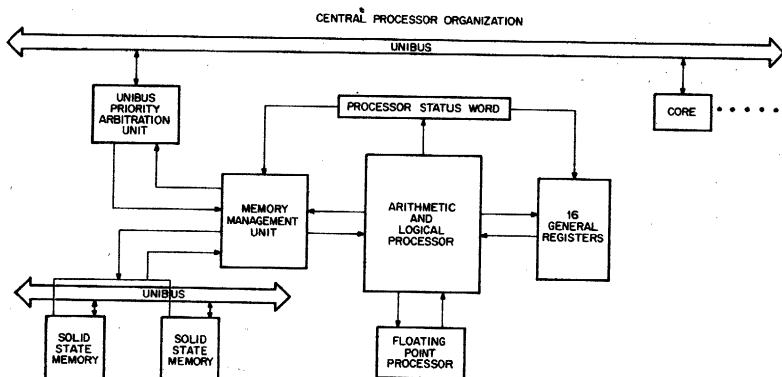


Figure 2-4 Central Processor Data Paths

The CPU performs all of the computer's computation and logic operations in a parallel binary mode through step by step execution of individual instructions. The instructions are stored in either core or solid state memory.

2.3.1 General Registers

The general registers can be used for a variety of purposes; the uses varying with requirements. The general registers can be used as accumulators, index registers, autoincrement registers, autodecrement registers, or as stack pointers for temporary storage of data. Chapter 3 on Addressing describes these uses of the general registers in more detail. Arithmetic operations can be from one general register to another, from one memory or device register to another, or between memory or a device register and a general register.

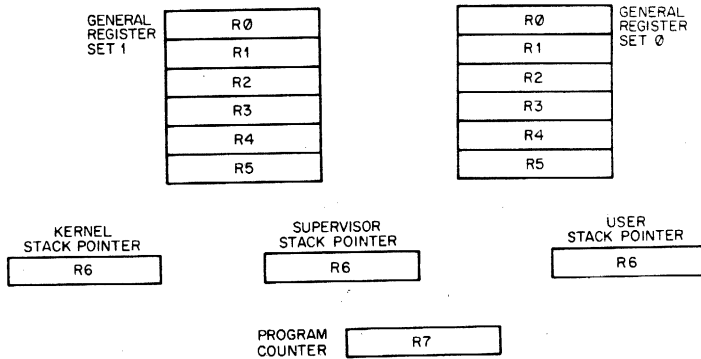


Figure 2-5 The General Registers

R7 is used as the machine's program counter (PC) and contains the address of the next instruction to be executed. It is a general register normally used only for addressing purposes and not as an accumulator for arithmetic operations.

The R6 register is normally used as the Processor Stack Pointer indicating the last entry in the appropriate stack (a common temporary storage area with "Last-In First-Out" characteristics). (For information on the programming uses of stacks, please refer to Chapter 5.) The three stacks are called the Kernel Stack, the Supervisor Stack, and the User Stack. When the Central Processor is operating in Kernel mode it uses the Kernel Stack, in Supervisor mode, the Supervisor Stack, and in User mode, the User Stack. When an interrupt or trap occurs, the PDP-11/45 automatically saves its current status on the Processor Stack selected by the service routine. This stack-based architecture facilitates reentrant programming.

The remaining 12 registers are divided into two sets of unrestricted registers, R0-R5. The current register set in operation is determined by the Processor Status Word.

The two sets of registers can be used to increase the speed of real-time data handling or facilitate multi-programming. The six registers in General Register Set 0 could each be used as an accumulator and/or index register for a real-time task or device, or as general registers for a Kernel or Supervisor mode program. General Register Set 1 could be used by the remaining programs or User mode programs. The Supervisor can therefore protect its general registers and stack from User programs, or other parts of the Supervisor.

2.3.2 Processor Status Word




Figure 2-6 Processor Status Word

The Processor Status Word, located at location 777776, contains information on the current status of the PDP-11/45. This information includes the register set currently in use; current processor priority; current and previous operational modes; the condition codes describing the results of the last instruction; and an indicator for detecting the execution of an instruction to be trapped during program debugging.

Modes

Mode information includes the present mode, either User, Supervisor, or Kernel (bits 15, 14); the mode the machine was in prior to the last interrupt or trap (bits 13, 12); and which register set (General Register Set 0 or 1) is currently being used (bit 11).

The three modes permit a fully protected environment for a multi-programming system by providing the user with three distinct sets of Processor Stacks and Memory Management Registers for memory mapping. In all modes except Kernel a program is inhibited from executing a "HALT" instruction and the processor will trap through location 4 if an attempt is made to execute this instruction. Furthermore, the processor will ignore the "RESET" and "SPL" instructions. In Kernel mode, the processor will execute all instructions.

A program operating in Kernel mode can map users' programs anywhere in core and thus explicitly protect key areas (including the devices registers and the Processor Status Word) from the User operating environment.

Processor Priority

The Central Processor operates at any of eight levels of priority, 0-7. When the CPU is operating at level 7 an external device cannot interrupt it with a request for service. The Central Processor might be operating at a lower priority than the priority of the external device's request in order for the interruption to take effect. The current priority is maintained in the processor status word (bits 5-7). The 8 processor levels provide an effective interrupt mask, which can be dynamically altered through use of the Set Priority Level (SPL) instruction which is described in Chapter 4 and which can only be used by the Kernel. This instruction allows a Kernel mode program to alter the Central Processor's priority without affecting the rest of the Processor Status Word.

Condition Codes

The condition codes contain information on the result of the last CPU operation. They include; a carry bit (C), which is set by the previous operation if the operation caused a carry out of its most significant bit; a negative bit (N) set if the result of the previous operation was negative; a zero bit (Z), set if the result of the previous operation was zero; and an overflow bit (V), set if the result of the previous operation resulted in an arithmetic overflow.

Trap

The trap bit (T) can be set or cleared under program control. When set, a processor trap will occur through location 14 on completion of instruction execution and a new Processor Status Word will be loaded. This bit is especially useful for debugging programs as it provides an efficient method of installing breakpoints.

Interrupts and trap instructions both automatically cause the previous Processor Status Word and Program Counter to be saved and replaced by the new values corresponding to those required by the routine servicing the interrupt or trap. The user can, thus, cause the central processor to automatically switch modes (context switching), register sets, alter the CPU's priority, or disable the Trap Bit whenever a trap or interrupt occurs.

2.3.3 Stack Limit Register

All PDP-11's have a Stack Overflow Boundary at location 400. The Kernel Stack Boundary, in the PDP-11/45 is a variable boundary set through the Stack Limit Register found in location 777774.

Once the Kernel stack exceeds its boundary, the Processor will complete the current instruction and then trap to location 4 (Yellow or Warning Stack Violation). If, for some reason, the program persists beyond the 16-word limit, the processor will abort the offending instruction, set the stack pointer (R6) to 4 and trap to location 4 (Red or Fatal Stack Violation). A description of these traps is contained in Appendix C.

2.4 FLOATING POINT PROCESSOR

The PDP-11/45 Floating Point Processor fits integrally into the Central Processor. It provides a supplemental instruction set for performing single and double precision floating point arithmetic operations and floating-integer conversions in parallel with the CPU. It is fully described in Chapter 7.

2.5 MEMORY

Memory is the primary storage medium for instructions and data. Three types are available for the PDP-11/45:

SOLID STATE:

Bipolar Memory with a cycle time of 300 nsec.

MOS Memory with a cycle time of 450 nsec.

CORE:

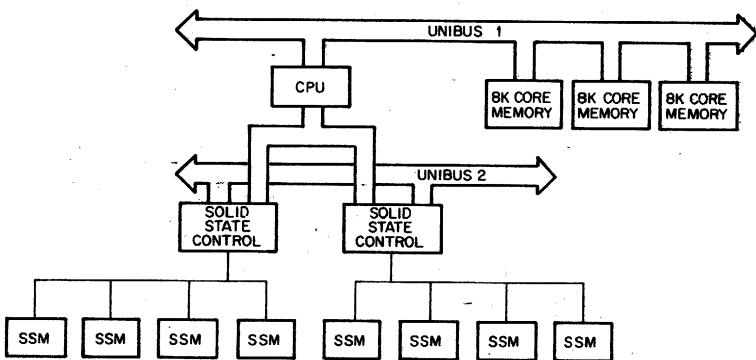
Magnetic Core Memory with a cycle time of 850 ns, access at 350 ns (450 ns at the UNIBUS).

Any system can be expanded from the basic 4,096 words to 126,976 words in increments of 4,096 words. The system can be configured with various mixtures of the three types of memory, with a maximum limit of 32,768 words of Solid State Memory.

2.5.1 Solid State Memory

The Central Processor communicates directly with the MOS and Bipolar memories through a very high speed data path which is internal to the PDP-11/45 processor system. The CPU can control up to two independent Solid State Memory controllers, each of which can have from one to four 4,096 word increments of MOS memory (16,384 words) per controller, or one 4,096 word increment of Bipolar memory per controller. Each controller can handle MOS or Bipolar memory but not a mixture of the two. The user can therefore have a total of 32K of MOS, or 8K of Bipolar, or 16K of MOS and 4K of Bipolar.

Each controller has dual ports and provides one interface to the CPU and another to a second UNIBUS.



SSM= SOLID STATE MEMORY MATRIX (4K MOS OR 1K BIPOLAR)

Figure 2-7 Memory Configuration

There are two UNIBUSes on the PDP-11/45 but in a single processor environment the second UNIBUS is generally connected into the first

and become part of it. The existence of a second UNIBUS becomes significant where a high speed device would like to directly access the solid state memory. A device using the second UNIBUS must include a UNIBUS Priority Arbitration Unit, and the bus thus lends itself to multi-processor environments.

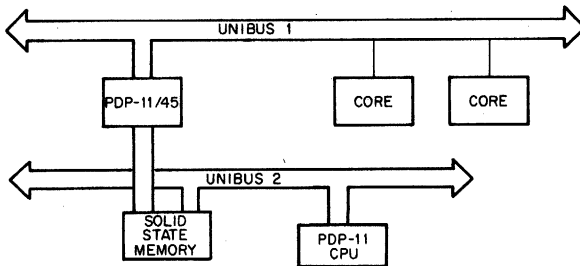


Figure 2-8 Multiprocessor Use of the Second UNIBUS

The UNIBUS and data path to the Solid State Memory are independent. While the Central Processor is operating on data in one Solid State Memory controller through the direct data path, any device could be using the UNIBUS to transfer information to core, to another device, or to the other Solid State Memory Controller. This autonomy significantly increases the throughput of the system.

2.5.2 Memory Retention

MOS memory bits have a capacitance which is charged to denote a 1 and uncharged to denote a 0. The entire MOS memory must be refreshed periodically, or the data will be lost. On the PDP-11/45, 1/32nd of the memory is refreshed every 60 microseconds. This process consumes only one solid state memory cycle.

The power required to refresh MOS memory is significantly less than that required for operation of the memory. Bipolar memory, on the other hand, does not require a refresh cycle but does require the same power to retain information as to operate.

2.5.3 Core Memory

The Central Processor communicates with core memory through the UNIBUS.

Each memory bank operates independently from other banks through its own controller which interfaces directly to the UNIBUS. Core memory can be continuously attached to the UNIBUS until the system contains a total of 124K (126,976 words) of memory.

An external device may use the UNIBUS to read or write core memory completely independent of, and simultaneously with the Central Processor's access of solid state memory. Furthermore, core memory and solid state memory may be used by the processor interchangeably.

2.5.4 Memory Interleaving

Generally, memory locations are numbered consecutively in a memory

bank. Thus, when the address register is incremented on successive memory cycles the same bank is addressed and a full memory cycle must be completed before the new address can be used. The maximum data transfer rate for a device using memory is therefore limited by memory cycle time. Memory interleaving is a technique that places consecutive word addresses in different memory banks and thus allows the write cycle in one memory to be overlapped with the read cycle in another.

The PDP-11/45's architecture with a controller inherent to each memory bank lends itself to memory interleaving, and mixing memories with different cycle times. The standard core memory has a basic cycle time of 850 nsec and is interleaved in pairs of 8,192 word units. A 16K system would be fully interleaved, whereas a 24K system (3 controllers) would only have 16K interleaved. Interleaving this memory provides the user with an effective cycle time of 650 nsec for consecutive accesses to sequential word locations.

The MOS memories are interleaved on an equal number of MOS blocks on each controller. Bipolar is not interleaved.

Memory interleaving is completely transparent to the computer programmer.

2.5.5 Mixing Memory

The PDP-11/45 can be used with the memory mix best suited to a user's needs. He can not only mix high speed solid state memories with magnetic core memories, but he can also choose core memories of different speeds. This is possible because of the independent nature of the core memory controllers.

The PDP-11/45 provides the user with an additional degree of freedom in mixing memories. The programmer need not address all of his solid state memory consecutively, but can intermix solid state and core physical addresses. Each solid state memory can address a 16K consecutive segment (32K when MOS is interleaved) beginning on a 16K boundary. If the controller contains a full 16K complement of MOS then the MOS will use up the full 16K address space; however, if there is less than 16K of MOS on the controller the user can intermix 4K blocks of core in with the 4K blocks of MOS.

When a program is using Memory Management (Chapter 6) this manipulation of physical addresses is unnecessary as it may be done in the mapping of virtual space into physical space.

The user can reduce the cost of his system by buying only as much high speed memory as required; and he can increase system performance through the independence of data transfers on the UNIBUS and the CPU connection to the Solid State Memories.

2.5.6 Memory Parity

Memory Parity is optional for core and solid state memory. Parity words are extended to 18 bits and the last two bits (17, 18) contain the parity indicators for the two bytes. Parity is generated when a word is written and checked when the word is read. Parity errors cause the Central Pro-

cessor either to trap through location 4 or to halt. The user is referred to Appendix E for more information on memory parity.

2.5.7 Memory Organization

PDP-11 memories and instructions are designed to handle both 16-bit words, and 8-bit bytes. Therefore, since every word contains two bytes, a 4,096 word block contains 8,192 byte locations. Consecutive words are therefore found in even numbered addresses.




Figure 2-9 Memory Organization

Certain memory locations have been reserved by the system for interrupt and trap handling, and processor stacks, general registers, and peripheral device registers. Kernel virtual addresses from 0 to 370₈ are always reserved and those to 777₈ are reserved on large system configurations for traps and interrupt handling. The top 4,096 word addresses from 770000₈ up) have been reserved for general registers and peripheral devices. Appendix C presents a detailed memory address map.

A 16-bit word used for byte addressing can address a maximum of 32K words. However, the top 4,096 word locations are traditionally reserved for peripheral and register addresses and the user therefore has 28K of core to program. To expand above 28K the user must use the Memory Management Unit. This device, described in detail in Chapter 6, provides the programmer with an 18-bit effective memory address which permits him to address up to 126,976 words of actual memory. The unit also provides a memory management facility which permits individual user programs up to 64K in length (32K of instructions, 32K of data) and provides a relocation and protection facility through three sets of 16 registers.

2.6 SYSTEM INTERACTION

System intercommunication is carried out through the UNIBUS.

A device will request the UNIBUS for one of two purposes:

To make a non-processor (NPR) transfer of data. (Direct Data Transfers such as DMA), or

To interrupt program execution and force the processor to branch to a service routine.

There are two sources of interrupts, hardware and software.

2.6.1 Hardware Interrupt Requests

A hardware interrupt occurs when a device wishes to indicate to the program, or Central Processor, that a condition has occurred (such as transfer completed, end of tape, etc.). The interrupt can occur on any one of the four Bus Request levels and the processor responds to the interrupt through a service routine.

2.6.2 Program Interrupt Requests

Hardware interrupt servicing is often a two-level process. The first level is directly associated with the device's hardware interrupt and consists of retrieving the data. The second, is a software task that manipulates the raw information. The second process can be run at a lower priority than the first, because the PDP-11/45 provides the user with the means of scheduling his software servicing through seven levels of Program Interrupt Requests. The Program Interrupt Request Register is located at address 777772. An interrupt is generated by the programmer setting a bit in the high order byte of this register.

The reader is referred to Appendix D for more detailed information.

2.6.3 Priority Structure

When a device capable of becoming bus master requests use of the bus, handling of the request depends on the hierarchical position of that device in the priority structure.

The relative priority of the request is determined by the Processor's priority and the level at which the request is made.

The processor's priority is set under program control to one of eight levels using bits 7-5 in the processor Status Word. Bus requests are inhibited on the same or lower levels.

Bus requests from external devices can be made on any one of the five request lines. A non-processor request (NPR) has the highest priority, and its request is granted between bus cycles of an instruction execution. Bus Request 7 (BR 7) is the next highest priority and Bus Request 4 (BR 4) is the lowest. The four lower priority level requests (BR 7-BR 4) are granted by the processor between instructions providing that they occur on higher levels than the processor's. Therefore an interrupt may only occur on a Bus Request Level and not on a Non Processor Request level.

Any number of devices can be connected to a specific BR or NPR line.

If two devices with the same priority request the bus, the device physically closest to the processor on the UNIBUS has the higher priority.

Program Interrupt Requests can be made on any one of 7 levels (PIR 7-PIR 1). Requests are granted by the processor between instructions providing that they occur on higher levels than the processor's.

Program Interrupt Requests take precedence over equivalent level Bus Requests.

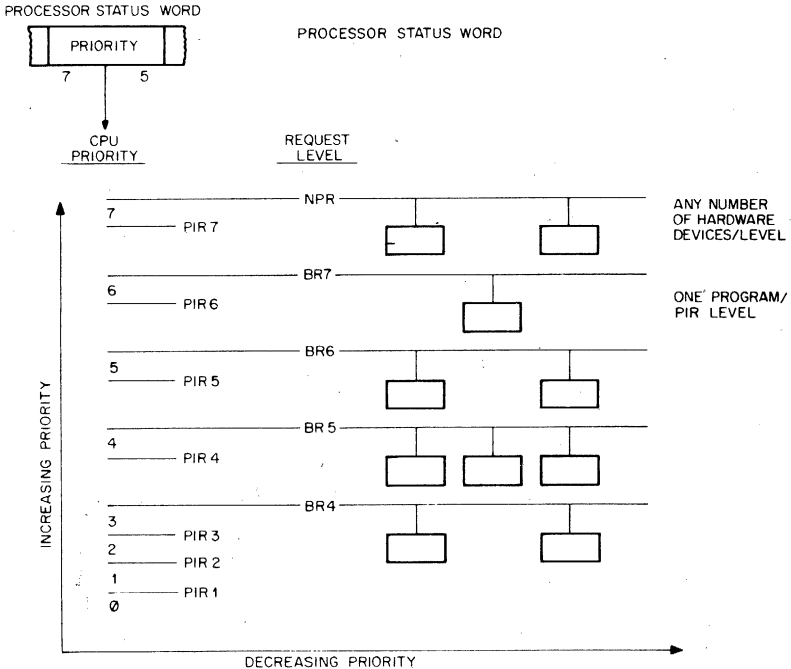


Figure 2-10 UNIBUS Priority Structure

2.6.4 Non-Processor Data Transfers

Direct memory or direct data transfers can be accomplished between any two peripherals without processor supervision. These Non-Processor transfers, called NPR level data transfers, are usually made for Direct Memory Access (memory to/from mass storage) or direct device transfers (disk refreshing a CRT display).

A NPR device provides extremely fast access to the UNIBUS and can transfer data at high rates once it gains control of the bus. The state of the processor is not affected by this type of transfer, and, therefore, the processor can relinquish bus control while an instruction is still in progress. The bus can be released at the end of any bus cycle, except during a read-modify-write cycle sequence. (This occurs for example in destructive read-out devices such as core memory for certain instructions.)

In the PDP-11/45 an NPR device can gain bus control in 3.5 microseconds or less (depending on the number of devices on the UNIBUS), and can transfer 16-bit words to memory at the same speed as the effective cycle time of the memory being addressed.

2.6.5 Using the Interrupts

Devices that gain bus control with one of the Bus Request Lines (BR 7-BR 4), can take full advantage of the Central Processor by requesting an interrupt. In this way, the entire instruction set is available for manipulating data and status registers.

When a service routine is to be run, the current task being performed by the central processor is interrupted, and the device service routine is initiated. Once the request has been satisfied, the Processor returns to its former task. Interrupts may also be used to schedule program execution by using the Program Interrupt Request.

2.6.6 Interrupt Procedure

Interrupt handling is automatic in the PDP-11/45. No device polling is required to determine which service routine to execute. The operations required to service an interrupt are as follows:

1. Processor relinquishes control of the bus, priorities permitting.
2. When a master gains control, it sends the processor an interrupt command and a unique memory address which contains the address of the device's service routine in Kernel virtual address space, called the interrupt vector address. Immediately following this pointer address is a word (located at vector address +2) which is to be used as a new Processor Status Word.
3. The processor stores the current Processor Status Word (PS) and the current Program Counter (PC) into CPU temporary registers.
4. The new PC and PS (the interrupt vector) are taken from the specified address. The old PS and PC are then pushed onto the current stack as indicated by bits 15,14 of the new PS and the previous mode in effect is stored in bits 13,12 of the new PS. The service routine is then initiated.

These operations are performed in 2.5 μ sec from the time the control processor receives the interrupt command until the time it starts executing the first instruction of the service routine. This time interval assumes no NPR transfer occurred during this time interval.

5. The device service routine can cause the processor to resume the interrupted process by executing the Return from Interrupt (RTI or RTT) instruction, described in Chapter 4, which pops the two top words from the current processor stack and uses them to load the PC and PS registers.

This instruction requires 1.5 μ sec providing there is no NPR request.

A device routine can be interrupted by a higher priority bus request any time after the new PC and PS have been loaded. If such an interrupt occurs, the PC and the PS of the service routine are automatically stored

in the temporary registers and then pushed onto the new current stack, and the new device routine is initiated.

2.6.7 Interrupt Servicing

Every hardware device capable of interrupting the processor has a unique set of locations (2 words) reserved for its interrupt vector. The first word contains the location of the device's service routine, and the second, the Processor Status Word that is to be used by the service routine. Through proper use of the PS, the programmer can switch the operational mode of the processor, alter the General Register Set in use (context switching), and modify the Processor's Priority level to mask out lower level interrupts.

There is one interrupt vector for the Program Interrupt Request. It will generally be necessary in a multi-processing environment to determine which program generated the PIR and where it is located in memory. Appendix D provides an example of how this is done.

2.7 PROCESSOR TRAPS

There are a series of errors and programming conditions which will cause the Central Processor to trap to a set of fixed locations. These include Power Failure, Odd Addressing Errors, Stack Errors, Timeout Errors, Memory Parity Errors, Memory Management Violations, Floating Point Processor Exception Traps, Use of Reserved Instructions, Use of the T bit in the Processor Status Word, and use of the IOT, EMT, and TRAP instructions.

Stack Errors, Memory Parity Errors, and the T bit Trap have already been discussed in this chapter. Segmentation Violations and Floating Point Exception Traps are described in Chapters 6 and 7 respectively. The IOT, EMT, and TRAP instructions are described in Chapter 4.

2.7.1 Power Failure

Whenever AC power drops below 95 volts for 110v power (190 volts for 220v) or outside a limit of 47 to 63 Hz, as measured by DC power, the power fail sequence is initiated. The Central Processor automatically traps to location 24 and the power fail program has 2 msec. to save all volatile information (data in registers), and to condition peripherals for power fail.

When power is restored the processor traps to location 24 and executes the power up routine to restore the machine to its state prior to power failure.

2.7.2 Odd Addressing Errors

This error occurs whenever a program attempts to execute a word instruction on an odd address (in the middle of a word boundary). The instruction is aborted and the CPU traps through location 4.

2.7.3 Time-out Errors

These errors occur when a Master Synchronization pulse is placed on the UNIBUS and there is no slave pulse within 5 μ sec. This error usually occurs in attempts to address non-existent memory or peripherals.

The offending instruction is aborted and the processor traps through location 4.

2.7.4 Reserved Instructions

There is a set of illegal and reserved instructions which cause the processor to trap through Location 10. The set is fully described in Appendix C.

2.7.5 Trap Handling

Appendix C includes a list of the reserved Trap Vector locations, and System Error Definitions which cause processor traps. When a trap occurs, the processor follows the same procedure for traps as it does for interrupts (saving the PC and PS on the new Processor Stack etc....).

In cases where traps and interrupts occur concurrently, the processor will service the conditions according to the priority sequence illustrated in Figure 2-11.

- Odd Addressing Error
- Fatal Stack Violations (Red)
- Memory Management Violations
- Timeout Errors
- Parity Errors
- Floating Point Processor Transfer Request
- Memory Management Traps
- Warning Stack Violation (Yellow)
- Power Failure
- Processor Priority level 7
- Floating Point Exception Trap
- PIR 7
- BR 7
- .
- .
- .
- .
- PIR 1
- Processor 0

Figure 2-11 Processor Service Hierarchy

Appendix C includes more details on the Trap sequence and Trap/Interrupt interaction.

2.8 MULTIPROGRAMMING

The PDP-11/45's architecture with its three modes of operation, its two sets of general registers, its Memory Management capability and its Program Interrupt Request facility provides an ideal environment for multi-programming systems.

In any multi-programming system there must be some method of transferring information and control between programs operating in the same or different modes. The PDP-11/45 provides the user with these communication paths.

2.8.1 Control Information

Control is passed inwards (User, Supervisor, Kernel) by all traps and interrupts. All trap and interrupt vectors are located in Kernel virtual space. Thus all traps and interrupts pass through Kernel space to pick up their new PC and PS and determine the new mode of processing.

Control is passed outwards (Kernel, Supervisor, User) by the RTI and RTT instructions (described in Chapter 4).

2.8.2 Data

Data is transferred between modes by four instructions Move From Previous Instruction space (MFPI), Move From Previous Data space (MFPD), Move To Previous Instruction space (MTPI) and Move To Previous Data space (MTPD). There are four instructions rather than two as Memory Management distinguishes between instructions and data (Chapter 6). The instructions are fully described in Chapter 4. However, it should be noted that these instructions have been designed to allow data transfers to be under the control of the innermost mode (Kernel, Supervisor, User) program and not the outermost, thus providing protection of an inner program from an outer.

2.8.3 Processor Status Word

The PDP 11/45 protects the PS from implicit references by Supervisor and User programs which could result in damage to an inner level program.

A program operating in Kernel mode can perform any manipulation of the PS. Programs operating at outer levels (Supervisor and User) are inhibited from changing bits 5-7 (the Processor's Priority). They are also restricted in their treatment of bits 15, 14 (Current Mode), bits 13, 12 (Previous Mode), and bit 11 Register Set); these bits may only be set, they are only cleared by an interrupt or trap.

Thus, a programmer can pass control outwards through the RTI and RTT instructions to set bits in the mode fields of his PS. To move inwards, however, bits must be cleared and he must, therefore, issue a trap or interrupt.

The Kernel can further protect the PS from explicit references (Move data to location 777776—the PS) through Memory Management.

ADDRESSING MODES

Data stored in memory must be accessed, and manipulated. Data handling is specified by a PDP-11 instruction (MOV, ADD etc.) which usually indicates:

- the function (operation code);

- a general purpose register to be used when locating the source operand and/or a general purpose register to be used when locating the destination operand;

- an addressing mode (to specify how the selected register(s) is/are to be used.

Since a large portion of the data handled by a computer is usually structured (in character strings, in arrays, in lists etc.), the PDP-11 has been designed to handle structured data efficiently and flexibly. The general registers may be used with an instruction in any of the following ways:

- as accumulators. The data to be manipulated resides within the register.

- as pointers. The contents of the register are the address of the operand, rather than the operand itself.

- as pointers which automatically step through core locations. Automatically stepping forward through consecutive core locations is known as autoincrement addressing; automatically stepping backwards is known as autodecrement addressing. These modes are particularly useful for processing tabular data.

- as index registers. In this instance the contents of the register, and the word following the instruction are summed to produce the address of the operand. This allows easy access to variable entries in a list.

PDP-11's also have instruction addressing mode combinations which facilitate temporary data storage structures for convenient handling of data which must be frequently accessed. This is known as the "stack." (see Chapter 5)

In the PDP-11 any register can be used as a "stack pointer" under program control; however, certain instructions associated with subroutine linkage and interrupt service automatically use Register 6 as a "hardware stack pointer." For this reason R6 is frequently referred to as the "SP."

R7 is used by the processor as its program counter (PC). It is recommended that R7 not be used as a stack pointer.

An important PDP-11/45 feature, which must be considered in conjunction with the addressing modes, is the register arrangement;


- Two sets of general purpose registers (R0-R5)
- three hardware stack pointers (R6)
- a single Program Counter (PC) register (R7).

Register R7 is used as a common program counter (PC). At any point in time only one register set is active. Thus a programmer need only concern himself with the existence of multiple register sets for those special supervisory tasks which involve Kernel, Supervisor, User communications (e.g. MTPX, MFPX); otherwise he need never worry about which R3 or R6 an instruction will reference, the choice is automatic and transparent to his program.

Instruction mnemonics and address mode symbols are sufficient for writing machine language programs. The programmer need not be concerned about conversion to binary digits; this is accomplished automatically by the PDP-11/45 assembler.

3.1 SINGLE OPERAND ADDRESSING

The instruction format for all single operand instructions such as clear, increment, test) is:



Bits 15 through 6 specify the operation code that defines the type of instruction to be executed.


Bits 5 through 0 form a six-bit field called the destination address field. This consists of two subfields:

- a) Bits 0 through 2 specify which of the eight general purpose registers is to be referenced by this instruction word.
- b) Bits 4 and 5 specify how the selected register will be used (address mode). Bit 3 indicates direct or deferred (indirect) addressing.

3.2 DOUBLE OPERAND ADDRESSING

Operations which imply two operands (such as add, subtract, move and compare) are handled by instructions that specify two addresses. The

first operand is called the source operand, the second the destination operand. Bit assignments in the source and destination address fields may specify different modes and different registers. The Instruction format for the double operand instruction is:



The source address field is used to select the source operand, the first operand. The destination is used similarly, and locates the second operand and the result. For example, the instruction ADD A,B adds the contents (source operand) of location A to the contents (destination operand) of location B. After execution B will contain the result of the addition and the contents of A will be unchanged.

Examples in this section and further in this chapter use the following sample PDP-11 instructions:

Mnemonic	Description	Octal Code
CLR	clear (zero the specified destination)	0050nn
CLRB	clear byte (zero the byte in the specified destination)	1050nn
INC	increment (add 1 to contents of destination)	0052nn
INCB	increment byte (add 1 to the contents of destination byte)	1052nn
COM	complement (replace the contents of the destination by their logical complement; each 0 bit is set and each 1 bit is cleared)	0051nn
COMB	complement byte (replace the contents of the destination byte by their logical complement; each 0 bit is set and each 1 bit is cleared).	1051nn
ADD	add (add source operand to destination operand and store the result at destination address)	06mmnn

3.3 DIRECT ADDRESSING

The following table summarizes the four basic modes used with direct addressing.

DIRECT MODES

Binary	Name	Assembler Syntax	Function
0 0 0	Register	Rn	Register contains operand
0 1 0	Autoincrement	(Rn)+	Register is used as a pointer to sequential data then incremented.
1 0 0	Autodecrement	-(Rn)	Register is decremented and then used as a pointer.
1 1 0	Index	X(Rn)	Value X is added to (Rn) to produce address of operand. Neither X nor (Rn) are modified.

3.3.1 Register Mode

OPR Rn

With register mode any of the general registers may be used as simple accumulators and the operand is contained in the selected register. Since they are hardware registers, within the processor, the general registers operate at high speeds and provide speed advantages when used for operating on frequently-accessed variables. The PDP-11 assembler interprets and assembles instructions of the form OPR Rn as register mode operations. Rn represents a general register name or number and OPR is used to represent a general instruction mnemonic. Assembler syntax requires that a general register be defined as follows:

R0 = %0 (% sign indicates register definition)

R1 = %1

R2 = %2, etc.


Registers are typically referred to by name as R0, R1, R2, R3, R4, R5, R6 and R7. However R6 and R7 are also referred to as SP and PC, respectively.

Register Mode Examples

(all numbers in octal)


	Symbolic	Octal Code	Instruction Name
1.	INC R3	005203	Increment

Operation: Add one to the contents of general register 3




2. **ADD R2,R4** 060204 Add

Operation: Add the contents of R2 to the contents of R4.



3. **COMB R4** 105104 Complement Byte

Operation: One's complement bits 0-7 (byte) in R4.
 (When general registers are used, byte instructions only operate on bits 0-7; i.e. byte 0 of the register)



3.3.2 Autoincrement Mode

OPR (Rn)+


This mode provides for automatic stepping of a pointer through sequential elements of a table of operands. It assumes the contents of the selected general register to be the address of the operand. Contents of registers are stepped (by one for bytes, by two for words, always by two for R6 and R7) to address the next sequential location. The autoincrement mode is especially useful for array processing and stacks. It will access an element of a table and then step the pointer to address the next operand in the table. Although most useful for table handling, this mode is completely general and may be used for a variety of purposes.

Autoincrement Mode Examples

Symbolic Octal Code Instruction Name


1. **CLR (R5)+** 005025 Clear

Operation: Use contents of R5 as the address of the operand. Clear selected operand and then increment the contents of R5 by two.




2. **CLRB (R5)+** 105025 Clear Byte

Operation: Use contents of R5 as the address of the operand. Clear selected byte operand and then increment the contents of R5 by one.



3. **ADD (R2)+, R4 062204 Add**

Operation: The contents of R2 are used as the address of the operand which is added to the contents of R4. R2 is then incremented by two.



3.3.3 Autodecrement Mode

OPR-(Rn)


This mode is useful for processing data in a list in reverse direction. The contents of the selected general register are decremented (by two for word instructions, by one for byte instructions) and then used as the address of the operand. The choice of postincrement, predecrement features for the PDP-11 were not arbitrary decisions, but were intended to facilitate hardware/software stack operations (See Chapter 5 for complete discussions of stacks).

Autodecrement Mode Examples

Symbolic Octal Code Instruction Name


1. **INC-(R0)** 005240 Increment

Operation: The contents of R0 are decremented by two and used as the address of the operand. The operand is increased by one.



2. **INCB-(R0)** 105240 Increment Byte


Operation: The contents of R0 are decremented by one then used as the address of the operand. The operand byte is increased by one.



3. **ADD—(R3),R0** 064300 Add

Operation:

The contents of R3 are decremented by 2 then used as a pointer to an operand (source) which is added to the contents of R0 (destination operand).



3.3.4 Index Mode

OPR X(Rn)

The contents of the selected general register, and an index word following the instruction word, are summed to form the address of the operand. The contents of the selected register may be used as a base for calculating a series of addresses, thus allowing random access to elements of data structures. The selected register can then be modified by program to access data in the table. Index addressing instructions are of the form OPR X(Rn) where X is the indexed word and is located in the memory location following the instruction word and Rn is the selected general register.

Index Mode Examples

Symbolic

Octal Code

Instruction Name

1. CLR 200(R4)


005064

Clear

000200

Operation:


The address of the operand is determined by adding 200 to the contents of R4. The location is then cleared.



2. COMB 200(R1) 105161 Complement Byte
000200

Operation:

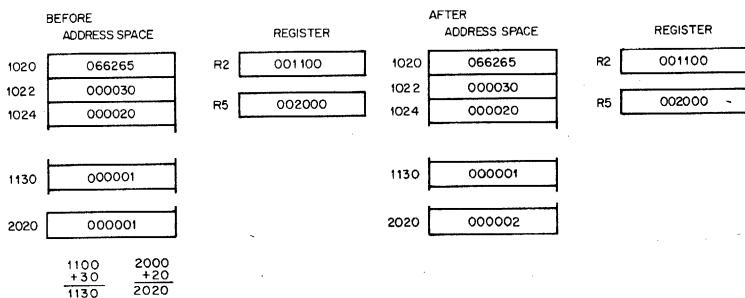
The contents of a location which is determined by adding 200 to the contents of R1 are one's complemented (i.e. logically complemented).



3. ADD 30(R2), 20(R5) 066265 Add
000030
000020

Operation:

The contents of a location which is determined by adding 30 to the contents of R2 are added to the contents of a location which is determined by adding 20 to the contents of R5. The result is stored at the destination address, i.e. 20(R5)



3.4 DEFERRED (INDIRECT) ADDRESSING

The four basic modes may also be used with deferred addressing. Whereas in the register mode the operand is the contents of the selected register, in the register deferred mode the contents of the selected register is the address of the operand.

In the three other deferred modes, the contents of the register selects the address of the operand rather than the operand itself. These modes are therefore used when a table consists of addresses rather than operands. Assembler syntax for indicating deferred addressing is "@". The following table summarizes the deferred versions of the basic modes:


Binary Code	Name	Assembler Syntax	Function
0 0 1	Register Deferred	@Rn or (Rn)	Register contains the address of the operand
0 1 1	Autoincrement Deferred	@(Rn)+	Register is first used as a pointer to a word containing the address of the operand, then incremented (always by 2; even for byte instructions)
1 0 1	Autodecrement Deferred	@-(Rn)	Register is decremented (always by two; even for byte instructions) and then used as a pointer to a word containing the address of the operand
1 1 1	Index Deferred	@X(Rn)	Value X (stored in a word following the instruction) and (Rn) are added and the sum is used as a pointer to a word containing the address of the operand. Neither X nor (Rn) are modified.

Since each deferred mode is similar to its basic mode counterpart, separate descriptions of each deferred mode are not necessary. However, the following examples illustrate the deferred modes.

Register Deferred Mode Example

Symbolic	Octal Code	Instruction Name
CLR @R5	005015	Clear


Operation: The contents of location specified in R5 are cleared.



Autoincrement Deferred Mode Example

Symbolic	Octal Code	Instruction Name
INC @(R2)+	005232	Increment

Operation: The contents of R2 are used as the address of the address of the operand. Operand is increased by one. Contents of R2 is incremented by 2.




Autodecrement Deferred Mode Example

Symbolic	Octal Code	Complement
COM @—(R0)	005150	

Operation:

The contents of R0 are decremented by two and then used as the address of the address of the operand. Operand is one's complemented. (i.e. logically complemented)

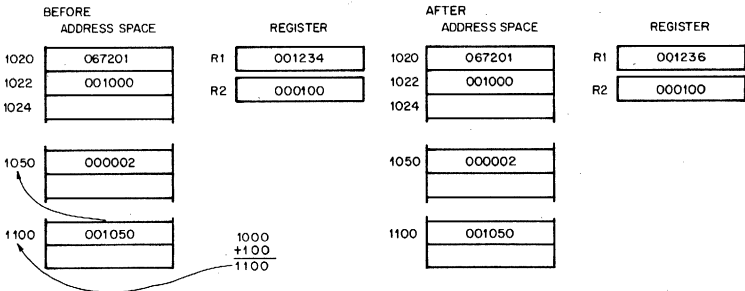


Index Deferred Mode Example

Symbolic	Octal Code	Instruction Name
ADD @1000(R2),R1	067201 001000	Add

Operation:

1000 and contents of R2 are summed to produce the address of the address of the source operand the contents of which are added to contents of R1; the result is stored in R1.



3.5 USE OF THE PC AS A GENERAL REGISTER

Although Register 7 is a general purpose register, it doubles in function as the Program Counter for the PDP-11. Whenever the processor uses the program counter to acquire a word from memory, the program counter is automatically incremented by two to contain the address of the next word of the instruction being executed or the address of the next instruction to be executed. (When the program uses the PC to locate byte data, the PC is still incremented by two.)

The PC responds to all the standard PDP-11 addressing modes. However, there are four of these modes with which the PC can provide advantages for handling position independent code (PIC—see Chapter 5) and unstructured data. When regarding the PC these modes are termed immediate, absolute (or immediate deferred), relative and relative deferred, and are summarized below:

Binary Code	Name	Assembler Syntax	Function
0 1 0	Immediate	# n	Operand follows instruction.
0 1 1	Absolute	@ # A	Absolute Address follows instruction.
1 1 0	Relative	A	Address of A, relative to the instruction, follows the instruction.
1 1 1	Relative Deferred	@A	Address of location containing address of A, relative to the instruction follows the instruction.

The reader should remember that the special effect modes are the same as modes described in 3.3 and 3.4, but the general register selected is R7, the program counter.

When a standard program is available for different users, it often is helpful to be able to load it into different areas of core and run it there. PDP-11's can accomplish the relocation of a program very efficiently through the use of position independent code (PIC) which is written by using the PC addressing modes. If an instruction and its objects are moved in such a way that the relative distance between them is not altered, the same offset relative to the PC can be used in all positions in memory. Thus, PIC usually references locations relative to the current location. PIC is discussed in more detail in Chapter 5.

The PC also greatly facilitates the handling of unstructured data. This is particularly true of the immediate and relative modes which are discussed more fully in Paragraphs 3.5.1 and 3.5.2.

3.5.1 Immediate Mode

OPR #n,DD

Immediate mode is equivalent to using the autoincrement mode with the PC. It provides time improvements for accessing constant operands by


including the constant in the memory location immediately following the instruction word.

Immediate Mode Example

Symbolic	Octal Code	Instruction Name
ADD # 10,R0	062700 000010	Add

Operation:

The value 10 is located in the second word of the instruction and is added to the contents of R0. Just before this instruction is fetched and executed, the PC points to the first word of the instruction. The processor fetches the first word and increments the PC by two. The source operand mode is 27 (autoincrement the PC). Thus, the PC is used as a pointer to fetch the operand (the second word of the instruction) before being incremented by two to point to the next instruction.



3.5.2 Absolute Addressing

OPR @ # A


This mode is the equivalent of immediate deferred or autoincrement deferred using the PC. The contents of the location following the instruction are taken as the address of the operand. Immediate data is interpreted as an absolute address (i.e., an address that remains constant no matter where in memory the assembled instruction is executed).

Absolute Mode Examples

	Symbolic	Octal Code	Instruction Name
1.	CLR @ #1100	005037 001100	Clear

Operation:


Clear the contents of location 1100.



2. ADD @ # 2000, R3 063703
 002000

Operation:

Add contents of location 2000 to R3.



3.5.3 Relative Addressing

OPR A or

OPR X(PC), where X is the location of A relative to the instruction.


This mode is assembled as index mode using R7. The base of the address calculation, which is stored in the second or third word of the instruction, is not the address of the operand, but the number which, when added to the (PC), becomes the address of the operand. This mode is useful for writing position independent code (see Chapter 5) since the location referenced is always fixed relative to the PC. When instructions are to be relocated, the operand is moved by the same amount.

Relative Addressing Example

Symbolic	Octal Code	Instruction Name
INC A	005267 000054	Increment

Operation:

To increment location A, contents of memory location immediately following instruction word are added to (PC) to produce address A. Contents of A are increased by one.



3.5.4 Relative Deferred Addressing

OPR@ or


OPR@X(PC), where x is location containing address of A, relative to the instruction.

This mode is similar to the relative mode, except that the second word of the instruction, when added to the PC, contains the address of the address of the operand, rather than the address of the operand.

Relative Deferred Mode Example

Symbolic	Octal Code	Instruction Name
CLR @A	005077 000020 000020	Clear

Operation: Add second word of instruction to PC to produce address of address of operand. Clear operand.



3.6 USE OF STACK POINTER AS GENERAL REGISTER

The processor stack pointer (SP, Register 6) is in most cases the general register used for the stack operations related to program nesting. Autodecrement with Register 6 "pushes" data on to the stack and autoincrement with Register 6 "pops" data off the stack. Index mode with the SP permits random access of items on the stack. Since the SP is used by the processor for interrupt handling, it has a special attribute: autoincrements and autodecrements are always done in steps of two. Byte operations using the SP in this way simply leave odd addresses unmodified. Use of stacks is explained in detail in Chapter 5.

On the PDP-11/45 there are three R6 registers selected by the PS; but at any given time there is only one in operation.

The following table is a concise summary of the various PDP-11 addressing modes

DIRECT MODES			
Binary Code	Name	Assembler Syntax	Function
000	Register	Rn	Register contains operand.
010	Autoincrement	(Rn) +	Register contains address of operand. Register contents incremented after reference.
100	Autodecrement	-(Rn)	Register contents decremented before reference register contains address of operand.
110	Index	X(Rn)	Value X (stored in a word following the instruction) is added to (Rn) to produce address of operand. Neither X nor (Rn) are modified.

DEFERRED MODES

Binary Code	Name	Assembler Syntax	Function
001	Register Deferred	@Rn or (Rn)	Register contains the address of the operand
011	Autoincrement Deferred	@(Rn) +	Register is first used as a pointer to A word containing the address of the operand, then incremented (always by 2; even for byte instructions)
101	Autodecrement Deferred	@-(Rn)	Register is decremented (always by two; even for byte instructions) and then used as a pointer to a word containing the address of the operand
111	Index Deferred	@X(Rn)	Value X (stored in a word following the instruction) and (Rn) are added and the sum is used as a pointer to a word containing the address of the operand. Neither X nor (Rn) are modified

PC ADDRESSING

010	Immediate	#n	Operand follows instruction
011	Absolute	@#A	Absolute address follows instruction
110	Relative	A	Address of A, relative to the instruction, follows the instruction.
111	Relative Deferred	@A	Address of location containing address of A, relative to the instruction follows the instruction.

INSTRUCTION SET

4.1 INTRODUCTION

This chapter describes the PDP-11/45 instructions in the following order:

Single Operand (4.4)

General

Shifts

Multiple Precision

Rotates

Double Operand (4.5)

Arithmetic Instructions

General Register Destination

Logical Instructions

Program Control Instructions (4.6)

Branches

Subroutines

Traps

Miscellaneous (4.7)**Condition Code Operators (4.8)**

The specification for each instruction includes the mnemonic, octal code, binary code, a diagram showing the format of the instruction, a symbolic notation describing its execution and the effect on the condition codes, timing information, a description, special comments, and examples.

MNEMONIC: This is indicated at the top corner of each page. When the word instruction has a byte equivalent, the byte mnemonic is also shown.

INSTRUCTION FORMAT: A diagram accompanying each instruction shows the octal op code, the binary op code, and bit assignments. (Note that in byte instructions the most significant bit (bit 15) is always a 1.)

OPERATION: The operation of each instruction is described with a single notation. The following symbols are used:

() = contents of

src = source address

dst = destination address

loc = location

← = becomes

↑ = "is popped from stack"

↓ = "is pushed onto stack"

∧ = boolean AND

v = boolean OR

∨ = exclusive OR

~ = boolean not

Reg or R = register

B = Byte


INSTRUCTION TIMING: The minimum execution time, including the fetch of the next instruction, is specified for each instruction. The instruction is assumed to reside in bipolar memory; both source and destination are assumed to be general purpose registers. For example, MOV is assumed to be from a general register to a general register, and JMP is assumed to be in register deferred mode. For detailed timing information consult Appendix B.

ISP—The Instruction Set Processor (ISP) notation has been used with each instruction. It is a precise notation for defining the action of any instruction set and is described in detail in Appendix A. It was included for the benefit of PDP-11 users who wish to gain an in depth understanding of each instruction. However, understanding ISP is not essential to understanding PDP-11 instructions.


4.2 INSTRUCTION FORMATS

The major instruction formats are:


Single Operand Group




Double Operand Group




Condition Code Operators




Register-Source or Destination



Subroutine Return




Branch



4.3 BYTE INSTRUCTIONS

The PDP-11 processor includes a full complement of instructions that manipulate byte operands. Since all PDP-11 addressing is byte-oriented, byte manipulation addressing is straightforward. Byte instructions with autoincrement or autodecrement direct addressing cause the specified register to be modified by one to point to the next byte of data. Byte operations in register mode access the low-order byte of the specified register. These provisions enable the PDP-11 to perform as either a word or byte processor. The numbering scheme for word and byte addresses in core memory is:



The most significant bit (Bit 15) of the instruction word is set to indicate a byte instruction.

Example:

	Symbolic	Octal
CLR		0050DD
CLRB		1050DD

4.4 SINGLE OPERAND INSTRUCTIONS

4.4.1 Single Operand Arithmetic Instructions


General:	CLR	DEC	INC	NEG	TST	COM
	CLRB	DECB	INCB	NEGB	TSTB	COMB
Shifts:	ASR	ASL	ASH	ASHC		
	ASRB	ASLB				
Multiple Precision:	ADC	SBC	SXT			
	ADCB	SBCB				
Rotates:	ROL	ROR	SWAB			
	ROLB	RORB				

300 ns

CLR CLRB

Clear dst

n050DD



Operation: (dst) ← 0

Condition Codes: N: cleared
Z: set
V: cleared
C: cleared

Description: Word: Contents of specified destination are replaced with zeroes.

Byte: Same

Example:

CLR R1

Before
(R1) = 177777

After
(R1) = 000000

N Z V C
1 1 1 1


N Z V C
0 1 0 0

300 ns

**DEC
DECB**

Decrement dst

n053DD



Operation: (dst) ← (dst) - 1

Condition Codes: N: set if result is <0; cleared otherwise
Z: set if result is 0; cleared otherwise
V: set if (dst) was 100000; cleared otherwise
C: not affected

Description: Word: Subtract 1 from the contents of the destination
Byte: Same

Example: DEC R5


Before	After
(R5) = 000001	(R5) = 000000
N Z V C	N Z V C
1 0 0 0	0 1 0 0

300 ns

INC INCB

Increment dst

n052DD



Operation: $(dst) \leftarrow (dst) + 1$

Condition Codes: N: set if result is <0 ; cleared otherwise
Z: set if result is 0; cleared otherwise
V: set if (dst) held 077777; cleared otherwise
C: not affected

Description: Word: Add one to contents of destination
Byte: Same

Example:

INC R2


Before	After
(R2) = 000333	(R2) = 000334
N Z V C	N Z V C
0 0 0 0	0 0 0 0

750 ns

NEGB
NEG

Negate dst

n0054DD



Operation: (dst) ← -(dst)

Condition Codes: N: set if the result is <0; cleared otherwise
Z: set if result is 0; cleared otherwise
V: set if the result is 100000; cleared otherwise
C: cleared if the result is 0; set otherwise

Description: Word: Replaces the contents of the destination address by its two's complement. Note that 100000 is replaced by itself (in two's complement notation the most negative number has no positive counterpart).
Byte: Same

Example:

NEG R0


	Before		After
(R0) =	000010	(R0) =	177770
	N Z V C		N Z V C
	0 0 0 0		1 0 0 1

300 ns

TST TSTB

Test dst

n057DD



Operation: (dst) ← -(dst)

Condition Codes: N: set if the result is <0; cleared otherwise
Z: set if result is 0; cleared otherwise
V: cleared
C: cleared


Description: Word: Sets the condition codes N and Z according to the contents of the destination address
Byte: Same

Example: TST R1

	Before		After
(R1) =	012340	(R1) =	012340
	N Z V C		N Z V C
	0 0 1 1		0 0 0 0

300 ns

COM COMB



Operation: (dst) ← (dst)

Condition Codes: N: set if most significant bit of result is set; cleared otherwise
Z: set if result is 0; cleared otherwise
V: cleared
C: set

Description: Replaces the contents of the destination address by their logical complement (each bit equal to 0 is set and each bit equal to 1 is cleared)
Byte: Same

Example:

		COM R0
	Before	After
(R0) =	013333	(R0) = 164444
	N Z V C	N Z V C
	0 1 1 0	1 0 0 1

4.4.2 Shifts

Scaling data by factors of two is accomplished by the shift instructions:

ASR—Arithmetic shift right	ASC—Multiple shift one word
ASL—Arithmetic shift left	ASC—Multiple shift one word


The sign bit (bit 15) of the operand is replicated in shifts to the right. The low order bit is filled with 0 in shifts to the left. Bits shifted out of the C bit, as shown in the following examples, are lost.

300 ns


ASR ASRB

Arithmetic Shift Right dst

n062DD



- Operation:** (dst) ← (dst) shifted one place to the right
- Condition Codes:** N: set if the high-order bit of the result is set (result < 0); cleared otherwise
Z: set if the result = 0; cleared otherwise
V: loaded from the Exclusive OR of the N-bit and C-bit (as set by the completion of the shift operation)
C: loaded from low-order bit of the destination
- Description:** Word: Shifts all bits of the destination right one place. Bit 15 is replicated. The C-bit is loaded from bit 0 of the destination. ASR performs signed division of the destination by two.
Word:




300 ns

ASL ASLB


Arithmetic Shift Left dst

n063DD




- Operation:** $(dst) \leftarrow (dst)$ shifted one place to the left
- Condition Codes:** N: set if high-order bit of the result is set (result < 0); cleared otherwise
Z: set if the result = 0; cleared otherwise
V: loaded with the exclusive OR of the N-bit and C-bit (as set by the completion of the shift operation)
C: loaded with the high-order bit of the destination

Description: Word: Shifts all bits of the destination left one place. Bit 0 is loaded with an 0. The C-bit of the status word is loaded from the most significant bit of the destination. ASL performs a signed multiplication of the destination by 2 with overflow indication.
Byte: Same
Word:



Byte:




750 ns + 150 ns x absolute value shift count

ASH

Shift Arithmetically


072RSS



Operation: $R \leftarrow R$ Shifted arithmetically NN places to right or left
Where NN = (src)

Condition Codes: N: set if result < 0 ; cleared otherwise
Z: set if result = 0; cleared otherwise
V: set if sign of register changed during shift; cleared otherwise
C: loaded from last bit shifted out of register

Description: The contents of the register are shifted right or left the number of times specified by the source operand. The shift count is taken as the low order 6 bits of the source operand. This number ranges from -32 to $+31$. Negative is a right shift and positive is a left shift.




750 ns + 150 ns x absolute value shift count

ASHC

Arithmetic Shift Combined


073RSS



Operation: R, Rv1 ←R, Rv1 The double word is shifted NN places to the right or left, where NN = (src)

Condition Codes: N: set if result <0; cleared otherwise
 Z: set if result =0; cleared otherwise
 V: set if sign bit changes during the shift; cleared otherwise
 C: loaded with high order bit when SC>0; loaded with low order bit when SC<0 (loaded with the last bit shifted out of the 32-bit operand)


Description: The contents of the register and the register ORed with one are treated as one 32 bit word, R + 1 (bits 0-15) and R (bits 16-31) are shifted right or left the number of times specified by the shift count. The shift count is taken as the low order 6 bits of the source operand. This number ranges from -32 to +31. Negative is a right shift and positive is a left shift. When the register chosen is an odd number the register and the register OR'ed with one are the same. In this case the right shift becomes a rotate. The 16 bit word is rotated right the number of bits specified by the shift count.



4.4.3 Multiple Precision

It is sometimes necessary to do arithmetic on operands considered as multiple words or bytes. The PDP-11 makes special provision for such operations with the instructions ADC (Add Carry) and SBC (Subtract Carry) and their byte equivalents.

For example two 16-bit words may be combined into a 32-bit double precision word and added or subtracted as shown below:



Example:

The addition of -1 and -1 could be performed as follows:

$$-1 = 3777777777$$

$$(R1) = 177777 \quad (R2) = 177777 \quad (R3) = 177777 \quad (R4) = 177777$$

ADD R1, R2


ADC R3

ADD R4, R3

1. After (R1) and (R2) are added, 1 is loaded into the C bit
2. ADC instruction adds C bit to (R3); (R3) = 0
3. (R3) and (R4) are added
4. Result is 37777777776 or -2

300 ns

ADC ADCB



Operation: $(dst) \leftarrow (dst) + (C)$

Condition Codes: N: set if result < 0 ; cleared otherwise
Z: set if result $= 0$; cleared otherwise
V: set if (dst) was 077777 and (C) was 1; cleared otherwise
C: set if (dst) was 177777 and (C) was 1; cleared otherwise

Description: Adds the contents of the C-bit into the destination. This permits the carry from the addition of the low-order words to be carried into the high-order result.
Byte: Same


Example: Double precision addition may be done with the following instruction sequence:
ADD A0,B0 ; add low-order parts
ADC B1 ; add carry into high-order
ADD A1,B1 ; add high order parts

300 ns

SBC
SBCB

Subtract Carry dst

n056DD



Operation: $(dst) \leftarrow (dst) - (C)$

Condition Codes: N: set if result < 0 ; cleared otherwise
Z: set if result 0; cleared otherwise
V: set if result is 100000; cleared otherwise
C: cleared if result is 0 and $C = 1$; set otherwise


Description: Word: Subtracts the contents of the C-bit from the destination. This permits the carry from the subtraction of two low-order words to be subtracted from the high order part of the result.
Byte: Same

Example: Double precision subtraction is done by:

```
SUB  A0,B0
SBC  B1
SUB  A1,B1
```

300 ns

SXT



Operation: (dst) ← 0 if N bit is clear
(dst) ← -1 N bit is set

Condition Codes: N: unaffected
Z: set if N bit clear
V: cleared
C: unaffected

Description: If the condition code bit N is set then a -1 is placed in the destination operand; if N bit is clear, then a 0 is placed in the destination operand. This instruction is particularly useful in multiple precision arithmetic because it permits the sign to be extended through multiple words.

Example:

SXT A


Before	After
(A) = 012345	(A) = 177777
N Z V C	N Z V C
1 0 0 0	1 0 0 0

4.4.4 Rotates

The rotate instructions operate on the destination word and the C bit as though they formed a 17-bit "circular buffer." These instructions facilitate sequential bit testing and detailed bit manipulation.

300 ns

ROL ROLB




Condition Codes:

- N: set if the high-order bit of the result word is set (result < 0); cleared otherwise
- Z: set if all bits of the result word = 0; cleared otherwise
- V: loaded with the Exclusive OR of the N-bit and C-bit (as set by the completion of the rotate operation)
- C: loaded with the high-order bit of the destination


Description: Word: Rotate all bits of the destination left one place. Bit 15 is loaded into the C-bit of the status word and the previous contents of the C-bit are loaded into Bit 0 of the destination.
Byte: Same

Example:



300 ns

ROR RORB




Condition Codes:

- N: set if the high-order bit of the result is set (result < 0); cleared otherwise
- Z: set if all bits of result = 0; cleared otherwise
- V: loaded with the Exclusive OR of the N-bit and C-bit (as set by the completion of the rotate operation)
- C: loaded with the low-order bit of the destination


Description: Rotates all bits of the destination right one place. Bit 0 is loaded into the C-bit and the previous contents of the C-bit are loaded into bit 15 of the destination.
Byte: Same


Example:

Word:



Byte:



SWAB

Operation: Byte 1/Byte 0 \leftarrow Byte 0/Byte 1

Condition Codes: N: set if high-order bit of low-order byte (bit 7) of result is set; cleared otherwise
 Z: set if low-order byte of result = 0; cleared otherwise
 V: cleared
 C: cleared

Description: Exchanges high-order byte and low-order byte of the destination word (destination must be a word address).

Example:

SWAB R1

	Before		After
(R1) =	077777	(R1) =	177577
	N Z V C		N Z V C
	1 1 1 1		0 0 0 0

ISP:

4.5 DOUBLE OPERAND INSTRUCTIONS


Double operand instructions provide an instruction (and time) saving facility since they eliminate the need for "load" and "save" sequences such as those used in accumulator-oriented machines.

General:	MOV MOVB	ADD	SUB	CMP CMPB
Register Destination:	MUL	DIV	XOR	
Logical:	BIS BISB	BIT BITB	BIC BICB	

4.5.1 Double Operand General Instructions

300 ns

MOV MOVB



Operation: (dst) ← (src)

Condition Codes: N: set if (src) < 0; cleared otherwise

Z: set if (src) = 0; cleared otherwise

V: cleared

C: not affected

Description: Word: Moves the source operand to the destination location. The previous contents of the destination are lost. The contents of the source address are not affected.

Byte: Same as MOV. The MOV_B to a register (unique among byte instructions) extends the most significant bit of the low order byte (sign extension). Otherwise MOV_B operates on bytes exactly as MOV operates on words.

Example: MOV XXX,R1 ; loads Register 1 with the contents of memory location; XXX represents a programmer-defined mnemonic used to represent a memory location


MOV #20,R0 ; loads the number 20 into Register 0; “#” indicates that the value 20 is the operand

MOV @ #20, -(R6) ; pushes the operand contained in location 20 onto the stack

MOV (R6)+, @ #177566 ; pops the operand off a stack and moves it into memory location 177566 (terminal print buffer)

MOV R1,R3 ; performs an interregister transfer

MOVB @ #177562, @ #177566 ; moves a character from terminal keyboard buffer to terminal buffer

ADD

Operation: (dst) ← (src) + (dst)


Condition Codes: N: set if result < 0; cleared otherwise
 Z: set if result = 0; cleared otherwise
 V: set if there was arithmetic overflow as a result of the operation; that is both operands were of the same sign and the result was of the opposite sign; cleared otherwise
 C: set if there was a carry from the most significant bit of the result; cleared otherwise

Description: Adds the source operand to the destination operand and stores the result at the destination address. The original contents of the destination are lost. The contents of the source are not affected. Two's complement addition is performed.

Examples: Add to register: ADD 20,R0
 Add to memory: ADD R1,XXX
 Add register to register: ADD R1,R2
 Add memory to memory: ADD @#17750,XXX
 XXX is a programmer-defined mnemonic for a memory location.

300 ns

SUB



Operation: $(dst) \leftarrow (dst) - (src)$ [in detail $(dst) \leftarrow (dst) + (src) + 1$]


Condition Codes:
N: set if result < 0 ; cleared otherwise
Z: set if result $= 0$; cleared otherwise
V: set if there was arithmetic overflow as a result of the operation, that is if operands were of opposite signs and the sign of the source was the same as the sign of the result; cleared otherwise
C: cleared if there was a carry from the most significant bit of the result; set otherwise

Description: Subtracts the source operand from the destination operand and leaves the result at the destination address. The original contents of the destination are lost. The contents of the source are not affected. In double-precision arithmetic the C-bit, when set, indicates a "borrow"

Example:

SUB R1, R2

	Before		After
(R1) = 011111		(R1) = 011111	
(R2) = 012345		(R2) = 001234	
	N Z V C		N Z V C
	1 1 1 1		0 0 0 0

**CMP
CMPB**

Operation: (src)−(dst) [in detail, (src) + ~ (dst) + 1]


Condition Codes: N: set if result < 0; cleared otherwise
 Z: set if result = 0; cleared otherwise
 V: set if there was arithmetic overflow; that is, operands were of opposite signs and the sign of the destination was the same as the sign of the result; cleared otherwise
 C: cleared if there was a carry from the most significant bit of the result; set otherwise

Description: Compares the source and destination operands and sets the condition codes, which may then be used for arithmetic and logical conditional branches. Both operands are unaffected. The only action is to set the condition codes. The compare is customarily followed by a conditional branch instruction.

Note that unlike the subtract instruction the order of operation is (src)−(dst), not (dst)−(src).

3.3 μ s

MUL



Operation: R, Rv1 \leftarrow R x(src)


Condition Codes: N: set if product is <0 ; cleared otherwise
Z: set if product is 0; cleared otherwise
V: cleared
C: set if the result is less than -2^{15} or greater than or equal to $2^{15}-1$.

Description: The contents of the destination register and source taken as two's complement integers are multiplied and stored in the destination register and the succeeding register (if R is even). If R is odd only the low order product is stored. Assembler syntax: MUL S,R.
(Note that the actual destination is R,Rv1 which reduces to just R when R is odd.)

Example: 16-bit product (R is odd)

```
CLC ;Clear carry condition code
MOV #400,R1
MUL #10,R1
BCS ERROR ;Carry will be set if
           ;product is less than
           ; $-2^{15}$  or greater than or equal
           ;to  $2^{15}$ 
           ;no significance lost

Before After
(R1) = 000400 (R1) = 004000
```

DIV

Operation: R, Rv1 \leftarrow RA, Rv1 / (src)

Condition Codes: N: set if quotient < 0 ; cleared otherwise
 Z: set if quotient $= 0$; cleared otherwise
 V: set if source $= 0$ or if the absolute value of the register is larger than the absolute value of the source. (In this case the instruction is aborted because the quotient would exceed 15 bits.)
 C: set if divide 0 attempted; cleared otherwise


Description: The 32-bit two's complement integer in R and Rv1 is divided by the source operand. The quotient is left in R; the remainder in Rv1. Division will be performed so that the remainder is of the same sign as the dividend. R must be even.

Example: CLR R0
 MOV #20001,R1
 DIV #2,R0

Before	After	
(R0) = 000000	(R0) = 010000	Quotient
(R1) = 020001	(R1) = 000001	Remainder

300 ns

XOR



Operation: (dst) ← Rv(dst)

Condition Codes: N: set if the result < 0; cleared otherwise
Z: set if result = 0; cleared otherwise
V: cleared
C: unaffected

Description: The exclusive OR of the register and destination operand is stored in the destination address. Contents of register are unaffected. Assembler format is: XOR R,D

Example: XOR R0,R2


	Before		After
(R0) = 001234		(R0) = 001234	
(R2) = 001111		(R2) = 000325	

4.5.2 Logical Instructions

These instructions have the same format as the double operand arithmetic group. They permit operations on data at the bit level.

300 ns

BIS BISB



Operation: (dst) ← (src) v (dst)

Condition Codes: N: set if high-order bit of result set, cleared otherwise
Z: set if result = zero; cleared otherwise
V: cleared
C: not affected

Description: Performs "Inclusive OR" operation between the source and destination operands and leaves the result at the destination address; that is, corresponding bits set in the source are set in the destination. The content of the destination are lost.


Example:

BIS R0,R1

	Before		After
(R0) =	001234	(R0) =	001234
(R1) =	001111	(R1) =	001335
	N Z V C		N Z V C
	0 0 0 0		0 0 0 0

300 ns

BIT BITB



Operation: (dst) \wedge (src)


Condition Codes: N: set if high-order bit of result set; cleared otherwise
Z: set if result =0; cleared otherwise
V: cleared
C: not affected

Description: Performs logical “and” comparison of the source and destination operands and modifies condition codes accordingly. Neither the source nor destination operands are affected. The BIT instruction may be used to test whether any of the corresponding bits that are set in the destination are also set in the source or whether all corresponding bits set in the destination are clear in the source.

Example: BIT #30,R3 ; test bits 3 and 4 of R3
; to see if both are off
BEQ HELP ; BEQ to HELP will occur
; if both are off

300 ns

BIC BICB



Operation: $(dst) \leftarrow \sim(src) \wedge (dst)$

Condition Codes: N: set if high order bit of result set; cleared otherwise
Z: set if result = 0; cleared otherwise
V: cleared
C: not affected

Description: Clears each bit in the destination that corresponds to a set bit in the source. The original contents of the destination are lost. The contents of the source are unaffected.

Example:

BIC R3,R4

Before	After
(R3) = 001234	(R3) = 001234
(R4) = 001111	(R4) = 000101
N Z V C	N Z V C
1 1 1 1	0 0 0 1

4.6 PROGRAM CONTROL INSTRUCTIONS

4.6.1 Branches

The instruction causes a branch to a location defined by the sum of the offset (multiplied by 2) and the current contents of the Program Counter if:

- a) the branch instruction is unconditional
- b) it is conditional and the conditions are met after testing the condition codes (status word).

The offset is the number of words from the current contents of the PC. Note that the current contents of the PC point to the word following the branch instruction.

Although the PC expresses a byte address, the offset is expressed in words. The offset is automatically multiplied by two to express bytes before it is added to the PC. Bit 7 is the sign of the offset. If it is set, the offset is negative and the branch is done in the backward direction. Similarly if it is not set, the offset is positive and the branch is done in the forward direction.

The 8-bit offset allows branching in the backward direction by 200_8 words (400_8 bytes) from the current PC, and in the forward direction by 177_8 words (376_8 bytes) from the current PC.

The PDP-11 assembler handles address arithmetic for the user and computes and assembles the proper offset field for branch instructions in the form:

Bxx loc


Where "Bxx" is the branch instruction and "loc" is the address to which the branch is to be made. The assembler gives an error indication in the instruction if the permissible branch range is exceeded. Branch instructions have no effect on condition codes.

600 ns

BR

Branch (unconditional)

0004 loc



Operation: $PC \leftarrow PC + (2 \times \text{offset})$

Description: Provides a way of transferring program control within a range of -128 to $+127$ words with a one word instruction.

Example:

```
001000    BR xxx
001002
001004
xxx:      001006
          001010
```

Simple Conditional Branches

BEQ

BNE

BMI

BPL

BCS

BCC


BVS

BVC

300 ns—no branch

600 ns—branch

BEQ



Operation: $PC \leftarrow PC + (2 \times \text{offset})$ if $Z = 1$

Condition Codes: Unaffected

Description: Tests the state of the Z-bit and causes a branch if Z is set. As an example, it is used to test equality following a CMP operation, to test that no bits set in the destination were also set in the source following a BIT operation, and generally, to test that the result of the previous operation was zero.

Example:

```
CMP  A,B          ; compare A and B
BEQ  C            ; branch if they are equal

will branch to C if A = B      (A - B = 0)
and the sequence

ADD  A,B          ; add A to B
BEQ  C            ; branch if the result = 0


will breach to C if A + B = 0.
```

300 ns—no branch
600 ns—branch

BNE

Branch Not Equal (Zero)

0010 offset



Operation: $PC \leftarrow PC + (2 \times \text{offset})$ if $Z = 0$

Condition Codes: Unaffected

Description: Tests the state of the Z-bit and causes a branch if the Z-bit is clear. BNE is the complementary operation to BEQ. It is used to test inequality following a CMP, to test that some bits set in the destination were also in the source, following a BIT, and generally, to test that the result of the previous operation was not zero.

Example: `CMP A,B ; compare A and B`
`BNE C ; branch if they are not equal`


will branch to C if $A = B$
and the sequence

`ADD A,B ; add A to B`
`BNE C ; branch if the result not equal to 0`

will branch to C if $A + B = 0$

300 ns—no branch
600 ns—branch

BMI




Operation: $PC \leftarrow PC + (2 \times \text{offset})$ if $N = 1$

Condition Codes: Unaffected

Description: Tests the state of the N-bit and causes a branch if N is set. It is used to test the sign (most significant bit) of the result of the previous operation).

300 ns—no branch
600 ns—branch

BPL




Operation: $PC \leftarrow PC + (2 \times \text{offset})$ if $N = 0$

Description: Tests the state of the N-bit and causes a branch if N is clear. BPL is the complementary operation of BMI.

300 ns—no branch
600 ns—branch

BCS



Operation: $PC \leftarrow PC + (2 \times \text{offset})$ if $C = 1$


Description: Tests the state of the C-bit and causes a branch if C is set. It is used to test for a carry in the result of a previous operation.

300 ns—no branch
600 ns—branch

BCC

Branch on Carry Clear

1030 offset




Operation: $PC \leftarrow PC + (2 \times \text{offset})$ if $C = 0$

Description: Tests the state of the C-bit and causes a branch if C is clear. BCC is the complementary operation to BCS

300 ns—no branch
600 ns—branch

BVS



Operation: $PC \leftarrow PC + (2 \times \text{offset})$ if $V = 1$


Description: Tests the state of V bit (overflow) and causes a branch if the V bit is set. BVS is used to detect arithmetic overflow in the previous operation.

300 ns—no branch
600 ns—branch

BVC

Branch on Overflow Clear

1020 offset



Operation: $PC \leftarrow PC + (2 \times \text{offset})$ if $V = 0$

Description: Tests the state of the V bit and causes a branch if the V bit is clear. BVC is complementary operation to BVS.

Signed Conditional Branches

Particular combinations of the condition code bits are tested with the signed conditional branches. These instructions are used to test the results of instructions in which the operands were considered as a signed (two's complement) values.

Note that the sense of signed comparisons differs from that of unsigned comparisons in that in signed 16-bit, two's complement arithmetic the sequence of values is as follows:

largest	077777
	077776
positive	.
	.
	000001
	000000
	177777
	177776

negative	.
	.
	100001
smallest	100000

whereas in unsigned 16-bit arithmetic the sequence is considered to be highest


	177777
	.
	.
	.
	.
	.
	000002
	000001
lowest	000000

The signed conditional branch instructions are:

BLT	BGE
BLE	BGT

300 ns—no branch
600 ns—branch

BLT




Operation: $PC \leftarrow PC + (2 \times \text{offset})$ if $N \vee V = 1$

Description: Causes a branch if the "Exclusive Or" of the N and V bits are 1. Thus BLT will always branch following an operation that added two negative numbers, even if overflow occurred. In particular, BLT will always cause a branch if it follows a CMP instruction operating on a negative source and a positive destination (even if overflow occurred). Further, BLT will never cause a branch when it follows a CMP instruction operating on a positive source and negative destination. BLT will not cause a branch if the result of the previous operation was zero (without overflow).

300 ns—no branch
600 ns—branch

BGE



Operation: $PC \leftarrow PC + (2 \times \text{offset})$ if $N \vee V = 0$


Description: Causes a branch if N and V are either both clear or both set. BGE is the complementary operation to BLT. Thus BGE will always cause a branch when it follows an operation that caused addition to two positive numbers. BGE will also cause a branch on a zero result.

300 ns—no branch
600 ns—branch

BLE

Branch on Less than or Equal (zero)

0034 offset




Operation: $PC \leftarrow PC + (2 \times \text{offset})$ if $Z \vee (N \vee V) = 1$

Description: Operation is similar to BLT but in addition will cause a branch if the result of the previous operation was zero.

300 ns—no branch
600 ns—branch

BGT



Operation $PC \leftarrow PC + (2 \times \text{offset})$ if $Z \vee (N \vee V) = 0$

Description: Operation of BGT is similar to BGE, except BGT will not cause a branch on a zero result.


Unsigned Conditional Branches

The Unsigned Conditional Branches provide a means for testing the result of comparison operations in which the operands are considered as unsigned values.

BHI
BLOS
BHIS
BLO

300 ns—no branch
600 ns—branch

BHI



Operation: $PC \leftarrow PC + (2 \times \text{offset})$ if $C = 0$ and $Z = 0$


Description: Causes a branch if the previous operation caused neither a carry nor a zero result. This will happen in comparison (CMP) operations as long as the source has a higher unsigned value than the destination.

300 ns—no branch
600 ns—branch

BLOS

Branch on Lower or Same

1014 offset



Operation:


$PC \leftarrow PC + (2 \times \text{offset})$ if $C \vee Z = 1$

Description:

Causes a branch if the previous operation caused either a carry or a zero result. BLOS is the complementary operation to BHI. The branch will occur in comparison operations as long as the source is equal to, or has a lower unsigned value than the destination.

300 ns—no branch
600 ns—branch

BLO



Operation: $PC \leftarrow PC + (2 \times \text{offset})$ if $C = 1$


Description: BLO is same instruction as BCS. This mnemonic is included only for convenience.

300 ns—no branch
600 ns—branch

BHIS

Branch on Higher or Same

1030 offset



Operation: $PC \leftarrow PC + (2 \times \text{offset})$ if $C = 0$

Description: BHIS is the same instruction as BCC. This mnemonic is included only for convenience.


4.6.2 Subroutine Instructions

The subroutine call in the PDP-11 provides for automatic nesting of subroutines, reentrancy, and multiple entry points. Subroutines may call other subroutines (or indeed themselves) to any level of nesting without making special provision for storage or return addresses at each level of subroutine call. The subroutine calling mechanism does not modify any fixed location in memory, thus providing for reentrancy. This allows one copy of a subroutine to be shared among several interrupting processes. For more detailed description of subroutine programming see Chapter 5.

JSR

Jump to Sub Routine

004 reg. dst



- Operation:** (tmp) \leftarrow (dst) (tmp is an internal processor register)
 \downarrow (SP) \leftarrow reg (push reg contents onto processor stack)
 reg \leftarrow PC (PC holds location following JSR; this address)
 PC \leftarrow (tmp) (now put in reg)

- Description:** In execution of the JSR, the old contents of the specified register (the "LINKAGE POINTER") are automatically pushed onto the processor stack and new linkage information placed in the register. Thus subroutines nested within subroutines to any depth may all be called with the same linkage register. There is no need either to plan the maximum depth at which any particular subroutine will be called or to include instructions in each routine to save and restore the linkage pointer. Further, since all linkages are saved in a reentrant manner on the processor stack execution of a subroutine may be interrupted, the same subroutine reentered and executed by an interrupt service routine. Execution of the initial subroutine can then be resumed when other requests are satisfied. This process (called nesting) can proceed to any level.

In both JSR and JMP instructions the destination address is used to load the program counter, R7. Thus for example a JSR in destination mode 1 for general register R1 (where (R1) = 100), will access a subroutine at location 100. This is effectively one level less of deferral than operate instructions such as ADD.


A subroutine called with a JSR reg,dst instruction can access the arguments following the call with either autoincrement addressing, (reg) +, (if arguments are accessed sequentially) or by indexed

addressing, X(reg), (if accessed in random order) These addressing modes may also be deferred, @(reg) + and @X(reg) if the parameters are operand addresses rather than the operand themselves.

JSR PC, dst is a special case of the PDP-11 subroutine call suitable for subroutine calls that transmit parameters through the general registers. The SP and the PC are the only registers that may be modified by this call.

Another special case of the JSR instruction is JSR PC, @(SP) + which exchanges the top element of the processor stack and the contents of the program counter. Use of this instruction allows two routines to swap program control and resume operation when recalled where they left off. Such routines are called "co-routines."

Return from a subroutine is done by the RTS instruction. RTS reg loads the contents of reg into the PC and pops the top element of the processor stack into the specified register.

MARK

Operation: $SP \leftarrow PC + 2xnn$ $nn = \text{number of parameters}$
 $PC \leftarrow R5$
 $R5 \leftarrow (SP) \uparrow$

Condition Codes: unaffected

Description: Used as part of the standard PDP-11 subroutine return convention. MARK facilitates the stack clean up procedures involved in subroutine exist. Assembler format is: MARK N

Example:

```

MOV R5, -(SP)           ;place old R5 on stack
MOV P1, -(SP)          ;place N parameters
MOV P2, -(SP)          ;on the stack to be
                       ;used there by the
                       ;subroutine

MOV PN, -(SP)
MOV =MARKN, -(SP)     ;places the instruction
                       ;MARK N on the stack
MOV SP, R5            ;set up address at Mark N
                       ;instruction
JSR PC, SUB           ;jump to subroutine
  
```

At this point the stack is as follows:

OLD R5
P1
PN
MARK N
OLD PC

And the program is at the address SUB which is the beginning of the subroutine.

SUB: ;execution of the subroutine itself

RTS R5 ;the return begins: this causes

the contents of R5 to be placed in the PC which then results in the execution of the instruction MARK N. The contents old PC are placed in R5

MARK N causes: (1) the stack pointer to be adjusted to point to the old R5 value; (2) the value now in R5 (the old PC) to be placed in the PC; and (3) contents of the old R5 to be popped into R5 thus completing the return from subroutine.


Note: If Memory Management is in use a stack must be in I and D spaces (Chapter 6) to execute the MARK instruction.

1.2 μ s

RTS

Return from Subroutine

00020 Reg



Operation: PC \leftarrow reg
reg \leftarrow (SP) \uparrow

Description: Loads contents of reg into PC and pops the top element of the processor stack into the specified register.

Return from a non-reentrant subroutine is typically made through the same register that was used in its call. Thus, a subroutine called with a JSR PC, dst exits with a RTS PC and a subroutine called with a JSR R5, dst, may pick up parameters with addressing modes (R5)+, X(R5), or @X(R5) and finally exits, with an RTS R5.

4.6.3 Program Control Instructions


SPL
JMP
SOB

600 ns

SPL

Set Priority Level

00023N



Operation: PS (bits 7-5) ← Priority

Condition Codes: not affected

Description The least significant three bits of the instruction are loaded into the Program Status Word (PS) bits 7-5 thus causing a changed priority. The old priority is lost.

Assembler syntax is: SPL N


Note: This instruction is a no op in User and Supervisor modes.

600 ns

JMP

Jump

0001DD



Operation: PC ← (dst)

Condition Codes: not affected

Description: JMP provides more flexible program branching than provided with the branch instructions. Control may be transferred to any location in memory (no range limitation) and can be accomplished with the full flexibility of the addressing modes, with the exception of register mode 0. Execution of a jump with mode 0 will cause an "illegal" instruction condition. (Program control cannot be transferred to a register.) Register deferred mode is legal and will cause program control to be transferred to the address held in the specified register. Note that instructions are word data and must therefore be fetched from an even-numbered address. A "boundary error" trap condition will result when the processor attempts to fetch an instruction from an odd address.


Deferred index mode JMP instructions permit transfer of control to the address contained in a selectable element of a table of dispatch vectors.

750 ns—no branch
600 ns—branch

SOB

Subtract One and Branch

077R offset



Operation: $R \leftarrow R - 1$ if this result = 0 then $PC \leftarrow PC - (2x \text{ offset})$

Condition Codes: unaffected

Description: The register is decremented. If it is not equal to 0, twice the offset is subtracted from the PC (now pointing to the following word). The offset is interpreted as a six bit positive number. This instruction provides a fast, efficient method of loop control. Assembler syntax is:

SOB R,A

Where A is the address to which transfer is to be made if the decremented R is not equal to 0. Note that the SOB instruction can not be used to transfer control in the forward direction.

4.6.4 Traps

Trap instructions provide for calls to emulators, I/O monitors, debugging packages, and user-defined interpreters. A trap is effectively an interrupt generated by software. When a trap occurs the contents of the current Program Counter (PC) and Program Status Word (PS) are pushed onto the processor stack and replaced by the contents of a two-word trap vector containing a new PC and new PS. The return sequence from a trap involves executing an RTI or RTT instruction which restores the old PC and old PS by popping them from the stack. Trap vectors are located permanently assigned fixed address.


TRAP
EMT
BPT
IOT
RTI
RTT

2.25 μ s

EMT

Emulator Traps

104000-104377



Operation: $\downarrow(\text{SP}) \leftarrow \text{PS}$
 $\downarrow(\text{SP}) \leftarrow \text{PC}$
 $\text{PC} \leftarrow (30)$
 $\text{PS} \leftarrow (32)$


Condition Codes: N: loaded from trap vector
 Z: loaded from trap vector
 V: loaded from trap vector
 C: loaded from trap vector

Description: All operation codes from 104000 to 104377 are EMT instructions and may be used to transmit information to the emulating routine (e.g., function to be performed). The trap vector for EMT is at address 30. The new PC is taken from the word at address 30; the new central processor status (PS) is taken from the word at address 32.

Caution: EMT is used frequently by DIGITAL system software and is therefore not recommended for general use.

2.25 μ s

TRAP



Operation:

- $\downarrow(\text{SP}) \leftarrow \text{PS}$
- $\downarrow(\text{SP}) \leftarrow \text{PC}$
- $\text{PC} \leftarrow (34)$
- $\text{PS} \leftarrow (36)$

Condition Codes:

- N: loaded from trap vector
- Z: loaded from trap vector
- V: loaded from trap vector
- C: loaded from trap vector

Description: Operation codes from 104400 to 104777 are TRAP instructions. TRAPs and EMTs are identical in operation, except that the trap vector for TRAP is at address 34.


Note: Since DEC software makes frequent use of EMT, the TRAP instruction is recommended for general use.

2.25 μ S

BPT

Breakpoint Trap

000003




Operation: \downarrow (SP) \leftarrow PS
 \downarrow (SP) \leftarrow PC
 PC \leftarrow (14)
 PC \leftarrow (16)

Condition Codes: N: loaded from trap vector
 Z: loaded from trap vector
 V: loaded from trap vector
 C: loaded from trap vector

Description: Performs a trap sequence with a trap vector address of 14. Used to call debugging aids. The user is cautioned against employing code 000003 in programs run under these debugging aids.
(no information is transmitted in the low byte.)

2.25 μ s

IOT



Operation: \downarrow (SP) \leftarrow PS
 \downarrow (SP) \leftarrow PC
 PC \leftarrow (20)
 PS \leftarrow (22)

Condition Codes: N: loaded from trap vector
 Z: loaded from trap vector
 V: loaded from trap vector
 C: loaded from trap vector


Description: Performs a trap sequence with a trap vector address of 20. Used to call the I/O Executive routine IOX in the paper tape software system, and for error reporting in the Disk Operating System. (no information is transmitted in the low byte)

1.5 μ s

RTI

Return from Interrupt

000002



Operation: PC \leftarrow (SP) \uparrow
PS \leftarrow (SP) \uparrow

Condition Codes: N: loaded from processor stack
Z: loaded from processor stack
V: loaded from processor stack
C: loaded from processor stack


Description: Used to exit from an interrupt or TRAP service routine. The PC and PS are restored (popped) from the processor stack.

1.5 μ s

RTT

Return from Trap

000006



Operation: PC \leftarrow (SP) \uparrow
PS \leftarrow (SP) \uparrow

Condition Codes: N: loaded from processor stack
Z: loaded from processor stack
V: loaded from processor stack
C: loaded from processor stack

Description: This is the same as the RTI instruction except that it inhibits a trace trap, while RTI permits a trace trap. If a trace trap is pending, the first instruction after the RTT will be executed prior to the next "T" trap. In the case of the RTI instruction the "T" trap will occur immediately after the RTI.

Reserved Instruction Traps—These are caused by attempts to execute instruction codes reserved for future processor expansion (reserved instructions) or instructions with illegal addressing modes (illegal instructions). Order codes not corresponding to any of the instructions described are considered to be reserved instructions. JMP and JSR with register mode destinations are illegal instructions. Reserved and illegal instruction traps occur as described under EMT, but trap through vectors at addresses 10 and 4 respectively.

Stack Overflow Trap

Bus Error Traps—Bus Error Traps are:

1. **Boundary Errors**—attempts to reference instructions or word operands at odd addresses.
2. **Time-Out Errors**—attempts to reference addresses on the bus that made no response within 5 μ s in the PDP-11/45. In general, these are caused by attempts to reference non-existent memory, and attempts to reference non-existent peripheral devices.

Bus error traps cause processor traps through the trap vector address 4.

Trace Trap—Trace Trap enables bit 4 of the PS and causes processor traps at the end of instruction executions. The instruction that is executed after the instruction that set the T-bit will proceed to completion and then cause a processor trap through the trap vector at address 14. Note that the trace trap is a system debugging aid and is transparent to the general programmer.

The following are special cases and are detailed in subsequent paragraphs.

1. The traced instruction cleared the T-bit.
2. The traced instruction set the T-bit.
3. The traced instruction caused an instruction trap.
4. The traced instruction caused a bus error trap.
5. The traced instruction caused a stack overflow trap.
6. The process was interrupted between the time the T-bit was set and the fetching of the instruction that was to be traced.
7. The traced instruction was a WAIT.
8. The traced instruction was a HALT.
9. The traced instruction was a Return from Trap.

Note: The traced instruction is the instruction after the one that sets the T-bit.

An instruction that cleared the T-bit—Upon fetching the traced instruction an internal flag, the trace flag, was set. The trap will still occur at the end of execution of this instruction. The stacked status word, however, will have a clear T-bit.

An instruction that set the T-bit—Since the T-bit was already set, setting it again has no effect. The trap will occur.

An instruction that caused an Instruction Trap—The instruction trap is sprung and the entire routine for the service trap is executed. If the service routine exits with an RTI or in any other way restores the stacked status word, the T-bit is set again, the instruction following the traced instruction is executed and, unless it is one of the special cases noted above, a trace trap occurs.

An instruction that caused a Bus Error Trap—This is treated as an Instruction Trap. The only difference is that the error service is not as likely to exit with an RTI, so that the trace trap may not occur.

An instruction that caused a stack overflow—The instruction completes execution as usual—the Stack Overflow does not cause a trap. The Trace Trap Vector is loaded into the PC and PS, and the old PC and PS are pushed onto the stack. Stack Overflow occurs again, and this time the trap is made.

An interrupt between setting of the T-bit and fetch of the traced instruction—The entire interrupt service routine is executed and then the T-bit is set again by the exiting RTI. The traced instruction is executed (if there have been no other interrupts) and, unless it is a special case noted above, causes a trace trap.

Note that interrupts may be acknowledged immediately after the loading of the new PC and PS at the trap vector location. To lock out all interrupts, the PS at the trap vector should raise the processor priority to level 7.

A WAIT—The trap occurs immediately.

A HALT—The processor halts. When the continue key on the console is pressed, the instruction following the HALT is fetched and executed. Unless it is one of the exceptions noted above, the trap occurs immediately following execution.

A Return from Trap—The return from trap instruction either clears or sets the T-bit. It inhibits the trace trap. If the T-bit was set and RTI is the traced instruction the trap is delayed until completion of the next instruction.

Power Failure Trap—is a standard PDP-11 feature. Trap occurs whenever the AC power drops below 95 volts or outside 47 to 63 Hertz. Two milliseconds are then allowed for power down processing. Trap vector for power failure is at locations 24 and 26.

Trap priorities—in case multiple processor trap conditions occur simultaneously the following order of priorities is observed (from high to low):

1. Odd Address
2. Fatal Stack Violation

3. Segment Violation
4. Timeout
5. Parity Error
6. Console Flag
7. Segment Management Trap
8. Warning Stack Violation
9. Power Failure

The details on the trace trap process have been described in the trace trap operational description which includes cases in which an instruction being traced causes a bus error, instruction trap, or a stack overflow trap.

If a bus error is caused by the trap process handling instruction traps, trace traps, stack overflow traps, or a previous bus error, the processor is halted.

If a stack overflow is caused by the trap process in handling bus errors, instruction traps, or trace traps, the process is completed and then the stack overflow trap is sprung.

4.7 MISCELLANEOUS

HALT

WAIT

RESET


MTPD

MTPI

MFPD

MFPI

HALT



Condition Codes: not affected


Description: Causes the processor operation to cease. The console is given control of the bus. The console data lights display the contents of R0; the console address lights display the address after the halt instruction. Transfers on the UNIBUS are terminated immediately. The PC points to the next instruction to be executed. Pressing the continue key on the console causes processor operation to resume. No INIT signal is given.

Note: A halt issued in Supervisor or User Mode will generate a trap.

WAIT

Wait for Interrupt

000001



Condition Codes: not affected


Description: Provides a way for the processor to relinquish use of the bus while it waits for an external interrupt. Having been given a WAIT command, the processor will not compete for bus use by fetching instructions or operands from memory. This permits higher transfer rates between a device and memory, since no processor-induced latencies will be encountered by bus requests from the device. In WAIT, as in all instructions, the PC points to the next instruction following the WAIT operation. Thus when an interrupt causes the PC and PSW to be pushed onto the processor from the interrupt routine (i.e. execution of an RTI instruction) will cause resumption of the interrupted process at the instruction following the WAIT.

10 ms

RESET

Reset External Bus

000005



Condition Codes: not affected


Description: Sends INIT on the UNIBUS for 10 ms. All devices on the UNIBUS are reset to their state at power up.

900 ns

MTPI

Move to Previous Instruction Space

0066DD



Operation: (temp) \leftarrow (SP) \uparrow
(dst) \leftarrow (temp)

Condition Codes: N: set if the source <0 ; otherwise cleared
Z: set if the source $=0$; otherwise cleared
V: cleared
C: unaffected


Description: The address of the destination operand is determined in the current address space. MTPI then pops a word off the current stack and stores that word in the destination address in the previous mode's I space (bits 13, 12 of PS).

900 ns

MTPD

Move to Previous Data Space

1066DD



Operation: (temp) ← (SP)↑
(dst) ← (temp)

Condition Codes: N: set if the source < 0; otherwise cleared
Z: set if the source = 0; otherwise cleared
V: cleared
C: unaffected


Description: The address of the destination operand is determined in the current address space as in MTPI. MTPD then pops a word off the current stack and stores that word in the destination address in the previous mode's D space.

1.2 μ s

MFPI

Move from Previous Instruction Space

0065SS



Operation: (temp) \leftarrow (src)
 \downarrow (SP) \leftarrow (temp)

Condition Codes: N: set if the source <0 ; otherwise cleared
Z: set if the source $=0$; otherwise cleared
V: cleared
C: unaffected


Description: This instruction is provided in order to allow inter-address space communication when the PDP11/45 is using the Memory Management unit. The address of the source operand is determined in the current address space. That is, the address is determined using the SP and memory pages determined by PS \langle 15:14 \rangle . The address itself is then used in the previous I space (as determined by PS \langle 13:12 \rangle) to get the source operand. This operand is then pushed on to the current R6 stack.

1.2 μ s

MFPD

Move from Previous Data Space

1065SS



Operation: (temp) \leftarrow (src)
 \downarrow (SP) \leftarrow (temp)

Condition Codes: N: set if the source <0 ; otherwise cleared
Z: set if the source $=0$; otherwise cleared
V: cleared
C: unaffected

Description: This instruction is provided in order to allow inter-address space communication when the PDP-11/45 is using the Memory Management unit. The address of the source operand is determined in the current address space. That is, the address is determined using the SP and memory pages determined by PS \langle 15:14 \rangle . The address itself is then used in the previous D space (as determined by PS \langle 13:12 \rangle) to get the source operand. This operand is then pushed on to the current R6 stack.


4.8 Condition Code Operators

600 ns

CLN	SEN
CLZ	SEZ
CLV	SEV
CLC	SEC

Condition Code Operators

0002XX



Description:

Set and clear condition code bits. Selectable combinations of these bits may be cleared or set together. Condition code bits corresponding to bits in the condition code operator (Bits 0-3) are modified according to the sense of bit 4, the set/clear bit of the operator. i.e. set the bit specified by bit 0, 1, 2 or 3, if bit 4 is a 1. Clear corresponding bits if bit 4 = 0.

Mnemonic	Operation	OP Code
CLC	Clear C	000241
CLV	Clear V	000242
CLZ	Clear Z	000244
CLN	Clear N	000250
SEC	Set C	000261
SEV	Set V	000262
SEZ	Set Z	000264
SEN	Set N	000270
SCC	Set all CC's	000277
CCC	Clear all CC's	000257
	Clear V and C	000243
	No Operation	000240

Combinations of the above set or clear operations may be ORed together to form combined instructions.

ADVANCED PROGRAMMING TECHNIQUES

In order to produce programs which fully utilize the power and flexibility of the PDP-11/45, the reader should become familiar with the various programming techniques which are part of the basic design philosophy of the PDP-11. Although it is possible to program the PDP-11/45 along traditional lines such as "accumulator orientation" this approach does not fully exploit the architecture and instruction set of the PDP-11/45.

5.1 THE STACK

A "stack," as used on the PDP-11, is an area of memory set aside by the programmer for temporary storage or subroutine/interrupt service linkage. The instructions which facilitate "stack" handling are useful features not normally found in low-cost computers. They allow a program to dynamically establish, modify, or delete a stack and items on it. The stack uses the "last-in, first-out" concept, that is, various items may be added to a stack in sequential order and retrieved or deleted from the stack in reverse order. On the PDP-11, a stack starts at the highest location reserved for it and expands linearly downward to the lowest address as items are added to the stack.




Figure 5-1: Stack Addresses

To keep track of the last item added to the stack (or "where we are" in the stack) a General Register always contains the memory address where the last item is stored in the stack. In the PDP-11 any register except Register 7 (the Program Counter-PC) may be used as a "stack pointer" under program control; however, instructions associated with subroutine linkage and interrupt service automatically use Register 6 (R6) as a hardware "Stack Pointer." For this reason R6 is frequently referred to as the system "SP".

Stacks in the PDP-11 may be maintained in either full word or byte units. This is true for a stack pointed to by any register except R6, which must be organized in full word units only.

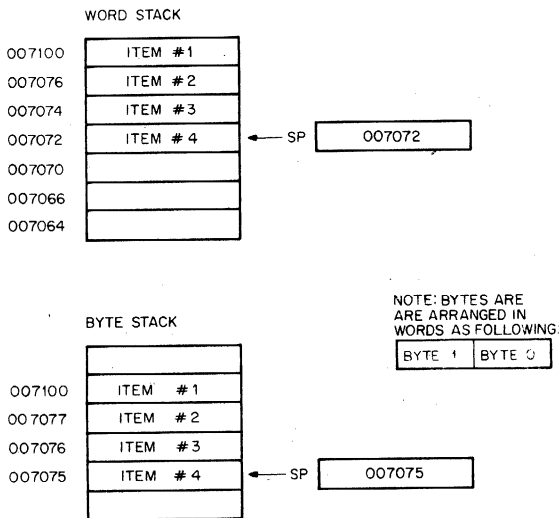


Figure 5-2: Word and Byte Stacks

Items are added to a stack using the autodecrement addressing mode with the appropriate pointer register. (See Chapter 3 for description of the autoincrement/decrement modes).

This operation is accomplished as follows:

MOV Source,—(SP) ;MOV Source Word onto the stack
or

MOVB Source,—(SP) ;MOVB Source Byte onto the stack

This is called a “push” because data is “pushed onto the stack.”

To remove an item from a stack the autoincrement addressing mode with the appropriate SP is employed. This is accomplished in the following manner:

MOV(SP)+,Destination ;MOV Destination Word off the stack
or

MOVB(SP)+,Destination ;MOVB Destination Byte off the stack

Removing an item from a stack is called a “pop” for “popping from the stack.” After an item has been “popped,” its stack location is considered free and available for other use. The stack pointer points to the last-used location implying that the next (lower) location is free. Thus a stack may represent a pool of shareable temporary storage locations.

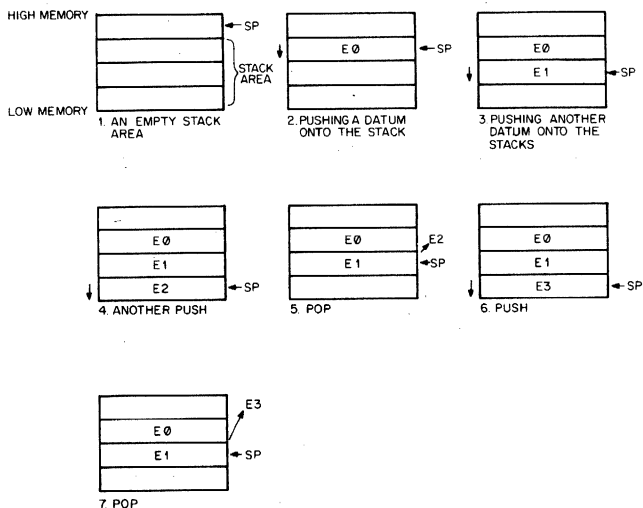


Figure 5-3: Illustration of Push and Pop Operations

As an example of stack usage consider this situation: a subroutine (SUBR) wants to use registers 1 and 2, but these registers must be returned to the calling program with their contents unchanged. The subroutine could be written as follows:

Address	Octal Code	Assembler Syntax
076322	010167	SUBR: MOV R1,TEMP1 ; save R1
076324	000074	*
076326	010267	MOV R2,TEMP2 ;save R2
076330	000072	*
.	.	.
.	.	.
.	.	.
076410	016701	MOV TEMP1,R1 ;Restore R1
076412	000006	*
076414	016702	MOV TEMP2,R2 ; Restore R2
076416	000004	*
076420	000207	RTS PC
076422	000000	TEMP1: 0
076424	000000	TEMP2: 0

*Index Constants

Figure 5-4: Register Saving Without the Stack

OR: Using the Stack

Address	Octal Code	Assembler Syntax
010020	010143	SUBR: MOV R1, -(R3) ;push R1
010022	010243	MOV R2, -(R3) ;push R2
.	.	.
.	.	.
.	.	.
010130	012301	MOV (R3)+,R2 ;pop R2
010132	012302	MOV (R3)+,R1 ;pop R1
010134	000207	RTS PC

Note: In this case R3 was used as a Stack Pointer

Figure 5-5: Register Saving using the Stack

The second routine uses four less words of instruction code and two words of temporary 'stack' storage. Another routine could use the same stack space at some later point. Thus, the ability to share temporary storage in the form of a stack is a very economical way to save on memory usage.

As a further example of stack usage, consider the task of managing an input buffer from a terminal. As characters come in, the terminal user may wish to delete characters from his line; this is accomplished very easily by maintaining a byte stack containing the input characters. Whenever a backspace is received a character is "popped" off the stack and eliminated from consideration. In this example, a programmer has the choice of "popping" characters to be eliminated by using either the MOV_B (MOVE BYTE) or INC (INCREMENT) instructions.




Figure 5-6: Byte Stack used as a Character Buffer

NOTE that in this case using the increment instruction (INC) is preferable to MOV_B since it would accomplish the task of eliminating the unwanted character from the stack by readjusting the stack pointer without the need for a destination location. Also, the stack pointer (SP) used in this example cannot be the system stack pointer (R6) because R6 may only point to word (even) locations.

5.2 SUBROUTINE LINKAGE

5.2.1 Subroutine Calls

Subroutines provide a facility for maintaining a single copy of a given routine which can be used in a repetitive manner by other programs located anywhere else in memory. In order to provide this facility, generalized linkage methods must be established for the purpose of control transfer and information exchange between subroutines and calling programs. The PDP-11 instruction set contains several useful instructions for this purpose.

PDP-11 subroutines are called by using the JSR instruction which has the following format.

a general register (R) for linkage
JSR R,SUBR
 an entry location (SUBR) for the subroutine

When a JSR is executed, the contents of the linkage register are saved on the system R6 stack as if a MOV reg,—(SP) had been performed. Then the same register is loaded with the memory address following the JSR instruction (the contents of the current PC) and a jump is made to the entry location specified.

Address	Assembler Syntax	Octal Code
001000	JSRR5,SUBR	004567
001002	index constant for SUBR	000064
001064	SUBR: MOV A,B	0Innm

Figure 5-7: JSR using R5




Figure 5-8: JSR

Note that the instruction JSR R6,SUBR is not normally considered to be a meaningful combination.

5.2.2 Argument Transmission

The memory location pointed to by the linkage register of the JSR instruction may contain arguments or addresses of arguments. These arguments may be accessed from the subroutine in several ways. Using Register 5 as the linkage register, the first argument could be obtained by using the addressing modes indicated by (R5),(R5)+,X(R5) for actual data, or @(R5)+, etc. for the address of data. If the autoincrement

mode is used, the linkage register is automatically updated to point to the next argument.

Figures 5-9 and 5-10 illustrate two possible methods of argument transmission.

Address Instructions and Data

010400		JSR R5, SUBR	
010402		Index constant for SUBR	SUBROUTINE CALL
010404		arg #1	ARGUMENTS
010406		arg #2	
.	.	.	.
.	.	.	.
.	.	.	.
020306	SUBR:	MOV (R5)+,R1	;get arg #1
020310		MOV (R5)+,R2	;get arg #2 Retrieve Arguments from SUB

Figure 5-9: Argument Transmission-Register Autoincrement Mode

Address Instructions and Data

010400		JSR R5, SUBR	
010402		Index constant for SUBR	SUBROUTINE CALL
010404	077722		Address of arg #1
010406	077724		Address of arg #2
010410	077726		Address of arg #3
.	.	.	.
.	.	.	.
.	.	.	.
077722	arg #1		
077724	arg #2		arguments
077726	arg #3		
.	.	.	.
.	.	.	.
020306	SUBR:	MOV @(R5)+,R1	;get arg #1
020301		MOV @(R5)+,R2	;get arg #2

Figure 5-10: Argument Transmission-Register Autoincrement Deferred Mode

Another method of transmitting arguments is to transmit only the address of the first item by placing this address in a general purpose register. It is not necessary to have the actual argument list in the same general area as the subroutine call. Thus a subroutine can be called to work on data located anywhere in memory. In fact, in many cases, the operations performed by the subroutine can be applied directly to the data located on or pointed to by a stack without the need to ever actually move this data into the subroutine area.

Calling Program: MOV POINTER, R1
 JSR PC, SUBR

SUBROUTINE ADD (R1) + ,(R1) ;Add item #1 to item #2, place result in item #2, R1 points to item #2 now

etc.
 or

ADD (R1), 2(R1) ;Same effect as above except that R1 still points to item #1 etc.




Figure 5-11: Transmitting Stacks as Arguments

Because the PDP-11 hardware already uses general purpose register R6 to point to a stack for saving and restoring PC and PS (processor status word) information, it is quite convenient to use this same stack to save and restore intermediate results and to transmit arguments to and from subroutines. Using R6 in this manner permits extreme flexibility in nesting subroutines and interrupt service routines.

Since arguments may be obtained from the stack by using some form of register indexed addressing, it is sometimes useful to save a temporary copy of R6 in some other register which has already been saved at the beginning of a subroutine. In the previous example R5 may be used to index the arguments while R6 is free to be incremented and decremented in the course of being used as a stack pointer. If R6 had been used directly as the base for indexing and not "copied," it might be difficult to keep track of the position in the argument list since the base of the stack would change with every autoincrement/decrement which occurs.




Figure 5-12: Shifting Indexed Base

However, if the contents of R6 (SP) are saved in R5 before any arguments are pushed onto the stack, the position relative to R5 would remain constant.



Figure 5-13: Constant Index Base Using "R6 Copy"

5.2.3 Subroutine Return

In order to provide for a return from a subroutine to the calling program an RTS instruction is executed by the subroutine. This instruction should specify the same register as the JSR used in the subroutine call. When executed, it causes the register specified to be moved to the PC and the top of the stack to be then placed in the register specified. Note that if an RTS PC is executed, it has the effect of returning to the address specified on the top of the stack.

Note that the JSR and the JMP Instructions differ in that a linkage register is always used with a JSR; there is no linkage register with a JMP and no way to return to the calling program.

When a subroutine finishes, it is necessary to "clean-up" the stack by eliminating or skipping over the subroutine arguments. One way this can be done is by insisting that the subroutine keep the number of arguments as its first stack item. Returns from subroutines would then involve calculating the amount by which to reset the stack pointer, resetting the stack pointer, then restoring the original contents of the register which was used as the copy of the stack pointer. The PDP-11/45, however, has a much faster and simpler method of performing these tasks. The MARK instruction which is stored on a stack in place of "number of argument" information may be used to automatically perform these "clean-up" chores. (For more information on the MARK instruction refer to Chapter 4.)

5.2.4 PDP-11 Subroutine Advantages

There are several advantages to the PDP-11 subroutine calling procedure.

- arguments can be quickly passed between the calling program and the subroutine.
- if the user has no arguments or the arguments are in a general register or on the stack the JSR PC,DST mode can be used so that none of the general purpose registers are taken up for linkage.
- many JSR's can be executed without the need to provide any saving procedure for the linkage information since all linkage information is automatically pushed onto the stack in sequential order. Returns can simply be made by automatically popping this information from the stack in the opposite order of the JSR's.

Such linkage address bookkeeping is called automatic "nesting" of subroutine calls. This feature enables the programmer to construct fast,

efficient linkages in a simple, flexible manner. It even permits a routine to call itself in those cases where this is meaningful (e.g. SQRT in FORTRAN SQRT(SQRT(X))). Other ramifications will appear after we examine the PDP-11 interrupt procedures.

5.3 INTERRUPTS

5.3.1 General Principles

Interrupts are in many respects very similar to subroutine calls. However, they are forced, rather than controlled, transfers of program execution occurring because of some external and program-independent event (such as a stroke on the teleprinter keyboard). Like subroutines, interrupts have linkage information such that a return to the interrupted program can be made. More information is actually necessary for an interrupt transfer than a subroutine transfer because of the random nature of interrupts. The complete machine state of the program immediately prior to the occurrence of the interrupt must be preserved in order to return to the program without any noticeable effects. (i.e. was the previous operation zero or negative, etc.) This information is stored in the Processor Status Word (PS). Upon interrupt, the contents of the Program Counter (PC) (address of next instruction) and the PS are automatically pushed onto the R6 system stack. The effect is the same as if:

```
MOV PS ,-(SP)           ;Push PS
MOV R7, -(SP)          ;Push PC
```

had been executed.


The new contents of the PC and PS are loaded from two preassigned consecutive memory locations which are called an "interrupt vector." The actual locations are chosen by the device interface designer and are located in low memory addresses of Kernel virtual space (see interrupt vector list, Appendix C). The first word contains the interrupt service routine address (the address of the new program sequence) and the second word contains the new PS which will determine the machine status including the operational mode and register set to be used by the interrupt service routine. The contents of the interrupt service vector are set under program control.

After the interrupt service routine has been completed, an RTI (return from interrupt) is performed. The two top words of the stack are automatically "popped" and placed in the PC and PS respectively, thus resuming the interrupted program.


5.3.2 Nesting

Interrupts can be nested in much the same manner that subroutines are nested. In fact, it is possible to nest any arbitrary mixture of subroutines and interrupts without any confusion. By using the RTI and RTS instructions, respectively, the proper returns are automatic.


1. Process 0 is running;
SP is pointing to location P0.




2. Interrupt stops process 0 with PC = PC0, and status = PSO; starts process 1.




3. Process 1 uses stack for temporary storage (TE0, TE1).




4. Process 1 interrupted with PC = PC1 and status = PS1; process 2 is started




5. Process 2 is running and does a JSR R7,A to Subroutine A with PC = PC2.




6. Subroutine A is running and uses stack for temporary storage.




7. Subroutine A releases the temporary storage holding TA1 and TA2.




8. Subroutine A returns control to process 2 with an RTS R7, PC is reset to PC2.



9. Process 2 completes with an RTI instruction (dismisses interrupt) PC is reset to PC1 and status is reset to PS1; process 1 resumes.



10. Process 1 releases the temporary storage holding TE0 and TE1.



11. Process 1 completes its operation with an RTI is reset to PC0 and status is reset to PS0.




Figure 5-14: Nested Interrupt Service Routines and Subroutines


Note that the area of interrupt service programming is intimately involved with the concept of CPU and device priority levels. For a full discussion of the uses of the PDP-11/45 priority structure, refer to Chapter 2, System Architecture.

5.4 REENTRANCY

Further advantages of stack organization becomes apparent in complex situations which can arise in program systems that are engaged in the concurrent handling of several tasks. Such multi-task program environ-

ments may range from relatively simple single-user applications which must manage an intermix of I/O interrupt service and background computation to large complex multi-programming systems which manage a very intricate mixture of executive and multi-user programming situations. In all these applications there is a need for flexibility and time/memory economy. The use of the stack provides this economy and flexibility by providing a method for allowing many tasks to use a single copy of the same routine and a simple, unambiguous method for keeping track of complex program linkages.

The ability to share a single copy of a given program among users or tasks is called reentrancy. Reentrant program routines differ from ordinary subroutines in that it is unnecessary for reentrant routines to finish processing a given task before they can be used by another task. Multiple tasks can be in various stages of completion in the same routine at any time. Thus the following situation may occur:



Programs 1, 2, and 3 can share Subroutine A.

A separate copy of Subroutine A must be provided for each program.

Figure 5-15: Reentrant Routines

The chief programming distinction between a non-shareable routine and a reentrant routine is that the reentrant routine is composed solely of "pure code," i.e., it contains only instructions and constants. Thus, a section of program code is reentrant (shareable) if and only if it is "non self-modifying," that is it contains no information within it that is subject to modification.

Using reentrant routines, control of a given routine may be shared as illustrated in Figure 5-16.




Figure 5-16: Reentrant Routine Sharing

1. Task A has requested processing by Reentrant Routine Q.
2. Task A temporarily relinquishes control (is interrupted) of Reentrant Routine Q before it finishes processing.
3. Task B starts processing in the same copy of Reentrant Routine Q.
4. Task B relinquishes control of Reentrant Routine Q at some point in its processing.
5. Task A regains control of Reentrant Routine Q and resumes processing from where it stopped.

The use of reentrant programming allows many tasks to share frequently used routines such as device interrupt service routines, ASCII-Binary conversion routines, etc. In fact, in a multi-user system it is possible, for instance, to construct a reentrant FORTRAN compiler which can be used as a single copy by many user programs.

As an application of reentrant (shareable) code, consider a data processing program which is interrupted while executing a ASCII-to-Binary subroutine which has been written as a reentrant routine. The same conversion routine is used by the device service routine. When the device servicing is finished, a return from interrupt (RTI) is executed and execution for the processing program is then resumed where it left off inside the same ASCII-to-Binary subroutine.

Shareable routines generally result in great memory saving. It is the hardware implemented stack facility of the PDP-11 that makes shareable or reentrant routines reasonable.

A subroutine may be reentered by a new task before its completion by the previous task as long as the new execution does not destroy any linkage information or intermediate results which belong to the previous programs. This usually amounts to saving the contents of any general purpose registers, to be used and restoring them upon exit. The choice of whether to save and restore this information in the calling program or the subroutine is quite arbitrary and depends on the particular application. For example in controlled transfer situations (i.e. JSR's) a main program which calls a code-conversion utility might save the contents of registers which it needs and restore them after it has regained control, or the code conversion routine might save the contents of registers which it uses and restore them upon its completion. In the case of interrupt service routines this save/restore process must be carried out by the service routine itself since the interrupted program has no warning of an impending interrupt. The advantage of using the stack to save and restore (i.e. "push" and "pop") this information is that it permits a program to isolate its instructions and data and thus maintain its reentrancy.

In the case of a reentrant program which is used in a multi-programming environment it is usually necessary to maintain a separate R6 stack for each user although each such stack would be shared by all the tasks of a given user. For example, if a reentrant FORTRAN compiler is to be shared between many users, each time the user is changed,

R6 would be set to point to a new user's stack area as illustrated in Figure 5-17.




Figure 5-17: Multiple R6 Stack

5.5 POSITION INDEPENDENT CODE—PIC

Most programs are written with some direct references to specific addresses, if only as an offset from an absolute address origin. When it is desired to relocate these programs in memory, it is necessary to change the address references and/or the origin assignments. Such programs are constrained to a specific set of locations. However, the PDP-11 architecture permits programs to be constructed such that they are not constrained to specific locations. These Position Independent programs do not directly reference any absolute locations in memory. Instead all references are "PC-relative" i.e. locations are referenced in terms of offsets from the current location (offsets from the current value of the Program Counter (PC)). When such a program has been translated to machine code it will form a program module which can be loaded anywhere in memory as required.

Position Independent Code is exceedingly valuable for those utility routines which may be disk-resident and are subject to loading in a dynamically changing program environment. The supervisory program may load them anywhere it determines without the need for any relocation parameters since all items remain in the same positions relative to each other (and thus also to the PC).

Linkages to program routines which have been written in position independent code (PIC) must still be absolute in some manner. Since these routines can be located anywhere in memory there must be some fixed or readily locatable linkage addresses to facilitate access to these routines. This linkage address may be a simple pointer located at a fixed address or it may be a complex vector composed of numerous linkage information items.

5.6 RECURSION

It is often meaningful for a program routine to call itself as in the case of calculating a fourth root in FORTRAN with the expression $SQRT(SQRT(X))$. The ability to nest subroutine calls to the same subroutine is called recursion. The use of stack organization permits easy unambiguous recursion. The technique of recursion is of great use to the mathematical analyst as it also permits the evaluation of some otherwise non-computable mathematical functions. Although it is beyond the scope of this chapter to discuss the concept of recursive routines in detail, the reader should realize that this technique often permits very significant memory and speed economies in the linguistic operations of compilers and other higher-level software programs.

5.7 CO-ROUTINES

In some situations it happens that two program routines are highly interactive. Using a special case of the JSR instruction i.e., JSR PC, @(R6)+ which exchanges the top element of the Register 6 processor stack and the contents of the Program Counter (PC), two routines may be permitted to swap program control and resume operation where they stopped, when recalled. Such routines are called "co-routines." This control swapping is illustrated in Figure 5-18.

Routine # 1 is operating, it then executes:

```
MOV #PC2,-(R6)
```

```
JSR PC,@(R6)+
```

with the following results:

- 1) PC2 is popped from the stack and the SP autoincremented
- 2) SP is autodecremented and the old PC (i.e. PC1) is pushed
- 3) control is transferred to the location PC2 (i.e. routine # 2)

Routine # 2 is operating, it then executes:

```
JSR PC,@(R6)+
```

with the result the PC2 is exchanged for PC1 on the stack and control is transferred back to routine # 1.




Figure 5-18—Co-Routine Interaction

MEMORY MANAGEMENT

The PDP-11/45 Memory Management Unit provides the hardware facilities necessary for complete memory management and protection. It is designed to be a memory management facility for systems where the system memory size is greater than 28K words and for multi-user, multi-programming systems where memory protection and relocation facilities are necessary.

In order to most effectively utilize the power and efficiency of the PDP-11/45 in medium and large scale systems it is necessary to run several programs simultaneously. In such multi-programming environments several user programs would be resident in memory at any given time. The task of the supervisory program would be: control the execution of the various user programs, manage the allocation of memory and peripheral device resources, and safeguard the integrity of the system as a whole by careful control of each user program.

In a multi-programming system, the Memory Management Unit provides the means for assigning memory pages to a user program and preventing that user from making any unauthorized access to these pages outside his assigned area. Thus, a user can effectively be prevented from accidental or willful destruction of any other user program or the system executive program.

The basic characteristics of the PDP-11/45 Memory Management Unit are:

- 16 User mode memory pages
- 16 Supervisor mode memory pages
- 16 Kernel mode memory pages
- 8 pages in each mode for instructions
- 8 pages in each mode for data
- page lengths from 32 to 4096 words
- each page provided with full protection and relocation
- transparent operation
- 6 modes of memory access control
- memory extension to 124K words (248K bytes)

6.1 PDP-11 FAMILY BASIC ADDRESSING LOGIC

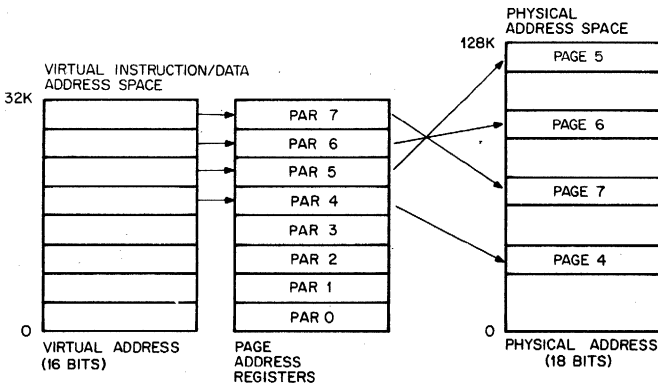
The addresses generated by all PDP-11 Family Central Processor Units (CPUs) are 18-bit direct byte addresses. Although the PDP-11 Family word length and operational logic is all 16-bit length, the UNIBUS and CPU addressing logic actually is 18-bit length. Thus, while the PDP-11 word can only contain address references up to 32K words (64K bytes)

the CPU and UNIBUS can reference addresses up to 128K words (256K bytes). These extra two bits of addressing logic provide the basic framework for expanded memory operation.

In addition to the word length constraint on basic memory addressing space, the uppermost 4K words of address space is always reserved for UNIBUS I/O device registers. In a basic PDP-11/45 memory configuration (without the Memory Management Option) all address references to the uppermost 4K words of 16 bit address space (170000-177777) are converted to full 18-bit references with bits 17 and 16 always set to 1. Thus, a 16 bit reference to the I/O device register at address 173224 is automatically internally converted to a full 18-bit reference to the register at address 773224. Accordingly, the basic PDP-11/45 configuration can directly address up to 28K words of true memory, and 4K words of UNIBUS I/O device registers. Memory configurations beyond this require the PDP-11/45 Memory Management Unit.

6.2 VIRTUAL ADDRESSING

When the PDP-11/45 Memory Management Unit is operating, the normal 16 bit direct byte address is no longer interpreted as a direct Physical Address (PA) but as a Virtual Address (VA) containing information to be used in constructing a new 18-bit physical address. The information contained in the Virtual Address (VA) is combined with relocation information contained in the Page Address Register (PAR) (see 6.4) to yield an 18-bit Physical Address (PA). Using the Memory Management Unit, memory can be dynamically allocated in pages each composed of from 1 to 128 integral blocks of 32 words.



PAR = Page Address Register

Figure 6-1 Virtual Address Mapping into Physical Address

The starting physical address for each page is an integral multiple of 32 words, and each page has a maximum size of 4096 words. Pages may be located anywhere within the 128K Physical Address space. The determination of which set of 16 page registers is used to form a Physical

Address is made by the current mode of operation of the CPU, i.e., Kernel, Supervisor or User mode.

6.3 INTERRUPT CONDITIONS UNDER MEMORY MANAGEMENT CONTROL

The Memory Management Unit relocates all addresses. Thus, when it is enabled, all trap, abort, and interrupt vectors are considered to be in Kernel mode Virtual Address Space. When a vectored transfer occurs, control is transferred according to a new Program Counter (PC) and Processor Status Word (PS) contained in a two-word vector relocated through the Kernel Page Address Register Set. Relocation of trap addresses means that the hardware is capable of recovering from a failure in the first physical bank of memory.

When a trap, abort, or interrupt occurs the "push" of the old PC, old PS is to the User/Supervisor/Kernel R6 stack specified by CPU mode bits 15,14 of the new PS in the vector (bits 15,14: 00 = Kernel, 01 = Supervisor, 11 = User). The CPU mode bits also determine the new PAR set. In this manner it is possible for a Kernel mode program to have complete control over service assignments for all interrupt conditions, since the interrupt vector is located in Kernel space. The Kernel program may assign the service of some of these conditions to a Supervisor or User mode program by simply setting the CPU mode bits of the new PS in the vector to return control to the appropriate mode.

6.4 CONSTRUCTION OF A PHYSICAL ADDRESS

All addresses with memory relocation enabled either reference information in instruction (I) Space or Data (D) Space. I Space is used for all instruction fetches, index words, absolute addresses and immediate operands, D Space is used for all other references. I Space and D Space each have 8 PAR's in each mode of CPU operation, Kernel, Supervisor, and User. Using Status Register #3 (6.6.4), the operating system may select to disable D space and map all references (Instructions and Data) through I space, or to use both I and D space.

The basic information needed for the construction of a Physical Address (PA) comes from the Virtual Address (VA), which is illustrated in Figure 6-2, and the appropriate PAR set.




Figure 6-2: Interpretation of a Virtual Address

The Virtual Address (VA) consists of:

1. The Active Page Field (APF). This 3-bit field determines which of eight Page Address Registers (PAR0-PAR7) will be used to form the Physical Address (PA).
2. The Displacement Field (DF). This 13-bit field contains an address relative to the beginning of a page. This permits page lengths up to

4K words ($2_{13} = 8K$ bytes). The DF is further subdivided into two fields as shown in Figure 6-3.




Figure 6-3: Displacement Field of Virtual Address

The Displacement Field (DF) consists of:

1. The Block Number (BN). This 7-bit field is interpreted as the block number within the current page.
2. The Displacement in Block (DIB). This 6-bit field contains the displacement within the block referred to by the Block Number (BN).

The remainder of the information needed to construct the Physical Address comes from the 12-bit Page Address Field (PAF) (part of the Page Address Register (PAR)) and specifies the starting address of the memory page which that PAR describes. The PAF is actually a block number in the physical memory, e.g. $PAF = 3$ indicates a starting address of 96 (3×32) words in physical memory.

The formation of a physical address (PA) takes 90 ns. Thus in situations which do not require the facilities of the Memory Management Unit, it should be disabled to permit time savings.

The formation of the Physical Address (PA) is illustrated in Figure 6-4.

The logical sequence involved in constructing a Physical Address (PA) is as follows:

1. Select a set of Page Address Registers depending on the space being referenced.
2. The Active Page Field (APF) of the Virtual Address is used to select a Page Address Register (PAR0-PAR7).
3. The Page Address Field (PAF) of the selected Page Address Register (PAR) contains the starting address of the currently active page as a block number in physical memory.
4. The Block Number (BN) from the Virtual Address (VA) is added to the block number from the Page Address Field (PAF) to yield the number of the block in physical memory (PBN-Physical Block Number) which will contain the Physical Address (PA) being constructed.
5. The Displacement in Block (DIB) from the Displacement Field (DF) of the Virtual Address (VA) is joined to the Physical Block Number (PBN) to yield a true 18-bit PDP-11/45 Physical Address (PA).

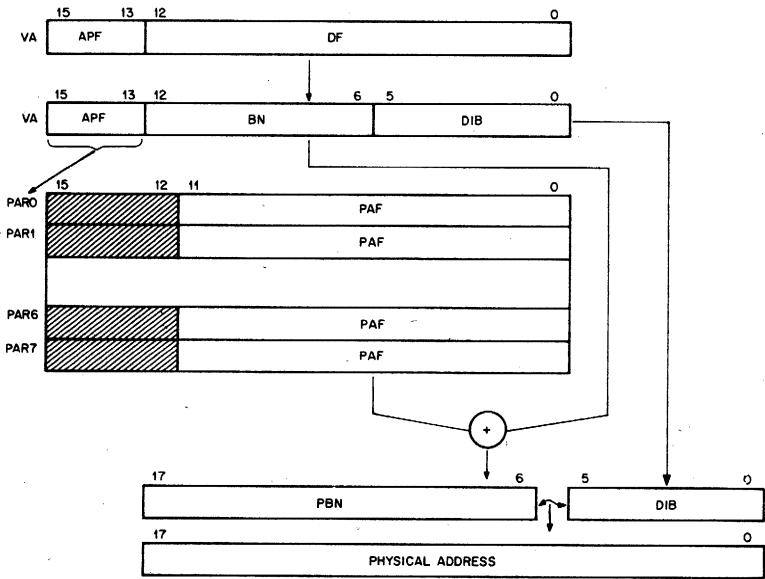


Figure 6-4: Construction of a Physical Address

6.5 MANAGEMENT REGISTERS

The PDP-11/45 Memory Management Unit implements three sets of 32 sixteen bit registers. One set of registers is used in Kernel mode, another in Supervisor, and the other in User mode. The choice of which set is to be used is determined by the current CPU mode contained in the Processor Status word. Each set is subdivided into two groups of 16 registers. One group is used for references to Instruction (I) Space, and one to Data (D) Space. The I Space group is used for all instruction fetches, index words, absolute addresses and immediate operands. The D Space group is used for all other references, providing it has not been disabled by Status Register #3 (6.6.4). Each group is further subdivided into two parts of 8 registers. One part is the Page Address Register (PAR) whose function has been described in previous paragraphs. The other part is the Page Descriptor Register (PDR). PARs and PDRs are always selected in pairs by the top three bits of the virtual address. A PAR/PDR pair contain all the information needed to describe and locate a currently active memory page.

The various Memory Management Registers are located in the uppermost 4K of PDP-11 physical address space along with the UNIBUS I/O device registers. For the actual addresses of these registers refer to Paragraph 6.9, Memory Management Unit—Register Map.

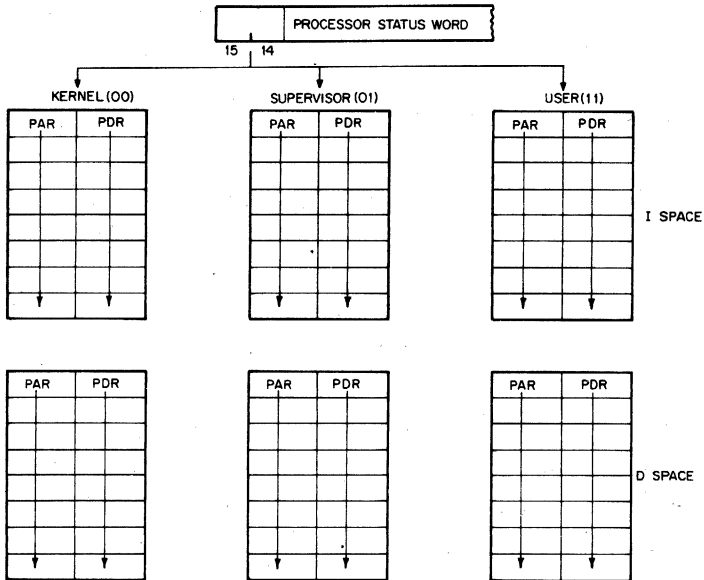


Figure 6-5: Active Page Registers

6.5.1 Page Address Registers (PAR)

The Page Address Register (PAR) contains the Page Address Field (PAF), a 12-bit field, which specifies the starting address of the page as a block number in physical memory.




Figure 6-6: Page Address Register

Bits 15-12 of the PAR are unused and reserved for possible future use.

The Page Address Register (PAR) which contains the Page Address Field (PAF) may be alternatively thought of as a relocation register containing a relocation constant, or as a base register containing a base address. Either interpretation indicates the basic importance of the Page Address Register (PAR) as a relocation tool.

6.5.2 Page Descriptor Register

The Page Descriptor Register (PDR) contains information relative to page expansion, page length, and access control.




Figure 6-7: Page Description Register

6.5.2.1 Access Control Field (ACF)

This three-bit field, occupying bits 2-0 of the Page Descriptor Register (PDR) contains the access rights to this particular page. The access codes or "keys" specify the manner in which a page may be accessed and whether or not a given access should result in a trap or an abort of the current operation. A memory reference which causes an abort is not completed while a reference causing a trap is completed. In fact, when a memory reference causes a trap to occur, the trap does not occur until the entire instruction has been completed. Aborts are used to catch "missing page faults," prevent illegal access, etc.; traps are used as an aid in gathering memory management information.

In the context of access control the term "write" is used to indicate the action of any instruction which modifies the contents of any addressable word. "Write" is synonymous with what is usually called a "store" or "modify" in many computer systems.

The modes of access control are as follows:

000	non-resident	abort all accesses
001	read-only	abort on write attempt memory management trap on read
010	read-only	abort on write attempt
011	unused	abort all accesses—reserved for future use
100	read/write	memory management trap upon completion of a read or write
101	read/write	memory management trap upon completion of a write
110	read/write	no system trap/abort action
111	unused	abort all accesses—reserved for future use

It should be noted that the use of I Space provides the user with a further form of protection, execute only.

6.5.2.2 Access Information Bits

A Bit (bit 7)—This bit is used by software to determine whether or not any accesses to this page met the trap condition specified by the Access Control Field (ACF). (A = 1 is Affirmative) The A Bit is used in the process of gathering memory management statistics.

W Bit (bit 6)—This bit indicates whether or not this page has been modified (i.e. written into) since either the PAR or PDR was loaded. (W = 1 is Affirmative) The W Bit is useful in applications which involve disk swapping and memory overlays. It is used to determine which pages have been modified and hence must be saved in their new form and which pages have not been modified and can be simply overlaid.

Note that A and W bits are “reset” to “0” whenever either PAR or PDR is modified (written into).

6.5.2.3 Expansion Direction (ED)

This one-bit field, located at bit 3 of the Page Descriptor Register (PDR), specifies whether the page expands upward from relative zero (ED = 0) or downwards toward relative zero (ED = 1). Relative zero, in this case, is the PAF (Page Address Field). Expansion is done by changing the Page Length Field (6.5.2.4). In expanding upwards, blocks with higher relative addresses are added; in expanding downwards, blocks with lower relative addresses are added to the page. Upward expansion is usually used to add more program space, while downward expansion is used to add more stack space.

6.5.2.4 Page Length Field (PLF)

The seven-bit field, occupying bits 14-8 of the Page Descriptor Register (PDR), specifies the number of blocks in the page. A page consists of at least one and at most 128 blocks, and occupies contiguous core locations. If the page expands upwards, this field contains the length of the page minus one (in blocks). If the page expands downwards, this field contains 128 minus the length of the page (in blocks).

A Length Error occurs when the Block Number (BN) of the virtual address (VA) is greater than the Page Length Field (PLF), if the page expands upwards, or if the page expands downwards, when the BN is less than the PLF.

6.5.2.5 Reserved Bits

Bits 15, 4 and 5 are reserved for future use, and are always 0.

6.6 FAULT RECOVERY REGISTERS

Aborts and traps generated by the Memory Management hardware are vectored through Kernel virtual location 250, Status Registers #0, #1, #2 and #3 are used in order to differentiate an abort from a trap, determine why the abort or trap occurred, and allow for easy program restarting. Note that an abort or trap to a location which is itself an invalid address will cause another abort or trap. Thus the Kernel program must insure that Kernel Virtual Address 250 is mapped into a valid address, otherwise a loop will occur which will require console intervention.

6.6.1 Status Register #0 (SR0) (status and error indicators)

SR0 contains error flags, the page number whose reference caused the abort, and various other status flags. The register is organized as shown in Figure 6-8.




Figure 6-8: Format of Status Register #0 (SR0)

Bits 15-12 are the error flags. They may be considered to be in a "priority queue" in that "flags to the right" are less significant and should be ignored. That is, a "non-resident" fault service routine would ignore length, access control, and memory management flags. A "page length" service routine would ignore access control and memory management faults, etc.

Bits 15-13 when set (error conditions) cause Memory Management to freeze the contents of bits 1-7 and Status Registers #1 and #2. This has been done to facilitate error recovery (discussed in 6.6.5).

Bits 15-12 are enabled by a signal called "RELOC." "RELOC." is true when an address is being relocated by the Memory Management unit. This implies that either SR0, bit 0 is equal to 1 (relocation operating) or that SR0, bit 8 (MAINTENANCE) is equal to 1 and the memory reference is the final one of a destination calculation (maintenance/destination mode).

Note that Status Register #0 (SR0) bits 0, 8, and 9 can be set under program control to provide meaningful control information. However, information written into all other bits is not meaningful. Only that information which is automatically written into these remaining bits as a result of hardware actions is useful as a monitor of the status of the Memory Management Unit. Setting bits 15-12 under program control will not cause traps to occur; these bits however must be reset to 0 after an abort or trap has occurred in order to resume status monitoring.

6.6.1.1 Abort—Non-Resident

Bit 15 is the "Abort—Non-Resident" bit. It is set by attempting to access a page with an Access Control Field (ACF) key equal to 0, 3, or 7. It is also set by attempting to use Memory Relocation with a processor mode of 2.

6.6.1.2 Abort—Page Length

Bit 14 is the "Abort Page Length" bit. It is set by attempting to access a location in a page with a block number (Virtual Address bits, 12-6) that is outside the area authorized by the Page Length Field (PLF) of the Page Descriptor Register (PDR) for that page. Bits 14 and 15 may be set simultaneously by the same access attempt.

6.6.1.3 Abort—Read Only

Bit 13 is the "Abort—Read Only" bit. It is set by attempting to write in a "Read-Only" page. "Read-Only" pages have access keys of 1 or 2.

6.6.1.4 Trap—Memory Management

Bit 12 is the "Trap—Memory Management" bit. It is set by a read operation which references a page with an Access Control Field (ACF) of 1 or 4, or by a write operation to a page with an ACF key of 4 or 5.

6.6.1.5 Bits 11, 10

Bits 11 and 10 are spare locations and are always equal to 0. They are unused and reserved for possible future expansion.

6.6.1.6 Enable Memory Management Traps

Bit 9 is the "Enable Memory Management Traps" bit. It can be set or cleared by doing a direct write into SRO. If bit 9 is 0, no Memory Management traps will occur. The A and W bits will, however, continue to log potential Memory Management Traps. When bit 9 is set to 1, the next "potential" Memory Management trap will cause a trap, vectored through Kernel Virtual Address 250.

Note that if an instruction which sets bit 9 to 0 (disable Memory Management Trap) causes a potential Memory Management trap in the course of any of its memory references prior to the one actually changing SRO, then the trap will occur at the end of the instruction anyway.

6.6.1.7 Maintenance/Destination Mode

Bit 8 specifies Maintenance use of the Memory Management Unit. It is provided for diagnostic purposes only and must not be used for other purposes.

6.6.1.8 Instruction Completed

Bit 7 indicates that the current instruction has been completed. It will be set to 0 during T bit, Parity, Odd Address, and Time Out traps and interrupts. This provides error handling routines with a way of determining whether the last instruction will have to be repeated in the course of an error recovery attempt. Bit 7 is Read-Only (it cannot be written). It is initialized to a 1. Note that EMT, TRAP, BPT, and IOT do not set bit 7.

6.6.1.9 Processor Mode

Bits 5, 6 indicate the CPU mode (User/Supervisor/Kernel) associated with the page causing the abort. (Kernel = 00, Supervisor = 01, User = 11). If an illegal mode (10) is specified, bit 15 will be set and an abort will occur.

6.6.1.10 Page Address Space

Bit 4 indicates the type of address space (I or D) the Unit was in when a fault occurred (0 = I Space, 1 = D Space). It is used in conjunction with bits 3-1, Page Number.

6.6.1.11 Page Number

Bits 3-1 contain the page number of a reference causing a Memory Management fault. Note that pages, like blocks, are numbered from 0 upwards.

6.6.1.12 Enable Relocation

Bit 0 is the "Enable Relocation" bit. When it is set to 1, all addresses are relocated by the unit. When bit 0 is set to 0 the Memory Management Unit is inoperative and addresses are not relocated or protected.

6.6.2 Status Register #1 (SR1)

SR1 records any autoincrement/decrement of the general purpose registers, including explicit references through the PC. SR1 is cleared at the beginning of each instruction fetch. Whenever a general purpose register is either autoincremented or autodecremented the register number and the amount (in 2s complement notation) by which the register was modified, is written into SR1.

The information contained in SR1 is necessary to accomplish an effective recovery from an error resulting in an abort. The low order byte is written first and it is not possible for a PDP-11 instruction to autoincrement/decrement more than two general purpose registers per instruction before an "abort-causing" reference. Register numbers are recorded "MOD 8"; thus it is up to the software to determine which set of registers (User/Supervisor/Kernel—General Set 0/General Set 1) was modified, by determining the CPU and Register modes as contained in the PS at the time of the abort. The 6-bit displacement on R6(SP) that can be caused by the MARK instruction cannot occur if the instruction is aborted.




Figure 6-9: Format of Status Register #1 (SR1)

6.6.3 Status Register #2

SR2 is loaded with the 16-bit Virtual Address (VA) at the beginning of each instruction fetch, or with the address Trap Vector at the beginning of an interrupt, "T" Bit trap, Parity, Odd Address, and Timeout traps. Note that SR2 does not get the Trap Vector on EMT, TRAP, BPT and IOT instructions. SR2 is Read-Only; it can not be written. SR2 is the Virtual Address Program Counter.

6.6.4 Status Register #3

The Status Register #3 (SR3) enables or disables the use of the D space PAR's and PDR's. When D space is disabled, all references use the I space registers; when D space is enabled, both the I space and D space registers are used. Bit 0 refers to the User's Registers, Bit 1 to the Supervisor's, and Bit 2 to the Kernel's. When the appropriate bits are set D space is enabled; when clear, it is disabled. Bits 3-15 are unused. On initialization this register is set to 0 and only I space is in use.




Figure 6-10: Format of Status Register #3 (SR3)

6.6.5 Instruction Back-Up/Restart Recovery

The process of "backing-up" and restarting a partially completed instruction involves:

1. Performing the appropriate memory management tasks to alleviate the cause of the abort (e.g. loading a missing page, etc.)
2. Restoring the general purpose registers indicated in SR1 to their original contents at the start of the instruction by subtracting the "modify value" specified in SR1.
3. Restoring the PC to the "abort-time" PC by loading R7 with the contents of SR2, which contains the value of the Virtual PC at the time the "abort-generating" instruction was fetched.

Note that this back-up/restart procedure assumes that the general purpose register used in the program segment will not be used by the abort recovery routine. This is automatically the case if the recovery program uses a different general register set.

6.6.6 Clearing Status Registers Following Trap/Abort

At the end of a fault service routine bits 15-12 of SR0 must be cleared (set to 0) to resume error checking. On the next memory reference following the clearing of these bits, the various Status Registers will resume monitoring the status of the addressing operations (SR2), will be loaded with the next instruction address, SSR1 will store register change information and SR0 will log Memory Management Status information.

6.7 EXAMPLES

6.7.1 Normal Usage

The Memory Management Unit provides a very general purpose memory management tool. It can be used in a manner as simple or complete as desired. It can be anything from a simple memory expansion device to a very complete memory management facility.

The variety of possible and meaningful ways to utilize the facilities offered by the Memory Management Unit means that both single-user and multi-programming systems have complete freedom to make whatever memory management decisions best suit their individual needs. Although a knowledge of what most types of computer systems seek to achieve may indicate that certain methods of utilizing the Memory Management Unit will be more common than others, there is no limit to the ways to use these facilities.

In most normal applications, it is assumed that the control over the actual memory page assignments and their protection resides in a supervisory type program which would operate at the nucleus of a CPU's executive (Kernel mode). It is further assumed that this Kernel mode program would set access keys in such a way as to protect itself from willful or accidental destruction by other Supervisor mode or User mode programs. The facilities are also provided such that the nucleus can dynamically assign memory pages of varying sizes in response to system needs.

6.7.2 Typical Memory Page

When the Memory Management Unit is enabled, the Kernel mode program, a Supervisor mode program and a User mode program each have eight active pages described by the appropriate Page Address Registers and Page Descriptor Registers for data, and eight, for instructions. Each segment is made up of from 1 to 128 blocks and is pointed to by the Page Address Field (PAF) of the corresponding Page Address Register (PAR) as illustrated in Figure 6-11.




Figure 6-11: Typical Memory Page

The memory segment illustrated in Figure 6-11 has the following attributes:

1. Page Length: 40 blocks.
2. Virtual Address Range: 140000—144777.
3. Physical Address Range: 312000—316777.

4. No trapped access has been made to this page.
5. Nothing has been modified (i.e. written) in this page.
6. Read-Only Protection.
7. Upward Expansion.

These attributes were determined according to the following scheme:

1. Page Address Register (PAR6) and Page Descriptor Register (PDR6) were selected by the Active Page Field (APF) of the Virtual Address (VA). (Bits 15-13 of the VA = 6_8 .)
2. The initial address of the page was determined from the Page Address Field (PAF) of APR6 ($312000 = 3120_8 \text{ blocks} \times 40_8 \text{ (} 32_{10} \text{) words per block} \times 2 \text{ bytes per word}$).

Note that the PAR which contains the PAF constitutes what is often referred to as a base register containing a base address or a relocation register containing relocation constant.

3. The page length ($47_8 + 1 = 40_{10}$ blocks) was determined from the Page Length Field (PLF) contained in Page Descriptor Register PDR6. Any attempts to reference beyond these 40_{10} blocks in this page will cause a "Page Length Error," which will result in an abort, vectored through Kernel Virtual Address 250.
4. The Physical Addresses were constructed according to the scheme illustrated in Figure 6-4.
5. The Access bit (A-bit) of PDR6 indicates that no trapped access has been made to this page (A bit = 0). When an illegal or trapped reference, (i.e. a violation of the Protection Mode specified by the Access Control Field (ACF) for this page), or a trapped reference (i.e. Read in this case), occurs, the A-bit will be set to a 1.
6. The Written bit (W-bit) indicates that no locations in this page have been modified (i.e. written). If an attempt is made to modify any location in this particular page, an Access Control Violation Abort will occur. If this page were involved in a disk swapping or memory overlay scheme, the W-bit would be used to determine whether it had been modified and thus required saving before overlay.
7. This page is Read-Only protected; i.e. no locations in this page may be modified. In addition, a memory management trap will occur upon completion of a read access. The mode of protection was specified by the Access Control Field (ACF) of PDR6.
8. The direction of expansion is upward (ED = 0). If more blocks are required in this segment, they will be added by assigning blocks with higher relative addresses.

Note that the various attributes which describe this page can all be determined under software control. The parameters describing the page are all loaded into the appropriate Page Address Register (PAR) and Page Descriptor Register (PDR) under program control. In a normal applica-

tion it is assumed that the particular page which itself contains these registers would be assigned to the control of a supervisory type program operating in Kernel mode.

6.7.3 Non-Consecutive Memory Pages

It should be noted at this point that although the correspondence between Virtual Addresses (VA) and PAR/PDR pairs is such that higher VAs have higher PAR/PDR's, this does not mean that higher Virtual Addresses (VA) necessarily correspond to higher Physical Addresses (PA). It is quite simple to set up the Page Address Fields (PAF) of the PAR's in such a way that higher Virtual Address blocks may be located in lower Physical Address blocks as illustrated in Figure 6-12.




Figure 6-12: Non-Consecutive Memory Pages

Note that although a single memory page must consist of a block of contiguous locations, memory pages as macro units do not have to be located in consecutive Physical Address (PA) locations. It also should be realized that the assignment of memory pages is not limited to consecutive non-overlapping Physical Address (PA) locations.

6.7.4 Stack Memory Pages

When constructing PDP-11/45 programs it is often desirable to isolate all program variables from "pure code" (i.e. program instructions) by placing them on a register indexed stack. These variables can then be "pushed" or "popped" from the stack area as needed (see Chapter 3, Addressing Modes). Since all PDP-11 Family stacks expand by adding

locations with lower addresses, when a memory page which contains "stacked" variables needs more room it must "expand down," i.e. add blocks with lower relative addresses to the current page. This mode of expansion is specified by setting the Expansion Direction (ED) bit of the appropriate Page Descriptor Register (PDR) to a 1. Figure 6-13 illustrates a typical "stack" memory page. This page will have the following parameters:

PAR6: PAF = 3120

PDR6: PLF = 175_8 or 125_{10} ($128_{10}-3$)

ED = 1

A = 0 or 1

W = 0 or 1

ACF = nnn (to be determined by programmer as the need dictates).

note: the A, W bits will normally be set by hardware.




Figure 6-13: Typical Stack Memory Page

In this case the stack begins 128 blocks above the relative origin of this memory page and extends downward for a length of three blocks. A "PAGE LENGTH ERROR" abort vectored through Kernel Virtual Address (VA) 250 will be generated by the hardware when an attempt is made to reference any location below the assigned area, i.e. when the Block Number (BN) from the Virtual Address (VA) is less than the Page Length Field (PLF) of the appropriate Page Descriptor Register (PDR).

6.8 TRANSPARENCY

It should be clear at this point that in a multiprogramming application it is possible for memory pages to be allocated in such a way that a particular program seems to have a complete 32K basic PDP-11/45 memory configuration. Using Relocation, a Kernel Mode supervisory-type program can easily perform all memory management tasks in a manner entirely transparent to a Supervisor or User mode program. In effect, a PDP-11/45 System can utilize its resources to provide maximum throughput and response to a variety of users each of which seems to have a powerful system "all to himself."

6.9 MEMORY MANAGEMENT UNIT—REGISTER MAP

REGISTER	ADDRESS
Status Register #0(SR0)	777572
Status Register #1(SR1)	777574
Status Register #2(SR2)	777576
Status Register #3(SR3)	772516
User I Space Descriptor Register (UISDR0)	777600
.	.
User I Space Descriptor Register (UISDR7)	777616
User D Space Descriptor Register (UDSDR0)	777620
.	.
User D Space Descriptor Register (UDSDR7)	777636
User I Space Address Register (UISAR0)	777640
.	.
User I Space Address Register (UISAR7)	777656
User D Space Address Register (UDSAR0)	777660
.	.
User D Space Address Register (UDSAR7)	777676
Supervisor I Space Descriptor Register (SISDR0)	772200
.	.
Supervisor I Space Descriptor Register (SISDR7)	772216
Supervisor D Space Descriptor Register (SDSDR0)	772226
.	.
Supervisor D Space Descriptor Register (SDSDR7)	772236
Supervisor I Space Address Register (SISAR0)	772240
.	.
Supervisor I Space Address Register (SISAR7)	772256

REGISTER	ADDRESS
Supervisor D Space Address Register (SDSAR0)	772260
.	.
.	.
Supervisor D Space Address Register (SDSDR7)	772276
Kernel I Space Descriptor Register (KISDR0)	772300
.	.
.	.
Kernel I Space Descriptor Register (KIDSR7)	772316
Kernel D Space Descriptor Register (KDSDR0)	772320
.	.
.	.
Kernel D Space Descriptor Register (KDSDR7)	772336
Kernel I Space Address Register (KISAR0)	772340
.	.
.	.
Kernel I Space Address Register (KISAR7)	772356
Kernel D Space Address Register (KDSAR0)	772360
.	.
.	.
Kernel D Space Address Register (KDSAR7)	772376

FLOATING POINT PROCESSOR

7.1 INTRODUCTION

The PDP-11 Floating Point Processor is an optional arithmetic processor which fits integrally into the PDP-11/45 Central Processor. It performs all floating point arithmetic operations and converts data between integer and floating point formats.

The hardware provides a time and money saving alternative to the use of software floating point routines. Its use can result in many orders of magnitude improvement in the execution of arithmetic operations.

The features of the unit are:

- Overlapped operation with central processor
- High speed
- Single and double precision (32 or 64 bit) floating point modes
- Flexible addressing modes
- Six 64-bit floating point accumulators
- Error recovery aids.

7.2 OPERATION

The Floating Point Processor is an integral part of the Central Processor. It operates using similar address modes, and the same memory management facilities provided by the Memory Management Option, as the Central Processor. Floating Point Processor instructions can reference the floating point accumulators, the Central Processor's general registers, or any location in memory.

When, in the course of a program, an FPP Instruction is fetched from memory, the FPP will execute that instruction in parallel with the CPU continuing with its instruction sequence. The CPU is delayed a very short period of time during the FPP instruction's Fetch operation, and then is free to proceed independently of the FPP. The interaction between the two processors is automatic, and a program can take full advantage of the parallel operation of the two processors by intermixing Floating Point Processor and Central Processor instructions.

Interaction between Floating Point Processor and Central Processor instructions is automatically taken care of by the hardware. When an FPP instruction is encountered in a program, the machine first checks the status of the Floating Point Processor. If the FPP is "busy," the CPU will wait until it is "done" before continuing execution of the program.

LDD (R3)+,AC3 ;Pick up constant operand and place it in AC3.

ADDLP: LDD (R3)+,AC0 ;Load AC0 with next value in table

```

MUL AC3,AC0      ;and multiply by constant in AC3
ADDD AC0,AC1     ;and add the result into AC1
SOB R5,ADDLP    ;check to see whether done
STCDI AC1,@R4   ;done, convert double to integer and store

```

In the above example the Floating Point Processor would execute the next three instructions. After the "ADDD" was fetched into the FPP, the CPU would execute the "SOB" and then wait for the FPP to be "done" with the "ADDD" before giving it the "LDD" or "STCDI" instruction.

As can be seen from this example, autoincrement and autodecrement addressing automatically adds or subtracts the correct amount to the contents of the register depending on the modes represented by the instruction.

7.3 ARCHITECTURE

The Floating Point Processor contains scratch registers, a Floating Exception Address pointer (FEA), a Program Counter, a set of Status and Error Registers, and six general purpose accumulators (AC0-AC5).

Each accumulator is interpreted to be 32 or 64 bits long depending on the instruction and the status of the Floating Point Processor. For 32-bit instructions only the left-most 32 bits are used, while the remaining 32 bits remain unaffected.

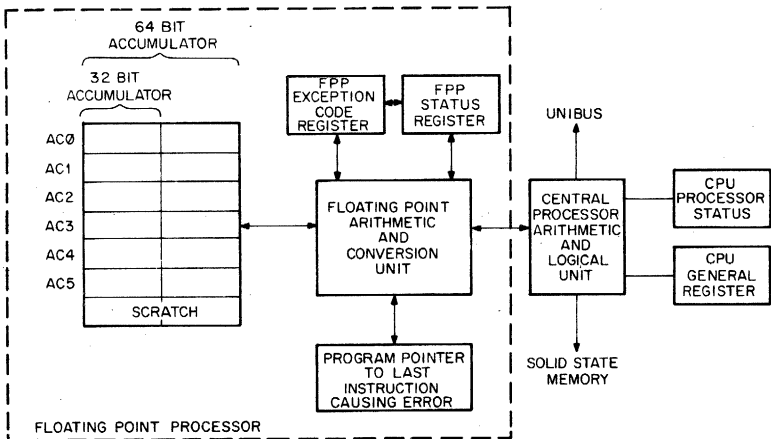



Figure 7-1: Floating Point Processor

The six Floating Point Accumulators are used in numeric calculations and interaccumulator data transfers; the first four (AC0-AC3) are also used for all data transfers between the FPP and the General Registers or Memory.


7.4 FLOATING POINT DATA FORMATS

The FPP handles two types of floating point data: Single Precision or Floating Mode (F) which is 32 bits long, and Double Precision (D) which is 64 bits long. The exponent is stored in excess 128 (200_8) notation. Exponents from -128 to $+127$ are therefore represented by the binary equivalent of 0 to 255 ($0-377_8$). Fractions are represented in sign-magnitude notation with the binary radix point to the left. Numbers are assumed to be normalized and, therefore, the most significant bit is not stored because it is redundant. It is always a 1 except where the exponent is zero, then the complete number is declared to be 0.

F Formats:



D Formats:



S = Sign of Fraction

EXP = Exponent in excess 200_8 notation


FRACTION = 23 bits in F Format, 55 bits in D Format, + one hidden bit (normalization). Binary Radix point to the left.

The results of a Floating Point operation may be either truncated or rounded off. "Rounding" rounds away from zero and thus increases the absolute value of the number.


The FPP provides for conversion of Floating Point to Integer Format and vice-versa. The processor thus recognizes single precision integer (I) and double precision integer long (L) numbers.

The numbers are stored in standard two's complement form.

I Format:



L Format:



S = Sign of Number

NUMBER = 15 bits in I Format, 31 bits in L Format.

7.5 FLOATING POINT UNIT STATUS REGISTER

This register provides mode control for the floating point unit, as well as the condition code and error recovery information from the execution of the previous instruction.

Four bits control the modes of operation:

Single/Double—Floating Point numbers can be either single or double precision.

Long/Short—Integer numbers can be 16 bits or 32 bits long.

Truncate/Round—The result of Floating Point operation can be either truncated or rounded off.

Normal/Maintenance—a special maintenance mode is available.

There are four condition codes:

Carry, overflow, zero, and negative, which are equivalent to the CPU condition codes, and five error interrupts which can be disabled individually or as a group.

FER	FID	UNUSED	FIUV	FIU	FIV	FIC	FD	FL	FT	FMM	FN	FZ	FV	FC	
15	14	13	12	11	10	9	8	7	6	5	4	3	2	1	0

BIT	NAME	DESCRIPTION
15	Floating Error (FER)	Floating Point Error flag. The result of the last operation resulted in a Floating Point Exception and the individual interrupt (FIUV, FIU, FIV, FIC) was enabled.
14	Interrupt Disable (FID)	All FPP interrupts disabled when this bit is set.
13	Not Used	
12	Not Used	
11	Interrupt on Undefined Variable (FIUV)	When set and a -0 is obtained from memory, an interrupt will occur. When clear, -0 can be loaded and used in any arithmetic operation.
10	Interrupt on Underflow (FIU)	When set, Floating Underflow will cause an interrupt. The result of the operation, causing the interrupt, will be correct except for the exponent which will be off by $+400_8$. If the bit is reset and the underflow occurs, the result will be set to zero.

BIT	NAME	DESCRIPTION
9	Interrupt on Overflow (FIV)	When set, Floating Overflows will cause an interrupt. The result of the operation causing the interrupt will be correct except for the exponent which will be off by $+400_8$. If the bit is reset, the result of the operation will be the same as detailed above but no interrupt will occur.
8	Interrupt on Integer Conversion Error (FIC)	When set, and the STCFI (Store and Convert Floating to Integer) instruction causes FC to be set, an interrupt will occur. If the interrupt occurs, the destination is set to 0 and all other registers are left untouched. If the bit is reset, the result of the operation will be the same as detailed above, but no interrupt will occur.
7	Floating Double Precision Mode (FD)	Determines the precision that is used for Floating Point calculations. When set, Double precision is assumed; when reset Floating precision is used.
6	Floating Long Integer Mode (FL)	Active in conversion between Integer and Floating Point format. When set, the Integer format assumed is Double Precision two's complement (i.e. 31 bits + sign). When reset, the integer format is assumed to be Single Precision two's complement (i.e. 15 bits + sign).
5	Floating Truncate Mode (FT)	When set, causes the result of any arithmetic operation to be truncated. When reset, the results are rounded.
4	Floating Maintenance Mode (FMM)	
3	Floating Negative (FN)	The result of the last operation was negative.
2	Floating Zero (FZ)	The result of the last operation was zero.
1	Floating Overflow (FV)	The result of the last operation resulted in an arithmetic overflow.

BIT	NAME	DESCRIPTION
0	Floating Carry (FC)	The result of the last operation resulted in a carry of the most significant bit. This can only occur in integer-Floating conversions.

7.6 FEC REGISTER: ERROR DETECTION

One Interrupt vector is assigned to take care of all floating point exceptions (location 244). The eight possible errors causing the trap are coded in a four bit register, the FPP's Exception Code, "FEC," Register.

The error assignments are as follows:

0	Not used
2	Floating OP Code Error
4	Floating Divide by Zero
6	Floating Integer Conversion Error
8	Floating Overflow
10	Floating Underflow
12	Floating Undefined Variable
14	Maintenance Trap

7.7 FLOATING POINT PROCESSOR INSTRUCTION ADDRESSING

Floating Point Processor instructions use the same type of addressing as the Central Processor instructions. A source or destination operand is specified by designating one of eight addressing modes and one of eight central processor general registers to be used in the specified mode. The modes of addressing are the same as those of the central processor except for mode 0. In mode 0 the operand is located in the designated Floating Point Processor Accumulator, rather than in a Central processor general register. The modes of addressing:

- 0 = Direct Accumulator
- 1 = Deferred
- 2 = Auto-increment
- 3 = Auto-increment deferred
- 4 = Auto-decrement
- 5 = Auto-decrement deferred
- 6 = Indexed
- 7 = Indexed deferred

Autoincrement and autodecrement operate on increments and decrements of 4 for F Format and 10_8 for D Format.

In mode 0, the user can make use of all six FPP accumulators (AC0—AC5) as his source or destination. In all other modes, which involve transfer of data from memory or the general register, the user is restricted to the first four FPP accumulators (AC0—AC3).

In immediate addressing (Mode 2, R7) only 16 bits are loaded or stored.

7.8 FLOATING POINT PROCESSOR INSTRUCTION TIMING

The following table represents execution times of the Floating Point Processor for AC—AC operations (Address Mode 0).

Instruction	Execution Time in Microseconds			
	Single Precision		Double Precision	
	min	max	min	max
ADDX	3.2	4.6	3.8	6.2
SUBX	3.2	4.6	3.8	6.2
MULX	4.6	6.6	6.6	12.0
DIVX	4.6	10.0	6.6	18.4

Minimum and maximum times for FPP instruction execution are given in this chapter using the definition Mode 0-AC operations operating in bipolar memory.


For more detailed information on FPP instruction timing and FPP/CPU interaction consult Appendix B.

7.9 FLOATING POINT INSTRUCTIONS


Each instruction that references a floating point number can operate on floating or double precision numbers depending on the state of the FD mode bit. In a similar fashion, there is a mode bit FL that determines whether a 32-bit integer (FL = 1) or a 16-bit integer (FL = 0) is used in conversion between integer and floating point representation. FSRC and FDST use floating point addressing modes, SRC and DST use CPU addressing Modes.

Floating Point Instruction Format

Double Operand Addressing



Single Operand Addressing



OC = Op Code = 17

FOC = Floating Op Code

AC = Accumulator

FSRC, FDST use FPP Address Modes

SRC, DST use CPU Address Modes

General Definitions:

XL = largest fraction that can be represented:

$$1-2^{-24}; \text{FD} = 0$$

$$1-2^{-56}; \text{FD} = 1$$

XLL = smallest number that is not identically zero 2^{-128}

XUL = largest number that can be represented: $2^{127} * XL$

JL = largest integer that can be represented:

$$2^{15}-1 \text{ If FL} = 0$$


$$2^{31}-1 \text{ If FL} = 1$$

2.77 μ s
2.97 μ s

LDF
LDD

Load Floating/Double

172(AC + 4)FSRC



Operation: AC \leftarrow (FSRC)

Condition Codes: FC \leftarrow 0
FV \leftarrow 0
FZ \leftarrow 1 if (AC) = 0 else FZ \leftarrow 0
FN \leftarrow 1 if (AC) < 0 else FN \leftarrow 0


Description: Load Single or Double Precision Number into Accumulator

2.15 μ s
2.15 μ s

STF STD

Store Floating/Double

174ACFDST



Operation: FDST \leftarrow (AC)

Condition Codes:
FC \leftarrow FC
FV \leftarrow FV
FZ \leftarrow FZ
FN \leftarrow FN


Description: Store Single or Double Precision Number from Accumulator

5.17 μ s
6.47 μ s

ADDF
ADDD

Add Floating/Double

172ACFSRC



Operation: $AC \leftarrow (AC) + (FSRC)$ If $[(AC) + (FSRC)] > XLL$ or $FIU = 1$, else $AC \leftarrow 0$

Condition Codes: $FC \leftarrow 0$
 $FV \leftarrow 1$ If $(AC) > XUL$ else $FV \leftarrow 0$
 $FZ \leftarrow 1$ If $(AC) = 0$ else $FZ \leftarrow 0$
 $FN \leftarrow 1$ If $(AC) < 0$ else $FN \leftarrow 0$


Description: Add the contents of FSRC to the contents of accumulator. In Single or Double Precision result is in accumulator unless Underflow occurs and the interrupt is not enabled; in this case AC is set to 0.

5.17 μ s
 6.47 μ s

SUBF SUBD

Subtract Floating/Double

173ACFSRC



Operation: $AC \leftarrow (AC) - (FSRC)$ If $[(AC) - (FSRC)] - XLL$ or $FIU = 1$ else $AC \leftarrow 0$

Condition Codes: $FC \leftarrow 0$
 $FV \leftarrow 1$ If $(AC) > XUL$ else $FV \leftarrow 0$
 $FZ \leftarrow 1$ If $(AC) = 0$ else $FZ \leftarrow 0$
 $FN \leftarrow 1$ If $(AC) < 0$ else $FN \leftarrow 0$


Description: Subtract the contents of FSRC from the accumulator in Single or Double Precision. Result is in accumulator unless Floating Underflow occurs and the interrupt is not enabled, in this case AC is set to 0.

2.97 μ s
2.97 μ s

NEGF
NEGD

Negate Floating/Double

1707FDST



Operation: $FDST \leftarrow \neg(FDST)$

Condition Codes: $FC \leftarrow 0$
 $FV \leftarrow 0$
 $FZ \leftarrow 1$ If $(FDST) = 0$ else $FZ \leftarrow 0$
 $FN \leftarrow 1$ If $(FDST) < 0$ else $FN \leftarrow 0$


Description: Negate Floating or Double Precision number, store result in same location. (FDST)

7.17 μ s
10.97 μ s

MULF MULD

Multiply Floating/Double.

171ACFSRC



Operation: $AC \leftarrow (AC) * (FSRC)$ If $[(AC) * (FSRC)] > XLL$ or $FIU=1$,
else $AC \leftarrow 0$

Condition Codes: $FC \leftarrow 0$
 $FV \leftarrow 1$ If $AC > XUL$ else $FV \leftarrow 0$
 $FZ \leftarrow 1$ If $(AC) = 0$ else $FZ \leftarrow 0$
 $FN \leftarrow 1$ If $(AC) < 0$ else $FN \leftarrow 0$


Description: Multiply the contents of the selected accumulator by the contents of FSRC. Store result in accumulator unless Floating Underflow occurs without the interrupt enabled in this case AC is set to 0.

8.17 μs
11.87 μs

DIVF
DIVD

Divide Floating/Double

174(AC + 4)FSRC




- Operation:** If (FSRC) = 0
AC \leftarrow (AC)/(FSRC) If [(AC)/(FSRC)] > XLL or FIU=1,
else AC \leftarrow 0
If (FSRC) = 0 registers, including AC, untouched
- Condition Codes:** FC \leftarrow 0
FV \leftarrow 1 If (AC) > XUL else FV \leftarrow 0
FZ \leftarrow 1 If (AC) = 0 else FZ \leftarrow 0
FN \leftarrow 1 If (AC) < 0 else FN \leftarrow 0
- Description:** If the contents of FSRC are not equal to zero, divide the contents of the accumulator by (FSRC) and store the result in the accumulator unless Floating Underflow occurs and the interrupt is not enabled in this case the AC is set to 0. If attempt is made to divide by zero, accumulator is left unchanged and FEC Register is set to 4.

4.17 μ s
4.37 μ s

CMPF CMPD

Compare Floating/Double

173(AC + 4)FSRC



Operation: (FSRC) - (AC)


Condition Codes: FC \leftarrow 0
FV \leftarrow 0
FZ \leftarrow 1 if (FSRC) - (AC) = 0 else FZ \leftarrow 0
FN \leftarrow 1 if (FSRC) - (AC) < 0 else FN \leftarrow 0

Description: Compare the contents of FSRC with the accumulator. Set the appropriate floating point condition codes: FSRC and the accumulator are left unchanged.

7.87 μ S
15.27 μ S

MODF
MODD

Multiply and Integerize Floating/Double 171(AC + 4)FSRC



Operation: $AC \vee 1 \leftarrow \text{Int}[(AC) * (FSRC)]$ If $[(AC) * (FSRC)] > XLL$
or $FIU = 1$, else $AC \vee 1 \leftarrow 0$
 $AC \leftarrow (AC) * (FSRC) - (AC \vee 1)$ If $[(AC) * (FSRC)] > XLL$
or $FIU = 1$, else $AC \leftarrow 0$

Condition Codes: $FC \leftarrow 0$
 $FV \leftarrow 1$ If $(AC) > XUL$ else $FV \leftarrow 0$
 $FZ \leftarrow 1$ If $(AC) = 0$ else $FZ \leftarrow 0$
 $FN \leftarrow 1$ If $(AC) < 0$ else $FN \leftarrow 0$

Description: The product of (AC) and (FSRC) is produced to 48 bits in Floating Mode and 59 bits in Double Mode. The integer part $\text{Int}[(AC) * (FSRC)]$ of the product is then found and stored in $AC \vee 1$. $AC \vee 1$ is the FPP Accumulator OR'd with 1. The fractional part is then obtained and stored in AC. Thus if even-numbered Accumulators (0 or 2) are used this instruction uses two accumulators (0 and 1; 2 and 3); whereas if odd-numbered accumulators are used only one Accumulator is used (1:3) and all that is left is the fractional part of the operation. If underflow occurs and the interrupt is not enabled, AC and $AC \vee 1$ are loaded with zero.

NOTE: Multiplication by 10 can be done with zero error allowing decimal digits to be "stripped off" with no loss in precision.


3.27 μ s

3.27 μ s

LDCDF LDCFD

Load and convert from Double to Floating
or from Floating to Double

177(AC + 4)FSRC



Operation: $AC \leftarrow C_x(FSRC)$ If $[(FSRC)] > XLL$ or $FIU = 1$, else $AC \leftarrow 0$ Where C_x specifies conversion from floating mode x to floating mode y , and $x = F$ and $y = D$ If $FD = 0$, or $x = D$ and $y = F$ If $FD = 1$.

Condition Codes: $FC \leftarrow 0$
 $FV \leftarrow 1$ If $(AC) > XUL$ else $FV \leftarrow 0$
 $FZ \leftarrow 1$ If $(AC) = 0$ else $FZ \leftarrow 0$
 $FN \leftarrow 1$ If $(AC) < 0$ else $FN \leftarrow 0$


Description: If the current mode is Floating Mode ($FD = 0$) the source is assumed to be a double-precision number and is converted to single precision. If the Floating Truncate bit is set the number is truncated, otherwise the number is rounded. If the current mode is Double Mode ($FD = 1$) the source is assumed to be a single-precision number and is loaded left justified in the AC. The lower half of the AC is cleared.

2.97 μ s
3.57 μ s

STCFD STCDF

Store and convert from Floating to Double
or from Double to Floating

176ACFDST




- Operation:** $FDST \leftarrow C_{xy}(AC)$ where C_{xy} specifies conversion from floating mode x to floating mode y and $x = F$ and $y = D$ if $FD = 0$, or $x = D$ and $y = F$ if $FD = 1$
- Condition Codes:** $FC \leftarrow 0$
 $FV \leftarrow 1$ if $(AC) > XUL$ else $FV \leftarrow 0$
 $FZ \leftarrow 1$ if $(AC) = 0$ else $FZ \leftarrow 0$
 $FN \leftarrow 1$ if $(AC) < 0$ else $FN \leftarrow 0$
- Description:** If the current mode is Floating, the Accumulator is stored left justified in FDST and the lower half is cleared; otherwise in Double Precision, the contents of the accumulator are converted to single precision, truncated or rounded depending on the state of FT and stored in FDST.

5.37 μ s
 5.57 μ s
 5.97 μ s
 6.27 μ s

LDCIF
LDCID
LDCLF
LDCLD

Load and Convert Integer or Long Integer to
 Floating or Double Precision

177ACSRC



Operation: $AC \leftarrow C_{ix}(SRC)$ where C_{ix} specifies conversion from integer mode j to floating mode x and $j = I$ if $FL = 0$ or L if $FL = 1$ and $x = F$ if $FD = 0$, or D if $FD = 1$.

Condition Codes:
 $FC \leftarrow 0$
 $FV \leftarrow 0$
 $FZ \leftarrow 1$ If $(AC) = 0$ else $FZ \leftarrow 0$
 $FN \leftarrow 1$ If $(AC) < 0$ else $FN \leftarrow 0$


Description: Conversion is performed on the contents of SRC from a 2's complement Integer with precision j to a floating point number of precision x . Note that j and x are determined by the state of the mode bits FL and FD: i.e. $J = I$ or L , and $X = F$ or D .

When a 32 bit Integer is specified (L mode) and (SRC) has an addressing mode of 0, or immediate addressing mode is specified, the 16 bits of the source register are left justified and the remaining 16 bits loaded with zeros before conversion. In the case of LDCLF the fraction is truncated and only the highest 24 significant bits are used.

5.17 μ S
 5.97 μ S
 5.97 μ S
 5.97 μ S

STCFI
STCFL
STCDI
STCDL

Store and Convert from Floating or Double to
 Integer or Long Integer 175(AC + 4)DST



Operation: $DST \leftarrow C_{xi}(AC)$ If $-JL-1 < C_{xi}(AC) < JL$, else $DST \leftarrow 0$
 where C_{xi} specifies conversion from floating mode x to integer mode j and $j = I$ if $FL = 0$ or L if $FL = 1$ and $x = F$ if $FD = 0$, or D if $FD = 1$

Condition Codes: $C \leftarrow FC \leftarrow 0$ If $-JL-1 < C_{xi}(AC) < JL$ else $FC \leftarrow 1$
 $V \leftarrow FV \leftarrow 0$
 $Z \leftarrow FZ \leftarrow 1$ If $(DST) = 0$ else $FZ \leftarrow 0$
 $N \leftarrow FN \leftarrow 1$ If $(DST) < 0$ else $FN \leftarrow 0$


Description: Conversion is performed from a floating point representation of the data in the accumulator to an integer representation. When the conversion is to a 32 bit word (L mode) and an address mode of 0, or immediate addressing mode, is specified, only the most significant 16 bits are stored in the destination register. If the operation is out of the integer range selected by FL, FC is set to 1 and the contents of the DST are set to 0. Numbers to be converted are always truncated (rather than rounded) before conversion. This is true even when the truncate mode bit is cleared in the Floating Point Status Register.

2.77 μ s

LDEXP

Load Exponent

176(AC + 4)SRC



Operation: AC SIGN \leftarrow (AC SIGN)
AC EXP \leftarrow (SRC) + 200

Condition Codes: FC \leftarrow 0
FV \leftarrow 0
FZ \leftarrow 1 If (AC) = 0 else FZ \leftarrow 0
FN \leftarrow 1 If (AC) < 0 else FN \leftarrow 0


Description: Load Exponent Word from SCR into Accumulator. Convert (SRC) from 2's complement to excess 200₈ notation.

3.67 μ s

STEXP

Store Exponent

175ACDST




Operation: $DST \leftarrow AC \text{ EXPONENT} - 200$

Condition Codes: $C \leftarrow FC \leftarrow 0$
 $V \leftarrow FV \leftarrow 0$
 $Z \leftarrow FZ \leftarrow 1$ if $(DST) = 0$ else $FZ \leftarrow 0$
 $N \leftarrow FN \leftarrow 1$ if $(DST) < 0$ else $FN \leftarrow 0$

Description: Store accumulator's exponent in DST, convert it from excess 200_8 notation to 2's complement.

2.37 μ s
2.57 μ s

CLRF CLRD



Operation: FDST \leftarrow 0

Condition Codes: FC \leftarrow 0
FV \leftarrow 0
FZ \leftarrow 1
FN \leftarrow 0


Description: Set FDST to 0. Set FZ condition code.

2.97 μ s
2.97 μ s

ABSF
ABSD

Make Absolute Floating/Double

1706FDST



Operation: $FDST \leftarrow -(FDST)$ If $(FDST) < 0$ else $FDST \leftarrow (FDST)$

Condition Codes: $FC \leftarrow 0$
 $FV \leftarrow 0$
 $FZ \leftarrow 1$ If $(FDST) = 0$ else $FZ \leftarrow 0$
 $FN \leftarrow 0$


Description: Set the contents of FDST to its absolute value.

2.77 μ s
2.77 μ s

TSTF TSTD

Test Floating/Double

1705FDST



Operation: FDST \leftarrow (FDST)

Condition Codes: FC \leftarrow 0
FV \leftarrow 0
FZ \leftarrow 1 IF (FDST) = 0 else FZ \leftarrow 0
FN \leftarrow 1 IF (FDST) < 4 else FN \leftarrow 0


Description: Set the Floating Point Processor's Condition Codes according to the contents of FDST.

2.37 μ s

SETF

Set Floating Mode

170001



Operation: FD \leftarrow 0


Description: Set the FPP in Single Precision Mode

2.37 μ s

SETD

Set Floating Double Mode

170011



Operation: FD \leftarrow 1


Description: Set the FPP in Double Precision Mode

2.37 μ s

SETI

Set Integer Mode

170002



Operation: FL \leftarrow 0


Description: Set the FPP for Integer Data

2.37 μ s

SETL

Set Long Integer Mode

170012




Operation: FL \leftarrow 1

Description: Set the FPP for Long Integer Data

LDFPS

Load FPPs Program Status

1701SRC



Operation: FPS ← (SRC)


Description: Load FPP's Status from SRC.

2.15 μs

STFPS

Store FPPs Program Status

1702DST



Operation: DST ← (FPS)


Description: Store FPP's Status in DST

2.57 μ s

STST

Store FPPs Status

1703DST



Operation: DST \leftarrow (FEC)
DST + 2 \leftarrow (FEA)

Description: Store the FEC and then the FPP's Exception Address Pointer in DST and DST + 2


Note: If destination mode specifies a general register or immediate addressing, only the FEC is saved.

2.37 μ s

CFCC

Copy Floating Condition Codes

170000



Operation: C \leftarrow FC
V \leftarrow FV
Z \leftarrow FZ
N \leftarrow FN

Description: Copy FPP Condition Codes into the CPU's Condition Codes.




Figure 8-1 System Operator's Console

THE SYSTEM OPERATOR'S CONSOLE

The PDP-11/45 System Operator's Console is designed for convenient system control. A complete set of function switches and display indicators provide comprehensive status monitoring and control facilities.

The System Operator's Console is illustrated in Figure 8.1.

8.1 CONSOLE ELEMENTS

The PDP-11/45 System Operator's Console provides the following facilities:

- 1) A System Key Switch (OFF/ON/LOCK)
- 2) A bank of 7 indicator lights, indicating the following Central Processor states: RUN, PAUSE, MASTER(UNIBUS), USER, SUPERVISOR, KERNEL, DATA.
- 3) An 18-bit Address Register display
- 4) An Addressing Error indicator light (ADRS ERR)
- 5) A 16-bit Data Register display
- 6) An 18-bit Switch Register
- 7) Control knobs
 - a) Address Display Select
 1. USER I VIRTUAL
 2. USER D VIRTUAL
 3. SUPERVISOR I VIRTUAL
 4. SUPERVISOR D VIRTUAL
 5. KERNEL I VIRTUAL
 6. KERNEL D VIRTUAL
 7. PROGRAM PHYSICAL
 8. CONSOLE PHYSICAL
 - b) Data Display Select
 1. DATA PATHS
 2. BUS REGISTER
 3. FPP μ ADRS.CPU μ ADRS.
 4. DISPLAY REGISTER

8) Control Switches

- a) LOAD ADRS (Load Address)
- b) EXAM (Examine)
- c) REG EXAM (Register Examine)
- d) CONT (Continue)
- e) ENABLE/HALT
- f) S-INST/S-BUS CYCLE (Single Instruction/Single Bus Cycle)
- g) START
- h) DEPOSIT
- i) REG DEPOSIT (Register Deposit)

8.2 SYSTEM POWER SWITCH

The System Power Switch controls Central Processor power as follows:

OFF	Power off for CPU. Solid-State Memory still receives power in order to insure data retention.
POWER	Power ON for CPU—normal use all console controls operable.
PANEL LOCK	Power ON for CPU. All console controls not operable except switch register.

Note: Since the theory of operation of high speed solid state memory involves the retention of a capacitive charge, it is essential that power be continually supplied in order to insure full data retention during those periods when the CPU power is OFF. When this facility is not required, Memory Power may be discontinued by flipping the Master Power switch in the rear of the CPU mounting cabinet to OFF.

8.3 CENTRAL PROCESSOR STATE INDICATORS

This bank of indicator lights shows the current major system state as follows:

RUN	The CPU is executing program instructions. If the instruction being executed is a WAIT instruction, the RUN light will be on. The CPU will proceed from the WAIT on receipt of an external interrupt, or on console intervention.
PAUSE	The CPU is inactive because: 1) The current instruction execution has been completed as far as possible without more data from the UNIBUS and the CPU is waiting to regain con-

trol of the UNIBUS (UNIBUS master-ship) (see MASTER state.)

OR

2) The CPU has been HALTed from the System Operator's Console.

MASTER

The CPU is in control of the UNIBUS (UNIBUS Master). The CPU relinquishes control of the UNIBUS during DMA and NPR data transfers.

USER

The CPU is executing program instructions in USER mode. When the Memory Management Unit is enabled all address references are in USER Virtual Address Space

SUPERVISOR

The CPU is executing program instructions in SUPERVISOR mode. When the Memory Management Unit is enabled, all address references are in SUPERVISOR Virtual Addressing space.

KERNEL

The CPU is executing program instructions in KERNEL mode. When the Memory Management Unit is enabled, all address references are in KERNEL Virtual Addressing space.

DATA

If on, the last memory reference was to D address space in the current CPU mode. If a 0, the last memory reference was to I address space in the current CPU mode.

8.4 ADDRESS DISPLAY REGISTER

The Address Display Register is primarily a software development and maintenance aid. The contents of this 18-bit indicator are controlled by the Address Select knob as follows:

VIRTUAL

The Address Display Register indicates the current address reference as a 16-bit Virtual Address when the Memory Management Unit is enabled, otherwise it indicates the true 16-bit Physical Address. Bits 17 and 16 will be off unless the Memory Management Unit is disabled AND the current address references some UNIBUS device register in the uppermost 4K of basic address space (i.e. 28K-32K).

PROGRAM PHYSICAL

The Address Display Register indicates the current address reference as a true 18-bit Physical Address.

CONSOLE PHYSICAL

The Address Display Register indicates the current address reference as a 16-bit Virtual Address when the Memory Management Unit is enabled otherwise it indicates the true 16-bit Physical Address.

Bits 17 and 16 indicate the contents of corresponding bits of the Switch Register as of the last LOAD ADRS console operation.

8.5 ADDRESSING ERROR DISPLAY

This 1-bit display indicates the occurrence of any addressing errors. The following address references are invalid:

1. Non-existent memory
2. Access Control violations
3. Unassigned memory pages

(See chapter 6: Memory Management)

8.6 DATA DISPLAY REGISTER

The Data Display Register is primarily a hardware maintenance facility. The contents of this 16-bit indicator are controlled by the Data Display Select knob as follows:

DATA PATHS	The Data Display Register indicates the current output of the PDP-11/45 Arithmetic/Logical Unit subsystem.
BUS REGISTER	The Data Display Register indicates the current output of the PDP-11/45 CPU (UNIBUS I, II and the EXPRESS BUS).
FPP μ ADRS.CPU μ ADRS.	The Data Display Register indicates the current ROM address, FPP control micro-program (bits 15-8), and the CPU control micro-program (bits 7-0).
DISPLAY	The Data Display Register indicates the current contents of the 16-bit write-only "Switch Register" located at Physical Address 777570. This register is generally used to display diagnostic information, although it can be used for any meaningful purpose.

8.7 SWITCH REGISTER

The functions of this 18-bit bank of switches are determined by:

- 1) Control Switches
- 2) Address Display Select knob

These functions will be described in the next section along with the appropriate control switch.

Note that the current setting of the Switch Register may be read under program control from a read-only register at Physical Address 777570.

8.8 CONTROL SWITCHES

8.8.1 LOAD ADRS (Load Address)

When the LOAD ADRS switch is depressed the contents of the Switch Register are loaded into the CPU Bus Address Register and displayed in the Address Display Register lights. If the Memory Management Unit is disabled the address displayed is the true Physical Address.

If the Memory Management Unit is enabled the interpretation of the address indicated by the Switch Register is determined by the Address Display Select knob.

Note that the LOAD ADRS function does not distinguish between PROGRAM PHYSICAL and CONSOLE PHYSICAL.

8.8.2 EXAM (Examine)

Depressing the EXAM switch causes the contents of the current location specified in the CPU Bus Address Register to be displayed in the DATA Display Register.

Depressing the EXAM switch again causes a EXAM-STEP operation to occur. The result is the same as the EXAM except that the contents of the CPU Bus Address Register are incremented by two before the current location has been selected for display. An EXAM-STEP will not cross a 32K memory block boundary.

An EXAM operation which causes an ADRS ERR (Addressing Error) must be corrected by performing a new LOAD ADRS operation with a valid address.

8.8.3 REG EXAM (Register Examine)

Depressing the REG EXAM switch causes the contents of the General Purpose Register specified by the low order five bits of the bus-address register to be displayed in the Data Display Register.

The Switch Register is interpreted as follows:

Contents	Register Displayed
0-5	General Registers 0-5 (set 0)
6	Kernel Mode Register 6
7	Program Counter
10 ₈ —15 ₈	General Register 0-5 (set 1)
16 ₈	Supervisor Mode Register 6
17 ₈	User Mode Register R6

8.8.4 CONT (Continue)

Depressing the CONT switch causes the CPU to resume executing in-

structions or bus cycles at the address specified in the Program Counter (Register 7-(PC)). The CONT switch has no effect when the CPU is in RUN state.

The function of the CONT switch is modified by the setting of the ENABLE/HALT and S/INST-S/BUS cycles switches as follows:

ENABLE (up)	CPU resumes normal operation under program control.
HALT (down)	A. S/INST (up)—CPU executes next instruction then stops.
	B. S/BUS cycle (down)—CPU executes next address reference then stops.

8.8.5 ENABLE/HALT

The ENABLE/HALT switch⁴ is a two-position switch with the following functions:

ENABLE (up)	The CPU is able to perform normal operations under program control.
HALT (down)	The CPU is stopped and is only operable by the console switches.

The setting of the ENABLE/HALT switch modifies the function of the CONTINUE (8.8.4) and START (8.8.7) switches.

8.8.6 S/INST—S/BUS CYCLE (Single Instruction/Single Bus Cycle)

The S/INST-S/BUS CYCLE switch effects only the operation of the CONTINUE switch as described in section 8.8.4. This switch has no effect on any switches when the ENABLE/HALT switch is set to ENABLE.

8.8.7 START

The functions of the START switch depend upon the setting of the ENABLE/HALT switch as follows:

ENABLE	Depressing the START switch causes the CPU to start executing program instructions at the address specified by the current contents of the CPU Bus Address Register. The START switch has no effect when the CPU is in RUN state.
HALT	Depressing the START switch causes a console reset to occur.

8.8.8 DEP (Deposit)

Raising the DEP switch causes the current contents of the Switch Register to be deposited into the address specified by the current contents of the CPU Bus Address Register.

Raising the DEP switch again causes a DEP-STEP operation to occur. The result is the same as the DEP except that the contents of the CPU Bus Address Register are incremented by two before the current location

has been selected for the deposit operation. A DEP-STEP will not cross a 32K memory block boundary.

A DEP operation which causes an ADRS ERR (addressing Error) is aborted and must be corrected by performing a new LOAD ADRS operation with a valid address.

8.8.9 REG DEPOSIT (Register Deposit)

Raising the REG DEP causes the contents of the Switch Register to be deposited into the General Purpose Register specified by the current contents of the CPU Bus Address Register.

The CPU Bus Address Register should have been previously loaded by a LOAD ADRS operation according to the Switch register settings described in REG EXAM (8.8.3).

NOTE: The EXAM and DEP switches are coupled to enable an EXAM-DEP-EXAM sequence to be carried out on a location, without having to do a LOAD ADRS. The following sequence is possible:

```
EXAM
DEP   ADDRESS A
EXAM
STEP EXAM
DEP   ADDRESS A + 1
EXAM
```

8.8.10 ADDRESS SELECT

The ADDRESS SELECT knob is used for two functions. It provides an interpretation for the ADDRESS DISPLAY REGISTER as explained in section 8.4. It also determines for EXAM, STEP EXAM, DEP and STEP DEP, what set of Page Address Registers, if any, will be used to relocate the address loaded by the LD ADRS function.

KERNEL I, KERNEL D, SUPER I, SUPER D, USER I and USER D positions cause the address loaded into the switch register to be relocated if the Memory Management Option is installed and operating. Which set of the 6 sets of Page Address Registers (PARs) is used is determined by the ADDRESS SELECT switch. EXAMs, STEP EXAMs, DEPs and STEP DEPs, under these conditions, are relocated to the physical address specified by the appropriate PAR. If the action attempted from the console is not allowed (for example—attempting to DEP into a READ ONLY page) the ADRS ERROR indicator will come on. A new LD ADRS must be done to clear this condition. Note that, in the general case, the physical location accessed is different from the virtual address loaded into the switch register. The ADDRESS DISPLAY REGISTER will always, in these 6 positions, show exactly what was loaded from the switch register. These positions make it convenient to examine and change programs which are subject to relocation, without requiring any knowledge of where they have actually been relocated in physical memory.

PROGRAM PHYSICAL. This position is provided to allow one, when "single cycling" through a program, to monitor the physical addresses being accessed by the program. It is most useful when the accesses are being relocated by the Memory Management Option. In this case the Address shown in Address Display Register is different than that shown in the other positions. This position should *not* be used to perform EXAM, STEP-EXAM, DEP or STEP DEP functions.

CONSOLE PHYSICAL—This position is provided to allow EXAM, STEP EXAM, DEP and STEP DEP Functions to physical memory locations whether or not the Memory Management option is installed or operating. In this position the ADDRESS DISPLAY register indicates the physical address loaded from the switch register.

APPENDIX A

INSTRUCTION SET PROCESSOR

A DESCRIPTION OF THE PDP-11 USING THE INSTRUCTION SET PROCESSOR (ISP) NOTATION¹

ISP is a language (or notation) which can be used to define the action of a computer's instruction set. It defines a computer, including console and peripherals, as seen by a programmer. It has two goals: to be precise enough to constitute the complete specification for a computer and to still be highly readable by a human user for purposes of reference, such as this manual. The main part of the manual contained an English language description of the PDP-11, using ISP expressions as support in defining each instruction. This appendix contains an ISP description of the PDP-11, using a few English language comments as support.

The following brief introduction to the notation is given using examples from the PDP-11 Model 20 ISP description. The complete PDP-11 description follows the introduction.

A processor is completely defined at the programming level by giving its instruction set and its interpreter in terms of basic operations, data types and the system's memory. For clarity the ISP description is usually given in a fixed order:

Declare the system's memory:

Processor state (the information necessary to restart the processor if stopped between instructions, e.g., general registers, PC, index registers)

Primary memory state (the memory directly addressable from the processor)

Console state (any external keys, switches, lights, etc., that affect the interpretation process)

Secondary memory (the disks, drums, decapes, magnetic tapes, etc.)

Transducer state (memory available in any peripheral devices that is assumed in the instructions of the processor)

Declare the instruction format

Define the operand address calculation process

Declare the data types

Declare the operations on the data types

Define the instruction interpretation process including interrupts, traps, etc.

Define the instruction set and the instruction execution process (provides an ISP expression for each instruction)

Thus, the computer system is described by first declaring memory, data-types and primitive data operations. The instruction interpreter and the instruction-set is then defined in terms of these entities.

The ISP notation is similar to that used in higher level programming languages. Its statements define entities by means of expressions involving other entities in the system. For example, an instruction to increment (add-one) to memory would be

Increment := (M[x] - M[x] + 1); *add one to memory, x*

This defines an operation, called "increment", that takes the contents of memory M at an address, x, and replaces it with a value one higher. The := symbol simply assigns a name (on the left) to stand for the expression (on the right). English language comments are given in italics. Table 1 gives a reference list of notations, which are illustrated below.

ISP expressions are inherently interpreted in parallel, reflecting the underlying parallel nature of hardware operations. This is an important difference between ISP and standard programming languages, which are inherently serial. For example, in

¹The notation derived and used in the book, Computer Structures: Readings and Examples, McGraw-Hill, 1971 by C. Gordon Bell and Allen Newell. The book contains ISP's of 14 computers.

$Z := (M[x] \leftarrow S'D'; M[y] \leftarrow M[x]);$

both righthand sides of the data transmission operator (\leftarrow) are evaluated in the current memory state in parallel and then transmission occurs. Thus the old value of $M[x]$ would go into $M[y]$. Serial ordering of processing is indicated by using the term "next". For example,

$Z := (M[x] \leftarrow S'D'; \text{next } M[y] \leftarrow M[x]);$

performs the righthand data transmission after the lefthand one. Thus, the new value of $M[x]$ would be used for $M[y]$ in this latter case.

Memory Declarations

Memory is defined by giving a memory declaration as shown in Table 1. For example,

$Mp[0:2^k - 1] < 15:0 >$

declares a memory named Mp , of 2^k words (where k has been given a value). The addresses of the words in memory are $0, 1, \dots, 2^k - 1$. Each word has 16 bits and the bits are labeled $15, 14, \dots, 0$. Some other examples of memory declarations are:

Boundary-error ₂ } Boundary-error ₂ } Activity ₃ N/Negative CC<3> M[0:2 ¹⁸ -1]<7:0> M[0:15][0:4095]<7:0> brop<1:0> ₂ ¹⁶ } brop<7:0> ₂ ¹⁶ }	<i>boolean memories; scalar bit alternatives</i> <i>ternary digit, holding value 0, 1, or 2</i> <i>alias, N and Negative are synonymous</i> <i>bit 3 of a register</i> <i>vector of 2¹⁸ 8-bit words</i> <i>array of 16 x 4096 8-bit words</i> <i>alternative ways of defining a register</i> <i>using base 16 and base 2</i>
--	--

Renaming and Restructuring of Previously Defined Registers

Registers can be defined in terms of existing registers. In effect, each time the name to the left of the $:=$ symbol is encountered, the value is computed according to the expression to the right of $:=$. A process can be evoked to form the value and side-effects are possible when the value is computed.

Examples of simple renaming in part or whole of existing memory

N/Negative := CC<3> *N is name of bit 3 of register CC*
SP<15:0> := R[6]<15:0> *SP is the same as register R[6]*

Examples of register formed by concatenation

LAC<L,0:11> := LGAC<0:11>
AB<0:47> := A<0:23>□B<0:23>
Mword[0]<15:0> := Mbyte[0]<7:0>□Mbyte[1]<7:0>

Examples of values and registers formed by evaluation of a process

ai/address-increment<1:0> := (*value of ai is 2 if - byte op,*
 ¬ byte-op ⇒ 2; *else value is 1*
 byte-op ⇒ 1)
Run := (Activity = 0) *Run=1 or 0 depending on value of Activity*
 being 0 or not 0

Instruction Format

Instruction formats are declared in the same fashion as memory and are not distinguishable as special non-memory entities. The instructions are carried in a register; thus it is natural to declare them by giving names to the various parts of the instruction register. Usually only a single declaration is made, the instruction/ i , followed by the declarations of the parts of the instruction; the operation code, the address fields, indirect bit, etc.

Example

This declaration would correspond to the usual box diagram:


Table 1. ISP Character-Set and Expression Forms

A,...,Z,a,...,z,-,.,',",0,...,9	name alphabet. This character set is used for names.
$M\{a:b \dots \underbrace{\{v:w\}}_n : x:y\}_z$	<p>comments. Italics are used for comments.</p> <p>memory declaration. An n-dimensional memory array of words where a:b ... v:w are the range of values for the first and last dimensions. The values of the first dimension are, for example, a, a+1, ..., b for $a \leq b$ (or a, a-1, ..., b for $a > b$). The word length base, z, is normally 2 if not specified. The digits of the word are x, x+1, ..., y.</p>
$a := f(\text{expression})$	definition. The operator, :=, defines memory, names, process, or operations in terms of existing memory and operations. Each occurrence of "a" causes the in place substitution by f(expression).
$b(c, \dots, e) := g(\text{expression})$	The definition b, may have dummy parameters, c, ..., e, which are used in g(expression).
name' := h(expression)	side effects naming convention. In this description we have used ' to indicate that a reference to this name will cause other registers to change.
$a \leftarrow f(\text{expression})$ $f(\text{expression}) \rightarrow a$	transmission operator. The contents in register a are replaced by the value of the function.
()	parentheses. Defines precedence and range of various operations and definitions (roughly equivalent to begin, and end).
{data-type}	operator and data-type modifier
boolean = expression;	conditional expression; equivalent to ALGOL <u>if</u> boolean <u>then</u> expression
boolean = (expression-1 <u>else</u> expression-2);	equivalent to Algol <u>if</u> boolean <u>then</u> expression-1 <u>else</u> expression-2
; next	sequential delimiter interpretation is to occur
□	concatenation. Consider the registers to the left and right of □ to be one.
;	statement delimiter. Separates statements.
,	item delimiter. Separates lists of variables.
a/b	division and synonym. Used in two contexts: for division and for defining the name, a, to be an alias (synonym) of the name, b.
?	unknown or unspecified value
‡	set value. Takes on all values for a digit of the given base, e.g., $1‡_2$ specifies either 10_2 or 11_2
X(:= boolean) = expression;	instruction value definition. The name X is defined to have the value of the boolean. When the boolean is true, the expression will be evaluated.

Table I. cont'd.

Common Arithmetic, Logical and Relational Operators

<u>Arithmetic</u>	<u>Logical</u>	<u>Relational</u>
+ add	¬ not	≡ identical
- subtract, also negative	∧ and	≠ not identical
x multiply	∨ or	= equal
/ divide	⊕ exclusive-or	≠ not equal
mod, modulo (remainder)	≡ equivalence	> greater than
() ² squared		≥ greater than or equal
() ^a exponentiation		< less than
() ^a _b exponentiation		≤ less than or equal
() _b base		
() _b base		
sqrt () square root		
abs () absolute value		
sign-extend ()		



<p>/instruction<15:0> bop<3:0> := i<15:12> sf<5:0> := i<11:6> df<5:0> := i<5:0></p>	<p><i>the instruction</i> <i>specifies binary (dyadic) operations</i> <i>specifies source (first) operand</i> <i>specifies second operand and destination</i></p>
--	--

Operand Address Calculation Process

In all processors, instructions make use of operands. In most conventional processors, the operand is usually in memory or in the processor, defined as M[z], where z is the effective address. In PDP-11, a destination address, Address, is used in this fashion for only two instructions. It is defined in ISP by giving the process that calculates it. This process may involve only accesses to primary memory (possibly indexed), but it may also involve side effects, i.e., the modification of either of primary memory or processor memory (e.g., by incrementing a register). Note that the effective address is calculated whenever its name is encountered in evaluating an ISP expression (either in an instruction or in the interpretation expression). That is, it is evaluated on demand. Consequently, any side effects may be executed more than once.

Operation Determination Processes

Instead of effective-address, the operands are usually determined directly. For example, the 16-bit destination register is just the register selected by the dr field of an instruction, i.e.,

Rd := R[dr] *the destination register*

In one other case, the operand is just the next word following an instruction. This next word can be defined,

nw<15:0>/next-word := (Mw[PC]; PC ← PC + 2) *the next word is selected and PC is moved*

Here, the ' shows that a reference to nw will cause side effects, in this case, PC ← PC + 2. For calculating the source operand, S, the process is:

<p>S'<15:0> := ((sm=0) ⇒ R[st]; (sm=1) ⇒ Mw[R[sr]] (sm=2) ∧ (sr=7) ⇒ nw;</p>	<p><i>value for source operand</i> <i>if mode=0 then S' is the Register addressed</i> <i>by instruction field sr</i> <i>if mode=1 the S' is indirect via R sr</i> <i>if mode=2 and source register=PC then the</i> <i>next word is the operand; this can be</i> <i>seen by substituting the expression for nw'</i></p>
--	--

An expression is also needed for the operand, S, which does not cause the side effects, and assuming the effects have taken place, counteracts them. Thus, S would be:

```

S<15:0> := (
    (sm=0) => R[sr];           no side effects
    (sm=1) => Mw[R[sr]];      no side effects
    (sm=2) ^ (sr=7) => Mw[PC-2] counteract previous side effects
    :
    :
    :

```

In the ISP description a general process is given which determines operands for Source-Destination, word-byte, and with-without side-effects. In order to clarify what really happens, the source operand calculation, for words, with side effects, is given below.

```

Sf<5:0> := i<11:6>           source field (6-bits) of instruction
sm8 := sf<5:3>              source mode control field
sd := sf<3>                 deferred address control
sr8 := sf<2:0>             register specification for source

nw'<15:0> := (Mw[PC]; PC - PC+2) next word; used as operand
Rs<15:0> := R[sr]          source register specification

S'<15:0>/Source := ((      value for the source--direct addressing
    (sm=0) = Rs;           use the register Rs as operand
    (sm=2) ^ (sr≠7) => (Mw[Rs] direct auto-increment (increment
        Rs - Rs + 2);      Rs); usually used as POP
    (sm=2) ^ (sr=7) => nw;  direct; actually immediate operand
    (sm=4) => (Rs - Rs - 2; next direct; auto-decrement (decrement
        Mw[Rs]);          Rs); usually used as PUSH
    (sm=6) ^ (sr≠7) => Mw[nw' + Rs]; direct; indexed via Rs--uses next-word
    (sm=6) ^ (sr=7) => Mw[nw' + PC]; direct; relative to PC; uses next-word
                                        value for the source-defined addressing
    (sm=1) = Mw[Rs];          defer through Rs
    (sm=3) ^ (sr≠7) => (Mw[Mw[Rs]]); defer through stack; auto
        Rs - Rs + 2);        increment
    (sm=3) ^ (sr=7) => M[nw']; defer via next word; absolute addressing
    (sm=5) => (Rs - Rs - 2; next defer through stack after auto
        Mw[Mw[Rs]]);        decrement
    (sm=7) ^ (sr≠7) => Mw[Mw[nw' + Rs]]; defer, indexed via Rs
    (sm=7) ^ (sr=7) => Mw[Mw[nw' + PC]]; defer relative to PC
    );                       end calculation process;
    (sr=6) ^ ((sm=4) ∨ (sm=5)) ^ checks if stack overflowed for several
    (SP<4008) => (Stack overflow - 1) modes
    )                       end source calculation

```

Data-Types

A data-type specifies the encoding of a meaning into an information medium. The meaning of the data-type (what it designates or refers to) is called its referent (or value). The referent may be anything ranging from highly abstract (the uninterpreted bit) to highly concrete (the payroll account for a specific type of employee).

Every data-type has a carrier, into which all its component data-types can be mapped. The carrier is used in storing the data-type in memories and is usually a word or multiple thereof. It must be extensive enough to hold all the component data-types, but may be a larger (having error checking and correcting bits, or

even unused bits). The mapping of the component data-types into the carrier is called the format. It is given as a list which associates to each component an expression involving the carrier (e.g., as in the instruction format).

ISP provides a way of naming data-types, which also serves as a basis for abbreviations. Some data-types simply have conventional names (e.g., character/ch, floating point numbers/f); others are named by their value (e.g., integer/i). Data-types which are iterates of a basic component can be named by the component suffixed by a length-type. The length-type can be array/a, implying a multi-dimensional array of fixed, but unspecified dimensions; a string/st, implying a single sequence, of variable length (on each occurrence); or a vector/v, implying a one dimensional array of a fixed but unspecified number of components. The length-type need not exist, and then this form of the name is not applicable. Thus, iv is the abbreviation for an integer vector. It is also possible to name a data-type by simply listing its components.

Data-types are often of a given precision and it has become customary to measure this in terms of the number of components that are used, e.g., triple precision integers. In ISP this is indicated by prefixing the precision symbol to the basic data-type name, e.g., di for double precision integer. Note that a double precision integer, while taking two words, is not the same thing as a two integer vector, so that the precision and the length-type, though both implying something about the size of the carrier, do not express the same thing.

A list of common data-types and their abbreviations is given in Table 2.

Operations on Data-types

Operations produce results of specific data-types from operands of specific data-types. The data-types themselves determine by and large the possible operations that apply to them. No attempt will be made to define the various operations here, as they are all familiar. A reasonably comprehensive list is given in Table 1. An operation-modifier, enclosed in braces, { }, can be used to distinguish variant operations. The operation-modifier is usually the name of a data-type, e.g., A+B{f} is a floating point addition. Modifiers can also be a description name applying to the operation, e.g., a X2 {rotate}.

New operations can be defined by means of forms. For example, the various add operations on differing data-types are specified by writing {data-type} after the operation.

Instruction Interpretation Process

The instruction interpretation expression and the instruction set constitute a single ISP expression that defines the processor's action. In effect, this single expression is evaluated and all the other parts of the ISP description of a processor are evoked as indirect consequences of this evaluation. Simple interpreter without interrupt facilities show the familiar cycle of fetch-the-instruction and execute-the instruction.

Example:

```
Run = (instruction ← M[PC]; PC ← PC + 1; next
      Instruction-execution; next)
```

This is a simple interpreter, not the one for the FDP-11

In more complex processors the conditions for trapping and interrupting must also be described. The effective address calculation may also be carried out in the interpreter, prior to executing the instruction, especially if it is to be calculated only once and will have a fixed value independent of anything that happens while executing instructions. Console activity can also be described in the interpreter, e.g., the effect of a switch that permits stepping through the program under manual control, or interrogating and changing memory.

The normal statement for PDP-11 interpretation is just:

```
← Interrupt-rq ∧ Run = (instruction ← Mw[PC]; PC ← PC + 2; next
  Instruction-execution; next
  T-flag = (State-change(148); T-flag ← 0)
```

*fetch
execute
trace mode*

Table 2. Common Data-Types Abbreviations

<u>Primitive</u>	<u>String and Vector</u>
b bit or boolean	bv bit.vector
by byte	by.st byte.string
ch character	ch.st character.string
cx complex	
df double precision floating	
dw double word	
d digit	jd j-digit number
f floating	
fr fraction	
hw half word	
i integer	
mx mixed number	
qw quadruple length word	
tw triple length word	
w word	

Instruction-Set and Instruction Execution Process

The instruction set and the process by which each instruction is executed are usually given together in a single definition; this process is called instruction-execution in most ISP descriptions. This usually includes the definition of the conditions for execution, i.e., the operation code, value, the name of the instruction, a mnemonic alias, and the process for its execution. Thus, an individual instruction typically has the form:

```

MOV (:= bop = 00012) = (           move word
  r ← S'; next                    move source to intermediate register
  N ← r<15>;                       negative?
  (r<15:0> = 0) = (Z ← 1 else Z ← 0); zero?
  V ← 0;                            overflow cleared
  D ← r);                            transmit result to destination

```

With this format for the instruction, the entire instruction set is simply a list of all the instructions. On any particular execution, as evoked by the interpretation expression, typically one and only one operation code correlation will be satisfied, hence one and only one instruction will be executed.

In the case of PDP-11, the text carries the definition of the individual instructions, hence they are not redefined in the appendix. Instead, the appendix defines the condition for executing the instructions. For example,

```
MOV := (bop = 00012)
```

is given in the appendix, and the action of MOV is defined (in ISP) in the text.

THE PDP-11 ISF

PDP-11's Primary (Program) Memory and Processor State

The declaration of this memory includes all the state (bits, words, etc.) that a program (programmer) has access to in this part of the computer. The console is not included. The various secondary memories (e.g., disks, tapes) and input-output device state declarations are included in a following section.

Primary (program) Memory

Mp[0:2^k-1]<15:0> actual physical, 16-bit memory of a particular system; k = 12, ..., 17

Mw/Mword[x<15:0>]<15:0> := (word-accessed memory
 → x<0> = Mp[x<15:1>]; word on even byte boundary, all right
 x<0> = (?value ; Boundary-error - 1)) word on odd byte boundary, trap

Mb/Mbyte[x<15:0>]<7:0> := (byte-accessed memory
 → x<0> = Mp[x<15:1>]<7:0>; take low-order bits if even
 x<0> = Mp[x<15:1>]<15:8>) take hi-order bits if odd

Processor State

R[0:7]<15:0> eight, 16-bit General-Registers, used for accumulators, indexing and stacks

SP<15:0>/Stack-Pointer := R[6] special stack, controlled by R[6]
 PC<15:0>/Program-Counter := R[7] location next instruction, also R[?]

PS<15:0>/Processor-State-Word 16-bit register giving rest of state
 Unused<7:0>/Undefined := PS<15:8> mapping of bits into PS
 P<2:0>/Priority := PS<7:5> interrupt level control of processor
 T/Trace := PS<4> denotes whether trap is to occur after each instruction

CC<3:0>/Condition-Codes := PS<3:0> set as a function of instruction and results
 N/Negative := CC<3> if result = -
 Z/Zero := CC<2> if result = 0
 V/Overflow := CC<1> if result overflows
 C/Carry := CC<0> if result carried into/borrowed from most significant bit

Processor-Controlled Error Flags (resulting from instruction-execution)

Boundary-Error set if word is accessed on odd byte boundary
 Stack-Overflow set if word accessed, via SP < 400₈
 Time-Out-Error set if non-existent memory or device is referenced
 Illegal-Instruction set if a particular class of instructions is executed

Processor-activity

Activity₃ ternary, specifying state of processor
 Run := (Activity = 0) normal instruction interpretation
 Wait := (Activity = 1) waiting for interrupt
 Off := (Activity = 2) off, dormant

Error-Flags (resulting from without the processor)

Power-Fail-Flag set if power is low
 Power-Up-Flag set when power comes on

Instruction format field declarations

1<15:0>/instruction

bop<3:0> := 1<15:12>
sf<5:0> := 1<11:6>
 sm_g := sf<5:3>
 sd := sf<3>
 sr_g := sf<2:0>
df<5:0> := 1<5:0>
 dm_g := df<5:3>
 dd := df<3>
 dr_g := df<2:0>

binary opcode format
source field
source mode - 3 bits
source defer bit
source register - 3 bits
destination field
destination mode - 3 bits
destination defer bit
destination register - 3 bits

uop<3:0>_g := 1<15:6>
df

unary op code (arith., logical, shifts)
see binary op format

jsop<7:0> := 1<15:9>
 sr; df

jsr format
see binary op format

brop<1:0>₁₆ := 1<15:8>
offset<7:0> := sign-extend(1<7:0>)

branch format
offset value

trop<1:0>₁₆ := 1<15:8>
unused-trop<1:0>₁₆ := 1<7:0>


trap format

eop<6:0> := 1<15:9>
er<3:0> := 1<8:6>
esf<5:0> := 1<5:0>
 esm_g := esf<5:3>
 esd := esf<3>
 esr_g := esf<2:0>


extended opcode format
extended register
extended source field
mode
defer
register

fop<7:0> := 1<15:8>
fr<7:0> := 1<7:6>
fsf<5:0> := 1<5:0>


floating op format
register destination
source




binary operand (2 operands) format



unary operand (1 operand), JMP format




JSR format




branch format
value := sign-extend (offset)



trap format



extended operation format



floating op format

si/address-increment<1:0> := (

 ¬ Byte-op = 2;

 Byte-op = 1)

Byte-op := (MOVB ∨ BICB ∨ BISB ∨ BITB ∨ CLRB ∨

 COMB ∨ INCB ∨ DECB ∨ NEGB ∨ ADCB ∨

 SBCB ∨ TSTB ∨ ROVB ∨ ROLB ∨ ASRR ∨

 ASLB ∨ SWAB)

Reserved-instruction := ((i =) ∨ (i =) ∨ ... ∨ (i =)) unused instructions

Registers and Data Addressed via Instruction Format Specifications

nw'/next-word<15:0> := Mw[PC]

used in operand determination

nw'/next-word<15:0> := (Mw[PC]; PC ← PC + 2)

with side effects

lw'/last-word<15:0> := Mw[PC - 2]

undoes side effects

Rs<15:0> := R[sr]<15:0>

the source register

Rd<15:0> := R[dr]<15:0>

the destination register

Operand Determination for Source and Destination

Two types of operands are used: S', D', Sb' and Db' - for operands that cause side-effects (i.e., other registers are changed; and S, D, Sb and Db for operands that do not cause side effects. Two general procedures Wo' and Wo are used to determine these operands for side effects and no side effects, respectively

S'<15:0> := Oprd'<15:0>(Mw, 2, sm, sr)

source word operand side-effects

S<15:0> := Oprd'<15:0>(Mw, 2, sm, sv)

source word operands no side-effects

Sb'<7:0> := Oprd'<7:0>(Mb, 2, sm, sr)

source byte

Sb<7:0> := Oprd'<7:0>(Mb, 1, sm, sr)

D'<15:0> := Oprd'<15:0>(Mw, 2, dm, dr)

Destination operands

D<15:0> := Oprd'<15:0>(Mw, 2, dm, dr)

Db'<7:0> := Oprd'<7:0>(Mb, 1, dm, dr)

Db<7:0> := Oprd'<7:0>(Mb, 1, dm, dr)

General Operand Calculation Process (with Side Effects)

Oprd'<wl:0>(M, si, m, rg) := ((

value for word or byte operand; direct addressing: wl indicates length; m mode, and rg register

Rr<15:0> := R[rg]

secondary definition for register

(m=0) ⇒ Rr<wl:0>;

0, use the register, Rr, as operand

(m=2) ∧ (rg≠7) ⇒ (M[Rr]; next

2, direct auto-increment (increment

 Rr ← Rr + si);

Rr); usually used in pop stack

(m=2) ∧ (rg=7) ⇒ nw'<wl:0>;

2, direct; next-word is immediate operand

(m=4) ⇒ (Rr ← Rr - si; next

4, direct; after auto decrement

 M[Rr]);

usually used as PUSH stack

(m=6) ∧ (rg≠7) ⇒ M[nw' + Rr];

6, direct; indexed via Rr uses next-word

(m=6) ∧ (rg=7) ⇒ M[nw' + PC];

6, direct; relative to PC; uses next-word value for word operand defer addressing

(m=1) ⇒ M[Rr];

1, defer through Rr

(m=3) ∧ (rg≠7) ⇒ (M[Mw[Rr]]; next

3, defer through Mw[Rr] (usually stack),

 Rr ← Rr + 2);

auto-increment

(m=3) ∧ (rg=7) ⇒ M[nw'];

3, defer via next-word; absolute addressing

(m=5) ⇒ (Rr ← Rr - si; next

5, defer through stack after auto

 M[Mw[Rr]]);

decrement

```

(m=7)  $\wedge$  (rgd7) = M[Mw[nw' + Rr]];
(m=7)  $\wedge$  (rg=7) = M[Mw[nw' + PC]];
);
(m=6)  $\wedge$  ((m=4)  $\vee$  (m=5))  $\wedge$ 
(SP < 4008) = (Stack-overflow - 1)
)

```

?, defer indexed via Rr
 ?, defer relative to PC
 end calculation process
 check if stack overflows

end operand calculation process

General Operand Calculation Process (without Side Effects)

Oprd<w1:0> (M,ai,m,rg) := (

```

Rr<15:0> := R[rg]
(m=0)  $\Rightarrow$  Rr<w1:0>;
(m=2)  $\wedge$  (rgd7) = Mw[Rr - ai];
(m=2)  $\wedge$  (rg=7) = lw<w1:0>;
(m=4) = M[Rr];
(m=6)  $\wedge$  (rgd7) = M[lw + Rr];
(m=6)  $\wedge$  (rg=7) = M[lw + PC];

(m=1) = M[Rr];
(m=3)  $\wedge$  (rgd7) = M[Mw[Rr - 2]];
(m=3)  $\wedge$  (rg=7) = M[lw];
(m=5) = M[Mw[Rr]];
(m=7)  $\wedge$  (rgd7) = M[Mw[lw + Rr]];
(m=7)  $\wedge$  (rg=7) = M[Mw[lw + PC]]

```

undo previous side-effects
 undo previous side-effects
 undo previous side-effects
 undo previous side-effects
 undo previous side-effects
 undo previous side-effects
 undo previous side-effects
 undo previous side-effects

Destination addresses for JMP and JSR

Da<15:0> := ((

```

(dm=0)  $\Rightarrow$  (?; Illegal-instruction - 1);
(dm=2)  $\wedge$  (drd7) = (Rd; Rd - Rd + 2);
(dm=2)  $\wedge$  (dr=7) = (PC; PC - PC + 2);
(dm=4)  $\Rightarrow$  (Rd - Rd - 2; next Rd);
(dm=6)  $\wedge$  (drd7) = (nw' + Rd);
(dm=6)  $\wedge$  (dr=7) = (nw' + PC);

```

directs:

illegal register address
 auto-increment
 null
 auto-decrement
 indexed
 relative

```

(dm=1) = Mw[Rd];
(dm=3)  $\wedge$  (drd7) = (Mw[Rd]; Rd - Rd + 2);
(dm=3)  $\wedge$  (dr=7) = nw';
(dm=5) = (Rd - Rd - 2; next Mw[Rd]);
(dm=7)  $\wedge$  (drd7) = Mw[nw' + Rd];
(dm=7)  $\wedge$  (dr=7) = Mw[nw' + PC]; next

```

defers:

via register
 via auto-increment
 absolute address
 auto-decrement
 via index
 relative to PC

```

(dr=6)  $\wedge$   $\neg$  ((dm=0)  $\vee$  (dm=3)  $\vee$  (dm=7))  $\wedge$  (SP < 4008) = (check for stack overflow
stack-overflow - 1)

```

Data Type Formats

by/byte<7:0>

w/word<15:0>

wi/word.integer<15:0>

bybv/byte.boolean-vector<7:0>

wbv/word.boolean-vector<15:0>

d/d.w/double.word<31:0>

f/d.f/double.word.floating<31:0>
fs/floating.sign := *f*<31>
fe/floating.exponent<7:0> := *f*<30:23>
fm/floating.mantissa<22:0> := *f*<22:0>
t/triple.word<47:0>
q/quadruple.word<63:0>
qf/quadruple.word.floating-point<63:0>
qfs := *qf*<63>
qfe := *qf*<62:55>
qfm := *qf*<54:0>

I/O Devices and Interrupts, State Information

Device[0:N-1]

Device-name[*J*]<15:0> := *J*

Device-interrupt-location[*J*]<15:0> := *K*

dob/device-output-buffer[*J*]<15:0>

dib/device-input-buffer[*J*]<15:0>

ds/device-status[*J*]<15:0>

derr/device-error-flags[*J*]<3:0> := *ds*[*J*]<15:12>

dbusy/device-busy[*J*] := *ds*[*J*]<11>

dunit/device-unit-selection[*J*]<2:0> := *ds*[*J*]<10:8>

ddone[*J*] := *ds*[*J*]<7>

denb/device-done-interrupt-enable := *ds*[*J*]<6>

derrenb/device-error-interrupt-enable := *ds*[*J*]<5>

dme/device-memory-extension[*J*]<4:3> := *ds*[*J*]<4:3>

dfnc/device-function[*J*]<2:0> := *ds*[*J*]<2:0>

dintrq/device-interrupt-request[*J*] := (

(*ddone*[*J*] ∧ *denb*[*J*] ∨ ((*derr*[*J*] ≠ 0) ∧ *derrenb*[*J*]))

dil/device-interrupt-level[*J*]<7:4>

N I/O devices - assume device *J*

number to which device responds and is addressed

each device has a value, *K*, which it uses as an address to interrupt processor program controlled device data

a register with device control state

common

status

assignments

each device is assigned to 1 of 4 levels

Mapping of Devices into M. Each device's registers are mapped into primary word memory, e.g., Teletype

M'[177560]₈ := *tks/ds*[TTY-keyboard]

keyboard status

M'[177562]₈ := *tkb/dib*[TTY-keyboard]

keyboard input data

M'[177564]₈ := *tps/ds*[TTY-printer]

teletypewriter status

M'[177566]₈ := *tpb/dob*[TTY-printer]

teletypewriter data to print

Interrupt Requests

br[bus-request-for-interrupt<7:4>] := (

(*dintrr*[0] = *dil*[0]) ∨

(*dintrq*[1] = *dil*[1]) ∨ ...

(*dintrq*[*J*] = *dil*[*J*]) ∨ ...

(*dintrq*[*N*] = *dil*[*N*]))

OR of all device requests

Interrupt-rq := (*intrql* ≥ *p*)

interrupt if a request is ≥ priority/*P*

intrql/interrupt-request-level<2:0> := (

br<7> = 7;

¬ *br*<7> ∧ *br*<6> = 6;

¬ *br*<7> ∧ ¬ *br*<6> ∧ ¬ *br*<5> ∧ *br*<4> = 4)

Instruction Interpretation Process

```
Interrupt-rq  $\wedge$  Run  $\Rightarrow$  (Normal-interpretation);
Normal-interpretation := (I  $\leftarrow$  Mw[PC]; PC  $\leftarrow$  PC + 2; next)  fetch
Instruction-execution; next                                  execute
T-flag  $\Rightarrow$  (State-change(14g); T-flag = 0)              trace
Interrupt-rq  $\wedge$  Off  $\Rightarrow$  (
  State-change(Device-interrupt-location[J]);                assume device J interrupts
  P  $\leftarrow$  intrql);
off = ( );
 $\neg$  Interrupt-rq  $\wedge$  Wait  $\Rightarrow$  ( );
State-change(x) := (                                         for stacking state and restore
  SP  $\leftarrow$  SP - 2; next
  Mw[SP]  $\leftarrow$  PS;
  SP  $\leftarrow$  SP - 2; next
  Mw[SP]  $\leftarrow$  PC;
  PC  $\leftarrow$  Mw[x];
  PS  $\leftarrow$  Mw[x+2])
Boundary-Error  $\Rightarrow$  (State-change(4g); Boundary-error = 0)
Time-Out-Error  $\leftarrow$  (State-change(4g); Time-Out-Error = 0)
Power-Fail-Flag  $\Rightarrow$  (state-change(24g); Power-Fail-Flag = 0); program must turn off computer
Power-Up-Flag  $\Rightarrow$  (PC  $\leftarrow$  24g; Power-Up-Flag = 0; Activity = 0) Start Up on power-up
```

Instruction-Set Definition

Each instruction is defined in ISP in the text, therefore, it will not be repeated here.

ISP for Floating Point Processor/FPP

Device-interrupt-location [FPP] := M'[244_g]

FEC<15:0>

FOCE := (FEC=2)

FDZE := (FEC=4)

FICE := (FEC=6)

FVE := (FEC=8)

FUE := (FEC=10)

FUVE := (FEC=12)

floating point processor error code register

floating op code error

floating divide by zero

floating integer conversion error

floating overflow

floating underflow

floating undefined variable

FAC[0:5]<63:0>

F<63:0>

FPC<15:0>

FPSR<15:0>

FER := FPSR<15>

FIE := FPSR<14>

FIUV := FPSR<11>

FIU := FPSR<10>

FIV := FPSR<9>

FIC := FPSR<8>

FD := FPSR<7>

FL := FPSR<6>

FT := FPSR<5>

FMM := FPSR<4>

6 floating point accumulators

temporary floating point register

floating point PC

floating point processor status register

floating error

interrupt enable

interrupt on undefined variable

interrupt on underflow

interrupt on overflow

interrupt on integer conversion error

floating double precision mode

floating long integer mode

floating truncate mode

floating maintenance mode

FN := FPSR<3> floating negative
 FZ := FPSR<2> floating zero
 FV := FPSR<1> floating overflow
 FC := FPSR <0> floating carry

Instruction format

OC<3:0> := i<15:12> op code
 FOC<3:0> := i<11:8> floating op code
 AC<1:0> := i<7:6> accumulator

General Definitions

XL := ((FD=0) = $1-2^{-24}$;
 (FD=1) = $1-2^{-56}$) largest fraction
 XLL := 2^{-128} smallest non-zero number
 XUL := $2^{127} * XL$ largest number
 JL := ((FL=0) = $2^{15}-1$;
 (FL=1) = $2^{31}-1$) largest integer

Address Calculation

FPS<63:0> := { floating point processor source
 (dm=0) = FAC(dr);
 (dmp=0) = (
 (FD=0) = D<15:0>[Mw[PC+2]];
 (FD=1) = D<15:0>[Mw[PC+2]]
 Mw[PC+4][Mw[PC+6]])
 FPS'<63:0> := (floating point processor source with
 (dm=0) = FAC(dr); side effects
 (dmp=0) = (
 (FD=0) = D'<15:0>[Cnw'
 (FD=1) = D'<15:0>[Cnw' Cnw' Cnw'])
 FPD<63:0> := FPS<63:0> floating point processor destination
 FPD'<63:0> := FPS'<63:0> floating point processor destination with
 side effects
 FS<15:0> := D<15:0> floating source, CPU mode
 FS'<15:0> := D'<15:0> floating source with side effects,
 CPU mode
 FD<15:0> := D<15:0> floating destination, CPU mode
 FD'<15:0> := D'<15:0> floating destination with side effects,
 CPU mode
 Fac := FAC(AC) destination floating register

¹a 17 bit result, r, used only for descriptive purposes

²A prime is used in S (e.g., S') and D (e.g., D') to indicate that when a word is accessed in this fashion, side effects may occur. That is, registers of R may be changed.

³if all 16 bits of result, r = 0, then Z is set to 1 else Z is set to 0.

⁴The 8 least significant bits are used to form a 16-bit positive or negative number by extending bit 7 into 15:8.

⁵a = b means: if boolean a is true then b is executed.

⁶Mw means the memory taken as a work-organized memory.

APPENDIX B

INSTRUCTION TIMING

B.1 INSTRUCTION EXECUTION TIME

The execution time for an instruction depends on the instruction itself, the modes of addressing used, and the type of memory being referenced. In the most general case, the Instruction Execution Time is the sum of a Source Address Time, and an Execute, Fetch Time.

$$\text{Instr Time} = \text{SRC Time} + \text{DST Time} + \text{EF Time}$$

Some of the instructions require only some of these times, and are so noted. Times are typical; processor timing, with core memory, may vary +15% to -10%.

B.1.1 BASIC INSTRUCTION SET TIMING

Double Operand

all instructions,

$$\text{except MOV: Instr Time} = \text{SRC Time} + \text{DST Time} + \text{EF Time}$$

$$\text{MOV Instruction: Instr Time} = \text{SRC Time} + \text{EF Time}$$

Single Operand

$$\text{all instructions: Instr Time} = \text{DST Time} + \text{EF Time} \quad \text{or} \\ \text{Instr Time} = \text{SRC Time} + \text{EF Time}$$

Branch, Jump, Control, Trap & Misc

$$\text{all instructions: Instr Time} = \text{EF Time}$$

B.1.2 USING THE CHART TIMES

To compute a particular instruction time, first find the instruction "EF" Time. Select the proper EF Time for the SRC and DST modes. Observe all "NOTES" to the EF Time by adding the correct amount to basic EF number.

Next, note whether the particular instruction requires the inclusion of SRC and DST Times, if so, add the appropriate amounts to correct EF number.

B.1.3 NOTES

1. The times specified generally apply to Word instructions. In most cases Even Byte instructions have the same times, with some Odd Byte instructions taking longer. All exceptions are noted.
2. Timing is given without regard for NRP or BR servicing. Memory types MM11-S, MF11-L, and ML11 are assumed with memory within the CPU mounting assembly.
3. If the Memory Management (KT11-C) option is installed and operating, instruction execution times increase by .09 μ sec for each memory cycle used.
4. When MM11-S, MM11-L, MF11-L or ML11 are used, due to overlap of central processor operation with memory cycles, there is no advantage as far as processor instruction timing is concerned by interleaving memory.
5. All times are in microseconds.

B.1.4 SOURCE ADDRESS TIME

Instruction	Source Mode	SRC Time			Memory Cycles
		Bipolar	MOS	Core	
Double Operand	0	0.00	.00	0.00	0
	1	.30	.45	.83	1
	2	.30	.45	.83	1
	3	.75	1.05	1.81	2
	4	.45	.60	.98	1
	5	.90	1.20	1.96	2
	6	.60	.90	1.73	2
	7	1.05	1.50	2.71	3

B.1.5 DESTINATION ADDRESS TIME

Instruction	Mode Destination	DST Time (A)			Memory Cycles
		Bipolar	MOS	Core	
Single Operand and Double Operand (except MOV, MTP, JMP, JSR)	0	.00	.00	.00	0
	1	.30	.45	.83(B)	1
	2	.30	.45	.83(B)	1
	3	.75	1.05	1.81(B)	2
	4	.45	.60	.98	1
	5	.90	1.20	1.96	2
	6	.60	.90	1.73(B)	2
	7	1.05	1.50	2.71(B)	3

NOTE (A): Add .15 μ sec for odd byte instructions, except DST Mode 0.

NOTE (B): Add .07 μ sec if SRC Mode = 1-7.

B.1.6 EXECUTE, FETCH TIME Double Operand

Instruction (Use with SRC Time and DST Time)	SRC Mode 0 DST Mode 0			Mem Cyc	SRC Mode 1-7 DST Mode 0			Mem Cyc	SRC Mode 0 to 7 DST Mode 1 to 7			Mem Cyc
	EF Time				ET Time				EF Time			
	Bipolar	MOS	Core		Bipolar	MOS	Core		Bipolar	MOS	Core	
ADD, SUB, BIC, BIS	.30 (D)	.45 (D)	.90 (C)	1	.45 (D)	.60 (D)	1.05 (E)	2	.75	1.05	2.00	2
CMP, BIT	.30 (D)	.45 (D)	.90 (C)	1	.45 (D)	.60 (D)	1.05 (E)	1	.45	.60	1.13	1
XOR	.30 (D)	.45 (D)	.90 (C)	1	—	—	—		.75	1.05	2.00	2

NOTE (C): Add .23 μ sec if DST is R7.

NOTE (D): Add .3 μ sec if DST is R7.

NOTE (E): Add .23 μ sec if DST is R7, add .08 μ sec if DST is odd byte and not R7.

Double Operand (Cont.)

Instruction (Use with SRC Time)	DST Mode	DST Register	EF Time (SRC MODE = 0)			EF Time (SRC MODE = 1-7)			Memory Cycles
			Bipolar	MOS	Core	Bipolar	MOS	Core	
MOV	0	0-6	.30	.45	.9	.45	.60	1.05	1
	0	7	.60	.75	1.13	.75	.90	1.28	1
	1	0-7	.75	1.05	2.00	.75	1.05	1.95	2
	2	0-7	.75	1.05	2.00	.75	1.05	1.95	2
	3	0-7	1.20	1.65	2.98	1.20	1.65	3.05	3
	4	0-7	.90	1.20	2.15	.90	1.20	2.03	2
	5	0-7	1.25	1.80	3.13	1.25	1.80	3.13	3
	6	0-7	1.05	1.50	2.90	1.20	1.65	3.05	3
7	0-7	1.50	2.10	3.88	1.65	2.25	3.96	4	

Single Operand

Instruction (Use with DST Time)	DST MODE = 0			Memory Cycles	DST MODE 1 TO 7			Memory Cycles
	EF Time				EF Time			
	Bipolar	MOS	Core		Bipolar	MOS	Core	
CLR COM, INC, DEC, ADC, ABC, ROL, ASL, SWAB, SXT	.30 (J)	.45 (J)	.90 (G)	1	.75	1.05	2.0	2
NEG	.75	.90	1.28	1	1.05	1.40	2.18 (F)	2
TST	.30 (J)	.45 (J)	.90 (G)	1	.45	.60	1.13	1
ROR, ASR	.30 (J)	.45 (J)	.90 (G)	1	.75	1.05	2.0 (H)	2
ASH, ASHC	.75 (I)	.9 (I)	1.28 (I)	1	.90 (I)	1.05 (I)	1.43 (I)	1

NOTE (F): Add .12 μ sec if odd byte.

NOTE (G): Add .23 μ sec if DST is R7.

NOTE (H): Add .15 μ sec if odd byte.

NOTE (I): Add .15 μ sec per shift.

NOTE (J): Add .30 μ sec if DST is R7.

Single Operand (Cont.)

Instruction (Use with SRC Times)	Bipolar	MOS	Core	Memory Cycles
MUL	3.30	3.45	3.83	1
DIV				
by zero	.90	1.05	1.43	1
shortest	6.90	7.05	7.83	1
longest	8.55	8.70	9.08	1

Instruction	Bipolar	MOS	Core	Memory Cycles	
MFPI	1.05	1.35	2.18	2	(use with
MFPD	1.05	1.35	2.18	2	SRC Times)

Instruction	DST Mode	Instruction Time			Memory Cycles
		Bipolar	MOS	Core	
MTPI MTPD	0	.90	1.20	2.03	2
	1	1.20	1.65	2.93	3
	2	1.20	1.65	2.93	3
	3	1.65	2.25	4.03	4
	4	1.35	1.80	3.01	3
	5	1.80	2.40	4.11	4
	6	1.65	2.25	4.03	4
	7	2.10	2.85	5.01	5

Branch Instructions

Instruction	Instr Time (Branch)			Instr Time (No Branch)			Memory Cycles
	Bipolar	MOS	Core	Bipolar	MOS	Core	
BR, BNE, BEQ, BPL, BMI, BVC, BVS, BCC, BCS, BGE, BLT, BGT, BLE, BHI, BLOS, BHS, BLO	.60	.90	1.13	.30	.45	.90	1
SOB	.75	.90	1.13	.60	.75	1.28	1

Jump Instructions

Instruction	Destination Mode	Instr Time			Memory Cycles
		Bipolar	MOS	Core	
JMP	1	.90	1.15	1.43	1
	2	.90	1.15	1.43	1
	3	1.20	1.50	2.26	2
	4	.90	1.15	1.43	1
	5	1.35	1.65	2.41	2
	6	1.05	1.35	2.18	2
	7	1.50	1.95	3.16	3
JSR	1	1.50	1.80	2.63	2
	2	1.50	1.80	2.63	2
	3	1.80	2.25	3.46	3
	4	1.50	1.80	2.63	2
	5	1.95	2.35	3.61	3
	6	1.50	1.95	3.38	3
	7	2.10	2.70	4.36	4

Control, Trap, & Miscellaneous Instructions

Instruction	Instr Time			Memory Cycles	
	Bipolar	MOS	Core		
RTS	1.05	1.45	2.11	2	
MARK	.90	1.20	2.03	2	
RTI, RTT	1.50	1.95	3.16	3	
SET N, Z, V, C					
CLR N, Z, V, C	.60	.75	1.13	1	
HALT	1.05	1.05	1.05	0	
WAIT	.45	.45	.45	0	WAIT Loop for a BR is .3 μ sec.
RESET	10ms	10ms	10ms	1	
IOT, EMT, TRAP, BPT	2.40	3.15	5.26	5	
SPL	.60	.75	1.13	1	
INTERRUPT	2.25	2.85	4.95	4	First Device

B.2 LATENCY

Interrupts (BR requests) are acknowledged at the end of the current instruction. For a typical instruction execution time of 3 μ sec, the average time to request acknowledgement would be one-half this or 1.5 μ sec. The worst case (longest) instruction time (Negative Divide with SRC Mode 7) and, hence, the longest request acknowledgement would be 11.79 μ sec max with core (10.2 μ sec with MOS and 9.00 μ sec with Bipolar).

The Interrupt service time, which is the time from BR request acknowledgement to the fetch of the first subroutine instruction, is 4.95 μ sec max with core, 2.85 μ sec with MOS and 2.25 μ sec with Bipolar.

Hence, the total worst case time from BR request to begin the fetch of the first service routine instruction is:

	Bipolar	MOS	Core
Normal	11.25	12.87	16.74
Memory Management Operating	11.70	13.32	17.90

The total average time for BR request to begin the fetch of the first service routine instruction is:

	Bipolar	MOS	Core
Normal	3.95	4.85	8.45
Memory Management Operating	4.40	5.30	8.90

NPR Latency is 3.5 μ sec worst case.

B.3 FLOATING POINT INSTRUCTION TIMING

Floating point times are calculated in a similar manner to the CPU Instruction times. The times involved are preexecution Interaction time, source or destination time, execution time, CPU displacement time, and the time taken to fetch the next Instruction.

With the floating point Instructions the CPU and the FPP operate in parallel and hence, the Instruction time includes a CPU time and a parallel FPP time. These times do not coincide, each unit is free to continue at a different time with its next operation.

Instruction Time (CPU) = pre-interaction + source + disengage + fetch of the next instruction.

Instruction Time (FPP) = pre-interaction + source + execution.

Pre-execution Interaction time: This involves the passing of information between the CPU and the FPP. The total time is 600 ns. The floating point unit Interacts only during the last 150 ns.

Therefore, the CPU becomes active at the time 0, and remains so until time 600 ns, while the FPP becomes active at time 450 and remains active until the time 600 ns. The FPP could have been active from a previous task during the initial 450 ns.

B.3.1 Source or Destination Times

These times are the same whether the calculation is for source or destination. The times given are for the address calculation. To this must be added memory access time to actually fetch the operands:
Integer—1 word; Long integer and Floating—2 words;
Double Precision—4 words

Therefore: SOURCE/DST = Calculation Time + MEMORY ACCESS TIME.

Calculation Time

	300 ns Bipolar (MS11-C)	450 ns MOS (MS11-B)	850 ns Core (MM11-S)
Reg mode 0	1120	1120	1120
Floating mode 0	1120	1120	1120
1	1260	1260	1260
2	1260	1260	1260
3	1650	1760	2100
4	1260	1260	1260
5	1800	1920	2250
6	1650	1760	2100
7	2080	2300	3100

MEMORY ACCESS TIME: to be added to all modes except 0.

ADD: 1600 ns for core OR

ADD: 1100 ns for MOS OR

ADD: 990 ns for BIPOLAR

for every memory Access required to fetch the data.

B.3.2 Floating Point Execution

Floating Point Instructions			
Mnemonic	Instruction	Time (μ s)	
Floating AC—Floating Source Group: OPR FSRC, AC			
		MIN	MAX
LDF	Load floating	1.5	1.5
LDD	Load floating double	1.7	1.7
ADDF	Add floating	2.4	5.5
ADDD	Add floating double	2.6	7.9
SUBF	Subtract floating	2.4	5.5
SUBD	Subtract floating double	2.6	7.9
MULF	Multiply floating	4.7	7.1
MULD	Multiply floating double	6.6	12.8
DIVF	Divide floating	5.4	8.4
DIVD	Divide Floating double	7.5	13.8
MODF	Multiply and Integerize floating	5.3	7.9
MODD	Multiply and Integerize floating double	7.8	20.2
LDCDF	Load and convert from double to floating	1.7	2.4
LDCFD	Load and convert from floating to double	1.7	2.4
STF	Store floating	.88	.88
STD	Store floating double	.88	.88
Floating AC—Floating Destination Group: OPR AC, FDST			
CMPF	Compare floating	2.6	3.2
CMPD	Compare double	2.8	3.5
STCFD	Store and convert from floating to double	1.7	
STCDF	Store and convert from double to floating	1.7	3.0
Floating AC—Source Group: OPR SRC, AC			
LDCIF	Load and convert from integer to floating	3.6	4.6
LDCID	Load and convert from integer to double	3.8	4.8
LDCLF	Load and convert from long integer to floating	3.8	5.7
LDCLD	Load and convert from long integer to double	4.1	5.9
LDEXP	Load Exponent	1.5	

Mnemonic	Instruction	Time (μ s)	
		MIN	MAX
Floating AC—Destination Group: OPR AC, DST			
STCFI	Store and convert from floating to integer	3.4	4.4
STCFL	Store and convert from floating to long integer	4.2	5.2
STCDI	Store and convert from double to integer	4.2	5.2
STCDL	Store and convert from double to long integer	4.2	5.2
STEXP	Store exponent	2.4	
Floating Destination Group: OPR FDST			
CLRF	Clear floating	1.1	
CLRD	Clear double	1.3	
NEGF	Negate floating	1.7	
NEGD	Negate double	1.7	
ABSF	Make absolute floating	1.7	
ABSD	Make absolute double	1.7	
TSTF	Test floating	1.5	
TSTD	Test double	1.5	
Operate Group: OPR SRC			
LDFS	Load floating program status OPR DST	.88	
STFPS	Store floating program status	.88	
STST	Store floating status (exception & code and program counter)	1.3	
	Copy Condition codes	1.1	
SETF	Set floating mode	1.1	
SETI	Set integer mode	1.1	
SETD	Set double mode	1.1	
SETL	Set long integer mode	1.1	

B.3.3 Disengage and Fetch Next Instruction

This is the time required by the CPU to disengage from the FPP and fetch the next instruction. If the instruction is a Floating Point instruction then the CPU restarts the cycle described in this section (B.2), otherwise, it returns to the CPU cycle described in section B.1.

CPU time only:

990 ns Core

600 ns MOS

450 ns Bipolar

Example

ADDF A(R), AC0 :a floating point add instruction in indexed source mode (6), and in single (32 bit precision).

Pre-Interaction Time:

CPU 600 μ s FPP 150 μ s

Source:	300 μ s Bipolar (MS11-C)	450 μ s MOS (MS11-B)	850 μ s Core (MM11-S)
Calculation Time:	1650	1760	2100
Memory Access (2 words)	<u>1930</u>	<u>2200</u>	<u>3200</u>
	3630	3960	5300
Execution Average	3950	3950	3950
Disengage Fetch Next	450	600	990


CPU Instruction Time: Starting at T = 0

Bipolar	$600 + 3630 + 450 = 4880$
MOS	$600 + 3960 + 600 = 5160$
Core	$600 + 530 + 990 = 6890$

FPP Instruction Time: Starting at T = 450

Bipolar	$150 + 3630 + 3950 = 7730$
MOS	$150 + 3960 + 3950 = 8060$
Core	$150 + 5300 + 3950 = 9400$

MEMORY MAP AND RESERVED LOCATIONS



The location of trap vectors for processor conditions is as follows:

Note: all locations are located in Kernel Virtual Address Space.

Vector Condition

- 4 Odd address error—an attempt has been made to reference a word at an odd address. The instruction causing the error has been aborted (at the microinstruction causing the “bus pause” associated with the bad bus cycle).
- 4 Fatal stack violation (red), warning stack violation (yellow)—an attempt has been made to modify (DATIP, DATO, OR DATOB) a word whose address is below the final stack boundary (red), or 16 words before the boundary (yellow) with an R6-related operation.* The instruction causing the error is aborted (as above), R6 is set equal to the quantity 4 and a trap is taken. Odd address errors and timeout associated with R6-related operations will also lead to fatal stack violations (instead of their normal sequence).

* The final stack boundary is 16 words beyond the stack limit.

- 4 Timeout—no SSYN was seen within 5 μ sec of processor transmission of MSYN. The instruction causing the error is aborted.
- 4 Parity error—incorrect data parity has been detected by a slave device during a processor initiated DATA or DATIP. Note that this will be acted upon only at the next bus cycle initiated by the processor.
- 10 Illegal and reserved instructions—JMP R;JSR M,R; “HALT” in user mode; FPP instructions when FPP not available; and the reserved instructions are:

000007-000077
 000210-000227
 007000-007777
 075000-076777
 106400-107777

NOTE: If FPP is available, illegal FPP instructions trap to Location 244 with an exception code of 2.

- 14 Opcode 000003 and the T bit trap through this vector. The T bit causes this trap whenever it is set and there is no RTT instruction in the instruction register.
- 20 IOT Trap
- 24 Power Fail
- 30 Emulator Trap (EMT)
- 34 Trap Instruction (TRAP)

Interrupt Vectors

- 240 Programming Interrupt Request
- 244 Floating Point Exception
- 250 Memory Management violations and memory management traps

Register Locations

Memory Management: (All Memory Management Registers use 2 Word Locations)

- 777572 through 777577 Memory Management Status Registers 0-2
- 772516 through 772517 Memory Management Status Register 3
- 777600 through 777617 User Instruction Descriptor Registers 0-7
- 777620 through 777637 User Data Descriptor Registers 0-7
- 777640 through 777657 User Instruction Address Registers 0-7
- 777660 through 777677 User Data Address Registers 0-7
- 772200 through 772217 Supervisor Instruction Descriptor Registers 0-7
- 772220 through 772237 Supervisor Data Descriptor Registers 0-7
- 772240 through 772257 Supervisor Instruction Active Registers 0-7
- 772260 through 772277 Supervisor Data Address Registers 0-7
- 772300 through 772317 Kernel Instruction Descriptor Registers 0-7
- 772320 through 772337 Kernel Data Descriptor Registers 0-7
- 772340 through 772357 Kernel Instruction Address Registers 0-7
- 772360 through 772377 Kernel Data Address Registers 0-7

Memory Parity Status Registers:

772110 Memory Parity Status Register 0-8K
772112 Memory Parity Status Register 8K-16K
772114 Memory Parity Status Register 16K-24K
772116 Memory Parity Status Register 24K-32K
772120 Memory Parity Status Register 32K-40K
772122 Memory Parity Status Register 40K-48K
772124 Memory Parity Status Register 48K-56K
772126 Memory Parity Status Register 56K-64K
772130 Memory Parity Status Register 64K-72K
772132 Memory Parity Status Register 72K-80K
772134 Memory Parity Status Register 80K-88K
772136 Memory Parity Status Register 88K-96K
772140 Memory Parity Status Register 96K-104K
772142 Memory Parity Status Register 104K-112K
772144 Memory Parity Status Register 112K-120K
772146 Memory Parity Status Register 120K-124K

Processor:

777570 Console Switch and Display Register
777772 Program Interrupt Register
777774 Stack Limit Register
777776 Processor Status Word

ORDER OF SERVICE

In the case of concurrent trap/interrupt conditions. The PDP-11/45 service requests in the following order:

Order	Condition	Action
0	Odd Address	Trap (4)
0	Fatal Stack Violation (Red)	SP4; Trap (4)
0	Page Violation	Trap (250)
0	Timeout (NXM)	Trap (4)
0	Parity Error	Trap (4)
1	FPP Data XFER Request	DO Bus Cycle
2	Console Flag	Console Control
3	Memory Management Trap	Trap (250)
4	Warning Stack Violation (Yel)	Trap (4)
5	Power/Fail	Trap (24)
	Processor Priority Level 7	
6	FPP Exception Trap	Trap (224)
7	PIRQ 7	Trap (240)
8	BR7,INTR	Interrupt

Processor Priority Level 6

- | | | |
|----|----------|------------|
| 9 | PIRQ 6 | Trap (240) |
| 10 | BR6,INTR | Interrupt |

Processor Priority Level 5

- .
- .
- .
- .
- ***Processor Priority Level 1***
- | | | |
|----|--------|------------|
| 17 | PIRQ 1 | Trap (240) |
|----|--------|------------|

*** Processor Priority Level 0***

- | | | |
|----|-----------------------|-----------|
| 18 | T-Bit Set and not RTT | Trap (14) |
|----|-----------------------|-----------|

NOTES

1. Order 0 conditions are mutually exclusive and take immediate precedence over all other conditions. They force immediate interruption (abort) of instruction execution when they occur.
2. Other order conditions permit the instruction in progress (including traps) to proceed to completion before being serviced.
3. Lower order conditions depend on the Processor Status setup by previously serviced conditions. Thus power/fail service should (and generally would) raise the processor priority to level 7 to lock out interrupts.
4. The T-bit condition is not locked out by any processor priority level.
5. Fatal Stack Violations supersede warning stack violations. In case both occur during an instruction only the fatal violation is acknowledged.
6. If a Memory Management Trap and a Fatal Stack Violation occur together, the PDP-11/45 will loop:
 - A. The Fatal Stack Violation sequence is exceeded.
 - B. The first instruction of the "Trap 4" Service Routine is executed setting the SP to 0.
 - C. The Memory Management Trap begins but causes a Fatal Stack Violation.
 - D. Step A, Step B, Step C,...
7. If power fails before execution of a Fatal Stack Violation Trap Sequence:
 - A. The PS and PC are taken from the power fail sequence.
 - B. The SP is set to 4.

C. The previous PC and PS (from the routine causing the stack violation) are saved and (pushed into) locations 0 and 2.

D. Execution of the Power/Fail Service Routine begins—a suggested beginning is:

```
TST SP
BNE SPOK
MOV #N,SP ;set up emergency stack
```

8. If power failed during (or after) execution of a Fatal Stack Violation Trap Sequence:
 - A. The Stack Violation Trap Sequence completes—saving the PC and PS associated with the violations 0 and 2 and establishing a new PS and PC as determined by the vector at 4.
 - B. The first instruction of the “Trap 4” service routine is executed.
 - C. The Power/Fail Vector is loaded into the PS and PC.
 - D. An attempt is made to push the old PC and PS into the stack. The stack pointer however is probably at 0 so a Fatal Stack Violation occurs.
 - E. Operation proceeds as in note #7. The PC and PS of the routine causing the Fatal Stack Violation are overwritten (they were saved in locations 0 and 2). The PC and PS for the “Trap 4” Service Routine are the same as the contents of the Power/Fail Vector—also the same as the contents of the current PS and PC.
9. The vectors for Power/Fail and Stack Violation Traps (at 24 and 4) must raise the processor Priority to level 7 to lock out interrupts. In addition they must not set the T-bit. Since there is no stack left, any interrupts or T-Bit Traps will cause Fatal Stack Violation and either:
 - A. Loop if the vector at 4 is mis-set.
 - B. Loss of the Power/Fail Trap if the vector at 24 is mis-set.
10. If an interrupt or trap causes a Fatal Stack Violation, the PS and PC loaded from the Interrupt Vector are saved in locations 0 and 2. The PC and PS of the Interrupted (Main Line) Program are lost.
11. If any order 0 conditions occur in the Fatal Stack Violation Trap Sequence before a new stack is established. The processor will enter the loop described in Note 6 above.

12. If a particular error condition is endangered by the routine to service that condition the stack will grow to its limit and then cause a Fatal Stack Violation.
13. Reserved and illegal instructions (including the Floating Point instructions when Floating Point Hardware is not implemented as well as attempts to halt in User Mode) are treated just as instruction traps (e.g., trap, EMT). Note however that in the case of Floating Point Instructions the PC pushed on the stack points at the instruction word.

PROGRAM INTERRUPT REQUESTS

Program Interrupt Requests

A request is booked by setting one of the bits 15 through 9 (for PIR 7—PIR 1) in the Program Interrupt Register at location 777772. The hardware sets bits 7—5 and 3—1 to the encoded value of the highest PIR bit set. This Program Interrupt Active (PIA) should be used to set the Processor Level and also index through a table of interrupt vectors for the seven software priority levels. Figure D-1 shows the layout of the PIR Register.




Figure D-1 Program Interrupt Request Register

When the PIR is granted, the Processor will Trap to location 240 and pick up PC in 240 and the PSW in 242. It is the interrupt service routine's responsibility to queue requests within a priority level and to clear the PIR bit before the interrupt is dismissed.

The actual interrupt dispatch program should look like:

```

MOVB PIR,PS           ; places Bits 5—7 in PSW Priority Level
                       Bits
MOV  R5,—(SP)         ; save R5 on the stack
MOV  PIR,R5
BIC  #177761,R5       ; Gets Bits 1—3
JMP  @DISPAT(R5)      ; use to index through table

```


which requires 15 core locations.

APPENDIX E

MEMORY PARITY

When memory parity is installed on a system, the user can keep track of the status of his memory through memory status registers.

There are 16 memory status registers on the PDP-11/45 each one associated with an 8K section of memory. (See Appendix C) The status register uses the format:



Bit 15 Parity Error

Bits 11, 10 Type of Parity

Bits 11 and 10 are associated with the high order 4K and low order 4K of this memory address bank. When set to a 1 they specify odd parity for their respective half banks; when clear, even parity.

Bit 9 Halt Enable

When bit 9 is set the machine will execute a halt if a parity error occurs; when clear, the machine will perform an effective timeout and 'interrupt' through location 4.

Bit 8 Parity Disable

When this bit is clear, a parity error will cause an interrupt (or halt as specified in bit 9); if it is set, no action will be taken on a parity error.

When the machine is powered-up, the status registers have bit 15 set to 0, and the remaining bits set to 1: halt, odd parity enable; parity disable.

INDEX

Absolute Addressing	36	Operator's Console	194
Addressing	3, 25	Options	5
Architecture	9	PC	26
Autodecrement Mode	30	PDP-11 Family	1
Autoincrement Mode	29	Peripherals	5
Branch Instructions	77	PIC	140
Bus Options	6	Position Independent Code ...	140
Byte Instructions	43	Power Fail And Restart	4
Central Processor	2, 11	Power Failure	22
Co-Routines	141	Priority Interrupts	3, 19
Communications	7	Processor Status Word	13, 24
Core Memory	16	Processor Traps	22
Condition Codes	14	Processor Priority	14
Condition Code Instructions ..	126	Program Counter	26, 35
Console	194	Program Control Instructions ..	77
CPU	2, 11	Programming Techniques	127
Data Acquisition and Control ...	7	PS	13, 24
Deferred Addressing	32	Recursion	140
Direct Addressing	27	Reentrancy	137
Direct Memory Access	4	Reentrant Code	3
DMA	4	Register Mode	28
Double Operand Instructions ..	65	Registers	12
Floating Point Instructions ..	168	Relative Addressing	37
Floating Point Processor 4, 14,	163	Rotate Instructions	61
FORTRAN	6	RSTS-11	6
FPP	4	Shift Instructions	52
General Registers	12	Single Operand Instructions ...	45
Immediate Mode	35	Software	6
Index Mode	31	Solid State Memory	15
Indirect Addressing	32	SP	25
Instructions	2, 41	Stack Limit Register	14
Instruction Timing	217	Stack Pointer	25
Instruction Set Processor 42,	203	Stacks	4, 14, 127
Interleaving	16	Storage Devices	5
ISP	42, 203	Subroutine Instructions	98
Logical Instructions	73	Subroutines	131
Memories	4, 15	System Interaction	19
Memory Management	5, 143	T Bit	14
Memory Parity	17	Timing	217
Memory Retention	16	Traps	14, 22
Modes, Addressing	25	Trap Instructions	108
Nesting	135	UNIBUS	1, 10

INSTRUCTION INDEX

ADC(B)	58	HALT	119
ADD	67	INC(B)	48
ASL(B)	54	IOT	112
ASH	55	JMP	106
ASHC	56	JSR	99
ASR(B)	53	MARK	101
BCC	85	MFPD	125
BCS	84	MFPI	124
BEQ	80	MOV(B)	66
BGE	90	MTPD	123
BGT	92	MTPI	122
BHI	94	MUL	70
BHIS	97	NEG(B)	49
BIC(B)	76	RESET	121
BIS(B)	74	ROL(B)	62
BIT(B)	75	ROR(B)	63
BLT	89	RTI	113
BLE	91	RTS	103
BLO	96	RTT	114
BLOS	95	SBC(B)	59
BMI	82	SOB	107
BNE	81	SPL	105
BPL	83	SUB	68
BPT	111	SWAB	64
BR	78	SXT	60
BVC	87	TRAP	110
BVS	86	TST(B)	50
CLR(B)	46	WAIT	120
CMP(B)	69	XOR	72
COM(B)	51		
COND. CODES	126		
DEC(B)	47		
DIV	71		
EMT	109		

FPP INSTRUCTIONS

ABSD	187	MODF	179
ABSF	187	MULD	176
ADDD	173	MULF	176
ADDF	173		
		NEGD	175
CFCC	192	NEGF	175
CLRD	186		
CLRF	186	SETD	189
CMPD	178	SETF	189
CMPF	178	SETI	190
		SETL	190
		STCDF	181
DIVD	177	STCDI	183
DIVF	177	STCDL	183
		STCFD	181
LDCDF	180	STCFI	183
LDCFD	180	STCFL	183
LDCID	182	STD	172
LDCIF	182	STEXP	185
LDCLD	182	STF	172
LDCLF	182	STFPS	191
LDD	171	STST	192
LDEXP	184	SUBD	174
LDF	171	SUBF	174
LDFPS	191		
		TSTD	188
MODD	179	TSTF	188

DIGITAL EQUIPMENT CORPORATION **digital** WORLDWIDE SALES AND SERVICE

MAIN OFFICE AND PLANT

146 Main Street, Maynard, Massachusetts, U.S.A. 01754 • Telephone: From Metropolitan Boston, 646-8600 • Elsewhere: (617)-997-5111
 TWX: 710-347-0212 Cable: DIGITAL MAYN Telex: 94-8457

UNITED STATES

NORTHEAST

REGIONAL OFFICE:

275 Wyman Street, Waltham, Massachusetts 02154
 Telephone: (617)-890-0320/0330 TWX: 710-324-6919

WALTHAM

15 Lunds Street, Waltham, Massachusetts 02154
 Telephone: (617)-891-1030 TWX: 710-324-6919

CAMBRIDGE/BOSTON

899 Main Street, Cambridge, Massachusetts 02139
 Telephone: (617)-491-8130 TWX: 710-320-1187

ROCHESTER

130 Allens Creek Road, Rochester, New York 14618
 Telephone: (716)-461-1700 TWX: 710-253-3078

SYRACUSE

5858 East Molloy Road, Rm. 142, Picard Building
 Syracuse, New York 13211

Telephone: (315)-455-5987/88

CONNECTICUT

240 Pomeroy Avenue, Meriden, Connecticut 06450
 Telephone: (203)-237-8441/7466 TWX: 710-461-0054

MID-ATLANTIC — SOUTHEAST

REGIONAL OFFICE:

U.S. Route 1, Princeton, New Jersey 08540
 Telephone: (609)-452-2940 TWX: 510-885-2338

MANHATTAN

810 7th Ave.,
 New York, N.Y. 10019
 Telephone: (212)-582-1300

NW YORK

9 Cedar Lane, Englewood, New Jersey 07631
 Telephone: (201)-871-4694, (212)-594-6955, (212)-36-0447
 TWX: 710-991-9721

NEW JERSEY

1259 Route 46, Parsippany, New Jersey 07054
 Telephone: (201)-335-3300 TWX: 710-987-8319

PRINCETON

U.S. Route 1
 Princeton, New Jersey 08540
 Telephone: (609)-452-2940 TWX: 510-885-2338

METUCHEN

195 Main St.,
 Metuchen, N.J. 08840
 Telephone: (201)-549-4100/2000

MID-ATLANTIC — SOUTHEAST (cont.)

LONG ISLAND

1 Huntington Quadrangle
 Suite 1507 Huntington Station, New York 11746
 Telephone: (516)-894-4131, (212)-695-8095

PHILADELPHIA

Station Square Three, Paoli, Pennsylvania 19301
 Telephone: (215)-647-4900/4410 Telex: 510-668-8309

WASHINGTON

Executive Building
 6811 Kenilworth Ave., Riverdale, Maryland 20840
 Telephone: (301)-779-1600/752-8787 TWX: 710-826-9662

DURHAM/CHAPEL HILL

Executive Park
 3700 Chapel Hill Blvd.
 Durham, North Carolina 27707
 Telephone: (919)-489-3347 TWX: 510-927-0912

ORLANDO

Suite 130, 7001 Lake Ellenor Drive, Orlando, Florida 32809
 Telephone: (305)-851-4450 TWX: 810-850-0180

ATLANTA

2815 Clearview Place, Suite 100,
 Atlanta, Georgia 30340
 Telephone: (404)-458-3133/3134/3135 TWX: 810-755-4223

KNOXVILLE

6311 Kingston Pike, Suite 21E
 Knoxville, Tennessee 37919
 Telephone: (615)-588-6571 TWX: 810-583-0123

CENTRAL

REGIONAL OFFICE:

1850 Frontage Road, Northbrook, Illinois 60062
 Telephone: (312)-498-2500 TWX: 910-885-0655

PITTSBURGH

400 Penn. Center Boulevard
 Pittsburgh, Pennsylvania 15225
 Telephone: (412)-243-9404 TWX: 710-797-3657

CHICAGO

1850 Frontage Road, Northbrook, Illinois 60062
 Telephone: (312)-498-2500 TWX: 910-885-0655

ANN ARBOR

230 Huron View Boulevard, Ann Arbor, Michigan 48103
 Telephone: (313)-761-1150 TWX: 810-223-6053

DETROIT

2377 Greenfield Road, Suite 189
 Southfield, Michigan 48075
 Telephone: (313)-559-6565

CENTRAL (cont.)

INDIANAPOLIS

21 Beachway Drive — Suite G
 Indianapolis, Indiana 46224
 Telephone: (317)-243-8341 TWX: 810-341-3436

MINNEAPOLIS

Suite 111, 8030 Cedar Avenue South,
 Minneapolis, Minnesota 55420
 Telephone: (612)-854-6562-3-4-5 TWX: 910-576-2618

CLEVELAND

25000 Euclid Ave.,
 Euclid, Ohio 44117
 Telephone: (216)-946-8484 TWX: 810-427-2608

KANSAS CITY

532 East 42nd St., Independence, Missouri 64055
 Telephone: (816)-461-3440 TWX: 816-461-3100

ST. LOUIS

Suite 110, 115 Progress Parkway, Maryland Heights,
 Missouri 63043
 Telephone: (314)-878-4310 TWX: 910-784-0831

DAYTON

3101 Kettering Boulevard, Dayton, Ohio 45439
 Telephone: (513)-294-3323 TWX: 810-459-1678

MILWAUKEE

8531 W. Capitol Drive, Milwaukee, Wisconsin 53222
 Telephone: (414)-463-9110 TWX: 910-262-1199

DALLAS

8855 North Stemmons Freeway, Dallas, Texas 75247
 Telephone: (214)-638-4680 TWX: 910-861-4000

HOUSTON

3417 Milam Street, Suite A, Houston, Texas 77002
 Telephone: (713)-524-2961 TWX: 910-881-1851

NEW ORLEANS

3100 Ridgeland Drive, Suite 108
 Metairie, Louisiana 70002
 Telephone: (504)-837-0257

ROCKFORD

500 South Wyma St.
 Rockford, Illinois 61101
 Telephone: (815)-965-5557

TULSA

3140 S. Winston
 Winston Sq. Bldg.
 Tulsa, Oklahoma 74135
 Telephone: (918)-749-4476

WEST

REGIONAL OFFICE:

310 Soquel Way, Sunnyvale, California 94086
 Telephone: (408)-735-9200

SANTA ANA

2110 S. Anne St.,
 Santa Ana, Calif. 92704
 Telephone: (714)-979-2460

F.S.

714-979-2464 TWX: 910-381-1189

WEST LOS ANGELES

1510 Cotner Avenue, Los Angeles, California 90025
 Telephone: (213)-479-3191/4318 TWX: 910-342-6999

SAN DIEGO

6154 Mission Gorge Road, Suite 110
 San Diego, California 92120
 Telephone: (619)-280-7880, 7970 TWX: 910-335-1230

SAN FRANCISCO

1400 Terra Bella, Mountain View, California 94040
 Telephone: (415)-964-6200 TWX: 910-373-1266

OAKLAND

7850 Edgewater Drive, Oakland, California 94621
 Telephone: (415)-635-5453/7830 TWX: 910-366-7238

ALBUQUERQUE

6303 Indian School Road, N.E., Albuquerque, N.M. 87110
 Telephone: (505)-296-5411/5428 TWX: 910-989-0614

DENVER

2305 South Colorado Boulevard, Suite #5
 Denver, Colorado 80222
 Telephone: (303)-757-5332/758-1656/758-1659
 TWX: 910-931-2650

SEATTLE

1521 130th N.E., Bellevue, Washington 98005
 Telephone: (206)-454-4058/455-5404 TWX: 910-443-2306

SALT LAKE CITY

431 South 3rd East, Salt Lake City, Utah 84111
 Telephone: (801)-328-9838 TWX: 910-925-5834

PHOENIX

4358 East Broadway Road, Phoenix, Arizona 85040
 Telephone: (602)-268-3488 TWX: 910-950-4691

PORTLAND

Suite 168
 5319 S.W. Westgate Drive, Portland, Oregon 97221
 Telephone: (503)-297-3761/3765

digital

digital equipment corporation

67 • 00473 • 0743
JN-09 30

Printed in U.S.A.

pdp11/45