Book 1

Programming with the PDP-10 Instruction Set

PDP-10
System Reference Manual

Changes are indicated by a triangle ($\triangle$) in the outside margin
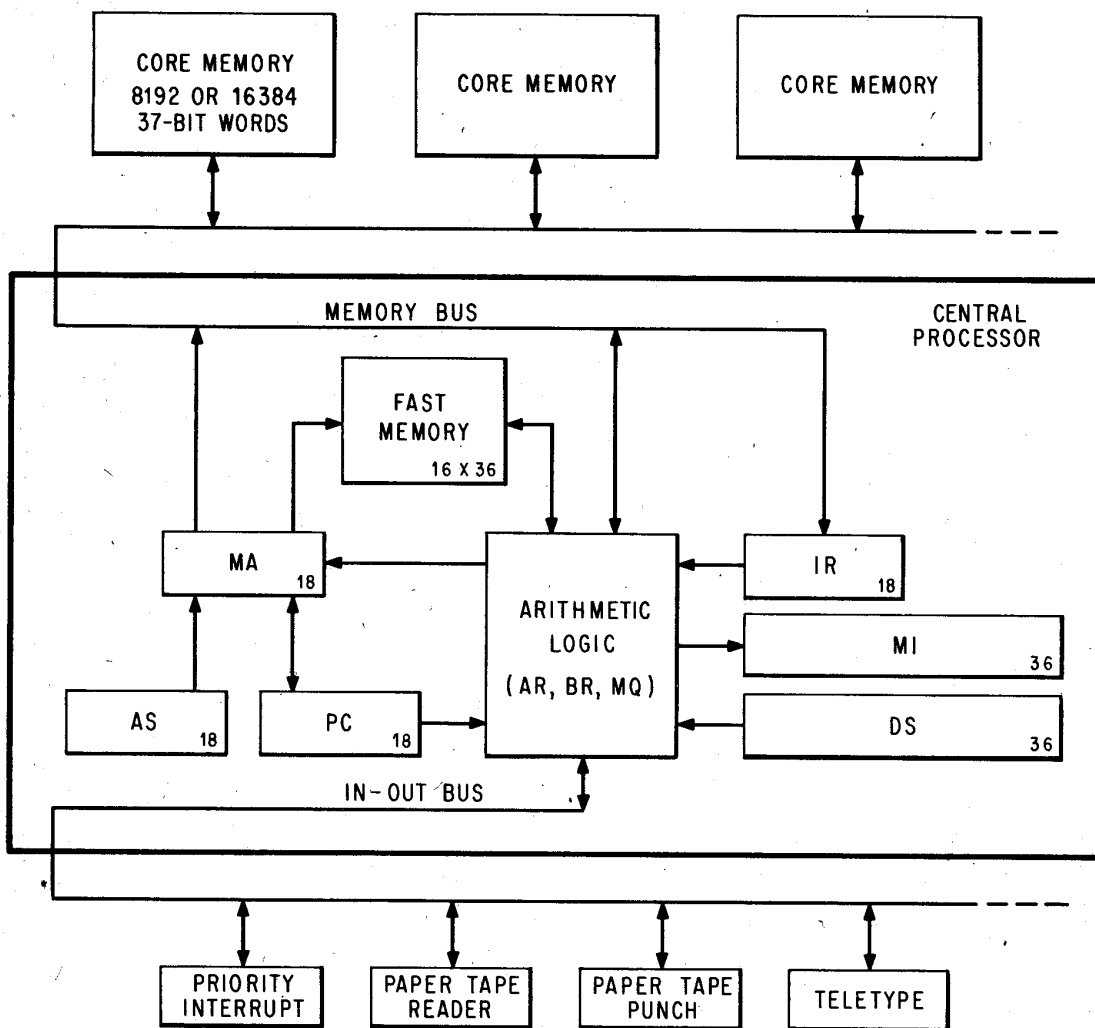
Contents

1

Introduction

The PDP-10 is a general purpose, stored program computer that includes a central processor, a memory, and a variety of peripheral equipment such as paper tape reader and punch, teletype, card reader, line printer, DECtape, magnetic tape, disk file and display. The central processor is the control unit for the entire system: it governs all peripheral in-out equipment, sequences the program, and performs all arithmetic, logical and data handling operations. The processor is connected to one or more memory units by a memory bus and to the peripheral equipment by an in-out bus. The fastest devices, such as the disc file, although controlled by the processor over the in-out bus, have direct access to memory over a second memory bus.

The processor handles words of thirty-six bits, which are stored in a memory with a maximum capacity of 262,144 words. Storage in memory is usually in the form of 37-bit words, the extra bit producing odd parity for the word. The bits of a word are numbered 0-35, left to right, as are the bits in the registers that handle the words. The processor can also handle half words, wherein the left half comprises bits 0-17, the right half, bits 18-35. Optional hardware is available for byte manipulation — a byte is any contiguous set of bits within a word. Registers that hold addresses have eighteen bits, numbered 18-35 according to the position of the address in a word. Words are used either as computer instructions in the program, as addresses, or as operands (data for the program).

Of the internal registers shown in the illustration on the next page, only PC, the 18 bit program counter, is directly relevant to the programmer. The processor performs a program by executing instructions retrieved from the locations addressed by PC. At the beginning of each instruction PC is incremented by one so that it normally contains an address one greater than the location of the current instruction. Sequential program flow is altered by changing the contents of PC, either by incrementing it an extra time in a skip instruction or by replacing its contents with the value specified by a jump instruction. Also of importance to the programmer is the 36-bit data switch register DS on the processor console: through this register the program can read data supplied by the operator. The processor also contains flags that detect various types of errors, including several types of overflow in arithmetic and pushdown operations, and provide other information of interest to the programmer.

The processor has other registers but the programmer is not usually concerned with them except when manually stepping through a program to debug it. By means of the address switch register AS, the operator can

PDP-10 SIMPLIFIED

examine the contents of, or deposit information into, any memory location; stop or interrupt the program whenever a particular location is referenced; and through AS the operator can supply a starting address for the program. Through the memory indicators MI the program can display data for the operator. The instruction register IR contains the left half of the current instruction word, *ie* all but the address part. The memory address register MA supplies the address for every memory access. The heart of the processor is the arithmetic logic, principally the 36-bit arithmetic register AR.

This register takes part in all arithmetic, logical and data handling operations; all data transfers to and from memory, peripheral equipment and console are made via AR. Associated with AR are an extremely fast full adder, a buffer register BR that holds a second operand in many arithmetic and logical instructions, a multiplier-quotient register MQ that serves primarily as an extension of AR for handling double length operands, and smaller registers that handle floating point exponents and control shift operations and byte manipulation.

From the point of view of the programmer however the arithmetic logic can be regarded as a black box. It performs almost all of the operations necessary for the execution of a program, but it never retains any information from one instruction to the next. Computations performed in the black box either affect control elements such as PC and the flags, or produce results that are always sent to memory and must be retrieved by the processor if they are to be used as operands in other instructions.

An instruction word has only one 18-bit address field for addressing any location throughout all of memory. But most instructions have two 4-bit fields for addressing the first sixteen memory locations. Any instruction that requires a second operand has an accumulator address field, which can address one of these sixteen locations as an accumulator; in other words as though it were a result held over in the processor from some previous instruction (the programmer usually has a choice of whether the result of the instruction will go to the location addressed as an accumulator or to that addressed by the 18-bit address field, or to both). Every instruction has a 4-bit index register address field, which can address fifteen of these locations for use as index registers in modifying the 18-bit memory address (a zero index register address specifies no indexing). Although all computations on both operands and addresses are performed in the single arithmetic register AR, the computer actually has sixteen accumulators, fifteen of which can double as index registers. The factor that determines whether one of the first sixteen locations in memory is an accumulator or an index register is not the information it contains nor how its contents are used, but rather how the location is addressed. There need be no difference physically between these locations and other memory locations, but an optional, fast flip-flop memory contained in the processor can be substituted for the bottom sixteen locations in core. This allows much quicker access to these locations whether they are addressed as accumulators, index registers or ordinary memory locations. They can even be addressed from the program counter, gaining faster execution for a short but oft-repeated subroutine.

Besides the registers that enter into the regular execution of the program and its instructions, the processor has a priority interrupt system and can contain optional equipment to facilitate time sharing. The interrupt system facilitates processor control of the peripheral equipment by means of a number of priority-ordered channels over which external signals may interrupt the normal program flow. The processor acknowledges an interrupt request by executing the instruction contained in a particular location assigned to the channel. Assignment of channels to devices is entirely under program control. One of the devices to which the program can assign a channel is the processor itself, allowing internal conditions such as overflow or a parity

error to signal the program.

The time share hardware provides memory protection and relocation. Without time sharing, all instructions and all memory are available to the program. Otherwise a number of programs share processor time, with each program relocated and restricted to a specific area in core, and certain instructions are usually illegal. An attempt by any user to execute an illegal instruction or address a memory location outside of his area results in a transfer of control back to the time-sharing monitor.

1.1 NUMBER SYSTEM

The program can interpret a data word as a 36-digit, unsigned binary number, or the left and right halves of a word can be taken as separate 18-bit numbers. The PDP-10 repertoire includes instructions that effectively add or subtract one from both halves of a word, so the right half can be used for address modification when the word is addressed as an index register, while the left half is used to keep a control count.

The standard arithmetic instructions in the PDP-10 use twos complement, fixed point conventions to do binary arithmetic. In a word used as a number, bit 0 (the leftmost bit) represents the sign, 0 for positive, 1 for negative. In a positive number the remaining 35 bits are the magnitude in ordinary binary notation. The negative of a number is obtained by taking its twos complement. If x is an n-digit binary number, its twos complement is $2^n - x$, and its ones complement is $(2^n - 1) - x$, or equivalently $(2^n - x) - 1$. Subtracting a number from $2^n - 1$ (ie, from all 1s) is equivalent to performing the logical complement, ie changing all 0s to 1s and all 1s to 0s. Therefore, to form the twos complement one takes the logical complement (usually referred to merely as the complement) of the entire word including the sign, and adds 1 to the result. In a negative number the sign bit is 1, and the remaining bits are the twos complement of the magnitude.

$$+153_{10} = +231_8 = \boxed{000\ 000\ 000\ 000\ 000\ 000\ 000\ 000\ 000\ 010\ 011\ 001}$$
$$\phantom{+153_{10} = +231_8 = }\ {}^0 \ {}^{35}$$

$$-153_{10} = -231_8 = \boxed{111\ 111\ 111\ 111\ 111\ 111\ 111\ 111\ 111\ 101\ 100\ 111}$$
$$\phantom{-153_{10} = -231_8 = }\ {}^0 \ {}^{35}$$

Zero is represented by a word containing all 0s. Complementing this number produces all 1s, and adding 1 to that produces all 0s again. Hence there is only one zero representation and its sign is positive. Since the numbers are symmetrical in magnitude about a single zero representation, all even numbers both positive and negative end in 0, all odd numbers in 1 (a number all 1s represents −1). But since there are the same number of positive and negative numbers and zero is positive, there is one more negative number than there are nonzero positive numbers. This is the most negative number and it cannot be produced by negating any positive number (its octal representa-

tion is 400000 000000$_8$ and its magnitude is one greater than the largest positive number).

If ones complements were used for negatives one could read a negative number by attaching significance to the 0s instead of the 1s. In twos complement notation each negative number is one greater than the complement of the positive number of the same magnitude, so one can read a negative number by attaching significance to the rightmost 1 and attaching significance to the 0s at the left of it (the negative number of largest magnitude has a 1 in only the sign position). In a negative integer, 1s may be discarded at the left, just as leading 0s may be dropped in a positive integer. In a negative fraction, 0s may be discarded at the right. So long as only 0s are discarded, the number remains in twos complement form because it still has a 1 that possesses significance; but if a portion including the rightmost 1 is discarded, the remaining part of the fraction is now a ones complement.

The computer does not keep track of a binary point — the programmer must adopt a point convention and shift the magnitude of the result to conform to the convention used. Two common conventions are to regard a number as an integer (binary point at the right) or as a proper fraction (binary point at the left); in these two cases the range of numbers represented by a single word is -2^{35} to $2^{35} - 1$ or -1 to $1 - 2^{-35}$. Since multiplication and division make use of double length numbers, there are special instructions for performing these operations with integral operands.

Floating Point Arithmetic. Optional PDP-10 hardware is available for processing floating point numbers. A floating point instruction interprets bit 0 of a word as the sign, but interprets the rest of the word as an 8-bit exponent and a 27-bit fraction. For a positive number the sign is 0, as before. But the contents of bits 9-35 are now interpreted only as a binary fraction, and the contents of bits 1-8 are interpreted as an integral exponent in excess 128 (200$_8$) code. Exponents from -128 to $+127$ are therefore represented by the binary equivalents of 0 to 255 (0-377$_8$). Floating point zero and negatives are represented in exactly the same way as in fixed point: zero by a word containing all 0s, a negative by the twos complement. A negative number has a 1 for its sign and the twos complement of the fraction, but since every f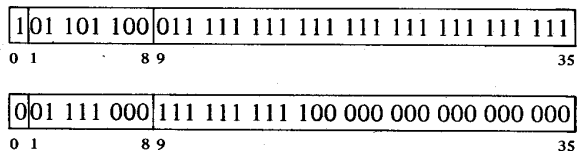raction must ordinarily contain a 1 unless the entire number is zero (see below), it has the ones complement of the exponent code in bits 1-8. Since the exponent is in excess 128 code, an actual exponent x is represented in a positive number by $x + 128$, in a negative number by $127 - x$. The programmer, however, need not be concerned with these representations as the hardware compensates automatically. *Eg*, for

Multiplication produces a double length product, and the programmer must remember that discarding the low order part of a double length negative leaves the high order part in correct twos complement form only if the low order part is null.

$$+153_{10} = +231_8 = +.462_8 \times 2^8 =$$

0	10 001 000	100 110 010 000 000 000 000 000 000

0 1   8 9   35

$$-153_{10} = -231_8 = -.462_8 \times 2^8 =$$

1	01 110 111	011 001 110 000 000 000 000 000 000

0 1   8 9   35

the instruction that scales the exponent, the hardware interprets the integral scale factor in standard twos complement form but produces the correct ones complement result for the exponent.

Except in special cases the floating point instructions assume that all non-zero operands are normalized, and they normalize a nonzero result. A floating point number is considered normalized if the magnitude of the fraction is greater than or equal to ½ and less than 1. These numbers thus have a fractional range in magnitude of ½ to $1 - 2^{-27}$ and an exponent range of -128 to $+127$. The hardware may not give the correct result if the program supplies an operand that is not normalized or that has a zero fraction with a nonzero exponent.

The precaution about truncation given for fixed point multiplication applies to all floating point operations as they all produce extra length results; but here the programmer may request rounding, which automatically restores the high order part to twos complement form if it is negative. In division the two words of the result are quotient and remainder, but in the other operations they form a double length number which is stored in two accumulators if the instruction is executed in "long" mode. This number contains a 54-bit fraction, half of which is in bits 9–35 of each word. The sign and exponent are in bits 0 and 1–8 respectively of the word containing the more significant half, and the standard twos complement is used to form the negative of the entire 63-bit string. In the remaining part of the less significant word, bit 0 is 0, and bits 1–8 contain a number 27 less than the exponent, but this is expressed in positive form even though bits 9–35 may be part of a negative fraction. *Eg* the number $2^{18} + 2^{-18}$ has this two-word representation:

```
| 0 | 10 010 011 | 100 000 000 000 000 000 000 000 000 |
  0  1            8 9                                 35
```

```
| 0 | 01 111 000 | 000 000 000 100 000 000 000 000 000 |
  0  1            8 9                                 35
```
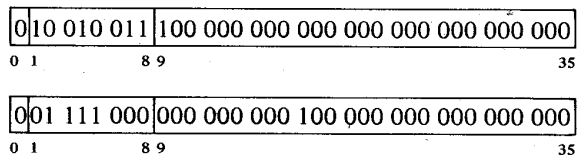
whereas its negative is

```
| 1 | 01 101 100 | 011 111 111 111 111 111 111 111 111 |
  0  1            8 9                                 35
```

```
| 0 | 01 111 000 | 111 111 111 100 000 000 000 000 000 |
  0  1            8 9                                 35
```

1.2 INSTRUCTION FORMAT

In all but the input-output instructions, the nine high order bits (0–8) specify the operation, and bits 9–12 usually address an accumulator but are sometimes used for special control purposes, such as addressing flags. The

rest of the instruction word usually supplies information for calculating the effective address, which is the actual address used to fetch the operand or alter program flow. Bit 13 specifies the type of addressing, bits 14–17 specify an index register for use in address modification, and the remaining eighteen bits (18–35) address a memory location. The instruction codes


BASIC INSTRUCTION FORMAT
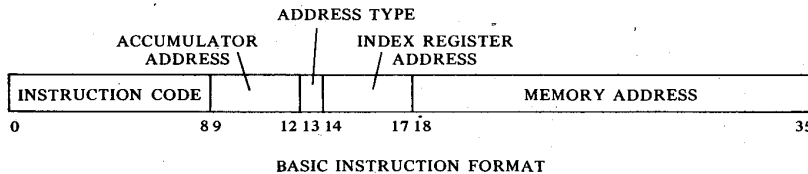
that are not assigned as specific instructions are executed by the processor as so-called "unimplemented operations", as are the codes for floating point and byte manipulation in any PDP–10 that does not have the optional hardware for these instructions. When the processor encounters one of these unimplemented codes in a program, it stores bits 0–12 of the instruction word and the calculated effective address in a particular memory location and then executes the instruction contained in a second location.

An input-output instruction is designated by three 1s in bits 0–2. Bits 3–9 address the in-out device to be used in executing the instruction, and bits 10–12 specify the operation. The rest of the word is the same as in other instructions.

IN-OUT INSTRUCTION FORMAT
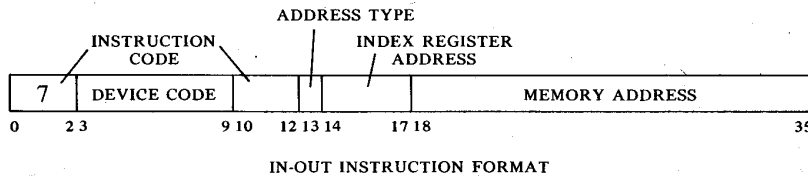
Effective Address Calculation. Bits 13–35 have the same format in *every* instruction whether it addresses a memory location or not. Bit 13 is the
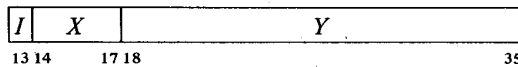
indirect bit, bits 14–17 are the index register address, and if the instruction must reference memory, bits 18–35 are the memory address Y. The effective address E of the instruction depends on the values of I, X and Y. If X is nonzero, the contents of index register X are added to Y to produce a modified address. If I is 0, addressing is direct, and the modified address is the effective address used in the execution of the instruction; if I is 1, addressing is indirect, and the processor retrieves another address word from the location specified by the modified address already determined. This new word is processed in exactly the same manner: X and Y determine the effective address if I is 0, otherwise they are used for yet another level of address

retrieval. This process continues until some referenced location is found with a 0 in bit 13; the 18-bit number calculated from the X and Y parts of this location is the effective address E.

The calculation outlined above is carried out for *every* instruction even if it need not address a memory location. If the indirect bit in the instruction word is 0 and no memory reference is necessary, then Y is not an address. It may be a mask in some kind of test instruction, conditions to be sent to an in-out device, or part of it may be the number of places to shift in a shift or rotate instruction or the scale factor in a floating scale instruction. Even when modified by an index register, bits 18-35 do not contain an address when I is 0. But when I is 1, the number determined from bits 14-35 is an indirect address no matter what type of information the instruction requires, and the word retrieved in any step of the calculation contains an indirect address so long as I remains 1. When a location is found in which I is 0, bits 18-35 (perhaps modified by an index register) contain the desired effective mask, effective conditions, effective shift number, or effective scale factor. Many of the instructions that usually reference memory for an operand even have an "immediate" mode in which the result of the effective address calculation is itself used as a half word operand instead of a word taken from the memory location it addresses.

The important thing for the programmer to remember is that the same calculation is carried out for every instruction regardless of the type of information that must be specified for its execution, or even if the result is ignored. In the discussion of any instruction, E refers to the actual quantity derived from I, X and Y and used in the execution of the instruction, be it the entire half word as in the case of an address, immediate operand, mask or conditions, or only part of it as in a shift number or scale factor.

1.3 MEMORY

All timing in the PDP-10 is asynchronous. The internal timing for each in-out device and each memory is entirely independent of the central processor. Because core memory readout is destructive, every word read must be written back in unless new information is to take its place. The basic read-write cycle time of the standard core memory is either 1.00 or 1.65 microseconds, but the processor need never wait the entire cycle time. To read, it waits only until the information is available and then continues its operations while the memory performs the write portion of the cycle; to write, it waits only until the data is accepted, and the memory then performs an entire cycle to clear and write. To save time in an instruction that fetches an operand and then writes new data into the same location, the memory executes a read-pause-write cycle in which it performs only the read part initially and then completes the cycle when the processor supplies the new data.

Access times for the accumulator-index register locations are decreased considerably by substitution of a fast memory (contained in the processor) for the first sixteen core locations. Readout is nondestructive, so the fast memory has no basic cycle: the processor reads a word directly, but to write

it must first clear the location and then load it. Access times in nanoseconds (including 20 feet of cable delay) for the three memories are as follows.

	Read	Write
MA10 or MA10A Core Memory (1.00 μs)	580	200 ▲
MB10 Core Memory (1.65 μs)	600 (700)*	200 (300)
KM10 Fast Memory (18-bit address)	210	210

*Numbers in parentheses are the longer times required in a multiprocessor system.

NOTE: When a fast memory location is addressed as an accumulator or index register, the access time is usually considerably shorter than that listed here.

From the simple addressing point of view, the entire memory is a set of contiguous locations whose addresses range from zero to a maximum dependent upon the capacity of the particular installation. In a system with the greatest possible capacity, the largest address is octal 777777, decimal 262,143. (Addresses are always in octal notation unless otherwise specified.) But the whole memory would usually be made up of a number of core memories each having a capacity of 8192 or 16,384 words. Hence a single 18-bit address actually selects a particular memory and a specific location within it. For an 8K memory the high order five address bits select the memory, the remaining thirteen bits address a single location in it; selecting a 16K memory takes four bits, leaving fourteen for the location. The times given above assume the addressed memory is idle when access is requested. To avoid waiting for a previously requested memory cycle to end, the program can make consecutive requests to different memories by taking instructions from one memory and data from another. The hardware also allows pairs of memories to be interleaved in such a way that consecutive addresses actually alternate between the two memories in the pair (thus increasing the probability that consecutive references are to different memories). Appropriate switch settings at the memories interchange the least significant address bits in the memory and location parts, so that in any two memories numbered n and $n + 1$ where n is even, all even addresses are locations in the first memory, all odd addresses are locations in the second. Hence memories 0 and 1 can be interleaved as can 6 and 7, but not 3 and 4 or 5 and 7.

Memory Allocation. The use of certain memory locations is defined by the hardware.

0	Holds a pointer word during a bootstrap readin
0–17	Can be addressed as accumulators
1–17	Can be addressed as index registers
40–41	Trap for unimplemented user operations (UUOs)
42–57	Priority interrupt locations
60–61	Trap for remaining unimplemented operations: these include the unassigned instruction codes that are reserved for future use, and also the byte manipulation and floating point instructions when the hardware for them is not installed
140–161	Allocated to second processor if connected (same use as 40–61 for first processor)

All information given in this manual about memory locations 40–61 applies instead to locations 140–161 for programming a second central processor connected to the same memory.

The initial control word address for the DF10 Data Channel must be less than 1000.

1.4 PROGRAMMING CONVENTIONS

The computer has five instruction classes: data transmission, logical, arithmetic, program control and in-out. The instructions in the in-out class control the peripheral equipment, and also control the priority interrupt and time sharing, control and read the processor flags, and communicate with the console.   The next chapter describes all instructions mentioned above, presents a general description of input-output, and describes the effects of the in-out instructions on the processor, priority interrupt and time share hardware. Effects of in-out instructions on particular peripheral devices are discussed with the devices.

The MACRO−10 assembly program recognizes a number of mnemonics and other initial symbols that facilitate constructing complete instruction words and organizing them into a program. In particular there are mnemonics for the instruction codes (Appendix A), which are six bits in in-out instructions, otherwise nine or thirteen bits. *Eg* the mnemonic

The assembler translates every statement into a 36-bit word, placing 0s in all bits whose values are unspecified.

> MOVNS

assembles as 213000 000000, and

> MOVNS   2570

assembles as 213000 002570. This latter word, when executed as an instruction, produces the twos complement negative of the word in memory location 2570.

NOTE

Throughout this manual all numbers representing instruction words, register contents, codes and addresses are always octal, and any numbers appearing in program examples are octal unless otherwise indicated. On the other hand, the ordinary use of numbers in the text to count steps in an operation or to specify word or byte lengths, bit positions, exponents, etc employs standard decimal notation.

The initial symbol @ preceding a memory address places a 1 in bit 13 to produce indirect addressing. The example given above uses direct addressing, but

> MOVNS   @2570

assembles as 213020 002570, and produces indirect addressing. Placing the number of an index register (1−17) in parentheses following the memory address causes modification of the address by the contents of the specified register. Hence

> MOVNS   @2570(12)

which assembles as 213032 002570, produces indexing using index register 12, and the processor then uses the modified address to continue the effective address calculation.

An accumulator address (0−17) precedes the memory address part (if any)

and is terminated by a comma. Thus

         MOVNS   4,@2570(12)

assembles as 213232 002570, which negates the word in location E and stores the result in both E and in accumulator 4. The same procedure may be used to place 1s in bits 9–12 when these are used for something other than addressing an accumulator, but mnemonics are available for this purpose.

   The device code in an in-out instruction is given in the same manner as an accumulator address (terminated by a comma and preceding the address part), but the number given must correspond to the octal digits in the word (000–774). Mnemonics are however available for all standard device codes. To control the priority interrupt system whose code is 004, one may give

         CONO   4,1302

which assembles as 700600 001302, or equivalently   ▲

         CONO   PI,1302

   The programming examples in this manual use the following addressing conventions:
◆ A colon following a symbol indicates that it is a symbolic location name.

A: ADD   6,5704

indicates that the location that contains ADD 6,5704 may be addressed symbolically as A.
◆ The period represents the current address, eg

         ADD   5,.+2

is equivalent to

A: ADD   5,A+2

◆ Square brackets specify the contents of a location, leaving the address of the location implicit but unspecified. Eg

         ADD   12,[7256004]

and

         ADD   12,A

          :
          :
A: 7256004

are equivalent.
   Anything written at the right of a semicolon is commentary that explains the program but is not part of it.

2

Central Processor

This chapter describes all PDP-10 instructions but does not discuss the effects of those in-out instructions that address specific peripheral devices. In the description of each instruction, the mnemonic and name are at the top, the format is in a box below them. The mnemonic assembles to the word in the box, where bits in those parts of the word represented by letters assemble as 0s. The letters indicate portions that must be added to the mnemonic to produce a complete instruction word.

For many of the non-IO instructions, a description applies not to a unique instruction with a single code in bits 0–8, but rather to an instruction set defined as a basic instruction that can be executed in a number of modes. These modes define properties subsidiary to the basic operation; *eg* in data transmission the mode specifies which of the locations addressed by the instruction is the source and which the destination of the data, in test instructions it specifies the condition that must be satisfied for a jump or skip to take place. The mnemonic given at the top is for the basic mode; mnemonics for the other forms of the instruction are produced by appending letters directly to the basic mnemonic. Following the description is a table giving the mnemonics and octal codes (bits 0–8) for the various modes.

The processor execution time for each instruction is also given at the top unless the time differs from one mode to another. The time listed is that required for direct addressing without indexing (*ie* with no effective address calculation), assuming the instruction and location E are both in the same 1.00 microsecond core memory, and that an accumulator is addressed only if necessary and is in fast memory. The time that can be saved (if any) by interleaving or keeping instructions and operands in different memories is indicated either with the description or with the discussion of the modes preceding a group of instructions. To determine the exact time required for an instruction under any circumstances, refer to the timing chart in Appendix C.

In a description E refers to the effective address, half word operand, mask, conditions, shift number or scale factor calculated from the I, X and Y parts of the instruction word. In an instruction that ordinarily references memory, a reference to E as the source of information means that the instruction retrieves the word contained in location E; as a destination it means the instruction stores a word in location E. In the immediate mode of these instructions, the effective half word operand is usually treated as a full word that contains E in one half and zero in the other, and is represented either as $0,E$ or $E,0$ depending upon whether E is in the right or left half.

Letters representing modes are suffixes, which produce new mnemonics that are recognized as distinct symbols by the assembler.

The times listed should be re- ▲ garded as good approximations. For more exact times with the conditions given here (*ie* 1.00 microsecond core, etc) add 60 nanoseconds to the listed time, plus an additional 30 nanoseconds for each core memory access for retrieval of an operand and another 30 nanoseconds if the instruction does not write a result in core.

Most of the non-IO instructions can address an accumulator, and in the box showing the format this address is represented by A; in the description, "AC" refers to the accumulator addressed by A. "AC left" and "AC right" refer to the two halves of AC. If an instruction uses two accumulators, these have addresses A and $A+1$, where the second address is 0 if A is 17. In some cases an instruction uses an accumulator only if A is nonzero: a zero address in bits 9–12 specifies no accumulator.

It is assumed throughout that time sharing is not in effect, and the program is unrestricted. For completeness, however, the effects of restrictions on particular instructions are noted; and execution times are given both for unrestricted operation and including relocation in a user program (the latter time is given in parentheses). §2.15 lists all restrictions on user programs and explains the special effects produced by certain instructions when executed under control of the monitor while the processor is in user mode.

Some simple examples are included with the instruction descriptions, but more complex examples using a variety of instructions are given in §2.11.

2.1 HALF WORD DATA TRANSMISSION

These instructions move a half word and may modify the contents of the other half of the destination location. There are sixteen instructions determined by which half of the source word is moved to which half of the destination, and by which of four possible operations is performed on the other half of the destination. The basic mnemonics are three letters that indicate the transfer

HLL	Left half of source to left half of destination
HRL	Right half of source to left half of destination
HRR	Right half of source to right half of destination
HLR	Left half of source to right half of destination

plus a fourth, if necessary, to indicate the operation.

Operation	Suffix	Effect on Other Half of Destination
Do nothing		None
Zeros	Z	Places 0s in all bits of the other half
Ones	O	Places 1s in all bits of the other half
Extend	E	Places the sign (the leftmost bit) of the half word moved in all bits of the other half. This action extends a right half word number into a full word number but is valid arithmetically only for positive left half word numbers – the right extension of a number requires 0s regardless of sign (hence the Zeros operation should be used to extend a left half word number).

An additional letter may be appended to indicate the mode, which determines the source and destination of the half word moved.

Mode	Suffix	Source	Destination
Basic		E	AC
Immediate	I	The word 0, E	AC
Memory	M	AC	E
Self	S	E	E, but also AC if A is nonzero

Note that selecting the left half of the source in immediate mode merely clears the selected half of the destination.

Keeping instructions and operands in different memories saves .20 (.09) μs in self mode; in memory mode the same saving results if no action is taken on the other half, otherwise .47 (.36) μs is saved.

When E addresses a fast memory location, a half word transfer takes .34 μs less in basic mode, either .46 (.35) or .54 (.43) μs less in memory mode depending respectively on whether or not any action is taken on the other half, and .54 (.43) μs less in self mode.

HLL   Half Word Left to Left
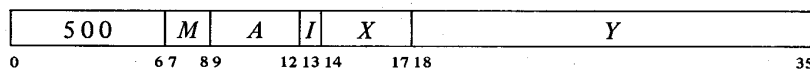
5 0 0	M	A	I	X	Y

0   6 7   8 9   12 13 14   17 18   35

Move the left half of the source word specified by M to the left half of the specified destination. The source and the destination right half are unaffected; the original contents of the destination are lost.

HLL	Half Left to Left	500	2.35 (2.57) μs
HLLI	Half Left to Left Immediate	501	1.50 (1.61) μs
HLLM	Half Left to Left Memory	502	2.90 (3.01) μs
HLLS	Half Left to Left Self	503	2.76 (2.87) μs

HLLI merely clears AC left. If A is zero, HLLS is a no-op, otherwise it is equivalent to HLL.

HLLZ   Half Word Left to Left, Zeros

5 1 0	M	A	I	X	Y

0   6 7   8 9   12 13 14   17 18   35

Move the left half of the source word specified by M to the left half of the specified destination, and clear the destination right half. The source is unaffected, the original contents of the destination are lost.

HLLZ	Half Left to Left, Zeros	510 2.21 (2.43) μs
HLLZI	Half Left to Left, Zeros, Immediate	511 1.36 (1.47) μs
HLLZM	Half Left to Left, Zeros, Memory	512 2.47 (2.58) μs
HLLZS	Half Left to Left, Zeros, Self	513 2.76 (2.87) μs

HLLZI merely clears AC. If A is zero, HLLZS merely clears the right half of location E.

HLLO   Half Word Left to Left, Ones
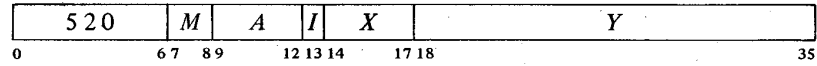
5 2 0	M	A	I	X	Y
0	6 7	8 9	12 13 14	17 18	35

Move the left half of the source word specified by M to the left half of the specified destination, and set the destination right half to all 1s. The source is unaffected, the original contents of the destination are lost.

HLLO	Half Left to Left, Ones	520
		2.21 (2.43) μs

HLLOI sets AC to all 0s in the left half, all 1s in the right.

HLLOI	Half Left to Left, Ones, Immediate	521
		1.36 (1.47) μs
HLLOM	Half Left to Left, Ones, Memory	522
		2.47 (2.58) μs
HLLOS	Half Left to Left, Ones, Self	523
		2.76 (2.87) μs

HLLE   Half Word Left to Left, Extend
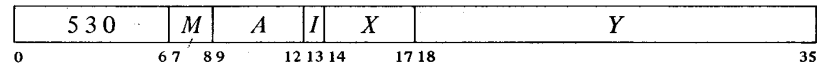
5 3 0	M	A	I	X	Y
0	6 7	8 9	12 13 14	17 18	35

Move the left half of the source word specified by M to the left half of the specified destination, and make all bits in the destination right half equal to bit 0 of the source. The source is unaffected, the original contents of the destination are lost.

HLLE	Half Left to Left, Extend	530
		2.21 (2.43) μs

HLLEI is equivalent to HLLZI (it merely clears AC).

HLLEI	Half Left to Left, Extend, Immediate	531
		1.36 (1.47) μs
HLLEM	Half Left to Left, Extend, Memory	532
		2.47 (2.58) μs
HLLES	Half Left to Left, Extend, Self	533
		2.76 (2.87) μs

HRL   Half Word Right to Left

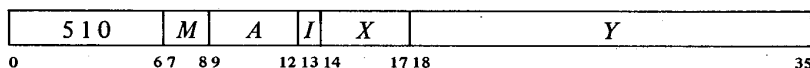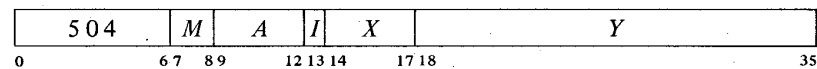| 5 0 4	M	A	I	X	Y
0	6 7	8 9	12 13 14	17 18	35

Move the right half of the source word specified by M to the left half of the specified destination. The source and the destination right half are unaffected; the original contents of the destination left half are lost.

HRL	Half Right to Left	504	2.70 (2.92) μs
HRLI	Half Right to Left Immediate	505	1.85 (1.96) μs

HRLM	Half Right to Left Memory	506	2.90 (3.01) μs
HRLS	Half Right to Left Self	507	2.76 (2.87) μs

HRLZ   Half Word Right to Left, Zeros
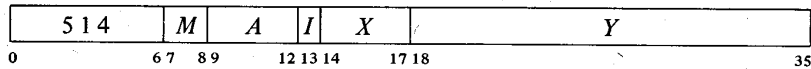
5 1 4	M	A	I	X	Y
0	6 7	8 9	12 13 14	17 18	35

Move the right half of the source word specified by M to the left half of the specified destination, and clear the destination right half. The source is unaffected, the original contents of the destination are lost.

HRLZ	Half Right to Left, Zeros	514
		2.21 (2.43) μs
HRLZI	Half Right to Left, Zeros, Immediate	515
		1.36 (1.47) μs
HRLZM	Half Right to Left, Zeros, Memory	516
		2.47 (2.58) μs
HRLZS	Half Right to Left, Zeros, Self	517
		2.76 (2.87) μs

HRLZI loads the word $E,0$ into AC.

HRLO   Half Word Right to Left, Ones
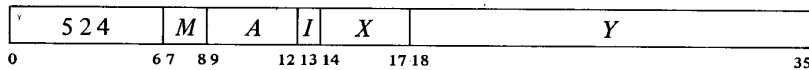
5 2 4	M	A	I	X	Y
0	6 7	8 9	12 13 14	17 18	35

Move the right half of the source word specified by M to the left half of the specified destination, and set the destination right half to all 1s. The source is unaffected, the original contents of the destination are lost.

HRLO	Half Right to Left, Ones	524
		2.21 (2.43) μs
HRLOI	Half Right to Left, Ones, Immediate	525
		1.36 (1.47) μs
HRLOM	Half Right to Left, Ones, Memory	526
		2.47 (2.58) μs
HRLOS	Half Right to Left, Ones, Self	527
		2.76 (2.87) μs

HRLE   Half Word Right to Left, Extend
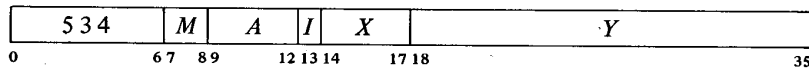
5 3 4	M	A	I	X	Y
0	6 7	8 9	12 13 14	17 18	35

Move the right half of the source word specified by M to the left half of the

specified destination, and make all bits in the destination right half equal to bit 18 of the source. The source is unaffected, the original contents of the destination are lost.

HRLE	Half Right to Left, Extend	534
		2.21 (2.43) μs
HRLEI	Half Right to Left, Extend, Immediate	535
		1.36 (1.47) μs
HRLEM	Half Right to Left, Extend, Memory	536
		2.47 (2.58) μs
HRLES	Half Right to Left, Extend, Self	537
		2.76 (2.87) μs

HRR   Half Word Right to Right
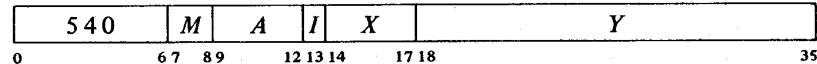
5 4 0	M	A	I	X	Y
0	6 7	8 9	12 13 14	17 18	35

Move the right half of the source word specified by M to the right half of the specified destination. The source and the destination left half are unaffected; the original contents of the destination right half are lost.

HRR	Half Right to Right	540	2.35 (2.57) μs
HRRI	Half Right to Right Immediate	541	1.50 (1.61) μs
HRRM	Half Right to Right Memory	542	2.90 (3.01) μs
HRRS	Half Right to Right Self	543	2.76 (2.87) μs

If A is zero, HRRS is a no-op; otherwise it is equivalent to HRR.

HRRZ   Half Word Right to Right, Zeros

5 5 0	M	A	I	X	Y
0	6 7	8 9	12 13 14	17 18	35

Move the right half of the source word specified by M to the right half of the specified destination, and clear the destination left half. The source is unaffected, the original contents of the destination are lost.

HRRZ	Half Right to Right, Zeros	550
		2.21 (2.43) μs
HRRZI	Half Right to Right, Zeros, Immediate	551
		1.36 (1.47) μs
HRRZM	Half Right to Right, Zeros, Memory	552
		2.47 (2.58) μs
HRRZS	Half Right to Right, Zeros, Self	553
		2.76 (2.87) μs

HRRZI loads the word $0,E$ into AC. If A is zero, HRRZS merely clears the left half of location E.

HRRO   **Half Word Right to Right, Ones**
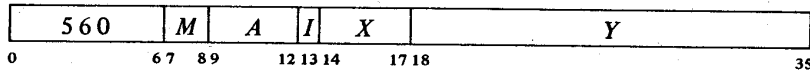
5 6 0	M	A	I	X	Y
0	6 7	8 9	12 13 14	17 18	35

Move the right half of the source word specified by M to the right half of the specified destination, and set the destination left half to all 1s. The source is unaffected, the original contents of the destination are lost.

HRRO	Half Right to Right, Ones	560
		2.21 (2.43) μs
HRROI	Half Right to Right, Ones, Immediate	561
		1.36 (1.47) μs
HRROM	Half Right to Right, Ones, Memory	562
		2.47 (2.58) μs
HRROS	Half Right to Right, Ones, Self	563
		2.76 (2.87) μs

HRRE   **Half Word Right to Right, Extend**

5 7 0	M	A	I	X	Y ·
0	6 7	8 9	12 13 14	17 18	35

Move the right half of the source word specified by M to the right half of the specified destination, and make all bits in the destination left half equal to bit 18 of the source. The source is unaffected, the original contents of the destination are lost.

HRRE	Half Right to Right, Extend	570
		2.21 (2.43) μs
HRREI	Half Right to Right, Extend, Immediate	571
		1.36 (1.47) μs
HRREM	Half Right to Right, Extend, Memory	572
		2.47 (2.58) μs
HRRES	Half Right to Right, Extend, Self	573
		2.76 (2.87) μs

HLR   **Half Word Left to Right**

5 4 4	M	A	I	X	Y
0	6 7	8 9	12 13 14	17 18	35

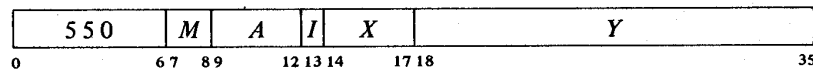Move the left half of the source word specified by M to the right half of the specified destination. The source and the destination left half are unaffected; the original contents of the destination right half are lost.

HLR	Half Left to Right	544	2.70 (2.92) μs	
HLRI	Half Left to Right Immediate	545	1.85 (1.96) μs	HLRI merely clears AC right.

| HLRM | Half Left to Right Memory | 546 | 2.90 (3.01) μs |
| HLRS | Half Left to Right Self | 547 | 2.76 (2.87) μs |

HLRZ   Half Word Left to Right, Zeros
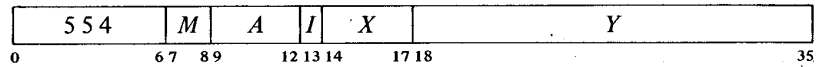
5 5 4	M	A	I	X	Y
0	6 7	8 9	12 13 14	17 18	35

Move the left half of the source word specified by M to the right half of the specified destination, and clear the destination left half. The source is unaffected, the original contents of the destination are lost.

| HLRZ | Half Left to Right, Zeros | 554 |
| | | 2.21 (2.43) μs |

HLRZI merely clears AC and is thus equivalent to HLLZI.

HLRZI	Half Left to Right, Zeros, Immediate	555
		1.36 (1.47) μs
HLRZM	Half Left to Right, Zeros, Memory	556
		2.47 (2.58) μs
HLRZS	Half Left to Right, Zeros, Self	557
		2.76 (2.87) μs

HLRO   Half Word Left to Right, Ones
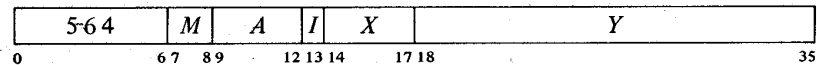
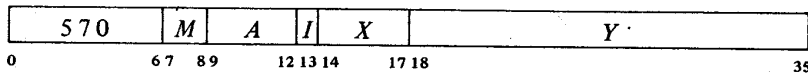| 5 6 4	M	A	I	X	Y
0	6 7	8 9	12 13 14	17 18	35

Move the left half of the source word specified by M to the right half of the specified destination, and set the destination left half to all 1s. The source is unaffected, the original contents of the destination are lost.

| HLRO | Half Left to Right, Ones | 564 |
| | | 2.21 (2.43) μs |

HLROI sets AC to all 1s in the left half, all 0s in the right.

HLROI	Half Left to Right, Ones, Immediate	565
		1.36 (1.47) μs
HLROM	Half Left to Right, Ones, Memory	566
		2.47 (2.58) μs
HLROS	Half Left to Right, Ones, Self	567
		2.76 (2.87) μs

HLRE   Half Word Left to Right, Extend

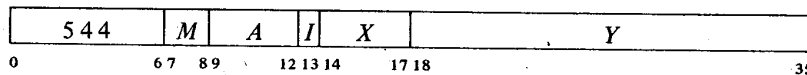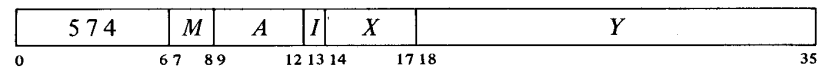| 5 7 4	M	A	I	X	Y
0	6 7	8 9	12 13 14	17 18	35

Move the left half of the source word specified by M to the right half of the specified destination, and make all bits in the destination left half equal to

bit 0 of the source. The source is unaffected, the original contents of the destination are lost.

HLRE	Half Left to Right, Extend	574
		2.21 (2.43) μs
HLREI	Half Left to Right, Extend, Immediate	575
		1.36 (1.47) μs
HLREM	Half Left to Right, Extend, Memory	576
		2.47 (2.58) μs
HLRES	Half Left to Right, Extend, Self	577
		2.76 (2.87) μs

HLREI is equivalent to HLRZI (it merely clears AC).

EXAMPLES. The half word transmission instructions are very useful for handling addresses, and they provide a convenient means of setting up an accumulator whose right half is to be used for indexing while a control count is kept in the left half. *Eg* this pair of instructions loads the 18-bit numbers *M* and *N* into the left and right halves respectively of an accumulator that is addressed symbolically as XR.

        HRLZI   XR,*M*
        HRRI   XR,*N*

Of course the source program must somewhere define the value of the symbol XR as an octal number between 1 and 17.

Suppose that at some point we wish to use the two halves of XR independently as operands (taken as 18-bit positive numbers) for computations. We can begin by moving XR left to the right half of another accumulator AC and leaving the contents of XR right alone in XR.

        HLRZM   XR,AC
        HLLI   XR,   ;Clear XR left

It is not necessary to clear the other half of XR when loading the first half word. But any instruction that modifies the other half is faster than the corresponding instruction that does not, as the latter must fetch the destination word in order to save half of it. (The difference does not apply to self mode, for here the source and destination are the same.)

2.2 FULL WORD DATA TRANSMISSION

These are the instructions whose basic purpose is to move one or more full words of data from one place to another, usually from an accumulator to a memory location or vice versa. In a few cases instructions may perform minor arithmetic operations, such as forming the negative or the magnitude of the word being processed.

EXCH   **Exchange**   2.90 (3.01) μs

2 5 0		A	I	X	Y
0		8 9	12 13 14	17 18	35

Keeping instructions and operands in different memories saves .20 (.09) μs.
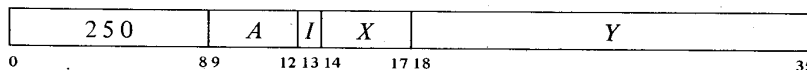
Move the contents of location *E* to AC and move AC to location *E*.

The time depends on the number and type of transfers. Assuming at least one word is moved a BLT takes .97 (1.08) μs plus 2.26 (2.48) μs per transfer from fast memory to core and 2.61 (2.83) μs per transfer from core to fast memory or from one core location to another.

BLT   **Block Transfer**

2 5 1		A	I	X	Y
0		8 9	12 13 14	17 18	35
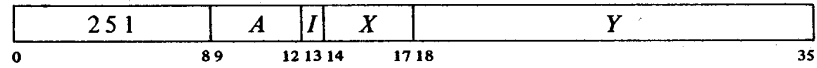
Beginning at the location addressed by AC left, move words to another area of memory beginning at the location addressed by AC right. Continue until a word is moved to location E. The total number of words in the block is thus $E - AC_R + 1$.

CAUTION

Priority interrupts are allowed during the execution of this instruction, following the processing of each word. If an interrupt occurs, the BLT stores the source and destination addresses for the next word in AC, so when the processor restarts upon the return to the interrupted program, it actually resumes at the correct point within the BLT. Therefore, unless the interrupt system is inactive, A and X must not address the same register as this would produce a different effective address calculation upon resumption should an interrupt occur; and the program must not attempt to load an accumulator addressed either by A or X unless it is the final location being loaded. Furthermore, the program cannot assume that AC is the same after the BLT as it was before.

EXAMPLES. This pair of instructions loads the accumulators from memory locations 2000–2017.

    HRLZI   17,2000   ;Put 2000 000000 in AC 17
    BLT   17,17

But to transfer the block in the opposite direction requires that one accumulator first be made available to the BLT:

    MOVEM   17,2017   ;Move AC 17 to 2017 in memory
    MOVEI   17,2000   ;Move the number 2000 to AC 17
    BLT   17,2016

If at the time the accumulators were loaded the program had placed in location 2017 the control word necessary for storing them back in the same block (2000), the three instructions above could be replaced by

    EXCH   17,2017
    BLT   17,2016

Move Instructions

Each of these instructions moves a single word, which may be changed in the process (*eg* its two halves may be swapped). There are four instructions,

each with four modes that determine the source and destination of the word moved.

Mode	Suffix	Source	Destination
Basic	.	E	AC
Immediate	I	The word 0,E	AC
Memory	M	AC	E
Self	S	E	E, but also AC if A is nonzero

Keeping instructions and operands in different memories saves .47 (.36) μs in memory mode, .20 (.09) μs in self mode.

When E addresses a fast memory location, a move instruction takes .34 μs less in basic mode, .46 (.35) μs less in memory mode, .54 (.43) μs less in self mode.

MOVE   Move

200	M	A	I	X	Y

0   6 7   8 9   12 13 14   17 18   35
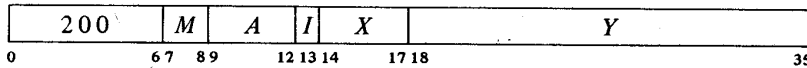
Move one word from the source to the destination specified by M. The source is unaffected, the original contents of the destination are lost.

MOVE	Move	200	2.21 (2.43) μs ▲
MOVEI	Move Immediate	201	1.36 (1.47) μs
MOVEM	Move to Memory	202	2.47 (2.58) μs ▲
MOVES	Move to Self	203	2.76 (2.87) μs

MOVEI loads the word 0,E into AC and is thus equivalent to HRRZI. If A is zero, MOVES is a no-op; otherwise it is equivalent to MOVE.

MOVS   Move Swapped

204	M	A	I	X	Y

0   6 7   8 9   12 13 14   17 18   35
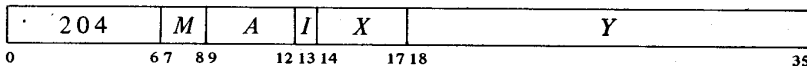
Interchange the left and right halves of the word from the source specified by M and move it to the specified destination. The source is unaffected, the original contents of the destination are lost.

MOVS	Move Swapped	204	2.21 (2.43) μs ▲
MOVSI	Move Swapped Immediate	205	1.36 (1.47) μs
MOVSM	Move Swapped to Memory	206	2.47 (2.58) μs ▲
MOVSS	Move Swapped to Self	207	2.76 (2.87) μs

Swapping halves in immediate mode loads the word E,0 into AC. MOVSI is thus equivalent to HRLZI.

MOVN   Move Negative

210	M	A	I	X	Y

0   6 7   8 9   12 13 14   17 18   35
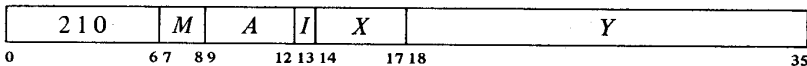
Negate the word from the source specified by M and move it to the specified destination. If the source word is fixed point -2^{35} (400000 000000) set the

Overflow and Carry 1 flags. (Negating the equivalent floating point -1×2^{127} sets the flags, but this is not a normalized number.) If the source word is zero, set Carry 0 and Carry 1. The source is unaffected, the original contents of the destination are lost.

MOVNI loads AC with the negative of the word $0,E$ and can set no flags.

▲ MOVN	Move Negative		210	2.39 (2.61) μs
MOVNI	Move Negative Immediate		211	1.54 (1.65) μs
▲ MOVNM	Move Negative to Memory		212	2.65 (2.76) μs
MOVNS	Move Negative to Self		213	2.94 (3.05) μs

MOVM   **Move Magnitude**

2 1 4	M	A	I	X	Y
0	6 7	8 9	12 13 14	17 18	35
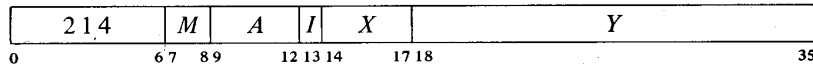
Take the magnitude of the word contained in the source specified by M and move it to the specified destination. If the source word is fixed point -2^{35} (400000 000000) set the Overflow and Carry 1 flags. (Negating the equivalent floating point -1×2^{127} sets the flags, but this is not a normalized number.) The source is unaffected, the original contents of the destination are lost.

The word $0,E$ is equivalent to its magnitude, so MOVMI is equivalent to MOVEI.

▲ MOVM	Move Magnitude		214	2.39 (2.61) μs
MOVMI	Move Magnitude Immediate		215	1.54 (1.65) μs
▲ MOVMM	Move Magnitude to Memory		216	2.65 (2.76) μs
MOVMS	Move Magnitude to Self		217	2.94 (3.05) μs

An example at the end of the preceding section demonstrates the use of a pair of immediate-mode half word transfers to load an address and a control count into an accumulator. The same result can be attained by a single move instruction. This saves time but still requires two locations. *Eg* if the number 200 001400 is stored in location M, the instruction

MOVE   AC,M

loads 200 into AC left and 1400 into AC right. If the same word, or its negative, or with its halves swapped, must be loaded on several occasions, then both time and space can be saved as each transfer requires only a single move instruction that references M.

Pushdown List

These two instructions insert and remove full words in a pushdown list. The address of the top item in the list is kept in the right half of a pointer in AC, and the program can keep a control count in the left half. There are also

two subroutine-calling instructions that utilize a pushdown list of jump addresses [§2.9].

PUSH **Push Down**   3.85 (4.07) μs

261	A	I	X	Y
0	8 9	12 13 14	17 18	35
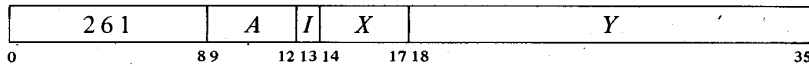
Add 1 000001$_8$ to AC to increment both halves by one, then move the contents of location E to the location now addressed by AC right. If the addition causes the count in AC left to reach zero, set the Pushdown Overflow flag. The contents of E are unaffected, the original contents of the location added to the list are lost.

Keeping instructions and the pushdown list in different memories saves .47 (.36) μs.

When the word added to the list is from fast memory, PUSH takes .34 μs less than the time given.

POP **Pop Up**   3.93 (4.15) μs

262	A	I	X	Y
0	8 9	12 13 14	17 18	35
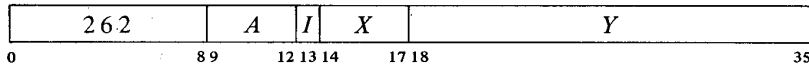
Move the contents of the location addressed by AC right to location E, then subtract 1 000001$_8$ from AC to decrement both halves by one. If the subtraction causes the count in AC left to reach −1, set the Pushdown Overflow flag. The original contents of E are lost.

Because of the order in which the operands are stored, the instruction POP AC,AC would load the contents of the location addressed by AC right into AC on top of the pushdown count, destroying it.

When the word taken from the list is placed in fast memory, POP takes .46 (.35) μs less than the time given.

The incrementing and decrementing of both halves of AC simultaneously is effected by adding and subtracting 1 000001$_8$. Hence a count of −2 in AC left is increased to zero if $2^{18} - 1$ is incremented in AC right, and conversely, 1 in AC left is decreased to −1 if zero is decremented in AC right.

A pushdown list is simply a set of consecutive memory locations from which words are read in the order opposite that in which they are written. In more general terms, it is any list in which the only item that can be removed at any given time is the last item in the list. This is usually referred to as "first in, last out" or "last in, first out". Suppose locations $a, b, c, ...$ are set aside for a pushdown list. We can deposit data in a, b, c, d, then read d, then write in d and e, then read e, d, c, etc.

Note that by using the Pushdown Overflow flag and a control count in AC left, the programmer can set a limit to the size of the list by starting the count negative, or he can prevent the program from extracting more words than there are in the list by starting the count at zero, but he cannot do both at once.

Pushdown storage is very convenient for a program that can use data stored in this manner as the pointer is initialized only once and only one accumulator is required for the most complex pushdown operations. To initialize a pointer P for a list to be kept in a block of memory beginning at BLIST and to contain at most N items, the following suffices.

              MOVSI   P,$-N$
              HRRI   P,BLIST-1

Of course the programmer must define BLIST elsewhere and set aside locations BLIST to BLIST $+ N - 1$.   Using MACRO to full advantage one could instead give

              MOVE   P,[IOWD N,BLIST]

where the pseudoinstruction

              ʼ   IOWD J, K

is replaced by a word containing $-J$ in the left half and $K - 1$ in the right. Elsewhere there would appear

BLIST: BLOCK N

which defines BLIST as the current contents of the location counter and sets aside the N locations beginning at that point.

In the PDP-10 the pushdown list is kept in a random access core memory, so the restrictions on order of entry and removal of items actually apply only to the standard addressing by the pointer in pushdown instructions — other addressing methods can reference any item at any time. The most convenient way to do this is to use the right half of the pointer as an index register. To move the last entry to accumulator AC we need simply give

              MOVE   AC,(P)

Of course this does not shorten the list — the word moved remains the last item in it.

One usually regards an index register as supplying an additive factor for a basic address contained in an instruction word, but the index register can supply the basic address and the instruction the additive factor. Thus we can retrieve the next to last item by giving

              MOVE   AC,-1(P)

and so forth. Similarly

              PUSH   P,-3(P)

adds the third to last item to the end of the list;

              POP   P,-2(P)

removes the last item and inserts it in place of the next to last item in the shortened list.

2.3 BYTE MANIPULATION

This set of five instructions allows the programmer to pack or unpack bytes of any length anywhere within a word. Movement of a byte is always between AC and a memory location: a deposit instruction takes a byte from the right end of AC and inserts it at any desired position in the memory location; a load instruction takes a byte from any position in the memory location and places it right-justified in AC.

The byte manipulation instructions have the standard memory reference format, but the effective address E is used to retrieve a pointer, which is used in turn to locate the byte or the place that will receive it. The pointer has the format

P	S		I	X	Y

0   5 6   11 12 13 14   17 18   35
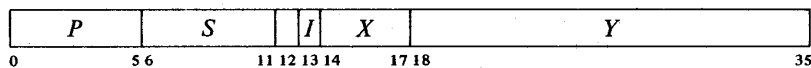
where S is the size of the byte as a number of bits, and P is its position as the number of bits remaining at the right of the byte in the word (eg if P is 3 the rightmost bit of the byte is bit 32 of the word). The rest of the pointer is interpreted in the same way as in an instruction: I, X and Y are used to calculate the address of the location that is the source or destination of the byte. Thus the pointer aims at a word whose format is

	S BITS	P BITS

0   35-P-S+1   35-P 35-P+1   35
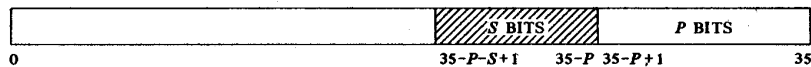
where the shaded area is the byte.

To facilitate processing a series of bytes, several of the byte instructions increment the pointer, ie modify it so that it points to the next byte position in a set of memory locations. Bytes are processed from left to right in a word, so incrementing merely replaces the current value of P by $P - S$, unless there is insufficient space in the present location for another byte of the specified size ($P - S < 0$). In this case Y is increased by one to point to the next consecutive location, and P is set to $36 - S$ to point to the first byte at the left in the new location.

CAUTION

Do not allow Y to reach maximum value. The whole pointer is incremented, so if Y is $2^{18} - 1$ it becomes zero and X is also incremented.

• The address calculation for the pointer uses the original X, but if a priority interrupt should occur before the calculation is complete, the incremented X is used when the instruction is repeated.

Among these five instructions one simply increments the pointer, the others load or deposit a byte with or without incrementing. Brackets enclose the additional time required when incrementing overflows the word boundary.

Keeping the pointer in fast memory saves .34 μs. Taking bytes from a fast memory location saves another .34 μs.

LDB **Load Byte** $4.02\ (4.35) + .15(P + S)\ [+.26]\ \mu s$
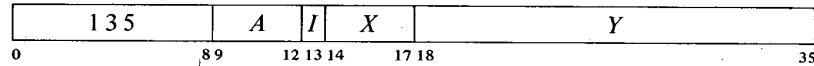
1 3 5		A	I	X	Y
0	8 9	12 13 14	17 18		35

Retrieve a byte of S bits from the location and position specified by the pointer contained in location E, load it into the right end of AC, and clear the remaining AC bits. The location containing the byte is unaffected, the original contents of AC are lost.

Keeping the pointer in fast memory saves .34 μs. Keeping instructions and the packing area in different memories saves .20 (.09) μs. Packing bytes in fast memory saves .54 (.43) μs.

DPB **Deposit Byte** $4.87\ (5.20) + .15(P + S)\ [+.26]\ \mu s$
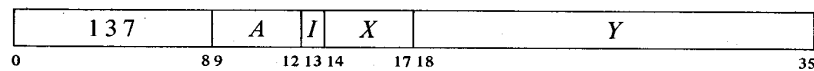
1 3 7	A	I	X	Y
0	8 9	12 13 14	17 18	35

Deposit the right S bits of AC into the location and position specified by the pointer contained in location E. The original contents of the bits that receive the byte are lost, AC and the remaining bits of the deposit location are unaffected.

Keeping the pointer in fast memory saves .54 (.43) μs; keeping it in a different memory from the instruction saves .20 (.09) μs

The *A* portion of this instruction is ignored.

IBP **Increment Byte Pointer** $2.87\ (2.98)\ [+.26]\ \mu s$

1 3 3	A	I	X	Y
0	8 9	12 13 14	17 18	35

Increment the byte pointer in location E as explained above.

Keeping the pointer in fast memory saves .34 μs. Taking bytes from a fast memory location saves another .34 μs.
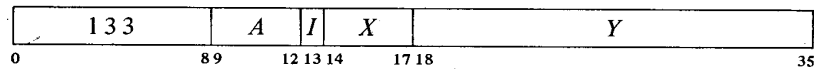
ILDB **Increment Pointer and Load Byte**

$4.24\ (4.57) + .15(P + S)\ [+.26]\ \mu s$

1 3 4	A	I	X	Y
0	8 9	12 13 14	17 18	35
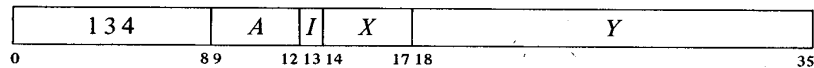
Increment the byte pointer in location E as explained above. Then retrieve a byte of S bits from the location and position specified by the newly incremented pointer, load it into the right end of AC, and clear the remaining AC bits. The location containing the byte is unaffected, the original contents of AC are lost.

Keeping the pointer in fast memory saves .34 μs. Keeping instructions and the packing area in different memories saves .20 (.09) μs. Packing bytes in fast memory saves .54 (.43) μs.

IDPB **Increment Pointer and Deposit Byte**

$5.29\ (5.51) + .15(P + S)\ [+.26]\ \mu s$

1 3 6	A	I	X	Y
0	8 9	12 13 14	17 18	35
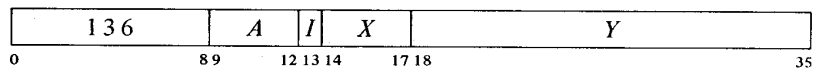
Increment the byte pointer in location E as explained above. Then deposit

the right S bits of AC into the location and position specified by the newly incremented pointer. The original contents of the bits that receive the byte are lost, AC and the remaining bits of the deposit location are unaffected.

Note that in the pair of instructions that both increment the pointer and process a byte, it is the *modified* pointer that determines the byte location and position. Hence to unpack bytes from a block of memory, the program should set up the pointer to point to a byte just *before* the first desired, and then load them with a loop containing an ILDB. If the first byte is at the left end of a word, this is most easily done by initializing the pointer with a P of 36 (44_8). Incrementing then replaces the 36 with $36 - S$ to point to the first byte. At any time that the program might inspect the pointer during execution of a series of ILDBs or IDPBs, it points to the last byte processed (this may not be true when the pointer is tested from an interrupt routine [§2.13]).

Special Considerations. If S is greater than P and also greater than 36, incrementing produces a new P equal to $100 - S$ rather than $36 - S$. For $S > 36$ the byte is at most the entire word; for $P \geq 36$ no byte is processed (loading merely clears AC). If both P and S are less than 36 but $P + S > 36$, a byte of size $36 - P$ is loaded from position P, or the right $36 - P$ bits of the byte are deposited in position P.

2.4   LOGIC

For logical operations the PDP-10 has instructions for shifting and rotating as well as for performing the complete set of sixteen Boolean functions of two variables (including those in which the result depends on only one or neither variable). The Boolean functions operate bitwise on full words, so each instruction actually performs thirty-six logical operations simultaneously. Thus in the AND function of two words, each bit of the result is the AND of the corresponding bits of the operands. The table on page 2-23 lists the bit configurations that result from the various operand configurations for all instructions.

Each Boolean instruction has four modes that determine the source of the non-AC operand, if any, and the destination of the result.

Mode	Suffix	Source of non-AC operand	Destination of result
Basic		E	AC
Immediate	I	The word 0, E	AC
Memory	M	E	E
Both	B	E	AC and E

Keeping instructions and operands in different memories saves .47 (.36) μs in memory and both modes in the first four of these instructions (those that have no operand or only an AC operand), .20 (.09) μs in memory and both modes in the remaining twelve (those that have a memory or immediate operand).

A Boolean instruction in which E addresses a fast memory location takes .46 (.35) μs less in memory or both mode if it has no operand or only an AC operand. If it has a memory operand, it takes .34 μs less in basic mode, .54 (.43) μs less in memory or both mode.

For an instruction without an operand (one that merely clears a location or sets it to all 1s) the modes differ only in the destination of the result, so basic and immediate modes are equivalent. The same is true also of an instruction that uses only an AC operand. When specified by the mode, the result goes to the accumulator addressed by A, even when there is no AC operand.

SETZ   Set to Zeros

400	M	A	I	X	Y
0	6 7	8 9	12 13 14	17 18	35
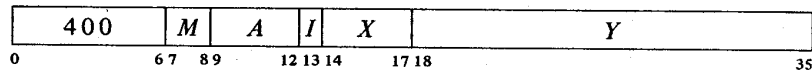
Change the contents of the destination specified by M to all 0s.

SETZ and SETZI are equivalent (both merely clear AC). MACRO also recognizes CLEAR, CLEARI, CLEARM and CLEARB as equivalent to the set-to-zeros mnemonics.

SETZ	Set to Zeros	400	1.36 (1.47) μs
SETZI	Set to Zeros Immediate	401	1.36 (1.47) μs
SETZM	Set to Zeros Memory	402	2.33 (2.44) μs
SETZB	Set to Zeros Both	403	2.33 (2.44) μs

SETO   Set to Ones

474	M	A	I	X	Y
0	6 7	8 9	12 13 14	17 18	35
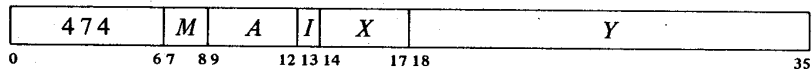
Change the contents of the destination specified by M to all 1s.

SETO and SETOI are equivalent.

SETO	Set to Ones	474	1.36 (1.47) μs
SETOI	Set to Ones Immediate	475	1.36 (1.47) μs
SETOM	Set to Ones Memory	476	2.33 (2.44) μs
SETOB	Set to Ones Both	477	2.33 (2.44) μs

SETA   Set to AC

424	M	A	I	X	Y
0	6 7	8 9	12 13 14	17 18	35
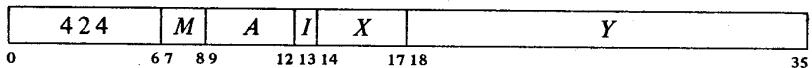
Make the contents of the destination specified by M equal to AC.

SETA and SETAI are no-ops. SETAM and SETAB are both equivalent to MOVEM (all move AC to location E).

SETA	Set to AC	424	1.50 (1.61) μs
SETAI	Set to AC Immediate	425	1.50 (1.61) μs
SETAM	Set to AC Memory	426	2.47 (2.58) μs
SETAB	Set to AC Both	427	2.47 (2.58) μs

SETCA   Set to Complement of AC

450	M	A	I	X	Y

0   6 7 8 9   12 13 14   17 18   35
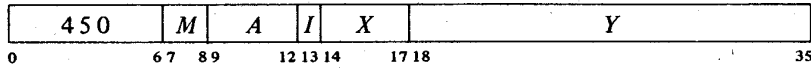
Change the contents of the destination specified by M to the complement of AC.

SETCA	Set to Complement of AC	450
		1.50 (1.61) μs
SETCAI	Set to Complement of AC Immediate	451
		1.50 (1.61) μs
SETCAM	Set to Complement of AC Memory	452
		2.47 (2.58) μs
SETCAB	Set to Complement of AC Both	453
		2.47 (2.58) μs

SETCA and SETCAI are equivalent (both complement AC).

SETM   Set to Memory

414	M	A	I	X	Y

0   6 7 8 9   12 13 14   17 18   35
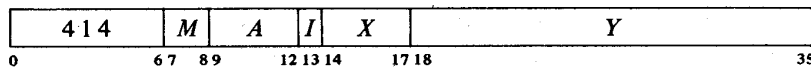
Make the contents of the destination specified by M equal to the specified operand.

SETM	Set to Memory	414	2.21 (2.43) μs
SETMI	Set to Memory Immediate	415	1.36 (1.47) μs
SETMM	Set to Memory Memory	416	2.76 (2.87) μs
SETMB	Set to Memory Both	417	2.76 (2.87) μs

SETM and SETMB are equivalent to MOVE. SETMI moves the word 0,E to AC and is thus equivalent to MOVEI. SETMM is a no-op that references memory.

SETCM   Set to Complement of Memory
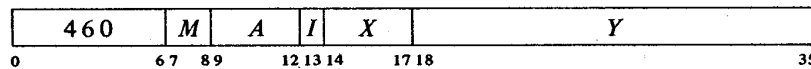
460	M	A	I	X	Y

0   6 7 8 9   12 13 14   17 18   35

Change the contents of the destination specified by M to the complement of the specified operand.

SETCM	Set to Complement of Memory	460
		2.21 (2.43) μs
SETCMI	Set to Complement of Memory Immediate	461
		1.36 (1.47) μs
SETCMM	Set to Complement of Memory Memory	462
		2.76 (2.87) μs
SETCMB	Set to Complement of Memory Both	463
		2.76 (2.87) μs

SETCMI moves the complement of the word 0,E to AC. SETCMM complements location E.

AND   **And with AC**
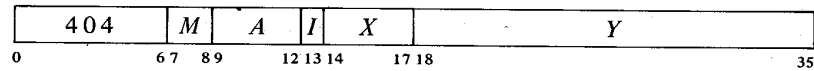
4 0 4	M	A	I	X	Y
0	6 7	8 9	12 13 14	17 18	35

Change the contents of the destination specified by M to the AND function of the specified operand and AC.

AND	And	404	2.35 (2.57) μs
ANDI	And Immediate	405	1.50 (1.61) μs
ANDM	And to Memory	406	2.90 (3.01) μs
ANDB	And to Both	407	2.90 (3.01) μs

ANDCA   **And with Complement of AC**
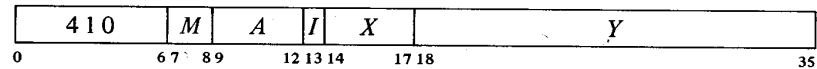
4 1 0	M	A	I	X	Y
0	6 7	8 9	12 13 14	17 18	35

Change the contents of the destination specified by M to the AND function of the specified operand and the complement of AC.

ANDCA	And with Complement of AC	410
		2.70 (2.92) μs
ANDCAI	And with Complement of AC Immediate	411
		1.85 (1.96) μs
ANDCAM	And with Complement of AC to Memory	412
		3.52 (3.63) μs
ANDCAB	And with Complement of AC to Both	413
		3.52 (3.63) μs

ANDCM   **And Complement of Memory with AC**
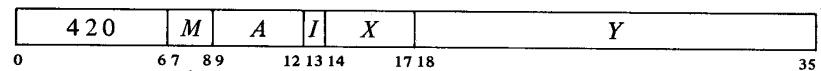
4 2 0	M	A	I	X	Y
0	6 7	8 9	12 13 14	17 18	35

Change the contents of the destination specified by M to the AND function of the complement of the specified operand and AC.

ANDCM	And Complement of Memory	420
		2.35 (2.57) μs
ANDCMI	And Complement of Memory Immediate	421
		1.50 (1.61) μs
ANDCMM	And Complement of Memory to Memory	422
		2.90 (3.01) μs
ANDCMB	And Complement of Memory to Both	423
		2.90 (3.01) μs
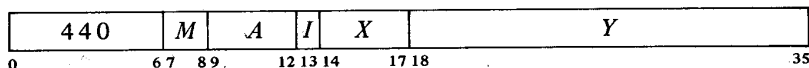
ANDCB   **And Complements of Both**

4 4 0	M	A	I	X	Y

0   6 7   8 9.   12 13 14   17 18   35

Change the contents of the destination specified by M to the AND function of the complements of both the specified operand and AC. The result is the NOR function of the operands.

ANDCB	And Complements of Both	440
		2.70 (2.92) μs
ANDCBI	And Complements of Both Immediate	441
		1.85 (1.96) μs
ANDCBM	And Complements of Both to Memory	442
		3.52 (3.63) μs
ANDCBB	And Complements of Both to Both	443
		3.52 (3.63) μs

IOR   **Inclusive Or with AC**
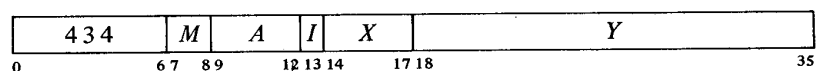
4 3 4	M	A	I	X	Y

0   6 7   8 9   12 13 14   17 18   35

Change the contents of the destination specified by M to the inclusive OR function of the specified operand and AC.

IOR	Inclusive Or	434	2.35 (2.57) μs
IORI	Inclusive Or Immediate	435	1.50 (1.61) μs
IORM	Inclusive Or to Memory	436	2.90 (3.01) μs
IORB	Inclusive Or to Both	437	2.90 (3.01) μs

MACRO also recognizes OR, ORI, ORM and ORB as equivalent to the inclusive OR mnemonics.

ORCA   **Inclusive Or with Complement of AC**
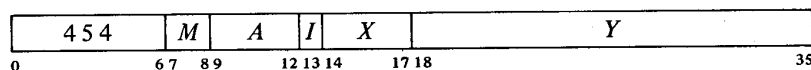
4 5 4	M	A	I	X	Y

0   6 7   8 9   12 13 14   17 18   35

Change the contents of the destination specified by M to the inclusive OR function of the specified operand and the complement of AC.

ORCA	Or with Complement of AC	454
		2.70 (2.92) μs
ORCAI	Or with Complement of AC Immediate	455
		1.85 (1.96) μs
ORCAM	Or with Complement of AC to Memory	456
		3.52 (3.63) μs
ORCAB	Or with Complement of AC to Both	457
		3.52 (3.63) μs

ORCM Inclusive Or Complement of Memory with AC

4 6 4	M	A	I	X	Y
0	6 7	8 9	12 13 14	17 18	35

Change the contents of the destination specified by M to the inclusive OR function of the complement of the specified operand and AC.

ORCM	Or Complement of Memory	464
		2.35 (2.57) μs
ORCMI	Or Complement of Memory Immediate	465
		1.50 (1.61) μs
ORCMM	Or Complement of Memory to Memory	466
		2.90 (3.01) μs
ORCMB	Or Complement of Memory to Both	467
		2.90 (3.01) μs

ORCB Inclusive Or Complements of Both

4 7 0	M	A	I	X	Y
0	6 7	8 9	12 13 14	17 18	35

Change the contents of the destination specified by M to the inclusive OR function of the complements of both the specified operand and AC. The result is the NAND function of the operands.

ORCB	Or Complements of Both	470
		2.70 (2.92) μs
ORCBI	Or Complements of Both Immediate	471
		1.85 (1.96) μs
ORCBM	Or Complements of Both to Memory	472
		3.52 (3.63) μs
ORCBB	Or Complements of Both to Both	473
		3.52 (3.63) μs

XOR Exclusive Or with AC

4 3 0	M	A	I	X	Y
0	6 7	8 9	12 13 14	17 18	35

Change the contents of the destination specified by M to the exclusive OR function of the specified operand and AC.

XOR	Exclusive Or	430	2.35 (2.57) μs
XORI	Exclusive Or Immediate	431	1.50 (1.61) μs
XORM	Exclusive Or to Memory	432	2.90 (3.01) μs
XORB	Exclusive Or to Both	433	2.90 (3.01) μs

The original contents of the destination can be recovered except in XORB, where both operands are replaced by the result. In the other three modes the replaced operand is restored by repeating the instruction in the same mode, ie by taking the exclusive OR of the remaining operand and the result.

EQV   **Equivalence with AC**

4 4 4	M	A	I	X	Y
0	6 7	8 9	12 13 14	17 18	35
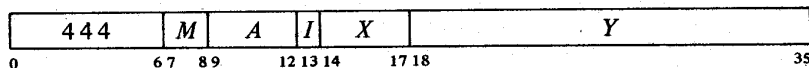
Change the contents of the destination specified by M to the complement of the exclusive OR function of the specified operand and AC (the result has 1s wherever the corresponding bits of the operands are the same).

EQV	Equivalence	444	2.35 (2.57) μs
EQVI	Equivalence Immediate	445	1.50 (1.61) μs
EQVM	Equivalence to Memory	446	2.90 (3.01) μs
EQVB	Equivalence to Both	447	2.90 (3.01) μs

The original contents of the destination can be recovered except in EQVB, where both operands are replaced by the result. In the other three modes the replaced operand is restored by repeating the instruction in the same mode, *ie* by taking the equivalence function of the remaining operand and the result.

For the four possible bit configurations of the two operands, the above sixteen instructions produce the following results. In each case the result as listed is equal to bits 3–6 of the instruction word.

	AC	0	1	0	1
Mode Specified Operand		0	0	1	1
SETZ		0	0	0	0
AND		0	0	0	1
ANDCA		0	0	1	0
SETM		0	0	1	1
ANDCM		0	1	0	0
SETA		0	1	0	1
XOR		0	1	1	0
IOR		0	1	1	1
ANDCB		1	0	0	0
EQV		1	0	0	1
SETCA		1	0	1	0
ORCA		1	0	1	1
SETCM		1	1	0	0
ORCM		1	1	0	1
ORCB		1	1	1	0
SETO		1	1	1	1

Shift and Rotate

The remaining logical instructions shift or rotate right or left the contents of AC or the contents of two accumulators, A and $A+1$ (mod 20_8), concatenated into a 72-bit register with A on the left. The illustration below shows the movement of information these instructions produce in the accu-

LSH

LSHC

ROT

ROTC

ASH

ASHC
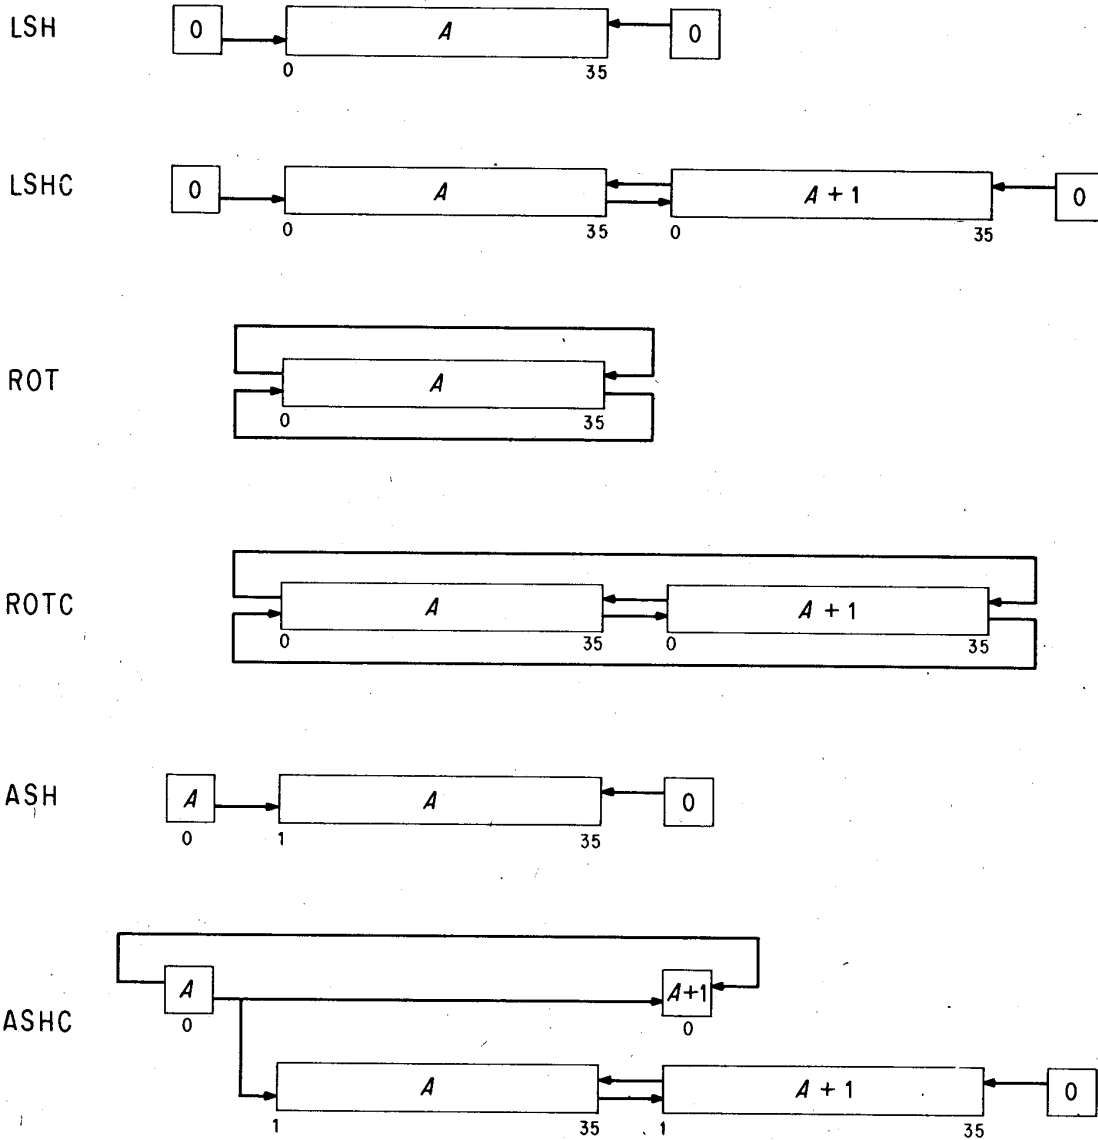
ACCUMULATOR BIT FLOW IN SHIFT AND ROTATE INSTRUCTIONS

mulators. In a (logical) shift the contents of a register are moved bit-to-bit with 0s brought in at the end being vacated; information shifted out at the other end is lost. [*For a discussion of arithmetic shifting see* §2.5.] In rotation the contents are moved cyclically such that information rotated out at one end is put in at the other.

The number of places moved is specified by the result of the effective address calculation taken as a signed number (in twos complement notation) modulo 2^8 in magnitude. In other words the effective shift E is the number composed of bit 18 (which is the sign) and bits 28–35 of the calculation result. Hence the programmer may specify the shift directly in the instruction (perhaps indexed) or give an indirect address to be used in calculating the shift. A positive E produces motion to the left, a negative E to the right; maximum movement is 255 places.

LSH **Logical Shift** Left: $1.62 (1.73) + .15|E| \mu s$
Right: $1.46 (1.57) + .15|E| \mu s$

242	A	I	X	Y

0   8 9   12 13 14   17 18   35
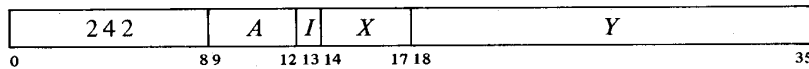
Shift AC the number of places specified by E. If E is positive, shift left bringing 0s into bit 35; data shifted out of bit 0 is lost. If E is negative, shift right bringing 0s into bit 0; data shifted out of bit 35 is lost.

LSHC **Logical Shift Combined** Left: $2.00 (2.11) + .15|E| \mu s$
Right: $1.84 (1.95) + .15|E| \mu s$

246	A	I	X	Y

0   8 9   12 13 14   17 18   35
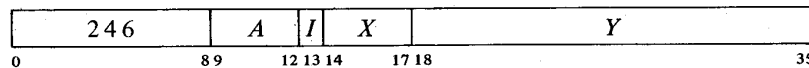
Concatenate accumulators A and $A+1$ with A on the left, and shift the 72-bit combination the number of places specified by E. If E is positive, shift left bringing 0s into bit 71 (bit 35 of AC $A+1$); bit 36 is shifted into bit 35; data shifted out of bit 0 is lost. If E is negative, shift right bringing 0s into bit 0; bit 35 is shifted into bit 36; data shifted out of bit 71 is lost.

ROT **Rotate** Left: $1.62 (1.73) + .15|E| \mu s$
Right: $1.46 (1.57) + .15|E| \mu s$

241	A	I	X	Y

0   8 9   12 13 14   17 18   35
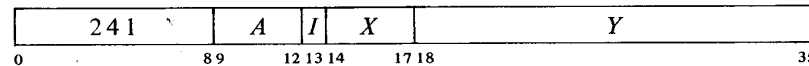
Rotate AC the number of places specified by E. If E is positive, rotate left; bit 0 is rotated into bit 35. If E is negative, rotate right; bit 35 is rotated into bit 0.

ROTC   **Rotate Combined**   Left: 2.00 (2.11) + .15|E| μs
                                                     Right: 1.84 (1.95) + .15|E| μs

2 4 5	A	I	X	Y
0	'8 9	12 13 14	17 18	35

Concatenate accumulators A and $A+1$ with A on the left, and rotate the 72-bit combination the number of places specified by E. If E is positive, rotate left; bit 0 is rotated into bit 71 (bit 35 of AC $A+1$) and bit 36 into bit 35. If E is negative, rotate right; bit 35 is rotated into bit 36 and bit 71 into bit 0.

2.5   FIXED POINT ARITHMETIC

For fixed point arithmetic the PDP–10 has instructions for arithmetic shifting (which is essentially multiplication by a power of 2) as well as for performing addition, subtraction, multiplication and division of numbers in fixed point format [§1.1]. In such numbers the position of the binary point is arbitrary (the programmer may adopt any point convention). The add and subtract instructions involve only single length numbers, whereas multiply supplies a double length product, and divide uses a double length dividend. The high and low order words respectively of a double length fixed point number are in accumulators A and $A+1$ (mod 20_8), where the magnitude is the 70-bit string in bits 1–35 of the two words and the signs of the two are identical. There are also integer multiply and divide instructions that involve only single length numbers and are especially suited for handling smaller integers, particularly those of eighteen bits or less such as addresses (of course they can be used for small fractions as well provided the programmer keeps track of the binary point). For convenience in the following, all operands are assumed to be integers (binary point at the right).

The processor has four flags, Overflow, Carry 0, Carry 1 and No Divide, that indicate when the magnitude of a number is or would be larger than can be accommodated. Carry 0 and Carry 1 actually detect carries out of bits 0 and 1 in certain instructions that employ fixed point arithmetic operations: the add and subtract instructions treated here, the move instructions that produce the negative or magnitude of the word moved [§2.2], and the arithmetic test instructions that increment or decrement the test word [§2.7]. In these instructions an incorrect result is indicated — and the Overflow flag set — if the carries are different, ie if there is a carry into the sign but not out of it, or vice versa. The Overflow flag is also set by No Divide being set, which means the processor has failed to perform a division because the magnitude of the dividend is greater than or equal to that of the divisor, or in integer divide, simply that the divisor is zero. In other overflow cases only Overflow itself is set: these include too large a product in multiplication, and loss of significant bits in left arithmetic shifting.

These flags can be read and controlled by certain program control instructions [§2.9], and Overflow is available as a processor condition (via in-out

> Overflow is determined directly from the carries, not from the carry flags, as their states may reflect events in previous instructions.

instructions [§2.14]) that can request a priority interrupt if enabled. The conditions detected can only set the flags and the hardware does not clear them, so the program must clear them before an instruction if they are to give meaningful information about the instruction afterward. However, the program can check the flags following a series of instructions to determine whether the entire series was free of the types of error detected.

All but the shift instructions have four modes that determine the source of the non-AC operand and the destination of the result.

Mode	Suffix	Source of non-AC operand	Destination of result
Basic		E	AC
Immediate	I	The word $0,E$	AC
Memory	M	E	E
Both	B	E	AC and E

ADD   Add

2 7 0	M	A	I	X	Y
0	6 7	8 9	12 13 14	17 18	35

Add the operand specified by M to AC and place the result in the specified destination. If the sum is $\geqslant 2^{35}$ set Overflow and Carry 1; the result stored has a minus sign but a magnitude in positive form equal to the sum less 2^{35}. If the sum is $< -2^{35}$ set Overflow and Carry 0; the result stored has a plus sign but a magnitude in negative form equal to the sum plus 2^{35}. Set both carry flags if both summands are negative, or their signs differ and their magnitudes are equal or the positive one is the greater in magnitude.

ADD	Add	270	2.53 (2.75) μs
ADDI	Add Immediate	271	1.68 (1.79) μs
ADDM	Add to Memory	272	3.08 (3.19) μs
ADDB	Add to Both	273	3.08 (3.19) μs

SUB   Subtract

2 7 4	M	A	I	X	Y
0	6 7	8 9	12 13 14	17 18	35

Subtract the operand specified by M from AC and place the result in the specified destination. If the difference is $\geqslant 2^{35}$ set Overflow and Carry 1; the result stored has a minus sign but a magnitude in positive form equal to the difference less 2^{35}. If the difference is $< -2^{35}$ set Overflow and Carry 0; the result stored has a plus sign but a magnitude in negative form equal to the difference plus 2^{35}. Set both carry flags if the signs of the operands are the same and AC is the greater or the two are equal, or the signs of the operands differ and AC is negative.

Besides indicating error types, the carry flags facilitate performing multiple precision arithmetic.

Keeping instructions and operands in different memories saves .20 (.09) μs in ADDM and ADDB.

When E addresses a fast memory location, ADD takes .34 μs less than the time given, ADDM and ADDB take .54 (.43) μs less.

SUB	Subtract	274	2.53 (2.75) μs
SUBI	Subtract Immediate	275	1.68 (1.79) μs
SUBM	Subtract to Memory	276	3.08 (3.19) μs
SUBB	Subtract to Both	277	3.08 (3.19) μs

Keeping instructions and operands in different memories saves .20 (.09) μs in SUBM and SUBB.

When E addresses a fast memory location. SUB takes .34 μs less than the time given. SUBM and SUBB take .54 (.43) μs less.

MUL   Multiply

224	M	A	I	X	Y
0	6 7	8 9	12 13 14	17 18	35
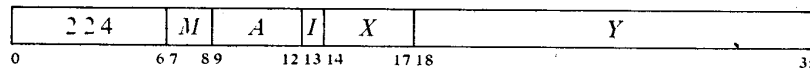
Multiply AC by the operand specified by M, and place the high order word of the double length result in the specified destination. If M specifies AC as a destination, place the low order word in accumulator $A+1$. If both operands are -2^{35} set Overflow: the double length result stored is -2^{70}.

MUL	Multiply	224	10.60 (10.82) μs
MULI	Multiply Immediate	225	8.58 (8.69) μs
MULM	Multiply to Memory	226	11.41 (11.63) μs
MULB	Multiply to Both	227	11.41 (11.63) μs

Keeping instructions and operands in different memories saves .47 (.36) μs in MULM. .31 (.20) μs in MULB.

When E addresses a fast memory location. MUL takes .34 μs less than the time given. MULM takes .80 (.69) μs less, and MULB takes .64 (.53) μs less.

Timing. The times given above are average. The algorithm modifies the running sum of partial products at each 1-0 or 0-1 transition scanning from one bit to the next in the multiplier, which is the operand specified by the mode: in other words the number of operations equals the number of pairs of adjacent bits that differ in the multiplier including the sign bit and taking the bit at the right of the LSB as 0 (an LSB of 1 is regarded as a transition). Minimum times with a zero multiplier are

MUL	8.26 (8.48) μs
MULI	7.41 (7.52) μs
MULM	9.07 (9.29) μs
MULB	9.07 (9.29) μs

These must be increased by .13 μs for each transition. The programmer can minimize the time by using as the multiplier the operand with fewer transitions.

IMUL   Integer Multiply

220	M	A	I	X	Y
0	6 7	8 9	12 13 14	17 18	35
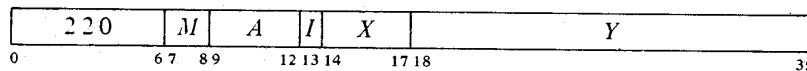
Multiply AC by the operand specified by M, and place the sign and the 35 low order magnitude bits of the product in the specified destination. Set Overflow if the product is $\geqslant 2^{35}$ or $< -2^{35}$ (*ie* if the high order word of the double length product is not null); the high order word is lost.

IMUL	Integer Multiply	220	9.59 (9.81) μs
IMULI	Integer Multiply Immediate	221	8.09 (8.20) μs
IMULM	Integer Multiply to Memory	222	10.56 (10.78) μs
IMULB	Integer Multiply to Both	223	10.56 (10.78) μs

Keeping instructions and operands in different memories saves .47 (.36) μs in IMULM and IMULB.

    When E addresses a fast memory location, IMUL takes .34 μs less than the time given, IMULM and IMULB take .80 (.69) μs less.

Timing. The times given above are average. Refer to the description of MUL for the timing effects of the multiplication algorithm. Minimum times with a zero multiplier are

IMUL	8.42 (8.64) μs
IMULI	7.57 (7.68) μs
IMULM	9.39 (9.61) μs
IMULB	9.39 (9.61) μs

These must be increased by .13 μs for each transition. The programmer can minimize the time by using as the multiplier the operand with fewer transitions.

DIV   **Divide**

234	M	A	I	X	Y
0	6 7	8 9	12 13 14	17 18	35
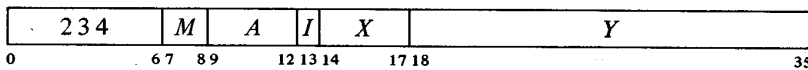
If the magnitude of the number in AC is greater than or equal to that of the operand specified by M, set Overflow and No Divide, and go immediately to the next instruction without affecting the original AC or memory operand in any way. Otherwise divide the double length number contained in accumulators A and $A+1$ by the specified operand, calculating a quotient of 35 magnitude bits including leading zeros. Place the unrounded quotient in the specified destination. If M specifies AC as a destination, place the remainder, with the same sign as the dividend, in accumulator $A+1$.

DIV	Divide	234	16.2 (16.4) μs
DIVI	Divide Immediate	235	15.4 (15.5) μs
DIVM	Divide to Memory	236	17.1 (17.3) μs
DIVB	Divide to Both	237	17.1 (17.3) μs

Keeping instructions and operands in different memories saves .5 (.4) μs in DIVM, .3 (.2) μs in DIVB.

    When E addresses a fast memory location, DIV takes .3 μs less than the time given, DIVM takes .8 (.7) μs less, and DIVB takes .6 (.5) μs less.

    If the division is not performed, only 2.5–3 μs are required.

IDIV   **Integer Divide**

230	M	A	I	X	Y
0	6 7	8 9	12 13 14	17 18	35

If the operand specified by M is zero, set Overflow and No Divide, and go immediately to the next instruction without affecting the original AC or memory operand in any way. Otherwise divide AC by the specified operand, calculating a quotient of 35 magnitude bits including leading zeros. Place

the unrounded quotient in the specified destination. If M specifies AC as the destination, place the remainder, with the same sign as the dividend, in accumulator $A+1$.

IDIV	Integer Divide	230	16.5 (16.7) μs
IDIVI	Integer Divide Immediate	231	15.7 (15.8) μs
IDIVM	Integer Divide to Memory	232	17.4 (17.6) μs
IDIVB	Integer Divide to Both	233	17.4 (17.6) μs

EXAMPLE. The integer multiply and divide instructions are very useful for computations on addresses or character codes, or performing any integral operations in which the result is small enough to be accommodated in a single register.

As an example suppose we wish to determine the parity of the 8-bit character *abcdefgh*, where the letters represent the bits of the character. Assuming the character is right-justified in AC, we first duplicate it twice to the left producing

abc def gha bcd efg hab cde fgh

where the bits (in positions 12–35) are grouped corresponding to the octal digits in the word. Anding this with

001 001 001 001 001 001 001 001

retains only the least significant bit in each 3-bit set, so we can represent the result by

cfadgbeh

where each letter represents an octal digit having the same value (0 or 1) as the bit originally represented by the same letter. Multiplying this by 11111111_8 generates the following partial products:

```
          c  f  a  d  g  b  e  h
       c  f  a  d  g  b  e  h
    c  f  a  d  g  b  e  h
   c  f  a  d  g  b  e  h
c  f  a  d  g  b  e  h
c  f  a  d  g  b  e  h
c  f  a  d  g  b  e  h
c  f  a  d  g  b  e  h
```

Since any digit is at most 1, there can be no carry out of any column with fewer than eight digits unless there is a carry into it. Hence the octal digit produced by summing the center column (the one containing all the bits of the character) is even or odd as the sum of the bits is even or odd. Thus its least significant bit (bit 14 of the low order word in the product) is the parity of the character, 0 if even, 1 if odd.

The above may seem a very complicated procedure to do something trivial, but it is effected by this quite simple sequence (with the character

right-justified in AC):

             IMULI   AC,200401
             AND   AC,ONES
             IMUL   AC,ONES
               .
               .
               .
ONES: 11111111

where the parity is indicated by AC bit 14. Of course, following the IMUL would be a test instruction to check the value of the bit.

Arithmetic Shifting

These two instructions produce an arithmetic shift right or left of the number in AC or the double length number in accumulators A and $A+1$. Shifting is the movement of the contents of a register bit-to-bit. The operation discussed here is similar to logical shifting [*see* §2.4 *and the illustration on page 2-24*], but in an arithmetic shift only the magnitude part is shifted — the sign is unaffected. In a double length number the 70-bit string made up of the magnitude parts of the two words is shifted, but the sign of the low order word is made equal to the sign of the high order word.

Null bits are brought in at the end being vacated: a left shift brings in 0s at the right, whereas a right shift brings in the equivalent of the sign bit at the left. In either case, information shifted out at the other end is lost. A single shift left is equivalent to multiplying the number by 2 (provided no bit of significance is shifted out); a shift right divides the number by 2.

The number of places shifted is specified by the result of the effective address calculation taken as a signed number (in twos complement notation) modulo 2^8 in magnitude. In other words the effective shift E is the number composed of bit 18 (which is the sign) and bits 28–35 of the calculation result. Hence the programmer may specify the shift directly in the instruction (perhaps indexed) or give an indirect address to be used in calculating the shift. A positive E produces motion to the left, a negative E to the right; E is thus the power of 2 by which the number is multiplied. Maximum movement is 255 places.

ASH   **Arithmetic Shift**   Left: 1.62 (1.73) + .15|E| μs
                                                         Right: 1.46 (1.57) + .15|E| μs

240	A	I	X	Y
0	8 9	12 13 14	17 18	35
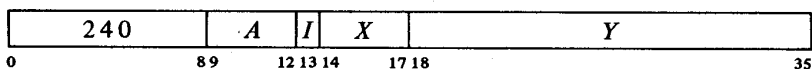
Shift AC arithmetically the number of places specified by E. Do not shift bit 0. If E is positive, shift left bringing 0s into bit 35; data shifted out of bit 1 is lost; set Overflow if any bit of significance is lost (a 1 in a positive number, a 0 in a negative one). If E is negative, shift right bringing 0s into bit 1 if AC is positive, 1s if negative; data shifted out of bit 35 is lost.

ASHC   **Arithmetic Shift Combined**   Left: 2.00 (2.11) + .15|E| µs
                                                Right: 1.84 (1.95) + .15|E| µs

244	A	I	X	Y

0   8 9   12 13 14   17 18   35
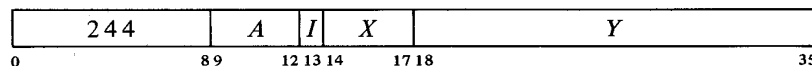
Concatenate the magnitude portions of accumulators A and $A+1$ with A on
the left, and shift the 70-bit combination in bits 1–35 and 37–71 the num-
ber of places specified by E. Do not shift AC bit 0, but make bit 0 of AC
$A+1$ equal to it if at least one shift occurs (*ie* if E is nonzero). If E is posi-
tive, shift left bringing 0s into bit 71 (bit 35 of AC $A+1$); bit 37 (bit 1 of AC
$A+1$) is shifted into bit 35; data shifted out of bit 1 is lost; set Overflow if
any bit of significance is lost (a 1 in a positive number, a 0 in a negative one).
If E is negative, shift right bringing 0s into bit 1 if AC is positive, 1s if nega-
tive; bit 35 is shifted into bit 37; data shifted out of bit 71 is lost.

2.6 FLOATING POINT ARITHMETIC

For floating point arithmetic the PDP–10 has instructions for scaling the
exponent (which is multiplication of the entire number by a power of 2)
and negating double length numbers as well as for performing addition, sub-
traction, multiplication and division of numbers in floating point format.
All instructions treated here interpret all operands as floating point numbers
in the format given in §1.1, and generate results in that format. The reader
is strongly advised to reread §1.1 if he does not remember the format in
detail.

For the four standard arithmetic operations the program can select wheth-
er or not the result shall be rounded. Rounding produces the greatest con-
sistent precision using only single length operands. Instructions without
rounding have a "long" mode, which supplies a two-word result for greater
precision; the other modes save time in one-word operations where rounding
is of no significance.

Actually the result is formed in a double length register in addition, sub-
traction and multiplication, wherein any bits of significance in the low order
part supply information for normalization, and then for rounding if re-
quested. Consider addition as an example. Before adding, the processor
right shifts the fractional part of the operand with the smaller exponent until
its bits correctly match the bits of the other operand in order of magnitude.
Thus the smaller operand could disappear entirely, having no effect on the
result ("result" shall always be taken to mean the information (one word or
two) stored by the instruction, regardless of the number of significant bits it
contains or even whether it is the correct answer). Long mode is likely to
retain information that would otherwise be lost, but in any given mode the
significance of the result depends on the relative values of the operands.
Even when both operands contain twenty-seven significant bits, a long addi-
tion may store two words that together contain only one significant bit. In
division the processor always calculates a one-word quotient that requires no

A subtraction involving two
like-signed numbers whose
exponents are equal and
whose fractions differ only in
the LSB gives a result con-
taining only one bit of signi-
ficance.

normalization if the original operands are normalized. An extra quotient bit is calculated for rounding when requested; long mode retains the remainder.

The processor has four flags, Overflow, Floating Overflow, Floating Underflow and No Divide, that indicate when the exponent is too large or too small to be accommodated or a division cannot be performed because of the relative values of dividend and divisor. Any of these circumstances sets Overflow and Floating Overflow. If only these two are set, the exponent of the answer is too large; if Floating Underflow is also set, the exponent is too small. No Divide being set means the processor failed to perform a division, an event that can be produced only by a zero divisor if all nonzero operands are normalized. These flags can be read and controlled by certain program control instructions [§2.9], and Overflow and Floating Overflow are available as processor conditions (via in-out instructions [§2.14]) that can request a priority interrupt if enabled. The conditions detected can only set the flags and the hardware does not clear them, so the program must clear them before a floating point instruction if they are to give meaningful information about the instruction afterward. However, the program can check the flags following a series of instructions to determine whether the entire series was free of the types of error detected.

The floating point hardware functions at its best if given operands that are either normalized or zero, and except in special situations the hardware normalizes a nonzero result. An operand with a zero fraction and a nonzero exponent can give wild answers in additive operations because of extreme loss of significance; eg adding $\frac{1}{2} \times 2^2$ and 0×2^{69} gives a zero result, as the first operand (having a smaller exponent) looks smaller to the processor and is shifted to oblivion. A number with a 1 in bit 0 and 0s in bits 9–35 is not simply an incorrect representation of zero, but rather an unnormalized "fraction" with value -1. This unnormalized number can produce an incorrect answer in any operation. Use of other unnormalized operands simply causes loss of significant bits, except in division where they can prevent its execution because they can satisfy a no-divide condition that is impossible for normalized numbers.

The processor normalizes the result by shifting the fraction and adjusting the exponent to compensate for the change in value. Each shift and accompanying exponent adjustment thus multiply the number both by 2 and by $\frac{1}{2}$ simultaneously, leaving its value unchanged.

Scaling

One floating point instruction is in a category by itself: it changes the exponent of a number without changing the significance of the fraction. In other words it multiplies the number by a power of 2, and is thus analogous to arithmetic shifting of fixed point numbers except that no information is lost, although the exponent can overflow or underflow. The amount added to the exponent is specified by the result of the effective address calculation taken as a signed number (in twos complement notation) modulo 2^8 in magnitude. In other words the effective scale factor E is the number composed of bit 18 (which is the sign) and bits 28–35 of the calculation result. Hence the programmer may specify the factor directly in the instruction (perhaps indexed) or give an indirect address to be used in calculating it. A positive E increases the exponent, a negative E decreases it; E is thus the power of 2 by which the number is multiplied. The scale factor lies in the range -256 to $+255$.

N is the number of left shifts needed to normalize the result.

This instruction can be used to float a fixed number with 27 or fewer significant bits. To float an integer contained within AC bits 9–35,

    FSC   AC,233

inserts the correct exponent to move the binary point from the right end to the left of bit 9 and then normalizes $(233_8 = 155_{10} = 128 + 27)$.

FSC   Floating Scale   2.75 (2.86) + .25N μs

1 3 2		A	I	X		Y	
0		8 9	12 13 14		17 18		35

If the fractional part of AC is zero, clear AC. Otherwise add the scale factor given by *E* to the exponent part of AC (thus multiplying AC by 2^E), normalize the resulting word bringing 0s into bit positions vacated at the right, and place the result back in AC.

NOTE

A negative *E* is represented in standard twos complement notation, but the hardware compensates for this when scaling the exponent.

If the exponent after normalization is > 127, set Overflow and Floating Overflow; the result stored has an exponent 256 less than the correct one. If < −128, set Overflow, Floating Overflow and Floating Underflow; the result stored has an exponent 256 greater than the correct one.

Operations with Rounding

In the hardware the rounding operation is actually somewhat more complex than stated here. If the result is negative, the hardware combines rounding with placing the high order word in twos complement form by decreasing its magnitude if the low order part is < ½LSB. Moreover an extra single-step renormalization occurs if the rounded word is no longer normalized.

Keeping instructions and operands in different memories saves .47 (.36) μs in memory and both modes.

When *E* addresses a fast memory location, a floating point instruction with rounding takes .34 μs less than the time listed in basic mode, .80 (.69) μs less in memory or both mode.

There are four instructions that use only one-word operands and store a single-length rounded result. Rounding is away from zero: if the part of the normalized answer being dropped (the low order part of the fraction) is greater than or equal in magnitude to one half the LSB of the part being retained, the magnitude of the latter part is increased by one LSB.

The rounding instructions have four modes that determine the source of the non-AC operand and the destination of the result. These modes are like those of logic and fixed point arithmetic, including an immediate mode that allows the instruction to carry an operand with it.

Mode	Suffix	Source of non-AC operand	Destination of result
Basic		*E*	AC
▲ Immediate	I	The word *E*,0	AC
Memory	M	*E*	*E*
Both	B	*E*	AC and *E*

Note however that floating point immediate uses *E*,0 as an operand, not 0,*E*. In other words the half word *E* is interpreted as a sign, an 8-bit exponent, and a 9-bit fraction.

The time required is a function of the number *N* of left shifts needed for normalization. Brackets enclose the additional time required when rounding actually changes the high order word.

In each of these instructions, the exponent that results from normaliza-

tion and rounding is tested for overflow or underflow. If the exponent is > 127, set Overflow and Floating Overflow; the result stored has an exponent 256 less than the correct one. If < −128, set Overflow, Floating Overflow and Floating Underflow; the result stored has an exponent 256 greater than the correct one.

FADR   Floating Add and Round

144	M	A	I	X	Y
0	6 7	8 9	12 13 14	17 18	35
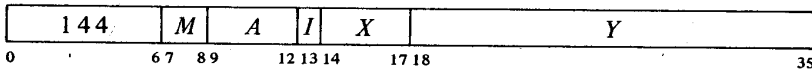
Floating add the operand specified by M to AC. If the double length fraction in the sum is zero, clear the specified destination. Otherwise normalize the double length sum bringing 0s into bit positions vacated at the right, round the high order part, test for exponent overflow or underflow as described above, and place the result in the specified destination.

FADR   Floating Add and Round   144
                        $4.46\ (4.68) + .15D + .25N\ [+.96]\ \mu s$

FADRI   Floating Add and Round Immediate   145
                        $3.70\ (3.81) + .15D + .25N\ [+.96]\ \mu s$

FADRM   Floating Add and Round to Memory   146
                        $5.43\ (5.65) + .15D + .25N\ [+.96]\ \mu s$

FADRB   Floating Add and Round to Both   147
                        $5.43\ (5.65) + .15D + .25N\ [+.96]\ \mu s$

D is the difference between the operand exponents provided that difference is $\leqslant 63$. Otherwise $D = 0$.

FSBR   Floating Subtract and Round

154	M	A	I	X	Y
0	6 7	8 9	12 13 14	17 18	35
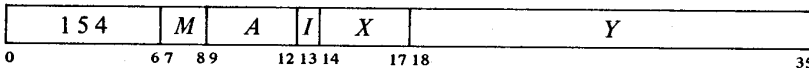
Floating subtract the operand specified by M from AC. If the double length fraction in the difference is zero, clear the specified destination. Otherwise normalize the double length difference bringing 0s into bit positions vacated at the right, round the high order part, test for exponent overflow or underflow as described above, and place the result in the specified destination.

FSBR   Floating Subtract and Round   154
                        $4.64\ (4.86) + .15D + .15N\ [+.96]\ \mu s$

FSBRI   Floating Subtract and Round Immediate   155
                        $3.88\ (3.99) + .15D + .15N\ [+.96]\ \mu s$

FSBRM   Floating Subtract and Round to Memory   156
                        $5.61\ (5.83) + .15D + .15N\ [+.96]\ \mu s$

FSBRB   Floating Subtract and Round to Both   157
                        $5.61\ (5.83) + .15D + .15N\ [+.96]\ \mu s$

D is the difference between the operand exponents provided that difference is $\leqslant 63$. Otherwise $D = 0$.

FMPR　　　　**Floating Multiply and Round**

1 6 4	M	A	I	X		Y
0	6 7	8 9		12 13 14	17 18	35
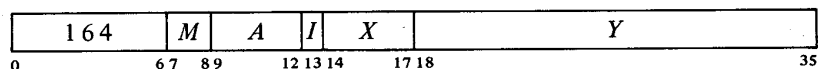
Floating Multiply AC by the operand specified by M. If the double length fraction in the product is zero, clear the specified destination. Otherwise normalize the double length product bringing 0s into bit positions vacated at the right, round the high order part, test for exponent overflow or underflow as described above, and place the result in the specified destination.

Use of normalized operands requires at most one normalization step for the result. If unnormalized operands are used, all times must be increased by .25N.

FMPR	Floating Multiply and Round	164
		10.29 (10.51) [+.96] μs
FMPRI	Floating Multiply and Round Immediate	165
		8.36 (8.47) [+.96] μs
FMPRM	Floating Multiply and Round to Memory	166
		11.26 (11.48) [+.96] μs
FMPRB	Floating Multiply and Round to Both	167
		11.26 (11.48) [+.96] μs

Timing. The times given above are average for normalized operands. Refer to the description of MUL [§2.5] for the timing effects of the multiplication algorithm. Minimum times with a zero multiplier are

FMPR	8.47 (8.69) [+.96] μs
FMPRI	7.71 (7.82) [+.96] μs
FMPRM	9.44 (9.66) [+.96] μs
FMPRB	9.44 (9.66) [+.96] μs

These must be increased by .13 μs for each transition. The programmer can minimize the time by using as the multiplier the operand with fewer transitions.

FDVR　　　　**Floating Divide and Round**
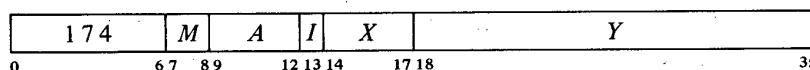
1 7 4	M	A	I	X		Y
0	6 7	8 9		12 13 14	17 18	35

Division fails if the divisor is zero, but the no-divide condition can otherwise be satisfied only if at least one operand is unnormalized.

If the magnitude of the fraction in AC is greater than or equal to twice that of the fraction in the operand specified by M, set Overflow, Floating Overflow and No Divide, and go immediately to the next instruction without affecting the original AC or memory operand in any way.

If the division can be performed, floating divide AC by the operand specified by M, calculating a quotient fraction of 28 bits (this includes an extra bit for rounding). If the fraction is zero, clear the specified destination.
▲ Otherwise round the fraction using the extra bit calculated. If the original operands were normalized, the single length quotient will already be normalized; if it is not, normalize it bringing 0s into bit positions vacated at the right. Test for

exponent overflow or underflow as described above, and place the result in the specified destination.

FDVR	Floating Divide and Round		174
			14.1 (14.3) μs
FDVRI	Floating Divide and Round Immediate		175
			13.3 (13.4) μs
FDVRM	Floating Divide and Round to Memory		176
			15.1 (15.3) μs
FDVRB	Floating Divide and Round to Both		177
			15.1 (15.3) μs

If unnormalized operands are used, all times must be increased by .25N. If the division is not performed, only 3.5–4 μs are required.

Operations without Rounding

Instructions that do not round are faster for processing floating point numbers with fractions containing fewer than 27 significant bits. On the other hand the long mode provides double precision or allows the programmer to use his own method of rounding. Besides the four usual arithmetic operations with normalization, there are two nonnormalizing instructions that facilitate double precision arithmetic [§2.11 *gives examples of double precision floating point routines*] . These two instructions have no modes.

DFN   **Double Floating Negate**   3.43 (3.54) μs

1 3 1	A	I	X	Y
0	8 9	12 13 14	17 18	35
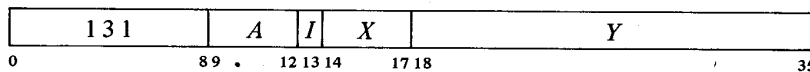
Negate the double length floating point number composed of the contents of AC and location E with AC on the left. Do this by taking the twos complement of the number whose sign is AC bit 0, whose exponent is in AC bits 1–8, and whose fraction is the 54-bit string in bits 9–35 of AC and location E. Place the high order word of the result in AC; place the low order part of the fraction in bits 9–35 of location E without altering the original contents of bits 0–8 of that location.

Usually the double length number is in two adjacent accumulators, and E equals A+1. In this case DFN takes ▲ only 2.89 (3.11) μs.

UFA   **Unnormalized Floating Add**   4.62 (4.84) + .15D μs
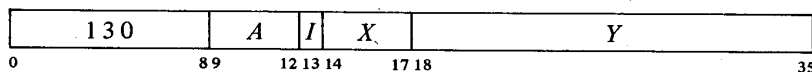
1 3 0	A	I	X	Y
0	8 9	12 13 14	17 18	35

Floating add the contents of location E to AC. If the double length fraction in the sum is zero, clear accumulator A+1. Otherwise normalize the sum only if the magnitude of its fractional part is $\geq$ 1, and place the high order part of the result in AC A+1. The original contents of AC and E are unaffected.

D is the difference between the operand exponents provided that difference is $\leq$ 63. Otherwise D = 0.

When E addresses a fast memory location, UFA takes .34 μs less than the time given.

The exponent of the sum is equal to that of the larger summand unless addition of the fractions overflows, in which case it is greater by 1. Exponent overflow can occur only in the latter case.

NOTE

The result is placed in accumulator *A+1*. This is the only arithmetic instruction that stores the result in a second accumulator, leaving the original operands intact.

If the exponent of the sum following the one-step normalization is > 127, set Overflow and Floating Overflow; the result stored has an exponent 256 less than the correct one.

The remaining floating point instructions perform the four standard arithmetic operations with normalization but without rounding. All use AC and the contents of location *E* as operands and have four modes.

Keeping instructions and operands in different memories saves .47 (.36) μs in memory and both modes.

When *E* addresses a fast memory location, a floating point instruction without rounding takes .34 μs less than the time listed in basic or long mode, .80 (.69) μs less in memory or both mode.

Mode	*Suffix*	*Effect*
Basic		High order word of result stored in AC.
Long	L	In addition, subtraction and multiplication, the two-word result (in the double length format described in §1.1) is stored in accumulators *A* and *A*+1. In division the dividend is the double length word in *A* and *A*+1; the single length quotient is stored in AC, the remainder in AC *A*+1.
Memory	M	High order word of result stored in *E*.
Both	B	High order word of result stored in AC and *E*.

In each of these instructions, the exponent that results from normalization is tested for overflow or underflow. If the exponent is > 127, set Overflow and Floating Overflow; the result stored has an exponent 256 less than the correct one. If < −128, set Overflow, Floating Overflow and Floating Underflow; the result stored has an exponent 256 greater than the correct one.

The time required is a function of the number *N* of left shifts needed for normalization.

FAD   Floating Add
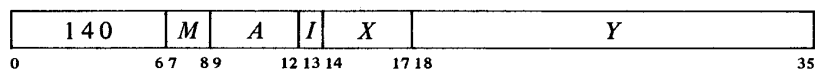
140	M	A	I	X	Y

0   6 7   8 9   12 13 14   17 18   35

Floating add the contents of location *E* to AC. If the double length fraction in the sum is zero, clear the destination specified by *M*, clearing both accu-

mulators in long mode. Otherwise normalize the double length sum bringing
0s into bit positions vacated at the right, test for exponent overflow or
underflow as described above, and place the high order word of the result in
the specified destination.

In long mode if the exponent of the sum is > 154 $(127 + 27)$ or < -101
$(-128 + 27)$ or the low order half of the fraction is zero, clear AC $A+1$.
Otherwise place a low order word for a double length result in $A+1$ by
putting a 0 in bit 0, an exponent in positive form 27 less than the exponent
of the sum in bits 1–8, and the low order part of the fraction in bits 9–35.

FAD	Floating Add	140	D is the difference between
		$4.46\ (4.68) + .15D + .25N\ \mu s$	the operand exponents pro-
FADL	Floating Add Long	141	vided that difference is $\leqslant 63$.
		$5.31\ (5.53) + .15D + .25N\ \mu s$	Otherwise $D = 0$.
FADM	Floating Add to Memory	142	
		$5.43\ (5.65) + .15D + .25N\ \mu s$	
FADB	Floating Add to Both	143	
		$5.43\ (5.65) + .15D + .25N\ \mu s$	

FSB   Floating Subtract

1 5 0	M	A	I	X	Y
0	6 7	8 9	12 13 14	17 18	35
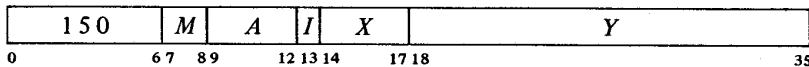
Floating subtract the contents of location E from AC. If the double length
fraction in the difference is zero, clear the destination specified by M, clear-
ing both accumulators in long mode. Otherwise normalize the double length
difference bringing 0s into bit positions vacated at the right, test for expo-
nent overflow or underflow as described above, and place the high order
word of the result in the specified destination.

In long mode if the exponent of the difference is > 154 $(127 + 27)$ or
< -101 $(-128 + 27)$ or the low order half of the fraction is zero, clear AC
$A+1$. Otherwise place a low order word for a double length result in $A+1$ by
putting a 0 in bit 0, an exponent in positive form 27 less than the exponent
of the difference in bits 1–8, and the low order part of the fraction in bits
9–35.

FSB	Floating Subtract	150	D is the difference between
		$4.64\ (4.86) + .15D + .25N\ \mu s$	the operand exponents pro-
FSBL	Floating Subtract Long	151	vided that difference is $\leqslant 63$.
		$5.49\ (5.71) + .15D + .25N\ \mu s$	Otherwise $D = 0$.
FSBM	Floating Subtract to Memory	152	
		$5.61\ (5.83) + .15D + .25N\ \mu s$	
FSBB	Floating Subtract to Both	153	
		$5.61\ (5.83) + .15D + .25N\ \mu s$	

FMP   Floating Multiply

1 6 0	M	A	I	X	Y
0	6 7	8 9	12 13 14	17 18	35
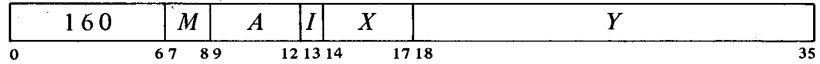
Floating multiply AC' by the contents of location E. If the double length fraction in the product is zero, clear the destination specified by M, clearing both accumulators in long mode. Otherwise normalize the double length product bringing 0s into bit positions vacated at the right, test for exponent overflow or underflow as described above, and place the high order word of the result in the specified destination.

In long mode if the exponent of the product is > 154 $(127 + 27)$ or < -101 $(-128 + 27)$ or the low order half of the fraction is zero, clear AC $A+1$. Otherwise place a low order word for a double length result in $A+1$ by putting a 0 in bit 0, an exponent in positive form 27 less than the exponent of the product in bits 1–8, and the low order part of the fraction in bits 9–35.

FMP	Floating Multiply	160	10.29 (10.51) μs
FMPL	Floating Multiply Long	161	11.14 (11.36) μs
FMPM	Floating Multiply to Memory	162	11.26 (11.48) μs
FMPB	Floating Multiply to Both	163	11.26 (11.48) μs

Use of normalized operands requires at most one normalization step for the result. If unnormalized operands are used, all times must be increased by .25N.

Timing. The times given above are average for normalized operands. Refer to the description of MUL [§2.5] for the timing effects of the multiplication algorithm. Minimum times with a zero multiplier are

FMP	8.47 (8.69) μs
FMPL	9.32 (9.54) μs
FMPM	9.44 (9.66) μs
FMPB	9.44 (9.66) μs

These must be increased by .13 μs for each transition. The programmer can minimize the time by using as the multiplier the operand with fewer transitions.

FDV   Floating Divide
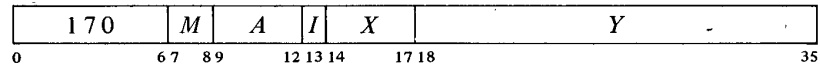
1 7 0	M	A	I	X	Y
0	6 7	8 9	12 13 14	17 18	35

Division fails if the divisor is zero, but the no-divide condition can otherwise be satisfied only if at least one operand is unnormalized.

If the magnitude of the fraction in AC is greater than or equal to twice that of the fraction in location E, set Overflow, Floating Overflow and No Divide, and go immediately to the next instruction without affecting the original AC or memory operand in any way.

If division can be performed, floating divide the AC operand by the contents of location E. In long mode the AC operand (the dividend) is the double length number in accumulators A and $A+1$; in other modes it is the single word in AC. Calculate a quotient fraction of 27 bits. If the fraction

is zero, clear the destination specified by *M*, clearing both accumulators in long mode if the double length dividend was zero. A quotient with a nonzero fraction will already be normalized if the original operands were normalized; if it is not, normalize it bringing 0s into bit positions vacated at the right. Test for exponent overflow or underflow as described above, and place the single length quotient part of the result in the specified destination.

In long mode calculate the exponent for the fractional remainder from the division according to the relative magnitudes of the fractions in dividend and divisor: if the dividend was greater than or equal to the divisor, the exponent of the remainder is 26 less than that of the dividend, otherwise it is 27 less. If the remainder exponent is > 127 or < −128 or the fraction is zero, clear AC *A*+1. Otherwise place the floating point remainder (exponent and fraction) with the sign of the dividend in AC *A*+1.

In long mode a nonzero unnormalized dividend whose entire high order fraction is zero produces a zero quotient. In this case the second AC receives rubbish.

FDV	Floating Divide	170	14.1 (14.3) μs
FDVL	Floating Divide Long	171	15.6 (15.8) μs
FDVM	Floating Divide to Memory	172	15.1 (15.3) μs
FDVB	Floating Divide to Both	173	15.1 (15.3) μs

If unnormalized operands are used, all times must be increased by .25*N*. If the division is not performed, only 4–4.5 μs are required.

2.7   ARITHMETIC TESTING

These instructions may jump or skip depending on the result of an arithmetic test and may first perform an arithmetic operation on the test word. Two of the instructions have no modes.

AOBJP   Add One to Both Halves of AC and Jump if Positive   1.68 (1.79) μs

2 5 2		*A*	*I*	*X*	*Y*
0		8 9	12 13 14	17 18	35

Add $1\,000001_8$ to AC and place the result back in AC. If the result is greater than or equal to zero (*ie* if bit 0 is 0, and hence a negative count in the left half has reached zero or a positive count has not yet reached 2^{17}), take the next instruction from location *E* and continue sequential operation from there.

AOBJN   Add One to Both Halves of AC and Jump if Negative   1.68 (1.79) μs

2 5 3		*A*	*I*	*X*	*Y*
0		8 9	12 13 14	17 18	35

Add $1\,000001_8$ to AC and place the result back in AC. If the result is less than zero (*ie* if bit 0 is 1, and hence a negative count in the left half has not yet reached zero or a positive count has reached 2^{17}), take the next instruction from location *E* and continue sequential operation from there.

The incrementing of both halves of AC simultaneously is effected by adding 1000001_8. A count of -2 in AC left is therefore increased to zero if $2^{18} - 1$ is incremented in AC right.

These two instructions allow the program to keep a control count in the left half of an index register and require only one data transfer to initialize. Problem: Add 3 to each location in a table of N entries starting at TAB. Only four instructions are required.

MOVSI	XR,$-N$	;Put $-N$ in XR left (clear XR right)
MOVEI	AC,3	;Put 3 in AC
ADDM	AC,TAB(XR)	;Add 3 to entry
AOBJN	XR,-1	;Update XR and go back unless all ;entries accounted for

The eight remaining instructions jump or skip if the operand or operands satisfy a test condition specified by the mode.

Mode	Suffix
Never	
Less	L
Equal	E
Less or Equal	LE
Always	A
Greater or Equal	GE
Not Equal	N
Greater	G

Instructions with one operand compare AC or the contents of location E with zero, those with two compare AC with E or the contents of location E. The processor always makes the comparison even though the result is used in only six of the modes. If the mnemonic has no suffix there is never any program control function, and the instruction may be a no-op; an A suffix produces an unconditional jump or skip — the action is always taken regardless of how the two quantities compare.

CAI   **Compare AC Immediate and Skip if Condition Satisfied**   1.68 (1.79) μs

30	M	A	I	X	Y
0	5 6	8 9	12 13 14	17 18	35
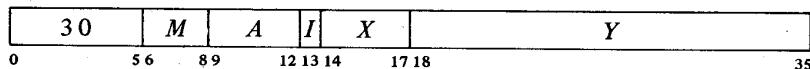
Compare AC with E (ie with the word $0,E$) and skip the next instruction in sequence if the condition specified by M is satisfied.

CAI	Compare AC Immediate but Do Not Skip	300
CAIL	Compare AC Immediate and Skip if AC Less than E	301
CAIE	Compare AC Immediate and Skip if Equal	302
CAILE	Compare AC Immediate and Skip if AC Less than or Equal to E	303
CAIA	Compare AC Immediate but Always Skip	304
CAIGE	Compare AC Immediate and Skip if AC Greater than or Equal to E	305
CAIN	Compare AC Immediate and Skip if Not Equal	306
CAIG	Compare AC Immediate and Skip if AC Greater than E	307

CAI is a no-op.

CAM   **Compare AC with Memory and Skip if Condition Satisfied**   2.53 (2.75) μs

3 1	M	A	I	X	Y
0	5 6	8 9	12 13 14	17 18	35

When E addresses a fast memory location, this instruction takes .34 μs less than the time given.

Compare AC with the contents of location E and skip the next instruction in sequence if the condition specified by M is satisfied. The pair of numbers compared may be either both fixed or both normalized floating point.

CAM	Compare AC with Memory but Do Not Skip	310
CAML	Compare AC with Memory and Skip if AC Less	311
CAME	Compare AC with Memory and Skip if Equal	312
CAMLE	Compare AC with Memory and Skip if AC Less or Equal	313
CAMA	Compare AC with Memory but Always Skip	314
CAMGE	Compare AC with Memory and Skip if AC Greater or Equal	315
CAMN	Compare AC with Memory and Skip if Not Equal	316
CAMG	Compare AC with Memory and Skip if AC Greater	317

CAM is a no-op that references memory.

JUMP   **Jump if AC Condition Satisfied**   1.68 (1.79) μs

3 2	M	A	I	X	Y
0	5 6	8 9	12 13 14	17 18	35

Compare AC (fixed or floating) with zero, and if the condition specified by M is satisfied, take the next instruction from location E and continue sequential operation from there.

JUMP	Do Not Jump	320
JUMPL	Jump if AC Less than Zero	321
JUMPE	Jump if AC Equal to Zero	322

JUMP is a no-op (instruction code 320 has this mnemonic for symmetry).

62

JUMPLE	Jump if AC Less than or Equal to Zero	323
JUMPA	Jump Always	324
JUMPGE	Jump if AC Greater than or Equal to Zero	325
JUMPN	Jump if AC Not Equal to Zero	326
JUMPG	Jump if AC Greater than Zero	327

When E addresses a fast memory location, this instruction takes .34 μs less than the time given.

SKIP   **Skip if Memory Condition Satisfied**   2.39 (2.61) μs

3 3	M	A	I	X	Y
0	5 6	8 9	12 13 14	17 18	35

If A is zero, SKIP is a no-op; otherwise it is equivalent to MOVE. (Instruction code 330 has mnemonic SKIP for symmetry.)
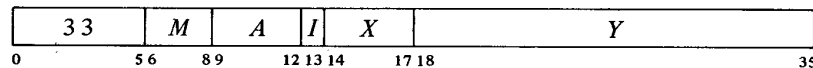
Compare the contents (fixed or floating) of location E with zero, and skip the next instruction in sequence if the condition specified by M is satisfied. If A is nonzero also place the contents of location E in AC.

SKIP	Do Not Skip	330
SKIPL	Skip if Memory Less than Zero	331
SKIPE	Skip if Memory Equal to Zero	332
SKIPLE	Skip if Memory Less than or Equal to Zero	333
SKIPA	Skip Always	334
SKIPGE	Skip if Memory Greater than or Equal to Zero	335
SKIPN	Skip if Memory Not Equal to Zero	336
SKIPG	Skip if Memory Greater than Zero	337

SKIPA is a convenient way to load an accumulator and skip over an instruction upon entering a loop.

AOJ   **Add One to AC and Jump if Condition Satisfied**   1.68 (1.79) μs

3 4	M	A	I	X	Y
0	5 6	8 9	12 13 14	17 18	35
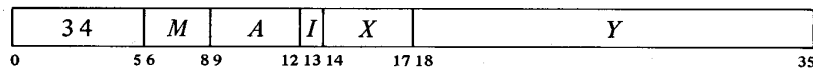
Increment AC by one and place the result back in AC. Compare the result with zero, and if the condition specified by M is satisfied, take the next instruction from location E and continue sequential operation from there. If AC originally contained $2^{35} - 1$, set the Overflow and Carry 1 flags; if -1, set Carry 0 and Carry 1.

AOJ	Add One to AC but Do Not Jump	340
AOJL	Add One to AC and Jump if Less than Zero	341
AOJE	Add One to AC and Jump if Equal to Zero	342
AOJLE	Add One to AC and Jump if Less than or Equal to Zero	343
AOJA	Add One to AC and Jump Always	344
AOJGE	Add One to AC and Jump if Greater than or Equal to Zero	345
AOJN	Add One to AC and Jump if Not Equal to Zero	346
AOJG	Add One to AC and Jump if Greater than Zero	347

AOS   **Add One to Memory and Skip if Condition Satisfied**   2.94 (3.05) μs

3 5	M	A	I	X	Y
0	5 6	8 9	12 13 14	17 18	35
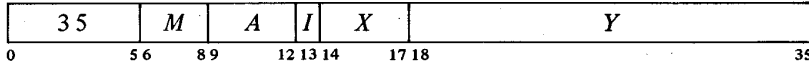
Increment the contents of location E by one and place the result back in E.
Compare the result with zero, and skip the next instruction in sequence if
the condition specified by M is satisfied. If location E originally contained
$2^{35} - 1$, set the Overflow and Carry 1 flags; if -1, set Carry 0 and Carry 1.
If A is nonzero also place the result in AC.

AOS	Add One to Memory but Do Not Skip	350
AOSL	Add One to Memory and Skip if Less than Zero	351
AOSE	Add One to Memory and Skip if Equal to Zero	352
AOSLE	Add One to Memory and Skip if Less than or Equal to Zero	353
AOSA	Add One to Memory and Skip Always	354
AOSGE	Add One to Memory and Skip if Greater than or Equal to Zero	355
AOSN	Add One to Memory and Skip if Not Equal to Zero	356
AOSG	Add One to Memory and Skip if Greater than Zero	357

Keeping the count in fast memory saves .54 (.43) μs; keeping it in a different memory from the instruction saves .20 (.09) μs.

SOJ   **Subtract One from AC and Jump if Condition Satisfied**   1.68 (1.79) μs

3 6	M	A	I	X	Y
0	5 6	8 9	12 13 14	17 18	35
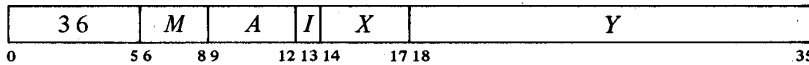
Decrement AC by one and place the result back in AC. Compare the result
with zero, and if the condition specified by M is satisfied, take the next in-
struction from location E and continue sequential operation from there. If
AC originally contained -2^{35}, set the Overflow and Carry 0 flags; if any other
nonzero number, set Carry 0 and Carry 1.

SOJ	Subtract One from AC but Do Not Jump	360
SOJL	Subtract One from AC and Jump if Less than Zero	361
SOJE	Subtract One from AC and Jump if Equal to Zero	362
SOJLE	Subtract One from AC and Jump if Less than or Equal to Zero	363
SOJA	Subtract One from AC and Jump Always	364
SOJGE	Subtract One from AC and Jump if Greater than or Equal to Zero	365
SOJN	Subtract One from AC and Jump if Not Equal to Zero	366
SOJG	Subtract One from AC and Jump if Greater than Zero	367

Keeping the count in fast memory saves .54 (.43) μs; keeping it in a different memory from the instruction saves .20 (.09) μs.

SOS   **Subtract One from Memory and Skip if Condition**   2.94 (3.05) μs
           Satisfied

37	M	A	I	X	Y
0	5 6	8 9	12 13 14	17,18	35
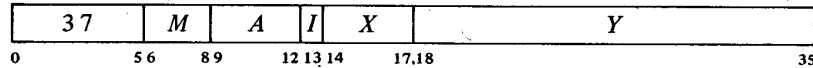
Decrement the contents of location E by one and place the result back in E. Compare the result with zero, and skip the next instruction in sequence if the condition specified by M is satisfied. If location E originally contained -2^{35}, set the Overflow and Carry 0 flags; if any other nonzero number, set Carry 0 and Carry 1. If A is nonzero also place the result in AC.

SOS	Subtract One from Memory but Do Not Skip	370
SOSL	Subtract One from Memory and Skip if Less than Zero	371
SOSE	Subtract One from Memory and Skip if Equal to Zero	372
SOSLE	Subtract One from Memory and Skip if Less than or Equal to Zero	373
SOSA	Subtract One from Memory and Skip Always	374
SOSGE	Subtract One from Memory and Skip if Greater than or Equal to Zero	375
SOSN	Subtract One from Memory and Skip if Not Equal to Zero	376
SOSG	Subtract One from Memory and Skip if Greater than Zero	377

Some of these instructions are useful for determining the relative values of fixed and floating point numbers; others are convenient for controlling iterative processes by counting. AOSE is especially useful in an interlock procedure in a multiprocessor system. Suppose memory contains a routine that must be available to two processors but cannot be used by both at once. When one processor finishes the routine it sets location LOCK to -1. Either processor can then test the interlock and make it busy with no possibility of letting the other one in, as AOSE cannot be interrupted once it starts to modify the addressed location.

This procedure is invalid if the programmer is making use of the drum split feature (which is not used by any DEC equipment).

```
AOSE   LOCK        ;Skip to interlocked code only if
JRST   .-1         ;LOCK is zero after addition
.                  ;Interlocked code starts here
.
.
SETOM  LOCK        ;Unlock
```

Since it takes several days to count to 2^{36}, it is alright to keep testing the lock.

2.8 LOGICAL TESTING AND MODIFICATION

These eight instructions use a mask to modify and/or test selected bits in
AC. The bits are those that correspond to 1s in the mask and they are
referred to as the "masked bits". The programmer chooses the mask, the
way in which the masked bits are to be modified, and the condition the
masked bits must satisfy to produce a skip.

The basic mnemonics are three letters beginning with T. The second letter
selects the mask and the manner in which it is used.

Mask	Letter	Effect	
Right	R	AC right is masked by E (AC is masked by the word $0,E$)	
Left	L	AC left is masked by E (AC is masked by the word $E,0$)	
Direct	D	AC is masked by the contents of location E	If a direct or swapped mask is taken from a fast memory location, a test instruction takes .34 μs less than the time listed.
Swapped	S	AC is masked by the contents of location E with left and right halves interchanged	

The third letter determines the way in which those bits selected by the mask
are modified.

Modification	Letter	Effect on AC
No	N	None
Zeros	Z	Places 0s in all masked bit positions
Complement	C	Complements all masked bits
Ones	O	Places 1s in all masked bit positions

An additional letter may be appended to indicate the mode, which spec-
ifies the condition the masked bits must satisfy to produce a skip.

Mode	Suffix	Effect	
Never		Never skip	These mode names are consistent with those for arithmetic testing and derive from the test method, which ands AC with the mask and tests whether the result is equal to zero or is not equal to zero. The programmer may find it convenient to think of the modes as Every and Not Every: every masked bit is 0 or not every masked bit is 0.
Equal	E	Skip if all masked bits equal 0	
Always	A	Always skip	
Not Equal	N	Skip if not all masked bits equal 0 (at least one bit is 1)	

If the mnemonic has no suffix there is never any skip, and the instruction is
a no-op if there is also no modification; an A suffix produces an uncondi-
tional skip — the skip always occurs regardless of the state of the masked
bits. Note that the skip condition must be satisfied by the state of the
masked bits *prior* to any modification called for by the instruction.

TRN   **Test Right, No Modification, and Skip if Condition**   1.85 (1.96) μs
           Satisfied
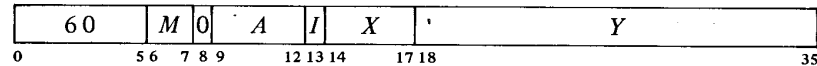
6 0	M	0	A	I	X	Y
0	5 6	7 8 9	12 13	14	17 18	35

If the bits in AC right corresponding to 1s in E satisfy the condition specified by M, skip the next instruction in sequence. AC is unaffected.

TRN is a no-op.

TRN	Test Right, No Modification, but Do Not Skip	600
TRNE	Test Right, No Modification, and Skip if All Masked Bits Equal 0	602
TRNA	Test Right, No Modification, but Always Skip	604
TRNN	Test Right, No Modification, and Skip if Not All Masked Bits Equal 0	606

TRZ   **Test Right, Zeros, and Skip if Condition Satisfied**   1.85 (1.96) μs
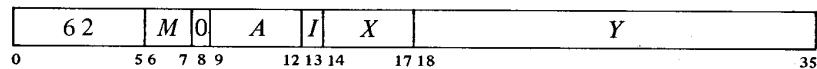
6 2	M	0	A	I	X	Y
0	5 6	7 8 9	12 13	14	17 18	35

If the bits in AC right corresponding to 1s in E satisfy the condition specified by M, skip the next instruction in sequence. Change the masked AC bits to 0s; the rest of AC is unaffected.

TRZ	Test Right, Zeros, but Do Not Skip	620
TRZE	Test Right, Zeros, and Skip if All Masked Bits Equaled 0	622
TRZA	Test Right, Zeros, but Always Skip	624
TRZN	Test Right, Zeros, and Skip if Not All Masked Bits Equaled 0	626

TRC   **Test Right, Complement, and Skip if Condition**   1.85 (1.96) μs
           Satisfied
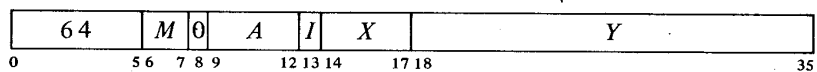
6 4	M	0	A	I	X	Y
0	5 6	7 8 9	12 13	14	17 18	35

If the bits in AC right corresponding to 1s in E satisfy the condition specified by M, skip the next instruction in sequence. Complement the masked AC bits; the rest of AC is unaffected.

TRC	Test Right, Complement, but Do Not Skip	640
TRCE	Test Right, Complement, and Skip if All Masked Bits Equaled 0	642
TRCA	Test Right, Complement, but Always Skip	644
TRCN	Test Right, Complement, and Skip if Not All Masked Bits Equaled 0	646

TRO   **Test Right, Ones, and Skip if Condition Satisfied**   1.85 (1.96) μs

6 6	M	0	A	I	X	Y
0	5 6	7 8 9	12 13 14	17 18		35

If the bits in AC right corresponding to 1s in E satisfy the condition specified by M, skip the next instruction in sequence. Change the masked AC bits to 1s; the rest of AC is unaffected.

TRO	Test Right, Ones, but Do Not Skip	660
TROE	Test Right, Ones, and Skip if All Masked Bits Equaled 0	662
TROA	Test Right, Ones, but Always Skip	664
TRON	Test Right, Ones, and Skip if Not All Masked Bits Equaled 0	666

TLN   **Test Left, No Modification, and Skip if Condition Satisfied**   1.85 (1.96) μs

6 0	M	1	A	I	X	Y
0	5 6	7 8 9	12 13 14	17 18		35

If the bits in AC left corresponding to 1s in E satisfy the condition specified by M, skip the next instruction in sequence. AC is unaffected.

TLN	Test Left, No Modification, but Do Not Skip	601	TLN is a no-op.
TLNE	Test Left, No Modification, and Skip if All Masked Bits Equal 0	603	
TLNA	Test Left, No Modification, but Always Skip	605	
TLNN	Test Left, No Modification, and Skip if Not All Masked Bits Equal 0	607	

TLZ   **Test Left, Zeros, and Skip if Condition Satisfied**   1.85 (1.96) μs

6 2	M	1	A	I	X	Y
0	5 6	7 8 9	12 13 14	17 18		35

If the bits in AC left corresponding to 1s in E satisfy the condition specified by M, skip the next instruction in sequence. Change the masked AC bits to 0s; the rest of AC is unaffected.

TLZ	Test Left, Zeros, but Do Not Skip	621
TLZE	Test Left, Zeros, and Skip if All Masked Bits Equaled 0	623
TLZA	Test Left, Zeros, but Always Skip	625
TLZN	Test Left, Zeros, and Skip if Not All Masked Bits Equaled 0	627

TLC   **Test Left, Complement, and Skip if Condition Satisfied**   1.85 (1.96) μs
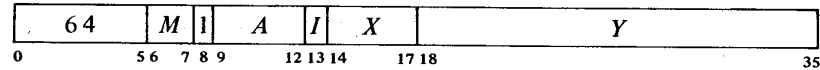
6 4	M	1	A	I	X	Y
0	5 6	7 8 9	12 13 14	17 18		35

If the bits in AC left corresponding to 1s in E satisfy the condition specified by M, skip the next instruction in sequence. Complement the masked AC bits; the rest of AC is unaffected.

TLC	Test Left, Complement, but Do Not Skip	641
TLCE	Test Left, Complement, and Skip if All Masked Bits Equaled 0	643
TLCA	Test Left, Complement, but Always Skip	645
TLCN	Test Left, Complement, and Skip if Not All Masked Bits Equaled 0	647

TLO   **Test Left, Ones, and Skip if Condition Satisfied**   1.85 (1.96) μs
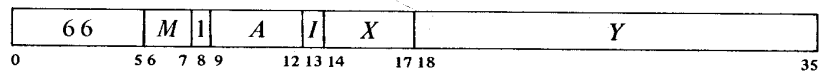
6 6	M	1	A	I	X	Y
0	5 6	7 8 9	12 13 14	17 18		35

If the bits in AC left corresponding to 1s in E satisfy the condition specified by M, skip the next instruction in sequence. Change the masked AC bits to 1s; the rest of AC is unaffected.

TLO	Test Left, Ones, but Do Not Skip	661
TLOE	Test Left, Ones, and Skip if All Masked Bits Equaled 0	663
TLOA	Test Left, Ones, but Always Skip	665
TLON	Test Left, Ones, and Skip if Not All Masked Bits Equaled 0	667

TDN   **Test Direct, No Modification, and Skip if Condition Satisfied**   2.70 (2.92) μs
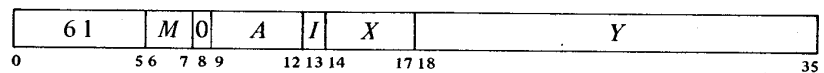
6 1	M	0	A	I	X	Y
0	5 6	7 8 9	12 13 14	17 18		35

If the bits in AC corresponding to 1s in the contents of location E satisfy the condition specified by M, skip the next instruction in sequence. AC is unaffected.

TDN is a no-op that references memory.

TDN	Test Direct, No Modification, but Do Not Skip	610
TDNE	Test Direct, No Modification, and Skip if All Masked Bits Equal 0	612
TDNA	Test Direct, No Modification, but Always Skip	614
TDNN	Test Direct, No Modification, and Skip if Not All Masked Bits Equal 0	616

TDZ   **Test Direct, Zeros, and Skip if Condition Satisfied**   2.70 (2.92) μs

6 3	M	0	A	I	X	Y
0	5 6	7 8 9	12 13 14	17 18		35
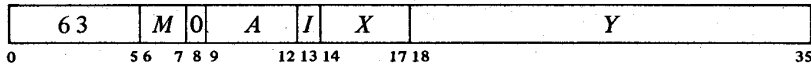
If the bits in AC corresponding to 1s in the contents of location E satisfy the condition specified by M, skip the next instruction in sequence. Change the masked AC bits to 0s; the rest of AC is unaffected.

TDZ	Test Direct, Zeros, but Do Not Skip	630
TDZE	Test Direct, Zeros, and Skip if All Masked Bits Equaled 0	632
TDZA	Test Direct, Zeros, but Always Skip	634
TDZN	Test Direct, Zeros, and Skip if Not All Masked Bits Equaled 0	636

TDC   **Test Direct, Complement, and Skip if Condition Satisfied**   2.70 (2.92) μs
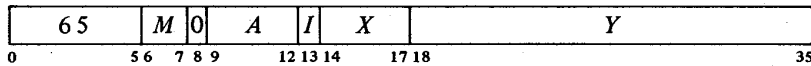
6 5	M	0	A	I	X	Y
0	5 6	7 8 9	12 13 14	17 18		35

If the bits in AC corresponding to 1s in the contents of location E satisfy the condition specified by M, skip the next instruction in sequence. Complement the masked AC bits; the rest of AC is unaffected.

TDC	Test Direct, Complement, but Do Not Skip	650
TDCE	Test Direct, Complement, and Skip if All Masked Bits Equaled 0	652
TDCA	Test Direct, Complement, but Always Skip	654
TDCN	Test Direct, Complement, and Skip if Not All Masked Bits Equaled 0	656

TDO   **Test Direct, Ones, and Skip if Condition Satisfied**   2.70 (2.92) μs
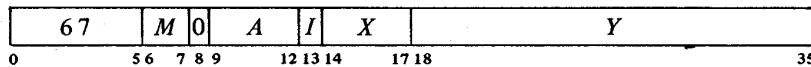
6 7	M	0	A	I	X	Y
0	5 6	7 8 9	12 13 14	17 18		35

If the bits in AC corresponding to 1s in the contents of location E satisfy the condition specified by M, skip the next instruction in sequence. Change the masked AC bits to 1s; the rest of AC is unaffected.

TDO	Test Direct, Ones, but Do Not Skip	670
TDOE	Test Direct, Ones, and Skip if All Masked Bits Equaled 0	672
TDOA	Test Direct, Ones, but Always Skip	674
TDON	Test Direct, Ones, and Skip if Not All Masked Bits Equaled 0	676

TSN   **Test Swapped, No Modification, and Skip if**   2.70 (2.92) μs
              Condition Satisfied
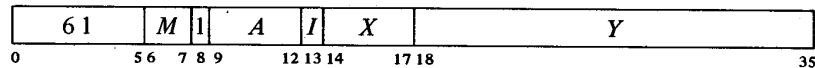
6 1		M	1	A	I	X	Y
0		5 6	7 8 9	12 13 14		17 18	35

If the bits in AC corresponding to 1s in the contents of location E with its left and right halves swapped satisfy the condition specified by M, skip the next instruction in sequence. AC is unaffected.

TSN is a no-op that references memory.

TSN   Test Swapped, No Modification, but Do Not Skip   611

TSNE   Test Swapped, No Modification, and Skip if All   613
              Masked Bits Equal 0

TSNA   Test Swapped, No Modification, but Always Skip   615

TSNN   Test Swapped, No Modification, and Skip if Not   617
              All Masked Bits Equal 0

TSZ   **Test Swapped, Zeros, and Skip if Condition Satisfied**   2.70 (2.92) μs
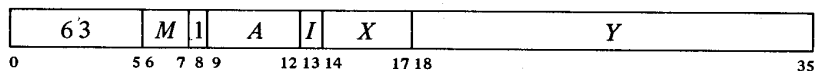
6 3		M	1	A	I	X	Y
0		5 6	7 8 9	12 13 14		17 18	35

If the bits in AC corresponding to 1s in the contents of location E with its left and right halves swapped satisfy the condition specified by M, skip the next instruction in sequence. Change the masked AC bits to 0s; the rest of AC is unaffected.

TSZ   Test Swapped, Zeros, but Do Not Skip   631

TSZE   Test Swapped, Zeros, and Skip if All Masked Bits   633
              Equaled 0

TSZA   Test Swapped, Zeros, but Always Skip   635

TSZN   Test Swapped, Zeros, and Skip if Not All Masked   637
              Bits Equaled 0

TSC   **Test Swapped, Complement, and Skip if Condition**   2.70 (2.92) μs
              Satisfied

6 5		M	1	A	I	X	Y
0		5 6	7 8 9	12 13 14		17 18	35
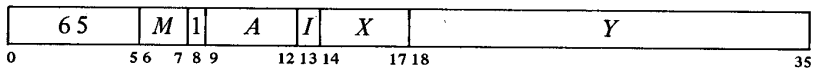
If the bits in AC corresponding to 1s in the contents of location E with its left and right halves swapped satisfy the condition specified by M, skip the next instruction in sequence. Complement the masked AC bits; the rest of AC is unaffected.

TSC   Test Swapped, Complement, but Do Not Skip   651

TSCE   Test Swapped, Complement, and Skip if All   653
              Masked Bits Equaled 0

TSCA   Test Swapped, Complement, but Always Skip   655
TSCN   Test Swapped, Complement, and Skip if Not   657
          All Masked Bits Equaled 0

TSO   **Test Swapped, Ones, and Skip if Condition Satisfied**   2.70 (2.92) μs
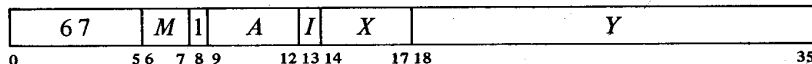
6 7	M	1	A	I	X	Y
0	5 6	7 8 9	12 13 14	17 18		35

If the bits in AC corresponding to 1s in the contents of location E with its
left and right halves swapped satisfy the condition specified by M, skip the
next instruction in sequence. Change the masked AC bits to 1s; the rest of
AC is unaffected.

TSO   Test Swapped, Ones, but Do Not Skip   671
TSOE   Test Swapped, Ones, and Skip if All Masked Bits   673
          Equaled 0
TSOA   Test Swapped, Ones, but Always Skip   675
TSON   Test Swapped, Ones, and Skip if Not All Masked   677
          Bits Equaled 0

With these instructions any bit throughout all of memory can be used as a
program flag, although an ordinary memory location containing flags must
be moved to an accumulator for testing or modification. The usual pro-
cedure, since locations 1–17 are addressable as index registers, is to use AC 0
as a register of flags (often addressed symbolically as F).

Unless one frequently tests flags in both halves of F simultaneously, it is
generally most convenient to select bits by 1s right in the address part of the
instruction word. A given bit selected by a half word mask M is then set by
one of these:

$$\text{TRO} \quad \text{F},M \qquad \text{TLO} \quad \text{F},M$$

and tested and cleared by one of these:

$$\text{TRZE} \quad \text{F},M \qquad \text{TRZN} \quad \text{F},M \qquad \text{TLZE} \quad \text{F},M \qquad \text{TLZN} \quad \text{F},M$$

Suppose we wish to skip if both bits 34 and 35 are 1 in location L. The
following suffices.

$$\text{SETCM} \quad \text{F,L}$$
$$\text{TRNE} \quad \text{F,3}$$

We can refer to a flag in a given bit position within a word as flag X, where X
is a binary number containing a single 1 in the same bit position as the flag.
This sequence determines whether flags X and Y in the right half of accumu-
lator F are both on:

```
TRC     F, X + Y        ;Complement flags X and Y
TRCE    F, X + Y        ;Test both and restore original states
. . .                   ;Do this if not both on
. . .                   ;Skip to here if both on
```

2.9 PROGRAM CONTROL

The program control class of instructions includes the unimplemented user operations [*discussed in the next section*] and the arithmetic and logical test instructions. Some instructions in this class are no-ops, as are a few of the instructions for performing logical operations. The most commonly used no-op is JFCL, which is discussed below. No-ops among the instructions previously discussed are SETA, SETAI, SETMM, CAI, CAM, JUMP, TRN, TLN, TDN, TSN. Of these, SETA, SETAI, CAI, JUMP, TRN and TLN do not use the calculated effective address to reference memory. Hence in these instructions one can store any information in bits 18–35 without fear of attempting to address a location outside a user block or in a memory that does not exist. The unassigned instruction codes 247 and 257 are used for instructions installed specially for a particular system. They execute as no-ops when run on a computer that contains no special hardware for them, but for program compatibility it is advised that they not be used regularly as no-ops.

The present section treats all program control instructions other than those mentioned above and in-out instructions that test input conditions [§2.12]. All but one of these are jumps, although the exception causes the processor to execute an instruction at an arbitrary location and may therefore be regarded as a jump with an immediate and automatic return. Also, all but two of the jumps are unconditional; one exception tests various flags, the other tests an accumulator.

Several of the jump instructions save the current contents of the program counter PC in the right half of an accumulator or memory location and save the states of various flags in the left half. The left bit positions that receive information are listed below; all other left bit positions are cleared. An X in a mnemonic indicates any letter (or none) that may appear in the given position to specify the mode, *eg* ADDX comprises ADD, ADDI, ADDM, ADDB.

Bit	Meaning of a 1 in the Bit
0	Overflow — any of the following has occurred:

A single instruction has set one of the carry flags (bits 1 and 2) without setting the other.

An ASH or ASHC has left shifted a 1 out of bit 1 in a positive number or a 0 out in a negative number.

An MULX has multiplied -2^{35} by itself (product 2^{70}).

An IMULX has multiplied two numbers with product $\geqslant 2^{35}$ or $< -2^{35}$.

As no-ops, code 247 takes 1.50 (1.61) μs, 257 takes 1.36 (1.47) μs.

Note that nothing is stored in bits 13–17, so when the PC word is addressed indirectly it can produce neither indexing nor further indirect addressing.

Floating Overflow has been set (bit 3).

No Divide has been set (bit 12).

1　　Carry 0 — if set without Carry 1 (bit 2) being set, causes Overflow to be set and indicates that one of the following has occurred:

An ADDX has added two negative numbers with sum $< -2^{35}$.

An SUBX has subtracted a positive number from a negative number with difference $< -2^{35}$.

An SOJX or SOSX has decremented -2^{35}.

But if set with Carry 1, indicates that one of these nonoverflow events has occurred:

In an ADDX both summands were negative, or their signs differed and their magnitudes were equal or the positive one was the greater in magnitude.

In an SUBX the signs of the operands were the same and AC was the greater or the two were equal, or the signs of the operands differed and AC was negative.

An AOJX or AOSX has incremented -1.

An SOJX or SOSX has decremented a nonzero number other than -2^{35}.

An MOVNX has negated zero.

2　　Carry 1 — if set without Carry 0 (bit 1) being set, causes Overflow to be set and indicates that one of the following has occurred:

An ADDX has added two positive numbers with sum $\geqslant 2^{35}$.

An SUBX has subtracted a negative number from a positive number with difference $\geqslant 2^{35}$.

An AOJX or AOSX has incremented $2^{35} - 1$.

An MOVNX or MOVMX has negated -2^{35}.

But if set with Carry 0, indicates that one of the nonoverflow events listed under Carry 0 has occurred.

3　　Floating Overflow — any of the following has set Overflow:

In a floating point instruction other than DFN, the exponent of the result was > 127.

Floating Underflow (bit 11) has been set.

No Divide (bit 12) has been set in an FDVX or FDVRX.

4　　Byte Interrupt — the processor is in a priority interrupt that interrupted a byte instruction after the processing of the pointer but before the processing of the byte. Hence if an ILDB or IDPB was interrupted, the pointer now points not to the last byte, but rather to the byte that should be handled upon the return to the interrupted program [§2.13].

5　　User — the processor is in user mode [§2.15].

Remember [§2.5], overflow is determined directly from the carries, not from the flags. The carry flags give meaningful information only if no more than one instruction that can set them occurs between clearing and reading them.

▲ 6   User In-out — even if the processor is in user mode, no instructions are illegal (but protection and relocation still apply) [§2.15].

11   Floating Underflow — in a floating point instruction other than DFN, the exponent of the result was < −128 and Overflow and Floating Overflow have been set.

12   No Divide — any of the following has set Overflow:

           In a DIVX the dividend was greater than or equal to the divisor.

           In an IDIVX the divisor was zero.

If normalized operands are used, only a zero divisor can cause floating division to fail.

           In an FDVX or FDVRX the divisor was zero, or the dividend fraction was greater than or equal to twice the divisor fraction in magnitude; in either case Floating Overflow has been set.

OVERFLOW											FLOATING UNDERFLOW					
	CARRY 0	CARRY 1			USER	USER IN-OUT					NO DIVIDE	0	0	0	0	0
0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17

Column labels: FLOATING OVERFLOW (at bit 3), BYTE INTERRUPT (at bit 4)

FLAG FORMAT, LEFT HALF OF PC WORD

The total time required is that listed plus the time for the instruction executed. If E addresses a fast memory location, the instruction executed takes .34 μs less than the time listed for it.

XCT   Execute   1.36 (1.47) μs

256		A	I	X	Y
0		8 9	12 13 14	17 18	35

The A portion of this instruction is ignored.

Execute the contents of location E as an instruction. Any instruction may be executed, including another XCT. If an XCT executes a skip instruction, the skip is relative to the location of the XCT (the first XCT if there are several in a chain). If an XCT executes a jump, program flow is altered as specified by the jump (no matter how many XCTs precede a jump instruction, when PC is saved it contains an address one greater than the location of the first XCT in the chain).

N is the number of leading 0s.

JFFO   Jump if Find First One   2.19 (2.30) + .20 (N mod 18) μs

243		A	I	X	Y
0		8 9	12 13 14	17 18	35

If AC contains zero, clear AC A+1 and go on to the next instruction in sequence.

If AC is not zero, count the number of leading 0s in it (0s to the left of the leftmost 1), and place the count in AC A+1. Take the next instruction

from location E and continue sequential operation from there.

In either case AC is unaffected, the original contents of AC $A+1$ are lost.

Note that when AC is negative, the second accumulator is cleared, just as it would be if AC were zero.

JFCL   Jump on Flag and Clear   1.36 (1.47) μs

2 5 5	F	I	X	Y

0   8 9   12 13 14   17 18   35

If any flag specified by F is set, clear it and take the next instruction from location E, continuing sequential operation from there. Bits 9-12 are programmed as follows.

Bit	Flag Selected by a 1
9	Overflow
10	Carry 0
11	Carry 1
12	Floating Overflow

To select one or a combination of these flags (which are among those described above) the programmer can specify the equivalent of an AC address that places 1s in the appropriate bits, but MACRO recognizes mnemonics for some of the 13-bit instruction codes (bits 0-12).

JFCL	JFCL 0,	No-op	25500
JOV	JFCL 10,	Jump on Overflow	25540
JCRY0	JFCL 4,	Jump on Carry 0	25520
JCRY1	JFCL 2,	Jump on Carry 1	25510
JCRY	JFCL 6,	Jump on Carry 0 or 1	25530
JFOV	JFCL 1,	Jump on Floating Overflow	25504

This instruction can be used simply to clear the selected flags by having the jump address point to the next consecutive location, as in

   JFCL   17,.+1

which clears all four flags without disrupting the normal program sequence. A JFCL that selects no flag is the fastest no-op as it neither fetches nor stores an operand, and bits 18-35 of the instruction word can be used to store information.

JSR   Jump to Subroutine   2.68 (2.79) μs ▲

2 6 4	A	I	X	Y

0   8 9   12 13 14   17 18   35

Interleaving memories saves .47 (.36) μs.
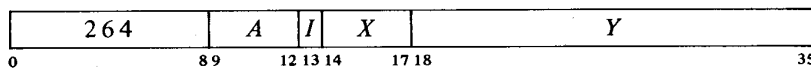
Place the current contents of the flags (as described above) in the left half of location E and the contents of PC in the right half (at this time PC contains an address one greater than the location of the JSR instruction). Take the next instruction from location $E+1$ and continue sequential operation from there. The flags are unaffected except Byte Interrupt, which is cleared.

While the processor is in user mode, if this instruction is executed as an in- ▲ terrupt instruction or in unrelocated 41 or 61, bit 5 of the PC word stored is 1 and the processor leaves user mode.

The A portion of this instruction is ignored.

JSP   **Jump and Save PC**   1.36 (1.47) μs

265	A	I	X	Y

0   89   12 13 14   17 18   35
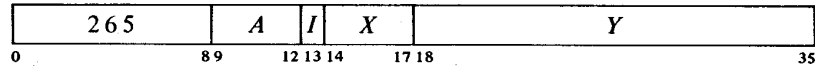
Place the current contents of the flags (as described above) in AC left and the contents of PC in AC right (at this time PC contains an address one greater than the location of the JSP instruction). Take the next instruction from location *E* and continue sequential operation from there. The flags are unaffected except Byte Interrupt, which is cleared.

▲   While the processor is in user mode, if this instruction is executed as an interrupt instruction or in unrelocated 41 or 61, bit 5 of the PC word stored is 1 and the processor leaves user mode.

JRST   **Jump and Restore**   1.36 (1.47) μs

254	F	I	X	Y

0   89   12 13 14   17 18   35
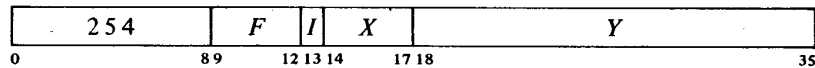
Perform the functions specified by *F*, then take the next instruction from location *E* and continue sequential operation from there. Bits 9–12 are programmed as follows.

Bit   *Function Produced by a 1*

9   Restore the channel on which the highest priority interrupt is currently being held [§2.13].

          Unless the User In-out flag is set, this function cannot be executed in a user program. Instead of restoring the channel, it stores its own instruction code, *F* and effective address *E* in bits 0–8, 9–12 and 18–35 respectively of unrelocated location 40 (clearing bits 13–17), and then executes the instruction contained in location 41, which is under control of the monitor [§2.15].

This is identical to UUO trapping [§2.10].

10   Halt the processor. When it stops, the MA lights on the console display an address one greater than that of the location containing the instruction that caused the halt, and PC displays the jump address (the location from which the next instruction will be taken if the operator causes the processor to resume operation without changing PC).

MA actually displays the address of the location that would have been executed next had the JRST been replaced by a no-op. So except for a JRST in a priority interrupt, MA points to the location one beyond that containing the instruction that caused the halt. This instruction is ordinarily the JRST or perhaps an XCT, but could even be a UUO.

          Unless the User In-out flag is set, this function cannot be executed in a user program. Instead of halting the processor, it stores its own instruction code, *F* and effective address *E* in Bits 0–8, 9–12 and 18–35 respectively of unrelocated location 40 (clearing bits 13–17), and then executes the instruction contained in location 41, which is under control of the monitor [§2.15].

11   Restore the flags listed above from the left half of the word in the last location referenced in the effective address calculation. Hence to restore flags requires that the JRST instruction use indexing or

indirect addressing.

Restoration of all but the user flags is directly according to the contents of the corresponding bits as given above: a flag is set by a 1 in the bit, cleared by a 0. A 1 in bit 5 sets User but a 0 has no effect, so the Monitor can restart a user program by restoring flags but the user cannot leave user mode by this method. A 0 in bit 6 clears User In-out, but a 1 sets it only if the JRST is being executed by the Monitor, *ie* if User is clear.

12 Enter user mode. The user program starts at relocated location *E*.

To produce one or a combination of these functions the programmer can specify the equivalent of an AC address that places 1s in the appropriate bits, but MACRO recognizes mnemonics for the most important 13-bit instruction codes (bits 0–12).

JRST	JRST 0,	Jump	25400
	JRST 10,	Jump and Restore Interrupt Channel	25440
HALT	JRST 4,	Halt	25420
JRSTF	JRST 2,	Jump and Restore Flags	25410
	JRST 1,	Jump to User Program	25404
JEN	JRST 12,	Jump and Enable	25450

In a JRSTF or JEN the flags are restored from bits 0–12 of the final word retrieved in the effective address calculation; hence any JRST with a 1 in bit 11 must use indirect addressing or indexing, which takes extra time. If the PC word was stored in AC (as in a JSP), a common procedure is to use AC to index a zero address (*eg*, JRSTF (AC)), so its right half becomes the effective (jump) address. If the PC word was stored in core (as in a JSR), one must address it indirectly (remember, bits 13–17 of the PC word are clear, so again its right half is the effective address). A JRSTF (AC) takes 1.64 (1.75) μs, a JRSTF @PCWORD takes 2.34 (2.56) μs.

CAUTION

Giving a JRSTF or JEN without indexing or indirect addressing restores the flags from the instruction code itself.

While the processor is in user mode, if this instruction is executed as an in- ▲ terrupt instruction or in unrelocated 41 or 61, bit 5 of the PC word stored is 1 and the processor leaves user mode.

JFCL is the only jump that can test any of the flags directly. In fact it is the only basic program control instruction that can do so — several of the flags can be tested as processor conditions by in-out instructions, but these are ordinarily illegal in user programs anyway. But JFCL can test only four

By manipulating the contents of the left half word used to restore the flags, the programmer can set them up in any desired way except that a user program cannot clear User or set User In-out. Setting Byte Interrupt prevents incrementing in the next ILDB or IDPB provided there is no intervening JSR, JSP or PUSHJ.

JEN completes an interrupt by restoring the channel and restoring the flags for the interrupted program. ▸

of the flags, and it saves no information for a subsequent return from a sub-routine. Hence it serves as a branch point for entry into either one of two main paths, which may or may not have a later point in common. *Eg*, it may test the carry flags simply to take appropriate action in a double precision fixed point routine.

JSR and JSP are regularly used to call subroutines. They are uncondi-tional, but the execution of such an instruction can be the result of a decision made by any conditional skip or jump. In the case of the flags, a basic overflow test and subroutine call can be made as follows.

> The fastest skip is CAIA.

```
JOV     .+2
JRST    .+2        ;Faster than skipping
JSR     OVRFLO     ;Jump over this if Overflow clear
.
.
.
```

If we wish to go to the DIVERR routine when No Divide is set, we must first read the flags into a test accumulator T and then use a test instruction.

```
JSP     T,.+1      ;Store flags but continue in sequence
TLNE    T,40       ;40 left selects bit 12
JSR     DIVERR     ;Skip this if No Divide clear
.
.
.
```

A subroutine called by a JSR must have its entry point reserved for the PC word. Hence it is nonreentrant: the JSR modifies memory so the subroutine cannot be shared with other programs. The JSP requires an accumulator, but it is faster and is convenient for argument passing. To return from a JSR-called subroutine one usually addresses the PC word indirectly, return-ing to the location following the JSR. But there are two ways to get back from a JSP. We can address the PC word indirectly with a JRST @AC (or JRSTF @AC if the flags must be restored), but we can also get it by addressing AC as an index register: JRST (AC). By using the second return method we can place N words of data for the subroutine immediately after the call, and return to the location following the data by giving a JRST N(AC).

Suppose we wish to call a print subroutine and supply the words from which the characters are to be taken. Our main program would contain the following:

```
JSP     T,PRINT    ;Put PC word in accumulator T
.                  ;Text inserted here by ASCIZ pseudo-
.                  ;instruction, which automatically
.                  ;places a zero (null) character at the
                   ;end
...                ;Next instruction here
```

The subroutine can use T as a byte pointer which already addresses the first word of data. For the print routine, characters are loaded into another accumulator CH.

```
PRINT:    HRLI    T,440700      ;Initialize left half of pointer
          ILDB    CH,T          ;Increment pointer and load byte
          JUMPE   CH,1(T)       ;Upon reaching zero character return
                                ;to one beyond last data word
          .                     ;Print routine
          .
          .
          JRST    PRINT+1       ;Get next character
```

JSA   **Jump and Save AC**   2.82 (2.93) μs

266	A	I	X	Y
0	8 9	12 13 14	17 18	35

Interleaving memories saves .47 (.36) μs.

Place AC in location E, the effective address E in AC left, and the contents of PC in AC right (at this time PC contains an address one greater than the location of the JSA instruction). Take the next instruction from location $E+1$ and continue sequential operation from there. The original contents of E are lost.

While the processor is in user mode, if this instruction is executed as an interrupt instruction or in unrelocated 41 or 61, bit 5 of the PC word stored is 1 and the processor leaves user mode.

JRA   **Jump and Restore AC**   2.92 (3.14) μs

267	A	I	X	Y
0	8 9	12 13 14	17 18	35

Place the contents of the location addressed by AC left into AC. Take the next instruction from location E and continue sequential operation from there.

A JSA combines advantages of the JSR and JSP. JSA does modify memory, but it saves PC in an accumulator without losing its previous contents (at a cost of not saving the flags). It is thus convenient for multiple-entry subroutines. In a subroutine called by a JSR, the returning JRST must refer to the (single) entry point. Since a JRA can retrieve the original PC by addressing AC as an index register, it is independent of any entry point without tying up an accumulator to the extent a JSP would.

The accumulator contents saved by a JSA are restored by a JRA paired with it despite intervening JSA-JRA pairs. Hence these instructions are especially useful for nesting subroutines, as shown by this example.

In FORTRAN IV, a CALL statement uses JSA with AC 16.

			;Main program
	JSA	17,S1	;Call to first subroutine (A)
S1:	0		;First subroutine starts here
	JSA	17,S2	;Call to second subroutine (B)
	JRA	17,(17)	;Return to A + 1 in main program
S2:	0		;Second subroutine starts here
	JSA	17,S3	;Call to third subroutine (C)
	JRA	17,(17)	;Return to B + 1 in first subroutine
S3:	0		;Third subroutine starts here
	JRA	17,(17)	;Return to C + 1 in second subroutine

To call the next deeper subroutine at any level, a JSA places E and PC in the left and right of AC 17, saves the previous contents of AC 17 in E (the first subroutine location), and jumps to $E + 1$. To return to the next higher level, a JRA restores the previous contents of AC 17 from the location addressed by AC 17 left (the first subroutine location) and jumps to the location addressed by AC 17 right (the location following the JSA in the higher subroutine). If N lines of data for the next subroutine follow a JSA, the return to the location following the data is made by giving a JRA 17,N(17).

Keeping instructions and the pushdown list in different memories saves .47 (.36) μs.

PUSHJ   **Push Down and Jump**   3.00 (3.11) μs

260	A	I	X	Y
0	8 9	12 13 14	17 18	35

Add 1000001_8 to AC to increment both halves by one and place the result back in AC. If the addition causes the count in AC left to reach zero, set the Pushdown Overflow flag. Then place the current contents of the flags (as described above) in the left half of the location now addressed by AC right and the contents of PC in the right half of that location (at this time PC contains an address one greater than the location of the PUSHJ instruction). Take the next instruction from location E and continue sequential operation from there.

The flags are unaffected except Byte Interrupt, which is cleared. The original contents of the location added to the list are lost.

▲   While the processor is in user mode, if this instruction is executed as an interrupt instruction or in unrelocated 41 or 61, bit 5 of the PC word stored is 1 and the processor leaves user mode.

POPJ   Pop Up and Jump   2.96 (3.18) μs
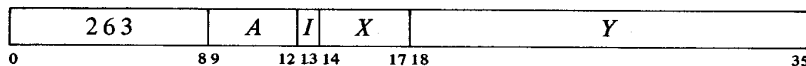
2 6 3	A	I	X	Y
0	8 9	12 13 14	17 18	35

Subtract 1000001_8 from AC to decrement both halves by one and place the result back in AC. If the subtraction causes the count in AC left to reach -1, set the Pushdown Overflow flag. Take the next instruction from the location addressed by the right half of the location that was addressed by AC right *prior* to the decrementing, and continue sequential operation from there.

The effective address E is ignored.

The address of the top item in the pushdown list is kept in the right half of the pointer in AC, and the program can keep a control count in the left half. The incrementing and decrementing of both halves of AC simultaneously is effected by adding and subtracting 1000001_8. Hence a count of -2 in AC left is increased to zero if $2^{18} - 1$ is incremented in AC right, and conversely, 1 in AC left is decreased to -1 if zero is decremented in AC right.

Since the pushdown list is independent of the subroutine called, PUSHJ-POPJ can be used like JSA-JRA for multiple entries. Moreover, ordering by level is inherent in the structure of a pushdown list [§2.2], so paired PUSHJ-POPJ instructions are excellent for nesting subroutines: there can be any number of subroutines at any level, each with more subroutines nested within it. Recursive subroutines are also possible.

Unlike JSA-JRA, the pushdown instructions tie up an accumulator, but the usual procedure is to keep both data and jump addresses in a single list so only one AC is required for the most complex pushdown operations. The programmer must keep track of whether a given entry in the list is data or a PC word; in other words, every item inserted by a PUSH should be removed by a POP, and every PUSHJ should be matched by a POPJ. If flag restoration is desired, the returning

POPJ   P,

can be replaced by

POP   P,AC
JRSTF   (AC)

which requires another accumulator. If the flags are not important, data may be stored in the left halves of the PC words in the stack, reducing the required pushdown depth.

By using the Pushdown Overflow flag and a control count in AC left, the programmer can set a limit to the size of the list by starting the count negative, or he can prevent the program from extracting more items than there are in the list by starting the count at zero, but he cannot do both at once. If only jump addresses are kept in the list, the first procedure limits the depth of nesting. A technique to catch extra POPJs is to put a PC word addressing an error routine at the bottom of the list.

2.10   UNIMPLEMENTED OPERATIONS

Many of the codes not assigned as specific instructions are executed as
unimplemented user operations, wherein the word given as an instruction is
trapped and must be interpreted by a routine included for this purpose by
the programmer. In time sharing, however, half of the codes are set aside for
user communication with the Monitor and are interpreted by it. Instructions
that are illegal in user mode also trap in this manner.

Unimplemented User Operation   2.33 (2.44) μs

000-077	A	I	X	Y
0	8 9	12 13 14	17 18	35
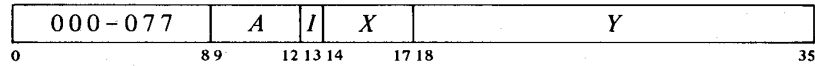
Store the instruction code, A and the effective address E in bits 0–8, 9–12
and 18–35 respectively of location 40; clear bits 13–17. Execute the
instruction contained in location 41. The original contents of location 40
are lost.

All of these codes are equivalent when they occur in the Monitor or when
time sharing is not in effect. But when a UUO appears in a user program, a
code in the range 001–037 uses relocated locations 40 and 41 (*ie* 40 and 41
in the user's block) and is thus entirely a part of and under control of the
user program. A code in the range 040–077 on the other hand uses
unrelocated 40 and 41, and the instruction in the latter location is under
control of the Monitor; these codes are thus specifically for user communica-
tion with the Monitor, which interprets them (refer to the Monitor manual
for the meanings of the various codes). The code 000 executes in the same
way as 040–077 but is not a standard communication code: it is included so
that control returns to the Monitor should a user program wipe itself out.

For a second processor connected to the same memory, the UUO trap is
locations 140-141 instead of 40-41.

The unimplemented operations also include the reserved (unassigned)
instruction codes 100–127, which execute like the Monitor-calling UUOs
but use unrelocated 60-61 (160-161 for a second processor); thus the
Monitor steps in when a user gives an incorrect code. The codes 130–177,
which are the floating point and byte manipulation instructions, are equiva-
lent to the unassigned codes if unimplemented, *ie* if the optional hardware
for them is not included. In this case all codes 100–177 trap to unrelocated
60-61. In general it is assumed that if software is available for floating point
and byte manipulation, the Monitor is responsible for calling the appropriate
routines.

2.11 PROGRAMMING EXAMPLES

Before continuing to input-output and related subjects, let us consider som simple programs that demonstrate the use of a variety of the instruction described thus far.

Suppose we wish to count the number of 1s in a word. We could of course check every bit in the word. But there is a quicker way if we remember that in any word and its twos complement the rightmost 1 is in the same position, both words are all 0s to the right of this 1, and no corresponding bits are the same to the left (the parts of both words at the left of the rightmost 1 are complements). Hence using the negative of a word as a mask for the word in a test instruction selects only the rightmost 1 for modification. The example uses three accumulators: the word being tested (which is lost) is in T, the count is kept in CNT, and the mask created in each step is stored in TEMP.

```
        MOVEI   CNT,0       ;Clear CNT
        MOVN    TEMP,T      ;Make mask to select rightmost 1
        TDZE    T,TEMP      ;Clear rightmost 1 in T
        AOJA    CNT,.-2     ;Increase count and jump back
        . . .               ;Skip to here if no 1s left in T
```

CNT is increased by one every time a 1 is deleted from T. After all 1s have been removed, the TDZE skips.

In the standard algorithm for converting a number N to its equivalent in base b, one performs the series of divisions

$$N/b \quad = \quad q_1 + r_1/b \qquad r_1 < b$$

$$q_1/b \quad = \quad q_2 + r_2/b \qquad r_2 < b$$

$$q_2/b \quad = \quad q_3 + r_3/b \qquad r_3 < b$$

$$\vdots$$

$$q_{n-1}/b \quad = \quad 0 + r_n/b \qquad r_n < b$$

The number in base b is then $r_n \ldots r_3 r_2 r_1$. Eg the octal equivalent of 61 decimal is 75:

$$61/8 \quad = \quad 7 + 5/8$$

$$7/8 \quad = \quad 0 + 7/8$$

The following decimal print routine converts a 36-bit positive integer in accumulator T to decimal and types it out. The contents of T and T + 1 are destroyed. The routine is called by a PUSHJ P, DECPNT where P is the pushdown pointer.

```
DECPNT:   IDIVI   T,12        ;12_8 = 10_10
          PUSH    P,T+1       ;Save remainder
          SKIPE   T           ;All digits formed?
          PUSHJ   P,DECPNT    ;No, compute next one
```

```
DECPN1:   POP    P,T        ;Yes, take out in opposite order
          ADDI   T,60       ;Convert to ASCII (60 is code for 0)
          JRST   TTYOUT     ;Type out
```

This routine repeats the division until it produces a zero quotient. Hence it suppresses leading zeros, but since it is executed at least once it outputs one "0" if the number is zero. The TTYOUT routine returns with a POPJ P, to DECPN1 until all digits are typed, then to the calling program.

Space can be saved in the pushdown stack by storing the computed digits in the left halves of the locations that contain the jump addresses. This is accomplished in the decimal print routine by making the following substitutions.

$$\text{PUSH} \quad \text{P,T+1} \quad \rightarrow \quad \text{HRLM} \quad \text{T+1,(P)}$$

$$\text{POP} \quad \text{P,T} \quad \rightarrow \quad \text{HLRZ} \quad \text{T,(P)}$$

The routine can handle a 36-bit unsigned integer if the IDIVI T,12 is replaced by

```
LSHC   T,-↑D35    ;Shift right 35 bits into T+1
LSH    T+1,-1     ;Vacate the T+1 sign bit
DIVI   T,12       ;Divide double length integer by 10
```

MACRO interprets a number following ↑D as decimal.

Many data processing situations involve searching for information in tables and lists of all kinds. Suppose we wish to find a particular item in a table beginning at location TAB and containing N items. Accumulator T contains the item. The right half of A is used to index through the table, while the left half keeps a control count to signal when a search is unsuccessful.

```
MOVSI   A,-N        ;Put -N, 0 in A
CAMN    T,TAB(A)    ;Skip if current item not the one
JRST    FOUND       ;Item found
AOBJN   A,.-2       ;Try next item until left count = 0
. . .               ;Item not in list
```

The location of the item (if found) is indicated by the number in the right half of A (its address is that quantity plus TAB). A slightly different procedure would be

```
HRLZI   A,-N
CAME    T,TAB(A)    ;Skip if current item is the one
AOBJN   A,.-1
JUMPL   A,FOUND     ;Jump if left count < 0
. . .               ;Item not found
```

Locations used for a list can be scattered throughout memory if data is kept in the left half of each location and the right half addresses the next location in the list. The final location is indicated by a zero right half. The following routine finds the last half word item in the list. It is entered at FIND with the first location in the list addressed by the right half of accumulator T. At the end the final item is in T right.

```
        MOVE    T,(T)           ;Move next item to T
FIND:   TRNE    T,777777        ;Skip if AC right = 0
        JRST    .-2
        HLRZS   T               ;Move final item to right
```

The following counts the length of the list in accumulator CNT.

```
        MOVEI   CNT,0           ;Clear CNT
        JUMPE   T,OUT           ;Jump out if T contains 0
        HRRZ    T,(T)           ;Get next address
        AOJA    CNT,.-2         ;Count and go back
```

Double Precision Floating Point. The following are straightforward routines for handling double precision floating point arithmetic [§2.6 *describes the floating point instructions*].

```
DFAD:   UFA     A+1,M+1         ;Sum of low parts to A+2
        FADL    A,M             ;Sum of high parts to A, A+1
        UFA     A+1,A+2         ;Add low part of high sum to A+2
        FADL    A,A+2           ;Add low sum to high sum
        POPJ    P,

DFSB:   DFN     A,A+1           ;Negate double length operand
        PUSHJ   P,DFAD          ;Call double floating add
        DFN     A,A+1           ;-(M - AC) = AC - M
        POPJ    P,

DFMP:   MOVEM   A,A+2           ;Copy high AC operand in A+2
        FMPR    A+2,M+1         ;One cross product to A+2
        FMPR    A+1,M           ;Other to A+1
        UFA     A+1,A+2         ;Add cross products into A+2
        FMPL    A,M             ;High product to A, A+1
        UFA     A+1,A+2         ;Add low part to cross sum in A+2
        FADL    A,A+2           ;Add low sum to high part of product
        POPJ    P,
```

A double precision division is of the form

$$\frac{A}{B} = \frac{a + c \times 2^{-27}}{b + d \times 2^{-27}}$$

Using the relationship

$$A/b = q + r \times 2^{-27}/b$$

where q and r are the quotient and remainder produced by FDVL, the following routine computes a double length quotient by the algorithm

$$\frac{A}{B} \cong q + \frac{(r - qd) \times 2^{-27}}{b}$$

which gives a result correct to the next-to-last bit in the low order half.

DFDV:	FDVL	A,M	;Get high part of quotient
	MOVN	A+2,A	;Copy negative of quotient in A+2
	FMPR	A+2,M+1	;Multiply by low part of divisor
	UFA	A+1,A+2	;Add remainder
	FDVR	A+2,M	;Divide sum by high part of divisor
	FADL	A,A+2	;Add result to original quotient
	POPJ	P,	

2.12   INPUT-OUTPUT

The input-output instructions govern all transfers of data to and from the peripheral equipment, and also perform many operations within the processor. An instruction in the in-out class is designated by 111 in bits 0–2, *ie* its left octal digit is 7. Bits 3–9 address the device that is to respond to the instruction. The format thus allows for 128 codes, two of which, 000 and 004 respectively, address the processor and priority interrupt, and are used for the console and time share hardware as well. A chart 'n Appendix A lists all devices for which codes have been assigned, and gives their mnemonics and DEC option numbers.

▲ Times are given for IO instructions when they occur alone. When two IO instructions are given consecutively, the second often takes longer (refer to the timing chart in Appendix C for details).

Bits 13–35 are the same as in all other instructions: they are the I, X, and Y parts, which are used to calculate an effective address, set of conditions, or mask to be used in the execution of the instruction. The remaining bits, 10–12, select one of the following eight IO instructions.

NOTE

All instructions described in the remainder of this manual are in-out instructions, which cannot be executed in user programs unless the User In-out flag is set. If an in-out instruction appears in a user program while User In-out is clear, it does not perform the functions given for it in the instruction description. Instead it stores its own instruction and device codes in bits 0–12 and its effective address E in bits 18–35 of unrelocated location 40 (clearing bits 13–17), and then executes the instruction contained in location 41. The latter location is under control of the Monitor [§2.15].

This is identical to UUO trapping [§2.10].

This user restriction will not be mentioned in the instruction descriptions, as it applies to *all* instructions from this point on.

E will always be regarded as being bits 18–35, even though it is actually placed on both halves of the bus and many devices receive the information from the left half.

CONO   **Conditions Out**   3.90 (4.01) μs

7		D		20	*I*	*X*			*Y*	
0	2 3			9 10	12 13 14	17 18				35
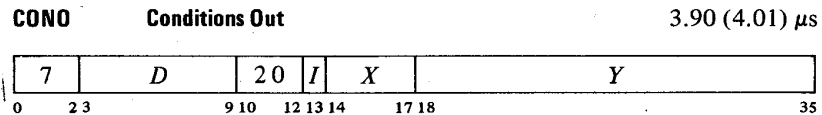
Set up device D with the effective initial conditions E. The number of condition bits in E that are actually used depends on the device.

CONI   **Conditions In**   . 4.87 (4.98) μs

7	D	24	I	X	Y
0 2 3		9 10 12 13 14 17 18			35
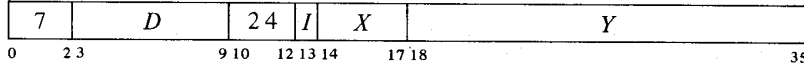
Read the input conditions from device D and store them in location E. The number of condition bits stored depends on the device; the remaining bits in location E are cleared.

Keeping instructions and operands in different memories saves .47 (.36) μs. Bringing conditions into fast memory saves .46 (.35) μs.

DATAO   **Data Out**   4.75 (4.97) μs

7	D	14	I	X	Y
0 2 3		9 10 12 13 14 17 18			35
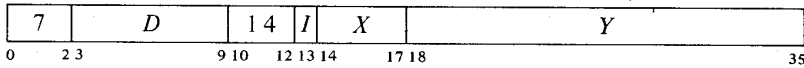
Send the contents of location E to the data buffer in device D, and perform whatever control operations are appropriate to the device.

The amount of data actually accepted by the device depends on the size of its buffer, its mode of operation, etc. The original contents of location E are unaffected.

Taking the output word from fast memory saves .34 μs.

DATAI   **Data In**   4.87 (4.98) μs

7	D	04	I	X	Y
0 2 3		9 10 12 13 14 17 18			35
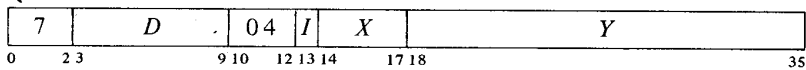
Move the contents of the data buffer in device D to location E, and perform whatever control operations are appropriate to the device.

The number of data bits stored depends on the size of the device buffer, its mode of operation, etc. Bits in location E that do not receive data are cleared.

Keeping instructions and operands in different memories saves .47 (.36) μs. Placing the input data in fast memory saves .46 (.35) μs.

CONSZ   **Conditions In and Skip if Zero**   4.11 (4.22) μs

7	D	30	I	X	Y
0 2 3		9 10 12 13 14 17 18			35
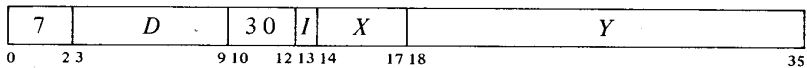
Test the input-conditions from device D against the effective mask E. If all condition bits selected by 1s in E are 0s, skip the next instruction in sequence.

If the device supplies more than 18 condition bits, only the right 18 are tested.

CONSO **Conditions In and Skip if One**   4.11 (4.22) μs

7	D	3 4	I	X	Y

0 2 3   9 10 12 13 14 17 18   35
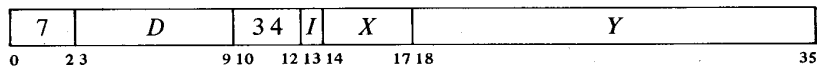
Test the input conditions from device *D* against the effective mask *E*. If any condition bit selected by a 1 in *E* is 1, skip the next instruction in sequence.

If the device supplies more than 18 condition bits, only the right 18 are tested.

Keeping the pointer in fast memory saves .43 (.34) μs.

▲ **BLKO** **Block Out**   6.49 (6.71) μs

7	D	1 0	I	X	Y

0 2 3   9 10 12 13 14 17 18   35
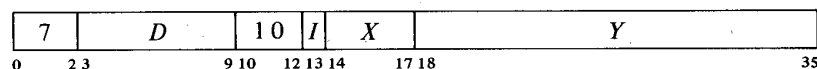
Keeping the pointer in fast memory saves .34 μs. Keeping the instruction and the data block in different memories saves .47 (.36) μs.

▲ **BLKI** **Block In**   6.49 (6.71) μs

7	D	0 0	I	X	Y

0 2 3   9 10 12 13 14 17 18   35
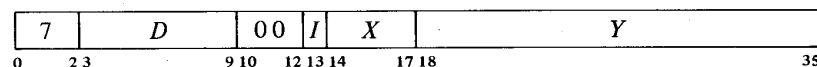
A block IO instruction is effectively a whole in-out data handling subroutine. It keeps track of the block location, transfers each data word, and determines when the block is finished.

Initially the left half of the pointer contains the negative of the number of words in the block, the right half contains an address one less than that of the first word in the block.

Add $1\,000001_8$ to a pointer in location *E* to increment both halves by one, and place the result back in *E*. Then perform a data IO instruction in the same direction as the block IO instruction, using the right half of the incremented pointer as the effective address. If the given instruction is a BLKO, perform a DATAO; if a BLKI, perform a DATAI.

The remaining actions taken by this instruction depend on whether it is executed as a priority interrupt instruction [§2.13].

◆ *Not as an Interrupt Instruction.* If the addition has caused the count in the left half of the pointer to reach zero, execute the next instruction in sequence. Otherwise skip the next instruction.

◆ *As an Interrupt Instruction.* If the addition has caused the count in the left half of the pointer to reach zero, execute the instruction in the second interrupt location for the channel. Otherwise dismiss the interrupt and return to the interrupted program.

The above eight instructions differ from one another in their total effect, but they are not all different with respect to any given device. A BLKO acts on a device in exactly the same way as a DATAO — the two differ only in counting and other operations carried out within the processor and memory. Similarly, no device can distinguish between a BLKI and a DATAI; and a device always supplies the same input conditions during a CONI, CONSZ or CONSO whether the program tests them or simply stores them.

Hence the eight instructions may be categorized as of four types, represented by the first four instructions described above. Moreover, a complete treatment of the programming of any device can be given in terms of these four instructions, two of which are for input and two for output. The four

exhaust the types of information transfer that occur in the IO system, at least three of which are applicable to any given device. Thus all instruction descriptions in the rest of this manual will be of the CONO, CONI, DATAO and DATAI instructions combined with the various device codes. The discussion of each device will present timing information pertinent to device operation, but no instruction times will be included as they are identical to those given above.

Every device requires initial conditions; these are sent by a CONO, which can supply up to eighteen bits of control information to the device control register. The program can determine the status of the device from up to thirty-six bits of input conditions that can be read by a CONI (but only the right eighteen can be tested by a CONSZ or CONSO). Some input bits simply reflect initial conditions sent by a previous CONO; others are set up by output conditions but are subject to subsequent adjustment by the device; and still others, such as status levels from a tape transport, have no direct connection with output conditions.

Data is moved in and out in characters of various sizes or in full 36-bit words. Each transfer between memory and a device data buffer requires a single DATAI or DATAO. Every device has a CONO and CONI, but it may have only one data instruction unless it is capable of both input and output. *Eg*, the paper tape reader has only a DATAI, the tape punch has only a DATAO, but the teletype has both. (A high speed device, such as a disc file, can be connected to the DF10 Data Channel, which in turn is connected directly to memory by a separate memory bus and handles data automatically. This eliminates the need for the program to give a DATAO or DATAI for each transfer.)

A Typical IO Device. Every device has a 7-bit device selection network, a priority interrupt assignment, and at least two flags, Busy and Done, or some equivalent. The selection network decodes bits 3–9 of the instruction so that only the addressed device responds to signals sent by the processor over the in-out bus. To use the device with the priority interrupt, the program must assign a channel to it. Then whenever an appropriate event occurs in the device, it requests an interrupt on the assigned channel.

The Busy and Done flags together denote the basic state of the device. When both are clear the device is idle. To place the device in operation, a CONO or DATAO sets Busy. If the device will be used for output, the program must give a DATAO that sends the first unit of data — a word or character depending on how the device handles information. When the device has processed a unit of data, it clears Busy and sets Done to indicate that it is ready to receive new data for output, or that it has data ready for input. In the former case the program would respond with a DATAO to send more data; in the latter, with a DATAI to bring in the data that is ready. If an interrupt channel has been assigned to the device, the setting of Done signals the program by requesting an interrupt; otherwise the program must keep testing Done to determine when the device is ready.

All devices function basically as described above even though the number of initial conditions varies considerably. Besides Busy and Done flags, the tape reader and punch have a Binary flag that determines the mode of operation of the device with respect to the data it processes — alphanumeric

The word "input" used without qualification always refers to the transfer of data from the peripheral equipment into the processor; "output" refers to the transfer in the opposite direction.

A DATAI that addresses an output-only device simply clears location E. DATAI PI, (code 70044) produces only this effect as the priority interrupt has no data for input. On the other hand a DATAO that addresses an input-only device is a no-op.

When the device code is undefined or the addressed device is not in the system, a DATAO, CONO or CONSO is a no-op, a CONSZ is an absolute skip, a DATAI or CONI clears location E.

Busy and Done both set is a meaningless situation.

Occasionally a device with a second code may use a DATAI or DATAO to transmit additional control or maintenance information.

or binary. The teletype has no binary flag, but it has two Busy flags and two Done flags — one pair for input, another for output. A complicated device, such as magnetic tape, may require two device codes to handle the large number of conditions associated with it. Initial conditions for a tape system include a transport address and an actual command the tape control is to perform; input conditions include error flags and transport status levels.

Most IO devices involve motion of some sort, usually mechanical (in a display only the electron beam moves). With respect to mechanical motion there are two types of devices, those that stay in motion and those that do not. Magnetic tape is an example of the former type. Here the device executes a command (such as read, write, space forward) and the done flag indicates when the entire operation is finished. A separate data flag signals each time the device is ready for the program to give a DATAI or DATAO, but the tape keeps moving until an entire record or file has been processed.

Paper tape, on the other hand, stops after each transfer, but the program need not give a new CONO every time. The reader logic is set up so that a DATAI not only reads the data, but also clears'Done and sets Busy. Hence if the instruction is given within a critical time, the tape moves continuously and only two CONOs are required for a whole series of transfers: one to start the tape, and one to stop it after the final DATAI.

Other devices operate in one or the other of these two ways but differ in various respects. The tape punch and teletype output are like the reader. Teletype input is initiated by the operator striking a key rather than by the program. The card reader reads an entire card on a single CONO, with a DATAI required for each column. The DECtape stays in motion, and the program must give a CONO to stop it or it will go all the way to the end zone.

Readin Mode

This mode of processor operation provides a means of placing information in memory without relying on a program already in memory or loading one word at a time manually. Its principal use is to read in a short loader program which is then used for loading other information. A loader program should ordinarily be used rather than readin mode, as a loader can check the validity of the information read.

Pressing the readin key on the console activates readin mode by starting the processor in a special hardware sequence that simulates a DATAI followed by a series of BLKI instructions, all of which address the device whose code is selected by the readin device switches on the small panel at the left of the paper tape reader. Various devices can be used, and for each there are special rules that must be followed. But the readin mode characteristics of any particular device are treated in the discussion of the device. Here we are concerned only with the general characteristics.

The information read is a block of data (such as a loader program) preceded by a pointer for the BLKI instructions. The left half of the pointer contains the negative of the number of words in the block, the right half contains an address one less than that of the location that is to receive the first word.

To read in, the operator must set up the device he is using, set its code into the readin device switches, and press the readin key. The processor places the device in operation, brings the first word (the pointer) into location 0, and then reads the data block, placing the words in the locations specified by the pointer. Data can be placed anywhere in memory (including fast memory) except in location 0. The operation affects none of memory except location 0 and the block area.

Upon completing the block, the processor halts only if the single instruction switch is on. Otherwise it leaves readin mode, and begins normal operation by executing the last word in the block as an instruction.

Console Data Transfers

Neither the processor nor the priority interrupt system require all four types of IO instructions, so the program can make use of their device codes for communicating with the console.

DATAI APR,   Data In, Console

70004	I	X	Y
0	12 13 14	17 18	35

Read the contents of the console data switches into location E.
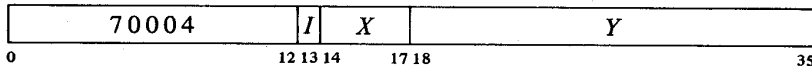
MACRO also recognizes the mnemonic RSW (Read Switches) as equivalent to DATAI APR,.

DATAO PI,   Data Out, Console

70054	I	X	Y
0	12 13 14	17 18	35
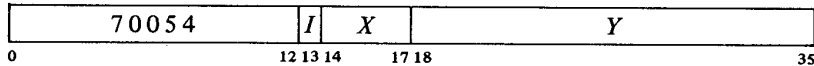
Unless the console MI program disable switch is on, display the contents of location E in the console memory indicators and turn on the triangular light beside the words PROGRAM DATA just above the indicators (turn off the light beside MEMORY DATA).

Once the indicators have been loaded by the program, no address condition selected from the console [§2.16] can load them until the operator turns on the MI program disable switch, executes a key function that references memory, or presses the reset key.

2.13 PRIORITY INTERRUPT

Most in-out devices must be serviced infrequently relative to the processor speed and only a small amount of processor time is required to service them, but they must be serviced within a short time after they request it. Failure to service within the specified time (which varies among devices) can often

result in loss of information and certainly results in operating the device below its maximum speed. The priority interrupt is designed with these considerations in mind, *ie* the use of interruptions in the current program sequence facilitates concurrent operation of the main program and a number of peripheral devices. The hardware also allows conditions internal to the processor to signal the program by requesting an interrupt.

Interrupt requests are handled through seven channels arranged in a priority chain, with assignment of devices to channels entirely at the discretion of the programmer. To assign a device to a channel, the program sends the number of the channel to the device control register as part of the conditions given by a CONO (usually bits 33–35). Channels are numbered 1–7, with 1 having the highest priority; a zero assignment disconnects the device from the interrupt channels altogether. Any number of devices can be connected to a single channel, and some can be connected to two channels (*eg* a device may signal that data is ready on one channel, that an error has occurred on another).

Interrupt Requests. When a device requires service it sends an interrupt request signal over the in-out bus to its assigned channel in the processor. If the channel is on, the processor accepts the request at the next memory access unless the processor is either starting an interrupt on any channel or holding an interrupt on the same channel. The request signal is a level, so it remains on the bus until turned off by the program (CONO, DATAO or DATAI). Thus if a request is not accepted because of the conditions given above, it will be accepted when those conditions no longer hold. A single channel will shut out all others of lower priority if every time its service routine dismisses the interrupt, a device assigned to it is already waiting with another request. The program can usually trigger a request from a device but delay its acceptance by turning on the channel later.

Starting an Interrupt. After a request is accepted the channel must wait for the interrupt to start. No interrupts can be started unless the priority interrupt system is active. Furthermore, the processor cannot start an interrupt if it is already holding an interrupt on a channel with priority higher than those on which requests have been accepted (in other words if the current program is a higher priority interrupt routine). If there is a higher priority channel waiting, the processor stops the current program to start an interrupt on the waiting channel that has highest priority. The interrupt starts following the retrieval of an instruction, following the retrieval of an address word in an effective address calculation (including the second calculation using the pointer in a byte instruction), or following a transfer in a BLT. When an interrupt starts, PC points to the interrupted instruction, so that a correct return can later be made to the interrupted program.

Two memory locations are assigned to each channel: unrelocated locations $40 + 2N$ and $41 + 2N$, where N is the channel number. Channel 1 uses locations 42 and 43, channel 2 uses 44 and 45, and so on to channel 7 which uses 56 and 57. The processor starts an interrupt on channel N by executing the instruction in location $40 + 2N$.

An instruction executed by the interrupt hardware in response to an interrupt request is referred to elsewhere in this manual as being executed "as an interrupt instruction". Some instructions, when so executed, perform

different functions than they do when executed in other circumstances. And the difference is not due merely to being executed in an interrupt location or in response (by the program) to an interrupt. To be an interrupt instruction, an instruction must be executed by the interrupt hardware, in location $40 + 2N$ or $41 + 2N$, because of a request on channel N. §2.12 describes the two ways a BLKO is performed. If a BLKO is contained in an interrupt routine called by a JSR, it is *not* executed "as an interrupt instruction" even if the routine is stored within the interrupt locations. There are two categories of interrupt instructions.

◆ *Non-IO Instructions.* After executing a non-IO interrupt instruction, the processor holds an interrupt on the channel and returns control to PC. Hence the instruction is usually a jump to a service routine. If the processor is in user mode and the interrupt instruction is a JSR, JSP, PUSHJ, JSA or JRST, the processor leaves user mode (the Monitor thus handles all interrupt routines [§2.15]).

If the interrupt instruction is not a jump, the processor continues the interrupted program while holding an interrupt — in other words it now treats the interrupted program as an interrupt routine. *Eg* the instruction might just move a word to a particular location. Such procedures are usually reserved for maintainence routines or very sophisticated programs.

◆ *Block or Data IO Instructions.* One or the other of two actions can result from executing one of these as an interrupt instruction.

If the instruction in $40 + 2N$ is a BLKI or BLKO and the block is not finished (*ie* the count does not cause the left half of the pointer to reach zero), the processor holds and immediately dismisses an interrupt on the channel, and returns to the interrupted program. The same action results if the instruction is a DATAI or DATAO.

If the instruction in $40 + 2N$ is a BLKI or BLKO and the count does reach zero, the processor continues to start the interrupt by executing the instruction in location $41 + 2N$. This *cannot* be an IO instruction and the actions that result from its execution as an interrupt instruction are those given above for non-IO instructions.

CAUTION

The execution, as an interrupt instruction, of a CONO, CONI, CONSO or CONSZ in location $40 + 2N$ or *any* IO instruction in location $41 + 2N$ hangs up the processor.

Dismissing an Interrupt. Automatic dismissal of an interrupt occurs only in a DATAI or DATAO, or in a BLKI or BLKO with an incomplete block. Following any non-IO interrupt instruction, the processor holds an interrupt until the program dismisses it, even if the interrupt routine is itself interrupted by a higher priority channel. Thus interrupts can be held on a number of channels simultaneously, but from the time an interrupt is started until it is dismissed, no interrupt can be started on that channel or any channel of lower priority (requests, however, can be accepted on lower priority channels).

A routine dismisses the interrupt by using a JEN (JRST 12,) to return to the interrupted program (the interrupt system must be active when the JEN is given). This instruction restores the channel on which the interrupt is being held, so it can again accept requests, and interrupts can be started on it and lower priority channels. JEN also restores the flags, whose states were saved in the left half of the PC word if the routine was called by a JSR, JSP, or PUSHJ [§2.9]. If flag restoration is not desired, a JRST 10, can be used instead.

CAUTION

> An interrupt routine must dismiss the interrupt when it returns to the interrupted program, or its channel and all channels of lower priority·will be disabled, and the processor will treat the new program as a continuation of the interrupt routine.

Priority Interrupt Conditions. The program can control the priority interrupt system by means of condition IO instructions. The device code is 004, mnemonic PI.

CONO PI,   Conditions Out, Priority Interrupt

70060	I	X	Y
0	12 13 14	17 18	35

Perform the functions specified by E as shown (a 1 in a bit produces the indicated function, a 0 has no effect).

							INITIATE INTERRUPT ON		DEACTIVATE PI		ACTIVATE PI						
CLEAR POWER FAILURE FLAG	CLEAR PARITY ERROR FLAG	DISABLE PARITY ERROR INTERRUPT	ENABLE		CLEAR PI SYSTEM		TURN ON	TURN OFF			SELECT CHANNELS FOR BITS 24, 25, 26						
						SELECTED CHANNELS					1	2	3	4	5	6	7
18	19	20	21	22	23	24	25	26	27	28	29	30	31	32	33	34	35

Notes.

Bits 18–21 are actually for processor conditions [§2.14].

20   Prevent the setting of the Parity Error flag from requesting an interrupt on the channel assigned to the processor.

21   Enable the setting of the Parity Error flag to request an interrupt on the channel assigned to the processor.

23   Deactivate the priority interrupt system, turn off all channels, eliminate all interrupt requests that have already been accepted but are still waiting, and dismiss all interrupts that are currently being held..

24   Request interrupts on channels selected by 1s in bits 29–35, and force the processor to accept them even on channels that are off.

A request is lost if it is made by this means to a channel on which an interrupt is already being held.

25   Turn on the channels selected by 1s in bits 29–35 so interrupt requests can be accepted on them.

26   Turn off the channels selected by 1s in bits 29–35, so interrupt requests cannot be accepted on them unless made by a CONO PI, with a 1 in bit 24.

27   Deactivate the priority interrupt system. The processor can then still accept requests, but it can neither start nor dismiss an interrupt.

28   Activate the priority interrupt system so the processor can accept requests and can start, hold and dismiss interrupts.

CONI PI,   **Conditions In, Priority Interrupt**

70064	I	X	Y
0	12 13 14	17 18	35

Read the status of the priority interrupt (and several bits of processor conditions) into the right half of location E as shown.

POWER FAILURE	PARITY ERROR	PARITY ERROR INTERRUPT ENABLED	INTERRUPT IN PROGRESS ON CHANNELS							PI ACTIVE	CHANNELS ON						
			1	2	3	4	5	6	7		1	2	3	4	5	6	7
18	19	20	21	22	23	24	25	26	27	28	29	30	31	32	33	34	35

Notes.

18   Ac power has failed. The program should save PC, the flags and fast memory in core, and halt the processor.

The setting of this flag requests an interrupt on the channel assigned to the processor. If the flag remains set for 5 ms, the processor is cleared.

19   A word with even parity has been read from core memory. If bit 20 is set, the setting of the Parity Error flag requests an interrupt on the channel assigned to the processor.

28   The priority interrupt system is active.

Channels that are on are indicated by 1s in bits 29–35; 1s in bits 21–27 indicate channels on which interrupts are currently being held.

Note that bits 18–20 actually read processor status conditions [§2.14].

Timing. The time a device must wait for an interrupt to start depends on the number of channels in use, and how long the service routines are for devices on higher priority channels. If only one device is using interrupts,

it need never wait longer than the time required for the processor to finish
the instruction that is being performed when the request is made. The
maximum time can be considered to be about 15 μs for FDVL, but a ridicu-
lously long shift could take over 35 μs.

Special Considerations. On a return to an interrupted program, the proc-
essor always starts the interrupted instruction over from the beginning. This
causes special problems in a BLT and in byte manipulation.

An interrupt can start following any transfer in a BLT. When one does,
the BLT puts the pointer (which has counted off the number of transfers
already made) back in AC. Then when the instruction is restarted following
the interrupt, it actually starts with the next transfer. This means that if
interrupts are in use, the programmer cannot use the accumulator that holds
the pointer as an index register in the same BLT, he cannot have the BLT
load AC except by the final transfer, and he cannot expect AC to be the
same after the instruction as it was before.

An interrupt can also start in the second effective address calculation in a
two-part byte instruction. When this happens, Byte Interrupt is set. This
flag is saved as bit 4 of a PC word, and if it is restored by the interrupt
routine when the interrupt is dismissed, it prevents a restarted ILDB or
IDPB from incrementing the pointer a second time. This means that the
interrupt routine must check the flag before using the same pointer, as it
now points to the next byte. Giving an ILDB or IDPB would skip a byte.
And if the routine restores the flag, the interrupted ILDB or IDPB would
process the same byte the routine did.

Programming Suggestions. The Monitor handles all interrupts for user
programs. Even if the User In-out flag is set, a user program generally cannot
reference the interrupt locations to set them up. Procedures for informing
the Monitor of the interrupt requirements of a user program are discussed in
the Monitor manual.

For those who do program priority interrupt routines, there are several
rules to remember.

♦ No requests can be accepted, not even on higher priority channels, while
a break is starting. Therefore do not use lengthy effective address calcula-
tions in interrupt instructions.

♦ The interrupt instruction that calls the routine must save PC if there is to
be a return to the interrupted program. Generally a JSR is used as it saves
both PC and the flags, and it uses no accumulator.

♦ The principal function of an interrupt routine is to respond to the situa-
tion that caused the interrupt. *Eg* computations that can be performed
outside the routine should not be included within it.

♦ The routine must dismiss the interrupt (with a JEN) when returning to the
interrupted program. The flags should be restored.

2.14 PROCESSOR CONDITIONS

There are a number of internal conditions that can signal the program by
requesting an interrupt on a channel assigned to the processor. Flags for

power failure and parity error are handled by the condition IO instructions that address the priority interrupt system [§2.13]. The remaining flags are handled by condition instructions that address the processor. Its device code is 000, mnemonic APR or CPA.

CONO APR,   Conditions Out, Arithmetic Processor

70020	I	X	Y
0	12 13 14	17 18	35

Perform the functions specified by E as shown (a 1 in a bit produces the indicated function, a 0 has no effect).

18	19	20	21	22	23	24	25	26	27	28	29	30	31	32	33	34	35
CLEAR PUSHDOWN OVERFLOW	CLEAR ALL IN-OUT DEVICES		CLEAR ADDRESS BREAK FLAG	CLEAR MEMORY PROTECTION FLAG	CLEAR NONEXISTENT MEMORY FLAG	DISABLE CLOCK INTERRUPT	ENABLE CLOCK INTERRUPT	CLEAR CLOCK FLAG	DISABLE FLOATING OVERFLOW INTERRUPT	ENABLE FLOATING OVERFLOW INTERRUPT	CLEAR FLOATING OVERFLOW	DISABLE OVERFLOW INTERRUPT	ENABLE OVERFLOW INTERRUPT	CLEAR OVERFLOW	PRIORITY INTERRUPT ASSIGNMENT		

Notes.

Enabling a particular flag to interrupt means that henceforth the setting of the flag will request an interrupt on the channel assigned (by bits 33–35) to the processor. Disabling prevents the flag from triggering a request.

A 1 in bit 19 produces the IO reset signal, which clears the control logic in all of the peripheral equipment (but affects neither the priority interrupt system, nor the processor flags cleared by this instruction or CONO PI,).

CONI APR,   Conditions In, Arithmetic Processor

70024	I	X	Y
0	12 13 14	17 18	35
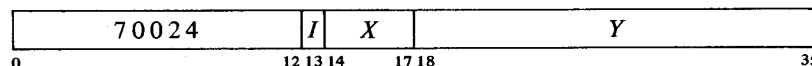
Read the status of the processor into the right half of location E as shown (all interrupt requests are made on the channel assigned to the processor).

18	19	20	21	22	23	24	25	26	27	28	29	30	31	32	33	34	35
PUSHDOWN OVERFLOW		USER IN-OUT	ADDRESS BREAK	MEMORY PROTECTION FLAG	NONEXISTENT MEMORY	CLOCK INTERRUPT ENABLED		CLOCK FLAG		FLOATING OVERFLOW INTERRUPT ENABLED	FLOATING OVERFLOW	TRAP OFFSET	OVERFLOW INTERRUPT ENABLED	OVERFLOW	PRIORITY INTERRUPT ASSIGNMENT		

Notes.

19   Pushdown Overflow — in a PUSH or PUSHJ the count in AC left
     reached zero; or in a POP or POPJ the count reached −1. The setting
     of this flag requests an interrupt.

▲ 20 User In-out — even if the processor is in user mode, no instructions
      are illegal (but protection and relocation still apply) [§2.15].

▲ PC bears no relation to the   21   Address Break — while the console address break switch was on, the
break if the access was re-   processor requested access to the memory location specified by the
quested for a console key   address switches and the memory reference was for the purpose
function.   selected by the address condition switches as follows:

          The instruction switch was on and access was for retrieval of an
          instruction (including an instruction executed by an XCT or con-
          tained in an interrupt location or a trap for an unimplemented
          operation) or an address word in an effective address calculation.

          The data fetch switch was on and access was for retrieval of an
          operand (other than in an XCT).

          The write switch was on and access was for writing a word in
          memory.

      The setting of this flag requests an interrupt, at which time PC points
      to the instruction that was being executed or to the one following it.

▲ This flag can also be set by ▲ 22 Memory Protection — a user program attempted to access a memory
an instruction executed from   location outside of its area or to write in a write-protected part of its
the console while the USER   area and the user instruction was terminated at that time. The setting
MODE light is on, in which   of this flag requests an interrupt, at which time PC points either to
case PC bears no relation to   the instruction that caused the violation or to the one following it.
the violation.

▲ PC bears no relation to the   23   Nonexistent Memory — the processor attempted to access a memory
unanswered reference if the   that did not respond within 100 μs. The setting of this flag requests
attempted access originated   an interrupt, at which time PC points either to the instruction con-
from a console key function.   taining the unanswered reference or to the one following it.

                                26   Clock — this flag is set at the ac power line frequency and can thus
                                     be used for low resolution timing (the clock has high long term
                                     accuracy). If bit 25 is set, the setting of the Clock flag requests an
                                     interrupt.

                                29   Floating Overflow — this is one of the flags saved in a PC word, and
                                     the conditions that set it are given at the beginning of §2.9. If bit 28
                                     is set, the setting of Floating Overflow requests an interrupt, at which
                                     time PC points to the instruction following that in which the over-
                                     flow occurred.

                                30   Trap Offset — the processor is using locations 140–161 for unimple-
                                     mented operation traps and interrupt locations.

                                32   Overflow — this is one of the flags saved in a PC word, and the condi-
                                     tions that set it are given at the beginning of §2.9. If bit 31 is set,
                                     the setting of Overflow requests an interrupt, at which time PC
                                     points to the instruction following that in which the overflow
                                     occurred.

2.15   TIME SHARING

Without time sharing the system has a single user and the program has no restrictions except those inherent in the hardware: the programmer must stay within the memory capacity, observe the restrictions placed on the use of certain memory locations by the hardware [§1.3], and observe the restrictions on interrupt instructions. Optional hardware can restrict processor operation to permit time sharing by a number of programs. Each user program is run with the processor in user mode, in which the program must operate within an assigned area in core and certain operations may be illegal. A program that runs unrestricted — the Monitor — is responsible for scheduling user programs, servicing interrupts, handling input-output needs, and taking action when control is returned to it from a user program.

Every user is assigned a core area and the rest of core is protected from him — he cannot gain access to the protected area for either storage or retrieval of information. The assigned area is divided into two parts. The low part is unique to a given user and can be used for any purpose. The high part may be for a single user, or it may be shared by several users. The Monitor can write-protect the high part so that the user cannot alter its contents, *ie* he cannot write anything in it. The Monitor would do this when the high part is to be a pure procedure to be used reentrantly by several users. One high pure segment may be used with any number of low impure segments. The user can request that the Monitor write-protect the high part of a single program, *eg* in order to debug a reentrant program. All users write programs beginning at address 0 for the low part, and beginning usually at 400000 for the high part. The programmed addresses are retained in the object program but are relocated by the hardware to the physical area assigned to the user as each access is made while the program is running.

The size and position of the user area are defined by specifying protection and relocation addresses for the low and high blocks. The protection address determines the maximum address the user can give; any address larger than the maximum is illegal. The relocation address is the address, as seen by the Monitor and the hardware, of the first location in the block. The Monitor defines these addresses by loading four 8-bit registers, each of which corresponds to the left eight bits (18–25) of an address whose right ten bits are all 0.

To determine whether an address is legal its left eight bits are compared with the appropriate protection register, so the maximum user address consists of the register contents in its left eight bits, 1777 in its right ten bits (*ie* it is equal to the protection address plus 1777). Since the set of all addresses begins at zero, a block is always an integral multiple of 1024_{10} (2000_8) locations. Relocation is accomplished simply by adding the contents of the appropriate relocation register to the user address, so the first address in a block is a multiple of 2000. The relative user and relocated address configurations are therefore as illustrated here, where P_l, R_l, P_h and R_h are respectively the protection and relocation addresses for the low and high parts as derived from the 8-bit registers loaded by the Monitor. If the low part is larger than 128K locations, *ie* more than half the maximum memory capacity ($P_l > 400000$), the high part starts at the first location after the low

Note that the relocated low part is actually in two sections with the larger beginning at $R_l + 20$. This is because addresses 0-17 are not relocated, all users having access to the accumulators. The Monitor uses the first sixteen locations in the low user block to store the user's accumulators when his program is not running.

Some systems have only the low pair of protection and relocation registers. In this case the user program is always nonreentrant and the assigned area comprises only the low part.

USER ADDRESSES
BEFORE RELOCATION

TYPICAL PHYSICAL ADDRESS
CONFIGURATION AFTER RELOCATION
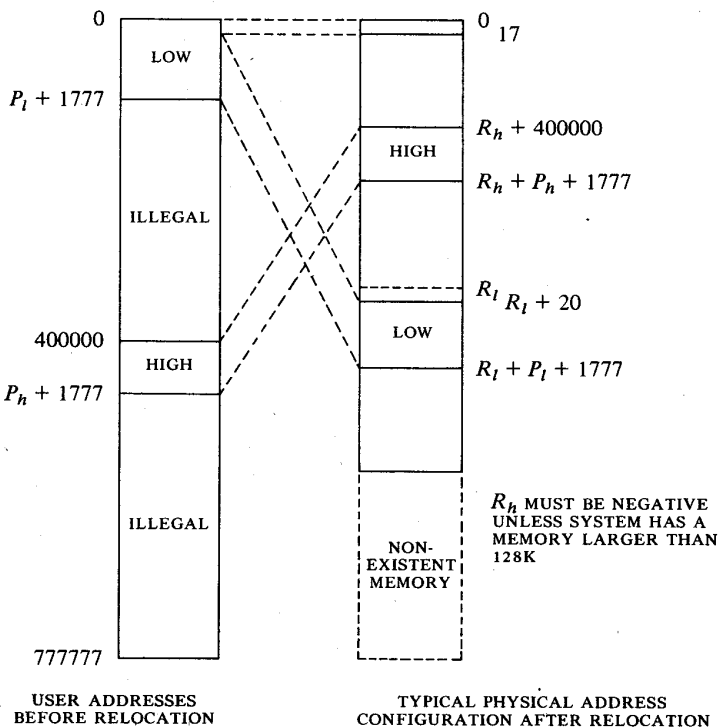
part (at location $P_l + 2000$). The high part is limited to 128K. If the Monitor defines two parts but does not write-protect the high part, the user has a two-part nonreentrant program.

If the user attempts to access a location outside of his assigned area, or if the high part is write-protected and he attempts to alter its contents, the current instruction terminates immediately, the Memory Protection flag is set (status bit 22 read by CONI APR,), and an interrupt is requested on the channel assigned to the processor [§2.14].

User Programming. The user must observe the following rules when programming on a time shared basis. [*Refer to the Monitor manual for further information including use of the Monitor for input-output.*]

♦ Use addresses only within the assigned blocks for all purposes – retrieval of instructions, retrieval of addresses, storage or retrieval of operands. The low part contains locations with addresses from 0 to the maximum; the high part contains from the greater of 400000 or $P_l + 2000$ to the maximum. Either part can address the other.

The user can actually write any size program: the Monitor will assign enough core for his needs. Basically the user must write a sensible program; if he uses absolute addresses scattered all over memory his program cannot be run on a time shared basis with others.

♦ If the high part is write-protected, do not attempt to store anything in it. In particular do not execute a JSR or JSA into the high part.

♦ Use instruction codes 000 and 040-127 only in the manner prescribed in the Monitor manual.

These instructions are illegal unless User In-out is set.

♦ Unless User In-out is set do not give any IO instruction, HALT (JRST 4,) or JEN (JRST 12, (specifically JRST 10,)). The program can determine if User In-out is set by examining bit 6 of the PC word stored by JSR, JSP or

PUSHJ.

The user can give a JRSTF (JRST 2,) but a 0 in bit 5 of the PC word does not clear User (a program cannot leave user mode this way); and a 1 in bit 6 does not set User In-out, so the user cannot void any of the restrictions himself. Note that a 0 in bit 6 will clear User In-out, so a user can discard his own special privileges.

UUOs 001–037 execute normally and are relocated to addresses 40 and 41 in the low block [§2.10].

Monitor Programming. The Monitor must assign the core area for each user program, set up trap and interrupt locations, specify whether the user can give IO instructions, transfer control to the user program, and respond appropriately when an interrupt occurs or an instruction is executed in unrelocated 41 or 61.

Core assignment is made by this instruction.

DATAO APR,   Data Out, Arithmetic Processor

70014	I	X	Y
0	12 13 14	17 18	35
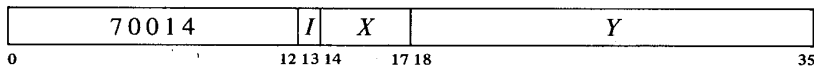
Load the protection and relocation registers from the contents of location E as shown, where P_l, P_h, R_l and R_h are the protection and relocation

$P_{l\,18-25}$		$P_{h\,18-25}$	P	$R_{l\,18-25}$		$R_{h\,18-25}$
0	7 8 9	16 17 18	25 26 27	34 35
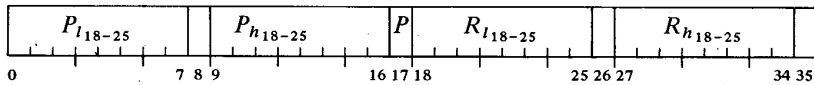
addresses defined above. If write-protect bit P (bit 17) is 1, do not allow the user to write in the high part of his area.

For a two part nonreentrant program, set $P = 0$. For a one-part nonreentrant program, make $P_h \leq P_l$. If the hardware has only one set of protection and relocation registers, the user area is defined by P_l and R_l, the rest of the word is ignored.

-Giving a JRSTF with a 1 in bit 6 of the PC word allows the user to handle his own input-output. The Monitor can also transfer control to the user with this instruction by programming a 1 in bit 5 of the PC word, or it may jump to the user program with a JRST 1, which automatically sets User. The set state of this flag implements the user restrictions.

While User is set, certain instructions are not part of the user program and are therefore completely unrestricted, namely those executed in the interrupt locations (which are not relocated) and in unrelocated trap locations 41 and 61. Illegal instructions and UUO codes 000 and 040–077 are trapped in unrelocated 40; codes 100–127 are trapped in unrelocated 60. BLKI and BLKO can be used in the even interrupt locations, and if there is no overflow, the processor returns to the interrupted user program. JSR should ordinarily be used in the remaining even interrupt locations, in odd interrupt locations following block IO instructions, and in 41 and 61. The JSR clears User and should jump to the Monitor. JSP, PUSHJ, JSA and JRST are acceptable in that they clear User, but the first two require an accumulator

(all accumulators should be available to the user) and the latter two do not save the flags.

After taking appropriate action, the Monitor can return to the user program with a JRSTF or JEN that restores the flags including User and User In-out.
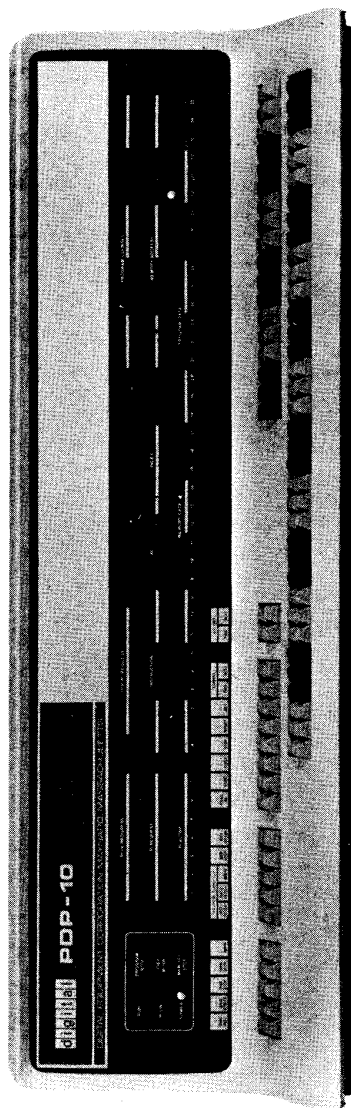
2.16   OPERATION

Most of the controls and indicators used for normal operation of the processor and for program debugging are located on the console operator panel shown here. The indicators are on the vertical part of the panel; in front of them are two rows of two-position keys and switches (keys are momentary contact, switches are alternate action). A key or switch is on or represents a 1 when the front part is down.

The thirty-six switches in the front row and the eighteen at the right in the back row are respectively the data and address switches through which the operator can supply words and addresses for the program and for use in conjunction with the operating keys and switches. The correspondence of switches to bit positions is indicated by the numbers at the bottom row of lights. At the left end of the back row are ten operating switches, which supply continuous control levels to the processor. At their right are ten operating keys, which initiate or terminate operations in the processor. The names of the operating keys and switches appear on the vertical part of the panel below the lights.

Also of interest to the operator is the small panel shown opposite, which is located above the main panel at the left of the tape reader. The upper section of this panel contains a total hours meter and the margin-check controls. The lower section contains the power switch, speed controls for slowing down the program, switches to select the device for readin mode (lower part in represents a 1), and four additional operating switches. The normal position for these last four is with the upper part in; in this position FM ENB (fast memory enable) is on, the others are all off.

Indicators

When any indicator is lit the associated flipflop is 1 or the associated function is true. Some indicators display useful information while the processor is running, but many change too frequently and can be discussed only in terms of the information they display when the processor is stopped. The program can stop the processor only at the completion of the HALT instruction; the operator can stop it at the end of

every instruction or memory reference, or for maintenance purposes, after every step in any operation that uses the shift counter (shifting, multiplication, division, byte manipulation).

Of the long rows of lights at the right on the operator panel, the top row displays the contents of PC, the middle row displays the instruction being executed or just completed, and the bottom row are the memory indicators. The right half of the middle row displays MA, the left half displays IR [see page 1-2]. In an IO instruction the left three instruction lights are on, the remaining instruction lights and the left AC light are the device code, and the remaining AC lights complete the instruction code. The I, index and MA lights change with each indirect reference in an effective address calculation; at the end of an instruction I is always off.

Above the memory indicators appear two pairs of words, PROGRAM DATA and MEMORY DATA. If the triangular light beside the former pair is on, the indicators display a word supplied by a DATAO PI,; if any other data is displayed the light beside MEMORY DATA is on instead. While the processor is running the physical addresses used for memory reference (the relocated address whenever relocation is in effect) are compared with the contents of the address switches. Whenever the two are equal the contents of the addressed location are displayed in the memory indicators. However, once the program loads the indicators, they can be changed only by the program until the operator turns on the MI program disable switch, executes a key function that references memory, or presses the reset key (see below).

The four sets of seven lights at the left display the state of the priority interrupt channels [see pages 2-74 and 2-75]. The PI ACTIVE lights indicate which channels are on. The IOB PI REQUEST lights indicate which channels are receiving request signals over the in-out bus; the PI REQUEST lights indicate channels on which the processor has accepted requests. Except in the case of a program-initiated interrupt, a REQUEST light can go on only if the corresponding ACTIVE light is on. The PI IN PROGRESS lights indicate channels on which interrupts are currently being held; the channel that is actually being serviced is the lowest-numbered one whose light is on. When a PROGRESS light goes on, the corresponding REQUEST goes off and cannot go on again until PROGRESS goes off when the interrupt is dismissed.

At the left end of the panel are a power light and these control indicators.

Above: Margin Check and Maintenance Panel
Opposite: Console Operator Panel

NOTE: If a REQUEST light stays on indefinitely with the associated PROGRESS light off and PC is static, check the PI CYC light on the indicator panel at the top of bay 2. If it is on, a faulty program has hung up the processor. Press STOP.

RUN
The processor is in normal operation with one instruction following another. When the light goes off, the processor stops.

PI ON
The priority interrupt system is active so interrupts can be started (this corresponds to CONI PI, bit 28).

USER MODE
The processor is in user mode (this corresponds to bit 5 of a PC word).

If RUN and PROGRAM STOP are both on, the processor is probably in an indirect address loop. Press STOP.

PROGRAM STOP
IR now contains a HALT instruction. If RUN is off, MA displays an address one greater than that of the location containing the instruction that caused the halt, and PC displays the jump address (the location from which the next instruction will be taken if the operator presses the continue key).

MEMORY STOP
The processor has stopped at a memory reference. This can be due to single cycle operation, satisfaction of an address condition selected at the console, reference to a nonexistent memory location, or detection of a parity error.

The remaining processor lights are on the indicator panels at the tops of the bays [*illustrated on page C8*]. Bay 2 displays AR, BR and MQ, the output of the AR adder, and the parity buffer which receives every word transmitted over the memory bus. The RL and PR lights at the lower right display the relocation and protection registers for the low part of the area assigned to a user program and the left eight bits of the relocated address for that part. The remaining lights are for maintenance.

The upper four rows on the bay 1 panel include the indicators for reader, punch and teletype, which are described in Chapter 3. The bottom row displays the information on the data lines in the IO bus. The AR lights at the upper right are the flags — FXU is Floating (exponent) Underflow, DCK is No Divide (divide check). OV COND is the condition that the 0 and 1 carries are different, *ie* the condition that indicates overflow. The Byte Interrupt flag is BYF6 in the MISC lights in the top row; User In-out is IOT USER in the EX lights at the center of the panel. The CPA lights in the top row and the five lights under them at the left are the processor conditions — PDL OV is Pushdown (list) Overflow. The AS= lights in the middle row indicate when the (relocated) core memory address or the fast memory address is the same as the address switches. The remaining lights are for maintenance.

The panels on the two types of memories are shown on page C9. These are almost exclusively for maintenance, and (as with most of the lights on the processor bays) if the operator must use them he should consult the maintenance manual and the flow charts. The ACTIVE lights indicate which processor currently has access to the memory.

Operating Keys

Each key except STOP turns on one of the key indicators at the upper right on the bay 2 panel. These are for flipflops that allow the key functions to be repeated automatically and also allow certain of them to be synchronized to the processor time chain so they can be performed while the processor is running.

READ IN

Clear all IO devices and all processor flags including User; turn on the RIM light in the upper right on bay 1 and the KEY RDI light in the upper right on bay 2. Execute DATAI $D,0$ where D is the device code specified by the readin device switches on the small panel at the left of the reader. Then execute a series of BLKI $D,0$ instructions until the left half of location 0 reaches zero, at which time turn off RIM and KEY RDI. Stop only if the single instruction switch is on; otherwise turn on RUN and execute the last word read as an instruction. [*For information on the data format refer to page 2-72.*]

If RUN is on, pressing this key has no effect.

The rightmost device switch is for bit 9 of the instruction and thus selects the least significant octal digit (which is always 0 or 4) in the device code.

CAUTION

Do not initiate any other key function while RIM is on. If read in must be stopped (*eg* because of a crumpled tape), press RESET (see below).

START

Load the contents of the address switches into PC, turn on RUN, and begin normal operation by executing the instruction at the location specified by PC.

This key function does not disturb the flags or the IO equipment; hence if USER MODE is lit a user program can be started.

If RUN is on, pressing this key has no effect.

CONT (Continue)

Turn on RUN (if it is off) and begin normal operation in the state indicated by the lights.

STOP

Turn off RUN so the processor stops before beginning the next instruction. Thus the processor usually stops at the end of the current instruction, which is displayed in the lights. However, if a key function that can be performed while RUN is on has been synchronized, the processor performs that function before stopping. In either case PC points to the next instruction.

If the processor does not reach the end of the instruction within 100 μs, inhibit further effective address calculation — it is assumed the processor is caught in an indirect addressing loop. Pressing CONT when the processor is stopped in an address loop causes it to start the same instruction over.

RESET

Clear all IO devices and clear the processor including all flags. Turn on the triangular light beside MEMORY DATA (turn off the light beside PROGRAM DATA). If RUN is on duplicate the action of the STOP key before clearing.

If STOP will not stop the processor, pressing this key will.

Note that an instruction exe-
cuted from the console can
alter the processor state just
like any instruction in the
program: it can change PC by
jumping or skipping, alter the
flags, or even cause a non-
existent-memory stop.

XCT

Execute the contents of the data switches as an instruction without incre-
menting PC. If RUN is on, insert this instruction between two instructions
in the program. Inhibit priority interrupts during its execution to guarantee
that it will be finished.

If USER MODE is lit all user restrictions apply to an instruction executed
from the console.

NOTE

The remaining key functions all reference memory.
They use an absolute address and all of memory is
available to them; in other words protection and
relocation are not in effect even if USER MODE is
lit. However they can set such flags as Address
Break and Nonexistent Memory.

EXAMINE THIS

Display the contents of the address switches in the MA lights and the con-
tents of the location specified by the address switches in the memory indica-
tors. Turn on the triangular light beside MEMORY DATA (turn off the
light beside PROGRAM DATA). If RUN is on, insert this function between
two instructions in the program.

EXAMINE NEXT

If RUN is on, pressing this
key has no effect.

Add 1 to the address displayed in the MA lights and display the contents of
the location specified by the incremented address in the memory indicators.
Turn on the triangular light beside MEMORY DATA (turn off the light
beside PROGRAM DATA).

DEPOSIT

Deposit the contents of the data switches in the location specified by the
address switches. Display the address in the MA lights and the word
deposited in the memory indicators. Turn on the triangular light beside
MEMORY DATA (turn off the light beside PROGRAM DATA). If RUN is
on, insert this function between two instructions in the program.

DEPOSIT NEXT

If RUN is on, pressing this
key has no effect.

Add 1 to the address displayed in the MA lights and deposit the contents of
the data switches in the location specified by the incremented address. Dis-
play the word deposited in the memory indicators. Turn on the triangular
light beside MEMORY DATA (turn off the light beside PROGRAM DATA).

CAUTION

Never press two keys simultaneously as the proc-
essor may attempt to perform both functions at
once.

Operating Switches

Whenever the processor references memory at the location specified by the
address switches (relocated if USER MODE is on), the contents of that loca-
tion are displayed in the memory indicators (unless the light beside
PROGRAM DATA is on). The group of five switches at the left of the keys
allows the operator to make the processor halt or request an interrupt when
reference is made to the specified location *in core memory* for a particular
purpose (no action is produced by fast memory reference). The purpose is
selected by the three address condition switches. INST FETCH selects the
condition that access is for retrieval of an instruction (including an instruc-
tion executed by an XCT or contained in an interrupt location or a trap for
an unimplemented operation) or an address word in an effective address cal-
culation. DATA FETCH selects access for retrieval of an operand (other
than in an XCT). WRITE selects access for writing in memory. Whenever
reference to the specified location satisfies any selected address condition,
the processor performs the action selected by the other two switches. ADR
STOP halts the processor with MEMORY STOP on (PC points to the instruc-
tion that was being executed, or if the MC WR light on bay 2 is on, PC may
point to the one following it); ADR BREAK turns on the CPA ADR BRK
light (Address Break flag, CONI APR, bit 21) on bay 1, requesting an inter-
rupt on the processor channel.

> AC and index register refer-
> ences can be included by
> turning off the FM ENB
> switch (see below).

The description of each switch relates the action it produces while it is on.

SING INST

Whenever the processor is placed in operation, clear RUN so that it stops at
the end of the first instruction. Hence the operator can step through a pro-
gram one instruction at a time, by pressing START for the first one and
CONT for subsequent ones. Each time the processor stops, the lights display
the same information as when STOP is pressed.

CLK FLAG (Clock flag) on bay 1 is held off to prevent clock interrupts
while SING INST is on. Otherwise interrupts would occur at a faster rate
than the instructions.

SING INST will not stop the processor if a hangup prevents it from getting
to the end of an instruction. Use STOP or RESET.

SING CYCLE

Whenever the processor is placed in operation, stop it with MEMORY STOP
on at the end of the first *core memory* reference. Hence the operator can
step through a program one memory reference at a time, by pressing START
for the first one and CONT for subsequent ones. To determine what infor-
mation is displayed in the lights, consult the flow charts.

> To stop at AC and index
> register references, turn off
> the FM ENB switch (see
> below).

CENTRAL PROCESSOR §2.16

PAR STOP

Stop with MEMORY STOP on at the end of any memory reference in which even parity is detected in a word read. A parity stop is indicated by the following: CPA PAR ERR (Parity Error flag) on bay 1 is on; and among the PAR lights in the bottom row on bay 2, IGN (ignore parity) and ODD are off, STOP is on, and BIT displays the parity bit for the word in the parity buffer at the left.

If IGN is on (it displays a signal from the memory), parity errors are not detected and no stop can occur.

NXM STOP

Stop with MEMORY STOP on if a memory reference is attempted but the memory does not respond within 100 μs. This type of stop is indicated by CPA NXM FLAG (Nonexistent Memory flag) on bay 1 being on.

REPT

If any key (except STOP) is pressed, then every time the key function is finished, wait a period of time determined by the setting of the speed control and repeat the given key function. If CONT is pressed and no switch is on that would stop the program (eg SING INST, SING CYCLE), then continue following the repeat delay whenever a HALT instruction is executed. Continue to repeat the key function until RESET is pressed or REPT is turned off.

The key function is repeated once after REPT is turned off, but this is noticeable only with very long repeat delays.

The end of a key function is equivalent to completion of all processor operations associated with the function only for read in, examine, examine next, deposit, and deposit next. In other cases the processor continues in operation. Eg the execute function is finished once the instruction to be executed is set up internally, but the processor then executes that instruction. Hence when using speed range 6, the operator must be careful not to allow the key function to restart before the processor is really finished.

The speed control includes a six-position switch that selects the delay range and a potentiometer for fine adjustment within the range. Delay ranges are as follows.

Position	Range
1	270 ms to 5.4 seconds
2	38 ms to 780 ms
3	3.9 ms to 78 ms
4	390 μs to 7.8 ms
5	27 μs to 540 μs
6	2.2 μs to 44 μs

MI PROG DIS

Turn on the triangular light beside MEMORY DATA (turn off the light beside PROGRAM DATA) and inhibit the program from displaying any information in the memory indicators. The indicators will thus continually display the contents of locations selected from the console.

REPT BYP

If REPT is on, trigger the repeat delay at the beginning of the key function. Hence the function is repeated even if it does not run to completion.

FM ENB

This switch is left on for normal operation with a fast memory. Turning it *off* (lower part in) substitutes the first sixteen core locations for the fast memory. The switch is left off if there is no fast memory, and it can be used to allow stopping or breaking at fast memory references.

SHIFT CNTR MAINT

Stop before each step in any shift operation. Pressing CONT resumes the operation. Once a shift has been stopped, the processor will continue to stop at each step throughout the rest of the given shift operation even if the switch is turned off.

At the right end of panel 1J behind the bay doors are two toggle switches. FP TRP causes the floating point and byte manipulation instructions (codes 130–177) to trap to locations 60-61. MA TRP OFFSET moves the trap and interrupt locations to 140–161 for a second processor connected to the same memory.

Inside each memory bay are switches for selecting the memory number and interleaving memories. Also in the memory are a power switch, a restart pushbutton, and a switch for single step operation (these three are located on the indicator panel for the MB10 memory).

3

Basic In-out Equipment

The PDP–10 contains three in-out devices as standard equipment: tape reader, tape punch, and teletype. These devices are used principally for communication between computer and operator using a paper medium, tape or form paper.

The punch supplies output in the form of 8-channel perforated paper tape in either of two modes. In alphanumeric mode, 8-bit characters are processed; in binary mode, 6-bit characters. The information punched in the tape can be brought into memory by the tape reader, which handles characters in the same two modes.

The program can type out characters on the teletype and can read characters that have been typed in at the keyboard. This device has the slowest transfer rate of any, but it provides a convenient means of man-machine interaction.

3.1 PAPER TAPE READER

The reader processes 8-channel perforated paper tape photoelectrically at a speed of 300 lines per second. The device can operate in alphanumeric or binary mode, as specified by the 0 or 1 state respectively of the Binary flag. In alphanumeric a single tape-moving command reads all eight channels from the first line encountered. In binary the device reads six channels from the first six lines in which hole 8 is punched and assembles the information into a 36-bit word. The interface contains a 36-bit buffer from which all data is retrieved by the processor. The reader device code is 104, mnemonic PTR.

CONO PTR,   **Conditions Out, Paper Tape Reader**

71060	I	X	Y
0	12 13 14	17 18	35

Set up the reader control register according to bits 30–35 of the effective conditions E as shown (a 1 in a flag bit sets the flag, a 0 clears it).

			BINARY	BUSY	DONE	PRIORITY INTERRUPT ASSIGNMENT		
27	28	29	30	31	32	33	34	35

CONI PTR,   **Conditions In, Paper Tape Reader**

71064	I	X	Y
0	12 13 14	17 18	35

Read the status of the reader into bits 27 and 30–35 of location E as shown.

TAPE			BINARY	BUSY	DONE	PRIORITY INTERRUPT ASSIGNMENT		
27	28	29	30	31	32	33	34	35

Placing the tape in motion sets the Tape flag and it remains set as long as the tape is in the read head. A 0 in bit 27 indicates that the last time an attempt was made to read, the reader was out of tape.

DATAI PTR,   **Data In, Paper Tape Reader**

71044	I	X	Y
0	12 13 14	17 18	35

Transfer the contents of the reader buffer into location E. Clear Done and set Busy.

TAPE CHANNELS

FEED
HOLE
⑧ ⑦ ⑥ ⑤ ④ ° ③ ② ①

↑
TAPE MOTION

Setting Busy clears the reader buffer, sets the Tape flag (if it is not already set) and places the reader in operation. If Binary is clear, all eight channels from the first line on tape are read into bits 28–35 of the buffer with channel 1 corresponding to bit 35 (the presence of a hole produces a 1 in the buffer). If Binary is set, the device reads only channels 1–6, but it reads the first six lines encountered in which channel 8 is punched (lines without a hole in channel 8 are skipped) and assembles them into a full word in the buffer. The first line is at the left in the word and channel 1 corresponds to the rightmost bit in each 6-bit byte.

After the specified number of lines has been read, the reader clears Busy and sets Done, requesting an interrupt on the assigned channel. A DATAI brings the data into memory and also causes the reader to continue in operation. The programmer must give a CONO to clear Busy if he does not want the reader to move the tape after the final DATAI is given.

If the tape runs out or malfunctions while a read operation is in progress, the Tape flag is cleared and the reader shuts down.

Timing. At 300 lines per second the reader takes 3.33 ms per alphanumeric character, 20 ms per binary word if the binary characters are contiguous. After Done is set, the program has 1.6 ms to give a DATAI and keep the tape in continuous motion. Waiting longer causes the reader to shut down for 40 ms. Thus start-stop operation is limited to 25 lines per second.

EXAMPLES. This program reads ten binary words (60 lines) from paper tape and stores them in memory beginning at location 4000. The block pointer is kept in accumulator PNT.

```
        MOVE    PNT,[IOWD  12,4000]     ;Put pointer in PNT
        CONO    PTR,60          ;Set up reader
NEXT:   CONSO   PTR,10          ;Watch Done
        JRST    .-1
        BLKI    PTR,PNT         ;Word ready, get it
        JRST    .+2             ;Got all data
        JRST    NEXT            ;Go gack for next word
        .
        .
        .
```

If instead of just waiting we wish to continue our program while the data is coming in, we can use the priority interrupt. The following uses channel 4 and signals the main program that the data is ready by setting bit 35 of accumulator F.

```
        MOVE    17,[BLKI  PTR,[IOWD   12,4000] ]
        MOVEM   17,50           ;Set up 50 and 51 for channel 4
        MOVE    17,[JSR   DONE]
        MOVEM   17,51
        CONO    PTR,64          ;Set up reader on channel 4
        CONO    PI,12210        ;Clear PI, then activate it and turn on
                                ;channel 4
        .                       ;Continue program
        .
        .
        TRZN    F,1             ;Check if data ready when needed
        JRST    .-1             ;Wait if necessary
        .
        .
        .
DONE:   0                       ;Interrupt routine, block done
        CONO    PTR,0           ;Stop tape
        TRO     F,1             ;Set F bit 35
        JEN     @DONE           ;Dismiss and restore flags
```

Operation. Tapes must be unoiled and opaque. The reader is located just above the console operator panel. To load it, place the fanfold tape stack vertically in the bin at the right, oriented so that the front end of the tape is nearer the read head and the feed holes are away from you. Lift the gate, take three or four folds of tape from the bin, and slip the tape into the reader from the front. Carefully line up the feed holes with the sprocket teeth to avoid damaging the tape, and close the gate. Make sure that the part of the tape in the left bin is placed to correspond to the folds, otherwise it will not stack properly. If the program requires that the Tape flag be set and it is not, briefly press the white feed button located on the face of the reader. After the program has finished reading the tape, run out the remaining trailer by pressing the feed button.

Indicators for the reader are on the panel at the top of bay 1 (the panel is

pictured in Appendix C). The paper tape reader lights in the second row from the bottom display the contents of the buffer. The PI assignment and flags are displayed in the PTR lights in the middle of the third row (EOT is the Tape flag). The remaining PTR lights are for maintenance.

Readin Mode

The only requirement (beyond those given in §2.12) for readin mode with paper tape is that the data must be in binary (hole 8 punched). To select the reader in the readin device switches, turn on the third from the left and the last on the right (104).

The program below is the RIM10B Loader, which is brought into the accumulators in readin mode, and then continues to read any number of blocks of binary data from the same tape. The tape is formatted as a series of blocks separated by a half-dozen lines of blank tape (tape with only feed holes punched). The first block is the loader in readin format. The rest of the tape contains any number of data blocks and ends with a transfer block. Each data block contains any number of words of program data, preceded by a standard IO block pointer for the data only, and followed by a checksum, which is the sum of all the data words and the pointer. It is recommended that the number of data words per block be limited to twenty for ease in repositioning the tape in case of error. The transfer block is a JRST to the starting location of the program, followed by a throw-away word to stop the reader.

This loader is written for minimum size and is quite complex. Do not approach it as a simple programming example.

```
          XWD    -16,0           ;14₁₀ words starting at location 1
ST:       CONO   PTR,60          ;Set up reader binary
ST1:      HRRI   A,RD+1          ;Put RD+1 in Y part of A
RD:       CONSO  PTR,10          ;Watch Done
          JRST   .-1
          DATAI  PTR,@TBL1-RD+1(A)    ;First and last words in
                                      ;ADR, data in block
          XCT    TBL1-RD+1(A)    ;TBL1+2 first word, +1 data,
                                 ;+0 checksum
          XCT    TBL2-RD+1(A)    ;TBL2+2 JRST, +1 data, +0
                                 ;bad checksum
A:        SOJA   A,              ;RD+1 first word, RD data, RD-1
                                 ;last word
TBL1:     CAME   CKSM,ADR        ;Compare computed checksum with
                                 ;one read
          ADD    CKSM,1(ADR)     ;Add word read to checksum
          SKIPL  CKSM,ADR        ;Put first word in CKSM, skip if
                                 ;pointer
TBL2:     JRST   4,ST            ;Halt if checksum bad
          AOBJN  ADR,RD          ;If data done, go to A; otherwise wait
                                 ;for next word
ADR:      JRST   ST1             ;Read in executes this. First and last
                                 ;word of each block also put here
CKSM=ADR+1
```

The processor halts if a computed checksum does not agree with the tape. To reread a block, move the tape back to the preceding blank area and press the continue key. A halt following the transfer block is not an error — many programs begin by halting.

3.2 PAPER TAPE PUNCH

The punch perforates 8-channel tape at speeds up to 50 lines per second. It can operate in alphanumeric or binary mode, as specified by the 0 or 1 state respectively of the Binary flag; but in either mode a single tape-moving command punches only one line. Alphanumeric mode punches an 8-bit character supplied by the program; binary mode always punches channel 8, never punches channel 7, and punches a 6-bit character in the remaining channels. The interface contains an 8-bit buffer that receives data from the processor. The punch device code is 100, mnemonic PTP.

CONO PTP,   **Conditions Out, Paper Tape Punch**

71020	I	X	Y
0	12 13 14	17 18	35

Set up the punch control register according to bits 30–35 of the effective conditions E as shown (a 1 in a flag bit sets the flag, a 0 clears it).

			BINARY	BUSY	DONE	PRIORITY INTERRUPT ASSIGNMENT		
27	28	29	30	31	32	33	34	35

CONI PTP,   **Conditions In, Paper Tape Punch**

71024	I	X	Y
0	12 13 14	17 18	35

Read the status of the punch into bits 29–35 of location E as shown.

		NO TAPE	BINARY	BUSY	DONE	PRIORITY INTERRUPT ASSIGNMENT		
27	28	29	30	31	32	33	34	35

A 1 in bit 29 indicates that the punch is out of tape.

DATAO PTP,   Data Out, Paper Tape Punch

71014	I	X	Y
0	12 13 14	17 18	35

Load the contents of bits 28–35 of location E into the punch buffer. Clear Done and set Busy.

A CONO need be given only to change Binary or the PI assignment; DATAO sets Busy while loading the buffer. Setting Busy places the punch in operation. If Binary is clear, one line is punched in tape from bits 28–35 of the buffer with bit 35 corresponding to channel 1 (a 1 in the buffer produces a hole in the tape). If Binary is set, channel 8 is punched, channel 7 is not punched, and the remaining channels are punched from bits 30–35 of the buffer with bit 35 corresponding to channel 1. After punching is complete, the device clears Busy and sets Done, requesting an interrupt on the assigned channel.

Timing. If Busy is set when the punch motor is off, punching is automatically delayed 1 second while the motor gets up to speed. While the motor is on, punching is synchronized to a punch cycle of 20 ms. After Done sets, the program has 10 ms within which to give a new DATAO to keep punching at the maximum rate; after 10 ms punching is delayed until the next cycle. If Busy remains clear for 5 seconds the motor turns off.

EXAMPLE. Suppose we wish to punch out the same information we read from tape in the examples of the previous section. We cannot use a BLKO as an interrupt instruction unless we first spread the 6-bit characters over sixty memory locations. The example uses channel 5 and assumes that other channels are already in use.

```
        MOVE   A,[JSR  PUNCH]
        MOVEM  A,52            ;Set up channel 5
        CONO   PTP,55          ;Request interrupt for first word
        CONO   PI,2004         ;Turn on channel 5
          .                    ;Continue program
          :
          :

PUNCH:  0                      ;Interrupt routine
        ILDB   A,BYPPNT        ;Put byte in A
        AOSL   CNT             ;Got all bytes?
        CONO   PTP,40          ;Yes, prevent interrupt after last word
        DATAO  PTP,A           ;Punch byte
        JEN    @PUNCH

BYPPNT: XWD    440600,4000     ;Generate pointer here
CNT:    ↑D-60                  ;Initialize count
```

Operation. The punch is located behind the reader; both are in a drawer that pulls out from the front of the console. Fanfold tape is fed from a box at the rear of the drawer. After it is punched, the tape moves into a storage

bin from which the operator may remove it through a slot on the front. Pushing the feed button beside the slot clears the punch buffer and punches blank tape as long as it is held in. Busy being set prevents the button from clearing the buffer, so pressing it cannot interfere with program punching.

To load tape, first empty the chad box behind the punch. Then tear off the top of a box of fanfold tape (the top has a single flap; the bottom of the box has a small flap in the center as well as the flap that extends the full length of the box). Set the box in the frame at the back and thread the tape through the punch mechanism. The arrows on the tape should be underneath and should point in the direction of tape motion. If they are on top, turn the box around. If they point in the opposite direction, the box was opened at the wrong end; remove the box, seal up the bottom, open the top, and thread the tape correctly.

To facilitate loading, tear or cut the end of the tape diagonally. Thread the tape under the out-of-tape plate, open the front guide plate (over the sprocket wheel), push the tape beyond the sprocket wheel, and close the front guide plate. Press the feed button long enough to punch about a foot and a half of leader. Make sure the tape is feeding and folding properly in the storage bin. Pushing the button labeled POWER sets No Tape, pushing it again clears the flag. It can be used to hold the program at bay while a tape is being loaded.

To remove a length of perforated tape from the bin, first press the feed button long enough to provide an adequate trailer at the end of the tape (and also leader at the beginning of the next length of tape). Remove the tape from the bin and tear it off at a fold within the area in which only feed holes are punched. Make sure that the tape left in the bin is stacked to correspond to the folds; otherwise, it will not stack properly as it is being punched. After removal, turn the tape stack over so the beginning of the tape is on top, and *label it* with *name, date*, and other appropriate information.

Indicators for the punch are the PTP lights in the top row of the panel at the top of bay 1. The numbered lights display the last line punched.

3.3 TELETYPE

Two teletypewriter models are regularly available with the PDP–10 for use at the console: the KSR 35, which is capable of speeds up to ten characters per second, and the KSR 37, which can handle up to fifteen characters per second. The program can type out characters and can read in the characters produced when keys are struck at the keyboard.

The teletype separates its input and output functions and in effect acts like two devices with a single device code: each has its own Busy and Done flags, but the two share a common interrupt channel. Placing the code for a character in the output buffer causes the teletype to print the character or perform the designated control function. Striking a key places the code for the associated character in the input buffer where it can be retrieved by the program, but it does nothing at the teletype unless the program sends the code back as output.

Character codes received from the teletype have eight bits wherein the most significant is an even parity bit. The Model 35 ignores the parity bit in characters transmitted to it. The Model 37 ignores the parity bit in a code for a printable character, but it performs no function when it receives a control code with incorrect parity.

The Model 37 has the entire character set listed in the table in Appendix B. Lower case characters are not available on the Model 35, but transmitting a lower case code to the teletype causes it to print the corresponding upper case character. To go to the beginning of a new line the program must send both a carriage return, which moves the type box to the left margin, and a line feed, which spaces the paper. The teletype device code is 120, mnemonic TTY.

CONO TTY,   Conditions Out, Teletype

71220	I	X	Y
0	12 13 14	17 18	35
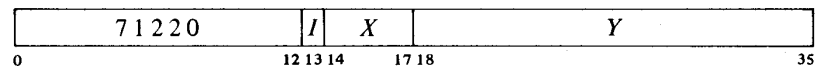
Set up the teletype control register according to bits 24–35 of the effective conditions E as shown (a 1 in bit 24 sets Test, a 0 clears it; all other flag functions are produced by 1s, 0s have no effect).

TEST	CLEAR INPUT BUSY	CLEAR INPUT DONE	CLEAR OUTPUT BUSY	CLEAR OUTPUT DONE	SET INPUT BUSY	SET INPUT DONE	SET OUTPUT BUSY	SET OUTPUT DONE	PRIORITY INTERRUPT ASSIGNMENT		
24	25	26	27	28	29	30	31	32	33	34	35

Setting Test connects the output buffer directly to the input buffer, allowing the program to check out the interface logic without the line and the device.

CONI TTY,   Conditions In, Teletype

71224	I	X	Y
0	12 13 14	17 18	35
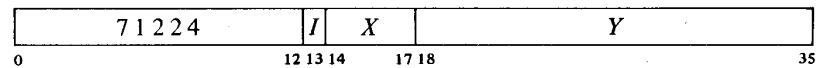
Read the status of the teletype into bits 24 and 29–35 of location E as shown.

TEST					INPUT BUSY	INPUT DONE	OUTPUT BUSY	OUTPUT DONE	PRIORITY INTERRUPT ASSIGNMENT		
24	25	26	27	28	29	30	31	32	33	34	35

DATAO TTY,   Data Out, Teletype

71214	I	X	Y
0	12 13 14	17 18	35
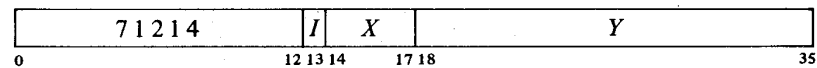
Load the contents of bits 28–35 of location E into the output buffer. Clear Output Done, set Output Busy, and enable the transmitter.

DATAI TTY,   Data In, Teletype

71204	I	X	Y
0	12 13 14	17 18	35
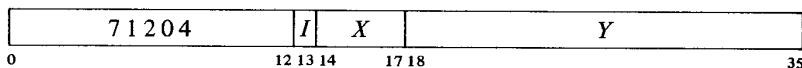
Transfer the contents of the input buffer into bits 28—35 of location E.
Clear Input Done.

Output. A CONO need be given only to change the PI assignment;
DATAO sets Output Busy and enables the transmitter while loading the
buffer. Enabling the transmitter causes it to send the contents of the output
buffer serially to the teletype. Completion of transmission clears Output
Busy and sets Output Done, requesting an interrupt on the assigned channel.

Input. Teletype reception requires no initiating action by the program
except to supply a PI assignment. Striking a key transmits the code for the
character serially to the input buffer. The beginning of reception sets Input
Busy; completion clears Input Busy and sets Input Done, requesting an
interrupt on the assigned channel. A DATAI brings the character into
memory and clears Input Done.

Timing. The Model 35 can type up to ten characters per second. After
Output Done is set, the program has 9.09 ms to give a DATAO to keep
typing at the maximum rate. After Input Done is set, the character is
available for retrieval by a DATAI for 22.73 ms before another key strike
can destroy it.

The 37 can handle fifteen characters per second, 66.7 ms per character.
After Output Done is set, the program has 6.67 ms to send a new character
to maintain the maximum typing rate. After Input Done is set, the character
is available for at least 10 ms.

The sequence carriage return-line feed, when given in that order, allows
sufficient time for the type box to get to the beginning of a new line. After
tabbing, the program must wait for completion of the mechanical function
by sending one or two rubouts. If the time is critical, the programmer
should measure the time required for his tabs. Tabs are normally set every
eight spaces (columns 9, 17, . . .) and require one rubout.

Operation. The illustrations on the following two pages show the two
teletype models. The teletype is actually two independent devices, keyboard
and printer, which can be operated simultaneously. Power must be turned
on by the operator. On the 35 the switch is beside the keyboard, and has an
unmarked third position (opposite ON) which turns on power but with the
machine off line so it can be used like a typewriter. A similar switch is
located beneath the stand on the 37.

The keyboard resembles that of a standard typewriter. Codes for printable
characters on the upper parts of the key tops on the 35 are transmitted by
using the shift key; most control codes require use of the control key. Those
familiar with the 35 who are using the 37 for the first time should take a
close look at the keyboard. On the 37 the shift is used for real upper case
characters. The control key is used for some control characters, but many

Teletype KSR 35

have separate keys. Note also that both the keyboard arrangement and the labels differ somewhat. On both, the line feed (labeled "new line" on the 37) spaces the paper vertically at six lines to the inch, and must be combined with a return to start a new line. The local advance (feed) and return keys affect the printer directly and do not transmit codes. Appendix B lists the complete teletype code, ASCII characters, key combinations, and differences between the two models.

Indicators for the teletype are the TTY lights in the second row of the

panel at the top of bay 1. The numbered lights display the last character
typed in from the keyboard (bit 8 is parity). The ACT lights indicate
activity in the transmitter and receiver. The remaining lights display the PI
assignment and flags (the Input and Output Done flags are labeled TTI
FLAG and TTO FLAG).

Teletype manuals supplied with the equipment give complete, illustrated
descriptions of the procedures for loading paper, changing the ribbon, and
setting horizontal and vertical tabs. The first two procedures are fairly

Teletype KSR 37

obvious: observe the paper or ribbon path and duplicate it. The other tasks are usually left for maintenance personnel. In any event, the best and easiest way to learn to do any of these things is to have someone who knows show you how.

4

Hardcopy Equipment

This chapter discusses the line printer, *XY* plotter, card reader, and card punch. Like the basic in-out equipment, these devices are primarily for communication between computer and operator using a paper medium: form paper, graph paper or cards.

The line printer provides text output at a relatively high rate. The program must effectively typeset each line; upon command the printer then prints the entire line. With the plotter, the program can produce ink drawings by controlling the incremental motion of pen on paper in a cartesian coordinate system. Curves and figures of any shape can be generated by proper combinations of motion in x and y.

The card equipment processes standard 12-row 80-column cards. Many programmers find cards a convenient medium for source program input and for supplying data that varies from one program run to another. Cards are convenient to prepare manually, input is much faster than paper tape, and simple changes are easy to make: individual cards can be repunched, and cards can be added or removed from the deck. The card reader cannot be used in readin mode, but a standard card-reading program in readin format can be kept on paper tape or DECtape. A possible consideration in using cards is that many installations do not include an online card punch.

These four devices are all run by the BA10 Hardcopy Control. Interface logic for a plotter can also be mounted in the TD10A DECtape Control.

4.1 LINE PRINTER

The line printer outputs hardcopy composed of lines 132 characters long at a nominal rate of 300,600 or 1000 lines per minute. The standard printer has sixty-four printing characters available to the program. The characters and codes are the figure and upper case sets, codes 040–137, in the teletype code [*Appendix B*]. When a lower case code (140–176) is given, the corresponding upper case code is loaded into the buffer. Besides accepting printing characters, the printer responds to ten control characters, HT, CR, LF, VT, FF, DLE and DC1–4. All other codes are ignored.

The printer has a 132-character buffer that holds the image of a single line; the program must first load the buffer up to five characters at a time, and then give a control character to print the entire line. The buffer is loaded from left to right, and only the portion filled produces a printout. Hence

for each line the program need send out characters (including spaces) only as far as the rightmost nonspace character. The characters are printed in the order that they pass the print hammers, and a given character is printed simultaneously in all positions that require it. In other words the drum has a row of 132 Ms, a row of Ns, etc; all Ms are printed together, all Ns together, and so forth. The first character printed depends only upon the position of the drum when the print command is given.

Printers having more than sixty-four characters are also available. The 96-character printer outputs 600 lines per minute and has the entire figure, upper case and lower case sets, codes 040–176. This is actually only ninety-five characters, but an option allows use of the delete code to select an extra character on the drum. A single delete code is ignored, but two consecutive 177s cause the code 177 to be loaded into the buffer. When the code for a printing character is the same as one for a nonprinting character and is loaded by giving it immediately after a delete code, the printing character is said to be "hidden" under the nonprinting one.

The 128-character printer outputs 500 lines per minute and uses the entire set of 7-bit codes for printing characters, with characters hidden under the ten control characters and also under null and delete.

Output Format. Paper motion is controlled by a format tape loop in the printer. The tape has eight columns and the loop corresponds to an integral number of pages of the fanfold form paper. With the exception of CR, every control character that prints a line from the contents of the buffer produces a different spacing by selecting a particular tape column. The paper then advances until a hole is encountered in the selected column.

The standard paper has 11-inch pages of sixty-six lines, and the standard tape for these generates the formats listed below. The fourth column gives the hole positions in terms of the numbered lines *on the tape*. The tape is usually installed at random and then positioned by pressing the top-of-form button on the printer. Then the paper is adjusted so that the desired line on the paper corresponds to line 0 on the tape. Ordinarily the paper is set with the print hammers at the fourth line, so all but one of these formats leaves a three-line margin at the top and a margin of at least three lines at the bottom of each page.

Virtually any character set can be had on any printer by special order. In any event characters after the first ninety-five are always special order.

Spacing other than the standard can be produced by using a different format tape. The length of the loop should correspond to one or more pages of the printer form used, with holes punched at the lines where paper spacing is to stop.

Programmers generally treat the data for the line printer and teletype identically, using the combination CR plus LF for printing and spacing. This way a given character string can be outputted on either device. CR is used alone only when the next print command will overprint, *ie* will print another character in a column position already printed. With this technique the program can produce a character such as "≠" by overprinting a slash on an equal sign (or vice versa).

Character	Column	Normal meaning	Hole positions
FF (014)	1	Top of form	Line 0
CR (015)	None	No spacing (paper motion inhibited)	
LF (012)	8	Single space with automatic top of form after every 60 impressions	Every line from 0 to 59
DC1 (021)	3	Double space with automatic top of form after every 30 impressions	Every even numbered line from 0 to 58
DC2 (022)	4	Triple space with automatic top of form after every 20 impressions	Every third line from 0 to 57

DC3 (023)	5	Single space	Every line
DC4 (024)	6	Space one sixth of a page	Lines 0, 10, 20 30, 40, 50
VT (013)	7	Space one third of a page	Lines 0, 20, 40
DLE (020)	2	Space half a page	Lines 0, 30

The actual printer action of advancing the paper to the next hole in the tape produces the "normal" format only if the program consistently selects the same tape column. Always using DC1 to print produces double spaced text from line 4 to line 62 on every page. But if the last print command spaced to an odd numbered line, DC1 moves the paper only one line.

Printing Speed. The printer is available in five models with differing printing speeds.

Printer	Nominal printing speed in lines per minute	Drum rotation in rpm	Time per revolution in ms
LP10A	300	333	180
LP10B	600	750	80
LP10C	1000	1250	48
96 Character	600	750	80
128 Character	500	550	109

Printing begins as soon as a print command is given and terminates when the last required character is printed, *ie* without necessarily waiting for a complete drum revolution. Therefore print time depends on the initial drum position and the number of characters that must pass the print head before the last is printed. No time is required for spaces: the printer produces spaces in a line by not printing anything in the columns corresponding to the buffer positions that hold space characters. As a given character is printed, space codes replace the codes for the character in all buffer positions that hold it, and printing ceases when the buffer is filled with spaces.

A complete print cycle consists of the print time plus the time required for advancing the paper; paper spacing begins immediately after printing terminates, and further printing is inhibited while the paper is moving. It takes about 12 ms to advance the paper one line, about 6–8 ms for each additional line. If the buffer is loaded only with spaces, the print cycle consists entirely of paper spacing.

Using an ordinary distribution of characters results in printing at or slightly above the nominal speed. Printing is faster however if paper spacing occurs while unused characters are passing the print head. *Eg* text that uses only the alphabet can be printed at the full drum rotation speed.

Instructions. The printer has the usual instructions for sending and reading conditions, but after initial setup it can be controlled entirely by the characters sent by a string of DATAOs. The program supplies five characters at a time to a 35-bit character buffer in the printer interface. The interface processes the characters from left to right loading valid data characters into the

line buffer, ignoring invalid characters, and sending control signals to the printer when a control character is encountered. The printer device code is 124, mnemonic LPT.

CONO LPT,   Conditions Out, Line Printer

7 1 2 6 0	I	X	Y
0	12 13 14	17 18	35
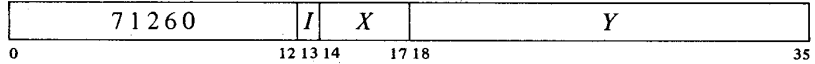
Perform the function given below if specified by a 1 in bit 25 and set up the printer control register according to bits 30–35 of the effective conditions E as shown (a 1 in a flag bit sets the flag, a 0 clears it).

	CLEAR PRINTER			BUSY	DONE	PRIORITY INTERRUPT ASSIGNMENT – ERROR			PRIORITY INTERRUPT ASSIGNMENT – DONE		
24	25	26	27	28	29	30	31	32	33	34	35
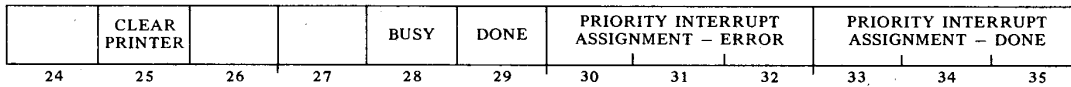
Power turnon and the IO reset signal generated by CONO APR,200000 duplicate this clear function.

If bit 25 is 1, clear Done, set Busy, clear the interface logic, and trigger a print cycle to clear the line buffer. The cycle clears the buffer by replacing the characters in it with spaces, and the time required is the same as would be required to print whatever is in it. Completion of the cycle clears Busy and sets Done, requesting an interrupt on the channel assigned by bits 33–35.

CONI LPT,   Conditions In, Line Printer

7 1 2 6 4	I	X	Y
0	12 13 14	17 18	35
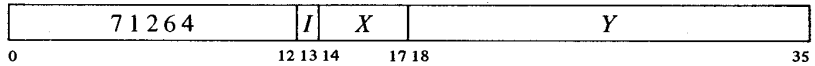
Read the status of the printer into bits 24–35 of location E as shown.

128	96		ERROR	BUSY	DONE	PRIORITY INTERRUPT ASSIGNMENT – ERROR			PRIORITY INTERRUPT ASSIGNMENT – DONE		
24	25	26	27	28	29	30	31	32	33	34	35

A 1 in bit 24 indicates that the printer has a 128-character drum; a 1 in bit 25 indicates that at least 95 characters are available to the program.

DATAO LPT,   Data Out, Line Printer

7 1 2 5 4	I	X	Y
0	12 13 14	17 18	35
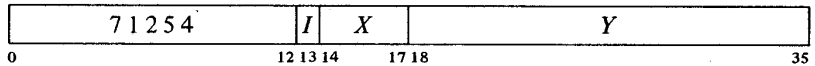
Load the contents of bits 0–34 of location E into the character buffer, clear Done, set Busy, and trigger the interface processing cycle. The format of the

data word and the order in which the characters are processed is as shown.

FIRST	SECOND	THIRD	FOURTH	FIFTH	
0	6 7	13 14	20 21	27 28	34

Characters are assembled into words in this manner by an IDPB loop or an ASCII or ASCIZ pseudoinstruction.

Following power turnon, the Error flag (CONI bit 27) is set if the printer cable is not connected or any other condition exists that makes the printer unavailable to the program [*these other conditions are given in the discussion of printer operation at the end of the section*]. If Error is set when a CONO gives an error PI assignment (with bits 30–32 of E), there is an immediate interrupt request on the error channel. Barring accident or hardware malfunction, an error interrupt is likely to occur during a printout run only when the printer is about to run out of paper or the operator stops it (in either case Error sets and the printer stops when the buffer is empty following the printing of a line).

At the beginning of a print run the program should give a CONO to clear the line buffer and assign the PI channels. After that a CONO need be given only to change the PI assignments; each DATAO starts the character-processing operations of the interface while loading the character buffer. The interface processes the characters from left to right, starting each character cycle when the line buffer is ready. Printing characters are simply sent to the buffer, with lower case codes translated to upper case for a 64-character printer. Unused codes are ignored. The interface responds as follows when a control character is encountered.

◆ A horizontal tab (HT) is simulated by sending a string of spaces to the line buffer. Tab stops are every eight columns (9, 17, . . .). The interface always sends at least one space, and then sends as many more as are necessary for the next character to be at a tab stop. Thus if a DATAO gives the sequence

These tabs are the same as the ones ordinarily used on the teletype.

<p align="center">A HT B</p>

where A is placed in column 7, B will go into column 9. But if A goes into column 8, B will go into column 17.

◆ Upon encountering any other printer control character, the interface signals the printer to print the contents of the line buffer, and unless the character is CR, it also selects a format tape column to space the paper as listed in the format discussion at the beginning of this section. When the buffer again becomes available, subsequent characters will be loaded starting in column 1. If printing is caused by a CR, the next line will overprint unless the paper is advanced before any nonspace characters are loaded into the buffer.

If the buffer is filled with 132 characters and the next character does not cause printing, the interface simulates a line feed to print and advance the paper, and then loads the next character at column 1 for the new line. If the program tabs to the end of a line, the interface simulates a line feed and also tabs at the beginning of the next line. In other words a printing character following the tab will be loaded at column 9 for the new line.

When the interface finishes processing the five characters supplied by a

DATAO, it clears Busy and sets Done, requesting an interrupt on the channel assigned by bits 33–35 of the conditions out.

Timing. The time from one DATAO to the next while the program is loading the buffer is simply the time required by the interface to process five characters. Loading each printing character, including each space in a horizontal tab, takes 10 μs. Skipping an illegal character takes 8 μs.

If the fifth character causes printing, Done is set immediately and the program can give a DATAO to send the first set of characters for the next line. However, the interface does not begin processing the new characters until the buffer becomes available after the printer finishes printing the previous line. If printing is produced by any character before the last, the print time elapses before the interface processes the next character in the current set.

The overall time required for a print run is the total printing and spacing time for all lines as given above in the discussion of the printing speed. The time required to process individual characters is a consideration in programming the DATAOs that load the buffer, but buffer loading time is not a factor in total printer operating time except when loading characters for overprinting (following a CR). This is because the buffer becomes available while the paper is moving, in plenty of time for the program to load it before the paper stops.

EXAMPLES. In the first example, which uses the line printer without the interrupt, we have simply filled in the missing part of the print subroutine given at the top of page 2-61 (it prints the characters that accompany the calling sequence given at the bottom of page 2-60).

```
PRINT:   HRLI    T,440700
         ILDB    CH,T
         JUMPE   CH,1(T)
         CONSZ   LPT,200     ;Skip when printer not busy
         JRST    .-1         ;Wait for Busy to clear
         LSH     CH,1        ;Shift character to bits 28-34
         DATAO   LPT,CH      ;Send character to printer
         JRST    PRINT+1
```

The same program could be used for output on the teletype by making the substitution

$$\text{CONSZ LPT,200} \rightarrow \text{CONSZ TTY,20}$$

and deleting the LSH CH,1.

The above is perhaps an overly simple example. It assumes the line buffer is clear initially and the printer is available. Moreover the processor spends most of its time waiting. Characters are processed individually in order to detect the null, but if the processor has anything else to do, it would be much more efficient to use the interrupt and send five characters at a time.

In the following example the main program sets up each print run by giving a JSR SETUP. The number of words printed and the starting location of the block containing them are determined by the contents of PNTR1. Once a run is set up, the program can change the contents of PNTR1 for the next one.

```
SETUP:   0
         SKIPGE  PNTR
         JRST    .-1           ;Wait for current IO to finish
         MOVE    T,[JSR  ERROR]
         MOVEM   T,42          ;Channel 1 for error
         MOVE    T,[JSR  DATA]
         MOVEM   T,44          ;Channel 2 for data
         MOVE    T,PNTR1
         MOVEM   T,PNTR        ;Set up new IO block pointer
         CONO    LPT,2012      ;Clear printer, assign channels
         CONO    PI,2340       ;Turn on PI and channels
         JRST    @SETUP

PNTR1:   0
PNTR:    0

ERROR:   0
         CONO    LPT,2         ;Drop error request by dropping error
                               ;PI assignment
                               ;Start typing error message
         .
         .
         .
         JEN     @ERROR

DATA:    0
         CONO    LPT,12        ;Reassign error channel
         BLKO    LPT,PNTR      ;Send out word
         CONO    LPT,0         ;Turn off printer
         JEN     @DATA
```

End of clear function sets Done, requesting a data interrupt.

Operation. The 600-line-per-minute printer is illustrated on the following page. At the left on the front of the printer are two round indicators and two columns of square buttons and indicators, some of which are not used. The round lights indicate whether the printer has power: green light for power on, red for off.

The buttons at the top of the columns operate the printer. Pushing START places the printer on line so it can respond to the program (the button is lit while the unit is on line). Pushing STOP takes the unit off line; the operator can then use the TOP OF FORM button to position the paper (or the format tape). If the program has left anything in the buffer, it can be printed by pressing MANUAL PRINT. The maintenance button TEST can also be used while STOP is lit. START, STOP and TOP OF FORM are duplicated at the rear of the printer.

At the bottom of the columns are four alarm lights that indicate when the paper supply is low, the printer is out of paper or the paper is broken, the yoke is open, or there is a circuit malfunction (ALARM STATUS). When the operator presses STOP or there is a paper low alert, START does not go out until the buffer is empty (in other words until the printer finishes printing a line currently being loaded or printed). START goes out immediately if any other alarm condition occurs or power fails. When START is out or the cable to the interface is not connected, the Error flag is set and the printer cannot respond to the program.

Line Printer LP10B   The lights for the interface are in the top two rows on the hardcopy
control indicator panel [*illustrated on the opposite page*]. The top row dis-
plays the contents of the character buffer; the 7-bit characters are shifted
left for processing. The shift and column counters at the left end of the
second row indicate the last character processed (0–4) and the last buffer
position loaded. The group of lights at the right display the status condi-

tions. Of the group in the center, BUFF AVAIL indicates the line buffer is ready for the next character; the remaining lights are for maintenance.
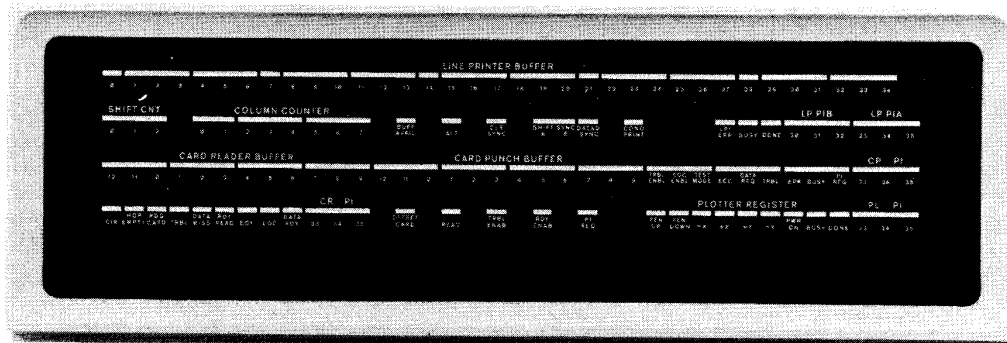
To load paper, press STOP. If START does not go out, the program probably left the last line in the buffer: press MANUAL. When printing is complete the light will go out. Open the printer cover. At the front are two toggle switches: switch both of them to OPEN. The printer yoke will slide forward. Lift the guide plates over the two pairs of tractors, pull out the remaining paper, and press TOP OF FORM to line up the spacing format tape. Bring the beginning of the paper up behind the yoke, over the top and through the rollers on the back. Move the paper until line 4 of a page is lined up with the print hammers (at most installations the point at which the fold should come is marked). Make sure the tractor wheels engage the holes at the edges of the paper, close the guide plates, switch the toggles to CLOSE, close the cover, and press START.

All of the larger and faster printers are as described above. On the slowest printer the lights are at the right, the single PAPER ALARM indicates the paper is either low or broken, and there are no buttons on the back. With the cover open the yoke is controlled by two unmarked plastic switches on either side at the top. Pressing them in at the end nearer the front opens the yoke. This printer has only one pair of tractors, but it has a pair of bars below the yoke. The paper must go over the stationary bar and under the movable one.

Indicator Panel,
Hardcopy Control

4.2 PLOTTER

The XY10 plotter control interfaces the PDP-10 central processor to various plotters that use cartesian coordinates. The models most frequently used are manufactured by Calcomp, but others can be accommodated. The following lists the type and paper size of the most commonly supplied Calcomp models.

Calcomp model	Type	Paper size in inches
502, 602	Bed	31 X 34
518, 618	Bed	54 X 72
563, 663	Drum	29½ X 1440
565, 665	Drum	11 X 1440

These are high accuracy, incremental digital plotters that produce fine quality ink plots of computer-generated data. Bidirectional stepping motors provide individual increments of motion in either coordinate or both at once. The program draws a continuous sequence of line segments by controlling the relative motion of pen and paper with the pen lowered, and it can raise the pen for repositioning.

Motion in y is movement of the pen carriage along a pair of rods. Motion in x is movement of the entire carriage-and-rod mechanism on a bed plotter, movement of the paper underneath the carriage on the drum type. On a bed plotter the coordinate directions are the standard ones when viewing the device from the front: positive x to the right, positive y to the back. The coordinate system on a drum is in the standard orientation when the viewer is standing at the right side, unrolling the paper from the drum with his left hand. In other words positive y is movement of the pen from right to left across the drum, positive x is drum rotation downward at the front (drawing a line toward the paper supply roll at the back).

The step sizes and plotting speeds available with the various Calcomp models are the following.

Calcomp plotters in the 600 series have two step sizes and two plotting speeds: a switch at the back selects the step size, delay settings in the plotter control determine the speed.

Model	Step size	Plotting speed in steps per second
502	*All sizes*	300
518	.005 inch	200
	.002 inch	450
	.1 mm	200
	.05 mm	400
563	.010 inch	200
	.005 inch	300
	.1 mm	300
565	*All sizes*	300
602	*All sizes*	450/900
618	.005/.0025 inch	200/400
	.002/.001 inch	450/900
	.1/.05 mm	200/400
	.05/.025 mm	450/900
663	.010/.005 inch	350/700
	.005/.0025 inch	450/900
	.0025/.00125 inch	450/900
665	*All sizes*	450/900

The program can draw any complete figure by giving a string of DATAOs, each of which supplies the information for one step. The plotter device code is 140, mnemonic PLT.

Calcomp Drum Plotte·
Model 565

CONO PLT,   Conditions Out, Plotter

71420	I	X	Y
0	12 13 14	17 18	35

Set up the plotter control register according to bits 31–35 of the effective conditions E as shown (a 1 in a flag bit sets the flag, a 0 clears it).
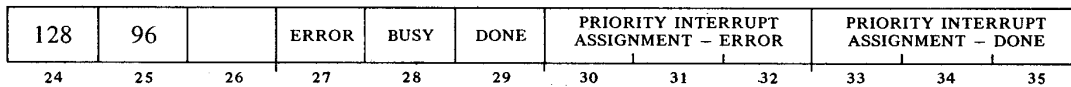
				BUSY	DONE	PRIORITY INTERRUPT ASSIGNMENT		
27	28	29	30	31	32	33	34	35

CONI PLT,   **Conditions In, Plotter**

72424	I	X	Y
0	12 13 14	17 18	35
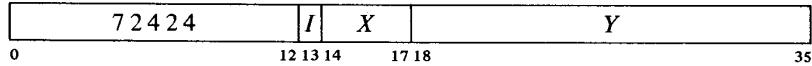
Read the status of the plotter into bits 30–35 of location E as shown.

			POWER ON	BUSY	DONE	PRIORITY INTERRUPT ASSIGNMENT		
27	28	29	30	31	32	33	34	35
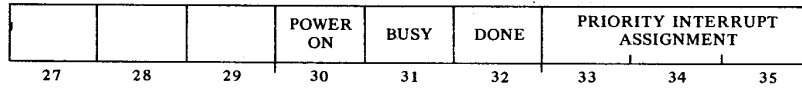
Power On is not available on all plotters.

DATAO PLT,   **Data Out, Plotter**

71414	I	X	Y
0	12 13 14	17 18	35
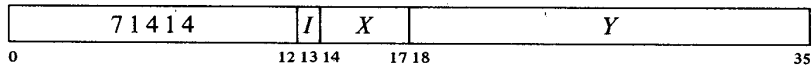
Clear Done, set Busy, and move the pen as specified by bits 30–35 of the contents of location E as shown (a 1 in a bit produces the indicated motion, a 0 has no effect).

RAISE PEN	LOWER PEN	$-\Delta X$ (DRUM UP)	$+\Delta X$ (DRUM DOWN)	$+\Delta Y$ (CARRIAGE LEFT)	$-\Delta Y$ (CARRIAGE RIGHT)
30	31	32	33	34	35
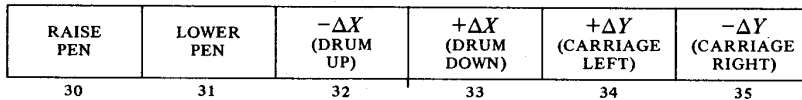
A CONO need be given only to change the PI assignment; DATAO places the plotter in operation by supplying plotting data. After sufficient time has elapsed for the device to carry out the specified action, the control clears Busy and sets Done, requesting an interrupt on the assigned channel.

To avoid drawing line segments shorter than one step, do not raise or lower the pen in the same DATAO that calls for xy motion. The consequences of specifying contradictory movements cannot be predicted.

Timing. Lowering the pen takes 60 ms, raising it takes 10 ms. The time required to move one step in either or both coordinates depends on the plotting speed as follows.

Plotting speed in steps per second	Time per step in ms
200	2.5
300	1.66
350	1.45
400	1.25
450	1.10
700	.70
900	.51

EXAMPLE. The plotting commands sent out by this program are contained six to a word in WC words beginning at location BUFFER. The interrupt routine uses one accumulator which is shared with the main program and other channels.

```
          CONSZ   PLT,7           ;Wait until previous run finished as
          JRST    .−1             ;indicated by no PI assignment
          MOVE    T,[JSR  DATA]
          MOVEM   T,50            ;Set up channel 4
          MOVEI   T,WC*6          ;Set up count for plotting commands
          MOVEM   T,COUNT
          MOVE    T,[POINT  6,BUFFER]     ;Initiate byte pointer
          MOVEM   T,CHARP
          CONO    PLT,4           ;Assign channel
          CONO    PI,2210         ;Turn on PI and channel
          DATAO   PLT,PUP         ;Raise pen to trigger first interrupt
          :
          :
DATA:     0
          SOSGE   COUNT           ;Is plot finished?
          JRST    DATA1           ;Yes
          MOVEM   T,TSAVE         ;Save T
          ILDB    T,CHARP         ;Get next plotting command
          DATAO   PLT,T           ;Plot point
          MOVE    T,TSAVE         ;Restore T
          JEN     @DATA

PUP:      40
TSAVE:    0
COUNT:    0
CHARP:    0

DATA1:    CONO    PLT,0           ;Disconnect plotter from interrupt
          DATAO   PLT,PUP         ;Raise pen
          JEN     @DATA
```

The asterisk is the sign for multiplication in MACRO. POINT is a pseudoinstruction that causes MACRO to generate a byte pointer from the three arguments that follow it. In order these arguments are the byte length in decimal, the address of the location containing the byte, and the position of the rightmost bit of the byte as the decimal number of the bit in the word. If the last argument is omitted, MACRO takes it as −1; in other words, after being incremented the pointer will point to the first byte. The left half of the pointer generated here is 440600.

Operation. On a drum plotter the supply roll is behind the drum. Bring the paper over the drum, down in front, and above and behind the pickup roll underneath the drum (use a piece of masking tape to attach the paper, or roll some onto the tube).

The controls are on the front [refer to the illustration on page 4-11]. To put the plotter on line simply turn on the power and the chart drive. The remaining controls are for manual operation: raising and lowering the pen, moving the carriage and drum in either direction, rapidly or single step. The switch that selects the step size on a 600-series plotter is on the back. The bed plotter has similar controls.

Lights for the plotter are the group at the right end in the bottom row on the hardcopy control indicator panel [page 4-9]. These display the status conditions and the plotting data supplied by the last DATAO. If the·plotter interface is mounted in a DECtape control, there are no lights.

4.3 CARD READER

The card reader handles standard 12-row 80-column cards at speeds up to 1000 cards per minute (833 if power is 50 Hz). Once started, an entire card is read column by column. The reader supplies each column to the processor as twelve bits, and the program can translate in any way it wishes; the standard DEC character representations and the translation to ASCII made by the Monitor are given in Appendix B. Of course the data can simply be in binary at three columns per word (a 7 and 9 punch in the first column is the standard indication that the rest of the card contains binary data).

The interface contains a 12-bit buffer from which each column is retrieved by the processor. The reader device code is 150, mnemonic CR.

CONO CR,　　　**Conditions Out, Card Reader**

71520	I	X	Y
0	12 13 14	17 18	35

Assign the interrupt channel specified by bits 33–35 of E and perform the functions specified by bits 23–32 as shown (in bits 27 and 29 a 1 enables the given flag to interrupt, a 0 disables it; in all other bits a 1 produces the indicated function, a 0 has no effect).

					CLEAR READER	OFFSET CARD		READ CARD	ENABLE TROUBLE INTERRUPTS /	CLEAR DATA MISSED	ENABLE READY TO READ INTERRUPTS /	CLEAR END OF FILE	CLEAR END OF CARD	CLEAR DATA READY	PRIORITY INTERRUPT ASSIGNMENT		
18	19	20	21	22	23	24	25	26	27	28	29	30	31	32	33	34	35

Notes.

23　　Dismiss the PI assignment (assign zero); clear flags Reading Card, Data Missed, End of File, End of Card, Data Ready, Trouble Interrupt Enabled, Ready to Read Interrupt Enabled; clear the card column buffer; and disable any read command given by a CONO if the reader has not yet started the card. If any action specified by the rest of the CONO bits conflicts with these actions, the clear function has precedence.

With the console model, offsetting a card places it in a separate stacker.

24　　If a card is currently being processed in the reader (Card In Reader, CONI bit 24, is 1), offset it when it is placed in the stacker. The card will actually stick out about a half inch from the rest of the stacked deck.

CONI CR,   **Conditions In, Card Reader**

7 1 5 2 4	I	X	Y
0	12 13 14	17 18	35

Read the status of the reader into the right half of location E as shown.

TROUBLE INTERRUPT ENABLED	READY TO READ INTERRUPT ENABLED	PICK FAILURE	PHOTO CELL ERROR	CARD MOTION ERROR	STOP	CARD IN READER	HOPPER EMPTY-STACKER FULL	READING CARD	* TROUBLE	* DATA MISSED	* READY TO READ	* END OF FILE	* END OF CARD	* DATA READY	PRIORITY INTERRUPT ASSIGNMENT
18	19	20	21	22	23	24	25	26	27	28	29	30	31•	32	33 34 35

Notes.   *These bits cause interrupts.

Interrupts are requested on the assigned channel by the setting of Data Ready, Data Missed, End of Card, End of File, and if enabled, Trouble and Ready to Read.

20   The reader has received a read command but has failed to bring in a card from the hopper.

21   The reader has failed to read a card properly and maintenance is probably required. The program should be dubious of any data taken from the card being read when the error occurred.

22   A card has failed to move properly through the reader (it has probably slipped). The program should be dubious of any data taken from the card being read when the error occurred.

23   Reader power is on but the reader is or soon will be unavailable to the program either because the operator has pressed the stop button or there is a trouble condition (bit 27). If Stop is set while a card is being read, the reader usually finishes it; only a power failure can stop the reader in the middle of a card.

24   The reader has brought a card in from the hopper and has not yet finished reading it. The program can give a CONO offset command while this bit is 1.

26   The reader has accepted a read command and has not yet finished reading the card.

27   Bit 20, 21, 22 or 25 is 1. If bit 18 is also 1, the setting of Trouble requests an interrupt on the assigned channel.

Any condition that sets Trouble also sets Stop (bit 23) and the reader will stop at the end of the current card (of course a pick failure prevents the reader from even starting a card). Although a 1 in bit 27 does not necessarily imply an error or malfunction, it always requires operator intervention. If bit 25 is 1 it is very likely that the only trouble is the hopper is empty or the stacker is full.

28   The program failed to retrieve a column of data before the next column was loaded into the buffer by the reader.

29   The reader is ready to accept a read command. If bit 19 is 1, the setting of Ready to Read requests an interrupt on the assigned channel.

The usual procedure is to put an end-of-file card at the end of the deck rather than use the button. Actually the button can be used to signal the program for any purpose provided the reader is off line (stopped).

30   The reader has stopped (probably because the hopper is empty) and the operator has pressed the end-of-file button.

DATAI CR,   Data In, Card Reader

71504	I	X	Y
0	12 13 14	17 18	35

Clear Data Ready, and transfer the contents of the card column buffer into bits 24–35 of location E where the correspondence of card rows to bit positions is as shown.

ROW 12	ROW 11	ROW 0	ROW 1	ROW 2	ROW 3	ROW 4	ROW 5	ROW 6	ROW 7	ROW 8	ROW 9
24	25	26	27	28	29	30	31	32	33	34	35

The program must give a CONO with a 1 in bit 26 to start every card. This read card command waits until the reader is ready, at which time Reading Card sets and the reader card cycle begins. Movement of a card in from the hopper sets Card in Reader. As each column is loaded into the buffer, Data Ready sets, requesting an interrupt on the assigned channel. The program must respond with a DATAI to transfer the column to memory and clear Data Ready. If Data Ready is still set when the next column is loaded into the buffer, Data Missed is set, requesting a second interrupt.

If the program does not retrieve the final column and a CONO that starts a new card does not clear Data Ready, Data Missed will be set by the first column in the new card.

After all eighty columns have been read, Card in Reader goes off, clearing Reading Card and setting End of Card, which requests an interrupt. The card then moves out to the stacker, and when the device is ready to begin a new card cycle, Ready to Read goes on, but only if no new read card command has been given. If a read card command is already waiting when the reader becomes ready, it simply accepts the command and Ready to Read remains off. If no command is waiting, Ready to Read goes on, requesting an interrupt if enabled (CONI bit 19 is 1), and it goes off automatically when a new command is given.

Timing. After Reading Card sets, 18 ms elapse before Card in Reader goes on. The first Data Ready occurs 1.8 ms later. Subsequent columns are ready every 370 μs — the program must give a DATAI within 350 μs after each setting of Data Ready. Total time from first to last Data Ready is 29.2 ms. After the final Data Ready, 1.8 ms elapse before Card in Reader and Reading Card clear and End of Card sets. The program then has 9.2 ms within which to give a new CONO read card command to keep the reader going at the maximum rate. Ready to Read goes on at the end of this period if no new command appears.

If the reader operates on 50 Hz power, all times must be increased by 20 per cent.

When the last card in a deck is read, the hopper empty signal is simul-

taneous with End of Card.

Operation. The reader has a hopper and stacker capacity of 1000 cards. To load a deck, first fan the cards and jog them on the reader shelf. Turn the deck over and put the first hundred cards (about an inch of the deck) into the hopper (upper right) with the 9 edge against the back so column 1 is read first. Place the rest of the deck on top of the first part. Cards can be added to the hopper while the reader is running, but always stop the reader before removing cards from the stacker.

The reader is operated by the buttons at the left. The alternate-action POWER switch lights green when power is on. Pushing START places the reader on line so the program can read cards. Pushing STOP turns off the reader, taking it off line.

Card Reader

The lights at the right indicate an empty hopper, a full stacker, a pick failure, a card motion error, and a photocell output that is too weak or too strong. When one of these lights goes on the STOP light also goes on (the reader always finishes the current card before stopping). Do not attempt to reread a worn or damaged card that has caused a pick failure or motion error — duplicate it first. If any trouble light remains on after the problem is corrected press the CLEAR button; this turns off both the lights and the corresponding status signals read by a CONI. Press START to allow the program to continue reading the deck. If the trouble persists, enter it in the system log and notify maintenance personnel.

Pressing the END OF FILE button (at the right) when the reader is off line, as when the hopper is empty, sets the End of File flag. When the TEST MODE light is on, the reader processes cards off line (the test switch is behind the panel under the shelf).

Lights for the interface are in the bottom two rows on the hardcopy control indicator panel [*page 4-9*]. The left section of the upper row displays the contents of the card column buffer; the lights are marked by card row. The left section of the bottom row displays bits 24–35 of the status conditions. (The second light from the left is labeled HOP EMPTY, but it goes on when the hopper is empty or the stacker is full.) Of the five lights in the center, the left one is the momentary offset signal. READ is on when a read command has been given but the reader is not yet ready. The next two lights display bits 18 and 19 of the status conditions, and the last light is on while an interrupt is being requested whatever the cause.
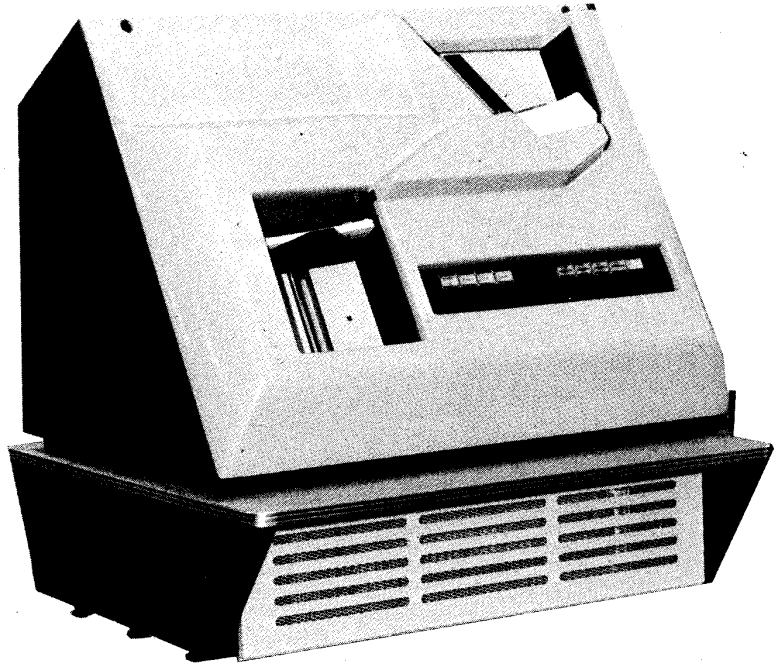
Also available is a console model reader that has a 2000-card hopper and two 2000-card stackers. In use it differs from the compact model only in that offsetting a card places it in the second stacker (the one on the right), and cards can be removed from the stackers while the reader is running.

4.4 CARD PUNCH

The card punch handles standard 12-row 80-column cards at speeds up to 200 cards per minute if all eighty columns are punched, 365 cards per minute if only the first sixteen columns are punched. The processor must supply each column to the punch as twelve bits, and the program can generate this data by any procedure it wishes; the standard DEC character representations and the translation from ASCII made by the Monitor are given in Appendix B. Of course the data can simply be in binary at three columns per word (punching rows 7 and 9 in the first column is the standard procedure for indicating that the rest of the card contains binary data).

A card is taken from the hopper only when the program supplies data for the first column. In the interface is a 12-bit buffer to which the processor sends each column, but the punch has a 48-bit buffer, and it punches four columns at a time from each set of four 12-bit bytes sent through the interface. The program can send a card to the stacker after punching any number of columns. The punch device code is 110, mnemonic CDP.

CONO CDP,   **Conditions Out, Card Punch**

7 1 1 2 0	I	X	Y
0	12 13 14 17 18		35

Assign the interrupt channel specified by bits 33–35 of the effective conditions E and perform the functions specified by bits 20–32 as shown (a 1 in a bit produces the indicated function, a 0 has no effect).

		CLEAR PUNCH	OFFSET CARD		EJECT CARD	DISABLE	ENABLE TROUBLE INTERRUPTS	CLEAR ERROR	DISABLE	ENABLE END OF CARD	CLEAR END OF CARD	SET PUNCH ON	CLEAR	SET DATA REQUEST	PRIORITY INTERRUPT ASSIGNMENT		
18	19	20	21	22	23	24	25	26	27	28	29	30	31	32	33	34	35

Notes.

20   Clear flags Trouble Interrupt Enabled, Error, End of Card Enabled, End of Card, Punch On, Busy, Data Request; clear the card column buffer. If any action specified by the rest of the CONO bits conflicts with these actions, the other bits have precedence.

21   If a card is currently being processed in the punch (Card in Punch, CONI bit 27, is 1) or was ejected less than 3 ms ago, offset it when it is placed in the stacker. The card will actually stick out about a half inch from the rest of the stacked deck.

23   If a card is currently being processed (Card in Punch, CONI bit 27, is 1), punch whatever data is in the 4-column buffer and then eject the card. Ejection moves a card through the punch head assembly four times as fast as punching blank columns.

> With the console model, off-setting a card places it in a separate stacker.

CONI CDP,   Conditions In, Card Punch

7 1 1 2 4	I	X	Y
0	12 13 14	17 18	35

Read the status of the punch into the right half of location E as shown.

TEST		HOPPER LOW	NEED OPERATOR SERVICE /	PICK FAILURE-STACK FAILURE	EJECT FAILURE	TROUBLE	TROUBLE INTERRUPT ENABLED /	ERROR	CARD IN PUNCH	END OF CARD ENABLED	END OF CARD	PUNCH ON	BUSY	DATA REQUEST	PRIORITY INTERRUPT ASSIGNMENT		
18	19	20	21	22	23	24	25	26	27	28	29	30	31	32	33	34	35

*These bits cause interrupts

Notes.

Interrupts are requested on the assigned channel by the setting of Data Request, End of Card, and if enabled, Trouble.

18   The operator has turned on the test switch, taking the punch off line.

20   Less than a hundred cards are left in the hopper.

21   The hopper is empty or the stacker or chip box is full.

22   The punch has received data for the first column but has failed to bring in a card from the hopper; or it has received an eject command but has failed to place the card properly in the stacker.

23   The punch has received an eject command but has failed to move the card out of the punch head assembly.

24   Bit 18, 22 or 23 is 1, or bit 21 is 1 because the hopper is empty or the stacker is full, or the operator has taken the punch off line. If bit 25 is also 1, the setting of Trouble requests an interrupt on the assigned channel.

     Ordinarily a trouble condition allows the punch to finish a card but prevents it from starting another; only a power failure or the operator turning on the test switch (bit 18) can take the punch off line in the middle of a card. A full chip box does not stop the punch at all as there is actually enough room left for the chips from a whole

hopper full of cards. Although a 1 in bit 24 does not necessarily imply a malfunction, it always requires operator intervention. If bit 21 is 1 it is very likely that the only trouble is the hopper is empty or the stacker is full.

26   A column punched in a card does not agree with the data sent by the processor.

27   A card is in the punch head assembly. The program can give a CONO offset or eject command while this bit is 1 (the offset can also be given within 3 ms after Card in Punch clears).

29   Bit 28 is 1 and the program has given either an eject command or data for column 80. The setting of End of Card requests an interrupt on the assigned channel.

DATAO CDP,   Data Out, Card Punch

7 1 1 1 4	I	X	Y
0	12 13 14	17 18	35

Clear Data Request, set Punch On and Busy, and load the contents of bits 24–35 of location E into the interface column buffer where the correspondence of bit positions to card rows is as shown.

ROW 12	ROW 11	ROW 0	ROW 1	ROW 2	ROW 3	ROW 4	ROW 5	ROW 6	ROW 7	ROW 8	ROW 9
24	25	26	27	28	29	30	31	32	33	34	35

Setting Punch On turns on the punch motor, but only a DATAO can pick a card. Since DATAO also sets Punch On, the program can initiate punch operations while supplying data, but the usual procedure is to set Punch On while giving other initial conditions.

When the punch is ready to take a card from the hopper it sends a ready signal to the interface. This sets Data Request, which requests an interrupt on the assigned channel. To pick a card the program must respond with a DATAO, which supplies the first column, clears Data Request, and sets Punch On and Busy. The interface then sends the column to the 4-column punch buffer and clears Busy. While the punch is picking a card it also makes three more data requests to each of which the program must respond with a DATAO. When the card is properly registered in the punch head assembly, Card in Punch sets. When this flag has set and the program has supplied the first four columns, the device punches the four columns simultaneously (a 1 sent to the column buffer produces a hole in the card). The punch then continues in this fashion making four data requests for each set of four columns.

Punch On clears when the program gives an eject command. This causes the device to punch whatever is in its 4-column buffer and the card then

If the program gives a DATAO to turn on the motor, the initial ready from the punch takes the first column from the column buffer but does not set Data Request. When that flag does set, the punch is ready for the second column.

moves out to the stacker. If the punch has already sent a ready signal, the CONO that ejects should also clear Data Request. If End of Card has been enabled by a 1 in CONO bit 28, the eject command sets it, requesting an interrupt. If no eject command has been given by the time data is supplied for column 80, End of Card sets anyway if it is enabled (producing an interrupt request), but the card remains in the punch head assembly until an eject command is given. The actual ejection of a card clears Card in Punch.

Timing. If Punch On is set when the punch motor is off, the first ready signal is delayed about 120 ms while the motor gets up to speed. When the motor is on, Card in Punch sets about 60 ms after the DATAO that sends the first column for a card. While the card is in the head assembly, punching is synchronized to a punch cycle of 11.1 ms. About 30 μs elapse from each DATAO to the next Data Request, but after the first request the program has the full punch cycle time to supply all four columns and keep punching at the maximum rate; after that punching is delayed until the next cycle.

Giving an eject command clears Punch On and sets End of Card after 5 μs, but Card in Punch does not clear until the card leaves the head assembly; this takes about 25 ms plus 2.8 ms for each set of four columns skipped over. After Card in Punch clears, about 30 μs elapse before the punch indicates that it is ready to pick another card from the hopper, at which time the program should give a DATAO to pick another card at the maximum rate. (Of course the first DATAO can be given right after the eject command, and the punch will then pick another card automatically without setting Data Request for the first column.) When the final card is punched, the hopper empty signal is simultaneous with End of Card. If Punch On remains clear for about 30 seconds, the motor turns off.

Operation. The punch has a hopper and stacker capacity of 1000 cards. To load the hopper, first fan the cards and jog them on the punch shelf. Turn the deck over and put the first hundred cards (about an inch of the deck) into the hopper (upper right) with the 9 edge against the back so column 1 is punched first. Hold the right end higher so the leading edge of the bottom card rests against the picker throat, and drop the cards in place. Put the rest of the deck on top of the first part. Cards can be added to the hopper while the punch is running, but always stop the punch before removing cards from the stacker. To remove cards, push down the elevator and lift the stack out.

If the program does not eject before the punch starts punching columns 77–80, it makes another data request. The program can then supply two more columns, which will be punched in the margin of the card.

CAUTION

Any data that is given but not punched (*eg* the first column(s) when there is a pick failure) is usually lost when the punch goes off line. Hence the program should always start with the first column of a card when the punch is restarted.
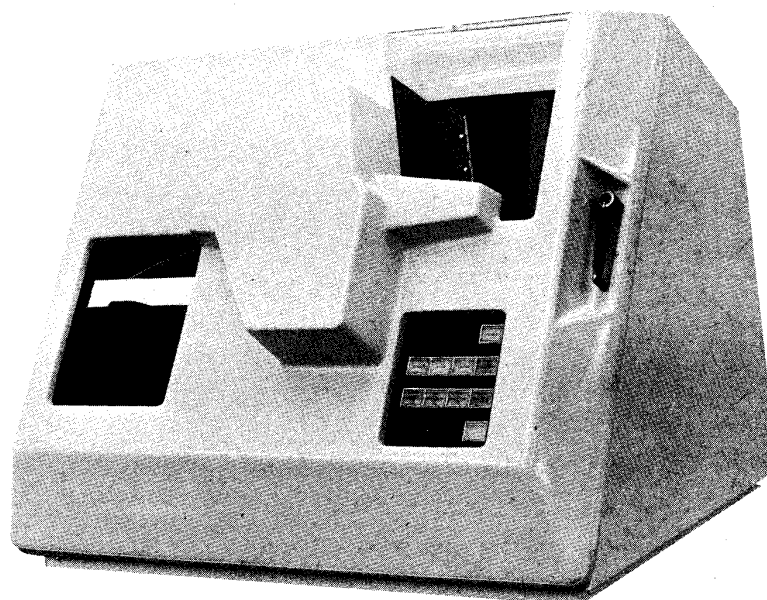
Card Punch

The punch is operated by the buttons in the upper part of the panel at the right. The alternate-action POWER switch lights green when power is on. Pushing START places the punch on line so the program can punch cards. The OPERATE indicator at the lower right lights green when the punch motor is up to speed. Pushing STOP takes the reader off line but does not stop the motor; the motor is turned off only by pressing CLEAR.

The lights in the bottom row indicate an empty hopper or full stacker, a full chip box, a pick failure, an eject failure, and a stack failure. When one of these lights other than CHIP BOX goes on, the STOP light also goes on (the punch always finishes the current card before stopping). If any trouble light remains on after the problem is corrected, press the CLEAR button; this turns off both the lights and the corresponding status signals read by a CONI. Pressing CLEAR also ejects a card if one is in the head assembly, and the button glows red when clear action is required (*eg* when a card has gotten stuck). For a pick failure, empty the hopper, throw out the bottom card, and reload. Press START to allow the program to continue punching. If the trouble persists enter it in the system log and notify maintenance personnel.

A full chip box does not stop the punch, but once it has been stopped by some other condition (such as an empty hopper), pressing START will not place the unit on line until the box has been emptied.

At the right is a light for the Card in Punch flag. The ERROR light displays the signal that sets the Error flag; it goes off when CLEAR is pressed. The CHECKOFF light is not used. When the TEST light is on, the device punches cards off line in a test pattern (the test switch is behind the panel under the shelf).

Lights for the interface are in the second row from the bottom on the hardcopy control indicator panel [*page 4-9*]. The middle section of the row displays the contents of the card column buffer; the lights are marked by card row. Among the lights in the right section, PI REQ is on while an interrupt is being requested whatever the cause. The remaining lights display some of the status conditions read by a CONI.

Also available is a console model punch that has a 2000-card hopper and two 2000-card stackers. In use it differs from the compact model only in that offsetting a card places it in the second stacker (the one on the right), and cards can be removed from the stackers while the punch is running.

Appendices

APPENDIX A

INSTRUCTION AND DEVICE MNEMONICS

The illustration on the next page shows the derivation of the instruction mnemonics. The two tables following it list all instruction mnemonics and their octal codes both numerically and alphabetically. When two mnemonics are given for the same octal code, the first is the preferred form, but the assembler does recognize the second. For completeness, UUOs are listed for user mode (an asterisk indicates a UUO mnemonic recognized by MACRO for communication with the PDP-10 Time Sharing Monitor). All UUOs 000-077 are identical when the processor is not in user mode.

In-out device codes are included only in the alphabetic listing and are indicated by a dagger (†). Following the tables is a chart that lists the devices with their mnemonic and octal codes and DEC option numbers for both PDP-10 and PDP-6. A device mnemonic ending in the numeral 2 is the recommended form for the second of a given device, but such codes are not recognized by MACRO — they must be defined by the user.

A1

A2   MNEMONICS

MOV { E / e Negative / e Magnitude / e Swapped } ──── { to AC / Immediate to AC / to Memory / to Self } Half word { Right / Left } to { Right / Left } { no effect / Ones / Zeros / Extend sign } BLock Transfer EXCHange AC and memory	ADD / SUBtract / MULtiply / Integer MULtiply / DIVide / Integer DIVide } ┐ └─and Round─┘ { ~ / Immediate / to Memory / to Both } Floating AdD / Floating SuBtract / Floating MultiPly / Floating DiVide } { ~ / Long / to Memory / to Both } Floating SCale Double Floating Negate Unnormalized Floating Add
use present pointer / Increment pointer } and { LoaD Byte into AC / DePosit Byte in memory } Increment Byte Pointer	
	Arithmetic SHift / Logical SHift / ROTate } { ~ / Combined }
PUSH down / POP up } { ~ / and Jump }	
SET to { Zeros / Ones / Ac / Memory / Complement of Ac / Complement of Memory } AND / inclusive OR } { ~ / with Complement of Ac / with Complement of Memory / Complements of Both } ─to { AC / AC Immediate / Memory / Both } Inclusive OR / eXclusive OR / EQuiValence }	Jump { to SubRoutine / and Save Pc / and Save Ac / and Restore Ac / if Find First One / on Flag and CLear it / on OVerflow (JFCL 10,) / on CaRrY 0 (JFCL 4,) / on CaRrY 1 (JFCL 2,) / on CaRrY (JFCL 6,) / on Floating OVerflow (JFCL 1,) / and ReSTore / and ReSTore Flags (JRST 2,) / and ENable PI channel (JRST 12,) } HALT (JRST 4,) eXeCuTe
SKIP if memory / JUMP if AC } ──── if { never / Less / Equal / Less or Equal / Always / Greater / Greater or Equal / Not equal } Add One to / Subtract One from } { memory and Skip / AC and Jump } if Compare Ac { Immediate / with Memory } and skip if AC─ Add One to Both halves of AC and Jump if { Positive / Negative }	DATA / BLocK } { In / Out } CONditions─ / └─in and Skip if { all masked bits Zero / some masked bit One }
Test AC { with Direct mask / with Swapped mask / Right with E / Left with E } { No modification / set masked bits to Zeros / set masked bits to Ones / Complement masked bits } and skip { never / if all masked bits Equal 0 / if Not all masked bits equal 0 / Always }	

INSTRUCTION MNEMONICS

NUMERIC LISTING

000	ILLEGAL	132	FSC	206	MOVSM
001		133	IBP	207	MOVSS
.	USER	134	ILDB	210	MOVN
:	UUO'S	135	LDB	211	MOVNI
037		136	IDPB	212 .	MOVNM
040	*CALL	137	DPB	213	MOVNS
041	*INIT	140	FAD	214	MOVM
042		141	FADL	215	MOVMI
043	RESERVED	142	FADM	216	MOVMM
044	FOR SPECIAL	143	FADB	217	MOVMS
045	MONITORS	144	FADR	220	IMUL
046		145	FADRI ▲	221	IMULI
047	*CALLI	146	FADRM	222	IMULM
050	*OPEN	147	FADRB	223	IMULB
051	*TTCALL ▲	150	FSB	224	MUL
052		151	FSBL	225	MULI
053	RESERVED FOR DEC	152	FSBM	226	MULM
054		153	FSBB	227	MULB
055	*RENAME	154	FSBR	230	IDIV
056	*IN	155	FSBRI ▲	231	IDIVI
057	*OUT	156	FSBRM	232	IDIVM
060	*SETSTS	157	FSBRB	233	IDIVB
061	*STATO	160	FMP	234	DIV
062	*STATUS	161	FMPL	235	DIVI
062	*GETSTS	162	FMPM	236	DIVM
063	*STATZ	163	FMPB	237	DIVB
064	*INBUF	164	FMPR	240	ASH
065	*OUTBUF	165	FMPRI ▲	241	ROT
066	*INPUT	166	FMPRM	242	LSH
067	*OUTPUT	167	FMPRB	243	JFFO
070	*CLOSE	170	FDV	244	ASHC
071	*RELEAS	171	FDVL	245	ROTC
072	*MTAPE	172	FDVM	246	LSHC
073	*UGETF	173	FDVB	247	
074	*USETI	174	FDVR	250	EXCH
075	*USETO	175	FDVRI ▲	251	BLT
076	*LOOKUP	176	FDVRM	252	AOBJP
077	*ENTER	177	FDVRB	253	AOBJN
100		200	MOVE	254	JRST
.	UNASSIGNED	201	MOVEI	25410	JRSTF
:	CODES	202	MOVEM	25420	HALT
127		203	MOVES	25450	JEN
130	UFA	204	MOVS	255	JFCL
131	DFN	205	MOVSI	25504	JFOV

25510	JCRY1	333	SKIPLE	410	ANDCA		
25520	JCRY0	334	SKIPA	411	ANDCAI		
25530	JCRY	335	SKIPGE	412	ANDCAM		
25540	JOV	336	SKIPN	413	ANDCAB		
256	XCT	337	SKIPG	414	SETM		
257		340	AOJ	415	SETMI		
260	PUSHJ	341	AOJL	416	SETMM		
261	PUSH	342	AOJE	417	SETMB		
262	POP	343	AOJLE	420	ANDCM		
263	POPJ	344	AOJA	421	ANDCMI		
264	JSR	345	AOJGE	422	ANDCMM		
265	JSP	346	AOJN	423	ANDCMB		
266	JSA	347	AOJG	424	SETA		
267	JRA	350	AOS	425	SETAI		
270	ADD	351	AOSL	426	SETAM		
271	ADDI	352	AOSE	427	SETAB		
272	ADDM	353	AOSLE	430	XOR		
273	ADDB	354	AOSA	431	XORI		
274	SUB	355	AOSGE	432	XORM		
275	SUBI	356	AOSN	433	XORB		
276	SUBM	357	AOSG	434	IOR		
277	SUBB	360	SOJ	434	OR		
300	CAI	361	SOJL	435	IORI		
301	CAIL	362	SOJE	435	ORI		
302	CAIE	363	SOJLE	436	IORM		
303	CAILE	364	SOJA	436	ORM		
304	CAIA	365	SOJGE	437	IORB		
305	CAIGE	366	SOJN	437	ORB		
306	CAIN	367	SOJG	440	ANDCB		
307	CAIG	370	SOS	441	ANDCBI		
310	CAM	371	SOSL	442	ANDCBM		
311	CAML	372	SOSE	443	ANDCBB		
312	CAME	373	SOSLE	444	EQV		
313	CAMLE	374	SOSA	445	EQVI		
314	CAMA	375	SOSGE	446	EQVM		
315	CAMGE	376	SOSN	447	EQVB		
316	CAMN	377	SOSG	450	SETCA		
317	CAMG	400	SETZ	451	SETCAI		
320	JUMP	400	CLEAR	452	SETCAM		
321	JUMPL	401	SETZI	453	SETCAB		
322	JUMPE	401	CLEARI	454	ORCA		
323	JUMPLE	402	SETZM	455	ORCAI		
324	JUMPA	402	CLEARM	456	ORCAM		
325	JUMPGE	403	SETZB	457	ORCAB		
326	JUMPN	403	CLEARB	460	SETCM		
327	JUMPG	404	AND	461	SETCMI		
330	SKIP	405	ANDI	462	SETCMM		
331	SKIPL	406	ANDM	463	SETCMB		
332	SKIPE	407	ANDB	464	ORCM		

465	ORCMI	546	HLRM	627	TLZN
466	ORCMM	547	HLRS	630	TDZ
467	ORCMB	550	HRRZ	631	TSZ
470	ORCB	551	HRRZI	632	TDZE
471	ORCBI	552	HRRZM	633	TSZE
472	ORCBM	553	HRRZS	634	TDZA
473	ORCBB	554	HLRZ	635	TSZA
474	SETO	555	HLRZI	636	TDZN
475	SETOI	556	HLRZM	637	TSZN
476	SETOM	557	HLRZS	640	TRC
477	SETOB	560	HRRO	641	TLC
500	HLL	561	HRROI	642	TRCE
501	HLLI	562	HRROM	643	TLCE
502	HLLM	563	HRROS	644	TRCA
503	HLLS	564	HLRO	645	TLCA
504	HRL	565	HLROI	646	TRCN
505	HRLI	566	HLROM	647	TLCN
506	HRLM	567	HLROS	650	TDC
507	HRLS	570	HRRE	651	TSC
510	HLLZ	571	HRREI	652	TDCE
511	HLLZI	572	HRREM	653	TSCE
512	HLLZM	573	HRRES	654	TDCA
513	HLLZS	574	HLRE	655	TSCA
514	HRLZ	575	HLREI	656	TDCN
515	HRLZI	576	HLREM	657	TSCN
516	HRLZM	577	HLRES	660	TRO
517	HRLZS	600	TRN	661	TLO
520	HLLO	601	TLN	662	TROE
521	HLLOI	602	TRNE	663	TLOE
522	HLLOM	603	TLNE	664	TROA
523	HLLOS	604	TRNA	665	TLOA
524	HRLO	605	TLNA	666	TRON
525	HRLOI	606	TRNN	667	TLON
526	HRLOM	607	TLNN	670	TDO
527	HRLOS	610	TDN	671	TSO
530	HLLE	611	TSN	672	TDOE
531	HLLEI	612	TDNE	673	TSOE
532	HLLEM	613	TSNE	674	TDOA
533	HLLES	614	TDNA	675	TSOA
534	HRLE	615	TSNA	676	TDON
535	HRLEI	616	TDNN	677	TSON
536	HRLEM	617	TSNN	70000	BLKI
537	HRLES	620	TRZ	70004	DATAI
540	HRR	621	TLZ	70004	RSW
541	HRRI	622	TRZE	70010	BLKO
542	HRRM	623	TLZE	70014	DATAO
543	HRRS	624	TRZA	70020	CONO
544	HLR	625	TLZA	70024	CONI
545	HLRI	626	TRZN	70030	CONSZ
				70034	CONSO

A6   MNEMONICS

INSTRUCTION MNEMONICS

ALPHABETIC LISTING

†ADC	024	BLT	251	DIVM	236
ADD	270	CAI	300	†DLS	240
ADDB	273	CAIA	304	DPB	137
ADDI	271	CAIE	302	▲ †DPC	250
ADDM	272	CAIG	307	†DSK	170
AND	404	CAIGE	305	†DTC	320
ANDB	407	CAIL	301	†DTS	324
ANDCA	410	CAILE	303	*ENTER	077
ANDCAB	413	CAIN	306	EQV	444
ANDCAI	411	*CALL	040	EQVB	447
ANDCAM	412	*CALLI	047	EQVI	445
ANDCB	440	CAM	310	EQVM	446
ANDCBB	443	CAMA	314	EXCH	250
ANDCBI	441	CAME	312	FAD	140
ANDCBM	442	CAMG	317	FADB	143
ANDCM	420	CAMGE	315	FADL	141
ANDCMB	423	CAML	311	FADM	142
ANDCMI	421	CAMLE	313	FADR	144
ANDCMM	422	CAMN	316	FADRB	147
ANDI	405	†CCI	014	FADRI	145
ANDM	406	†CDP	110	FADRM	146
AOBJN	253	†CDR	114	FDV	170
AOBJP	252	CLEAR	400	FDVB	173
AOJ	340	CLEARB	403	FDVL	171
AOJA	344	CLEARI	401	FDVM	172
AOJE	342	CLEARM	402	FDVR	174
AOJG	347	*CLOSE	070	FDVRB	177
AOJGE	345	CONI	70024	FDVRI	175
AOJL	341	CONO	70020	FDVRM	176
AOJLE	343	CONSO	70034	FMP	160
AOJN	346	CONSZ	70030	FMPB	163
AOS	350	†CPA	000	FMPL	161
AOSA	354	†CR	150	FMPM	162
AOSE	352	DATAI	70004	FMPR	164
AOSG	357	DATAO	70014	FMPRB	167
AOSGE	355	†DC	200	FMPRI	165
AOSL	351	†DCSA	300	FMPRM	166
AOSLE	353	†DCSB	304	FSB	150
AOSN	356	†DF	270	FSBB	153
†APR	000	DFN	131	FSBL	151
ASH	240	†DIS	130	FSBM	152
ASHC	244	DIV	234	FSBR	154
BLKI	70000	DIVB	237	FSBRB	157
BLKO	70010	DIVI	235	FSBRI	155

FSBRM	156	HRLZI	515	JSA	266
FSC	132	HRLZM	516	JSP	265
*GETSTS	062	HRLZS	517	JSR	'264
HALT	25420	HRR	540	JUMP	320
HLL	500	HRRE	570	JUMPA	324
HLLE	530	HRREI	571	JUMPE	322
HLLEI	531	HRREM	572	JUMPG	,327
HLLEM	532	HRRES	573	JUMPGE	325
HLLES	533	HRRI	541	JUMPL	321
HLLI	501	HRRM	542	JUMPLE	323
HLLM	502	HRRO	560	JUMPN	326
HLLO	520	HRROI	561	LDB	135 ▲
HLLOI	521	HRROM	562	*LOOKUP	076
HLLOM	522	HRROS	563	†LPT	124
HLLOS	523	HRRS	543	LSH	242
HLLS	503	HRRZ	550	LSHC	246
HLLZ	510	HRRZI	551	†MDF	260
HLLZI	511	HRRZM	552	MOVE	200
HLLZM	512	HRRZS	553	MOVEI	201
HLLZS	513	IBP	133	MOVEM	202
HLR	544	IDIV	230	MOVES	203
HLRE	574	IDIVB	233	MOVM	214
HLREI	575	IDIVI	231	MOVMI	215
HLREM	576	IDIVM	232	MOVMM	216
HLRES	577	IDPB	136	MOVMS	217
HLRI	545	ILDB	134	MOVN	210
HLRM	546	IMUL	220	MOVNI	211
HLRO	564	IMULB	223	MOVNM	212
HLROI	565	IMULI	221	MOVNS	213
HLROM	566	IMULM	222	MOVS	204
HLROS	567	*IN	056	MOVSI	205
HLRS	547	*INBUF	064	MOVSM	206
HLRZ	554	*INIT	041	MOVSS	207
HLRZI	555	*INPUT	066	*MTAPE	072
HLRZM	556	IOR	434	†MTC	220
HLRZS	557	IORB	437	†MTM	230
HRL	504	IORI	435	†MTS	224
HRLE	534	IORM	436	MUL	224
HRLEI	535	JCRY	25530	MULB	227
HRLEM	536	JCRY0	25520	MULI	225
HRLES	537	JCRY1	25510	MULM	226
HRLI	505	JEN	25460	*OPEN	050
HRLM	506	JFCL	255	OR	434
HRLO	524	JFFO	243	ORB	437
HRLOI	525	JFOV	25504	ORCA	454
HRLOM	526	JOV	25540	ORCAB	457
HRLOS	527	JRA	267	ORCAI	455
HRLS	507	JRST	254	ORCAM	456
HRLZ	514	JRSTF	25410	ORCB	470

A8

ORCBB	473	SETZM	402	TLCA	645
ORCBI	471	SKIP	330	TLCE	643
ORCBM	472	SKIPA	334	TLCN	647
ORCM	464	SKIPE	332	TLN	601
ORCMB	467	SKIPG	337	TLNA	605
ORCMI	465	SKIPGE	335	TLNE	603
ORCMM	466	SKIPL	331	TLNN	607
ORI	435	SKIPLI	333	TLO	661
ORM	436	SKIPN	336	TLOA	665
*OUT	057	SOJ	360	TLOE	663
*OUTBUF	065	SOJA	364	TLON	667
*OUTPUT	067	SOJE	362	TLZ	621
†PI	004	SOJG	367	TLZA	625
†PLT	140	SOJGE	365	TLZE	623
POP	262	SOJL	361	TLZN	627
POPJ	263	SOJLE	363	†TMC	340
†PTP	100	SOJN	366	†TMS	344
†PTR	104	SOS	370	TRC	640
PUSH	261	SOSA	374	TRCA	644
PUSHJ	260	SOSE	372	TRCE	642
*RELEAS	071	SOSG	377	TRCN	646
*RENAME	055	SOSGE	375	TRN	600
ROT	241	SOSL	371	TRNA	604
ROTC	245	SOSLE	373	TRNE	602
RSW	70004	SOSN	376	TRNN	606
SETA	424	*STATO	061	TRO	660
SETAB	427	*STATUS	062	TROA	664
SETAI	425	*STATZ	063	TROE	662
SETAM	426	SUB	274	TRON	666
SETCA	450	SUBB	277	TRZ	620
SETCAB	453	SUBI	275	TRZA	624
SETCAI	451	SUBM	276	TRZE	622
SETCAM	452	TDC	650	TRZN	626
SETCM	460	TDCA	654	TSC	651
SETCMB	463	TDCE	652	TSCA	655
SETCMI	461	TDCN	656	TSCE	653
SETCMM	462	TDN	610	TSCN	657
SETM	414	TDNA	614	TSN	611
SETMB	417	TDNE	612	TSNA	615
SETMI	415	TDNN	616	TSNE	613
SETMM	416	TDO	670	TSNN	617
SETO	474	TDOA	674	TSO	671
SETOB	477	TDOE	672	TSOA	675
SETOI	475	TDON	676	TSOE	673
SETOM	476	TDZ	630	TSON	677
*SETSTS	060	TDZA	634	TSZ	631
SETZ	400	TDZE	632	TSZA	635
SETZB	403	TDZN	636	TSZE	633
SETZI	401	TLC	641	TSZN	637

ALPHABETIC LISTING

*TTCALL	051 ▲	*USETO	075	XOR	430
UFA	130	†UTC	210	XORB	433
*UGETF	073	†UTS	214	XORI	431
*USETI	074	XCT	256	XORM·	432

DEVICE MNEMONICS

FIRST OCTAL DIGIT	00	04	10	14	20	24	30	34	40	44	50	54	60	64	70	74
0	APR CPA (6,10) 761 CENTRAL PROCESSOR	PI (6,10) PRIORITY INTERRUPT	DRUM (6) 167 DRUM PROCESSOR	CCI (10) DA10 PDP-8,9 INTERFACE	CCI2 (10) DA10 PDP-8,9 INTERFACE	ADC (10) AD10 ANALOG-DIGITAL CONVERTER	ADC2 (10) AD10 ANALOG-DIGITAL CONVERTER								DSK (10) RC10 SMALL DISK	DSK2 (10) RC10 SMALL DISK
1	PTP (6,10) 136 PAPER TAPE PUNCH	PTR (6,10) 760 PAPER TAPE READER	CDP (10) CP10 CARD PUNCH	CDR (6) 461 CARD READER	TTY (6) 626 CONSOLE TELETYPE	LPT (6,10) 646 LP10 LINE PRINTER	DIS (6,10) 340 VP10 516 DISPLAY	DIS2 (340 VP10) DISPLAY	PLT (10) XY10 PLOTTER	PLT2 (10) XY10 PLOTTER	CR (10) CR10 CARD READER	CR2 (10) CR10 CARD READER	MDF (10) RA10 165 PDP-7,8 INTERFACE / MASS DISK FILE	MDF2 (10) RA10 MASS DISK FILE	DF (6) 270 DISK FILE	
2	DC (6) DATA CONTROL	DC2 (6) 136 DATA CONTROL	UTC DECTAPE	UTS 551 DECTAPE	MTC (10) MAGNETIC TAPE	MTS MAGNETIC TAPE	MTM MAGNETIC TAPE		DLS (10) DC10 DATA LINE SCANNER	DLS2 (10) DC10 DATA LINE SCANNER	DPC (10) RP10 DISK PACK SYSTEM	DPC2 (10) RP10 DISK PACK SYSTEM				
3	DCSA (6) DATA COMMUNICATION	DCSB 630 DATA COMMUNICATION	DTC (10) DECTAPE	DTS DECTAPE	DTC2 DECTAPE	DTS2 DECTAPE	TMC (10) TM10 MAGNETIC TAPE	TMS TM10 MAGNETIC TAPE	TMC2 TM10 MAGNETIC TAPE	TMS2 TM10 MAGNETIC TAPE						
4																
5																
6																
7																

SECOND AND THIRD OCTAL DIGITS (column headers)

Rows 0–3: DEC STANDARD DEVICES
Rows 4–7: USER SPECIAL DEVICES

IN-OUT INSTRUCTION WORD

INSTRUCTION CODE	DEVICE CODE			INSTRUCTION CODE
	FIRST OCTAL DIGIT	SECOND OCTAL DIGIT	THIRD OCTAL DIGIT (1/58 ONLY)	
1 1 1	1 1 1	1 1 1	1	0
0 1 2	3 4 5	6 7 8	9	10 11 12

Legend (example device code **24**):

- 646 — Option number for PDP-6
- LP10 — Option number for PDP-10 (No number indicates device is part of central processor)
- LPT — Mnemonic for device code 124
- 6 — Used with PDP-6
- 10 — Used with PDP-10
- LPT / LINE PRINTER — Device whose code is 124

APPENDIX B

INPUT-OUTPUT CODES

The table beginning on the next page lists the complete teletype code. The lower case character set (codes 140–176) is not available on the Model 35, but giving one of these codes causes the teletype to print the corresponding upper case character. Other differences between the 35 and 37 are mentioned in the table. The definitions of the control codes are those given by ASCII. Most control codes, however, have no effect on the console teletype, and the definitions bear no necessary relation to the use of the codes in conjunction with the PDP–10 software.

The line printer has the same codes and characters as the teletype. The 64-character printer has the figure and upper case sets, codes 040–137 (again, giving a lower case code prints the upper case character). The "96"-character printer has these plus the lower case set, codes 040–176. The latter printer actually has only ninety-five characters unless a special character is "hidden" under the delete code, 177. A hidden character is printed by sending its code prefixed by the delete code. Hence a character hidden under DEL is printed by sending the printer two 177s in a row.

Besides printing characters, the line printer responds to ten control characters, HT, CR, LF, VT, FF, DLE and DC1–4. The 128-character printer uses the entire set of 7-bit codes for printable characters, with characters hidden under the ten control characters that affect the printer and also under null and delete. In all cases, prefixing DEL causes the hidden character to be printed. The extra thirty-three characters that complete the set are ordered special for each installation.

The first page of the table of card codes [*pages B6–8*] lists the column punch required to represent any character in the two DEC codes. The octal codes listed are those used by the PDP–10 software. In other words, when reading cards, the Monitor translates the column punch into the octal code shown; when punching cards, it produces the listed column punch when given the corresponding code. The remaining pages of the table show the relationship between the DEC card codes and several IBM card punches. Each of the column punches is produced by a single key on any punch for which a character is listed, the character being that which is printed at the top of the card.

INPUT-OUTPUT CODES

TELETYPE CODE

Even Parity Bit	7-Bit Octal Code	Character	Remarks
0	000	NUL	Null, tape feed. Repeats on Model 37. Control shift P on Model 35.
1	001	SOH	Start of heading; also SOM, start of message. Control A.
1	002	STX	Start of text; also EOA, end of address. Control B.
0	003	ETX	End of text; also EOM, end of message. Control C.
1	004	EOT	End of transmission (END); shuts off TWX machines. Control D.
0	005	ENQ	Enquiry (ENQRY); also WRU, "Who are you?" Triggers identification ("Here is . . . ") at remote station if so equipped. Control E.
0	006	ACK	Acknowledge; also RU, "Are you . . . ?" Control F.
1	007	BEL	Rings the bell. Control G.
1	010	BS	Backspace; also FEO, format effector. Backspaces some machines. Repeats on Model 37. Control H on Model 35.
0	011	HT	Horizontal tab. Control I on Model 35.
0	012	LF	Line feed or line space (NEW LINE); advances paper to next line. Repeats on Model 37. Duplicated by control J on Model 35.
1	013	VT	Vertical tab (VTAB). Control K on Model 35.
0	014	FF	Form feed to top of next page (PAGE). Control L.
1	015	CR	Carriage return to beginning of line. Control M on Model 35.
1	016	SO	Shift out; changes ribbon color to red. Control N.
0	017	SI	Shift in; changes ribbon color to black. Control O.
1	020	DLE	Data link escape. Control P (DC0).
0	021	DC1	Device control 1, turns transmitter (reader) on. Control Q (X ON).
0	022	DC2	Device control 2, turns punch or auxiliary on. Control R (TAPE, AUX ON).
1	023	DC3	Device control 3, turns transmitter (reader) off. Control S (X OFF).
0	024	DC4	Device control 4, turns punch or auxiliary off. Control T (TAPE, AUX OFF).
1	025	NAK	Negative acknowledge; also ERR, error. Control U.
1	026	SYN	Synchronous idle (SYNC). Control V.
0	027	ETB	End of transmission block; also LEM, logical end of medium. Control W.
0	030	CAN	Cancel (CANCL). Control X.
1	031	EM	End of medium. Control Y.
1	032	SUB	Substitute. Control Z.
0	033	ESC	Escape, prefix. This code is generated by control shift K on Model 35, but the Monitor translates it to 175.
1	034	FS	File separator. Control shift L on Model 35.
0	035	GS	Group separator. Control shift M on Model 35.

Even Parity Bit	7-Bit Octal Code	Character	Remarks
0	036	RS	Record separator. Control shift N on Model 35.
1	037	US	Unit separator. Control shift O on Model 35.
1	040	SP	Space.
0	041	!	
0	042	"	
1	043	#	
0	044	$	
1	045	%	
1	046	&	
0	047	'	Accent acute or apostrophe.
0	050	(	
1	051	)	
1	052	*	Repeats on Model 37.
0	053	+	
1	054	,	
0	055	–	Repeats on Model 37.
0	056	.	Repeats on Model 37.
1	057	/	
0	060	Ø	
1	061	1	
1	062	2	
0	063	3	
1	064	4	
0	065	5	
0	066	6	
1	067	7	
1	070	8	
0	071	9	
0	072	:	
1	073	;	
0	074	<	
1	075	=	Repeats on Model 37.
1	076	>	
0	077	?	
1	100	@	
0	101	A	
0	102	B	

B4

Even Parity Bit	7-Bit Octal Code	Character	Remarks
1	103	C	
0	104	D	
1	105	E	
1	106	F	
0	107	G	
0	110	H	
1	111	I	
1	112	J	
0	113	K	
1	114	L	
0	115	M	
0	116	N	
1	117	O	
0	120	P	
1	121	Q	
1	122	R	
0	123	S	
1	124	T	
0	125	U	
0	126	V	
1	127	W	
1	130	X	Repeats on Model 37.
0	131	Y	
0	132	Z	
1	133	[	Shift K on Model 35.
0	134	\	Shift L on Model 35.
1	135	]	Shift M on Model 35.
1	136	↑	
0	137	←	Repeats on Model 37.
0	140	`	Accent grave.
1	141	a	
1	142	b	
0	143	c	
1	144	d	
0	145	e	
0	146	f	
1	147	g	

Even Parity Bit	7-Bit Octal Code	Character	Remarks
1	150	h	
0	151	i	
0	152	j	
1	153	k	
0	154	l	
1	155	m	
1	156	n	
0	157	o	
1	160	p	
0	161	q	
0	162	r	
1	163	s	
0	164	t	
1	165	u	
1	166	v	
0	167	w	
0	170	x	Repeats on Model 37.
1	171	y	
1	172	z	
0	173	{	▲
1	174	|	
0	175	}	This code generated by ALT MODE on Model 35. ▲
0	176	~	This code generated by ESC key (if present) on Model 35, but the Monitor translates it to 175.
1	177	DEL	Delete, rub out. Repeats on Model 37.

Keys That Generate No Codes

REPT	Model 35 only: causes any other key that is struck to repeat continuously until REPT is released.
PAPER ADVANCE	Model 37 local line feed.
LOCAL RETURN	Model 37 local carriage return.
LOC LF	Model 35 local line feed.
LOC CR	Model 35 local carriage return.
INTERRUPT, BREAK	Opens the line (machine sends a continuous string of null characters).
PROCEED, BRK RLS	Break release (not applicable).
HERE IS	Transmits predetermined 21-character message.

B6

CARD CODES

▲ Character	PDP–10 ASCII	DEC 029	DEC 026	Character	PDP–10 ASCII	DEC 029	DEC 026
Space	040	None	None	@	100	8 4	8 4
!	041	11 8 2	12 8 7	A	101	12 1	12 1
"	042	8 7	0 8 5	B	102	12 2	12 2
#	043	8 3	0 8 6	C	103	12 3	12 3
$	044	11 8 3	11 8 3	D	104	12 4	12 4
%	045	0 8 4	0 8 7	E	105	12 5	12 5
&	046	12	11 8 7	F	106	12 6	12 6
'	047	8 5	8 6	G	107	12 7	12 7
(	050	12 8 5	0 8 4 ▲	H	110	12 8	12 8
)	051	11 8 5	12 8 4 ▲	I	111	12 9	12 9
*	052	11 8 4	11 8 4	J	112	11 1	11 1
+	053	12 8 6	12	K	113	11 2	11 2
,	054	0 8 3	0 8 3	L	114	11 3	11 3
–	055	11	11	M	115	11 4	11 4
.	056	12 8 3	12 8 3	N	116	11 5	11 5
/	057	0 1	0 1	O	117	11 6	11 6
0	060	0	0	P	120	11 7	11 7
1	061	1	1	Q	121	11 8	11 8
2	062	2	2	R	122	11 9	11 9
3	063	3	3	S	123	0 2	0 2
4	064	4	4	T	124	0 3	0 3
5	065	5	5	U	125	0 4	0 4
6	066	6	6	V	126	0 5	0 5
7	067	7	7	W	127	0 6	0 6
8	070	8	8	X	130	0 7	0 7
9	071	9	9	Y	131	0 8	0 8
:	072	8 2	11 8 2 or 11 0	Z	132	0 9	0 9
;	073	11 8 6	0 8 2	[	133	12 8 2	11 8 5
<	074	12 8 4	12 8 6	\	134	11 8 7	8 7
=	075	8 6	8 3	]	135	0 8 2	12 8 5
>	076	0 8 6	11 8 6	↑	136	12 8 7	8 5
?	077	0 8 7	12 8 2 or 12 0	←	137	0 8 5	8 2

Binary	7 9
Mode Switch	12 0 2 4 6 8
End of File	12 11 0 1

The octal codes given above are those generated by the Monitor from the column punches. The card reader interface actually supplies a direct binary equivalent of the column punch, as listed in the following two pages.

Column Punch	Character	Octal
None	Space	0000
0	0	1000
1	1	0400
2	2	0200
3	3	0100
4	4	0040
5	5	0020
6	6	0010
7	7	0004
8	8	0002
9	9	0001
12 1	A	4400
12 2	B	4200
12 3	C	4100
12 4	D	4040
12 5	E	4020
12 6	F	4010
12 7	G	4004
12 8	H	4002

Column Punch	Character	Octal
12 9	I	4001
11 1	J	2400
11 2	K	2200
11 3	L	2100
11 4	M	2040
11 5	N	2020
11 6	O	2010
11 7	P	2004
11 8	Q	2002
11 9	R	2001
0 1	/	1400
0 2	S	1200
0 3	T	1100
0 4	U	1040
0 5	V	1020
0 6	W	1010
0 7	X	1004
0 8	Y	1002
0 9	Z	1001

Column Punch	026 Data Processing	026 Fortran	029	DEC 026	DEC 029	Octal
12	&	+	&	+	&	4000
11	−	−	−	−	−	2000
12 0				?		5000
11 0				:		3000
8 2			:	←	:	0202
8 3	#	=	#	=	#	0102
8 4	@	−	@	@	@	0042
8 5			'	↑	'	0022
8 6			=	'	=	0012
8 7			"	\	"	0006
12 8 2			¢	?	[	4202
12 8 3	.	.	.	.	.	4102
12 8 4	¤	)	<	)	<	4042
12 8 5			(	]	(	4022
12 8 6			+	<	+	4012

B8 INPUT-OUTPUT CODES

Column Punch	026 Data Processing	026 Fortran	029	DEC 026	DEC 029	Octal
12 8 7			|	!	↑	4006
11 8 2			!	:	!	2202
11 8 3	$	$	$	$	$	2102
11 8 4	*	*	*	*	*	2042
11 8 5 .			)	[	)	2022
11 8 6			;	>	;	2012
11 8 7			¬	&	\	2006
0 8 ≥			*See note*	;	]	1202
0 8 3	,	,	,	,	,	1102
0 8 4	%	(	%	(	%	1042
0 8 5			←	"	←	1022
0 8 6			>	#	>	1012
0 8 7			?	%	?	1006
12 11 0 1				*End of File*	*End of File*	7400
12 0 2 4 6 8				*Mode Switch*	*Mode Switch*	5252
7 9				*Binary*	*Binary*	xx05

NOTE: There is a single key for the 0 8 2 punch on the 029 but printing is suppressed.

The Monitor translates the octal code for the 12 0 punch in DEC 026 to 4202 (which corresponds to a 12 8 2 punch), and the code for 11 0 to 2202 (11 8 2).

APPENDIX C

MISCELLANY

START

INSTRUCTION
FETCH

INTERRUPT
REQUEST
?   YES

NO

INDIRECT   ADDRESS
CALCULATION

DATA
FETCH

POINTER DONE IN   INSTRUCTION   BLT   INTERRUPT   NO
BYTE, BLKI, BLKO   EXECUTION   REQUEST
                                                            ?
                                                                     YES

DATA
STORE

INSTRUCTION   FLOW   SIMPLIFIED

BASIC INSTRUCTIONS

INSTRUCTION CODE (INCLUDING MODE)	A, F	I	X	Y
0 8	9 12	13 14	17 18	35

IN-OUT INSTRUCTIONS

1 1 1	DEVICE CODE	INSTRUCTION CODE	I	X	Y
0 2 3	9	10 12	13 14	17 18	35

PC WORD

FLAGS	0 0 0 0 0	PC
0	12 13 17	18 35

OVERFLOW	CARRY 0	CARRY 1	FLOATING OVERFLOW	BYTE INTERRUPT	USER	USER IN-OUT				FLOATING UNDER-FLOW	NO DIVIDE
0	1	2	3	4	5	6	7	8	9	10	11	12

BLT POINTER [XWD]

SOURCE ADDRESS	DESTINATION ADDRESS
0 17	18 35
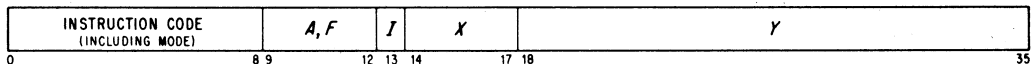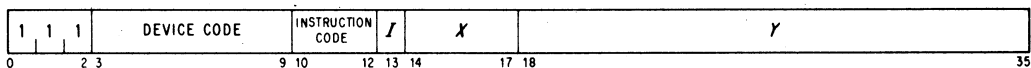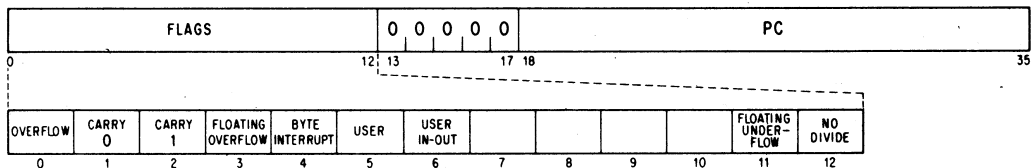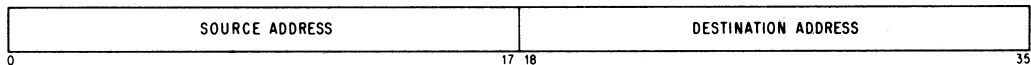
BLKI / BLKO POINTER, PUSHDOWN POINTER, DATA CHANNEL CONTROL WORD [IOWD]
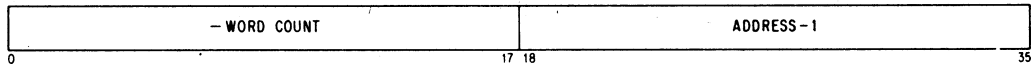
— WORD COUNT	ADDRESS − 1
0 17	18 35

BYTE POINTER

POSITION P	SIZE S		I	X	Y
0 5 6	11 12	13 14	17 18	35

BYTE STORAGE

	S BITS —	P BITS —	
	BYTE	NEXT BYTE	
0	35-P-S-1 35-P	35-P+1	35

FIXED POINT OPERANDS

SIGN 0+ 1−	BINARY NUMBER (TWOS COMPLEMENT)
0 1	35

FLOATING POINT OPERANDS

SIGN 0+ 1−	EXCESS 128 EXPONENT (ONES COMPLEMENT)	FRACTION (TWOS COMPLEMENT)
0 1	8 9	35
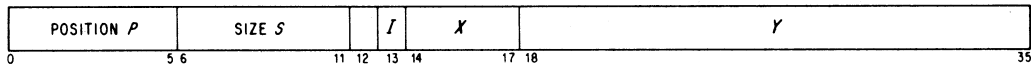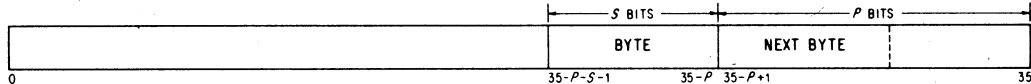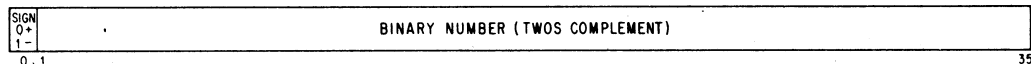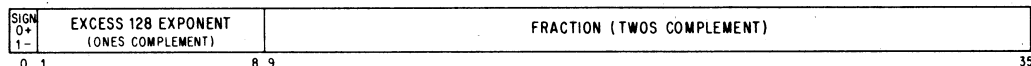
LOW ORDER WORD IN DOUBLE LENGTH FLOATING POINT OPERANDS

0	EXCESS 128 EXPONENT-27 IN POSITIVE FORM	LOW ORDER HALF OF FRACTION (TWOS COMPLEMENT)
0 1	8 9	35
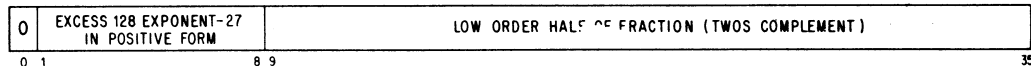
WORD FORMATS

C4
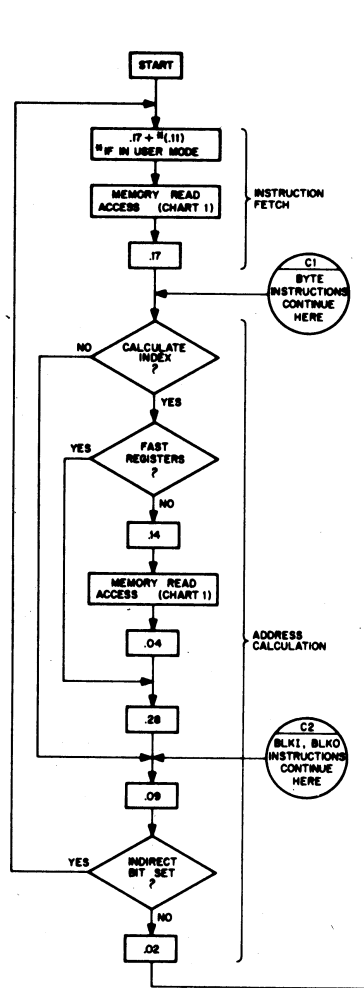
**INSTRUCTION TIMING
FLOW CHART**

INSTRUCTIONS THAT USE READ/MODIFY

All Boolean In Memory and Both Modes Except SETZ, SETA, SETCA, SETO
ADDM, ADDB, SUBM, SUBB
HRRM, HRLM, HLRM, HLLM and All Halt Words In Self Mode
MOVES, MOVNS, MOVMS, MOVSS
ILDB, IDPB (First Time Only)
IBP, BLKI, BLKO, DFN, EXCH
AOS, SOS in all modes

10-0129

INSTRUCTION EXECUTION

DATA STORE

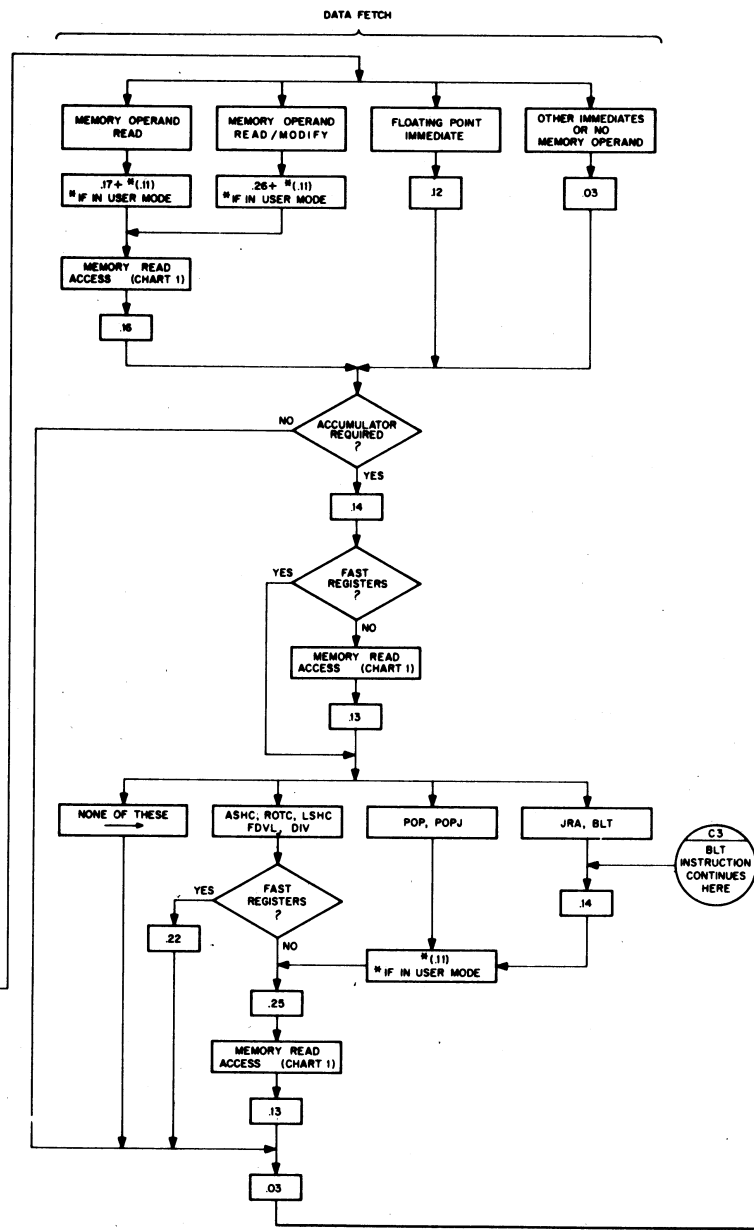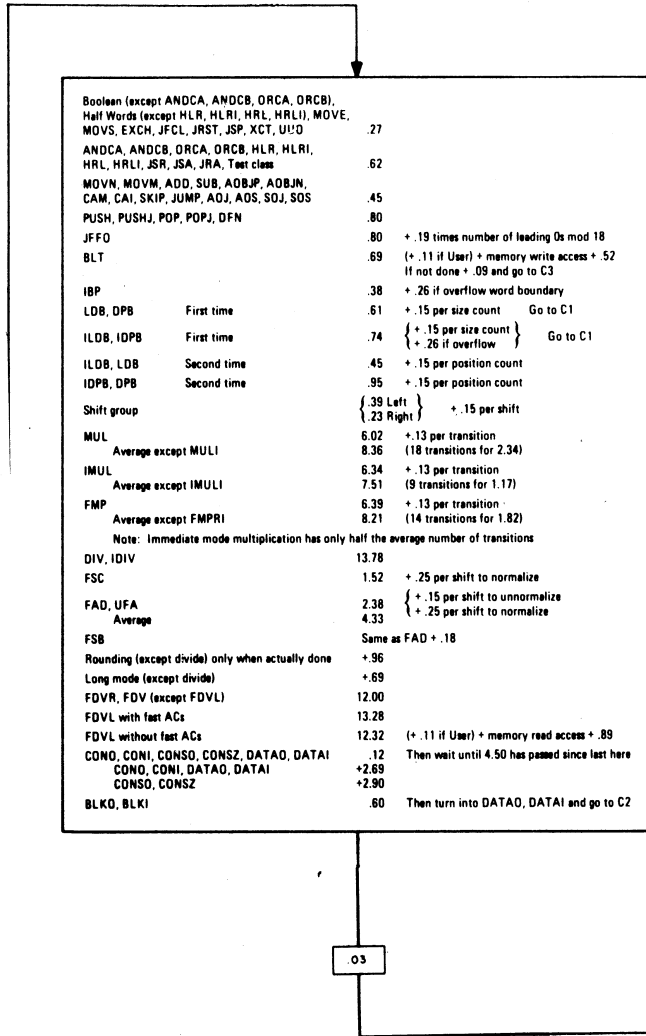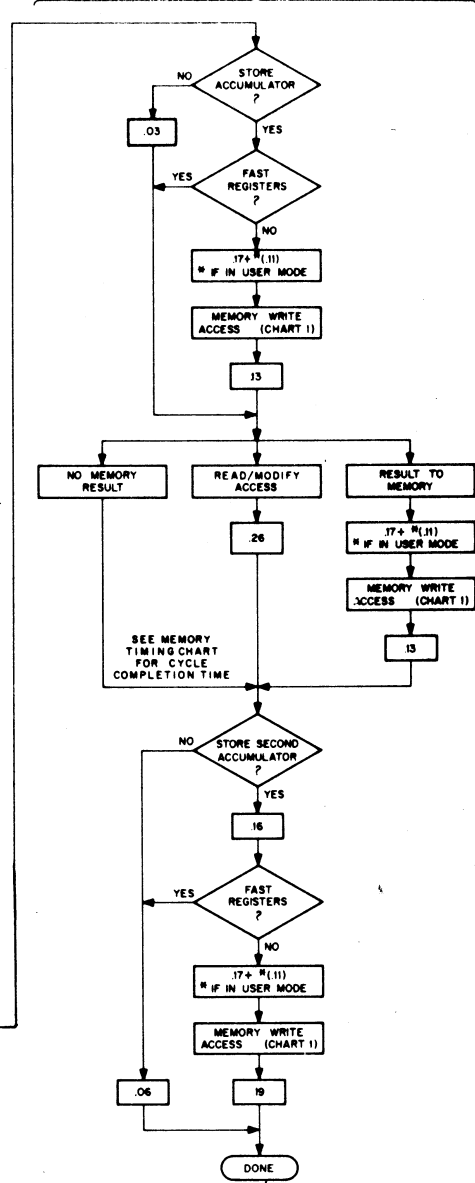|	
Boolean (except ANDCA, ANDCB, ORCA, ORCB), Half Words (except HLR, HLRI, HRL, HRLI), MOVE, MOVS, EXCH, JFCL, JRST, JSP, XCT, UUO	.27
ANDCA, ANDCB, ORCA, ORCB, HLR, HLRI, HRL, HRLI, JSR, JSA, JRA, Test class	.62
MOVN, MOVM, ADD, SUB, AOBJP, AOBJN, CAM, CAI, SKIP, JUMP, AOJ, AOS, SOJ, SOS	.45
PUSH, PUSHJ, POP, POPJ, DFN	.80
JFFO	.80 + .19 times number of leading 0s mod 18
BLT	.69 (+ .11 if User) + memory write access + .52 If not done + .09 and go to C3
IBP	.38 + .26 if overflow word boundary
LDB, DPB First time	.61 + .15 per size count Go to C1
ILDB, IDPB First time	.74 { + .15 per size count / + .26 if overflow } Go to C1
ILDB, LDB Second time	.45 + .15 per position count
IDPB, DPB Second time	.95 + .15 per position count
Shift group	{ .39 Left / .23 Right } + .15 per shift
MUL	6.02 + .13 per transition
Average except MULI	8.36 (18 transitions for 2.34)
IMUL	6.34 + .13 per transition
Average except IMULI	7.51 (9 transitions for 1.17)
FMP	6.39 + .13 per transition
Average except FMPRI	8.21 (14 transitions for 1.82)
Note: Immediate mode multiplication has only half the average number of transitions	
DIV, IDIV	13.78
FSC	1.52 + .25 per shift to normalize
FAD, UFA	2.38 { + .15 per shift to unnormalize
Average	4.33 + .25 per shift to normalize }
FSB	Same as FAD + .18
Rounding (except divide) only when actually done	+.96
Long mode (except divide)	+.69
FDVR, FDV (except FDVL)	12.00
FDVL with fast ACs	13.28
FDVL without fast ACs	12.32 (+ .11 if User) + memory read access + .89
CONO, CONI, CONSO, CONSZ, DATAO, DATAI	.12 Then wait until 4.50 has passed since last here
CONO, CONI, DATAO, DATAI	+2.69
CONSO, CONSZ	+2.90
BLKO, BLKI	.60 Then turn into DATAO, DATAI and go to C2

.03

STORE ACCUMULATOR ? —NO→ .03
↓ YES
FAST REGISTERS ? —YES→
↓ NO
.17 + *(.11) * IF IN USER MODE
↓
MEMORY WRITE ACCESS   (CHART 1)
↓
.13

NO MEMORY RESULT

READ/MODIFY ACCESS
↓
.26

RESULT TO MEMORY
↓
.17 + *(.11) * IF IN USER MODE
↓
MEMORY WRITE ACCESS   (CHART 1)
↓
.13

SEE MEMORY TIMING CHART FOR CYCLE COMPLETION TIME

STORE SECOND ACCUMULATOR ? —NO→
↓ YES
.16
↓
FAST REGISTERS ? —YES→
↓ NO
.17 + *(.11) * IF IN USER MODE
↓
MEMORY WRITE ACCESS   (CHART 1)

.06   .19

DONE

MEMORY TIMING

MEMORY	MA10	MB10	MB10	KM10
PROCESSORS	SINGLE OR MULTI	SINGLE	MULTI	SINGLE (BUILT IN)
CYCLE	1.00	1.65	1.65	—
READ ACCESS	.50	.60	.70	.21
WRITE ACCESS	.20	.20	.30	.21
MODIFY COMPLETION	.35	.35	1.20	—

NOTES :
    MEMORY ACCESS TIMES INCLUDE 20 FEET OF CABLE DELAY.
    ALL TIMES ±5%

IN-OUT DEVICE BIT ASSIGNMENTS

DEVICE	CODE	FUNCTION	0/18	1/19	2/20	3/21	4/22	5/23	6/24	7/25	8/26	9/27	10/28	11/29	12/30	13/31	14/32	15/33	16/34	17/35
APR	000	CONO	CLR PDL OV	I/O RESET		CLR ADR BREAK FLAG	CLR MEM PROT FLAG	CLR NXM FLAG	CLR CLOCK ENABLE	SET CLOCK ENABLE	CLR CLOCK FLAG	CLR FOV ENABLE	SET FOV ENABLE	CLR FOV FLAG	CLR AROV ENABLE	SET AROV ENABLE	CLR AROV FLAG	PIA		
		CONI		PDL OV FLAG	TGT USER FLAG	ADR BREAK FLAG	MEM PROT FLAG	NXM FLAG	CLOCK ENABLE	CLOCK ENABLE	CLOCK FLAG	FOV ENABLE	FOV ENABLE	FOV FLAG	AROV ENABLE	TRAP OFFSET	AROV FLAG	PIA		
		DATAO	PROTECTION REGISTER (LH 0-7), RELOCATION REGISTER (RH 18-25)																	
		DATAI	36 DATA SWITCHES																	
PI	004	CONO	CLEAR POWER FAIL FLAG	CLEAR PARITY FLAG	CLEAR PARITY ENABLE	SET PARITY ENABLE		CLEAR PI SYSTEM	REQUEST PI INTERRUPT ON CHANNELS	TURN ON CHANNELS	TURN OFF CHANNELS	TURN OFF PI ACTIVE	TURN ON PI ACTIVE	SELECT CHANNELS 1-7 FOR BITS 24,25,26						
		CONI	POWER FAILURE FLAG	PARITY ERROR FLAG	PARITY ENABLE			INTERRUPT IN PROGRESS ON CHANNELS 1-7					PI ACTIVE	CHANNEL ACTIVE 1-7						
		DATAO	36 BITS TO THE MEMORY INDICATORS																	
ADC A-D CONVERTER (AD10)	024																			
PTP	100	CONO												OUT OF TAPE	BINARY	BUSY	DONE FLAG	PIA		
		CONI										HOLE 8 UNLESS BINARY/BINARY			BINARY	BUSY	DONE FLAG	PIA		
		DATAO											HOLE 8 UNLESS BINARY	HOLE 7	HOLE 6	HOLE 5	HOLE 4	HOLE 3	HOLE 2	HOLE 1
PTR	104	CONI										TAPE FLAG			BINARY	BUSY	DONE FLAG	PIA		
		DATAI	36 BIT WORD IF BINARY, 8 BITS (28-35) IF NOT BINARY																	
TTY	120	CONO						TEST FLAG	TTI BUSY CLEAR	TTI FLAG CLEAR		TTO BUSY CLEAR	TTO FLAG CLEAR	TTI BUSY SET	TTI FLAG SET	TTO BUSY SET	TTO FLAG SET	PIA		
		CONI						TEST FLAG						TTI BUSY	TTI FLAG	TTO BUSY	TTO FLAG	PIA		
		DATAO	8 BIT CHARACTER TO TELETYPE																	
		DATAI	8 BIT CHARACTER FROM KEYBOARD																	
LPT (LP10)	124	CONO							128 CHAR	CLEAR PRINTER			BUSY				ERROR PIA	DONE PIA		
		CONI							128 CHAR	96 CHAR		ERROR	BUSY				ERROR PIA	DONE PIA		
		DATAO	CHARACTER	FIRST CHARACTER		SECOND CHARACTER			THIRD CHARACTER	FOURTH CHARACTER		FIFTH CHARACTER								
PLT (XY10)	140	CONO										BUSY			POWER ON	BUSY	DONE FLAG	DONE FLAG	PIA	
		CONI													POWER ON	PEN LOWER	PEN RAISE	PIA		
		DATAO													+Y DRUM DOWN	-X DRUM LOWER	+X DRUM	CLEAR DATA FLAG	DRUM LEFT	DRUM RIGHT
CR (CR10)	150	CONO	CLEAR READER					CLEAR READER	OFFSET CARD	READER STOP	CARD IN READER	HOPPER STACKER CARD	TROUBLE	READY TO READ	CLEAR END OF FILE	CLEAR END OF CARD	CLEAR DATA FLAG	PIA		
		CONI	READER TABLE ENAB	READER READY ENAB	PICK ERROR	PHOTO CELL ERROR	CARD MOTION ERROR					READ CARD	TROUBLE ENABLE	READING CARD	READER END OF FILE	END OF CARD	DATA FLAG	PIA		
		DATAI							ROW 12	ROW 11	ROW 0	ROW 1	ROW 2	ROW 3	ROW 4	ROW 5	ROW 6	ROW 7	ROW 8	ROW 9

C8

Indicator Panel, KA10 Arithmetic Processor, Bay 1

Indicator Panel, KA10 Arithmetic Processor, Bay 2

Indicator Panel, MB10 Core Memory (1.65 μs)

C10

POWERS OF TWO

2^N	N	2^{-N}
1	0	1.0
2	1	0.5
4	2	0.25
8	3	0.125
16	4	0.062 5
32	5	0.031 25
64	6	0.015 625
128	7	0.007 812 5
256	8	0.003 906 25
512	9	0.001 953 125
1 024	10	0.000 976 562 5
2 048	11	0.000 488 281 25
4 096	12	0.000 244 140 625
8 192	13	0.000 122 070 312 5
16 384	14	0.000 061 035 156 25
32 768	15	0.000 030 517 578 125
65 536	16	0.000 015 258 789 062 5
131 072	17	0.000 007 629 394 531 25
262 144	18	0.000 003 814 697 265 625
524 288	19	0.000 001 907 348 632 812 5
1 048 576	20	0.000 000 953 674 316 406 25
2 097 152	21	0.000 000 476 837 158 203 125
4 194 304	22	0.000 000 238 418 579 101 562 5
8 388 608	23	0.000 000 119 209 289 550 781 25
16 777 216	24	0.000 000 059 604 644 775 390 625
33 554 432	25	0.000 000 029 802 322 387 695 312 5
67 108 864	26	0.000 000 014 901 161 193 847 656 25
134 217 728	27	0.000 000 007 450 580 596 923 828 125
268 435 456	28	0.000 000 003 725 290 298 461 914 062 5
536 870 912	29	0.000 000 001 862 645 149 230 957 031 25
1 073 741 824	30	0.000 000 000 931 322 574 615 478 515 625
2 147 483 648	31	0.000 000 000 465 661 287 307 739 257 812 5
4 294 967 296	32	0.000 000 000 232 830 643 653 869 628 906 25
8 589 934 592	33	0.000 000 000 116 415 321 826 934 814 453 125
17 179 869 184	34	0.000 000 000 058 207 660 913 467 407 226 562 5
34 359 738 368	35	0.000 000 000 029 103 830 456 733 703 613 281 25
68 719 476 736	36	0.000 000 000 014 551 915 228 366 851 806 640 625
137 438 953 472	37	0.000 000 000 007 275 957 614 183 425 903 320 312 5
274 877 906 944	38	0.000 000 000 003 637 978 807 091 712 951 660 156 25
549 755 813 888	39	0.000 000 000 001 818 989 403 545 856 475 830 078 125
1 099 511 627 776	40	0.000 000 000 000 909 494 701 772 928 237 915 039 062 5
2 199 023 255 552	41	0.000 000 000 000 454 747 350 886 464 118 957 519 531 25
4 398 046 511 104	42	0.000 000 000 000 227 373 675 443 232 059 478 759 765 625
8 796 093 022 208	43	0.000 000 000 000 113 686 837 721 616 029 739 379 882 812 5
17 592 186 044 416	44	0.000 000 000 000 056 843 418 860 808 014 869 689 941 406 25
35 184 372 088 832	45	0.000 000 000 000 028 421 709 430 404 007 434 844 970 703 125
70 368 744 177 664	46	0.000 000 000 000 014 210 854 715 202 003 717 422 485 351 562 5
140 737 488 355 328	47	0.000 000 000 000 007 105 427 357 601 001 858 711 242 675 781 25
281 474 976 710 656	48	0.000 000 000 000 003 552 713 678 800 500 929 355 621 337 890 625
562 949 953 421 312	49	0.000 000 000 000 001 776 356 839 400 250 464 677 810 668 945 312 5
1 125 899 906 842 624	50	0.000 000 000 000 000 888 178 419 700 125 232 338 905 334 472 656 25
2 251 799 813 685 248	51	0.000 000 000 000 000 444 089 209 850 062 616 169 452 667 236 328 125
4 503 599 627 370 496	52	0.000 000 000 000 000 222 044 604 925 031 308 084 726 333 618 164 062 5
9 007 199 254 740 992	53	0.000 000 000 000 000 111 022 302 462 515 654 042 363 166 809 082 031 25
18 014 398 509 481 984	54	0.000 000 000 000 000 055 511 151 231 257 827 021 181 583 404 541 015 625
36 028 797 018 963 968	55	0.000 000 000 000 000 027 755 575 615 628 913 510 590 791 702 270 507 812 5
72 057 594 037 927 936	56	0.000 000 000 000 000 013 877 787 807 814 456 755 295 395 851 135 253 906 25
144 115 188 075 855 872	57	0.000 000 000 000 000 006 938 893 903 907 228 377 647 697 925 567 626 953 125
288 230 376 151 711 744	58	0.000 000 000 000 000 003 469 446 951 953 614 188 823 848 962 783 813 476 562 5
576 460 752 303 423 488	59	0.000 000 000 000 000 001 734 723 475 976 807 094 411 924 481 391 906 738 281 25
1 152 921 504 606 846 976	60	0.000 000 000 000 000 000 867 361 737 988 403 547 205 962 240 695 953 369 140 625
2 305 843 009 213 693 952	61	0.000 000 000 000 000 000 433 680 868 994 201 773 602 981 120 347 976 684 570 312 5
4 611 686 018 427 387 904	62	0.000 000 000 000 000 000 216 840 434 497 100 886 801 490 560 173 988 342 285 156 25
9 223 372 036 854 775 808	63	0.000 000 000 000 000 000 108 420 217 248 550 443 400 745 280 086 994 171 142 578 125
18 446 744 073 709 551 616	64	0.000 000 000 000 000 000 054 210 108 624 275 221 700 372 640 043 497 085 571 289 062 5
36 893 488 147 419 103 232	65	0.000 000 000 000 000 000 027 105 054 312 137 610 850 186 320 021 748 542 785 644 531 25
73 786 976 294 838 206 464	66	0.000 000 000 000 000 000 013 552 527 156 068 805 425 093 160 010 874 271 392 822 265 625
147 573 952 589 676 412 928	67	0.000 000 000 000 000 000 006 776 263 578 034 402 712 546 580 005 437 135 696 411 132 812 5
295 147 905 179 352 825 856	68	0.000 000 000 000 000 000 003 388 131 789 017 201 356 273 290 002 718 567 848 205 566 406 25
590 295 810 358 705 651 712	69	0.000 000 000 000 000 000 001 694 065 894 508 600 678 136 645 001 359 283 924 102 783 203 125
1 180 591 620 717 411 303 424	70	0.000 000 000 000 000 000 000 847 032 947 254 300 339 068 322 500 679 641 962 051 391 601 562 5
2 361 183 241 434 822 606 848	71	0.000 000 000 000 000 000 000 423 516 473 627 150 169 534 161 250 339 820 981 025 695 800 781 25
4 722 366 482 869 645 213 696	72	0.000 000 000 000 000 000 000 211 758 236 813 575 084 767 080 625 169 910 490 512 847 900 390 625

APPENDIX D

ALGORITHMS

All arithmetic operations on full and half words are performed in the 36-bit parallel adder. There are two sets of summand inputs to the adder, each set of 36 supplying one input to each adder stage. One set supplies the contents of AR, its complement, or zero; the other set supplies the contents of BR, its complement, or zero. Each stage also has a carry input, which is generated by the next less significant stage. Every stage has two outputs; the carry already mentioned, and a sum. The 36 sum outputs together form the sum of the two input words. The least significant stage has a carry input from the logic for performing twos complement arithmetic and incrementing by one. The negative of a number is formed at the sum outputs simply by supplying the complement of the number at one set of inputs and asserting the carry into stage 35. Adder stage 17 has extra input gating so that 1 can be added to or subtracted from both halves of AR simultaneously.

The adder produces a sum in the same way that one adds binary numbers using pencil and paper. Each adder stage has three inputs, two summand bits and a carry, and two outputs, sum and carry. The sum output of a given stage is 1 if any one or all three of the inputs are 1. The carry out is 1 if two or three of the inputs are 1. Calculations are performed as though the words represented 36-bit unsigned numbers, *ie* the signs are treated just like magnitude bits. In the absence of a carry into the sign stage, adding two numbers with the same sign produces a plus sign in the result. The presence of a carry gives a positive answer when the summands have different signs. The result has a minus sign when there is a carry into the sign bit and the summands have the same sign, or the summands have different signs and there is no carry.

Thus the program can interpret the numbers processed in fixed point arithmetic as signed numbers with 35 magnitude bits or as unsigned 36-bit numbers. A computation on signed numbers produces a result which is correct as an unsigned 36-bit number even if overflow occurs, but the hardware interprets the result as a signed number to detect overflow. Adding two positive numbers whose sum is greater than or equal to 2^{35} gives a negative result, indicating overflow; but that result, which has a 1 in the sign bit, is the correct answer interpreted as a 36-bit unsigned number in positive form. Similarly adding two negatives gives a result which is always correct as an unsigned number in negative form.

All operations discussed below have two operands, one of which is supplied to the adder from BR, which acts simply as a buffer and has no special input gating. MQ has shift gating so it can function as a low order extension of AR for handling double length operands. All actual computations take place in the single 36-bit adder, but the sum output can be placed in either AR or MQ, and all transfers to MQ from AR or BR are made through the adder. In multiplication MQ holds the multiplier and thus

controls the summation of partial products; as the multiplier is shifted out, the low order word of the product is shifted in. In division MQ supplies the low order part of the dividend to AR as the quotient is being constructed in MQ.

In any extended arithmetic operation, the requisite number of steps is counted in the 9-bit shift counter SC, which has a carry network for this purpose. SC also has a 9-bit adder for use in computations on floating point exponents and size and position calculations in byte manipulation.

FIXED POINT ALGORITHMS

Fixed point numbers are explained in detail in §1.1. For convenience let us take the computer representation of the positive number x as $+[x]$ where the brackets enclose the number in bits 1–35. Similarly the representation of $-x$ is $-[2^{35} - x]$ or $-[1 - x]$ depending on whether we are regarding numbers as integers or as proper fractions. The most negative number, -2^{35}, has the form $-[0]$, which is equivalent to the unsigned integer 2^{35}.

Addition. There are four cases of addition of two positive 35-bit numbers x and y.

$$\begin{array}{lll}
\text{I.} & x + y & \\
\text{II.} & (-x) + (-y) & \\
\text{III.} & x + (-y), & x \geqslant y \\
\text{IV.} & x + (-y), & x < y
\end{array}$$

The operands are held in AR and BR, but it makes no difference which one is in which register. The result appears in AR. For convenience in the exposition we shall regard the numbers as proper fractions; to view them as integers, simply substitute "2^{35}" for each occurrence of "1". Since the twos complement format allows a representation for -1, either x or y may be 1 in II, and y may be 1 in IV.

I. If $x + y < 1$ the adder output placed in AR is $+[x + y]$. If $x + y \geqslant 1$ the carry out of stage 1 changes the sign. Consequently if the addition of two positive numbers gives a negative result, it is apparent that the sum exceeds the capacity of the register. The processor detects the overflow by checking the sign carries: there is a carry into the sign stage but none out of it. AR then contains

$$- [x + y - 1]$$

II. Ignoring the carry into the sign bit in the addition of two negatives would give

$$\begin{array}{l}
- [1 - x] \\
- [1 - y] \\
\hline
+ [1 + 1 - x - y]
\end{array}$$

If $x + y \leqslant 1$ the carry changes the sign and the result is

$$- [1 - x - y]$$

which is the representation of $-(x + y)$. If $x + y > 1$ there is no carry into the sign, and its absence in the presence of a carry out indicates overflow. AR contains

$$+ [1 - (x + y - 1)]$$

III. Ignoring the carry into the sign in an addition where the signs are different would give

$$
\begin{array}{l}
+ [x] \\
- [1 - y] \\
\hline
- [1 + x - y]
\end{array}
$$

Since $x \geqslant y$, it follows that $1 + x - y \geqslant 1$. Hence the carry changes the sign and the result is

$$+ [x - y]$$

When the operand signs are different, the magnitude of the result cannot exceed the larger operand magnitude and there can be no overflow. Since in this case the positive number is at least as large in magnitude as the negative, there is always a carry into the sign, and this added to the operand minus sign produces a carry out.

IV. The addition of numbers of differing signs where the negative has the larger magnitude gives

$$
\begin{array}{l}
+ [x] \\
- [1 - y] \\
\hline
- [1 + x - y]
\end{array}
$$

Since $x < y$, then $1 + x - y < 1$. Hence there are no carries associated with the sign and no overflow. The above result is the twos complement representation of $x - y$, ie $-(y - x)$.

Subtraction. The minuend from AC is in AR, and the subtrahend, which is either $0, E$ or the word from location E, is in BR. Subtraction is done directly by adding the twos complement of BR to AR. The logic supplies the complement of BR to the adder and a carry into the adder LSB.

Let x be the absolute value of the number in AR, and y the absolute value of the number in BR. There are four cases.

I. $x - (-y)$
II. $(-x) - y$
III. $x - y,$ $x \geqslant y;$ $(-x) - (-y),$ $x \leqslant y$
IV. $x - y,$ $x < y;$ $(-x) - (-y),$ $x > y$

These correspond respectively to the four cases of addition discussed previously.

Multiplication. The multiplier, $0, E$ or the contents of location E, is in MQ, and the multiplicand from AC is in BR. AR is clear. The 36-step procedure is as follows.

If MQ35 (the multiplier LSB) is 1, subtract BR algebraically from AR, but put the result in AR shifted one place to the right, with the LSB of the result going into MQ0, and shift MQ right so a bit of the multiplier is dropped from MQ35. Put the sign of the result in AR0 and AR1 (as though the shift followed the subtraction and did not affect the sign but did move it to AR1). If MQ35 is 0, simply shift AR and MQ right one, with AR35 going into MQ0.

In each subsequent step perform only the shift if the bits moved in and out of MQ35 on the previous step were the same. If they were different, add or subtract along with the shift: if the shift moved a 0 in and a 1 out, add BR to AR; if a 1 in and a 0 out, subtract BR from AR.

Thus the low order bits of the running sum of partial products are shifted into MQ as the multiplier is shifted out. At each step the effect of the multiplicand in BR on the partial sum in AR is one binary order of magnitude greater than in the preceding step because the partial sum was shifted right. Therefore BR can be combined directly with AR. If MQ35 is initially 0, there is no subtraction until a 1 is shifted into it. Simple shifting then continues until the next transition (from 1 to 0), following which BR is added.

The process continues in this way, subtracting at every 0-1 transition, adding at every 1-0 transition. After 35 steps MQ0–34 contains the low half of the product magnitude, and MQ35 contains the sign of the multiplier. At the final step, add or subtract as required but put the result directly into AR; shift only MQ to move the low magnitude into the correct position and make MQ0 equal to the sign of the whole product.

If the original operands were both negative and the result is also negative, set Overflow; this can occur only when -2^{35} is squared. In IMUL, if the high word is not null (*ie* if AR is neither clear nor all 1s), set Overflow; move MQ to AR for storage of the low word.

To see that this procedure results in a correct product, consider the positive binary integer

$$1\ 0\ 0\ 1\ 1\ 1\ 0\ 1\ 1$$
$$8\ 7\ 6\ 5\ 4\ 3\ 2\ 1\ 0$$

where the decimal digits below the binary digits are the powers of 2 corresponding to the bit positions. This number is obviously equal to

$$\begin{array}{r} 1\ 0\ 0\ 0\ 0\ 0\ 0\ 0\ 0 \\ +\quad\quad 1\ 1\ 1\ 0\ 0\ 0 \\ +\quad\quad\quad\quad\quad 1\ 1 \end{array}$$

Now an *n*-bit string of 1s whose rightmost bit corresponds to 2^k is equal to $2^{k+n} - 2^k$, or equivalently $2^k(2^n - 2^0)$, *ie* $2^n - 2^0$ is a string of *n* 1s and the 2^k shifts the string left *k* places. Hence

$$\begin{array}{rcccl}
1\ 0\ 0\ 0\ 0\ 0\ 0\ 0\ 0 & = & 2^{8+1} - 2^8 & = & 2^9 - 2^8 \\
1\ 1\ 1\ 0\ 0\ 0 & = & 2^{3+3} - 2^3 & = & 2^6 - 2^3 \\
\underline{1\ 1} & = & 2^{2+0} - 2^0 & = & \underline{2^2 - 2^0} \\
1\ 0\ 0\ 1\ 1\ 1\ 0\ 1\ 1 & & & = & 2^9 - 2^8 + 2^6 - 2^3 + 2^2 - 2^0
\end{array}$$

In this last representation, each power of 2 that is subtracted corresponds to

a transition from 0 to 1 (scanning right to left), whereas each that is added corresponds to a 1-0 transition. The largest term corresponds to the transition to the sign bit, which is 0 for a positive number. The multiplication algorithm interprets the multiplier in this manner, alternately subtracting and adding the multiplicand to the partial sum in the order-of-magnitude positions corresponding to the transitions. If a multiplier of the same magnitude were negative, it would have the form

$$1\ 0\ 1\ 1\ 0\ 0\ 0\ 1\ 0\ 1$$
$$-\ 8\ 7\ 6\ 5\ 4\ 3\ 2\ 1\ 0$$

in which the extra bit at the left represents the sign. The number is now equivalent to

$$-2^9 + 2^8 - 2^6 + 2^3 - 2^2 + 2^1 - 2^0$$

wherein opposite signs correspond to opposite transitions. The algorithm may thus use exactly the same sequence for a negative multiplier: this time the subtraction of greatest magnitude is detected by the transition to the sign bit, which is now 1.

Division. The divisor, $0, E$ or the contents of location E, is in BR. In DIV the high and low halves of the dividend from two accumulators are in AR and MQ respectively. In IDIV the one-word dividend from AC is in AR. The two types of division differ mainly in setting up the dividend; in both cases the algorithm processes a positive dividend to get a positive quotient.

In DIV if the dividend is negative (AR0 = 1), make it positive and set the negative dividend flag. To negate the dividend, move the low word to AR and the complement of the high word to MQ. Then move the negative of the low word back to MQ and the complement of the high word back to AR. Now the double length negative of the original dividend is in AR and MQ unless MQ is clear; in this event add 1 to AR to give the twos complement negative of the high word. Once the dividend is in positive form shift MQ left one place to close the hole between the two halves; in other words drop the low sign and get the 70-bit magnitude into AR1–35, MQ0–34.

If the IDIV dividend in AR is negative, negate it and set the negative dividend flag. Move the one word dividend in positive form to MQ and clear AR. Shift MQ left, as the algorithm operates on a double length dividend in both types of division although the high part is null in this case.

After the dividend is set up, compare the divisor with it to determine whether the division can be performed. Subtract the absolute value of the divisor from the high half of the dividend (if the divisor is positive, subtract it; if negative, add it). Since the dividend is positive, the result is also positive if the magnitude of the divisor is less than or equal to the number in AR. For a fixed fraction, the divisor is subtracted from the actual dividend and no overflow is allowed. For a fixed integer, AR is clear and the result is positive only for a zero divisor; the worst possible case is the division of $2^{35} - 1$ by 1, whose integral result can be accommodated. (Placing the one word dividend in MQ effectively multiplies it by 2^{-35}, making it the fractional part of a two word dividend with the binary point in the middle. The quotient is then a proper fraction, which is multiplied by 2^{35} simply by interpreting it as an integer.) Thus if the result of this initial subtraction is

positive, set Overflow and No Divide, and terminate the procedure so the processor goes on to the next instruction. Dividing by zero is of course meaningless. The reason for prohibiting a fractional division where the result would be greater than 1 is that it is impossible to determine the position of the binary point in the quotient. So it is up to the programmer to shift the dividend to the correct position beforehand. If the result of the initial subtraction is negative, the division can be performed and the processor goes into the division loop.

In division on paper, one subtracts out the divisor the number of times it goes into the dividend, then shifts the dividend one place to the left (or the divisor to the right) and again subtracts out. In binary computations the divisor goes into the dividend either once or not at all. Each subtraction of the divisor thus generates a single bit of the quotient. If the subtraction leaves a positive difference, *ie* if the dividend is larger than the divisor, a 1 is entered into the quotient. If the difference is negative, a 0 is entered. To compensate for subtracting too much, the hardware could add the divisor back into the dividend before going to the next subtraction step. But the PDP-10 algorithm instead shifts first and adds the divisor back in at the new position. It then continues to shift and add putting 0s into the quotient until the result again becomes positive. This procedure generates the same quotient without ever going back a step.

The hardware procedure is as follows. As each addition or subtraction is formed in the adder, put the result in AR shifted one place to the left with AR35 receiving a new bit of the dividend from MQ0, and shift MQ left bringing in a bit of the quotient at MQ35. The bit brought in is the complement of the sign from the adder: if the divisor does not go into the dividend, the resulting minus sign (1) produces a 0 quotient bit; if the divisor does go in, the plus sign gives a 1. Each step loads one bit of the quotient into MQ35, and the low half of the dividend is shifted out of MQ as the quotient is shifted in.

The first step is the test subtraction. In each subsequent step, subtract the absolute value of the divisor if the quotient bit generated in the previous step is 1, but add it back in if the quotient bit is 0. Since the divisor may have either sign, subtract it algebraically if its sign differs from the quotient bit, add it if its sign is the same.

The hardware executes 36 steps to generate 35 magnitude bits. The initial test step must give a 0, which serves as the sign since we are producing a positive quotient. In the final step put the result of the addition or subtraction directly in AR without shifting so the remainder is in the correct position, but shift MQ left putting the sign from the first step in MQ0 and bringing in the last quotient bit. (The bit dropped out of MQ0 is superfluous; it was brought into MQ35 when the hole was closed between the dividend halves.)

To complete the division we must make sure the remainder is correct and determine the correct signs of the results. Since the operations were performed on positive operands, the remainder should also be positive. A negative remainder indicates that too much has been subtracted. To correct this add the absolute value of the divisor back in. If the negative dividend flag is set, negate AR so the remainder has the sign of the original dividend.

Now move the corrected remainder to MQ and move the quotient to AR.
If the negative dividend flag and the divisor sign are of opposite states,
negate AR to produce the correct quotient sign. The correct quotient and
remainder are now in AR and MQ ready for storage.

As an example of the way this algorithm operates, consider a division of
3-bit fixed fractions with a dividend of +.100100 and a divisor of +.101.
By paper computation we obtain the quotient this way.

```
              .111
      101⌐100.100
           10 1
           10 00
            1 01
             110
             101
               1
```

Taking the processor registers to be four bits in length, AR contains 0.100,
MQ has 0.100, and BR has 0.101. Before starting we close the hole changing
MQ to 1.000. The sequence has four steps.

```
              0.100    1.000
             -0.101
              1.111
     1 ←      1.111    0.000
            + 0.101
              0.100
     2 ←      1.000    0.001
             -0.101
              0.011
     3 ←      0.110    0.011
             -0.101
              0.001
     4        0.001   ←0.111
```

The quotient is in MQ at the right, the remainder in AR at the left.

FLOATING POINT ALGORITHMS

§1.1 explains floating point numbers and §2.6 discusses the general charac-
teristics of floating point arithmetic. Exponent computations are done in
the SC adder using the exponents and signs from the floating point operands.
Remember, the sign is that of the whole number, not of the exponent.
Although bits 1–8 of a floating point number represent an exponent in the
range −128 to +127, the discussion is entirely in terms of the excess 128
exponents in positive form, ie the set of numbers 0–255. Computations
generally use twos complement operations even though the exponent in a

ALGORITHMS

negative number is a ones complement. The SC sign bit is used to detect exponent overflow and underflow.

After exponent calculations are complete, operations on the fractions are done by the fixed point logic in AR, BR and MQ. Bits 1–8 of AR and BR are filled with null bits, 0s in a positive number, 1s in a negative. Double length operands are in AR and MQ with MQ8–35 forming a magnitude extension of AR. In almost all circumstances the logic treats AR0–35 and MQ8–35 as a single 64-bit register; in all two-word shifting AR35 is connected to MQ8 and MQ0–7 is ignored. Except in division the fixed point calculation generates a double length fraction, which is shifted arithmetically (in right shifting the sign goes into AR1; in left shifting the sign is unaffected and 0s enter MQ35). Almost all floating point instructions normalize the result, thus making use of the low order part even though the instruction may store only the high order word.

Addition, Subtraction. $E,0$ or the word from location E is in BR, and AC is in AR. For subtraction move the negative of the subtrahend from BR to AR and move the minuend from AR to BR. This reduces subtraction to addition, so the rest of the algorithm is the same for both.

The initial objective is to determine the difference between the exponents and to determine which exponent is the larger. If the signs of the operands differ, add the exponents into SC. If the signs are the same, subtract the BR exponent from the AR exponent by adding the twos complement. Let x and y be the AR and BR exponents in positive form. The table below shows the calculations as a function of the operand signs, and the sign of the result in SC as a function both of the operand signs and the relative values of x and y.

AR+, BR+	AR+, BR−	AR−, BR+	AR−, BR−
$+[x]$	$+[x]$	$-[255-x]$	$-[255-x]$
$-[256-y]$	$-[255-y]$	$+[y]$	$+[1+y]$
$-[256+x-y]$	$-[255+x-y]$	$-[255-x+y]$	$-[256-x+y]$

SC+	SC−	SC+	SC−	SC+	SC−	SC+	SC−
$x \geqslant y$	$x < y$	$x > y$	$x \leqslant y$	$x < y$	$x \geqslant y$	$x \leqslant y$	$x > y$

As can be seen from the above, if AR already contains the number with the smaller exponent, the SC and AR signs differ. Hence if the SC and AR signs are the same, switch BR and AR so the number with the smaller exponent can be shifted. If the exponents are equal, the signs may or may not be the same but it matters not whether the transfer takes place.

To control the shifting we must now get the negative of the difference between the exponents. Let d be $|x - y|$. There are four cases as a function of the SC sign and whether the AR and BR signs are equal. The second column lists the present contents of SC, the third tells what must be done to arrive at $-[256 - d]$ in SC.

SC+, AR0 = BR0	$+[d]$	Negate SC
SC+, AR0 ≠ BR0	$+[d-1]$	Complement SC

| SC−, AR0 = BR0 | −[256 − d] | Do nothing |
| SC−, AR0 ≠ BR0 | −[255 − d] | Add 1 to SC |

If $d < 64$ (indicated by a negative SC with a 0 in either SC1 or SC2) nullify AR1−8 and shift AR and MQ right d places so its bits correctly match the BR bits in order of magnitude. If $d > 64$ clear AR for its contents are of no significance.

Now move the larger exponent from BR to SC in positive form, nullify BR1−8, and add BR and AR into AR as fixed fractions. Finally enter the normalizing sequence.

This sequence first tests for a zero result. If AR and MQ8−35 are clear, bypass the rest of the procedure. If the fractional result has overflowed into AR8 (indicated by AR0 ≠ AR8 or AR8 = 1 and AR9−35 = 0), shift right and increase the exponent by one. The number is now normalized.

Complement the exponent in SC. If the instruction is not UFA and the number is not normalized go into the normalizing loop. In each step shift the double length fraction left and add 1 to the negative exponent (decreasing its magnitude by 1). Terminate the loop when the fraction is normalized, indicated by the sign and the MSB of the fraction being different (AR0 ≠ AR9) or the magnitude being ½ (AR9 = 1 and AR10−35 = 0).

If the instruction specifies rounding, adjust the high fraction so it is rounded and is in twos complement form if negative. The rounding is away from zero. For a positive result the high fraction must be increased if the low fraction is greater than half the value of the high fraction LSB. In a negative result the high fraction is a ones complement, which is one greater in magnitude than the twos complement. Hence it is already rounded and should be decreased in magnitude if the low fraction is < ½LSB. In either case add 2^{-27} into AR if MQ8 is 1 unless MQ9−35 is clear in a negative number. A 1 in MQ8 indicates a low fraction ⩾ ½LSB in a positive number, ⩽ ½LSB in a negative number. The condition that MQ9−35 not be zero in a negative number is the case where the low fraction is exactly ½LSB. If the high fraction is actually changed, renormalize it. A single normalizing shift is all that is required and it occurs in only two cases: a right shift when $1 − 2^{-27}$ is rounded, a left shift when −½ is changed to a correct twos complement.

Once the number has been normalized (and rounded if necessary) the exponent is in negative form. Thus if the SC sign bit is 0, set Overflow and Floating Overflow. If SC1 is also 0, the sign bit must have been changed by decreasing the exponent, so also set Floating Underflow (the maximum possible exponent overflow is 128 giving an SC contents of 777_8, and this can occur only in division). Insert the exponent in correct form into AR1−8.

The result is now ready to store from AR unless the instruction is in long mode. To ready the double length result subtract 27 from the positive exponent in SC. Save the high word in MQ, and move the low word to AR, but only if the decreased exponent is still positive. If the sign is 1, the true exponent of the low word is less than −128, so clear AR. (Note that this condition is also true if the low exponent is > 127, which can occur only if the high exponent is > 154.) If the low word is nonzero, shift AR right one place to put the fraction in bits 9−35 (remember that all shift operations

use MQ8–35), clear AR0 so the low word has a positive sign even if the double length fraction is negative, and insert the low exponent in positive form in bits 1–8. Finally switch AR and MQ so the high and low words are in correct position for storage.

Scaling. The 9-bit signed scale factor from bits 18 and 28–35 of E is in SC, and AC is in AR and BR. If the floating point number being scaled is positive, simply add the sign and exponent from BR0–8 to SC; if the number is negative, add the complement of BR0–8 to SC. Let x be the exponent in positive form and let y be the absolute value of the scale factor. There are only two cases,

$$
\begin{array}{ll}
+[x] & +[x] \\
+[y] & -[256-y] \\
\hline
+[x+y] & +[x-y]
\end{array}
$$

and in either the result is in positive form in SC.

Now enter the normalizing sequence described under floating addition. Only left shifting can occur bringing 0s in from MQ. The result can be zero, and exponent overflow or underflow can occur; but there is no rounding, and at the end the one-word result is in AR ready for storage.

Multiplication. $E,0$ or the word from location E is in BR, and AC is in AR. Place the AR exponent in positive form in SC, and add the positive form of the BR exponent to it. Since both are in excess-128 code, subtract 128. Save the result in the floating exponent register FE so SC can be used to control the multiplication of the fractions.

Nullify the exponent parts of AR and BR. Move the multiplier from BR to MQ and the multiplicand from AR to BR. Clear AR. Now multiply the fractions by the same procedure given for fixed point multiplication with the following differences:

♦ There are only 28 steps instead of 36.

♦ The shift register extension of AR for the construction of the product is MQ8–35. As the multiplier is shifted out, bits of the product come in at MQ8.

♦ In the final step place the adder output directly into AR but do not shift MQ – the low fraction is in MQ8–34, the correct position for normalization.

Clear MQ35, move the exponent back to SC, and enter the normalizing sequence described under floating addition. If the operands are normalized, at most one left shift is needed to normalize the result.

Division. The divisor, $E,0$ or the contents of location E, is in BR. The dividend from AC is in AR. In long mode the low half of the dividend from the second accumulator is in MQ; otherwise MQ is clear.

If the dividend is negative, make it positive and set the negative dividend flag. Except in long mode, negate the dividend simply by negating AR. For long mode follow the procedure given for DIV in the second paragraph of the fixed division algorithm. With a floating point operand the left MQ shift puts the low fraction in MQ8–34.

Place the AR exponent in positive form in SC. Subtract the magnitude of the BR exponent from it by adding the negative form of the exponent (ones complement) plus 1. Since the excess-128 factors cancel in the subtraction, add 128. Save the result in the floating exponent register FE so SC can be

used to control the division of the fractions.

Nullify the exponent parts of AR and BR. Subtract the absolute value of the divisor from the high half of the dividend. If the result is positive, indicating the divisor is less than or equal to the dividend, shift AR and MQ right and increase the exponent in SC by 1. Save the adjusted exponent in FE. The shift divides by 2, so if the operands are normalized, the dividend must now be less than the divisor.

Now divide the fractions by the same procedure given for fixed point fractional division with the following differences:

♦ Since the dividend has already been adjusted, the test in the first step stops the division only if the divisor is zero, or is unnormalized and less than the dividend. A normalized divisor cannot cause the quotient to overflow. If the result of the initial subtraction is positive, terminate the procedure and set Floating Overflow as well as Overflow and No Divide.

♦ Instead of 36 steps there are only 29 if the instruction specifies rounding, otherwise 28.

♦ The shift register extension of AR is MQ8−35. As quotient bits are brought in at MQ35, dividend bits are supplied to AR35 from MQ8. The shifting clears MQ0−7.

♦ The MQ shift in the final step places a 27-bit quotient fraction in MQ9−35 or a 28-bit fraction in MQ8−35.

♦ As in the fixed point algorithm generate the correct signed remainder, put it in MQ, and move the quotient to AR but leave it positive.

If the instruction specifies rounding, shift AR right placing the 27-bit fraction in the correct position, and if the bit shifted out of AR35 is 1, add it back into AR35 to round the positive quotient. If the quotient is zero bypass the rest of the procedure. The reaminder will also be zero except in an FDVL where the double length dividend is unnormalized and its high fraction is zero.

Complement the exponent in SC. If the instruction uses normalized operands the initial dividend adjustment guarantees that the quotient will be normalized. If it is not, shift AR left (bringing 0s into AR35) until a 1 appears in AR9, each time increasing the negative exponent by 1 (decreasing its magnitude).

Since the exponent is in negative form, if SC0 is 0, set Overflow and Floating Overflow. If SC1 is also 0, the sign bit must have been changed by decreasing the exponent, so also set Floating Underflow. Insert the exponent in correct form into AR1−8. If the negative dividend flag and the divisor sign (BR0) are of opposite states, negate AR to produce the correct quotient sign.

The quotient is now ready for storage from AR and the remaining operations are performed only for long mode. Save the quotient in BR and bring the high half of the original dividend from AC to AR. Put the dividend exponent in SC. Decrease its magnitude by 26 if the dividend was shifted right at the beginning to allow the division to be performed; otherwise decrease it by 27. Move the remainder to AR and insert the exponent in it provided the remainder is not zero and the exponent is within the proper range, −128 to 127 (the test is that the sign resulting from the exponent calculation is the same as the sign of the remainder). If the exponent is

outside that range clear AR; the assumption is that the remainder is of no significance (*ie* the exponent is too small). Move the remainder with its correct exponent from AR to MQ and put the quotient back in AR. The two words are now ready for storage.

Double Precision Division. The software routine that performs double precision floating point division and the algorithm it utilizes are given at the end of §2.11. FDVL performs the division

$$A/b \;=\; q + r2^{-27}/b$$

where q and r are the quotient and remainder. In a double precision division the divisor is of the form

$$B \;=\; b + d2^{-27}$$

Using the expansion

$$\frac{1}{x+y} \;=\; \frac{1}{x}\left[1 - \frac{y}{x} + \frac{y^2}{x^2} - \frac{y^3}{x^3} + \ldots\right] \qquad\qquad (y^2 < x^2)$$

and letting $x = b$ and $y = d2^{-27}$ gives

$$\frac{A}{B} \;=\; \left(q + \frac{r2^{-27}}{b}\right)\left[1 - \frac{d2^{-27}}{b} + \frac{d^2 2^{-54}}{b^2} - \frac{d^3 2^{-81}}{b^3} + \ldots\right]$$

Multiplying out and gathering like terms gives

$$\frac{A}{B} \;=\; q + \frac{1}{b}(r - qd)2^{-27} - \frac{d}{b^2}(r - qd)2^{-54} + \frac{d^2}{b^3}(r - qd)2^{-81} - \ldots$$

where the first two terms on the right are those in the equation at the bottom of page 2-67.

The ratio of adjacent terms is

$$\frac{T_{n+1}}{T_n} \;=\; \frac{-d2^{-27}}{b}$$

In an alternating convergent series, the error due to truncation is smaller than the first term dropped. Therefore

$$|Error| \;<\; \frac{d2^{-27}}{b}\,T_n$$

Since the maximum value of d is less than 1 and the minimum value of b (normalized) is ½,

$$|Error| \;<\; T_n 2^{-26}$$